

Linear-size autocatalytic sets are asymptotically absent in the sparse binary polymer model

Anonymous manuscript

14 September 2026

Abstract

In the binary polymer model of Kauffman, molecules are bit strings of length at most n , reactions are ligations and their reverse cleavages, and a random catalysis assignment decides which molecules catalyse which reactions. Hordijk and Steel (*J. Math. Chem.*, 2016) introduced a *sparse* catalysis law, in which a molecule is active with a small probability and an active molecule catalyses each reaction independently with probability $1/n$, and showed that when the expected number of reactions catalysed per molecule is λn for a large fixed λ , the system contains a reflexively autocatalytic and food-generated set (RAF) with fewer than $2\lambda n^2$ reactions with probability close to one. They asked whether the quadratic size can be reduced to linear, and in particular whether a single molecule is likely to catalyse a set of reactions that suffices to construct that molecule once cleavage reactions are allowed. We answer both questions negatively. For every fixed food horizon, every fixed nonnegative intensity λ and every fixed constant C , the probability that the model contains a nonempty RAF with at most Cn reversible reactions tends to zero as $n \rightarrow \infty$, at rate $O(1/n)$; the probability that some molecule catalyses a self-constructing set tends to zero as well, cleavage included. The proof separates two regimes by the number of catalysts a RAF uses. Boundedly many catalysts are excluded by a productive-prefix argument whose constants do not depend on n . Many catalysts are excluded by a counting lemma of independent interest: the number of m -reaction sets that can be built productively from the food, multiplied by $m!$, is at most $((f + 3m)^2 + n(f + 3m))^m$, where f is the food size. The factorial comes from an injective acyclic dependency encoding of arbitrarily labelled constructions and is exactly the entropy saving that the sparsity of active molecules cannot otherwise absorb. Consequences include the divergence of the minimum RAF size relative to n , the divergence of the minimum catalyst count, conditional versions on RAF existence, and a superlinear-to-quadratic window for the minimum RAF size in the regime of Hordijk and Steel. The main theorem, its constituent estimates and all stated consequences are verified in Lean 4 against the literal split-position model.

Keywords. autocatalytic sets; RAF theory; binary polymer model; sparse catalysis; random reaction networks; origin of life; formal verification.

MSC 2020. 92E20, 60C05, 05A16, 92D15, 68V20.

1 Introduction

A reflexively autocatalytic and food-generated set, or RAF, is a finite combinatorial abstraction of a collectively self-sustaining chemistry: a set of reactions all of whose reactants can be built from a freely available *food* set by reactions of the set, and each of whose reactions is catalysed by a molecule that is food or is produced by the set [2, 11]. The notion descends from Kauffman’s autocatalytic sets of proteins [7, 8], and the *binary polymer model* that Kauffman proposed to study them remains the standard random model. Molecules are the bit strings of length at most n ; food consists of all strings of length at most a fixed small t ; the reactions

are all ligations $u + v \rightarrow uv$ together with their reverse cleavages; and catalysis is assigned at random. When each molecule catalyses each reaction independently with the same probability (the *uniform* law), bounds on the RAF probability initiated by Kauffman and Steel [7, 10] were sharpened by Mossel and Steel [9], who showed that RAFs appear precisely when the expected number f of reactions catalysed per molecule grows linearly with n ; at this level of catalysis the RAFs that first appear are of exponential size [3, 5].

Hordijk and Steel [3, 4] then asked how these conclusions depend on the *distribution* of catalytic activity across molecule types, motivated by the observation that real metabolic networks are far from uniform. Among the laws they analysed is the *sparse* law: each molecule type independently is *active* with probability p , and an active molecule catalyses each reaction independently with probability $1/n$, so that $f = p|\mathcal{R}_n|/n$. They proved that the linear threshold $f \asymp n$ for RAF emergence is universal across such distribution-based laws, but that the size of the RAFs that appear is not. In the sparse law with $f = \lambda n$ and λ large, Chernoff’s bound gives at most $2\lambda n$ active molecules with high probability, and a chain of fewer than n ligations reaches every molecule, so a RAF with fewer than $2\lambda n^2$ reactions exists with probability at least $1 - \varepsilon - o(1)$ [3, Corollary 2]. Their concluding section [3, Section 7] poses the problem this paper resolves. Having shown that the sparse model is likely to have RAFs of at most quadratic size when RAFs first appear, they ask whether this can be improved to order n , and in particular whether it is likely that a single molecule type catalyses a sequence of reactions sufficient to construct that molecule. They note that a simple argument rules the latter out asymptotically when only ligation reactions are used, but that “the situation is less clear without this restriction”, that is, once cleavage reactions are allowed.

1.1 Main result

We work with the literal reversible split-position binary polymer model: a *channel* is a product word w together with a split position, its two orientations are the ligation of the two factors and the cleavage of the product, and both orientations are catalysed by the same random coordinate, exactly as in the bidirectional convention of [3]. The sparse law is parametrised so that the expected number of channels catalysed by a molecule is λn (Section 2).

Theorem 1.1 (absence of linear-size RAFs). *Fix a food horizon $t \in \mathbb{N}$, an intensity $\lambda \geq 0$ and a real constant C . In the sparse binary polymer model with maximum length n ,*

$$\Pr\{\exists T : T \text{ is a nonempty RAF and } |T| \leq Cn\} \longrightarrow 0 \quad (n \rightarrow \infty). \quad (1)$$

Moreover the probability is $O(1/n)$.

All three parameters are fixed before n tends to infinity. In particular no choice of a larger fixed intensity produces RAFs within a fixed linear budget with non-vanishing probability, so the quadratic bound of [3, Corollary 2] cannot be improved to linear in the sense they ask (Corollary 3.3). The theorem covers every number of catalysts, both orientations of every channel, and the actual dependence between catalytic events that share an active molecule. The single-molecule question is answered by a short separate argument that also handles cleavage (Proposition 3.2), and the same method shows that the minimum number of catalysts used by any RAF diverges (Theorem 3.1).

1.2 Method

Let T be a RAF and let K be a set of molecules, one catalyst per channel, chosen inside the food closure of T . Every member of K is active, an event of probability $p_n = \lambda n^2/|\mathcal{R}_n| \approx \lambda n/|\mathcal{X}_n|$ per molecule, and each channel of T is covered by a conditional coordinate of probability $1/n$. Two effects compete. Few catalysts are cheap in activity but must cover many channels each;

many catalysts cover cheaply but are individually rare. The proof treats the two regimes separately, and a fixed cutoff on $|K|$ that does not depend on n separates them (Figure 1).

