

Structural autocatalysis can be common while productive operation is rare: an exact separation theorem in a capped-Zipf binary polymer model

Author name to be inserted

Affiliation to be inserted

13 September 2026

Abstract

A reflexively autocatalytic and food-generated reaction set (RAF) certifies that a reaction network is closed under the resources and catalysts it needs. It does not say whether a reactor started from food alone will actually produce material. We compare the two questions in one explicit random model: the reversible binary polymer model with molecules of length at most n , food consisting of the monomers and dimers, and the capped-Zipf catalysis law of Hordijk and Steel at the critical exponent sequence $a_n = 2 - 2/n$, under which each molecule catalyzes on average $9n/\pi^2$ reactions. Three statements are proved for the same sampled network. First, the probability that a RAF exists converges to $\theta(q_*) \in (0, 1)$, the survival probability of an infinite independent reversible reaction-closure process at openness $q_* = 1 - e^{-9/\pi^2}$. Second, conditional on existence, the minimum RAF size M_n satisfies $\log M_n/n \rightarrow 0$ in probability, with an explicit subexponential cutoff attaining the full limiting mass. Third, in a fed, diluted stochastic mass-action reactor with all spontaneous reactions active, the probability that the reactor keeps its total mass bounded and exports a fixed positive amount of nonfood polymer during a fixed observation window is bounded above and below by constant multiples of the single-incidence probability $p_n \sim 9/(\pi^2 X_n)$, where $X_n = 2^{n+1} - 2$ is the number of molecular species, up to an explicit molecular-noise error that a sufficient quadratic volume scale makes negligible. The same order holds jointly with, and conditionally on, the existence of some RAF, and the same conclusions hold for the fraction of catalytic environments in which a reactor run succeeds with high probability. Removing catalysis makes the same output observable asymptotically negligible. Thus structural autocatalysis occurs with probability bounded away from zero and small structural certificates are typical, while the specified productive operation has probability of order $1/X_n$, decaying exponentially in n . The structural limits, the finite probability sandwich, and their asymptotic corollaries are formally verified in Lean 4; the remaining steps are conventional proofs given in full.

1 Introduction

Collectively autocatalytic sets were proposed by Kauffman [12] as a route by which a random chemistry could bootstrap a self-sustaining organization without a single self-replicating molecule. Their combinatorial formalization is the notion of a reflexively autocatalytic and food-generated set, or RAF, introduced by Steel [16] and developed with Hordijk [8, 17]: a nonempty set of reactions that can generate every molecule it needs from an ambient food set, and whose every reaction has a catalyst among the molecules generated. In the binary polymer model [8, 12, 16], where molecules are binary strings and reactions are ligations and cleavages, RAFs appear once the mean number of reactions catalyzed per molecule grows linearly in the maximum string length n [11, 14]. Hordijk

and Steel [10] showed that the distribution of catalytic activity across molecules affects both the appearance and the size of RAFs, and singled out power-law (capped Zipf) catalysis as a natural heterogeneous alternative to uniform catalysis.

A RAF is a statement about closure. It says nothing by itself about waiting times, competition between reactions, dilution, or the amount of material that a reactor actually exports; this gap between structural and dynamical autocatalysis has been discussed since the first kinetic simulations of the model [5] and is central to critiques of metabolism-first scenarios [9, 20]. This paper quantifies the gap for one fully specified random model. We fix the capped-Zipf sampler at the critical sequence $a_n = 2 - 2/n$, place the sampled network in a fed, diluted, well-mixed stochastic mass-action reactor whose only initial molecules are food, and measure a deliberately physical event: total polymer mass stays bounded and a fixed positive amount of nonfood polymer is exported during a fixed time window. Under one probability space that samples the catalytic environment, the kinetic coefficients, and the reactor trajectory, we prove the following.

- **Structural existence.** The probability that a RAF exists converges to $\theta(q_*)$, a number strictly between 0 and 1 (Theorem 3.1(i)).
- **Structural size.** Conditional on existence, the minimum RAF has subexponential size: for every $b > 0$, $\mathbb{P}(M_n \leq 2^{bn} \mid \text{RAF exists}) \rightarrow 1$, and one explicit stretched-exponential cutoff captures the whole limiting mass (Theorem 3.1(ii)).
- **Productive rarity.** At a sufficient quadratic volume scale, the probability of the productive event is $\Theta(p_n) = \Theta(1/X_n)$, where p_n is the probability that a fixed molecule catalyzes a fixed reaction and $X_n = 2^{n+1} - 2$ is the number of species; the same order holds jointly with, and conditionally on, RAF existence, and also for the fraction of catalytic environments in which a reactor run succeeds with high probability (Theorem 3.1(iii) and Corollary 3.2). With catalysis switched off, the same observable has probability $o(p_n)$ (Theorem 3.1(iv)).

So the structural certificate is common and typically small, while the productive requirement is exponentially rare in n , in the same sampled network and after conditioning on the certificate. The theorem is an asymptotic probability statement with explicit but very conservative constants; it is not a calibrated prediction of a transition, and it does not say that RAFs are chemically inactive. Section 8 states precisely what is and is not implied.

What is new and what is imported. The structural part sharpens the picture of [10] for the capped-Zipf law: existence has a nontrivial limit, and the limit is the survival probability of an infinite independent reversible closure process. The same survival function governs the uniform model in the companion critical-window paper [1], where the uniform RAF probability at linear intensity λ converges to $\Theta_{\text{unif}}(\lambda) = \theta(1 - e^{-\lambda})$; here the mean catalytic degree per unit length tends to $\Lambda = 9/\pi^2$ and the limit is $\theta(1 - e^{-\Lambda})$, so at the level of the limiting existence probability the heavy-tailed sampler behaves like the uniform sampler of the same intensity (Remark 3.3). The productive part is a kinetic probability law for a literal stochastic mass-action reactor, proved from the marked jump kernel rather than through a fluid-limit hypothesis. Stochastic and kinetic studies of autocatalytic polymer networks are classical [5–7]; our contribution is the two-sided probability law and its formal verification, not the introduction of stochastic autocatalysis. The lower bound constructs one self-catalyzed ligation and proves that it produces reliably in the presence of arbitrary additional nonfood catalysts; the upper bound applies to every mechanism that meets the output requirement.

Formal verification. The structural limits, the finite productive probability sandwich, its disabled-catalysis counterpart, and their asymptotic corollaries are formalized in Lean 4 [4] over Mathlib [18]. Appendix D lists the exact declarations. The Lean development for the reactor imports the Lean development for the sampler, so the compatibility of the two halves of the theorem is machine-checked: both halves refer to one configuration weight. Some consequences stated in prose (the environment-level reformulation, the interpretation of cutoff events as statements about M_n , and the elementary catalog identities) are conventional deductions from the verified statements; we say so where they occur.

Organization. Section 2 fixes the sampler, the RAF notion, the reactor, the observable, and the probability laws. Section 3 states the theorem and its environment-level corollary. Section 4 proves the structural statements, Sections 5–6 the productive lower and upper bounds, and Section 7 the source asymptotics and the conditional conclusion. Section 8 discusses scope. Appendices A–C contain the technical estimates, and Appendix D the formal concordance.

2 The shared model

Every statement in this paper uses the objects of this section. The sampler, the RAF definition, and the catalog conventions are literally those of the Lean development `PowerLawSmallRAF`; the reactor is that of `RandomViability`, which imports the former.

2.1 Molecules, channels, and food

For $n \geq 4$ let $\mathcal{X}_n$ be the set of nonempty binary words of length at most n , and write $\ell(x)$ for the length of x . The food set is $F = \mathcal{X}_2$, the two monomers and four dimers. A *channel* is an ordered pair $r = (u, w)$ of words with $\ell(u) + \ell(w) \leq n$; it represents both directions of the reversible reaction

$$u + w \rightleftharpoons uw.$$

Equivalently a channel is a product word uw of length between 2 and n together with a split position; pairs with different split positions are different channels even when the product word is repetitive. Let $\mathcal{R}_n$ be the set of channels. Then

$$X_n := |\mathcal{X}_n| = 2^{n+1} - 2, \quad R_n := |\mathcal{R}_n| = \sum_{j=2}^n (j-1)2^j = (n-2)2^{n+1} + 4. \quad (2.1)$$

2.2 The capped-Zipf catalysis sampler

Fix an exponent $a > 1$. For each molecule $x \in \mathcal{X}_n$, independently, draw an integer $K_x \geq 1$ with

$$\mathbb{P}(K_x = k) = \frac{k^{-a}}{\zeta(a)} \quad (k \geq 1), \quad (2.2)$$

put $D_x = \min(K_x, R_n) - 1 \in \{0, \dots, R_n - 1\}$, and, conditional on $D_x = d$, choose the set $C_x \subseteq \mathcal{R}_n$ of channels catalyzed by x uniformly among the d -subsets of $\mathcal{R}_n$. Thus the degree law is exactly

$$\mu_{a,R}(d) = \begin{cases} (d+1)^{-a}/\zeta(a), & 0 \leq d < R-1, \\ \sum_{k \geq R} k^{-a}/\zeta(a), & d = R-1, \end{cases} \quad (2.3)$$

with $R = R_n$; the cap atom at $d = R_n - 1$ carries the whole Zipf tail and is never discarded. A configuration $C = (C_x)_{x \in \mathcal{X}_n}$ has weight

$$w_{a,n}(C) = \prod_{x \in \mathcal{X}_n} \frac{\mu_{a,R_n}(|C_x|)}{\binom{R_n}{|C_x|}}, \quad (2.4)$$

and every source probability below is a finite sum of this weight against an event indicator. Rows C_x of different molecules are independent; coordinates within a row are not (the row has a prescribed size, sampled without replacement). Catalyzing a channel catalyzes both of its directions. Food molecules may catalyze. We write

$$p_n(a) = \frac{\mathbb{E}D_x}{R_n} = \mathbb{P}(r \in C_x) \text{ for any fixed } x, r, \quad q_{0,n}(a) = \mathbb{P}(D_x = 0) = \zeta(a)^{-1}, \quad (2.5)$$

and call p_n the single-incidence probability. The asymptotic statements use the critical sequence

$$a_n = 2 - \frac{2}{n}, \quad (2.6)$$

and omit a_n from the notation; the finite bounds of Theorem 3.1(iii)–(iv) hold for every $a > 1$. Section 7 shows that at a_n the mean degree satisfies $\mathbb{E}D_x/n \rightarrow 9/\pi^2$, so this is the linear-intensity regime of [11, 14]. This is the capped power-law sampling rule of Hordijk and Steel [10] with the cap, the shift, the without-replacement row, and the reversible split-position convention fixed as above; we make no claim about samplers with other conventions.

2.3 RAFs and their minimum size

For $S \subseteq \mathcal{R}_n$, the reversible food closure $\text{cl}_F(S)$ is the smallest set of words containing F and closed under both directions of every channel in S : if $(u, w) \in S$ and $u, w \in \text{cl}_F(S)$ then $uw \in \text{cl}_F(S)$, and if $uw \in \text{cl}_F(S)$ then $u, w \in \text{cl}_F(S)$. Catalysts are not prerequisites of this closure. A nonempty $S \subseteq \mathcal{R}_n$ is a *RAF* (for the configuration C) if every word on either side of every channel of S lies in $\text{cl}_F(S)$ and

$$\forall r \in S \quad \exists x \in \text{cl}_F(S) : r \in C_x. \quad (2.7)$$

This is reflexive autocatalysis and food generation for reversible channels. It does not require that a catalyst be present before its reaction first fires in some generation order, and it does not substitute basal accessibility for catalysis: the catalyst must belong to the closure generated by S itself.

Let G_n be the event that some RAF exists, and on G_n let M_n be the minimum number of channels of a RAF. Put

$$P_n = \mathbb{P}(G_n), \quad P_n(m) = \mathbb{P}(\exists S \text{ RAF} : |S| \leq m). \quad (2.8)$$

2.4 The reactor

The reactor is a continuous-time Markov jump process on molecule counts $N = (N_x)_{x \in \mathcal{X}_n} \in \mathbb{N}^{\mathcal{X}_n}$, in the standard stochastic mass-action form [2, 13], with the following explicit ingredients.

