

Positive Recurrence of Bimolecular Weakly Reversible Stochastic Reaction Networks with a Single Linkage Class

Alec Kriebel

Independent researcher; corresponding author: me@aleckriebel.com

ORCID: 0009-0001-9320-500X

Abstract

Stochastic reaction networks model interacting populations through random reaction events. Weak reversibility is widely expected to prevent such systems from escaping indefinitely, but this was not known even for second-order networks with a connected reaction graph unless every species appeared in a special pure-species complex. We remove that additional restriction. For all positive reaction rates, every closed component of the state space has finite mean return times and a unique stationary probability distribution; the construction also rules out explosion within each such component.

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

2020 Mathematics Subject Classification. Primary 60J27; secondary 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. They arise in chemical kinetics, systems biology, ecology, and related population models; see, for example, Anderson and Kurtz (2015). For a nonexplosive irreducible component, a basic stability question is whether it has a stationary probability distribution, or equivalently whether its states have finite mean return times.

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

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

Thus each species had to occur in a pure unary or pure double complex. 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 all second-order weakly reversible systems (Xu, 2026). A complementary two-species positive-recurrence theorem has been publicly announced in conference talks since 2022 (Agazzi, 2022, 2025), but no manuscript for that result was publicly available at the time of writing. It allows a broader network structure

in two species, whereas the present result allows arbitrarily many species under the binary, one-linkage assumptions. The May 2026 revision of Xu (2026) still records the general bimolecular positive-recurrence problem as open.

The present paper removes (1) while retaining the binary and one-linkage assumptions. 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.

Figure 1 summarizes the proof architecture and locates the sole structural use of bimolecularity.




Figure 1: The proof mechanism. The top-complex alternative either supplies the rare terminal source shown here or produces a stoichiometric invariant that precludes the divergent sequence. This schematic was generated during the disclosed AI-assisted manuscript preparation; it is expository and adds no claim beyond the surrounding proof outline.

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. The reaction graph has the complexes as vertices 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.

A subset $\Gamma \subseteq \mathbb{N}_0^d$ is a *closed communicating class* when its states communicate through enabled channels and no enabled channel exits Γ . The chain restricted to Γ is irreducible. In a nonabsorbing class, let J_1 be the first jump time and define the positive return time to x by $T_x^+ = \inf\{t \geq J_1 : X(t) = x\}$. The class is *positive recurrent* when one, equivalently every, state has $\mathbb{E}_x T_x^+ < \infty$. Absorbing singleton classes are positive recurrent by convention. For a nonexplosive irreducible countable-state CTMC, positive recurrence is equivalent to the existence of a stationary probability distribution.

A channel $y \rightarrow y$ has zero population increment and contributes zero to (6); discard it. 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 closed communicating 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. 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 infinite closed irreducible class Γ , a strongly connected graph with at least two complexes, and no self channels. Every complex then has a genuine outgoing channel. Write

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

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, the minimal population CTMC restricted to each closed communicating class is nonexplosive and positive recurrent.*

Corollary 2.2 (Stationary probability). *Under the hypotheses of Theorem 2.1, every closed communicating class Γ admits a stationary probability distribution π_Γ . It is unique on Γ .*

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. 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

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

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.

Lemma 3.1 (Exact target/source identity). *At (x, t) , if the next marked channel has source s and target u , then*

$$V(x - s + u, u) - V(x, t) = \log \frac{(x)_t}{(x)_s}. \quad (11)$$

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

$$p_x(y) = \frac{\bar{\kappa}_y(x)_y}{\Lambda(x)}. \quad (12)$$

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

$$d(x, t) = \sum_{s: x \geq s} p_x(s) \log \frac{(x)_t}{(x)_s}. \quad (13)$$

Lemma 3.2 (Source-probability bound). *For every reachable augmented state,*

$$\begin{aligned} d(x, t) &= \log p_x(t) - \sum_{s: x \geq s} p_x(s) \log p_x(s) \\ &\quad + \sum_{s: x \geq s} p_x(s) \log \bar{\kappa}_s - \log \bar{\kappa}_t, \end{aligned} \quad (14)$$

and consequently

$$d(x, t) \leq \log p_x(t) + C_0, \quad C_0 = \log |\mathcal{C}| + \log \frac{\bar{\kappa}_+}{\bar{\kappa}_-}. \quad (15)$$

Proof. For each enabled s , $(x)_s = p_x(s) \Lambda(x) / \bar{\kappa}_s$. Substitution into (13) gives (14). 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 $(x)_0 = 1$. □