For a bounded number k of catalysts we use a *productive-prefix* argument (Section 5). A catalyst that is not itself reachable by short constructions forces the RAF to contain a productive construction of depth d from the food, and the number of such depth- d constructions is bounded independently of n because a depth- d construction only sees molecules of length at most $t2^d$. Choosing $d = k + 1$ leaves the factor $(|\mathcal{X}_n|p_n/n)^k/n \rightarrow 0$.

For many catalysts the union over catalyst sets is too expensive and we must instead count the constructions themselves. The core of the paper is the following deterministic statement (Lemma 6.1). Let $\mathcal{A}_m$ be the family of m -channel sets that admit a *productive program* from the food, that is, an ordering in which every channel is enabled in one of its orientations when it fires and adds at least one new molecule. Then

$$|\mathcal{A}_m| \cdot m! \leq ((f + 3m)^2 + n(f + 3m))^m, \quad f = |F|. \quad (2)$$

Each channel of a labelled construction can be described by an *instruction* that refers to its inputs either as food molecules or as endpoints of earlier channels, and the resulting acyclic code determines the labelled channel assignment uniquely. The factorial counts the $m!$ arbitrary labellings of one support. It is *not* the number of productive orders of a support, which may be as small as one; a single productive order merely certifies that the code is acyclic. Under a linear budget $m \leq Cn$ the factorial exactly cancels the n^m from the cleavage split positions and leaves $|\mathcal{A}_m| \leq (Dn)^m$ with D depending only on t and C . The n^{-m} from the conditional coordinates then cancels the n^m from the support count, and the activity factor $((f + 3Cn)p_n)^k$ with k large is exponentially small (Section 6).

A RAF need not be a productive program: it may contain channels that no construction ever uses. A saturated productive extraction (Lemma 4.2) replaces the RAF by a productive sub-support whose final state contains the entire original food closure, so every original catalyst is still available and every extracted channel is still covered. This step is what allows the support count to control arbitrary RAFs.

1.3 Formal verification

The main theorem, the two probability estimates, the factorial count, the saturated extraction, the single-catalyst statement and all consequences in Section 9 have been checked in Lean 4 [1, 12] against the literal source model, with the sparse law realised as an infinite product measure. The verified statements depend only on the standard axioms of Lean’s kernel. Section 10 gives the concordance between the paper and the formal development and states exactly what is cited rather than proved.

1.4 Organisation

Section 2 defines the model and records the counts and parameters used throughout. Section 3 states all results. Section 4 develops productive programs, saturated extraction and prefix counting. Section 5 proves the bounded-rank theorem, Section 6 the factorial count and the large-rank tail, and Section 7 assembles the main theorem with its rate. Section 8 treats the single-molecule question, Section 9 the consequences, Section 10 the verification, and Section 11 the scope of the result and open problems.




Figure 1: Architecture of the proof of Theorem 1.1. A fixed catalyst cutoff separates the productive-prefix estimate, which has no size restriction, from the factorial support count, which uses the linear budget.

2 The model

2.1 Molecules, channels and closure

Let $\mathcal{X}_n$ be the set of nonempty binary words of length at most n . A *channel* is a pair (w, i) with $w \in \mathcal{X}_n$ and $1 \leq i < |w|$; writing $w = uv$ with $|u| = i$, the channel has the two orientations

$$u + v \longrightarrow w \quad (\text{ligation}), \quad w \longrightarrow u + v \quad (\text{cleavage}),$$

and its *endpoints* are $\text{end}(w, i) = \{u, v, w\}$. Let $\mathcal{R}_n$ be the set of channels. Distinct split positions of the same word are distinct channels even when they produce the same unordered factor pair, and equal factors are allowed; this is the literal reaction catalogue of the binary polymer model, in which each ligation and its reverse cleavage are counted as one bidirectional reaction [3]. Summing 2^ℓ words with $\ell - 1$ split positions over lengths $\ell \leq n$,

$$N_n := |\mathcal{X}_n| = 2^{n+1} - 2, \quad R_n := |\mathcal{R}_n| = (n - 2)2^{n+1} + 4 \quad (n \geq 2), \quad (3)$$

so that $R_n/(nN_n) \rightarrow 1$, in agreement with the relation $|\mathcal{R}_n| \sim n|\mathcal{X}_n|$ of [9]. We shall also use $R_n \geq 2^n$ for $n \geq 2$.

The food set is $F = \mathcal{X}_t$ for a fixed horizon t (the literature instance is $t = 2$), and $f := |\mathcal{X}_t| = 2^{t+1} - 2$. We assume $n \geq \max\{2, t\}$, which is harmless for limits.

For a set $T \subseteq \mathcal{R}_n$ of channels and a set $A \subseteq \mathcal{X}_n$ of available molecules, a channel is *enabled* at A if both its factors lie in A or its product lies in A . The *food closure* $\text{cl}_T(F)$ is the least set containing F and closed under both orientations of every channel of T ; concretely it is the union of the increasing sequence $A_0 = F$, $A_{j+1} = A_j \cup \bigcup_{r \in T, r \text{ enabled at } A_j} \text{end}(r)$. The set T is *food-generated* if every channel of T has all its endpoints in $\text{cl}_T(F)$, equivalently is enabled there.

Definition 2.1 (RAF). Given a catalysis relation $\mathcal{C} \subseteq \mathcal{X}_n \times \mathcal{R}_n$, a set $T \subseteq \mathcal{R}_n$ is a *RAF* if it is nonempty, food-generated, and every $r \in T$ has a catalyst $x \in \text{cl}_T(F)$ with $(x, r) \in \mathcal{C}$.

This is the reversible RAF notion of [3]: catalysis is assigned to the channel, so a catalyst of (w, i) catalyses both the ligation and the cleavage, and a catalyst may be produced by the set rather than being available before the construction starts.

