

Positive Recurrence for Single-Linkage Bimolecular Weakly Reversible Stochastic Reaction Networks

Alec Kriebel

Independent researcher; corresponding author: me@aleckriebel.com

ORCID: 0009-0001-9320-500X

22 August 2026

Abstract

Stochastic reaction networks model populations changed by random reaction events. A central stability question is whether they admit a stationary law rather than escape to infinity. Under weak reversibility, we prove that the states reachable from each initial population form a closed communicating class. If the finite network also has one linkage class and molecularity at most two, then every nonabsorbing reachable class is positive recurrent for every positive rate vector; an absorbing singleton has its point-mass law. Thus every initial population has a unique stationary law on its reachable class. This removes the previous theorem’s pure-species condition. Multiple linkage classes and higher molecularity remain open.

Keywords. Stochastic reaction network; systems biology; chemical master equation; positive recurrence; mass-action kinetics; weak reversibility; continuous-time Markov chain; Foster–Lyapunov method.

2020 Mathematics Subject Classification. Primary 60J27; secondary 60J28, 60J74, 92C42.

1 Introduction

Stochastic reaction networks are continuous-time Markov-chain models for interacting populations whose transition rates are determined by a finite reaction graph. In systems biology they describe random molecule counts in low-copy-number intracellular networks, including signalling and enzymatic systems, gene-regulatory circuits, and synthetic biochemical networks (Anderson and Kurtz, 2015). Such models are often called chemical-master-equation models. For a nonexplosive nonabsorbing irreducible component, a basic stability question is whether it has a stationary probability distribution, or equivalently whether its states have finite mean return times. A stationary law gives a long-run molecule-count distribution and expectations of bounded observables within that component; moments of unbounded count observables require separate integrability. Positive recurrence does not imply that sample paths remain in a bounded set. The result below gives a structural, rate-independent stability criterion for the low-copy-number biochemical models that satisfy its graph hypotheses.

The deterministic counterpart supplies additional motivation. For strictly positive initial conditions, weakly reversible single-linkage mass-action systems have bounded trajectories (Anderson, 2011a) and are permanent (Boros and Hofbauer, 2020) under their deterministic dynamics. Those results express a strong structural stability of concentrations, but deterministic boundedness or permanence does not by itself prove recurrence of the discrete stochastic model. Positive recurrence is the corresponding class-wise stochastic question addressed here.

Anderson and Kim (2018) formulated the positive-recurrence conjecture that weak reversibility should force this stability for stochastic mass-action systems. They established several structural sufficient conditions. Anderson et al. (2020a) subsequently proved the binary, one-linkage case under the additional condition

$$\{S_i, 2S_i\} \cap \mathcal{C} \neq \emptyset \quad \text{for every species } S_i. \quad (1)$$

Under binaryity, this is equivalent to the authors’ summary that each species has a complex that is a positive multiple of it: each species had to occur in a pure unary or pure double complex. The proof of Anderson et al. (2020a) applies an entropy-like function to an n -step embedded chain. In their published Section 6, a tier criterion, a finite reaction word, and a sampled-chain Foster argument are combined to prove positive recurrence. In the boundary case of Section 6.1, a D-top mixed complex $S_u + S_v$ is disabled because the u -coordinate is zero while the v -coordinate diverges. Their pure-species assumption forces S_v to be a complex after $2S_v$ is excluded by D-tier maximality; S_v then supplies the needed source-rate comparison (Anderson et al., 2020a, Section 6, especially Section 6.1). The marked-target construction here avoids the corresponding boundary-availability step: the carried target is automatically enabled at the post-jump state, and a directed target-following path propagates its reward until the top-complex alternative supplies a rare terminal source. It therefore does not require a unary or pure-double complex for every species.

Other progress includes complex-balanced product-form systems (Anderson et al., 2010), strongly endotactic structures (Anderson et al., 2020b), one-dimensional classifications (Wiuf and Xu, 2023), a path method for exponential ergodicity (Anderson et al., 2025), and nonexplosion of every bimolecular weakly reversible stochastic mass-action system (Xu, 2024). A complementary two-species positive-recurrence theorem has been publicly announced in conference talks since 2022 (Agazzi, 2022; Cappelletti, 2022; Agazzi, 2025). The current ConStRAINeD project page describes the proof as complete and sketches its approach, while listing the five-author manuscript as in preparation (ConStRAINeD project, 2026). We did not locate a public manuscript for that result as of 22 August 2026. The public descriptions concern two species and permit broader network structure; the present theorem permits arbitrarily many species but assumes one linkage class and molecularity at most two. Neither result should be presented as superseding the other without the announced manuscript. The May 2026 revision of Xu (2024) still records the bimolecular positive-recurrence problem as open.

The present paper removes (1) while retaining the binary and one-linkage assumptions. Its theorem contains the positive-recurrence conclusion of Anderson et al. (2020a) as a special case: the additional condition (1) is unnecessary here, and the states reachable from any initial population form one closed communicating class. The exact theorem is stated in Theorem 2.1. It does not address multiple linkage classes or complexes of molecularity above two, and therefore does not settle the full positive-recurrence conjecture.

A small motivating example is the directed cycle



Species A occurs in neither A nor $2A$, so (1) fails. At the boundary state $(n, 0)$, only the reaction $0 \rightarrow A + B$ is enabled and the one-step drift of total population is positive. The restoring effect is delayed: the created B enables a fast reaction consuming A . This is representative of why a one-step radial Lyapunov function is not the right universal object.

The delayed correction can be seen directly. Give the three channels in (2) rate constants $\kappa_0, \kappa_1, \kappa_2$, in the order displayed, and begin at the reachable marked state $((n, 0), 0)$. The two designated target-following reactions $0 \rightarrow A + B$ and $A + B \rightarrow B$ each have zero residual-potential increment. The second is selected with probability

$$\frac{\kappa_1(n+1)}{\kappa_0 + \kappa_2 + \kappa_1(n+1)} = 1 - O(n^{-1}).$$

At the resulting marked state $((n, 1), B)$, the expected next increment is

$$-\frac{\kappa_1 n}{\kappa_0 + \kappa_2 + \kappa_1 n} \log n = -\log n + o(\log n). \quad (3)$$

The exceptional branches at the preceding phase have probability $O(n^{-1})$ and increments of order at most $\log n$. Thus the complete target-following episode has a restoring drift of order $-\log n$, even though the first boundary jump increases the total population.

The proof instead uses four ideas. First, we augment the embedded jump chain by marking the actual reaction channel that fired and retaining its target complex. Second, after removing that target from the

population, the log-factorial potential has the exact increment

$$\log \frac{(x)_t}{(x)_s}$$

when the carried target is t and the next source is s . Following the carried target therefore costs exactly zero. Third, a finite episode follows a fixed directed path from the carried target toward a selected terminal complex. A sharp scalar envelope shows that a terminal source probability tending to zero forces the complete episode drift to tend to $-\infty$, no matter how the intermediate source rates are separated. Fourth, a normalized-log compactification of an arbitrary divergent sequence gives an exhaustive top-complex case split. Bimolecularity implies either the existence of a higher-weight source enabled over a lower terminal complex or a reaction-wise linear invariant that contradicts divergence inside a fixed communicating class. A finite-family random-time drift argument and a finite trace-chain construction complete the proof.