Kinetic marks. Independently of the source and of each other, let S_r, B_r ($r \in \mathcal{R}_n$) and H_{rx} ($r \in \mathcal{R}_n, x \in \mathcal{X}_n$) be uniform on $\{1, \frac{3}{2}, 2\}$. Fix

$$\varepsilon = \frac{1}{500\,000\,000}, \quad b_r = \varepsilon S_r B_r, \quad \gamma_{rx} = 4S_r H_{rx}. \quad (2.9)$$

Every channel $r = (u, w)$ has both *basal* directions $u + w \rightarrow uw$ and $uw \rightarrow u + w$ with coefficient b_r . If $r \in C_x$, it also has both *catalytic* directions



with coefficient γ_{rx} . Hence $\varepsilon \leq b_r \leq 4\varepsilon$ and $4 \leq \gamma_{rx} \leq 16$; the coefficients do not depend on n or on the volume. The shared factor S_r makes the positive vector $q_x = 4^{-\ell(x)}$ balance every internal pair ($q_u q_w = q_{uw}$), also after multiplication by a catalyst factor; this internal detailed balance is a modelling choice that keeps the chemistry thermodynamically consistent and is not used as an equilibrium assertion for the fed reactor.

Propensities. For a channel direction with input multiplicity vector α , output vector β and coefficient k , the rate in count state N at integer volume $V > 0$ is

$$\lambda(N) = kV^{1-|\alpha|} \prod_x (N_x)_{\alpha_x}, \quad (m)_j = m(m-1) \cdots (m-j+1), \quad (2.10)$$

with no factorial divisor; a catalyst that coincides with a substrate or product keeps its full input multiplicity in the propensity and cancels only in the net stoichiometric change.

Feed, dilution, initial state. Each food species $f \in F$ enters at rate V ; every species x leaves at rate N_x (unit turnover). Initially $N_f(0) = V$ for $f \in F$ and $N_x(0) = 0$ otherwise. Write

$$L(N) = V^{-1} \sum_x \ell(x) N_x, \quad M(N) = V^{-1} \sum_{\ell(x) > 2} \ell(x) N_x \quad (2.11)$$

for the total and the nonfood polymer mass per volume. Internal reactions conserve polymer mass, so the generator $\mathcal{G}$ of the process satisfies

$$\mathcal{G}L = 10 - L, \quad L(N(0)) = 10. \quad (2.12)$$

At fixed n and V every mass sublevel set is finite and the total rate is bounded on it, so the process is nonexplosive; this holds whether catalysis is enabled or not.

Disabled reactor. We write $\mathbb{P}_0$ for the same construction with every catalytic coefficient γ_{rx} replaced by 0, keeping the basal coefficients, feed, dilution, initial state and window.

2.5 The productive observable and the probability laws

Let $W^V((s, t])$ be the nonfood polymer mass exported by dilution during $(s, t]$: each outflow of a molecule x contributes $\ell(x) \mathbf{1}_{\ell(x) > 2}$. Let $B^V(t)$ be the positive basal nonfood input divided by V (the sum, over basal firings up to time t whose nonfood mass increment is positive, of that increment), and let $J_{\text{cat}}^V(t)$, $J_{\text{basal}}^V(t)$ be the signed nonfood mass increments, divided by V , of the catalytic and basal channels respectively, including negative reverse contributions. Pathwise,

$$M(t) + \frac{W^V((0, t])}{V} = J_{\text{cat}}^V(t) + J_{\text{basal}}^V(t), \quad J_{\text{basal}}^V(t) \leq B^V(t), \quad (2.13)$$

because there is no initial nonfood stock.

Definition 2.1 (Productive output). For $n \geq 4$ and integer $V > 0$, the productive event is

$$\mathcal{F}_{n,V} = \left\{ \sup_{0 \leq t \leq 100} L(t) \leq 11 \text{ and } W^V((1, 100])/V > \frac{1}{10} \right\}. \quad (2.14)$$

The event refers only to total mass and to exported nonfood mass. It does not mention a RAF, a selected molecule, a catalyst, or the labels distinguishing basal from catalytic firings. The window $(1, 100]$ and the thresholds 11 and $1/10$ are fixed once and for all; nothing in the event scales with n or V .

Two auxiliary events are used in the proofs. Let $U = 00$, $W = 11$, $X = 0011$ and $r_* = (U, W)$, the self-catalyzed ligation $00 + 11 \rightleftharpoons 0011$. The *witness event* on the source is

$$E_n = \{C_f = \emptyset \text{ for all } f \in F, \quad r_* \in C_X\}, \quad (2.15)$$

which says that no food molecule catalyzes anything and that the product of r_* catalyzes r_* . Every other row of the source is unrestricted on E_n . On E_n the singleton $\{r_*\}$ is a RAF: its reactants are food and its catalyst X is generated by r_* itself. Hence $E_n \subseteq G_n$. The *residence event* on trajectories is, with $\rho = (3 \cdot 10^{18})^{-1}$,

$$\begin{aligned} \mathcal{S}_{n,V} = \Big\{ \sup_{t \leq 100} L(t) \leq 11, \quad N_X(t)/V \geq \rho \text{ for all } t \in [1, 100], \\ E := W^V((1, 100])/V > \frac{1}{10}, \quad B^V(100) < E/4 \Big\}. \end{aligned} \quad (2.16)$$

Clearly $\mathcal{S}_{n,V} \subseteq \mathcal{F}_{n,V}$, and by (2.13), $J_{\text{cat}}^V(100) > 3/40$ on $\mathcal{S}_{n,V}$: at least three quarters of the exported nonfood mass has been supplied by catalytic firings. This is a collective whole-history current and is not attributed to r_* alone.

Probability laws. Let $\omega = (C_x)_x$ be the catalytic environment with law $\mathbb{P}_n$ given by (2.4), let m be the kinetic marks with their product law, and let $K_{n,V,\omega,m}$ be the law of the reactor trajectory at fixed ω, m . The joint law is

$$\mathbb{Q}_{n,V}(d\omega, dm, d\xi) = \mathbb{P}_n(d\omega) \mathbb{P}(dm) K_{n,V,\omega,m}(d\xi).$$

Probabilities written without a subscript, such as $\mathbb{P}(\mathcal{F}_{n,V})$ or $\mathbb{P}(\mathcal{F}_{n,V} \cap G_n)$, are $\mathbb{Q}_{n,V}$ -probabilities: they average over the environment, the marks and the trajectory. For a fixed environment and marks we write

$$\pi_{n,V}(\omega, m) = K_{n,V,\omega,m}(\mathcal{F}_{n,V}), \quad \pi_{n,V}^0(\omega, m) = K_{n,V,\omega,m}^0(\mathcal{F}_{n,V}) \quad (2.17)$$

for the enabled and disabled success probabilities of a single reactor run, so that $\mathbb{P}(\mathcal{F}_{n,V}) = \mathbb{E} \pi_{n,V}$. Two kinds of rarity statement will be made: about the average $\mathbb{E} \pi_{n,V}$, and about the fraction of (ω, m) for which $\pi_{n,V}(\omega, m)$ exceeds a threshold. Neither is substituted for the other; both are proved.

All trajectory events used are measurable in the jump representation (the mass condition is a countable intersection of prefix conditions; a windowed reward is specified through the two jump indices straddling the window endpoints), and nonexplosion with almost surely positive holding times identifies them with their physical-time versions.

3 Main theorem

3.1 The infinite closure process and the structural limit

Consider the infinite catalogue of all reversible binary channels (all product words of length ≥ 2 , all split positions), each declared *open* independently with probability $q \in [0, 1]$. Let $\mathcal{C}(q)$ be the closure of F under both directions of the open channels, and put

$$\theta(q) = \mathbb{P}(\mathcal{C}(q) \text{ contains words of unbounded length}), \quad \Lambda = \frac{9}{\pi^2}, \quad q_* = 1 - e^{-\Lambda}. \quad (3.1)$$

Numerically $\Lambda \approx 0.9119$ and $q_* \approx 0.5982$; q_* is a channel openness, not a RAF probability. No closed form for $\theta(q_*)$ is asserted.

3.2 Constants

Fix

$$d = \frac{1}{6 \cdot 10^{20}}, \quad c = \frac{d^2}{4 \cdot 10^7} \approx 6.94 \cdot 10^{-50}, \quad k = 4 \cdot 10^8, \quad C_k = X_k R_{k+2}, \quad \delta_{n,V} = e^{-cV/n}. \quad (3.2)$$

The constant C_k is finite and independent of n and V ; it is astronomically large and there is no proposal to enumerate the catalog it counts. Set also

$$h_n = \lfloor n^{3/4} \rfloor, \quad L_n = \lceil 8\sqrt{n} \rceil, \quad A_n = 2^{L_n+1} + (n^3 2^{h_n} + 2^{L_n+1})n, \quad B_n = A_n^2. \quad (3.3)$$

Theorem 3.1 (Separation of structural and productive autocatalysis). *In the model of Section 2 the following hold.*

(i) (Structural existence.) *At $a = a_n$,*

$$P_n = \mathbb{P}(G_n) \longrightarrow \theta(q_*), \quad 0 < \theta(q_*) \leq 1 - e^{-36\Lambda} < 1. \quad (3.4)$$

(ii) (Structural size.) *At $a = a_n$, for every $b > 0$, $P_n(\lfloor 2^{bn} \rfloor) \rightarrow \theta(q_*)$; consequently*

$$\mathbb{P}(M_n \leq 2^{bn} \mid G_n) \longrightarrow 1 \quad \text{for every } b > 0, \quad (3.5)$$

that is, $\log M_n/n \rightarrow 0$ in probability under the conditional laws given G_n . Moreover $\log B_n/n \rightarrow 0$ and $\mathbb{P}(M_n \leq B_n \mid G_n) \rightarrow 1$.

(iii) (Productive rarity.) *For every $a > 1$, $n \geq 4$ and integer $V \geq 2n/d$,*

$$q_{0,n}(a)^6 p_n(a) (1 - 24\delta_{n,V}) \leq \mathbb{P}(\mathcal{F}_{n,V}) \leq C_k p_n(a) + 4\delta_{n,V}, \quad (3.6)$$

$$q_{0,n}(a)^6 p_n(a) (1 - 24\delta_{n,V}) \leq \mathbb{P}(\mathcal{F}_{n,V} \cap G_n) \leq C_k p_n(a) + 4\delta_{n,V}, \quad (3.7)$$

and, uniformly over every environment $\omega \in E_n$ and every mark realization m ,

$$K_{n,V,\omega,m}(\mathcal{S}_{n,V}) \geq 1 - 24\delta_{n,V}. \quad (3.8)$$

Consequently, at $a = a_n$, for every integer sequence $V(n) \geq 10^{60}(n+1)^2$,

$$\mathbb{P}(\mathcal{F}_{n,V(n)}) = \Theta(p_n), \quad \mathbb{P}(\mathcal{F}_{n,V(n)} \cap G_n) = \Theta(p_n), \quad \mathbb{P}(\mathcal{F}_{n,V(n)} \mid G_n) = \Theta(p_n), \quad (3.9)$$

where $p_n \sim 9/(\pi^2 X_n)$ and $\frac{1}{n} \log p_n \rightarrow -\log 2$; each of the three probabilities therefore has logarithmic rate $-\log 2$.

(iv) (Removal of catalysis.) *Under the finite hypotheses of (iii), for every environment and marks, $\pi_{n,V}^0(\omega, m) \leq 2\delta_{n,V}$; hence $\mathbb{P}_0(\mathcal{F}_{n,V}) \leq 2\delta_{n,V}$, and along the sequences of (iii),*

$$\frac{\mathbb{P}_0(\mathcal{F}_{n,V(n)})}{\mathbb{P}(\mathcal{F}_{n,V(n)})} \longrightarrow 0. \quad (3.10)$$

The stronger event $\mathcal{S}_{n,V}$ has disabled probability exactly zero.

The finite lower bounds in (3.6)–(3.7) may be negative for a volume that is too small; they are then valid but empty. The envelope $V(n) \geq 10^{60}(n+1)^2$ is a sufficient condition that makes the noise term $\delta_{n,V(n)}$ smaller than e^{-3n} , hence than $e^{-n}p_n$ eventually; Section 7.5 states the smaller sufficient volume for a prescribed reliability. It is not a necessary volume and not a calibrated reactor size.

3.3 The environment-level formulation

The bounds in (iii) are proved pointwise in the environment and then averaged, which gives a statement about viable environments rather than about average success.

Corollary 3.2 (Fraction of productive environments). *Under the finite hypotheses of Theorem 3.1(iii), with $\mathbb{P}_n \otimes \mathbb{P}$ the law of (ω, m) ,*

$$(\mathbb{P}_n \otimes \mathbb{P})(\pi_{n,V} \geq 1 - 24\delta_{n,V}) \geq (\mathbb{P}_n \otimes \mathbb{P})(E_n) = q_{0,n}(a)^6 p_n(a), \quad (3.11)$$

$$(\mathbb{P}_n \otimes \mathbb{P})(\pi_{n,V} > 4\delta_{n,V}) \leq C_k p_n(a), \quad (3.12)$$

and $\pi_{n,V}^0 \leq 2\delta_{n,V}$ for every (ω, m) . In particular, along $V(n) \geq 10^{60}(n+1)^2$ and for every fixed $\eta \in (0, 1)$, the fraction of environment-and-mark realizations in which a single reactor run succeeds with probability at least η is $\Theta(p_n)$, and so is the fraction of such realizations that also contain a RAF.