2.2 The sparse catalysis law

For each molecule $x \in \mathcal{X}_n$ let A_x be a Bernoulli variable of parameter p_n (*activity*), and for each pair $(x, r) \in \mathcal{X}_n \times \mathcal{R}_n$ let $B_{x,r}$ be a Bernoulli variable of parameter q_n (*conditional channel bit*). All these variables are independent. The random catalysis relation is

$$(x, r) \in \mathbb{C} \iff A_x = 1 \text{ and } B_{x,r} = 1.$$

This is the sparse law of [3, Section 2]: with probability p_n a molecule “catalyzes each reaction independently with probability $1/n$ ”, otherwise it catalyses nothing. Following [3] we normalise the expected number of channels catalysed by a molecule, $f_n = p_n q_n R_n$, to be λn :

$$q_n = \frac{1}{n}, \quad p_n = \frac{\lambda n^2}{R_n}, \quad p_n q_n R_n = \lambda n. \quad (4)$$

For fixed $\lambda \geq 0$ we have $p_n \leq \lambda n^2 / 2^n \rightarrow 0$, so $p_n \leq 1$ for all large n ; the formal model clips p_n to $[0, 1]$ at the finitely many small n where the formula exceeds one, which does not affect any limit. We record the consequences of (3)–(4) that are used below:

$$np_n \rightarrow 0, \quad \frac{N_n p_n}{n} = \frac{\lambda n N_n}{R_n} \rightarrow \lambda, \quad (f + 3Cn) p_n \leq \left(\frac{9}{16}\right)^n \text{ for all large } n. \quad (5)$$

The last statement holds because $(f + 3Cn)p_n \leq \lambda(fn^2 + 3Cn^3)/2^n$ and a fixed polynomial is eventually dominated by $(9/8)^n$.

The events $(x, r) \in \mathbb{C}$ and $(x, r') \in \mathbb{C}$ for $r \neq r'$ are positively correlated, since both require $A_x = 1$; [3] point out that this correlation is exactly what distinguishes distribution-based laws from the uniform law. All probability estimates below are stated and proved for the product law of the underlying variables (A_x) and $(B_{x,r})$, and none replaces the sparse law by independent catalysis marginals.

2.3 Notation

Throughout, $\lambda \geq 0$, t and C are fixed. We write $p = p_n$ and $q = q_n$ when n is clear. For a finite set S of channels its *universe* is $\mathcal{U}(S) = F \cup \bigcup_{r \in S} \text{end}(r)$, of cardinality at most $f + 3|S|$.

3 Statement of results

Let $L_n(C)$ be the event in (1), that some nonempty RAF has at most Cn channels. A *catalyst cover* of a RAF T is a set $K \subseteq \text{cl}_T(F)$ of molecules such that every channel of T is catalysed by a member of K ; its size is the *rank* of the cover. Let $E_{n,k}$ be the event that some nonempty food-generated channel set has an active catalyst cover of rank exactly k whose members are all reachable, i.e. lie in its food closure. (Unused members of a cover may be deleted, so “rank at most k ” and “rank exactly k' for some $k' \leq k$ ” describe the same event.)

Theorem 3.1 (bounded catalyst rank). *For every fixed $k \geq 1$, $\Pr(E_{n,k}) = O(1/n)$. Consequently, for every fixed K , the probability that some RAF, of any size, has a catalyst cover of rank at most K tends to zero.*

Proposition 3.2 (single-molecule self-construction). *Let Self_n be the event that some molecule x catalyses every channel of a nonempty food-generated set T with $x \in \text{cl}_T(F)$. Then, cleavage included,*

$$\Pr(\text{Self}_n) \leq |\mathcal{X}_{2t}| p_n + N_n R_{2t} R_{4t} p_n q_n^2 \longrightarrow 0.$$

Corollary 3.3 (negation of the linear-size assertion). *For every fixed food horizon t , the following statement is false: for every $\varepsilon > 0$ there exist $\lambda \geq 0$ and $C \geq 0$ such that $\Pr(L_n(C)) \geq 1 - \varepsilon$ for all large n .*

Corollary 3.4 (rate). *For fixed $t, \lambda \geq 0$ and integer $C \geq 0$ there is a constant $A < \infty$ such that, for all large n ,*

$$\Pr(L_n(C)) \leq \frac{A}{n} + Cn(f + 3Cn)\left(\frac{3}{4}\right)^n. \quad (6)$$

An admissible A is given in (15).

Let M_n be the minimum number of channels of a nonempty RAF, and J_n the minimum rank of an active reachable catalyst cover of a RAF; both are $+\infty$ when no RAF exists. Let G_n be the event that a RAF exists.

Corollary 3.5 (minima and conditioning). *For every fixed real C and integer K , $\Pr(M_n \leq Cn) \rightarrow 0$ and $\Pr(J_n \leq K) \rightarrow 0$. If $\liminf_n \Pr(G_n) > 0$, then also $\Pr(M_n \leq Cn \mid G_n) \rightarrow 0$ and $\Pr(J_n \leq K \mid G_n) \rightarrow 0$.*

Corollary 3.6 (superlinear-to-quadratic window). *In the regime of [3, Corollary 2], i.e. t fixed, $0 < \varepsilon < 1$, and λ fixed and large enough that the RAF probability bound of [3, Theorem 1] is at least $1 - \varepsilon$, for every fixed C*

$$\Pr(Cn < M_n < 2\lambda n^2) \geq 1 - \varepsilon - o(1).$$

The quadratic upper event in Corollary 3.6 is a cited input; only the subtraction of the vanishing linear event is proved here. No lower bound of order n^2 and no concentration of M_n at a particular intermediate scale is claimed; see Section 11.

4 Productive programs

4.1 Programs and saturated extraction

Fix n and t . A channel r is *productive* at an available set A if it is enabled at A and $\text{end}(r) \not\subseteq A$. Firing r at A yields $A \cup \text{end}(r)$; since one side of r was already in A , this is exactly the closure update in the enabled orientation. A *productive program* from A with allowed set U is a finite sequence $r_1, \dots, r_m$ of channels of U such that each r_j is productive at A_{j-1} , where $A_0 = A$ and $A_j = A_{j-1} \cup \text{end}(r_j)$; its *support* is $\{r_1, \dots, r_m\}$ and its *final state* is A_m .

Lemma 4.1. *In a productive program the channels $r_1, \dots, r_m$ are pairwise distinct, $|A_j| \geq |A| + j$, and $A_m = A \cup \bigcup_j \text{end}(r_j)$. Every molecule of A_m has length at most $L2^m$ if every molecule of A has length at most L .*

Proof. After r_j fires, all its endpoints are available, so r_j is never productive again. Each firing adds at least one molecule, and adds exactly the endpoints of the fired channel. A ligation at most doubles the maximum length present and a cleavage does not increase it. $\square$

Lemma 4.2 (saturated extraction). *For every finite $T \subseteq \mathcal{R}_n$ there is a productive program from F with support $S \subseteq T$ whose final state equals $\text{cl}_T(F)$. In particular $\text{cl}_T(F) = \text{U}(S)$ and $|\text{cl}_T(F)| \leq f + 3|S|$.*

Proof. Starting from F , fire any productive channel of T while one exists. The available set strictly increases inside the finite set $\mathcal{X}_n$, so the procedure stops, at a state B . Every molecule added along the way lies in $\text{cl}_T(F)$, so $B \subseteq \text{cl}_T(F)$. At B no channel of T is productive: every channel of T enabled at B has all its endpoints in B . Hence B is closed under both orientations of every channel of T and contains F , so $\text{cl}_T(F) \subseteq B$. The formula for B and the cardinality bound are Lemma 4.1. $\square$

Now let T be a RAF and choose for every channel one catalyst in $\text{cl}_T(F)$; let K be the set of chosen molecules. Every member of K is active and $K \neq \emptyset$. Apply Lemma 4.2 to T and let S be the extracted support. Since $\text{cl}_S(F) = \text{cl}_T(F) \supseteq K$, every channel of S is still covered by K and every member of K is still reachable from S . Two cases arise.