4 Propagation along target-following paths

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

$$t = y_0 \longrightarrow y_1 \longrightarrow \cdots \longrightarrow y_L = c, \quad 0 \leq L \leq |\mathcal{C}| - 1. \quad (16)$$

Strong connectivity supplies the path. Starting from $(r + t, t)$, define the (t, c) -episode as follows. At phase y_k , $k < L$, 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

$$r + y_0 \longrightarrow r + y_1 \longrightarrow \cdots \longrightarrow r + c. \quad (17)$$

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

For the designated channel at phase k , put

$$q_k = \frac{\kappa_{y_k \rightarrow y_{k+1}}}{\bar{\kappa}_{y_k}} > 0, \quad p_k = p_{r+y_k}(y_k). \quad (18)$$

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

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

$$J_L(r) = d(r + c, c), \quad J_k(r) = d(r + y_k, y_k) + q_k p_k J_{k+1}(r). \quad (19)$$

Proof. At phase k , the first jump has expected increment $d(r + y_k, y_k)$. 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 3.1; the remaining conditional expectation is $J_{k+1}(r)$. $\square$

The following scalar calculation controls all possible intermediate source probabilities.

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

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

Then

$$F_q(M) = \begin{cases} C_0 + qM, & M \geq -1/q, \\ C_0 - 1 - \log(-qM), & M < -1/q. \end{cases} \quad (20)$$

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

Proof. 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 (20). The final assertion follows immediately. $\square$

Proposition 4.3 (Propagation of terminal drift). *Fix a target-following path (16). If along a sequence of residuals $r^{(n)}$ one has*

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

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

Proof. At the terminal phase, Lemma 3.2 gives $J_L \leq \log p_{r+c}(c) + C_0 \rightarrow -\infty$. If $J_{k+1} \leq M$, Lemmas 3.2 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 $r^{(n)} = x^{(n)} - t$. A diagonal extraction makes every coordinate of $r^{(n)}$ either one fixed nonnegative integer or divergent to infinity. Let I be the set of divergent coordinates and put

$$R_n = \sum_{i \in I} \log(r_i^{(n)} + 1). \quad (21)$$

After a further subsequence, the limits

$$w_i = \lim_{n \rightarrow \infty} \frac{\log(r_i^{(n)} + 1)}{R_n} \quad (i \in I), \quad w_i = 0 \quad (i \notin I) \quad (22)$$

exist. Thus $w \geq 0$ and $\sum_i w_i = 1$. A divergent coordinate may have $w_i = 0$; it remains in I , so all slower tiers are retained.

Lemma 5.1 (Falling-factorial asymptotics). *Let c and y be fixed complexes, and suppose y is enabled at $r^{(n)} + c$ for all sufficiently large n . Then*

$$\log(r^{(n)} + c)_y = R_n w \cdot y + o(R_n). \quad (23)$$

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

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

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

$$a = \max_{y \in \mathcal{C}} h(y), \quad T = \{y \in \mathcal{C} : h(y) = a\}. \quad (24)$$

The next lemma contains the decisive use of bimolecularity. When $T \neq \mathcal{C}$, its maximal weight satisfies $a > 0$, so every top complex contains at least one particle in a divergent coordinate. Because a binary complex contains at most two particles, only two configurations remain: a top complex contains two such particles, or every top complex contains exactly one. This elementary restriction makes the following alternative exhaustive; it is precisely what fails without the molecularity bound.

Lemma 5.2 (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 $r^{(n)} + c$ for all sufficiently large n .*

Proof. We follow the three exhaustive cases described above. First, if every complex is top, then $w \cdot (y' - y) = 0$ 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.

Second, assume $T \neq \mathcal{C}$ and 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 $r^{(n)}$, including the case $s = 2S_i$. Hence s is enabled at $r^{(n)} + c$ for every lower $c \notin T$, and $h(s) > h(c)$.

Third, suppose every top complex contains exactly one particle supported on a coordinate in I . Let $J \subseteq I$ be the species appearing in top complexes. If $i \in J$, the remaining possible particle in that top complex is supported on a bounded coordinate and therefore has weight zero; consequently $w_i = a$. No complex contains two J -particles, since its weight would be at least $2a > a$. Moreover,

$$y \in T \iff q_J(y) := \sum_{i \in J} y_i = 1. \quad (25)$$

The forward implication follows from the definition of J . Conversely, one J -particle already contributes a ; maximality forces all other weight to be zero and makes the complex top.

