Activated random walk: AI-generated results and proofs

A maintained collection from the research programme directed by Jacob Richey

8 November 2026

Abstract

This document collects the results and arguments developed in the AI research record on finite activated random walk. Its principal question is whether adding one initially active particle increases the final occupation probability of a specified site. The general conjecture remains open. The collection separates established parameter and graph regimes, structural tools, adjacent results, exact obstructions, and work whose proof is incomplete. It is organized by mathematical use, with a topic guide, linked contents, and a list of named results. Every result carries a reference to its supporting record and retains that record's hypotheses and verification status.

Overview and reading guide

The collection has four parts. Part I gives monotonicity and comparison theorems: low and high sleep regimes, stationary total-count order under changing sink strength, and directed graph constructions. Part II develops tools for the open conjecture, from history exchanges and root restoration to loop constructions and inverse stabilization. Part III collects related probability and algebra. Part IV records the remaining gaps, failed approaches, and supporting evidence. The contents give this subject map; the separate results index groups individual theorems and lemmas by subject.

The purpose is to make the AI-generated work separately readable. The source corpus is the maintained monotonicity notes, their explicitly AI-written evidence bundles, the separately identified sink-flow argument from the negative-association branch, and the reproducible verifiers listed at the end. The separate human-authored write-up supplied by Nicolas Forien is not reproduced here. Standard results and published theorems used as inputs are credited as background; their use is not an AI originality claim. "AI-generated" describes the provenance of this research record, not a completed literature-priority check or a claim of human verification.

Choose a subject

What to look for	What is covered	Where
I. Monotonicity and comparison theorems	Low sleep for all inputs; high sleep for two common particles; sink-strength order for stationary counts; directed gluing, small components, cycles, and drain inventories.	Part I
II. Proof techniques	History exchanges; backward restoration and root retention; creator laws and loops; inverse stabilization and releases; return-kernel certificates.	Part II
III. Related probability and algebra	Negative dependence; filling and matroids; currents; algebras and geometry; fluid growth; duality.	Part III
IV. Open questions and evidence	Unproved comparisons and failed approaches; source register, exact checks, and attribution.	Part IV

Status conventions. A theorem or proposition followed by a proof states a mathematical result at the scope written. A *computer-assisted theorem* relies on an exact symbolic identity or a finite exhaustive certificate for a stated parameter domain; this differs from testing finitely many parameter values. A *proof reference* points to a retained argument not reproduced in full here. A *conditional statement* names the missing assumption. Numerical experiments and unsuccessful proof strategies appear as evidence or obstructions, never as proofs of the general conjecture. Prior independent audits in these notes were AI audits unless explicitly stated otherwise.

Contents by subject

Overview and reading guide 1

Results by subject 4

I Monotonicity and comparison theorems 7

1 Model, target, and shared background 7

2 Sleep probabilities and sink strength 8

3 Directed gluing and drain inventories 19

II Proof techniques for the open conjecture 32

4 History exchanges and avalanche repairs 32

5 Root restoration and joint retention 38

6 Creator laws, loops, and activation 45

7 Inverse stabilization and release comparisons 54

8 Return kernels and boundary interpolation 68

III Related probability and algebra 71

9 Dependence, geometry, growth, and duality 71

IV Open questions and evidence 88

10 Open questions and failed approaches 88

11 Sources and verification 89

Results by subject

Sleep probabilities and sink strength

2.1 Theorem (Small-sleep cube) . . . 9
2.2 Corollary (Localization and a single-target weakening) . . . 12
2.4 Theorem (High-sleep N2 on the creation cone) . . . 12
2.5 Corollary (Arbitrary added batch) . . . 14
2.7 Proposition (Positive transport with fixed routing) . . . 14
2.8 Lemma (All-sleep tangent and one-coordinate interpolation) . . . 15
2.10 Theorem (All-rate stationary count order under sink reinjection) . . . 16
2.11 Corollary (A single sink port) . . . 18
2.12 Proposition (A Markov evolution in sink strength) . . . 18