- If $S \neq \emptyset$, then S is itself a RAF with $|S| \leq |T|$: its productive program supplies food generation and its final state $U(S)$ supplies the catalysts.
- If $S = \emptyset$, then $\text{cl}_T(F) = F$, so a catalyst of any channel of the nonempty set T is an *active food molecule*. This event has probability at most $fp_n \rightarrow 0$.

The empty case is genuine: a RAF consisting only of channels whose endpoints all lie in the food has an empty productive extraction, and the event is not covered by any statement that requires a nonempty support.

4.2 Productive prefixes and their counts

A productive program of length d from F with allowed set $\mathcal{R}_n$ is a *productive prefix of depth d* . Let $\mathcal{P}_d$ be the set of productive prefixes of depth d and $\mathcal{H}_d$ the union of the final states of all prefixes of depth less than d , the *shallow universe*. Both depend on n , but their sizes do not.

Lemma 4.3 (ambient-independent prefix counts). *Put $b_j = (f+3j)^2 + (f+3j)t^{2j}$, $P_d = \prod_{j=0}^{d-1} b_j$ (so $P_0 = 1$) and $h_d = \sum_{j=0}^{d-1} P_j(f+3j)$. Then $|\mathcal{P}_d| \leq P_d$ and $|\mathcal{H}_d| \leq h_d$ for every n .*

Proof. By Lemma 4.1, the state after j steps has at most $f+3j$ molecules, each of length at most t^{2j} . A productive ligation at that state is determined by its two ordered factors, both available: at most $(f+3j)^2$ choices, because the ordered factor pair determines the channel. A productive cleavage is determined by its available product and its split position, which is less than the product's length: at most $(f+3j)t^{2j}$ choices, because the product and split determine the channel. Hence at most b_j channels are productive at any state reached after j steps, and $|\mathcal{P}_d| \leq \prod_{j<d} b_j$ by induction. The shallow universe is the union over $j < d$ of at most P_j final states of size at most $f+3j$. $\square$

The bounds do not depend on n because depth- d constructions only ever see molecules of length at most t^{2d} . This is the only place where the exponential-in-length structure of the polymer model enters the bounded-rank argument.

Lemma 4.4 (far catalysts force deep prefixes). *Let $U \subseteq \mathcal{R}_n$ and let $x \in \text{cl}_U(F)$ with $x \notin \mathcal{H}_d$. Then some productive prefix of depth d has all its channels in U .*

Proof. By induction on d . Suppose a productive program $r_1, \dots, r_{d-1}$ with channels in U and final state B is given. Since its length is less than d , $B \subseteq \mathcal{H}_d$, so $x \notin B$. If no channel of U were productive at B , then B would contain $\text{cl}_U(F) \ni x$ by the argument of Lemma 4.2, a contradiction. Fire one productive channel of U . $\square$

5 Bounded catalyst rank

Lemma 5.1 (active-cover bound). *Let K be a set of k molecules and S a set of m distinct channels. The probability that every member of K is active and every channel of S is catalysed by some member of K is at most $p^k k^m q^m$.*

Proof. For each map $a: S \rightarrow K$ consider the cylinder event that $A_x = 1$ for all $x \in K$ and $B_{a(r),r} = 1$ for all $r \in S$. The $k+m$ coordinates involved are distinct, because the channels are distinct, so the cylinder has probability exactly $p^k q^m$. The cover event is contained in the union of these at most k^m cylinders. $\square$

The lemma keeps the shared activity of a molecule as a single factor p regardless of how many channels it covers. Replacing the sparse law by independent marginals $(x, r) \in \mathbb{C}$ of probability pq would give a smaller and *incorrect* bound for a molecule that covers two channels: the correct mass pq^2 exceeds $(pq)^2$.

Lemma 5.2 (fixed-depth bound). *For every $k \geq 1$ and $d \geq 1$,*

$$\Pr(E_{n,k}) \leq h_d p + N_n^k P_d k^d p^k q^d. \quad (7)$$

Proof. Let T be a witnessing food-generated set with active reachable cover K , $|K| = k$. If K meets $\mathcal{H}_d$, some fixed molecule of $\mathcal{H}_d$ is active: probability at most $|\mathcal{H}_d|p \leq h_d p$ by the union bound. Otherwise pick $x \in K$. It lies in $\text{cl}_T(F) \setminus \mathcal{H}_d$, so by Lemma 4.4 some productive prefix $\rho \in \mathcal{P}_d$ has all its d distinct channels in T , and all of them are covered by K . The event that a fixed pair (K, ρ) is active and covered has probability at most $p^k k^d q^d$ by Lemma 5.1. There are at most $\binom{N_n}{k} \leq N_n^k$ choices of K and at most P_d choices of ρ . $\square$

Proof of Theorem 3.1. Fix $k \geq 1$ and take $d = k + 1$ in (7). With $q = 1/n$ the second term equals

$$\frac{P_{k+1} k^{k+1}}{n} \left(\frac{N_n p_n}{n} \right)^k, \quad (8)$$

and $N_n p_n / n \rightarrow \lambda$ by (5). The first term $h_{k+1} p_n$ is exponentially small up to a polynomial factor. Hence $\Pr(E_{n,k}) = O(1/n)$. A cover of rank at most K can be pruned to one of rank exactly k for some $1 \leq k \leq K$ without losing coverage, so the event of the second statement is contained in $\bigcup_{k=1}^K E_{n,k}$. $\square$

The point of the choice $d = k + 1$ is that each additional level of depth costs a fixed factor b_d but gains a factor $q = 1/n$, while each additional catalyst costs only the bounded quantity $N_n p_n / n$. A fixed number of catalysts can therefore always be beaten by depth. This argument alone, however, cannot reach ranks growing with n : the constant $P_{k+1} k^{k+1}$ grows superexponentially in k , so (8) is useless once k is of order $\log n / \log \log n$. The remaining ranks require a different idea.