2 Reaction networks and the stochastic model

Let $\mathcal{S} = \{S_1, \dots, S_d\}$ be a finite species set. A *complex* is a vector $y = (y_1, \dots, y_d) \in \mathbb{N}_0^d$, identified with the formal sum $\sum_i y_i S_i$. Its molecularity is $|y| = \sum_i y_i$. A *reaction channel* is a labelled ordered pair $r : y_r \rightarrow y'_r$, together with a positive rate constant κ_r . Parallel channels are allowed. Let $\mathcal{C}$ be the finite set of complexes appearing as a source or target of at least one reaction channel. The reaction graph has vertex set $\mathcal{C}$ and the reaction channels as directed edges. Its linkage classes are the connected components of the underlying undirected graph. A network is *weakly reversible* when each linkage class is strongly connected. With one linkage class, weak reversibility is simply strong connectivity of the entire complex graph.

We call the network *binary* or *bimolecular* when

$$|y| \leq 2 \quad \text{for every } y \in \mathcal{C}. \quad (4)$$

Because every complex in a weakly reversible graph is a reaction source, this agrees here with the usual second-order/bimolecular terminology.

For $x \in \mathbb{N}_0^d$, define the multivariate falling factorial

$$(x)_y = \begin{cases} \prod_{i=1}^d \frac{x_i!}{(x_i - y_i)!}, & x \geq y, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

The stochastic mass-action propensity of channel r is $\lambda_r(x) = \kappa_r(x)_{y_r}$, and the formal generator is

$$\mathcal{L}f(x) = \sum_r \kappa_r(x)_{y_r} [f(x + y'_r - y_r) - f(x)]. \quad (6)$$

The convention using products of binomial coefficients is equivalent after rescaling each positive rate constant by a fixed factorial factor.

Write $x \rightsquigarrow z$ when a finite sequence of enabled reaction channels takes population x to population z , allowing the empty sequence, and put

$$\Gamma(x_0) = \{z \in \mathbb{N}_0^d : x_0 \rightsquigarrow z\}.$$

A *communicating class* is an equivalence class under mutual accessibility: x and z are equivalent when both $x \rightsquigarrow z$ and $z \rightsquigarrow x$. It is *closed* when no enabled channel from one of its states exits the class. In a nonabsorbing class, let τ_1 be the first jump time and define the positive return time to x by $T_x^+ = \inf\{t \geq \tau_1 : X(t) = x\}$. The class is *positive recurrent* when one, equivalently every, state has $\mathbb{E}_x T_x^+ < \infty$. Absorbing singleton classes are handled separately: their stationary probability is the point mass at the absorbing state. Henceforth positive return times are discussed only for nonabsorbing irreducible classes. For a nonexplosive irreducible countable-state CTMC, positive recurrence is equivalent to the existence of a stationary probability distribution (Norris, 1997).

Theorem 2.1 (Binary one-linkage positive recurrence). *Let a finite stochastic mass-action reaction network be weakly reversible, have one linkage class, and satisfy (4). For every positive rate vector and every initial population x_0 , the reachable set $\Gamma(x_0)$ is a closed communicating class. If it is nonabsorbing, the minimal CTMC on $\Gamma(x_0)$ is nonexplosive and positive recurrent. If it is an absorbing singleton, its stationary law is the point mass.*

Lemma 2.2 (Lifted state-return cycle). *In any weakly reversible reaction network, every enabled one-step population transition has a return path. Consequently accessibility is symmetric, every reachability set is a communicating class, every communicating class is closed, and every initial state already belongs to the closed irreducible component $\Gamma(x_0)$. These conclusions require neither one linkage class nor bimolecularity.*

Proof. Suppose the channel $y \rightarrow y'$ is enabled at x . Write $x = \rho + y$ with $\rho \in \mathbb{N}_0^d$, so firing it gives $x' = \rho + y'$. Weak reversibility supplies a directed complex path

$$y' = z_0 \longrightarrow z_1 \longrightarrow \cdots \longrightarrow z_m = y.$$

At population $\rho + z_j$, the source z_j is present, so the corresponding channel is enabled and produces $\rho + z_{j+1}$. The path therefore lifts to

$$x' = \rho + z_0 \longrightarrow \rho + z_1 \longrightarrow \cdots \longrightarrow \rho + z_m = x.$$

Thus every enabled edge can be reversed by a finite enabled path. Reversing the edges of any reachability path in reverse order proves $x \rightsquigarrow z \Rightarrow z \rightsquigarrow x$. Hence all states in $\Gamma(x_0)$ communicate. An enabled edge leaving a state reachable from x_0 has a reachable endpoint by definition, so $\Gamma(x_0)$ is closed. By symmetry it is exactly the communication class of x_0 , proving that every communicating class is closed. This property was recorded as “recurrence” of a discrete reaction network by Paulevé et al. (2014, Lemmas 4.5–4.6); that combinatorial term means symmetric reachability, not positive recurrence of a Markov chain. The lifted proof is included for the exact population-state interface used here. $\square$

The proof includes the zero complex and boundary states because the common residual ρ is merely nonnegative. It also applies channel by channel to parallel reactions or identical population displacements, and every lifted step respects any parity or lattice restriction. A self-channel creates no population edge and has the empty return path. If no population-changing channel is enabled, the reachable set is the absorbing singleton $\Gamma(x_0) = \{x_0\}$.

Corollary 2.3 (Stationary probability). *Under the hypotheses of Theorem 2.1, every reachable class $\Gamma(x_0)$ admits a unique stationary probability distribution.*

The corollary is proved by the regenerative occupation formula at the end of Section 7.

A channel $y \rightarrow y$ has zero population increment and contributes zero to (6); discard it. This deletion leaves the generator and the minimal population-valued CTMC unchanged. Exact parallel channels with the same source and target may be combined by adding their rates. Channels having the same population displacement but different sources or targets remain distinct, because the proof marks the actual channel. All reaction-indexed sums and rate quantities below, including Λ and $\bar{\kappa}_y$, refer to this reduced set of genuine channels.