Directed gluing, small components, cycles, and inventories

3.1 Theorem (Exact gluing in law) . . . 20
3.2 Theorem (Localization) . . . 20
3.3 Theorem (Three-site occupation and pair capacity) . . . 21
3.4 Theorem (Coherent DAG dual) . . . 22
3.5 Proposition (Weighted rows and small components) . . . 23
3.6 Lemma (Exact slices and zero-exit additions) . . . 24
3.7 Theorem (Static DAG bound) . . . 24
3.8 Lemma (Root entry into an outward forest) . . . 24
3.9 Theorem (Forest feedback) . . . 25
3.10 Corollary (Every directed cycle) . . . 25
3.11 Proposition (Static SCC bound and a conditional feedback criterion) . . . 26
3.12 Lemma (Finite forest enumeration and partial cancellation) . . . 27
3.13 Theorem (Endpoint and nonreturning mixed-site inventories) . . . 27
3.14 Lemma (Self-loop insertion) . . . 28
3.15 Theorem (One off-target sleeping site) . . . 28
3.16 Theorem (Unpadded directed-three-cycle inventories) . . . 28
3.17 Proposition (Permanent extra on general three-site kernels) . . . 29
3.18 Proposition (Positive homogeneous count blocks on every cycle) . . . 30

History exchanges and avalanche repairs

4.1 Proposition (First-contact cancellation) . . . 32
4.2 Theorem (Adapted activation with correct conditional blocks) . . . 32
4.3 Proposition (Discovered blocks retain their selection mass) . . . 33
4.4 Proposition (Creation-history factorization of a forced-jump defect) . . . 33
4.5 Lemma (Recoverable stopped-word exchange) . . . 34
4.6 Proposition (Connected-history reduction) . . . 34
4.7 Proposition (Joint success conditioning) . . . 35
4.8 Lemma (Released residents and suppressed final jumps) . . . 35
4.9 Lemma (Full avalanche block exchange) . . . 35
4.10 Theorem (Matched-background terminal repair) . . . 36
4.11 Corollary (Unmatched backgrounds with a specified avoidance test) . . . 36
4.12 Proposition (Creator and endpoint switches: recorded local families) . . . 37

Backward restoration and joint root retention

5.1 Theorem (Backward roots and pointwise minorization) . . . 38
5.2 Theorem (Canonical terminal-root restoration) . . . 39
5.3 Proposition (Entrance representation) . . . 40
5.4 Proposition (Root-block cancellation) . . . 41
5.5 Theorem (Boundary enlargement and joint retention) . . . 42
5.6 Theorem (N1, including occupied-set order) . . . 43

Creator laws, loops, and activation

6.1 Proposition (Three-particle contact renewal) . . . 45
6.2 Theorem (Positive creator mixtures) . . . 46
6.3 Theorem (Positive full-graph creator law) . . . 47
6.4 Theorem (Active-cycle erasure) . . . 48
6.5 Theorem (Inverse-diagonal hierarchy and thinning) . . . 49
6.6 Theorem (Periodic acceptance and coherent paired prefixes) . . . 50
6.7 Proposition (Promotion and creator resampling) . . . 51
6.8 Proposition (Surplus traversal and residual books) . . . 51
6.9 Theorem (Excursion dual before loss) . . . 52

Inverse stabilization, finite books, and release comparisons

7.1 Theorem (Complete-block inverse) . . . 54
7.2 Corollary (Waking transport) . . . 54
7.3 Lemma (First return and killed creation) . . . 55
7.4 Theorem (Invertibility of the signed response renewal) . . . 55
7.5 Proposition (Last-arrival deletion and its invariant cone) . . . 56
7.6 Theorem (Ordered normal form and rejected-loop currents) . . . 57
7.7 Theorem (Remainder factorization and complementary pairing) . . . 58
7.8 Theorem (Stable-cut inversion) . . . 59
7.9 Corollary (Boundary regularity and first cuts) . . . 60
7.10 Theorem (The greatest cut and its finite envelope) . . . 60
7.12 Proposition (Burning count and fixed-inventory ensembles) . . . 61
7.13 Theorem (Recurrent rotor correspondence) . . . 62
7.14 Proposition (Two exact full-minus-one reductions) . . . 63
7.15 Theorem (One-coordinate release resolvent) . . . 64
7.16 Theorem (Two-site prefix and forest-degree inequalities) . . . 64
7.17 Proposition (Rotation averages fail, with an exact repair family) . . . 66

Return kernels and boundary interpolation

8.1 Theorem (Three-cycle killed-return order; computer-assisted) . . . 68
8.2 Theorem (Positive boundary interpolation) . . . 69