6 The factorial support count and the large-rank tail

6.1 Acyclic dependency codes

Let $\mathcal{A}_m$ be the family of m -element channel sets that are supports of productive programs from F (with allowed set the support itself, equivalently $\mathcal{R}_n$). The following count is deterministic and makes no reference to catalysis.

Lemma 6.1 (factorial support count). *For $m \geq 1$,*

$$|\mathcal{A}_m| \cdot m! \leq B_m^m, \quad B_m = (f + 3m)^2 + n(f + 3m). \quad (9)$$

Proof. A *labelled assignment* is an injective map $a: \{1, \dots, m\} \rightarrow \mathcal{R}_n$; its image is a support. A *reference* is either a food molecule or a pair (j, e) with $j \in \{1, \dots, m\}$ a label and $e \in \{\text{left factor, right factor, product}\}$ an endpoint position; there are $f + 3m$ references, and relative to a a reference has a value in $\mathcal{X}_n$. A *node instruction* is either an ordered pair of references (a ligation instruction) or a reference together with a split position in $\{1, \dots, n\}$ (a cleavage instruction); there are B_m node instructions. A *code* is a list $c = (c_1, \dots, c_m)$ of node instructions, one per label; there are B_m^m codes.

Say that a *satisfies* c if there is a rank function $\varrho: \{1, \dots, m\} \rightarrow \mathbb{N}$ such that for every label i : every label referenced by c_i has rank less than $\varrho(i)$ (acyclicity), and c_i *matches* $a(i)$,

meaning that a ligation instruction's two references have values equal to the ordered factors of $a(i)$, and a cleavage instruction's reference has value equal to the product of $a(i)$ and its split position equals the split position of $a(i)$.

Existence. Let a be a labelled assignment whose image S lies in $\mathcal{A}_m$, and choose a productive program $r_1, \dots, r_m$ with support S ; let $\varrho(i)$ be the position of $a(i)$ in this program. When r_j fires, it is enabled in some orientation and every input molecule of that orientation is available, hence is food or an endpoint of an earlier channel $r_{j'}$, $j' < j$; reference it accordingly, choosing for $a(i) = r_j$ a ligation instruction if the ligation orientation was enabled and a cleavage instruction with the split of r_j otherwise. Every referenced label has smaller rank, and the instruction matches $a(i)$. So every such a satisfies at least one code.

Uniqueness. Suppose a and a' both satisfy c , with ϱ a rank certificate for a . We show $a(i) = a'(i)$ by strong induction on $\varrho(i)$. All references in c_i are food, which has the same value under both assignments, or endpoints of labels of smaller rank, whose channels agree by induction; so every reference in c_i has the same value under a and a' . If c_i is a ligation instruction, both $a(i)$ and $a'(i)$ have these ordered factors, and the ordered factor pair determines the channel. If c_i is a cleavage instruction, both have the same product and the same split position, which determine the channel. Only a 's certificate is used; a' need not have the same one.

Counting. Each support $S \in \mathcal{A}_m$ has exactly $m!$ labelled assignments with image S , and each of them satisfies some code. Assign to each labelled assignment one code it satisfies. By uniqueness this map is injective on the set of all labelled assignments with image in $\mathcal{A}_m$, which has $|\mathcal{A}_m| m!$ elements, into a set of B_m^m codes. $\square$

Example 6.2. Let $t = 2$, so $F = \{0, 1, 00, 01, 10, 11\}$, and consider the channels $r_a = (010, 2)$, $r_b = (01011, 3)$ and $r_c = (01011, 4)$, with orientations $01 + 0 \leftrightarrow 010$, $010 + 11 \leftrightarrow 01011$ and $0101 + 1 \leftrightarrow 01011$. The order r_a, r_b, r_c is a productive program: two ligations followed by a cleavage that produces the new molecule 0101. With labels $(1, 2, 3) = (a, b, c)$ a valid code is

$$c_1 = \text{lig}(\text{food } 01, \text{ food } 0), \quad c_2 = \text{lig}((1, \text{product}), \text{ food } 11), \quad c_3 = \text{clv}((2, \text{product}), 4).$$

With labels $(1, 2, 3) = (c, a, b)$ the same support has the code $c_1 = \text{clv}((3, \text{product}), 4)$, $c_2 = \text{lig}(\text{food } 01, \text{ food } 0)$, $c_3 = \text{lig}((2, \text{product}), \text{ food } 11)$, certified by the rank $\varrho(2) < \varrho(3) < \varrho(1)$. The support has a single productive order, yet six labelled codes. Note also that r_b and r_c share a product and an endpoint value 1 occurs both as food and as a factor of r_c ; neither coincidence affects decoding, since a node fixes either the ordered factors or the product and the split.

Lemma 6.3 (linear support entropy). *Let $C \geq 0$ be an integer and put $D = e(f + 3)((f + 3)C + 1)$. For $1 \leq m \leq Cn$,*

$$|\mathcal{A}_m| \leq (Dn)^m. \quad (10)$$

Proof. For $m \geq 1$ we have $f + 3m \leq (f + 3)m$ and $f + 3m + n \leq ((f + 3)C + 1)n$, so $B_m = (f + 3m)(f + 3m + n) \leq m(f + 3)((f + 3)C + 1)n$. Divide (9) by $m!$ and use $m^m/m! \leq e^m$. $\square$

Under a linear budget the factorial is therefore worth exactly a factor n^{-m} against the crude count $B_m^m \approx (mn)^m$. It is precisely this factor that the sparse law needs, as the next lemma shows.

6.2 The large-rank tail

Let $Q_{n,m,k}$ be the event that some productive support $S \in \mathcal{A}_m$ has an active catalyst cover of rank k contained in its universe $U(S)$. The universe is determined by S and has at most

$f + 3m$ molecules, so by Lemma 5.1 and the union bound over supports and over k -subsets of the universe,

$$\Pr(Q_{n,m,k}) \leq |\mathcal{A}_m| (f + 3m)^k k^m p^k q^m. \quad (11)$$

With $q = 1/n$, $m \leq Cn$ and (10) the factors n^m and n^{-m} cancel:

$$\Pr(Q_{n,m,k}) \leq (Dk)^m ((f + 3Cn) p_n)^k. \quad (12)$$

Lemma 6.4 (uniform large-rank bound). *There is an integer K_0 , depending only on t and C , such that for all large n , every $1 \leq m \leq Cn$ and every $k > K_0$,*

$$\Pr(Q_{n,m,k}) \leq \left(\frac{3}{4}\right)^{nk} \leq \left(\frac{3}{4}\right)^n. \quad (13)$$