Every finite nonabsorbing reachable class is positive recurrent: its finite irreducible jump chain has a stationary probability distribution, all state rates are finite, and the associated finite-state CTMC is nonexplosive (Norris, 1997). If the complex graph has one vertex, all closed classes are absorbing singletons. If all complexes have the same molecularity, every reaction preserves $|x|_1$, so each communicating class lies in a finite shell. We therefore fix henceforth an initial state whose reachable class $\Gamma = \Gamma(x_0)$ is infinite. By Lemma 2.2, it is closed and irreducible. The complex graph is strongly connected, has at least two complexes, and has no self channels. Every complex then has a genuine outgoing channel. Write

$$\bar{\kappa}_y = \sum_{r: y_r = y} \kappa_r, \quad \bar{\kappa}_- = \min_{y \in \mathcal{C}} \bar{\kappa}_y > 0, \quad \bar{\kappa}_+ = \max_{y \in \mathcal{C}} \bar{\kappa}_y. \quad (7)$$

3 The marked target-augmented chain

We first define the embedded chain without assuming nonexplosion. At a nonabsorbing population state x , sample the actual labelled reaction channel r with probability

$$\frac{\kappa_r(x)_{y_r}}{\Lambda(x)}, \quad \Lambda(x) = \sum_q \kappa_q(x)_{y_q}. \quad (8)$$

This discrete chain can be constructed independently of its holding times. After the sampled channel r fires, record its target $t = y'_r$. The mark is therefore obtained from the actual channel and is never inferred from the population displacement. Once exact parallel channels with the same source and target have been aggregated, retaining the target is the only part of the previous channel needed below. Write Z_n for this target-augmented embedded chain. A reachable augmented state is denoted (x, t) , where x is the post-jump population. If the next sampled channel is $q : s \rightarrow u$, then

$$(x, t) \mapsto (x - s + u, u). \quad (9)$$

The channel law depends on x , not on the historical mark, so the augmented process is Markov, and projection onto x is the ordinary population embedded chain.

Let $\tilde{\Gamma}$ be the augmented states reachable after genuine jumps within Γ . It is closed under (9). It is also irreducible, and its projection onto the population coordinate is exactly Γ . For the latter assertion, every $x \in \Gamma$ terminates a positive-length population path: irreducibility gives such a path directly, or, if necessary, one may take a genuine outgoing jump and a return path. The final channel supplies a reachable mark over x . For irreducibility, given $(x, t), (x', t') \in \tilde{\Gamma}$, choose a channel $s' \rightarrow t'$ and predecessor population $z = x' - t' + s' \in \Gamma$ witnessing reachability of the second pair. Population irreducibility supplies a positive-probability marked-channel path from x to z ; appending $s' \rightarrow t'$ reaches (x', t') .

The carried target is always enabled. If the preceding channel was $s \rightarrow t$, then its post-jump state has the form $x = z - s + t \geq t$. Put

$$\rho = x - t \in \mathbb{N}_0^d, \quad V(x, t) = \sum_{i=1}^d \log((x_i - t_i)!) = \sum_i \log(\rho_i!). \quad (10)$$

We use $\log(0!) = 0$. Thus $V \geq 0$. Every sublevel set of V in $\tilde{\Gamma}$ is finite, because $\log(n!) \rightarrow \infty$, each residual coordinate is bounded on a sublevel, and the target set is finite.

The entropy growth in (10) is classical. With $0 \log 0 = 0$, Stirling's formula gives


$$\sum_i \log(\rho_i!) = \sum_i (\rho_i \log \rho_i - \rho_i) + O\left(\sum_i \log(\rho_i + 1)\right). \quad (11)$$

This leading term is in the pseudo-Helmholtz or Horn–Jackson family

$$G_c(\rho) = \sum_i [\rho_i (\log(\rho_i/c_i) - 1) + c_i], \quad c_i > 0,$$

used in deterministic and stochastic chemical reaction network theory (Horn and Jackson, 1972; Anderson et al., 2015). Thus the residual log-factorial potential is a discrete, target-shifted analogue of that classical entropy, not the Horn–Jackson function itself. The target shift is the new feature in this argument: subtracting the complex actually produced by the preceding channel gives the exact identity below, rather than only an asymptotic Lyapunov estimate, and makes every target-following jump have exactly zero increment.

Remark 3.1 (The Anderson–Cappelletti–Kim boundary example). The target shift can be isolated on Example 4.1 of Anderson et al. (2020a),



For the entropy-like function used there,

$$V_{\text{ACK}}(x) = \sum_{i=1}^d [x_i(\log x_i - 1) + 1], \quad 0 \log 0 := 0,$$

the generator drift at $x_n = (n, 1, 0)$, $n \geq 2$, is $\kappa_1 n(2 \log 2 - 1) \rightarrow +\infty$. For the marked construction, choose the reachable mark A : starting at x_n , the first four channels reach $(n-1, 3, 0)$, and the last returns to x_n with target A . Follow the fixed path $A \rightarrow A+B \rightarrow A+C \rightarrow C$, then take the terminal ordinary jump. The immediate expected reward at (x_n, A) is zero, so its complete episode reward is exactly

$$J_n = p_0\{\delta_1 + p_1(\delta_2 + p_2\delta_3)\}, \quad (12)$$

where

$$\begin{aligned} p_0 &= \frac{\kappa_1}{\kappa_1 + \kappa_2}, \quad L_1 = n(\kappa_1 + 2\kappa_2) + 2\kappa_5, \quad p_1 = \frac{2\kappa_2 n}{L_1}, \\ \delta_1 &= \frac{\kappa_1 n}{L_1} \log 2 + \frac{2\kappa_5}{L_1} \log n, \\ L_2 &= n(\kappa_1 + \kappa_2 + \kappa_3) + \kappa_4, \quad p_2 = \frac{\kappa_3 n}{L_2}, \quad \delta_2 = \frac{\kappa_4}{L_2} \log n, \\ L_3 &= (n-1)(\kappa_1 + \kappa_2 + \kappa_3) + \kappa_4, \quad \delta_3 = -\frac{(n-1)(\kappa_1 + \kappa_2 + \kappa_3)}{L_3} \log(n-1). \end{aligned}$$

Thus

$$J_n = -\alpha \log n + O(1), \quad \alpha = \frac{\kappa_1}{\kappa_1 + \kappa_2} \frac{2\kappa_2}{\kappa_1 + 2\kappa_2} \frac{\kappa_3}{\kappa_1 + \kappa_2 + \kappa_3} > 0.$$

On the example used by Anderson, Cappelletti, and Kim to demonstrate failure of unshifted one-step entropy drift, the carried-target episode therefore has strictly negative logarithmic drift. This is precisely what the target augmentation adds.

Lemma 3.2 (Exact target/source identity). *At (x, t) , if the next marked channel has source s and target u , then*

$$V(x - s + u, u) - V(x, t) = \log \frac{(x)_t}{(x)_s}. \quad (13)$$

In particular, a channel sourced at the carried target has zero increment.

Proof. The new residual is $(x - s + u) - u = x - s$. Since both s and t are enabled,

$$\exp(\Delta V) = \prod_i \frac{(x_i - s_i)!}{(x_i - t_i)!} = \frac{(x)_t}{(x)_s}.$$

□

This identity gives a precise “complete-credit” interpretation: a product complex remains credited in full until the chain chooses a different source. The interpretation is not used in place of the algebraic identity.