Negative dependence of loss markers

9.1 Proposition (Determinantal mixture) . . . 72
9.2 Theorem (Scoped negative dependence) . . . 72

IDLA filling, matroids, and sleep asymptotics

9.4 Theorem (Sure-filling matroid) . . . 74
9.5 Theorem (Reachability and exact supports) . . . 74
9.6 Corollary (Circuit support of the ARW response) . . . 74
9.8 Lemma (Minimal sleep exponents) . . . 75

Currents and changes of routing

9.9 Theorem (Harmonic change of routing) . . . 76
9.10 Theorem (Current completion) . . . 76
9.11 Proposition (Internal currents and the full spectrum) . . . 77

Addition algebras, geometry, and input inequalities

9.12 Theorem (Addition algebra and its local factors) . . . 77
9.13 Proposition (Polytope realization and external geometric inputs) . . . 78
9.14 Corollary (Jordan blocks and inverse facets) . . . 78
9.15 Theorem (Full filling and exact-output input polynomials) . . . 79
9.16 Corollary (Posterior source counts) . . . 79

Poissonized IDLA and fluid growth

9.17 Theorem (Fluid representation) . . . 80
9.18 Theorem (Chambers and labeled exit counts) . . . 80
9.19 Theorem (Log-concavity and association) . . . 81
9.20 Proposition (Cover and first-hit times) . . . 82
9.21 Proposition (First filling and retirement) . . . 82
9.22 Proposition (Arrival-count moments) . . . 83
9.23 Proposition (Residual semigroup) . . . 83

Complementary-sleep duality

9.24 Theorem (Positive duality and stochastic reversal) . . . 84
9.25 Proposition (Support and impossibility of a positive target payoff) . . . 85

Part I
Monotonicity and comparison theorems

1 Model, target, and shared background

Let V be finite and let ∂ be an absorbing sink. The jump kernel K is stochastic on $V \cup \{\partial\}$, and every $z \in V$ can reach ∂ through positive jump probabilities. At z, an isolated active particle sleeps with probability $p_z \in (0,1)$ in the instruction representation. With two or more active particles present, a sleep attempt has no effect. An arriving particle wakes a sleeper. The sleep rate and jump odds are

$$\lambda_z = \frac{p_z}{1-p_z}, \qquad r_z = \frac{1-p_z}{p_z} = \lambda_z^{-1}.$$

Some records use $q_z = 1 - p_z$; this local notation is distinguished from the occupation function below. Self-loops are allowed unless explicitly removed.

For a finite all-active input $\eta \in \mathbb{N}^V$, let $\mathcal{S}(\eta)$ be its final sleeping support and write

$$q_v(\eta) = \mathbb{P}\{v \in \mathcal{S}(\eta)\}.$$

The central conjecture is

$$q_v(\eta + \delta_x) \geq q_v(\eta) \qquad (x, v \in V).$$

This is an order on a Bernoulli marginal. Stochastic inclusion of the entire sleeping support is stronger. Claims about an arbitrary initial configuration of sleeping particles, or about a stationary law when the sink changes, are different statements and are identified explicitly.

A binary input $C \subseteq V$ means one initially active particle at each site of C. The binary stabilization matrix is

$$P(C,B) = \mathbb{P}\{\mathcal{S}(\mathbf{1}_C) = B\}, \qquad C, B \subseteq V.$$

Matrices act on row measures from the right, and on payoff columns from the left. In particular $P\mathbf{1}_v$ is the column of target probabilities. Physical addition A_x adds one active particle to a sleeping support and stabilizes. "Fresh tails" refers to the unexposed independent instruction suffixes after a legal stopping line.

The abelian property, least-action principles, finite absorbing-walk potential theory, and ordinary IDLA are background inputs. The recorded surplus-IDLA reduction is especially useful: route surplus particles until there is at most one active particle per site, keeping fresh suffixes for later stabilization. IDLA can couple the resulting binary supports increasingly when input is added. Hence a binary add-one theorem on every nested pair implies the corresponding theorem for all finite all-active inputs. Applying a theorem only to binary supports of size at most two gives a bounded-particle statement; it does not by itself settle every input size.

Source keys. RN means `response-renewal.md`; DI means `drain-inventory.md`; SX, CX, and AX mean `stack-exploration.md`, `creator-exchange.md`, and `avalanche-exchange.md`. Older keys such as CB, R, VA, and DF are resolved in the record register. Equation numbers prefixed by letters, such as RN (J113), refer to that source, not to an equation number in this document.