There are four subcases.

- (a) If every complex has $q_J = 1$, then $M_J = \sum_{i \in J} X_i$ is a nonnegative invariant. It diverges along the sequence, a contradiction.
- (b) If a unary top complex S_i , $i \in J$, exists, it is enabled over every lower terminal complex, giving conclusion (ii).
- (c) 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 $r^{(n)} + c$, giving conclusion (ii).
- (d) If no lower complex contains any bounded companion species, let $\mathcal{D}$ be the set of distinct companion species. The signed linear functional

$$M_J - \sum_{D \in \mathcal{D}} X_D \quad (26)$$

has value zero on every complex: each top complex contains one 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; (26) 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 5.3 (Vanishing terminal source probability). *Under conclusion (ii) of Lemma 5.2,*

$$p_{r^{(n)}+c}(c) \longrightarrow 0. \quad (27)$$

Proof. Both s and c are enabled. Lemma 5.1 and the strict gap $h(s) > h(c)$ give

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

Consequently,

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

□

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

$$K = \left\{ (x, t) \in \tilde{\Gamma} : \min_{c \in \mathcal{C}} D_c(x, t) > -1 \right\}. \quad (28)$$

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

Proof. Suppose first that K were infinite. Properness of V supplies a sequence in K leaving every finite set. Apply the compactification of Section 5. An invariant conclusion from Lemma 5.2 contradicts membership in one communicating class. Otherwise, Lemma 5.3 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 $z_* \in K$. □

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

$$1 \leq \tau \leq L := |\mathcal{C}|, \quad \mathbb{E}_{(x,t)}[V(Z_\tau) - V(x, t)] \leq -1, \quad (x, t) \notin K. \quad (29)$$

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

7 Positive recurrence

We now convert (29) to a finite mean return time without invoking an unnamed random-time Foster theorem.

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 $\sigma_K = \inf\{n \geq 0 : Y_n \in K\}$. For each integer N , sum the conditional drift inequalities only through the bounded index $N \wedge \sigma_K$. All terms are integrable: after at most N episodes, each population coordinate is bounded above by its initial value plus $2LN$, by the deterministic overshoot bound following (29); the finite target set then bounds V . Conditional expectation and telescoping therefore give

$$\mathbb{E}_z V(Y_{N \wedge \sigma_K}) + \mathbb{E}_z(N \wedge \sigma_K) \leq V(z). \quad (30)$$

Because $V \geq 0$, monotone convergence yields

$$\mathbb{E}_z \sigma_K \leq V(z). \quad (31)$$

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

The following finite trace argument turns hitting K into positive return to one state. Details are included because hitting a finite set is not by itself a positive-recurrence proof.

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 (31). 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 L_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 $\mathbb{E}[L_j \mid \mathcal{F}_j] \leq B$. Tonelli's theorem and the tail-sum formula therefore yield

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

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_* : 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. Let N be a finite-mean positive return count to a recurrent embedded state x_* , 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 give

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

Thus the expected physical return time is finite, provided the construction is nonexplosive.

Finite mean return makes x_* recurrent for the embedded chain, so it is visited infinitely often almost surely. On each visit, the following holding time has the same exponential law of finite rate $\Lambda(x_*)$. The strong Markov construction makes these recurrent holding times independent copies. Their infinite sum diverges almost surely. Since total physical time dominates this subseries, jump times cannot accumulate at a finite time. The minimal CTMC is nonexplosive and has finite mean positive return time to x_* . $\square$

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,

$$\Lambda(x) \geq \kappa_r(x)_{y_r} \geq \kappa_{\min}, \quad x \in \Gamma. \quad (32)$$

Propositions 7.1 and 7.2 therefore prove Theorem 2.1. The same return-cycle construction yields the stationary distribution in Corollary 2.2; uniqueness follows from irreducibility. Explicitly, for the positive return time $T_{x_*}^+$ defined in Section 2,

$$\pi_\Gamma(y) = \frac{\mathbb{E}_{x_*} \left[\int_0^{T_{x_*}^+} \mathbf{1}_{\{X(t)=y\}} dt \right]}{\mathbb{E}_{x_*} T_{x_*}^+}, \quad y \in \Gamma, \quad (33)$$

is the normalized return-cycle occupation measure and is stationary.

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 5.2.