The corollary is a conventional consequence of the pointwise estimates of Sections 5–6; its proof is in Section 7.4. Table 1 summarizes the three quantities compared.

Quantity	What it certifies	Behaviour at a_n , $V(n) \geq 10^{60}(n+1)^2$
Structural existence $\mathbb{P}(G_n)$	A nonempty reaction set closed under its own resources and catalysts	$\rightarrow \theta(q_*) \in (0, 1)$
Structural size M_n given G_n	Number of channels in a smallest certificate	$\log M_n/n \rightarrow 0$; $M_n \leq B_n$ w.h.p.
Productive output $\mathbb{P}(\mathcal{F}_{n,V(n)})$	Bounded mass and export $> 1/10$ on $(1, 100]$ from food only, by any mechanism	$\Theta(p_n) = \Theta(1/X_n)$, rate $-\log 2$
conditional on G_n	Same output, in networks that contain a RAF	$\Theta(p_n)$
catalysis removed	Same output, basal chemistry only	$\leq 2e^{-cV/n} = o(p_n)$

Table 1: The three quantities compared by Theorem 3.1. The structural statements are about the sampled catalytic relation alone; the productive statements are about reactor trajectories started from food.

Remark 3.3 (Universality of the structural limit). Section 7 proves $\mathbb{E}D_x/n \rightarrow \Lambda = 9/\pi^2$ at a_n : the mean number of channels catalyzed per molecule is asymptotically linear with slope Λ . In the uniform model, where every molecule–channel incidence is present independently with probability f_n/R_n and $f_n/n \rightarrow \lambda$, the companion paper [1] proves that the RAF probability converges to $\theta(1 - e^{-\lambda})$ with the same survival function θ . Part (i) therefore says that the heavy-tailed sampler has the same limiting existence probability as the uniform sampler of equal intensity $\lambda = \Lambda$. The mechanism is Proposition 4.2: on any fixed finite catalogue, aggregation over X_n independent rows washes out the heavy tail. The paper of Hordijk and Steel [10] compares the samplers at finite n by simulation; no claim is made here about finite- n agreement or about other exponents.

Remark 3.4 (What is not claimed). Parts (i)–(iii) do not imply that productive environments typically contain small RAFs. Part (ii) gives $\mathbb{P}(G_n \cap \{M_n > 2^{bn}\}) = o(1)$, an error that need not be $o(p_n)$, so it says nothing about the $\Theta(p_n)$ -probability set of productive environments. The witness event E_n does contain the one-channel RAF $\{r_*\}$, so *some* productive environments of probability $\Theta(p_n)$ have $M_n = 1$; whether this holds for most productive environments is not addressed.




Figure 1: Proof map. The two halves share one source law; the conditional statement divides the joint lower bound, which comes from a witness event that is itself a RAF, by the nonvanishing structural probability.

4 Structural existence and size

This section proves parts (i) and (ii). Throughout, $a = a_n$. The proof has three components: a fixed-catalogue independent-channel limit, a constructive lower bound with a subexponential budget, and an upper bound through first escape from a fixed length. The finite identities, the record-word and nucleus arguments, the degree-band coupling and the target estimate are given in Appendix C.

4.1 A fixed-catalogue independent-channel limit

Call a channel *open* in the source if at least one molecule catalyzes it, and write $O_n = \bigcup_x C_x$. Forgetting which molecule catalyzes a channel is only used for the upper comparison and for local limits.

Lemma 4.1 (Critical degree moments). *For a single degree $D_n = D_x$ at $a = a_n$,*

$$\frac{\mathbb{E}D_n}{n} \rightarrow \Lambda, \quad \frac{X_n \mathbb{E}D_n}{R_n} \rightarrow \Lambda, \quad \frac{X_n \mathbb{E}D_n^2}{R_n^2} \rightarrow 0. \quad (4.1)$$

Proof. By (2.3), $\zeta(a_n)(\mathbb{E}D_n + 1) = \sum_{k=1}^{R_n-1} k^{1-a_n} + R_n \sum_{k \geq R_n} k^{-a_n}$. The cap term is $O(1)$ because $R_n \sum_{k \geq R_n} k^{-2+2/n} \leq R_n \cdot \frac{(R_n-1)^{-1+2/n}}{1-2/n} = O(R_n^{2/n})$ and $R_n^{2/n} \rightarrow 4$ since $\log R_n/n \rightarrow \log 2$. Integral comparison for the decreasing summands $k^{-1+2/n}$ gives

$$\frac{1}{n} \sum_{k=1}^{R_n-1} k^{-1+2/n} = \frac{1}{n} \left(\frac{n}{2} (R_n^{2/n} - 1) + O(1) \right) \rightarrow \frac{3}{2},$$

and $\zeta(a_n) \rightarrow \zeta(2) = \pi^2/6$, whence $\mathbb{E}D_n/n \rightarrow \frac{3}{2\zeta(2)} = \Lambda$. The second limit follows from $R_n/(nX_n) \rightarrow 1$, which is read off (2.1). For the second moment, integral bounds on the finite part and on the cap atom give $\mathbb{E}D_n^2 \leq 4R_n^{1+2/n}/\zeta(a_n)$ eventually; multiply by X_n/R_n^2 and use $X_n/R_n \rightarrow 0$. $\square$

Proposition 4.2 (Fixed-catalogue convergence). *Fix $N \geq 2$. The restriction of O_n to the literal copy of $\mathcal{R}_N$ inside $\mathcal{R}_n$ converges in distribution, as $n \rightarrow \infty$, to independent Bernoulli(q_*) coordinates. Consequently the probability of every event determined by $O_n \cap \mathcal{R}_N$ converges to its independent-channel probability.*

Proof. For a fixed set of s distinct channels, let $p_{n,s}$ be the probability that one row hits the set. Uniform subset sampling and the first two Bonferroni inequalities give

$$s \frac{\mathbb{E}D_n}{R_n} - \binom{s}{2} \frac{\mathbb{E}[D_n(D_n - 1)]}{R_n(R_n - 1)} \leq p_{n,s} \leq s \frac{\mathbb{E}D_n}{R_n}. \quad (4.2)$$

By Lemma 4.1, $X_n p_{n,s} \rightarrow s\Lambda$ and $p_{n,s} \rightarrow 0$. Independence of rows gives

$$\mathbb{P}(O_n \text{ misses all } s \text{ channels}) = (1 - p_{n,s})^{X_n} \rightarrow e^{-s\Lambda}. \quad (4.3)$$

For disjoint prescribed sets S (to be open) and C (to be closed), inclusion–exclusion,

$$\mathbb{P}(S \subseteq O_n, C \cap O_n = \emptyset) = \sum_{T \subseteq S} (-1)^{|T|} \mathbb{P}((T \cup C) \cap O_n = \emptyset) \rightarrow e^{-|C|\Lambda} (1 - e^{-\Lambda})^{|S|},$$

requires no independence within a row (Appendix C.1). Summing over the finitely many atoms of $\mathcal{R}_N$ proves the claim. $\square$

The second moment only excludes appreciable repeat hits by one molecule on a fixed catalogue. It does not make a whole row independent and gives no uniform approximation on the growing catalogue $\mathcal{R}_n$; the coupling used in the construction is a separate argument.

4.2 Survival, nuclei, and the constructive lower bound

Two facts about the infinite closure $\mathcal{C}(q)$ are used; both are proved in Appendix C.2. For $q > 0$, survival implies generation of every fixed word almost surely (a record-word argument). For $q > 1/2$, $\theta(q) > 0$, θ is continuous on $(1/2, 1]$, and for every m a nucleus containing all words of length at most m extends to survival except with probability at most

$$\tau_m(q) = \sum_{j>m} 2^j (1-q)^{j-1} = \frac{2[2(1-q)]^m}{1-2(1-q)} \rightarrow 0 \quad (m \rightarrow \infty), \quad (4.4)$$

uniformly for q in compact subsets of $(1/2, 1]$. Thus for every $c < \theta(q)$ there is a finite *seed cap* N such that the open channels of the fixed catalogue $\mathcal{R}_N$ generate the entire m -nucleus with probability greater than c .

Two degree bands. Besides (3.3), set

$$k_n = \lfloor n^{7/8} \rfloor, \quad d_0 = 2^{k_n}, \quad d_1 = 2^{n-h_n}, \quad d_2 = 2^{n-\lfloor h_n/20 \rfloor}, \quad \varepsilon_n = (n+1)^{-1/8}. \quad (4.5)$$

Condition on the complete degree vector $(D_x)_x$. For $d_0 \leq D_x < d_1$ (the *low band*) attach an auxiliary independent Bernoulli row in which each channel is present with probability $(1 - \varepsilon_n)D_x/R_n$; for $d_1 \leq D_x < d_2$ (the *high band*) attach a row with per-channel probability $\frac{19}{20}D_x/R_n$; all other auxiliary rows are empty. Given the degrees, the unions of the low and of the high auxiliary rows are independent-channel fields, independent of each other. Appendix C.3 proves the following. The auxiliary rows can be coupled into the exact uniform rows C_x (each auxiliary row is contained in C_x) except on an event of probability at most

$$X_n \exp(-\varepsilon_n^2 d_0/4) \leq e^{-n} \quad \text{eventually}; \quad (4.6)$$

since the bounded-RAF event is increasing in every row, lower bounds transfer with this error. The low union has openness at least $1 - e^{-t}$, for every fixed $t < \Lambda$, with probability tending to one; the high union has openness at least

$$\pi_n = \frac{29}{20}n^{-1/4} \quad (4.7)$$

with probability tending to one. The low band preserves the limiting ignition mass Λ ; the high band is reserved for generating catalyst labels.

Selecting and generating the owners. Fix $q \in (1/2, q_*)$ and $t < \Lambda$ with $q \leq 1 - e^{-t}$, then m , then a seed cap N as above. In the low field retain the whole finite witness inside $\mathcal{R}_N$ that generates the m -nucleus, and then one open split for each word of length between $m + 1$ and L_n . The retained support S_0 and its list of selected low catalyst labels (one low-band molecule per retained channel) each have size at most $R_N + 2^{L_n+1}$. The selector is a function of the degrees and the low rows only. Let W_H be the set of molecules with $D_x \geq d_1$. Outside an event of probability $o(1)$ (Appendix C.3),

$$|W_H| \leq n^3 2^{h_n}, \quad \ell(w) \geq n - 2h_n \ (w \in W_H), \quad \ell(w) \geq k_n - h_n \ (w \text{ a selected low label}). \quad (4.8)$$

The last bound excludes short molecules with low-band degree before any label is selected; it is not a claim that an adaptively selected label is a fresh uniform word.

The target estimate. For an independent-channel field with openness p , a word w of length ℓ can be generated from its substrings of length at most L using at most $\ell - 1$ additional channels, except with probability at most

$$e^{-p\ell/2} + 2\ell^2 e^{-pL^2/32}, \quad L \geq 4, \ \ell \geq L, \ pL \geq 32 \log 2 \quad (4.9)$$

(Appendix C.4). The estimate is uniform in w , including words with repeated substrings, and remains valid when the base support generating the substrings depends on the field, because failure is contained in a raw substring-exploration failure event. For the selected low labels it is applied after conditioning on the degrees and the low rows, which fixes the targets; the high field keeps its independent conditional law.

Apply (4.9) with $L = L_n$ and $p \geq \pi_n$ (so $pL \geq \frac{29}{20} \cdot 8n^{1/4} \geq 32 \log 2$ for $n \geq 1$) to all labels in W_H and to all selected low labels, with the union bounds from (4.8). For high labels the principal exponent is

$$h_n \log 2 - \frac{29}{40}n^{-1/4}(n - 2h_n) = (\log 2 - \frac{29}{40} + o(1))n^{3/4} \rightarrow -\infty,$$

since $\log 2 < 29/40$. For low labels, $\log(R_N + 2^{L_n+1}) = O(\sqrt{n})$ whereas $\pi_n(k_n - h_n)/2$ has order $n^{5/8}$. The second term of (4.9) has exponent $-\pi_n L_n^2/32 \leq -\frac{29}{10}n^{3/4}$, sufficient for both unions. Hence, with probability tending to one, every label in W_H and every selected low label has a bounded derivation from the retained support.