2 Sleep-parameter regimes and sink-strength comparisons

This section collects the parameter results proved in the AI working records. The main conclusions concern adding active particles at *fixed* routing and sleep parameters: a small-sleep theorem for every finite input, and a high-sleep theorem for at most two common particles. Parameter transport compares different sleep vectors with the jump kernel fixed. A separate sink-reinjection flow proves strict stationary total-count order at every positive sleep rate.

Result and location	**Scope and status**
Theorem 2.1	All finite all-active inputs; a graph-dependent heterogeneous small-sleep cube; proved.
Corollary 2.2	Target-component localization and a weaker target-specific graph condition; proved.
Theorem 2.4	Creation-cone measures with at most two sleepers; effective off-target sleep probabilities at least 1/2, or specified terminal exceptions; proved.
Corollary 2.5	Any finite added batch under the high-sleep condition, with the lower measure still in that two-particle cone; proved.
Proposition 2.7 and Lemma 2.8	Positive sleep transport, weighted occupation response, and the all-sleep boundary; proved identities with limited order consequences.
Theorem 2.10 and Proposition 2.12	Global sink reinjection: strict stationary total-count order at all positive site sleep rates, and a positive Markov evolution; proved.

The provenance is the AI-generated checkpoint, evidence index, maintained response-renewal notes, and the separately cited negative-association report for sink reinjection. Proof and AI-audit labels below report those records' status; neither means independent human verification or universal bibliographic novelty. Abelian stabilization and ordinary IDLA are background inputs. The stationary and intertwining interpretations credited in the notes to Levine–Liang are also background, not claimed AI originality.

Record: checkpoint, Discipline and evidence; evidence index, provenance; RN, opening conventions and Section 29; source policy.

2.1 Local notation and proof inputs

Let V be finite and s an absorbing sink. The transition kernel K on $V \cup \{s\}$ has a positive-probability path from every vertex to s; self-loops and zero transition probabilities are allowed. A lone active particle at i sleeps with probability $p_i \in (0,1)$; otherwise it jumps according to $K(i,\cdot)$. A sleep attempt with at least two active particles present is ineffective. Write

$$r_i = \frac{1-p_i}{p_i}, \qquad G = (I - K_V)^{-1}, \qquad h_i(j) = \mathbb{P}_i(T_j < T_s), \quad h_j(j) = 1.$$

Here K_V is the internal transition matrix, and the hitting time counts an initial visit. The input $\eta \in \mathbb{Z}^V_{\geq 0}$ is finite and entirely active. Let μ^p_η be its terminal sleeper-set law and $q_v(\eta;p) = \mu^p_\eta\{B : v \in B\}$. For a subset A, μ_A abbreviates $\mu_{\mathbf{1}_A}$. Point mass at a sleeping support B is δ_B.

The stochastic matrix A_x adds one active particle at x to a sleeping support and stabilizes; $A_s = I$. Put $J_x = K(x,s)I + \sum_{y\in V} K(x,y)A_y$. These operators commute and satisfy $A_x^2 = A_x J_x$. The binary stabilization matrix is $P_r(A,B) = \mu^p_A(B)$; its complete-block inverse is E_r.

In matrix products, $\mathbf{1}$ is the constant-one column. For a vertex set U, $K(a,U)=\sum_{u\in U}K(a,u)$; δ_x in an active input means one particle at x.

Surplus IDLA holds one active particle at each occupied site and moves only surplus particles to a vacancy or the sink. It represents $\mu_\eta^p=\nu_\eta P_r$, where ν_η is independent of p. The IDLA supports for η and $\eta+\delta_x$ have an inclusion coupling, and the subsequent ARW uses fresh suffixes. Thus target order for all binary inclusions implies target order for all finite active inputs. The full binary input V is stationary for every A_x: first drain the surplus from $V+\delta_x$, then stabilize V.

Record: CB, Sections 2–3 and 5; RN, Section 8, opening paragraph; proved algebra and reduction, using standard Abelian stabilization and IDLA.