For enabled source complexes define

$$p_x(y) = \frac{\bar{\kappa}_y(x)_y}{\Lambda(x)}. \quad (14)$$

Set $p_x(y) = 0$ for disabled sources, and use $0 \log 0 = 0$ in entropy expressions. Unless explicitly stated otherwise, sums in this subsection range only over enabled source complexes. Because the carried target is enabled, $p_x(t) > 0$. The expected one-jump increment is

$$\delta(x, t) = \sum_{s: x \geq s} p_x(s) \log \frac{(x)_t}{(x)_s}. \quad (15)$$

Lemma 3.3 (Source-probability bound). *For every reachable augmented state,*

$$\begin{aligned} \delta(x, t) &= \log p_x(t) - \sum_{s: x \geq s} p_x(s) \log p_x(s) \\ &\quad + \sum_{s: x \geq s} p_x(s) \log \bar{\kappa}_s - \log \bar{\kappa}_t, \end{aligned} \quad (16)$$

and consequently

$$\delta(x, t) \leq \log p_x(t) + C_0, \quad C_0 = \log |\mathcal{C}| + \log \frac{\bar{\kappa}_+}{\bar{\kappa}_-}. \quad (17)$$

Proof. For each enabled s , $(x)_s = p_x(s)\Lambda(x)/\bar{\kappa}_s$. Substitution into (15) gives (16). The entropy of a probability vector supported on at most $|\mathcal{C}|$ points is at most $\log |\mathcal{C}|$, and the last line contributes at most $\log(\bar{\kappa}_+/\bar{\kappa}_-)$. The zero complex creates no exception, since $(x)_0 = 1$. $\square$

4 Propagation along target-following paths

For every ordered pair $(t, c) \in \mathcal{C}^2$, fix a simple directed path of marked channels

$$t = y_0 \longrightarrow y_1 \longrightarrow \cdots \longrightarrow y_m = c, \quad 0 \leq m \leq |\mathcal{C}| - 1. \quad (18)$$

Strong connectivity supplies the path. Starting from $(\rho + t, t)$, define the (t, c) -episode as follows. At phase y_k , $k < m$, continue only when the exact designated channel $y_k \rightarrow y_{k+1}$ fires; after any other channel, stop immediately. If the path reaches c , take one final ordinary jump and stop. When $c = t$, the path has length zero and the episode is just that final jump.

This is a stopping time for the marked embedded filtration. On designated edges the lifted populations are

$$\rho + y_0 \longrightarrow \rho + y_1 \longrightarrow \cdots \longrightarrow \rho + c. \quad (19)$$

Every designated source is literally present, the residual remains ρ , and closedness of Γ keeps the entire lift in the class. The episode has between one and $|\mathcal{C}|$ jumps and only finitely many terminal branches.

For the designated channel at phase k , put

$$q_k = \frac{\kappa_{y_k \rightarrow y_{k+1}}}{\bar{\kappa}_{y_k}} > 0, \quad p_k = p_{\rho + y_k}(y_k). \quad (20)$$

Let $J_k(\rho)$ be the expected remaining V -increment from phase k .

Lemma 4.1 (Exact episode recursion). *The episode rewards satisfy*

$$J_m(\rho) = \delta(\rho + c, c), \quad J_k(\rho) = \delta(\rho + y_k, y_k) + q_k p_k J_{k+1}(\rho). \quad (21)$$

Proof. At phase k , the first jump has expected increment $\delta(\rho + y_k, y_k)$. Continuation occurs exactly when the designated channel fires, with probability $q_k p_k$. That channel is sourced at the carried target and has zero immediate increment by Lemma 3.2; the remaining conditional expectation is $J_{k+1}(\rho)$. $\square$

The following scalar calculation controls all possible intermediate source probabilities.

Lemma 4.2 (Scalar envelope). *For $q > 0$ and $M \in \mathbb{R}$, define*

$$F_q(M) = \sup_{0 < p \leq 1} \{\log p + C_0 + qpM\}.$$

The function F_q is nondecreasing in M . Then

$$F_q(M) = \begin{cases} C_0 + qM, & M \geq -1/q, \\ C_0 - 1 - \log(-qM), & M < -1/q. \end{cases} \quad (22)$$

Consequently, if $M_n \rightarrow -\infty$, then $F_q(M_n) \rightarrow -\infty$.

Proof. For each fixed $p \in (0, 1]$, the displayed objective is increasing in M , with slope $qp > 0$; taking its supremum proves monotonicity. The objective has derivative $p^{-1} + qM$ and second derivative $-p^{-2} < 0$. If $M \geq -1/q$, its derivative at $p = 1$ is nonnegative, so concavity places the maximum at $p = 1$. If $M < -1/q$, the unique maximizer is $p = (-qM)^{-1} \in (0, 1)$, which gives the second line of (22). The final assertion follows immediately. $\square$

Proposition 4.3 (Propagation of terminal drift). *Fix a target-following path (18). If along a sequence of residuals $\rho^{(n)}$ one has*

$$p_{\rho^{(n)}+c}(c) \longrightarrow 0,$$

then the expected increment of the complete (t, c) -episode tends to $-\infty$.

Proof. At the terminal phase, Lemma 3.3 gives $J_m \leq \log p_{\rho+c}(c) + C_0 \rightarrow -\infty$. If $J_{k+1} \leq M$, then the monotonicity just proved and Lemmas 3.3 and 4.1 give

$$J_k \leq \sup_{0 < p \leq 1} \{\log p + C_0 + q_k p M\} = F_{q_k}(M).$$

Apply Lemma 4.2 backward through the finite path. $\square$

The proof preserves every reaction-rate monomial: no edge is replaced and no unrelated rates are compared. Arbitrarily many nested source-rate scales are absorbed by the finite scalar recursion.

5 Logarithmic compactification and the top-complex alternative

We now show that every divergent augmented sequence supplies the terminal rarity required by Proposition 4.3, unless a linear stoichiometric invariant already makes such divergence impossible.

Take a sequence $(x^{(n)}, t^{(n)}) \in \tilde{\Gamma}$ leaving every finite set. Pass to a subsequence with $t^{(n)} = t$, and set $\rho^{(n)} = x^{(n)} - t$. A diagonal extraction makes every coordinate of $\rho^{(n)}$ either one fixed nonnegative integer or divergent to infinity. Let I be the set of divergent coordinates and put

$$R_n = \sum_{i \in I} \log(\rho_i^{(n)} + 1). \quad (23)$$

Because the sequence leaves every finite set while t is fixed, $I \neq \emptyset$ and $R_n \rightarrow \infty$. After a further subsequence, the limits

$$w_i = \lim_{n \rightarrow \infty} \frac{\log(\rho_i^{(n)} + 1)}{R_n} \quad (i \in I), \quad w_i = 0 \quad (i \notin I) \quad (24)$$

exist. Thus $w \geq 0$ and $\sum_i w_i = 1$.

This compactification is a coarse, tier-style encoding of relative monomial growth, in the lineage of deterministic tier partitions and their later stochastic refinements (Anderson, 2011a,b; Anderson and Kim, 2018; Anderson et al., 2020a,b). It records the leading logarithmic order $w \cdot y$, but need not separate complexes with equal leading order. The ingredients specific to the present argument are the carried-target exact identity and the scalar envelope, which propagates terminal rarity through a fixed path without comparing intermediate source-rate scales.

Remark 5.1 (Slower divergent tiers). A coordinate can diverge while having $w_i = 0$. It remains in I : zero normalized logarithmic weight means growth on a slower divergent tier, not a bounded species. This distinction is essential in the availability and signed-invariant branches of Lemma 5.3.