Assembly. Retain these derivations. Every channel of S_0 has its selected low catalyst generated; every added channel came from the high field and has a catalyst in W_H , also generated. All channels of the union have internally generated catalysts. Keep only the channels enabled at some stage of the food closure of the union; Appendix C.5 shows that this trimming preserves the closure and the catalysis witnesses and yields a RAF, nonempty because a word of length three is generated. Its size is at most

$$b_N(n) = (R_N + 2^{L_n+1}) + [n^3 2^{h_n} + R_N + 2^{L_n+1}]n = A_n + R_N(n + 1). \quad (4.10)$$

Proposition 4.3 (Constructive lower interface). *For every $r < \theta(q_*)$ there is a fixed N such that, eventually, $r < P_n(b_N(n))$, and $b_N(n) < 2^{bn}$ eventually for every fixed $b > 0$.*

Proof. For fixed q, t as above, choose m and N so that the nucleus mass, after the extension loss (4.4), exceeds a prescribed $r < \theta(q)$. Average over the true degree law, discarding the vanishing low-openness exception. The high-label admissibility error (4.8), the short-low-label error, the two target failure errors and the coupling overflow (4.6) all tend to zero; Appendix C.5 lists the five errors and the indicator inequality that adds them. The deterministic assembly then gives $P_n(b_N(n)) > r$ eventually. Since $\Lambda > \log 2$ we may take $t \uparrow \Lambda$ and $q \uparrow q_*$ within $(1/2, 1)$; continuity of θ gives every $r < \theta(q_*)$. Finally $h_n, L_n = o(n)$ and N is fixed, so $b_N(n)$ is subexponential. $\square$

4.3 The upper comparison: first length escape

Let $E_{n,K}$ be the event that $\text{cl}_F(O_n)$ contains a word longer than K , with $K \geq 2$ fixed before $n \rightarrow \infty$.

Lemma 4.4 (Finite escape interface). *For $n \geq 2K$, $E_{n,K}$ depends only on $O_n \cap \mathcal{R}_{2K}$.*

Proof. Consider the first stage at which the closure contains a word of length greater than K . It cannot be a cleavage, whose products are shorter than its input. It is a ligation of two words of length at most K , so its product has length at most $2K$, and every earlier productive step lies inside the same cap. Conversely, a derivation in $\mathcal{R}_{2K}$ is a derivation in $\mathcal{R}_n$. $\square$

Lemma 4.5 (RAF-to-escape comparison). *For $n \geq \max(4, 2K)$,*

$$P_n \leq \mathbb{P}(E_{n,K}) + 68X_K \frac{\mathbb{E}D_n}{R_n}. \quad (4.11)$$

Proof. If a RAF S exists, its food-generation derivation contains a channel enabled from food alone, or else every channel of S is already food-enabled; either way S contains a channel both of whose sides lie in $F \cup \{\text{products of length} \leq 4\}$. There are at most $R_4 = 68$ channels with product length at most four. If $E_{n,K}$ fails, then, since $\text{cl}_F(S) \subseteq \text{cl}_F(O_n)$, the catalyst of that channel has length at most K . For any specified catalyst and channel the incidence probability is exactly $\mathbb{E}D_n/R_n$; a union bound over at most $68X_K$ pairs gives (4.11). $\square$

By Proposition 4.2 and Lemma 4.4, $\mathbb{P}(E_{n,K}) \rightarrow \mathbb{P}(\mathcal{C}(q_*) \text{ escapes } K)$. The error in (4.11) vanishes at fixed K by Lemma 4.1. As $K \rightarrow \infty$ the independent-field escape events decrease to the survival event, so

$$\limsup_{n \rightarrow \infty} P_n \leq \theta(q_*). \quad (4.12)$$

The order of limits is: fixed boundary K , then $n \rightarrow \infty$, then $K \rightarrow \infty$. No growing-catalogue independent approximation is used.

4.4 Proof of parts (i) and (ii)

For every $b > 0$, Proposition 4.3 and integer rounding give $\liminf_n P_n(\lfloor 2^{bn} \rfloor) \geq \theta(q_*)$. Since $P_n(\lfloor 2^{bn} \rfloor) \leq P_n$, (4.12) gives both $P_n \rightarrow \theta(q_*)$ and $P_n(\lfloor 2^{bn} \rfloor) \rightarrow \theta(q_*)$. Positivity of $\theta(q_*)$ follows from $q_* > 1/2$ and the nucleus argument of Appendix C.2. For the upper bound in (3.4): a food-executable ligation has two food reactants, and there are exactly 36 such channels (four products of length two, sixteen of length three, and sixteen of length four with the split at position two); if all of them are closed, $\mathcal{C}(q_*) = F$, so $\theta(q_*) \leq 1 - (1 - q_*)^{36} = 1 - e^{-36\Lambda}$.

For (3.5), divide $P_n(\lfloor 2^{bn} \rfloor) \rightarrow \theta(q_*)$ by $P_n \rightarrow \theta(q_*) > 0$. For integer M_n , $M_n > 2^{bn}$ is equivalent to $M_n > \lfloor 2^{bn} \rfloor$, and on G_n these events are the positive tails of $\log_2 M_n/n$. For the explicit cutoff: $A_n \geq n + 1$ and $b_N(n) \leq (R_N + 1)A_n$; for fixed N , eventually $R_N \leq n$, so $b_N(n) \leq A_n^2 = B_n$, and Proposition 4.3 applied to every $r < \theta(q_*)$ gives $\liminf P_n(B_n) \geq \theta(q_*)$,

hence $P_n(B_n) \rightarrow \theta(q_*)$ and $\mathbb{P}(M_n \leq B_n \mid G_n) \rightarrow 1$. Direct estimation gives $A_n \leq 3n^4 2^{h_n + L_n + 1}$, so $\log_2 B_n \leq 2(h_n + L_n + 1) + 8 \log_2 n + 2 \log_2 3$ and $\log B_n/n \rightarrow 0$. $\square$

The explicit budget sharpens the zero-rate assertion but does not identify the actual order of M_n ; polynomial and much smaller subexponential scales remain possible. The statements are about probabilities, not expectations.

5 The productive lower bound

This section proves (3.8) and the lower bounds in (3.6)–(3.7). Fix $n \geq 4$, $V \geq 2n/d$, an environment $\omega \in E_n$ and marks m . The environment may contain arbitrarily many additional nonfood catalysts; the argument never removes them. All statements until Section 5.3 concern drifts, that is, the compensators of the count process, expressed in concentrations $x_z = N_z/V$.

5.1 Occupied-mass estimates

Let $Q = \sum_z x_z \leq L$. The basal forward monomials summed over ordered substrate pairs give total rate at most $4\varepsilon V Q^2$; the basal reverse channels count at most $\ell(z) - 1$ splits of each occupied word and give at most $4\varepsilon V L$. Each catalytic term carries one more occupied factor and coefficient at most 16: catalytic ligations contribute at most $16V Q^3$ and catalytic cleavages at most $16V Q L$. Falling factorials only lower these envelopes. Including feed ($6V$) and outflow (QV), on $\{L \leq 11\}$ the total jump rate satisfies

$$\lambda_{\text{total}} \leq 24000 V. \quad (5.1)$$

This bound counts occupied molecules, not possible reactions, and is valid for every source configuration, including coincidences and dense hosts.

On E_n every catalyst is a nonfood molecule, so the total catalyst concentration is at most $M/3$, and the speed of any channel, meaning the sum of its basal coefficient and of its catalytic coefficients weighted by catalyst concentrations, is at most

$$C(M) = 4\varepsilon + \frac{16}{3}M. \quad (5.2)$$

Write $u = x_U$, $w = x_W$, $x = x_X$. The coordinate drifts b_u, b_w, b_x obey

$$b_u \geq 1 - u - 23C(M)u, \quad b_w \geq 1 - w - 23C(M)w, \quad (5.3)$$

$$b_x \geq \varepsilon u w + x(4u w - 1 - 25C(M)). \quad (5.4)$$

Indeed, a word of length ℓ is consumed by ligations in two ordered positions, at total rate at most $2QC(M) \leq 22C(M)$ times its concentration, and by cleavages at rate at most $(\ell - 1)C(M)$ times its concentration; with $L \leq 11$ this gives the coefficients 23 for the dimers and 25 for X , together with the unit dilution. The gain terms retained are the feed for u, w , the basal forward firing of r_* (coefficient $\geq \varepsilon$) and the catalytic self-firing of r_* with catalyst X (coefficient ≥ 4), whose reactants are distinct so its propensity is exactly proportional to uwx . All other gains are discarded. No autonomous three-coordinate Markov model is claimed.

5.2 Corrected growth

Let $h(s) = (1 - s)_+^2$, $S = h(u) + h(w)$, and $\Phi = \log x - 4S$. For $u, w \geq 0$, with $\alpha = (1 - u)_+$ and $\beta = (1 - w)_+$,

$$uw \geq 1 - \alpha - \beta \geq \frac{3}{4} - 2(\alpha^2 + \beta^2) = \frac{3}{4} - 2S,$$

using $\alpha \leq 2\alpha^2 + \frac{1}{8}$. From (5.3) and $s(1-s)_+ \leq \frac{1}{4}$, $\frac{d}{dt}h(u) = -2(1-u)_+b_u \leq -2h(u) + \frac{46}{4}C(M)$, so $\dot{S} \leq -2S + 23C(M)$. Discarding the positive basal term in b_x/x ,

$$\dot{\Phi} \geq 4uw - 1 - 25C(M) + 8S - 92C(M) \geq 2 - 117C(M) = 2 - 468\varepsilon - 624M. \quad (5.5)$$

So if the collective nonfood mass stays small, the corrected logarithm grows at a fixed positive rate, while the mass bound limits its total growth. Food depletion cannot evade the argument by suppressing X while creating other nonfood material: that material is exactly what M measures.

Food-only startup. Since $C(M) < 59$ on $\{L \leq 11\}$,

$$b_u \geq 1 - 1358u, \quad b_w \geq 1 - 1358w, \quad b_x \geq \varepsilon uw - 1476x. \quad (5.6)$$

For the deterministic mass-action trajectory this gives $u, w \geq 1/1358$ for all $t \geq 0$ and

$$x(t) \geq \frac{\varepsilon}{1358^2 \cdot 1476} (1 - e^{-1476t}) \geq \rho \quad (t \geq 1).$$

This is startup at a fixed concentration scale from positive basal chemistry; no seed molecule is assumed.

5.3 Transfer to stochastic trajectories

Stop the process at the first exit of L above 11 or at time $T_* = 101$, whichever is first. For a coordinate or reward f write its increment as compensator plus centred noise A_f . The following simultaneous allowances are imposed:

quantity	$\sup_{t < T_*} A_f $ allowance
M	1/8
each of the six food coordinates (for mass control)	1/80
u, w (for startup and corrected growth)	10^{-5}
x	$\rho/100$
nonfood export W^V/V	1/40
positive basal input B^V	1/100

(5.7)

Appendix B proves from the marked jump kernel that, for $V \geq 2n/d$, each of these twelve two-sided bounds fails with probability at most $2\delta_{n,V}$, uniformly in the environment, the marks and the state history; their union costs at most $24\delta_{n,V}$. On the good event the following holds.

Mass control. The noise of L is A_M plus the six food noises weighted by length, and the food lengths sum to 10; so $|A_L| \leq \frac{1}{8} + \frac{10}{80} = \frac{1}{4}$. For a coordinate $y(t) = y(0) + \int_0^t b ds + A(t)$ with $|A| \leq \eta$ and drift $b \geq B - \alpha y$ ($\alpha > 0$), comparison of the continuous compensated coordinate $y - A$ gives

$$y(t) \geq \frac{B}{\alpha} + \left(y(0) - \frac{B}{\alpha}\right)e^{-\alpha t} - 2\eta, \quad (5.8)$$

and symmetrically from above; the loss 2η is paid once, not per jump. Applied to $GL = 10 - L$ with $L(0) = 10$, it gives $L \leq 10 + 2 \cdot \frac{1}{4} = 10.5 < 11$ up to the putative first exit, which therefore does not happen before T_* : on the good event, mass stays below 11 through time 101.

Startup. Applying (5.8) to (5.6) gives $u, w \geq \frac{1}{1358} - 2 \cdot 10^{-5}$ on $[0, T_*]$, and then, with $B = \varepsilon(\frac{1}{1358} - 2 \cdot 10^{-5})^2$ and $\alpha = 1476$,

$$x(t) \geq \frac{\varepsilon(\frac{1}{1358} - 2 \cdot 10^{-5})^2}{1477} - \frac{2\rho}{100} > \rho \quad (1 \leq t \leq 100),$$

since the left side is about $6.9 \cdot 10^{-19}$ and $\rho \approx 3.3 \cdot 10^{-19}$ (the factor 1477 absorbs the transient e^{-1476t} for $t \geq 1$).