For $i\in B$, let $W_i(B,C)$ be the subprobability of ending at C after removing the resident at i, forcing its first K-jump, and stabilizing the remaining sleeping background, discarding histories that visit i. In particular $i\notin C$. Define

$$L_i(B,C)=\begin{cases}\mathbf{1}_{\{B=C\}}, & i\notin B,\\ -r_iW_i(B,C), & i\in B,\end{cases}\qquad i\notin C,$$

and the creation cones

$$\mathcal{C}=\{\mu\geq 0:\mu L_i\geq 0 \text{ for every } i\},\qquad \mathcal{C}_2=\{\mu\in\mathcal{C}:\mu(B)=0 \text{ if } |B|>2\}.$$

Measures need not be normalized. Every all-active creation law belongs to $\mathcal{C}$, but membership is not equivalent to being a mixture of such laws. Let $\widehat{A}_x^i$ be addition killed on any visit to i, with $\widehat{A}_i^i=0$. The maintained identities are

$$\mu_\eta L_i=N_\eta^i\geq 0,\qquad A_xL_i=L_i\widehat{A}_x^i,\qquad \mu\geq\frac{\mu(C)}{\mu_C(C)}\,\mu_C\quad(\mu\in\mathcal{C}).\tag{1}$$

Here N_η^i is the unnormalized terminal law on histories never visiting i, and the last inequality is entrywise. The first identity changes the last target sleep to its waiting jump, with weight ratio r_i, retaining only a continuation avoiding i; changing the last target jump back to a sleep is the inverse. Intertwining follows by adding x before or after stabilization and cancelling P_r. Minorization follows by optional stopping of the excessive density μ/π under backward root restoration, where $\pi=\mu_V$. The construction and its proof are given in Theorems 5.1 and 5.2.

Record: RN, Section 6, (7); Section 7, (9); Section 8, (11)–(14); proved inputs; the restoration proof is reproduced in Section 5.

One further input is joint root retention:

$$\mu_{\mathbf{1}_R+\zeta}\{B:R\subseteq B\}\geq\mu_R(R)\quad(R\subseteq V,\ \zeta\in\mathbb{Z}_{\geq 0}^V).\tag{2}$$

Theorem 5.5 proves this using the fresh-book entrance-mass comparison. It compares to the binary-root baseline, not two arbitrary larger inputs.

Record: RN, Section 16, (47a)–(47c); proved, separate AI audit; root-retention proof reproduced in Theorem 5.5.

2.2 All-input monotonicity in a heterogeneous small-sleep cube

Theorem 2.1 (Small-sleep cube)**.** *Suppose K_V's positive-edge graph is strongly connected and*

$$h_i(j)<1\quad\textit{for every distinct } i,j\in V.\tag{3}$$

Equivalently, deleting any one vertex leaves every remaining vertex with a path to the sink. For $n = |V| \geq 2$, set

$$M = \max_i \mathbb{E}_i T_s, \quad h = \min_{i \neq j} h_i(j), \quad e = \min_{i \neq j}(1 - h_i(j)), \quad \kappa = h^3 e^{n-1}, \quad \delta(K) = \frac{\kappa}{768 n^4 M^3}.$$

For every $p \in (0, \delta(K))^V$, every finite all-active η, and every $x, v \in V$,

$$q_v(\eta + \delta_x; p) \geq q_v(\eta; p).$$

For every binary $A \subsetneq V$, $x \notin A$, and target v, the comparison between A and $A \cup \{x\}$ is strict. The probabilities vary independently: no comparability is required. The threshold is uniform over inputs and targets on this fixed graph, not over graphs. For $n = 1$, monotonicity holds at every interior sleep probability.

Record: RN, Section 26, A small-sleep theorem under a hitting hypothesis, (127t)–(127w), (127z); proved; two independent AI proof audits, with a separate audit of the quantitative bound.

Proof. For any fixed finite input, the raw absorption chain is finite. At zero sleep the potential $\sum_i n_i \mathbb{E}_i T_s$, counting active and sleeping particles, drops by one in conditional expectation at each jump. Thus its transient matrix has an inverse at $p = 0$, and terminal probabilities are analytic near zero. Occupying an exact terminal support C requires a sleep at each $c \in C$, so $\mu_\eta(C)$ is analytically divisible by $p_C = \prod_{c \in C} p_c$.