Lemma 5.2 (Falling-factorial asymptotics). *Let c and y be fixed complexes, and suppose y is enabled at $\rho^{(n)} + c$ for all sufficiently large n . Then*

$$\log(\rho^{(n)} + c)_y = R_n w \cdot y + o(R_n). \quad (25)$$

Proof. For $y_i = 1$ and a divergent coordinate, $\log(\rho_i^{(n)} + c_i) = \log(\rho_i^{(n)} + 1) + o(R_n)$; a fixed coordinate contributes $O(1)$. For $y_i = 2$ and $\rho_i^{(n)} \rightarrow \infty$,

$$\log((\rho_i^{(n)} + c_i)(\rho_i^{(n)} + c_i - 1)) = 2\log(\rho_i^{(n)} + 1) + o(R_n).$$

If that coordinate is fixed, enabledness makes its contribution finite and hence $O(1)$. Summing the coordinate contributions proves the claim. $\square$

For a complex y , write $h(y) = w \cdot y$, and define

$$h_\star = \max_{y \in \mathcal{C}} h(y), \quad \mathcal{T} = \{y \in \mathcal{C} : h(y) = h_\star\}. \quad (26)$$

The next lemma contains the decisive use of bimolecularity.

Lemma 5.3 (Top-complex availability or invariant). *Along the divergent sequence, at least one of the following conclusions is available:*

- (i) *a reaction-wise linear stoichiometric invariant has value tending to infinity along the sequence, contradicting membership in the fixed communicating class; or*
- (ii) *there are fixed complexes $s, c \in \mathcal{C}$ with $h(s) > h(c)$ such that both s and c are enabled at $\rho^{(n)} + c$ for all sufficiently large n .*

Proof. The exhaustive trichotomy is

- (A) $\mathcal{T} = \mathcal{C}$;
- (B) $\mathcal{T} \neq \mathcal{C}$ and some $s \in \mathcal{T}$ contains two particles supported on I ;
- (C) $\mathcal{T} \neq \mathcal{C}$ and every $s \in \mathcal{T}$ contains exactly one particle supported on I .

Indeed, $\mathcal{T} \neq \mathcal{C}$ implies $h_\star > 0$, so every top complex contains a positive-weight, hence divergent, particle; bimolecularity leaves only (B) or (C).

In case (A), $w \cdot (y' - y) = 0$ for every channel. Therefore $w \cdot X$ is a nonnegative linear invariant. A coordinate with positive weight diverges, so the invariant value tends to infinity, which is impossible inside one communicating class.

In case (B), a top complex s contains two particles supported on coordinates in I . Bimolecularity forces s to contain no particle supported on a bounded coordinate. Both required particles are eventually present in $\rho^{(n)}$, including the case $s = 2S_i$. Hence s is enabled at $\rho^{(n)} + c$ for every lower $c \notin \mathcal{T}$, and $h(s) > h(c)$.

In case (C), define the divergent species represented in top complexes by

$$\mathcal{J} := \{i \in I : y_i \geq 1 \text{ for some } y \in \mathcal{T}\}.$$

If $i \in \mathcal{J}$, the remaining possible particle in that top complex is supported on a bounded coordinate and therefore has weight zero; consequently $w_i = h_\star$. No complex contains two $\mathcal{J}$ -particles, since its weight would be at least $2h_\star > h_\star$. Moreover,

$$y \in \mathcal{T} \iff q_{\mathcal{J}}(y) := \sum_{i \in \mathcal{J}} y_i = 1. \quad (27)$$

The forward implication follows from the definition of $\mathcal{J}$. Conversely, one $\mathcal{J}$ -particle already contributes $h_\star$; maximality forces all other weight to be zero and makes the complex top.

If every complex had $q_{\mathcal{J}} = 1$, then (27) would give $\mathcal{T} = \mathcal{C}$, the all-top case already handled above. There remain three subcases.

- (a) If a unary top complex S_i , $i \in \mathcal{J}$, exists, it is enabled over every lower terminal complex, giving conclusion (ii).
- (b) Otherwise every top complex has the form $S_i + D$, where $D \notin I$; call D a bounded companion species. If a lower complex c contains such a D , the corresponding top source $S_i + D$ is enabled at $\rho^{(n)} + c$, giving conclusion (ii).

- (c) If no lower complex contains any bounded companion species, let $\mathcal{D}$ be the set of distinct companion species. The signed linear functional

$$\ell(x) := \sum_{i \in \mathcal{J}} x_i - \sum_{D \in \mathcal{D}} x_D \quad (28)$$

has value zero on every complex: each top complex contains one $\mathcal{J}$ -particle and one companion particle, whereas a lower complex contains neither. It is therefore reaction-wise invariant. Its positive part diverges, while every companion coordinate lies outside I and is fixed along the subsequence. Its value tends to infinity, contradicting the fixed communicating class.

This exhausts all binary complexes. We reserve “conservation law” for the nonnegative invariants above; (28) is a signed linear stoichiometric invariant. $\square$

The invariant contradiction is exact: if $\ell \cdot (y' - y) = 0$ on every channel, then $\ell \cdot x$ is constant along every path in Γ . In the signed branch, the coordinates with negative coefficients are fixed along the extracted sequence, so the invariant really does tend to $+\infty$. A species absent from every complex is itself reaction-wise invariant; thus an absent coordinate cannot diverge along a fixed communicating class.

Lemma 5.4 (Vanishing terminal source probability). *Under conclusion (ii) of Lemma 5.3,*

$$p_{\rho^{(n)}+c}(c) \longrightarrow 0. \quad (29)$$

Proof. Both s and c are enabled. Lemma 5.2 and the strict gap $h(s) > h(c)$ give

$$\frac{(\rho^{(n)} + c)_s}{(\rho^{(n)} + c)_c} \longrightarrow \infty.$$

Consequently,

$$p_{\rho^{(n)}+c}(c) \leq \frac{\bar{\kappa}_c(\rho^{(n)} + c)_c}{\bar{\kappa}_s(\rho^{(n)} + c)_s} \longrightarrow 0.$$

$\square$

6 Uniform random-time drift outside a finite set

For each $c \in \mathcal{C}$, let $D_c(x, t)$ denote the exact expected increment of the fixed (t, c) -episode. Define

$$K = \left\{ (x, t) \in \tilde{\Gamma} : \min_{c \in \mathcal{C}} D_c(x, t) > -1 \right\}. \quad (30)$$

Proposition 6.1 (Finite and nonempty exceptional set). *The set K is finite and nonempty.*

Proof. Suppose first that K were infinite. Since every sublevel set of V is finite, choose $z_n = (x^{(n)}, t^{(n)}) \in K$ with $V(z_n) > n$. Then z_n leaves every finite set. Apply the compactification of Section 5. An invariant conclusion from Lemma 5.3 contradicts membership in one communicating class. Otherwise, Lemma 5.4 and Proposition 4.3 supply a fixed terminal complex c for which $D_c(x^{(n)}, t) \rightarrow -\infty$. This contradicts the definition of K . Thus K is finite.