Corrected growth on the compensated coordinates. Put $\tilde{u} = u - A_u$, $\tilde{w} = w - A_w$, $\tilde{x} = x - A_x$. These are continuous and absolutely continuous on bounded intervals, with derivatives equal to the coordinate drifts. Since $x \geq \rho$ and $|A_x| \leq \rho/100$, we have $|A_x| \leq x/100$. Appendix A proves the shifted inequality

$$\frac{d}{dt} [\log \tilde{x} - 4h(\tilde{u}) - 4h(\tilde{w})] \geq 1.9 - 120C(M) = 1.9 - 480\varepsilon - 640M \quad (5.9)$$

and the endpoint bounds $\tilde{x} \in [\rho/2, 3]$, $\tilde{u}, \tilde{w} \geq -10^{-5}$, so that the bracket increases by less than 54 over $[1, 100]$. Integrating,

$$I := \int_1^{100} M(t) dt \geq \frac{99(1.9 - 480\varepsilon) - 54}{640} \geq 0.2095311015. \quad (5.10)$$

Export and basal budget. With unit turnover, $\int_1^{100} M dt$ is exactly the compensator of $W^V((1, 100])/V$ on the same window, and the two endpoint noise errors cost at most $2 \cdot \frac{1}{40} = \frac{1}{20}$; hence $E = W^V((1, 100])/V \geq I - \frac{1}{20} > \frac{1}{10}$. Each positive basal nonfood increment is at most $4/V$, and summing occupied ordered substrate pairs gives the uniform drift bound

$$\frac{d}{dt} \langle B^V \rangle \leq 16\varepsilon L^2 \leq 1936\varepsilon, \quad (5.11)$$

so on the good event $B^V(100) \leq 1936\varepsilon \cdot 100 + \frac{1}{100} = 0.0103872 < E/4$. Thus every condition of $\mathcal{S}_{n,V}$ holds on the good event, which proves (3.8).

5.4 Averaging over the witness event

Row independence gives exactly

$$\mathbb{P}_n(E_n) = q_{0,n}(a)^6 p_n(a) : \quad (5.12)$$

the six food rows are empty with probability $q_{0,n}(a)$ each, and the row of X contains r_* with probability $p_n(a)$; X is not a food molecule, and the rest of its row is unrestricted. Since $\mathcal{S}_{n,V} \subseteq \mathcal{F}_{n,V}$ and $E_n \subseteq G_n$, averaging (3.8) over E_n and all marks gives the lower bounds in (3.6) and (3.7). $\square$

6 The productive upper bound and the disabled reactor

The upper bound must exclude output by every mechanism the model allows, not only by the mechanism of Section 5. It rests on one structural event on the source and one capacity estimate that holds pointwise in the state.

6.1 Short incidences

Let H_k be the event that some molecule of length at most k catalyzes some channel with product length at most $k+2$. There are at most $X_k R_{k+2}$ such molecule-channel pairs in the catalogue, each present with probability exactly $p_n(a)$, so

$$\mathbb{P}_n(H_k) \leq C_k p_n(a), \quad C_k = X_k R_{k+2}. \quad (6.1)$$

No independence between these incidences is needed.

6.2 Catalytic capacity outside H_k

Let $J_+^V(t)$ be the positive catalytic nonfood input divided by V : the sum over catalytic firings up to t of the positive part of their nonfood mass increment. A firing has positive nonfood increment only if it is a ligation $u + w \rightarrow uw$ with at least one food reactant; its increment is then the total length of its food reactants, at most 4. Reverse cleavages never have positive nonfood increment.

Lemma 6.1 (Capacity bound). *Outside H_k , in every state with $L \leq 11$, the drift of J_+^V satisfies*

$$k \frac{d}{dt} \langle J_+^V \rangle \leq 4\kappa L^3, \quad \kappa = 16. \quad (6.2)$$

Proof. Write $\phi(z) = \ell(z)\mathbf{1}_{z \in F}$ for the food-mass weight and $\Phi_F = \sum_z \phi(z)x_z \leq L$. The drift of J_+^V is at most

$$\kappa \sum_z \sum_{(u,w) \in C_z} (\phi(u) + \phi(w))x_z x_u x_w,$$

using $\gamma \leq \kappa$ and the monomial envelope of the falling factorials. Outside H_k , for every incidence $(z, (u, w))$ in this sum either $\ell(z) > k$, or $\ell(u) + \ell(w) > k + 2$; in the second case the two reactants cannot both be food (their product would have length at most $4 \leq k + 2$), so exactly one is food and the other has length greater than k . Let $\Lambda_k = \sum_{\ell(z) > k} x_z \leq L/k$. The terms with $\ell(z) > k$ contribute at most $\Lambda_k \sum_{u,w} (\phi(u) + \phi(w))x_u x_w \leq 2\Lambda_k \Phi_F Q$. The remaining terms have a food reactant carrying the weight and a long nonfood reactant, in either order, and contribute at most $Q \cdot 2\Phi_F \Lambda_k$. Altogether the drift is at most $4\kappa \Lambda_k \Phi_F Q \leq 4\kappa L^3/k$. $\square$

For $k = 4 \cdot 10^8$ and $L \leq 11$, (6.2) gives $\frac{d}{dt} \langle J_+^V \rangle \leq 85184/k < 1/4000$. On the noise event for J_+^V with allowance $1/40$, therefore, $J_+^V(100) \leq \frac{100}{4000} + \frac{1}{40} = 0.05$. The bound is pointwise in the state, so it holds after the populations have adapted to the environment in any way whatsoever.

6.3 Completing the upper bound

On $\mathcal{F}_{n,V}$ the mass is at most 11 through time 100, so the stopped process of Section 5.3 has not been stopped before 100 and no separate mass-exit error is needed. Suppose $\omega \notin H_k$ and the two reward noises (positive basal input, allowance $1/100$; positive catalytic input, allowance $1/40$) stay within their allowances. Then $B^V(100) \leq 0.0103872 < 1/40$ by (5.11), and $J_+^V(100) \leq 0.05$. By (2.13), the total export through time 100 satisfies $W^V((0, 100])/V \leq B^V(100) + J_+^V(100)$, because $M \geq 0$ and each signed input is at most its positive part. Hence

$$\frac{W^V((1, 100])}{V} \leq \frac{W^V((0, 100])}{V} \leq 0.0103872 + 0.05 < \frac{1}{10},$$

contradicting $\mathcal{F}_{n,V}$. Each noise failure has probability at most $2\delta_{n,V}$ uniformly in ω and m (Appendix B). Therefore

$$\pi_{n,V}(\omega, m) \leq 4\delta_{n,V} \quad \text{for every } \omega \notin H_k \text{ and every } m, \quad (6.3)$$

and averaging with (6.1) gives $\mathbb{P}(\mathcal{F}_{n,V}) \leq C_k p_n(a) + 4\delta_{n,V}$. Intersecting with G_n can only decrease the probability, which gives the upper bound in (3.7). $\square$

The signed identity (2.13) is what makes the role of reverse reactions explicit: the upper bound uses only positive inputs, but the identity that justifies it includes the negative contributions of reverse firings. On $\mathcal{F}_{n,V} \cap \{B^V(100) < 1/40\}$ the signed catalytic input itself exceeds $3/40$.

6.4 The disabled reactor

With every catalytic coefficient zero, export is bounded by the positive basal input alone: $W^V((0, 100])/V \leq B^V(100)$. The basal noise event, of failure probability at most $2\delta_{n,V}$, gives $B^V(100) \leq 0.0103872 < 1/10$, contradicting $\mathcal{F}_{n,V}$. Hence $\pi_{n,V}^0(\omega, m) \leq 2\delta_{n,V}$ for every ω, m , and $\mathbb{P}_0(\mathcal{F}_{n,V}) \leq 2\delta_{n,V}$. On $\mathcal{S}_{n,V}$ the condition $B^V(100) < E/4$ contradicts (2.13) exactly when $J_{\text{cat}}^V \equiv 0$, so the disabled probability of $\mathcal{S}_{n,V}$ is zero. The comparison uses the same observable, the same feed, initial state, basal rates and window on both sides; it does not define enabled success through a catalytic label, and it does not say that the disabled reactor exports nothing. The ratio limit (3.10) follows from the lower bound in (3.6) once the asymptotics of Section 7 are in place.

7 Source asymptotics, the Θ law, and the conditional conclusion

7.1 Exact source asymptotics

At $a = a_n$, Lemma 4.1 and the identity $R_n/X_n = n - 2 + 2n/X_n$ from (2.1) give

$$\frac{\mathbb{E}D_n}{n} \rightarrow \frac{9}{\pi^2}, \quad X_n p_n \rightarrow \frac{9}{\pi^2}, \quad \frac{1}{n} \log p_n \rightarrow -\log 2, \quad q_{0,n} \rightarrow \frac{6}{\pi^2}. \quad (7.1)$$

Figure 2 shows the exact finite- n values of the first two normalizations, computed from (2.3) with the cap atom retained, using Hurwitz zeta evaluation at 100-digit precision rather than catalog enumeration; the approach to $9/\pi^2$ is slow, as the $O(1/n)$ -relative corrections in Lemma 4.1 suggest.




Figure 2: Exact evaluation of the capped source law at $a_n = 2 - 2/n$: $X_n p_n$ (left) and $\mathbb{E}D_n/n$ (right) against n , with the limit $9/\pi^2$. These are source-law values, not kinetic success probabilities.

7.2 Proof of the Θ law and the ratio limit

Let $V(n) \geq 10^{60}(n+1)^2$. Then $V(n) \geq 2n/d$ and $cV(n)/n \geq 3n$ for all $n \geq 1$ (with $c = d^2/(4 \cdot 10^7)$, $10^{60}c \geq 6 \cdot 10^{10}$), so $\delta_{n,V(n)} \leq e^{-3n}$. By (7.1), $p_n \geq e^{-2n}$ eventually, hence $\delta_{n,V(n)}/p_n \leq e^{-n} \rightarrow 0$. Since $q_{0,n} \rightarrow 6/\pi^2$, the sandwiches (3.6)–(3.7) give, eventually,

$$\frac{1}{2} \left(\frac{6}{\pi^2} \right)^6 p_n \leq \mathbb{P}(\mathcal{F}_{n,V(n)} \cap G_n) \leq \mathbb{P}(\mathcal{F}_{n,V(n)}) \leq (C_k + 1)p_n,$$

which is the first two statements of (3.9), and $\frac{1}{n} \log \mathbb{P}(\mathcal{F}_{n,V(n)}) \rightarrow -\log 2$. Dividing $\mathbb{P}_0(\mathcal{F}_{n,V(n)}) \leq 2\delta_{n,V(n)}$ by the lower bound proves (3.10). $\square$

7.3 The conditional conclusion

Corollary 7.1. *At $a = a_n$ and along $V(n) \geq 10^{60}(n+1)^2$,*

$$\mathbb{P}(\mathcal{F}_{n,V(n)} \mid G_n) = \frac{\mathbb{P}(\mathcal{F}_{n,V(n)} \cap G_n)}{P_n} = \Theta(p_n) \longrightarrow 0.$$

Proof. The numerator is the joint probability, bounded on both sides by (3.7); its lower bound comes from the witness event E_n , which is contained in G_n because $\{r_*\}$ is a RAF on E_n , and not from any independence between output and RAF existence. The denominator P_n is at most 1 and tends to $\theta(q_*) > 0$ by Theorem 3.1(i), so it lies eventually in $[\theta(q_*)/2, 1]$. Division gives the claim. $\square$

Thus, in the same sampled network, structural autocatalysis is present with probability bounded away from zero while the specified productive operation is present with probability of order $1/X_n$, and conditioning on the former does not change the order of the latter. This is tied to the fixed feed, bounded coefficients, observable and window; it does not declare any RAF inactive.

7.4 Proof of Corollary 3.2

The lower bound (3.11) restates (3.8) and (5.12): on $E_n \times \{\text{all marks}\}$, $\pi_{n,V} \geq K(\mathcal{S}_{n,V}) \geq 1 - 24\delta_{n,V}$. The upper bound (3.12) restates (6.3): $\{\pi_{n,V} > 4\delta_{n,V}\} \subseteq H_k \times \{\text{all marks}\}$, of probability at most $C_k p_n(a)$. The disabled statement is Section 6. For the last sentence, fix $\eta \in (0, 1)$; along $V(n)$ both 24δ and 4δ are eventually below $\min(\eta, 1 - \eta)$, so $\{\pi \geq 1 - 24\delta\} \subseteq \{\pi \geq \eta\} \subseteq \{\pi > 4\delta\}$, and the same inclusions hold after intersecting with $E_n \subseteq G_n$ on the left and with G_n on the right. $\square$

7.5 What the volume envelope buys

The envelope $10^{60}(n+1)^2$ is convenient for the joint statement. The finite theorem gives a smaller sufficient volume for a prescribed reliability: for $0 < \eta < 1$,

$$V \geq \max\left\{\frac{2n}{d}, \frac{n}{c} \log \frac{24}{\eta}\right\} \implies \inf_{\omega \in E_n, m} K_{n,V,\omega,m}(\mathcal{S}_{n,V}) \geq 1 - \eta, \quad (7.2)$$

because the right side is exactly $24e^{-cV/n} \leq \eta$. To make a noise term $Ce^{-cV/n}$ smaller than ηp_n , replace $24/\eta$ by $C/(\eta p_n)$; since $\log(1/p_n) \sim n \log 2$, a quadratic scale suffices to make molecular noise smaller than environmental rarity. Neither (7.2) nor the larger envelope is a necessary-volume statement. The constants d, c come from the concentration inequalities of Appendix B with the tolerances of (5.7), the smallest of which is $\rho/100$; they are sufficient proof bounds, not fitted or physically calibrated thresholds.