Proof. Since $(Dk)^C$ is a fixed polynomial in k , there is K_0 with $(Dk)^C \leq (4/3)^k$ for all $k > K_0$. By (5), $(f + 3Cn)p_n \leq (9/16)^n$ for all large n . Because $Dk \geq 1$ and $m \leq Cn$, (12) gives

$$\Pr(Q_{n,m,k}) \leq (Dk)^{Cn} ((9/16)^n)^k \leq ((4/3)^k)^n ((9/16)^n)^k = \left(\frac{4}{3} \cdot \frac{9}{16}\right)^{nk} = \left(\frac{3}{4}\right)^{nk}. \quad \square$$

The cutoff K_0 does not depend on n , on m or on the support. Since a cover inside a universe of size at most $f + 3Cn$ has rank at most $f + 3Cn$, the union over all pairs (m, k) with $1 \leq m \leq Cn$ and $K_0 < k \leq f + 3Cn$ costs at most the polynomial factor $Cn(f + 3Cn)$.

7 Proof of the main theorem and its rate

Proof of Theorem 1.1. First let C be a nonnegative integer and let K_0 be as in Lemma 6.4. Suppose $L_n(C)$ occurs: there is a RAF T with $|T| \leq Cn$. Choose an active cover $K \subseteq \text{cl}_T(F)$ as in Section 4 and extract a saturated productive support $S \subseteq T$ by Lemma 4.2. Exactly one of the following holds.

- (a) $S = \emptyset$. Then some food molecule is active.
- (b) $S \neq \emptyset$ and $|K| \leq K_0$. Then $E_{n,k}$ occurs for $k = |K|$, because T is food-generated with active reachable cover K .
- (c) $S \neq \emptyset$ and $|K| > K_0$. Then $m = |S|$ satisfies $1 \leq m \leq Cn$, the cover K lies in $\text{cl}_T(F) = \text{U}(S)$ by Lemma 4.2, so $k = |K| \leq f + 3Cn$, and every channel of $S \subseteq T$ is covered by K : the event $Q_{n,m,k}$ occurs.

Hence, for all n large enough that Lemma 6.4 applies,

$$\Pr(L_n(C)) \leq f p_n + \sum_{k=1}^{K_0} \Pr(E_{n,k}) + Cn (f + 3Cn) \left(\frac{3}{4}\right)^n. \quad (14)$$

The first term tends to zero by (5), the finite sum by Theorem 3.1, and the last term is a polynomial times a decaying exponential. This proves (1) for integer $C \geq 0$. For an arbitrary real C choose an integer $C_0 \geq \max\{C, 0\}$; the event with budget Cn is contained in the event with budget $C_0 n$. (For $C < 0$ the event is empty.) The rate $O(1/n)$ is Corollary 3.4, proved next. $\square$

Proof of Corollary 3.4. For all large n we have $np_n \leq 1$ and $N_n p_n / n \leq \lambda + 1$ by (5). Multiplying (7) with $d = k + 1$ by n gives, for such n ,

$$n \Pr(E_{n,k}) \leq h_{k+1} + P_{k+1} k^{k+1} (\lambda + 1)^k =: a_k.$$

Also $nfp_n \leq f$. Substituting into (14),

$$\Pr(L_n(C)) \leq \frac{A}{n} + Cn(f + 3Cn)\left(\frac{3}{4}\right)^n, \quad A = f + \sum_{k=1}^{K_0} a_k, \quad (15)$$

and $n^3(3/4)^n \rightarrow 0$, so the right side is $O(1/n)$. Neither A nor the threshold beyond which (15) holds has been optimised; the corollary establishes a rate, not a usable finite- n estimate. $\square$

Proof of Corollary 3.3. If the statement held, take $\varepsilon = 1/2$: some fixed λ, C would give $\Pr(L_n(C)) \geq 1/2$ for all large n , contradicting (1) for those parameters. $\square$

The two parts of (14) play complementary roles. The productive-prefix estimate controls a fixed number of catalysts without any restriction on the size of the RAF. The factorial count controls the entropy of a linear-size construction once the number of catalysts is large, using only that active molecules are exponentially sparse in $\mathcal{X}_n$ while the universe of a linear construction is only linear. Neither estimate alone gives the all-rank statement: the first fails at rank $\gg \log n$, and the second fails at bounded rank because $(Dk)^m$ is then not compensated by the activity factor.

8 A single molecule cannot construct itself

Theorem 3.1 with $K = 1$ already shows that a RAF catalysed by a single reachable molecule has vanishing probability. The question of [3] was posed for the more elementary self-construction event of Proposition 3.2, and for that event there is a direct argument with explicit constants that also exhibits why cleavage does not help. The argument uses only the lengths of molecules.

Lemma 8.1 (two short channels). *Let $T \subseteq \mathcal{R}_n$ be nonempty and food-generated, and let $x \in \text{cl}_T(F)$ have length greater than $2t$. Then T contains two distinct channels $r_0 \neq r_1$ with product lengths $|r_0| \leq 2t$ and $|r_1| \leq 4t$.*

Proof. Since T is nonempty and food-generated, some channel $r_0 \in T$ is enabled at F itself: otherwise the closure would never leave F , contradicting food generation of a nonempty set. Its factors or its product lie in F ; in either case its product has length at most $2t$. Suppose, for contradiction, that every channel of T with product length at most $4t$ equals r_0 . We claim that every molecule of $\text{cl}_T(F)$ has length at most $2t$, contradicting the assumption on x . Induct along the closure sequence. A ligation orientation enabled at a state of molecules of length at most $2t$ has a product of length at most $4t$, so the channel is r_0 and its product has length at most $2t$. A cleavage orientation enabled at such a state produces factors shorter than a product of length at most $2t$. Either way the new molecules have length at most $2t$. $\square$

Proof of Proposition 3.2. Suppose x catalyses every channel of a nonempty food-generated T with $x \in \text{cl}_T(F)$. In particular x is active. If $|x| \leq 2t$, then one of the $|\mathcal{X}_{2t}|$ short molecules is active, an event of probability at most $|\mathcal{X}_{2t}|p$. If $|x| > 2t$, Lemma 8.1 gives two distinct channels $r_0 \in \mathcal{R}_{2t}$, $r_1 \in \mathcal{R}_{4t}$, both catalysed by x ; the event $A_x = B_{x,r_0} = B_{x,r_1} = 1$ involves three distinct coordinates and has probability exactly pq^2 . Summing over $x \in \mathcal{X}_n$ and over the at most $R_{2t}R_{4t}$ ordered pairs of short channels yields the bound. The right side tends to zero because $N_n p_n q_n^2 = (N_n p_n / n) / n \rightarrow 0$ by (5), and $p_n \rightarrow 0$. $\square$

The bound is $O(1/n)$ with an explicit constant $\lambda R_{2t} R_{4t}$ in the leading term. The restriction to ligations mentioned in [3] is not needed: cleavage cannot enlarge the length frontier, so a molecule longer than $2t$ requires two distinct short channels whichever orientations are used, and two channels catalysed by one molecule cost $q^2 = n^{-2}$ against only $N_n p_n \approx \lambda n$ candidate molecules.