To prove nonemptiness, select any augmented state z_0 . The finite nonempty sublevel set $\{z : V(z) \leq V(z_0)\}$ contains a global minimizer z_* of V . Every endpoint of every target-following episode remains in $\tilde{\Gamma}$, so its potential is at least $V(z_*)$. Hence $D_c(z_*) \geq 0$ for every c , and $z_* \in K$. $\square$

Remark 6.2 (Rate dependence of the exceptional set). The finiteness proof is qualitative and cannot give a rate-independent bound on the location of K . Consider the directed cycle

$$0 \xrightarrow{\kappa_0} A \xrightarrow{\kappa_1} A + B \xrightarrow{\kappa_2} 0$$

and the augmented state $((m, 0), A)$, $m \geq 2$. The population chain is irreducible on $\mathbb{N}_0^2$, and each such marked state is reachable after a $0 \rightarrow A$ jump from $(m-1, 0)$; hence all these states lie in one fixed augmented class. For the target-following path from A through $A+B$ to terminal 0, the exact recursion gives

$$D_0(m, A) = a_m + p_m(b_m + q_m c_m), \quad (31)$$

where

$$\begin{aligned} a_m &= \frac{\kappa_0 \log m}{\kappa_0 + \kappa_1 m}, & p_m &= \frac{\kappa_1 m}{\kappa_0 + \kappa_1 m}, \\ b_m &= \frac{\kappa_0 \log m}{\kappa_0 + (\kappa_1 + \kappa_2)m}, & q_m &= \frac{\kappa_2 m}{\kappa_0 + (\kappa_1 + \kappa_2)m}, \\ c_m &= -\frac{\kappa_1(m-1)}{\kappa_0 + \kappa_1(m-1)} \log(m-1). \end{aligned}$$

The three terms are respectively the drift at carried targets A , $A+B$, and 0, with p_m and q_m the exact continuation probabilities. Hence

$$D_0(m, A) = -\frac{\kappa_2}{\kappa_1 + \kappa_2} \log m + O\left(\frac{\log m}{m}\right). \quad (32)$$

For the other two possible terminals, the same path library gives $D_A(m, A) = a_m > 0$ and $D_{A+B}(m, A) = a_m + p_m b_m > 0$. With κ_0, κ_1 fixed, one has $b_m \rightarrow a_m$ and $q_m c_m \rightarrow 0$, so

$$D_0(m, A) \rightarrow a_m(1 + p_m) > 0 \quad (\kappa_2 \downarrow 0) \quad (33)$$

for each fixed m . Thus, for every finite M , a sufficiently small positive κ_2 puts all the states $((m, 0), A)$, $2 \leq m \leq M$, in K . At the same time the negative coefficient in (32) can be arbitrarily small. Consequently, neither the location nor the diameter of K is bounded by the numbers of species and complexes alone, uniformly over positive rate vectors.

Severe attenuation in this particular construction is not confined to rate degeneration. For $d \geq 2$, consider the unit-rate cycle

$$0 \rightarrow y_1 \rightarrow \cdots \rightarrow y_d \rightarrow 0, \quad y_j = 2S_j.$$

For even N , the marked states $z_N = ((N, \dots, N) + y_1, y_1)$ are reachable in one augmented class. Along the unique path from y_1 to 0, the exact recursion gives

$$D_0(z_N) = -\frac{2}{d^d} \log N + O_d\left(\frac{\log N}{N}\right), \quad (34)$$

whereas the episode drift to every nonzero terminal is positive. Thus the logarithmic restoring coefficient supplied by the episode/Foster construction can decay super-exponentially in the number $d+1$ of complexes even when all rates equal one. This is a limitation of the construction, not a lower bound on the CTMC's true return or mixing rate.

Outside K , choose the lexicographically first terminal complex minimizing $D_c(x, t)$, before the episode begins. This is a deterministic measurable selector on the countable state space and is compatible with restarting under the strong Markov property. Its episode stopping time τ obeys

$$1 \leq \tau \leq N_C := |\mathcal{C}|, \quad \mathbb{E}_{(x,t)}[V(Z_\tau) - V(x, t)] \leq -1, \quad (x, t) \notin K. \quad (35)$$

The coordinate overshoot in one episode is deterministically bounded by $2N_C$, because each binary reaction changes each population coordinate by at most two.

7 Positive recurrence

The following argument is a short specialization of standard sampled-chain Foster theory (Meyn and Tweedie, 2009); we include it to record the exact stopping and return interfaces used here.

Let $Y_0 = z$. If $Y_n \notin K$, let Y_{n+1} be the endpoint of the episode selected at Y_n ; if $Y_n \in K$, set $Y_{n+1} = Y_n$. Thus the endpoint chain is stopped upon entering K . Put $\sigma_K = \inf\{n \geq 0 : Y_n \in K\}$. Let $\mathcal{G}_n$ be the endpoint-chain filtration and define

$$W_n = V(Y_{n \wedge \sigma_K}) + (n \wedge \sigma_K).$$

The drift bound (35) makes (W_n) a nonnegative supermartingale. Its terms are integrable: after at most N episodes, each population coordinate is bounded above by its initial value plus $2N_C N$, by the deterministic overshoot bound following (35). Each residual coordinate is nonnegative and no larger than the corresponding population coordinate, so the finite target set then bounds V . Applying the supermartingale inequality at the bounded time N gives

$$\mathbb{E}_z V(Y_{N \wedge \sigma_K}) + \mathbb{E}_z (N \wedge \sigma_K) \leq V(z). \quad (36)$$

Because $V \geq 0$, expectations and monotone convergence yield

$$\mathbb{E}_z \sigma_K \leq V(z). \quad (37)$$

The original augmented embedded chain therefore hits K in finite expected jump count, bounded by $N_C V(z)$.

The following finite trace argument turns hitting K into positive return to one state; the finite-chain step is standard (Norris, 1997), while the conversion from trace excursions to original jumps is recorded explicitly.

Proposition 7.1 (Finite trace-chain closure). *The augmented embedded chain has a state with finite expected positive return time.*

Proof. Let $T_K^+ = \inf\{n \geq 1 : Z_n \in K\}$. From $k \in K$, take one ordinary jump. There are finitely many channel successors, and from each successor the expected hitting time of K is finite by (37). Thus $\mathbb{E}_k T_K^+ < \infty$ for every $k \in K$.

Record successive visits to K . The resulting trace chain is finite and irreducible. Indeed, any augmented path between two points of K , with its intermediate K -visits retained, gives a positive-probability trace path. Fix $k_* \in K$. From every trace state choose a finite positive-probability trace path to k_* . Finiteness supplies common constants $m < \infty$ and $\varepsilon > 0$ such that, from every trace state, k_* is hit within the next m trace transitions with probability at least ε . Starting from k_* , first take one trace transition and then apply this estimate in successive blocks of m transitions. The number of blocks before the next visit to k_* is geometrically dominated, so the positive trace return time has finite mean.