Set all sleep probabilities outside C to zero. Trace ordinary excursions onto C, sending killed excursions to s. The trace kernel K^C satisfies $(I - K^C)^{-1} = G_{CC}$, with sleep probabilities on C unchanged. IDLA with capacity one on C and zero elsewhere produces a binary active subset of C. Only the full subset can finish with all of C occupied. If $H_C^K(\eta)$ is the probability that this IDLA fills C, processing every input even after losses, then on this face

$$\mu_\eta(C) = \frac{H_C^K(\eta)}{Z_C}.$$

Here Z_C is the trace graph's full inverse diagonal. The maintained determinant identity, with $d_D = \det(I - K_{DD}^C)$, $d_\varnothing = 1$, and $r_U = \prod_{i \in U} r_i$, is

$$Z_C = d_C \sum_{D \subseteq C} \frac{r_{C \setminus D}}{d_D}.$$

Hence $p_C Z_C \to d_C = 1/\det G_{CC}$. Analytic divisibility identifies the leading coefficient in all parameter directions:

$$\mu_\eta(C) = p_C\big[\det(G_{CC}) H_C^K(\eta) + O(\|p\|)\big]. \tag{4}$$

This is for each fixed input; uniformity over input sizes comes later.

Fix binary A, $x \notin A$, and v. The quotient

$$F(p) = \frac{q_v(A \cup \{x\}; p) - q_v(A; p)}{p_v}$$

extends analytically through $p_v = 0$. If $v \notin A$, singleton IDLA fills v exactly when at least one source walk hits it. Thus

$$F(0) = G_{vv} h_x(v) \prod_{a \in A} (1 - h_a(v)) > 0. \tag{5}$$

If $v \in A$, this constant cancels. Put

$$w_v(B) = W_v(B, \cdot)\mathbf{1}, \quad S_{vv} = \frac{1}{1 + r_v w_v(\{v\})}, \quad a_v(B) = r_v S_{vv}[w_v(\{v\}) - w_v(B)].$$

The released-resident comparison gives $w_v(B) \le w_v(\{v\})$, hence $a_v(B) \ge 0$ and $a_v(\{v\}) = 0$: replay the smaller background's forced history while holding any extra resident's emitted packet, then finish. Summing deletion gives

$$q_v(\eta) = S_{vv}[1 - N_\eta^v \mathbf{1}] + \sum_{\substack{B \ni v \\ |B| \ge 2}} a_v(B)\mu_\eta(B). \tag{6}$$

Both no-visit terms vanish for a common target root. Let $\rho_{v,c}$ be the probability that a forced departure from v hits $c \ne v$ before returning to v or hitting s. At zero sleep, a resident at c changes the forced packet's target-hitting probability by $\rho_{v,c}h_c(v)(1 - h_c(v))$: upon contact there are two ordinary active walks, whose avoidance probabilities multiply. The two-site trace gives $\rho_{v,c} \det G_{\{v,c\},\{v,c\}} = G_{vc}$, while $r_v S_{vv} \to G_{vv}$. Park the common particle at v first in target-capacity IDLA. The increase in filling c is $h_x(c) \prod_{a \in A \setminus \{v\}} (1 - h_a(c))$. The pair contribution in (6) consequently gives

$$\partial_{p_c} F(0) = G_{vc}G_{cv}(1 - h_c(v))h_x(c) \prod_{a \in A \setminus \{v\}} (1 - h_a(c)) > 0, \qquad c \notin A. \tag{7}$$

Larger supports have higher total degree.

An exact vanishing face handles coordinates of very different sizes. Set $p_{V \setminus A} = 0$, retaining small arbitrary p_A. Tracing onto A makes its full active law stationary. The extra walk is lost before reaching A or added there, with the same terminal law in either case. Thus $F = 0$ on this face. Continuity keeps every derivative in (7) positive on a sufficiently small cube. Turn on the outside coordinates one at a time and integrate to obtain $F > 0$. Equation (5) handles the other case.

For the explicit threshold, include all raw transient states with at most n particles. On $0 \le p_i \le 1/2$, the potential drops by at least $1/2$ per instruction. The fundamental matrix has row-sum norm at most $L = 2nM$. For a terminal payoff in $[0, 1]$, the first derivative solves a Poisson equation whose forcing is the difference of sleep and jump continuation values, of norm at most one. Each derivative of the transient matrix has absolute row-sum norm at most two. Successive differentiation bounds the first, second, and third mixed derivatives by $L, 4L^2, 24L^3$. Subtracting two initial rows doubles these bounds. If D denotes their payoff difference, then

$$F(p) = \int_0^1 \partial_{p_v} D(tp_v, p_{V \setminus \{v\}})\, dt.$$

On a cube of side $\delta \le 1/2$, integration gives

$$|F(p) - F(0)| \le 32n^3M^2\delta,$$
$$|\partial_{p_c} F(p) - \partial_{p_c} F(0)| \le 384n^4M^3\delta, \qquad c \ne v.$$

Using $G_{ij} = h_i(j)G_{jj}$ and $G_{ii} \ge 1$, each needed positive coefficient is at least κ. Our $\delta(K)$ makes the second error at most $\kappa/2$ and the first at most $\kappa/(24nM)$, preserving every sign.