Catalysis law	Size of the smallest RAF in the stated regime
Uniform	Sub-exponential RAFs are unlikely at the emergence transition [3, 5].
Sparse	Every fixed linear budget vanishes, at rate $O(1/n)$, for every fixed intensity (Theorem 1.1); quadratic witnesses exist at sufficiently large fixed intensity [3, Corollary 2]; hence the window of Corollary 3.6.
All-or-nothing	Linear constructions exist whenever a RAF exists [3].
Power law	No size theorem is asserted here; [3, 6] identify the question.

Table 1: Distribution-sensitive size of the smallest RAF in the binary polymer model. The sparse row combines the new lower-scale result with a published upper-scale result.

9 Consequences

Proof of Corollary 3.5. Because the model is finite, M_n and J_n are attained when finite. The event $\{M_n \leq Cn\}$ is $L_n(C)$, and $\{J_n \leq K\}$ is the union over $1 \leq k \leq K$ of the events $E_{n,k}$ after deleting unused cover members; Theorems 1.1 and 3.1 apply. Both events are contained in G_n , so their conditional probabilities given G_n are the unconditional probabilities divided by $\Pr(G_n)$. If $\liminf \Pr(G_n) \geq \delta > 0$ then $\Pr(G_n) \geq \delta/2$ for all large n and each quotient is at most $2/\delta$ times a quantity tending to zero. The same argument applies to any conditioning event of eventually positive probability, by bounding the intersection by the unconditioned event. $\square$

Corollary 3.5 says that $M_n/n \rightarrow \infty$ and $J_n \rightarrow \infty$ in probability, with $+\infty$ as specified on RAF-free instances. In words: whenever the sparse model has a RAF at all, its smallest RAF uses superlinearly many channels and unboundedly many distinct catalysts.

Proof of Corollary 3.6. By [3, Corollary 2], in the stated regime $\Pr(M_n < 2\lambda n^2) \geq 1 - \varepsilon - \eta_n$ with $\eta_n \rightarrow 0$. Removing the event $\{M_n \leq Cn\}$ costs at most its probability, so

$$\Pr(Cn < M_n < 2\lambda n^2) \geq 1 - \varepsilon - \eta_n - \Pr(L_n(C)),$$

and $\Pr(L_n(C)) \rightarrow 0$ by Theorem 1.1. $\square$

Corollary 3.6 also shows that the conditioning in Corollary 3.5 is not vacuous: in the regime of [3] RAF existence has probability bounded away from zero, and the smallest RAF is then, with probability at least $1 - \varepsilon - o(1)$, strictly superlinear and strictly subquadratic. The window is compatible with, for instance, a minimum of order $n \log n$ or of order $n^{3/2}$; identifying the true scale is open (Section 11).

Table 1 places the result next to the size statements known for the other catalysis laws of [3]. The rows describe their respective regimes and do not share a numerical threshold.

10 Formal verification

All results of Sections 3–9 have been formalised in Lean 4 (toolchain v4.30.0) with a pinned `mathlib` revision [1, 12]. The development lives in the directory `proofs/SparseLinearRAF/` of the accompanying repository, on top of a shared formalisation of the binary polymer source; Table 2 lists the concordance between the paper and the exported declarations.

What is formalised. Molecules are the dependent type $\Sigma_{k < n} \text{Fin}(2^{k+1})$ of bit strings of length $k + 1 \leq n$, and channels the type $\Sigma_{k < n} \text{Fin}(2^{k+1}) \times \text{Fin}(k)$ of a product word with a split position. The counts (3) are proved exactly. The reversible closure, food generation and RAF are the definitions of Section 2. The sparse law is the infinite product measure on

$\{0, 1\}^{\mathcal{X}_n \times (\{*\} \sqcup \mathcal{R}_n)}$ with the activity coordinate of each molecule and its conditional channel coordinates, and catalysis is the conjunction of the two coordinates; cylinder probabilities are computed from the independence of the coordinates. The parameters are those of (4), with the activity parameter clipped to $[0, 1]$ and a lemma showing that the clip is inactive for all large n . Productive programs are an inductive predicate on lists of channels; the saturated extraction, the prefix family and its ambient-independent bounds, the acyclic dependency codes with their uniqueness theorem, and the factorial count are all proved for the literal source, including repeated endpoint values and coinciding factor sets of different split positions. The two probability estimates, the decomposition of the linear-RAF event into the three cases of Section 7, and the limits are then assembled into the root declaration

```
SparseLinearRAF.sparse_linear_raf_literature_resolution
(t : Nat) (C : Real) {lambda : Real} (hl : 0 <= lambda) :
Tendsto (fun n => literatureRAFProbability n t lambda C)
atTop (nhds 0),
```

where `literatureRAFProbability n t lambda C` is the sparse-law probability of the event $L_n(C)$ with an arbitrary real budget constant, and into the explicit negation `no_high_probability_linear_raf` of Corollary 3.3.

What is trusted and what is cited. The receipt of the final audit target records that all twelve exported declarations of Table 2 compile with warnings treated as errors, that their axiom dependencies are exactly `propext`, `Classical.choice` and `Quot.sound`, and that no `sorry` and no additional mathematical axiom is present. The trusted base is the Lean 4 kernel and the pinned `mathlib` revision. Two items are *not* formalised and are cited from [3]: the positive RAF-probability bound of its Theorem 1 and the quadratic witness of its Corollary 2. The formal version of Corollary 3.6 (`size_window_lower_bound`) is the abstract scalar combination: from an eventual lower bound $1 - \varepsilon - \eta_n$ on the upper event, a vanishing lower event and a vanishing η_n , it produces the eventual bound on the window event. The conditional statements accept an eventual positive lower bound on the conditioning probability as a hypothesis. No numerical experiment is a premise of any theorem.