Remark 7.2 (Independent complete trials). Since $p_n \sim 9/(\pi^2 X_n)$, every probability in (3.9) is $\Theta(1/X_n) = \Theta(n/R_n)$. For independent complete trials, each resampling environment, marks and trajectory, the number of trials to the first productive one is geometric with mean $\Theta(X_n)$; with m trials the success probability is $1 - (1 - P)^m$, which tends to 0 when $m/X_n \rightarrow 0$ and to 1 when $m/X_n \rightarrow \infty$. Repeated runs in one fixed environment are not independent complete trials, and a Θ law gives no Poisson constant at $m \asymp X_n$. Exponential decay in n is polynomial decay in the explicitly enumerated catalog size and is not a computational hardness statement.

8 Discussion

What the separation says. Theorem 3.1 compares, in one probability space, a structural certificate and a productive requirement. The certificate is present with limiting probability $\theta(q_*) \in (0, 1)$ and, when present, is typically small. The requirement, bounded mass plus a fixed positive export from food alone in a fixed window, has probability of order $1/X_n$, and its conditional probability given the certificate has the same order. Removing catalysis makes the requirement asymptotically impossible at the same scale, so the productive events that do occur are catalytic. The two probabilities differ because they measure different things: the certificate is insensitive to rates, concentrations and competition, whereas the productive event needs a catalyst–substrate incidence that is *short* in the sense of Section 6, and short incidences are rare under the sampler. The lower bound shows that one such incidence, with food that catalyzes nothing, already suffices in the presence of arbitrary further nonfood catalysis.

What it does not say. The theorem is stated for one sampler, one exponent sequence, one reactor with fixed coefficients, and one observable with fixed window and thresholds. It does not say that RAF theory is wrong, that RAFs are dynamically inert, or that productive chemistry is universally rare. The constants are explicit but enormous, in both directions: C_k counts a catalog of length- $4 \cdot 10^8$ words, and the volume envelope contains 10^{60} . The result is an asymptotic order statement with a logarithmic rate, not a calibrated transition. Remark 3.4 records that nothing is claimed about the RAF size of typical productive environments. The reactor is an effective, unsaturated, well-mixed mass-action model; spatial structure, saturation, measured molecular energies and heredity are outside its scope.

Relation to the structural literature. Hordijk and Steel [10] introduced the capped power-law variant and reported by simulation that heterogeneous catalysis lowers the level of catalysis required for RAF emergence and changes RAF sizes. Part (i) and Remark 3.3 show that, at the level of the limiting existence probability at linear intensity, the capped-Zipf law at $a_n = 2 - 2/n$ is governed by the same survival function as the uniform law of the companion paper [1]; what the tail changes is the fluctuation structure on the growing catalogue, which is exactly what the two-band construction of Section 4.2 exploits to keep certificates small. The minimum RAF question is computationally hard in general [17]; part (ii) is a probabilistic statement about a typical sampled network, with no algorithmic content.

Relation to the kinetic literature. Kinetic simulations of the binary polymer model go back to [5]; stochastic small-number effects and incoming flux were studied in [6], and resource focusing with food-only initiation in [7]. The evolvability critique of [20] and the response of [9] concern selection between RAFs rather than the existence of production; the present theorem is a statement about one step earlier, whether a sampled RAF-containing network produces at all in a fixed window. The corrected growth function Φ belongs to the family of average Lyapunov and auxiliary functions used in permanence theory [15] and in bounding time averages [19]; its inequalities are established directly for this model. The transfer from drifts to trajectories is proved from the marked jump kernel (Appendix B) rather than through a fluid-limit theorem of the type surveyed in [3], because the error must be controlled at the scale $e^{-cV/n}$ where it is used, uniformly over the environment.

Numerical illustration. Figure 3 shows the deterministic mass-action trajectory of one specified environment in the complete $n = 4$ catalog (30 species, 68 channels), with all basal reactions at

the coefficients (2.9), the feed, the dilution, the self-catalyzed ligation r_* , and four further nonfood catalytic assignments chosen explicitly so that E_4 holds. It is a mechanism illustration, not a sample from the source law and not a count-process simulation; the probability theorem is analytic. The stored DOP853 computation uses absolute tolerance 10^{-18} and compares relative tolerances 10^{-10} and $2 \cdot 10^{-12}$: the maximum discrepancy over the saved arrays is below $1.6 \cdot 10^{-8}$ and the total mass differs from 10 by less than $8 \cdot 10^{-15}$. Enabled and disabled export over $(1, 100]$ are approximately 162.69 and $4.15 \cdot 10^{-5}$; export can exceed instantaneous mass because the reactor is continuously fed. These are numerical diagnostics, not rigorous integration certificates.




Figure 3: A specified finite host at the theorem’s kinetic parameters ($n = 4$). Left: the selected product $x = x_{0011}$ and the total nonfood mass M , with the catalysis-disabled trajectory. Right: signed catalytic input, positive basal input and export, accumulated over $(1, t]$, with the disabled export. The disabled trajectory keeps the same basal rates and feeds.

Further results left to a supplementary report. Three results proved in the same framework are omitted here because they concern different questions. A rare-event theorem attributes productive operation to one specified RAF at a much smaller volume $V = 40 + \lfloor \sqrt{n} \rfloor$, with a paid-path lower bound and a catalog upper bound giving the same logarithmic rate $-\log 2$ but no uniform Θ law. A deterministic proposition shows that the mass-action trajectory of any environment in E_n sustains a positive long-time mean catalytic input. A recurrence proposition shows that at any fixed finite volume the count process is positive recurrent and returns to the empty state, so that everlasting residence above a product floor is impossible; finite-window reliability as $V \rightarrow \infty$ and eventual loss as $t \rightarrow \infty$ at fixed V are different limits and do not conflict.

A The shifted growth inequality

This appendix proves (5.9). Put $D(v) = (1 - v)_+$ and $\eta = 10^{-5}$. Let a_u, a_w, a_x be the compensated coordinate errors, with $|a_u|, |a_w| \leq \eta$, $|a_x| \leq x/100$ and $C = C(M) \geq 0$. Between jumps the corrected coordinates $u - a_u$, $w - a_w$, $x - a_x$ are absolutely continuous with derivatives equal to the coordinate drifts; across jumps the same holds because the jumps of a coordinate and of its compensated error coincide. Hence

$$\tilde{\Phi} = \log(x - a_x) - 4D(u - a_u)^2 - 4D(w - a_w)^2$$

is absolutely continuous on any interval where x is bounded away from zero; the squared positive part is continuously differentiable, including at its corner.

The elementary inequalities $D(u - a_u) \leq 1 + \eta$ and $D(u)^2 \leq D(u - a_u)^2 + 3\eta$, together with (5.3), give

$$8D(u - a_u)b_u \geq 8D(u - a_u)^2 - 46C - 8(1 + \eta)\eta(1 + 23C),$$

where $D(u - a_u)u \leq (1 + \eta)^2/4$ bounds the depletion term; the unshifted coefficient is 46 and the explicit error term covers the perturbation. The same holds for w . Positivity of the basal term in (5.4) and $|a_x| \leq x/100$ give

$$\frac{b_x}{x - a_x} \geq 0.99(4uw) - 1.02(1 + 25C).$$

For $u, w \geq 0$,

$$0.99(4uw) + 8D(u)^2 + 8D(w)^2 \geq 2.97,$$

as one checks by minimizing on $[0, 1]^2$ and on its boundary; replacing a coordinate larger than 1 by 1 cannot increase the expression. Combining the three estimates, the constant term exceeds 1.9 and the coefficient of $-C$ is smaller than 120, which is (5.9).

On the good event $x \geq \rho$ and $L \leq 11$, so $x - a_x \geq \rho/2$ and $x - a_x \leq 3$, and the total possible increase of $\tilde{\Phi}$ on $[1, 100]$ is less than

$$\log(6/\rho) + 8(1 + \eta)^2 < 54;$$

for instance $e > 65/24$ and $(65/24)^{45} > 18 \cdot 10^{18} = 6/\rho$ verify the logarithm without transcendental evaluation. Integration of (5.9) over $[1, 100]$ gives (5.10). No logarithm of a jump process is treated as a drift-only quantity: the logarithm is applied to the continuous compensated coordinate.

B Stopped kernel bounds and physical time

We give the concentration argument behind $\delta_{n,V}$, including the final incomplete holding interval. Stop at the first exit from $\{L \leq 11\}$ or at time $T = 101$. On this region the total jump rate is at most $24000V$ by (5.1). Each concentration coordinate jumps by at most $2/V$, with quadratic-variation rate at most $96000/V$. For the nonfood mass and for the export and positive-input rewards, an internal jump has magnitude at most $4/V$, an outflow jump at most n/V , and the outflow quadratic-variation contribution is bounded using $V^{-2} \sum_x \ell(x)^2 N_x \leq nL/V$; so the reward quadratic-variation rate is at most $(384000 + 11n)/V \leq 96011n/V$ for $n \geq 4$.

Marked-kernel supermartingale. In a state with channel rates λ_j and total rate q , the next waiting time and mark have density $\lambda_j e^{-qs} ds$. If r_j are bounded reward increments, $\mu = \sum_j \lambda_j r_j$ and $Q \geq \sum_j \lambda_j r_j^2$, then

$$\sum_j \lambda_j (e^{\vartheta r_j} - 1) \leq \vartheta \mu + \vartheta^2 Q \quad (|\vartheta r_j| \leq 1),$$

by $e^v \leq 1 + v + v^2$ on $[-1, 1]$. Multiplied by the compensating waiting-time factor, the exponential of $\vartheta \times (\text{reward minus compensator})$ minus $\vartheta^2 \times (\text{integrated } Q)$ is a nonnegative supermartingale along the jump chain. At a deterministic endpoint the no-jump survival term e^{-qs} is included, so the compensator contains the elapsed part of the last holding interval; conditioning on completion of that interval would change the law and is not used.

Maximal inequality. Stop the supermartingale at the first threshold crossing, apply the expectation bound, and pass from finitely many jumps to the full stopped path by monotone limits.

Applying both signs of ϑ and optimizing within the admissible range gives, for a positive parameter d ,

$$\begin{aligned}\mathbb{P}(\sup |A_{\text{reward}}| > d + n/V) &\leq 2 \exp\left\{-\frac{d^2 V}{4n(96011T + d)}\right\}, \\ \mathbb{P}(\sup |A_{\text{coordinate}}| > d + 2/V) &\leq 2 \exp\left\{-\frac{d^2 V}{4(96000T + 2d)}\right\}.\end{aligned}$$

The added jump allowances permit one crossing jump and interpolation at physical-time endpoints. Rewards are frozen after the first mass exit and not restarted. The bounds need neither independent coordinates nor independent successive rewards; they hold conditionally on any environment, marks and history.

Constants. With d and c from (3.2) and $T = 101$, $\frac{d^2}{4n(96011 \cdot 101 + d)} \geq \frac{c}{n}$ and $\frac{d^2}{4(96000 \cdot 101 + 2d)} \geq \frac{c}{n}$, so each two-sided failure probability is at most $2\delta_{n,V}$. For $V \geq 2n/d$ the allowances $d + n/V \leq \frac{3}{2}d$ and $d + 2/V \leq \frac{5}{4}d$ (using $n \geq 4$) are smaller than every tolerance in (5.7), the smallest being $\rho/100 = 2d$. There are twelve applications: nonfood mass, six food coordinates for mass control, the two food coordinates u, w at the tighter startup tolerance, the product x , the export, and the positive basal input; their union costs at most $24\delta_{n,V}$. The upper bound uses two reward applications ($4\delta_{n,V}$) and the disabled bound one ($2\delta_{n,V}$).

Mass localization. The compensated total-mass error is bounded by $|A_L| \leq \frac{1}{8} + \sum_{f \in F} \ell(f) \frac{1}{80} = \frac{1}{4}$. The linear comparison (5.8) applied to $\mathcal{GL} = 10 - L$, $L(0) = 10$, gives $L \leq 10.5$ up to the putative first exit, and the crossing-jump convention excludes an exit through 11 on the good event. The food and product comparisons of Section 5.3 are therefore valid throughout $[0, T]$.