Finally, set

$$B = \max_{k \in K} \mathbb{E}_k T_K^+ < \infty.$$

If M is the positive trace return count and E_j is the number of original jumps in its j -th trace excursion, let $\mathcal{F}_j$ be the sigma-field at the start of that excursion. Continue the trace chain after its return to k_* , so these objects are defined for all $j \geq 0$. Then $\{j < M\} \in \mathcal{F}_j$, and the strong Markov property gives $\mathbb{E}[E_j \mid \mathcal{F}_j] \leq B$. Tonelli's theorem and the tail-sum formula therefore yield

$$\begin{aligned} \mathbb{E} \left[\sum_{j=0}^{M-1} E_j \right] &= \sum_{j \geq 0} \mathbb{E} [\mathbf{1}_{\{j < M\}} E_j] \\ &\leq B \sum_{j \geq 0} \mathbb{P}(j < M) = B \mathbb{E} M < \infty. \end{aligned}$$

Hence the augmented embedded chain has finite expected positive return time to k_* . $\square$

Projection away from the target mark shows that the population embedded chain has finite expected positive return to the population coordinate x_* of k_* . The projected transition law depends only on the population, so the projection started from k_* has exactly the law of the population embedded chain started from x_* ; its first positive population return occurs no later than the marked return to k_* . It remains to construct physical time and rule out explosion.

Proposition 7.2 (Embedded-chain-to-CTMC conversion). *Let a countable irreducible embedded chain be positive recurrent. Suppose a pure-jump construction assigns a finite rate $\Lambda(x) > 0$ at every state and that $\inf_x \Lambda(x) > 0$. Then the resulting minimal CTMC is nonexplosive and positive recurrent.*

Proof. Fix a state x_* . Since the embedded chain is irreducible and positive recurrent, it visits x_* infinitely often almost surely from every starting state. In the standard jump-chain construction, the holding times following those visits are independent exponential random variables of rate $\Lambda(x_*) < \infty$. Their sum diverges almost surely. Since total physical time dominates this subseries, jump times cannot accumulate at a finite time. Thus the minimal CTMC is nonexplosive.

Let N be the first positive return count to x_* , so $\mathbb{E}N < \infty$, and let H_j be the holding times. Conditional on the embedded path, H_j is exponential with mean $1/\Lambda(X_j)$. Tonelli's theorem and conditional expectation now give

$$\mathbb{E} \sum_{j=0}^{N-1} H_j = \mathbb{E} \sum_{j \geq 0} \mathbf{1}_{\{j < N\}} \frac{1}{\Lambda(X_j)} \leq \frac{\mathbb{E}N}{\inf_x \Lambda(x)} < \infty.$$

Hence the CTMC has finite mean positive return time to x_* and is positive recurrent. $\square$

Apply Proposition 7.2 directly to the population embedded chain on Γ . Let $\kappa_{\min} = \min_r \kappa_r > 0$, with the minimum taken over the finite reduced set of genuine channels. Every $x \in \Gamma$ enables at least one such channel r ; otherwise x would be absorbing and the infinite irreducible class Γ would be a singleton. Since a positive falling factorial is an integer,

$$\Lambda(x) \geq \kappa_r(x)_{y_r} \geq \kappa_{\min}, \quad x \in \Gamma. \quad (38)$$

Together with the finite and absorbing reductions in Section 2, Propositions 7.1 and 7.2 therefore prove the recurrence and nonexplosion assertions of Theorem 2.1; its closure assertion is Lemma 2.2. For context, non-explosion was previously obtained for complex-balanced stochastic systems through a stationary-integrability argument (Anderson et al., 2018). For the present subclass, the return-cycle construction recovers nonexplosion directly, consistently with the more general second-order endotactic nonexplosion theorem of Xu (2024).

The same return cycle yields Corollary 2.3. Explicitly, for the positive return time $T_{x_*}^+$ defined in Section 2,

$$\pi_\Gamma(y) = \frac{\mathbb{E}_{x_*} \left[\int_0^{T_{x_*}^+} \mathbf{1}_{\{X(t)=y\}} dt \right]}{\mathbb{E}_{x_*} T_{x_*}^+}, \quad y \in \Gamma, \quad (39)$$

is the normalized mean occupation measure of one regenerative cycle. Here $0 < \mathbb{E}_{x_*} T_{x_*}^+ < \infty$ by the preceding return-time argument. Each numerator is finite, and Tonelli's theorem, applied to $\sum_{y \in \Gamma} \mathbf{1}_{\{X(t)=y\}} = 1$, shows that the denominator normalizes (39) to a probability measure. The strong Markov property at successive returns to x_* makes these cycles regenerative, so the standard regenerative occupation theorem makes this measure stationary (Asmussen, 2003); irreducibility gives uniqueness. For an absorbing singleton, the unique stationary law is the point mass, as stated in Theorem 2.1.

8 Scope and further questions

The proof is class-wise. Coordinate faces, siphons, parity restrictions, and other lattice constraints require no separate interior argument: every constructed target-following path consists of actual enabled channels and therefore remains in the fixed closed class. Species absent from all complexes are constant on each class and appear automatically in the invariant branch of Lemma 5.3.

Lemma 2.2 shows that weak reversibility itself makes the reachability class of every initial state closed, even with multiple linkage classes or higher molecularity. The positive-recurrence argument is narrower: it still requires one linkage class and bimolecularity. Those extensions, and the full weakly reversible stochastic positive-recurrence conjecture, remain open.

The one-linkage hypothesis is used at one decisive point: for every carried target t and selected terminal complex c , a directed path (18) must exist. With several linkage classes, a carried target need not reach a terminal source belonging to another class, so the finite path library does not directly glue. That extension remains open here.

Bimolecularity is used in the exhaustive case split of Lemma 5.3: a top complex either contains two particles supported on divergent coordinates or exactly one such particle, leaving at most one bounded companion species. Complexes of molecularity three or higher admit additional top configurations not covered by the argument. No claim is made for them.

The potentially reusable part of the proof is the combination of marked reaction-channel augmentation, subtraction of the carried target, the exact residual-factorial increment, finite target-following propagation, and normalized-log top-complex compactification. The multiple-linkage obstruction is precise: a target in one linkage class need not have a directed path to a useful terminal complex in another linkage class. Overcoming that obstruction would require an additional mechanism and is not a routine extension of the present path argument.

The proof establishes positive recurrence and the existence of a stationary probability distribution. It does not provide explicit tail bounds, exponential ergodicity, or quantitative mixing rates. The finite set K is obtained by analytic compactness, not by enumeration, and the proof gives no practically useful universal estimate of its location or diameter. Its scale can depend arbitrarily strongly on positive rate ratios, and its restoring coefficient can deteriorate super-exponentially with path length even at unit rates, as the exact examples following Proposition 6.1 illustrate. Effective rate-dependent recurrence, tail, and mixing estimates remain open. Equation (39) gives a regenerative representation, but not a closed-form or product-form stationary distribution; unbounded molecule-count moments require separate integrability. Its structural hypotheses are not asserted to hold for all biological reaction networks, and no empirical or organism-specific claim is made.

Declarations

Author information. Alec Kriebel, Independent researcher.

Correspondence: me@aleckriebel.com.

ORCID: 0009-0001-9320-500X.

Data availability. No empirical or third-party dataset was used. All deterministic calibration inputs and generated outputs are included in the verification package described below.

Code and supporting materials availability. The dependency-free verifier, tests, canonical verification reports, manuscript source, release manifest, and supporting materials are available in the tagged Version 1.2.4 repository directory. The release records its exact Git commit and SHA-256 checksums. No computation is used in the universal proof; the finite checks are reproducibility and falsification aids and neither prove recurrence nor enumerate the analytic set K .