The one-linkage hypothesis is used at one decisive point: for every carried target t and selected terminal complex c , a directed path (16) 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 5.2: 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 proof establishes positive recurrence and the existence of a stationary probability distribution. It does not provide a product form, explicit tail bounds, exponential ergodicity, or quantitative mixing rates.

Declarations

Author information. Alec Kriebel, Independent researcher.

Correspondence: me@aleckriebel.com.

ORCID: 0009-0001-9320-500X.

Code availability. The self-contained verification package and supporting materials are available in the public project repository at https://github.com/AlecKriebel/Math/tree/bimolecular-positive-recurrence-v0.3/bimolecular_positive_recurrence. The submitted archive is also supplied with the manuscript so that its exact Git snapshot and checksums can be retained with the submission record. No computation is used in the universal proof.

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.

Declaration of generative-AI use

Generative-AI systems were used substantively throughout the research workflow, including mathematical exploration and derivation, adversarial proof analysis, symbolic and software work, literature support, and manuscript preparation. OpenAI ChatGPT, model GPT-5.6 Pro, was accessed through the ChatGPT interface at <https://chatgpt.com> on 5–6 August 2026 for exploratory proof search, attempted counterexample construction, verification-code development and testing, literature assistance, drafting, and editorial criticism. Anthropic Claude, model Opus 5 with Max effort, was accessed through <https://claude.ai> on 6 August 2026 for an adversarial review of the proof and manuscript. Anthropic Claude Code using Opus 5 was accessed through the Claude Code command-line application on 6 August 2026 to publish the manuscript on the project website and perform related site formatting; it did not supply an additional mathematical argument.

OpenAI Codex desktop was accessed on 8–9 August 2026 for a further publication-readiness adversarial review, manuscript and verification-package repair, and validation. It used an OpenAI GPT-5-family coding/research model; the exact backend deployment identifier was not exposed in the session. All of these systems are proprietary services of OpenAI or Anthropic and were used without model modification or fine-tuning. AI-generated literature leads were not treated

as scholarly authority; the release identifies the public or scholarly sources on which its literature claims rely.

The work was developed through an iterative, human-directed research process. The author determined the scope and claims released here, curated the resulting package, and assumes responsibility for the submitted manuscript and public verification materials. The proof and exact artifacts are supplied to make the AI-mediated research independently checkable; no claim of prior independent expert human verification is made. No AI system is listed as an author.

A Technical details for the trace chain and physical time

This appendix records several interfaces that can otherwise be obscured by standard Markov-chain shorthand.

First, the augmented state records the actual marked channel. If two channels have the same population displacement but different sources or targets, their post-jump populations can agree while their carried targets do not. They therefore remain separate. Only channels with identical source and target may be aggregated.

Second, K in (28) is nonempty, not merely finite. The global minimum argument in Proposition 6.1 is pathwise: every possible episode endpoint has potential at least the starting minimum. No optional stopping or integrability assertion is needed for this step because each episode has at most $|\mathcal{C}|$ jumps.

Third, the trace-chain proof distinguishes hitting from return. The first ordinary jump from $k \in K$ forces a positive return time. The finite expected time back to K follows by conditioning on finitely many successors. The subsequent geometric-block argument is applied to the finite trace chain, and only afterward are the random trace excursions expanded into original embedded jumps.

Fourth, nonexplosion is not inferred from recurrence without an argument. The embedded chain is constructed first. Its recurrent marked state is visited infinitely often. The holding times attached to those visits form an infinite sequence of independent exponentials with one fixed finite rate, whose sum diverges almost surely. This directly prevents finite-time accumulation.

Finally, singleton absorbing classes are handled before augmentation. In an infinite irreducible class every population state has an enabled genuine channel; otherwise it would be absorbing and its class would be a singleton.

B Supplementary computational checks

The accompanying dependency-free Python package performs exact-arithmetic checks of the falling-factorial identity and entropy rewrite, verifies the scalar-envelope branch conditions, and tests the top-complex classification on a deterministic finite atlas with separately validated witnesses. It also contains fixed-seed adversarial stress tests. The supplement documents the test cases, normalization rules, and reproduction command. These finite checks are regression tests and adversarial calibration; they are not used in place of the universal proof, in particular Lemma 5.2.

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. Conference announcement; a related project page lists joint work with David F. Anderson, Daniele Cappelletti, Lucie Laurence, and Jonathan C. Mattingly as in preparation.
- 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, 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.
- 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, 2026. URL <https://arxiv.org/abs/2409.05340>. Version 2, 9 May 2026.