All events are measurable in the jump representation, and nonexplosion with almost surely positive holding times identifies them with their physical-time counterparts. Integrating the finite source law, the marks and the jump law produces the probabilities in Theorem 3.1.

C Structural proof details

This appendix supplies the finite identities, the independent-closure facts, the degree-band coupling, the target estimate and the assembly used in Section 4. Throughout $a = a_n$.

C.1 Finite source weights and local convergence

For $A \subseteq \mathcal{R}_n$ the one-row probability $\mathbb{P}(C_x = A)$ is $\mu_{a, R_n}(|A|)/\binom{R_n}{|A|}$, and the configuration weight is (2.4); normalization follows by grouping subsets by cardinality. For s distinct channels the conditional one-row miss probability given $D_x = D$ is $\binom{R_n - s}{D}/\binom{R_n}{D}$, zero when $D > R_n - s$; the inclusion probabilities of one and two prescribed channels are $\mathbb{E}D_n/R_n$ and $\mathbb{E}[D_n(D_n - 1)]/[R_n(R_n - 1)]$, which give (4.2). The second factorial moment is at most the second moment and $R_n/(R_n - 1) \rightarrow 1$, so its scaled contribution vanishes by Lemma 4.1. For disjoint finite S, C the identity

$$\mathbf{1}\{S \subseteq O, C \cap O = \emptyset\} = \sum_{T \subseteq S} (-1)^{|T|} \mathbf{1}\{(T \cup C) \cap O = \emptyset\}$$

holds for every set O , by expanding $\prod_{r \in S} (1 - \mathbf{1}\{r \notin O\}) \prod_{r \in C} \mathbf{1}\{r \notin O\}$; multiplying by (2.4) and summing needs no independence within a row. The embedding of $\mathcal{R}_N$ into $\mathcal{R}_n$ preserves product word and split position, hence all cardinalities, which proves Proposition 4.2.

C.2 Independent closure: records, nuclei, continuity

Survival generates every target. For any cap $B \geq 2$, induction on a reversible derivation gives the alternative: the derivation fits inside cap B , or some accessible word w has $\ell(w) > B$ and a derivation entirely inside cap $\ell(w)$ (at a ligation leaving B the product is such a word; a cleavage inherits a record from its input). So survival supplies record words beyond every cap. Fix a target $v = v_1 \cdots v_\ell$. Given a record w , the sequence $w \rightarrow wv_1 \rightarrow \cdots \rightarrow wv$, $wv \rightarrow w + v$ uses food letters and at most $\ell + 1$ channels, all with products longer than w , whose coordinates lie outside the σ -algebra of channels up to length $\ell(w)$; their joint openness has probability at least $q^{\ell+1}$. Fix a cylinder A determined by a cap B and partition the occurrence of a longer record into disjoint first-record atoms, each determined by the coordinates up to its own record length; on each atom the displayed support is fresh. Summing gives, for $H_v = \{\text{survival and } v \text{ missing}\}$,

$$\mathbb{P}(H_v \cap A) \leq (1 - q^{\ell+1})\mathbb{P}(A) \quad (\text{C.1})$$

for every cylinder A . If $\mathbb{P}(H_v) > 0$, approximating its indicator by the increasing finite-coordinate σ -algebras would produce cylinders with conditional density arbitrarily close to one, contradicting (C.1). Hence $\mathbb{P}(H_v) = 0$, and a countable union over v shows that survival implies complete closure almost surely.

Finite witnesses and continuity above one half. Let $F_m(q)$ be the event that the closure generates all words of length at most m ; it is the increasing union over N of the events $F_{m,N}(q)$ that all these derivations exist inside cap N , so $\mathbb{P}(F_{m,N}(q)) \uparrow \mathbb{P}(F_m(q)) \geq \theta(q)$. If every word longer than m has an open split, $F_m(q)$ implies complete closure by induction on length; the probability that the split condition fails is at most $\tau_m(q)$ in (4.4), a union bound valid even when the finite witness uses channels beyond length m . Hence $\mathbb{P}(F_m(q)) \leq \theta(q) + \tau_m(q)$. For positivity of θ , require all channels through a fixed m open; channels with longer products are independent of this, and the sum of their word-level failure probabilities is finite, so the infinite product of success probabilities is positive and their intersection implies survival. For left continuity at $q > 1/2$: fix $q_0 \in (1/2, q)$, choose m with $\tau_m(q_0) < \eta$ and N with $\mathbb{P}(F_{m,N}(q)) > \theta(q) - \eta$; the finite-cap probability is a polynomial in the openness, so for $q' \uparrow q$ it eventually exceeds $\theta(q) - 2\eta$, and then $\theta(q') \geq \mathbb{P}(F_{m,N}(q')) - \tau_m(q') > \theta(q) - 3\eta$; monotonicity gives the other inequality. For right continuity, survival is the decreasing intersection of the fixed-boundary escape events, each determined by a finite catalogue with continuous probability, whose probabilities bound θ from above. This suffices at $q_* > 1/2$.

C.3 Degree bands and the source coupling

With $h_n, L_n, k_n, d_0, d_1, d_2, \varepsilon_n$ as in (3.3) and (4.5), eventually $d_0 \leq d_1 < d_2 < R_n$. Put $Y_L = D_n \mathbf{1}\{d_0 \leq D_n < d_1\}$ and $Y_H = D_n \mathbf{1}\{d_1 \leq D_n < d_2\}$; both intervals lie below the cap atom. Integral comparison for $(d+1)^{-a_n}d$, with negligible normalized endpoint errors, gives

$$\frac{\mathbb{E}Y_L}{n} \longrightarrow \frac{3}{2\zeta(2)} = \Lambda, \quad \frac{\mathbb{E}Y_H}{h_n} \longrightarrow \frac{19 \log 2}{5\zeta(2)}. \quad (\text{C.2})$$

For the first, $\log d_0/n \rightarrow 0$, $\log d_1/n \rightarrow \log 2$, and the primitive of $x^{-1+2/n}$ is $\frac{n}{2}x^{2/n}$. For the second, with $u = \log x/n$ the integral is $n \int e^{2u} du$ over an interval of width $[h_n - \lfloor h_n/20 \rfloor] \log 2/n$ near $u = \log 2$, whose integrand tends uniformly to 4, giving $4 \cdot \frac{19}{20} \log 2$ after division by h_n .

For either band variable $0 \leq Y \leq U$, independence of degrees gives $\text{Var}(\sum_x Y_x) \leq X_n \mathbb{E}Y^2 \leq X_n U \mathbb{E}Y$ with $U \leq X_n$ eventually, so a relative downward deviation δ has probability at most

$1/(\delta^2 \mathbb{E}Y)$; both band means diverge, so the relative deficits vanish. For the low union, conditional openness is

$$q_L(d) = 1 - \prod_{x: d_0 \leq d_x < d_1} \left(1 - \frac{(1 - \varepsilon_n)d_x}{R_n}\right) \geq 1 - \exp\left(-\frac{1 - \varepsilon_n}{R_n} \sum_x Y_{L,x}\right),$$

which by (C.2) and $R_n/(nX_n) \rightarrow 1$ is at least $1 - e^{-t}$ with probability tending to one, for each $t < \Lambda$. For the high union, with retention 19/20 and relative deficit 1/50, the normalized exponential intensity has lower limit

$$\frac{49}{50} \cdot \frac{19}{20} \cdot \frac{19 \log 2}{5\zeta(2)} > 1.47 > \frac{29}{20}, \quad (\text{C.3})$$

and since $1 - e^{-z} \sim z$ this gives the floor $\pi_n = \frac{29}{20}n^{-1/4}$ eventually, outside a vanishing degree exception. The strict margins in (C.3) are specific to $a_n = 2 - 2/n$.

Exact row marginals. Conditional on the degrees, generate independent auxiliary Bernoulli subsets A_x with the stated per-channel probabilities. A binomial Chernoff bound gives $\mathbb{P}(|A_x| > d_x \mid d) \leq e^{-\varepsilon_n^2 d_x/4}$ for low rows, and eventually $\varepsilon_n \leq 1/20$ covers the high rows. On $|A_x| \leq d_x$ complete A_x to a uniform d_x -subset; on overflow choose a uniform d_x -subset independently. The result is uniform by invariance under channel permutations, and on no overflow every auxiliary incidence is a source incidence. Doing this independently across rows and integrating the exact degree law gives the coupling, with total overflow probability at most $X_n e^{-\varepsilon_n^2 d_0/4} \leq e^{-n}$ eventually; inactive rows cost nothing. This is an increasing-event transfer with an explicit exceptional event, not a stochastic domination claim.

Number and length of labels. Since $\mathbb{E}D_n \leq 2n$ eventually, Markov's inequality gives $\mathbb{P}(\#\{x : D_x \geq d_1\} \geq n^3 2^{h_n}) \leq X_n \mathbb{E}D_n / (d_1 n^3 2^{h_n}) \leq 4/n^2$. For a length threshold s and degree threshold d , a union bound and Markov's inequality give $\mathbb{P}(\exists x : \ell(x) < s, D_x \geq d) \leq 2^s \mathbb{E}D_n / d$; at $(s, d) = (n - 2h_n, d_1)$ and $(k_n - h_n, d_0)$ these are $O(n2^{-h_n})$. They hold before any owner is selected, and every selected low label has degree at least d_0 , so the second exclusion covers the whole adaptive selection. This proves (4.8).

C.4 The uniform bounded target estimate

Let w be a fixed word of length ℓ and let the independent field have openness p . Work with the distinct nonempty substrings of w , ordered by length, with all substrings of length at most L available. For each distinct longer product word u expose its full row of split-position bits; rows of different product words are independent, and repeated occurrences of a substring use the same row and are never counted as independent trials.

Call u *cut-dense* for the current known set if at most $\ell(u)/8$ of its proper prefix cuts have a missing prefix and at most $\ell(u)/8$ have a missing suffix. For $\ell(u) \geq 4$, cut density gives at least $\ell(u)/2$ usable splits, so conditionally on the history of shorter words the probability that u is not generated is at most $e^{-p\ell(u)/2}$. If cut density fails somewhere after the exploration, take a shortest failing word u ; one side has more than $\ell(u)/8$ missing proper parts, with cut-index set J . Each such part has length greater than L , is cut-dense by minimality, and is a distinct word (for a fixed side, different cut indices give different lengths). For fixed u , side and J , summing over exposure histories, each selected part fails while cut-dense and charges at least $p/2$ times its length through its fresh row; iterated conditional expectation bounds the joint probability by $\exp(-\frac{p}{2} \sum_{i \in J} \ell(u_i)) \leq \exp(-pL\ell(u)/16)$. The density predicates refer to the final known set, but the shorter parts of each word are settled at its exposure, so the length ordering makes the charging valid. There are at most $2^{\ell(u)}$ sets J ; if $pL \geq 32 \log 2$ and $\ell(u) \geq L$ then $2^{\ell(u)} e^{-pL\ell(u)/16} \leq e^{-pL^2/32}$.

With two sides and at most ℓ^2 distinct substrings, any density failure has probability at most $2\ell^2 e^{-pL^2/32}$, and in its absence the failure of w itself costs at most $e^{-p\ell/2}$. This is (4.9).

On success, choose one successful split per needed product and trace a derivation of w to leaves in the nucleus; the tree has at most ℓ leaves (nonempty pieces partitioning w), hence at most $\ell - 1$ internal nodes and at most $\ell - 1$ distinct added channels. If a base support $S_0(H)$, possibly depending on the high field H , generates the nucleus, then raw exploration success supplies a bounded additional support in H regardless of that dependence; the guarded event that the base generates the nucleus but no bounded support exists is contained in the raw failure event. For the selected low targets, condition on the degrees and the low rows first, fixing targets and base; the high field then has the independent law just analysed, and a union bound over the fixed target list followed by averaging is legitimate.

C.5 Assembly and reversible trimming

Fix an admissible low openness $q \in (1/2, q_*)$ and $t < \Lambda$ with $q \leq 1 - e^{-t}$. Choose m with $\tau_m(q) < \eta$ and a seed cap N capturing mass greater than a chosen $c < \theta(q)$. Retain all channels of the seed witness inside cap N (at most R_N), require an open split for every word of length $m + 1, \dots, L_n$ and retain one such split; the extension loss is bounded by $\tau_m(q)$ uniformly in L_n . Select a low-row catalyst for each retained channel; support and label list have size at most $R_N + 2^{L_n+1}$. The high target set is W_H . The deterministic bounds and the guarded target estimates give a good event on which the union of the base and all target derivations has size at most $b_N(n)$ in (4.10). The loss from seed-and-extension mass to an actual source RAF is at most the sum of five vanishing terms: (1) high-target inadmissibility (too many high labels or a too-short high label); (2) a too-short active low label; (3) guarded failure of a bounded high-label derivation, including the high-openness exception; (4) guarded failure of a bounded selected-low-label derivation, including the same exception; (5) overflow in the Bernoulli-row coupling. Counting an openness exception in both envelopes is harmless, and no factorization of seed and target success is used: an indicator inequality and summation suffice.