A repair found by formalisation. The formalisation exposed one gap in an earlier informal version of the construction condition. A first draft characterised bounded-size RAFs by the existence of a catalyst set K and a food-generated channel set covered by K whose closure contains K , without requiring the channel set to be nonempty. An active food molecule with the empty channel set satisfies that condition but supplies no RAF, so the equivalence was false. The compiled equivalence (`bounded_raf_iff_build` in `CatalystRank`) requires a nonempty support, and the empty extraction case of Section 4 is handled separately as an active-food event.

11 Discussion

What the theorem says. In the sparse polymer model the probability that a RAF exists and the size of the smallest RAF behave very differently from the uniform model, but in the opposite direction from what the quadratic witness of [3] might suggest. A RAF built from a fixed number of catalysts is excluded because a catalyst that is not food must be reached by a productive construction whose depth costs a factor $1/n$ per step; a RAF built from many catalysts within a linear budget is excluded because the entropy of linear constructions, once divided by $m!$, is too small to compensate for the exponential rarity of active molecules. Both mechanisms are specific to the sparse law: the first uses the factor $q = 1/n$ per catalysed channel, the second uses that active molecules are $\approx \lambda n$ among $\approx 2^{n+1}$.

Paper statement	Lean declaration (module)
Theorem 1.1, limit	<code>sparse_linear_raf_literature_resolution</code> (<code>LiteratureResolution</code>)
Corollary 3.3	<code>no_high_probability_linear_raf</code> (<code>LiteratureResolution</code>)
Lemma 4.2	<code>saturated_support_extraction</code> (<code>SaturatedProgram</code>)
Lemma 4.3	<code>source_prefixes_card_le</code> , <code>source_shallow_card_le</code> (<code>PrefixCounting</code>)
Lemma 5.1	<code>active_cover_probability_le</code> (<code>CoverProbability</code>)
Lemma 5.2	<code>cover_rank_probability_le</code> (<code>LowRankBound</code>)
Theorem 3.1	<code>bounded_cover_rank_probability_tendsto_zero</code> (<code>BoundedRank</code>)
Lemma 6.1, uniqueness	<code>validAssembly_unique</code> (<code>AssemblyCode</code>)
Lemma 6.1, count	<code>support_count_mul_factorial_le</code> (<code>SupportCounting</code>)
Lemma 6.3	<code>linear_support_count_le</code> (<code>LinearSupportCount</code>)
Lemma 6.4	<code>large_rank_uniform_estimate</code> (<code>LargeRankEstimate</code>)
Case decomposition, Section 7	<code>linear_raf_decomposition</code> (<code>RAFReduction</code>)
Corollary 3.4	<code>linear_raf_quantitative_bound</code> (<code>Quantitative</code>)
Proposition 3.2	<code>single_catalyst_probability_tendsto_zero</code> (<code>SelfConstructionLimit</code>)
Lemma 8.1	<code>two_distinct_short_channels</code> (<code>TwoFrontier</code>)
Corollary 3.5, minima	<code>minimum_raf_size_superlinear</code> , <code>minimum_catalyst_rank_diverges</code> (<code>PublicationConsequences</code>)
Corollary 3.5, conditional	<code>linear_raf_conditional_vanishes</code> , <code>bounded_rank_conditional_vanishes</code> (same)
Corollary 3.6, combination	<code>size_window_lower_bound</code> (same)

Table 2: Concordance between the paper and the Lean development. All declarations are in the namespace `SparseLinearRAF`; module names are relative to `proofs/SparseLinearRAF/`.

What it does not say. The theorem is asymptotic in n with t , λ and C fixed. It does not identify the growth scale of M_n between n and n^2 ; the window of Corollary 3.6 is all that is proved. It does not cover an intensity λ_n growing with n , although the two scalar inputs of the large-rank branch, the activity decay $(f + 3Cn)p_n \leq (9/16)^n$ and the rank-mass identity behind (8), survive any fixed polynomial envelope $\lambda_n \leq Ln^a$ with $d = (a + 1)k + 1$, so we expect the conclusion to persist there. It is proved for the binary alphabet; the saturation lemma and the encoding are alphabet-independent in substance, but the formal source layer is binary. Finally, the constants A and K_0 are not optimised, and no claim is made about moderate values of n such as those of the simulations in [3].

Open problems.

1. Determine the order of M_n in the sparse model at fixed large λ . Is it $n \log n$, a power $n^{1+\alpha}$, or n^2 up to constants?
2. Prove the analogue of Theorem 1.1 for the power-law catalysis of [3, 6]. The factorial count of Lemma 6.1 is law-independent and applies verbatim; what is missing is an activity estimate replacing (5).
3. Extend the theorem to growing intensity λ_n and to κ -letter alphabets, as discussed above.
4. The sparse law is a two-level mixture. Which distribution-based laws in the sense of [3] admit linear-size RAFs at the emergence threshold, and which do not? The all-or-nothing law does; the sparse law does not.

Data and code availability

The Lean development (`proofs/SparseLinearRAF/`), the strict verification receipts with declaration interfaces and axiom profiles, the pinned toolchain and `mathlib` manifest, and the exact enumeration scripts used during the discovery phase (small exact checks of the encoding and of the parameter inequalities, none of which is a premise of any theorem) are distributed with the accompanying repository.

References

- [1] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CCADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi:[10.1007/978-3-030-79876-5_37](https://doi.org/10.1007/978-3-030-79876-5_37).
- [2] 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](https://doi.org/10.1016/j.jtbi.2003.11.020).
- [3] 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](https://doi.org/10.1007/s10910-016-0666-z).
- [4] Wim Hordijk and Mike Steel. Chasing the tail: The emergence of autocatalytic networks. *Biosystems*, 152:1–10, 2017. doi:[10.1016/j.biosystems.2016.12.002](https://doi.org/10.1016/j.biosystems.2016.12.002).
- [5] 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](https://doi.org/10.3390/ijms12053085).
- [6] Wim Hordijk, Leonard Hasenclever, Jie Gao, Dilyana Mincheva, and Jotun Hein. An investigation into irreducible autocatalytic sets and power law distributed catalysis. *Natural Computing*, 13(3):287–296, 2014. doi:[10.1007/s11047-014-9429-6](https://doi.org/10.1007/s11047-014-9429-6).
- [7] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi:[10.1016/S0022-5193\(86\)80047-9](https://doi.org/10.1016/S0022-5193(86)80047-9).
- [8] Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York, 1993.
- [9] 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](https://doi.org/10.1016/j.jtbi.2004.10.011).
- [10] 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](https://doi.org/10.1016/S0893-9659(99)00191-3).
- [11] 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](https://doi.org/10.1098/rsif.2018.0808).
- [12] 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](https://doi.org/10.1145/3372885.3373824).