Preprint status. This is an unrefereed preprint. A journal submission is planned. Earlier versioned repository copies remain publicly available and are superseded by Version 1.2.4.

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Competing interests. The author declares no competing interests.

A note on the use of AI

Generative-AI systems were used materially from 5–22 August 2026. The tools were OpenAI ChatGPT with GPT-5.6 Pro through the web interface; Anthropic Claude with Opus 5 at Max effort through claude.ai;

Anthropic Claude Code with Opus 5 through its command-line application; and OpenAI Codex desktop with a GPT-5-family coding/research deployment whose exact backend identifier was not exposed.

The systems were asked to assist with mathematical exploration and derivation, counterexample search, adversarial proof criticism, symbolic computation and verification code, literature-search support, drafting, editorial revision, publication preparation, and reproducibility review. Outputs retained in the released work were checked against the displayed proof, exact calculations and tests, or cited primary sources, as appropriate. Rejected approaches included universal one-step Lyapunov claims, rate-changing cycle comparisons without a uniform bound, and incomplete fast-slow hierarchy arguments; none is part of the released proof.

The author directed the research, reviewed and edited the outputs, determined the released scope and claims, and assumes responsibility for the manuscript and verification materials. No AI system is an author. Tool-by-tool dates, access routes, uses, and safeguards are recorded in `supplement/ai_use_full_statement.md`.

References

- Andrea Agazzi. Weakly reversible chemical reaction networks are recurrent in 2d. Talk at the Jean-Pierre Eckmann 75 Conference, University of Geneva, 2022. URL <https://www.unige.ch/jpe75conference/program.html>. Public program dated 10 June 2022; joint work with Jonathan C. Mattingly, David F. Anderson, and Daniele Cappelletti.
- Andrea Agazzi. Weakly reversible chemical reaction networks are recurrent in 2d. Talk abstract, SIAM Conference on Applied Algebraic Geometry (AG25), 2025. URL https://www.siam.org/media/13rgukxr/ag25_abstracts.pdf. Official abstract book, p. 79; conference held 7–11 July 2025.
- David F. Anderson. Boundedness of trajectories for weakly reversible, single linkage class reaction systems. *Journal of Mathematical Chemistry*, 49(10):2275–2290, 2011a. doi: 10.1007/s10910-011-9886-4.
- David F. Anderson. A proof of the global attractor conjecture in the single linkage class case. *SIAM Journal on Applied Mathematics*, 71(4):1487–1508, 2011b. doi: 10.1137/11082631X.
- David F. Anderson and Jinsu Kim. Some network conditions for positive recurrence of stochastically modeled reaction networks. *SIAM Journal on Applied Mathematics*, 78(5):2692–2713, 2018. doi: 10.1137/17M1161427.
- David F. Anderson and Thomas G. Kurtz. *Stochastic Analysis of Biochemical Systems*, volume 1 of *Mathematical Biosciences Institute Lecture Series*. Springer, Cham, 2015. doi: 10.1007/978-3-319-16895-1.
- David F. Anderson, Gheorghe Craciun, and Thomas G. Kurtz. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of Mathematical Biology*, 72(8):1947–1970, 2010. doi: 10.1007/s11538-010-9517-4.
- David F. Anderson, Gheorghe Craciun, Manoj Gopalkrishnan, and Carsten Wiuf. Lyapunov functions, stationary distributions, and non-equilibrium potential for reaction networks. *Bulletin of Mathematical Biology*, 77(9):1744–1767, 2015. doi: 10.1007/s11538-015-0102-8.
- David F. Anderson, Daniele Cappelletti, Masanori Koyama, and Thomas G. Kurtz. Non-explosivity of stochastically modeled reaction networks that are complex balanced. *Bulletin of Mathematical Biology*, 80(10):2561–2579, 2018. doi: 10.1007/s11538-018-0473-8.
- David F. Anderson, Daniele Cappelletti, and Jinsu Kim. Stochastically modeled weakly reversible reaction networks with a single linkage class. *Journal of Applied Probability*, 57(3):792–810, 2020a. doi: 10.1017/jpr.2020.28.
- David F. Anderson, Daniele Cappelletti, Jinsu Kim, and Tung D. Nguyen. Tier structure of strongly endotactic reaction networks. *Stochastic Processes and their Applications*, 130(12):7218–7259, 2020b. doi: 10.1016/j.spa.2020.07.012.

- David F. Anderson, Daniele Cappelletti, Wai-Tong Louis Fan, and Jinsu Kim. A new path method for exponential ergodicity of Markov processes on $\mathbb{Z}^d$, with applications to Stochastic Reaction Networks. *SIAM Journal on Applied Dynamical Systems*, 24(2):1668–1710, 2025. doi: 10.1137/24M1665933.
- Søren Asmussen. *Applied Probability and Queues*, volume 51 of *Stochastic Modelling and Applied Probability*. Springer, New York, 2 edition, 2003. ISBN 978-0-387-00211-8. doi: 10.1007/b97236.
- Balázs Boros and Josef Hofbauer. Permanence of weakly reversible mass-action systems with a single linkage class. *SIAM Journal on Applied Dynamical Systems*, 19(1):352–365, 2020. doi: 10.1137/19M1248431.
- Daniele Cappelletti. Solving the chemical recurrence conjecture in 2D. Talk at the Stochastic Networks Conference, Cornell University, 2022. URL https://vod.video.cornell.edu/media/t/1_hn76x7n9/261126952. Presented 22 June 2022; joint work with Andrea Agazzi, David F. Anderson, and Jonathan C. Mattingly; public Cornell recording.
- ConStRAINeD project. Results and publications. PRIN 2022 project website, 2026. URL <https://constrained.polito.it/publications/>. Accessed 22 August 2026; year is the access year; the overview describes the proof as complete and sketches its approach, while item 16 lists the five-author manuscript as in preparation.
- Friedrich J. M. Horn and Roy Jackson. General mass action kinetics. *Archive for Rational Mechanics and Analysis*, 47(2):81–116, 1972. doi: 10.1007/BF00251225.
- Sean P. Meyn and Richard L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, 2 edition, 2009. ISBN 978-0-521-73182-9. doi: 10.1017/CBO9780511626630.
- J. R. Norris. *Markov Chains*. Number 2 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1997. ISBN 978-0-521-63396-3. doi: 10.1017/CBO9780511810633.
- Loïc Paulevé, Gheorghe Craciun, and Heinz Koepl. Dynamical properties of Discrete Reaction Networks. *Journal of Mathematical Biology*, 69(1):55–72, 2014. doi: 10.1007/s00285-013-0686-2.
- Carsten Wiuf and Chuang Xu. Classification and threshold dynamics of stochastic reaction networks, 2023. URL <https://arxiv.org/abs/2012.07954>. Version 3, 24 January 2023.
- Chuang Xu. On the Regularity of Reaction Systems, 2024. URL <https://arxiv.org/abs/2409.05340>. arXiv:2409.05340v2, revised 9 May 2026.