Let S be the assembled support; every channel of S has a catalyst in $\text{cl}_F(S)$. Let $T \subseteq S$ consist of the channels at least one direction of which becomes enabled at some finite stage of the closure. Induction on closure stages gives $\text{cl}_F(T) = \text{cl}_F(S)$, since every productive step of S uses a channel retained in T . For a channel in T its enabled direction places both sides in the closure, and its catalyst remains generated because the closure is preserved. A word of length three is generated, so $T \neq \emptyset$; hence T is a RAF with $|T| \leq |S|$. For every $r < \theta(q)$ choose the seed and extension margins before letting $n \rightarrow \infty$; the five errors and the low-openness exception vanish, giving $P_n(b_N(n)) > r$ eventually. Continuity of θ and $q \uparrow q_*$ supply every $r < \theta(q_*)$, which is Proposition 4.3.

C.6 First-exit adapters

The source-open set O_n contains every RAF support, and monotonicity of reversible closure embeds all internally generated catalysts into $\text{cl}_F(O_n)$. A first food-enabled channel of a nonempty RAF has product length at most four, and the whole catalogue through length four has $R_4 = 68$ channels. For fixed $K \geq 2$, the first productive step escaping length K cannot be a cleavage; its two ligation inputs have length at most K and its product length lies in $(K, 2K]$, and induction on the preceding derivation gives exact equality of the escape event with its restriction to cap $2K$ for $n \geq 2K$, product words and split positions embedding literally. If escape fails but a RAF exists, a food-enabled channel has an internally generated catalyst of length at most K ; each specified pair has source probability exactly $\mathbb{E}D_n/R_n$, giving Lemma 4.5. The argument uses neither a bound on the number

of RAF witnesses nor independent catalysis by accessible molecules.

D Formal verification and reproducibility

The formal developments use Lean 4 (toolchain `leanprover/lean4:v4.30.0`) with the repository’s frozen Mathlib manifest (Mathlib commit `c5ea00351c28e24afc9f0f84379aa41082b1188f`; manifest SHA-256 `a8df0b60...`). Strict verification compiles the local import closure with warnings treated as errors, rejects `sorry`, and audits the axioms of the selected declarations; every selected declaration depends only on `Classical.choice`, `Quot.sound` and `propext`. No stochastic estimate is added as an axiom. Verification receipts record source hashes, dependency hashes, commands and axiom audits; their proof-file hashes are quoted below by their first sixteen hexadecimal digits.

D.1 Concordance

The structural development is the namespace `PowerLawSmallRAF`; the reactor development is `RandomViability`, which imports it. In the reactor development the source law is `sourcePowerLawConfigWeight`, the RAF event is `anySourceRAF` (existence of a reversible RAF for `binaryPolymerCRSn2` with the sampled catalysis), the productive event is `physicalOutputEvent` (mass at most $11V$ through time 100 and windowed export reward above $1/10$ on $(1, 100]$), the joint probability is `outputJointProbability` (source average, mark average and trajectory law), the incidence is `productiveSourceIncidence` $= p_n$, the noise rate is `markedNoiseRate` $= c$, and the volume envelope is `collectiveMinimalVolume` $= 10^{60}(n + 1)^2$.

Statement in this paper	Lean declaration (module; receipt hash prefix)
Lemma 4.1	<code>sourceExactCriticalFirstMoment_scaled_tendsto,</code> <code>sourceExactCriticalSecondMoment_scaled_tendsto_zero</code> (<code>SourceCriticalSecondMoment.lean</code>)
Proposition 4.2	<code>sourceFiniteCap_event_tendsto</code> (<code>SourceFiniteCapLimit.lean</code>)
Nucleus and survival (App. C.2)	<code>fixed_seed_captures_survival_with_small_extension_error</code> (<code>FixedSeedSurvivalApproximation.lean</code>)
Target estimate (4.9)	<code>finite_ligation_target_failure_bound</code> (<code>FiniteLigationTarget.lean</code>); <code>source_bounded_target_failure_bound</code> (<code>BoundedTargetProbability.lean</code>)
Source coupling (App. C.3)	<code>sourceFiniteSeedMass_eventually_le_source</code> (<code>FiniteSeedSourceProbability.lean</code>)
Proposition 4.3	<code>source_subexponential_RAF_endpoint_lower</code> (<code>SourceCriticalEndpoint.lean</code>)
Lemma 4.5, (4.12)	<code>sourceFullRAFProbability_le_escape_add_seed</code> (<code>SourceEscapeProbability.lean</code>); <code>sourceEscapeProbability_tendsto</code> (<code>SourceEscapeLimit.lean</code>); <code>source_RAF_survival_upper</code> (<code>SourceRAFSurvivalUpper.lean</code>)
Theorem 3.1(i), limit	<code>PowerLawSmallRAF.source_RAF_probability_tendsto</code> (<code>Main.lean</code> ; <code>425914e5b768a3d1</code>)
Theorem 3.1(ii), cutoff limits	<code>source_subexponential_cutoff_probability_tendsto</code> , <code>source_conditional_minimum_subexponential</code> (<code>Main.lean</code>)

Statement in this paper	Lean declaration (module; receipt hash prefix)
Theorem 3.1(ii), explicit B_n and tails	sourcePublicationBudget_subexponential, source_publication_cutoff_probability_tendsto, source_conditional_publication_cutoff_tendsto, source_large_minimum_exception_probability_tendsto_zero, source_conditional_exponential_rate_tail_tendsto_zero (PowerLawSmallRAF/PublicationCorollaries.lean; fb71101d4a92c251)
Capacity bound, Lemma 6.1	catalyticFoodEnvelope_mass_bound (FoodUptake.lean); positive_catalytic_drift_small (MarkedCatalyticInput.lean)
Startup floors, Section 5.3	physical_food_prefix_floor (PhysicalFoodStartup.lean); physical_product_interval_floor (PhysicalProductStartup.lean)
Output event; finite success inclusion; pointwise upper and disabled bounds	physicalOutputEvent, finite_success_implies_output (PhysicalOutputEvent.lean; 223f4f297ca30db7); physical_collective_upper_without_short_incidence (MarkedCatalyticUpper.lean); PhysicalOutputProbability.lean (07e443e1d91390dc)
Finite sandwiches (3.6)–(3.7)	RandomViability.averaged_output_bounds (SourceOutputProbability.lean; 91502845f4191378)
Disabled finite bound	RandomViability.averaged_output_disabled_bound (same module)
Θ law (3.9), first two	RandomViability.label_free_output_theta, RandomViability.output_and_some_raf_theta (RandomViability/PublicationCorollaries.lean; 1d10e1359bba5302); output_joint_probability_theta (OutputJointLimit.lean; 5d333feb48191ce3)
Ratio limit (3.10)	RandomViability.disabled_enabled_output_ratio
Corollary 7.1	RandomViability.output_conditioned_on_raf_theta (PublicationStructure.lean; 7126ee713efed312), which divides output_and_some_raf_theta by source_raf_probability_positive, itself derived from source_subexponential_RAF_endpoint_lower
$\mathbb{E}D_n/n \rightarrow 9/\pi^2$	RandomViability.mean_degree_per_length_limit (PublicationStructure.lean)
Finite reliability (7.2)	RandomViability.noise_error_below_tolerance, composed with (3.8)
Corollary 3.2; $0 < \theta(q_*) < 1$ upper bound; M_n interpretation; independent-trial remark	Conventional consequences of the verified statements, proved in the text; not separate Lean exports

The receipt for `PublicationStructure.lean` is the publication entry point for the productive half: its strict run compiled the whole import closure, including the structural development, in one pass. The structural terminal receipt covers the four root declarations of `Main.lean`; the corollary receipt covers the five explicit-cutoff statements. A source audit compared the current bytes of every file in both closures (566 distinct local sources for the reactor closure; 365 for the structural closure) with the hashes recorded in the receipts.

D.2 Scope of the formalization

The Lean statements are event-level: the minimum M_n is not a formal random variable, and the cutoff theorems are stated as convergence of $P_n(\lfloor 2^{bn} \rfloor)$ and of the conditional ratio; their reading as statements about $\log M_n/n$ is the elementary interpretation on the finite configuration space. The

productive theorems are stated for the averaged probability $\mathbb{P}(\mathcal{F}_{n,V})$ and for the pointwise upper bound outside the short-incidence event; the environment-level Corollary 3.2 combines them in prose. The catalog constant C_k appears in Lean as `collectiveSourceUpperConstant`. The results left to the supplementary report (attribution at small volume, deterministic sustained mean, finite-copy recurrence) are not used anywhere in this paper.

D.3 Reproduction

The manuscript source, bibliography, figure generator and figure data accompany the paper. The figure generator fixes the kinetic-mark seed and the four extra catalytic assignments of the $n = 4$ host, checks mass conservation, compares two integration tolerances, and evaluates the source law with 100-digit Hurwitz-zeta arithmetic; it performs no Monte Carlo estimate of any rare event. Strict verification of the terminal declarations is run through the repository’s verification bridge, which sets the Lean working directory to the frozen Mathlib project and enables warnings-as-errors; the commands are recorded with the receipts.

References

- [1] Author name to be inserted. A nontrivial critical window for RAF emergence in the binary polymer model. Companion manuscript, 2026; arXiv preprint in preparation, 2026.
- [2] David F. Anderson and Thomas G. Kurtz. Continuous time Markov chain models for chemical reaction networks. In Heinz Koepl, Gianluca Setti, Mario di Bernardo, and Douglas Densmore, editors, *Design and Analysis of Biomolecular Circuits*, pages 3–42. Springer, New York, 2011. doi: 10.1007/978-1-4419-6766-4_1.
- [3] R. W. R. Darling and J. R. Norris. Differential equation approximations for Markov chains. *Probability Surveys*, 5:37–79, 2008. doi: 10.1214/07-PS121.
- [4] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.
- [5] J. Doynne Farmer, Stuart A. Kauffman, and Norman H. Packard. Autocatalytic replication of polymers. *Physica D: Nonlinear Phenomena*, 22(1–3):50–67, 1986. doi: 10.1016/0167-2789(86)90233-2.
- [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] Varun Giri and Sanjay Jain. The origin of large molecules in primordial autocatalytic reaction networks. *PLoS ONE*, 7(1):e29546, 2012. doi: 10.1371/journal.pone.0029546.
- [8] 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.
- [9] Wim Hordijk and Mike Steel. Conditions for evolvability of autocatalytic sets: a formal example and analysis. *Origins of Life and Evolution of Biospheres*, 44(2):111–124, 2014. doi: 10.1007/s11084-014-9374-5.

- [10] Wim Hordijk and Mike Steel. Autocatalytic sets in polymer networks with variable catalysis distributions. *Journal of Mathematical Chemistry*, 54(10):1997–2021, 2016. doi: 10.1007/s10910-016-0666-z.
- [11] Wim Hordijk, Stuart A. Kauffman, and Mike Steel. Required levels of catalysis for emergence of autocatalytic sets in models of chemical reaction systems. *International Journal of Molecular Sciences*, 12(5):3085–3101, 2011. doi: 10.3390/ijms12053085.
- [12] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi: 10.1016/S0022-5193(86)80047-9.
- [13] 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.
- [14] Elchanan Mossel and Mike Steel. Random biochemical networks: the probability of self-sustaining autocatalysis. *Journal of Theoretical Biology*, 233(3):327–336, 2005. doi: 10.1016/j.jtbi.2004.10.011.
- [15] Swati Patel and Sebastian J. Schreiber. Robust permanence for ecological equations with internal and external feedbacks. *Journal of Mathematical Biology*, 77(1):79–105, 2018. doi: 10.1007/s00285-017-1187-5.
- [16] Mike Steel. The emergence of a self-catalysing structure in abstract origin-of-life studies. *Applied Mathematics Letters*, 13(3):91–95, 2000. doi: 10.1016/S0893-9659(99)00191-3.
- [17] 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.
- [18] 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.
- [19] Ian Tobasco, David Goluskin, and Charles R. Doering. Optimal bounds and extremal trajectories for time averages in nonlinear dynamical systems. *Physics Letters A*, 382(6):382–386, 2018. doi: 10.1016/j.physleta.2017.12.023.
- [20] 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.