There are finitely many binary comparisons. Successive inclusions therefore order each target. Coupled surplus IDLA then proves the non-strict assertion for every finite input with this same threshold. On one vertex, all nonempty active inputs share the stationary stabilization law, which proves the remaining case. □

Record: RN, Section 15, (44), for the trace inverse diagonal (proof reference, not reproduced); Section 26, (127), (127u)–(127w), (127z), for the small-sleep proof.

Corollary 2.2 (Localization and a single-target weakening)**.** *Let S be the strongly connected component of a fixed target v, and kill jumps on exit from S. It suffices that this local kernel obey*

$$h_i(d) < 1 \qquad (i \in S \setminus \{v\},\ d \in S \setminus \{i\}).$$

Then all-input monotonicity at v holds for sufficiently small sleep probabilities on S, with arbitrary interior probabilities outside S. For $|S| \geq 2$, the same explicit threshold applies with all constants computed locally and $e_v = \min_{i \neq v,\ d \neq i}(1 - h_i(d))$ in place of e. If $|S| = 1$, no smallness restriction is needed.

Proof. The weaker hypothesis keeps every needed escape factor in (5) and (7) positive: only the target itself may be forced through another vertex. Apply the proof locally. Stabilize the strict ancestors of S first, holding arrivals in S. Common upstream books give ordered emission prefixes for ordered inputs, hence ordered finite arrival vectors. Conditional on those books, the books in S remain fresh. Apply the local theorem and average. Neither upstream parameters nor the arrival count affect its threshold. □

Record: RN, Section 13, Localization to the target's strongly connected component; Section 26, (127y)–(127z); proved, separate AI audits.

Remark 2.3 (Strictness and the low-sleep boundary)*.* Strictness is for proper binary inclusions under the theorem's graph assumptions. It is not asserted for all multiplicities: $\mu_{V+\zeta} = \mu_V$, and an extra particle unable to reach a target component has no effect there. At $p_v = 0$ the target is always vacant. When a vertex is forced through another before the sink, coefficients in (7) can vanish. A positive coefficient on a fixed ray does not alone control a heterogeneous cube.

The saved single-leak-cycle derivation awaits an independent audit. Its scoped claim is: for $i \to i+1$ deterministically when $i < n-1$, and $n-1 \to 0$ with probability $c \in (0,1)$ and otherwise to s, each fixed vector $\lambda_i > 0$ admits a sufficiently small range $p_i = \varepsilon\lambda_i$ giving all-input monotonicity at target $n-1$. The threshold depends on that ray. This is an unaudited derivation, not part of Theorem 2.1; its proof is a *proof reference, not reproduced.* The accompanying fixed-root and pendant-leaf proposals remain unproved.

Record: RN, Section 26, discussion after (127y); Single-leak cycles along fixed positive sleep-rate rays, Fixed-root boundary extension, and One pendant leaf; proved boundary analysis distinguished from unaudited or open extensions.

2.3 High-sleep N2, terminal exceptions, and added batches

Here N2 means the add-one target comparison with at most two common active particles; the cone theorem below is stronger.

Theorem 2.4 (High-sleep N2 on the creation cone)**.** *Fix v, with strongly connected component S. Suppose every $a \in S \setminus \{v\}$ satisfies*

$$r_a(1 - K(a,a)) \leq 1 \quad \text{or} \quad K(a, S \setminus \{v,a\}) = 0. \tag{8}$$

Then $(\mu A_x)f_v \geq \mu f_v$ for every $\mu \in \mathcal{C}_2$ and every source x, where $f_v(B) = \mathbf{1}_{\{v \in B\}}$. In particular this gives the all-active add-one comparison with at most two common particles, including multiplicities. The target's own probability and probabilities outside S are unrestricted in $(0,1)$. Without self-loops the first alternative is $p_a \geq 1/2$.

Record: RN, Section 26, N2 from killed sleep insertion, (127a), (127c)–(127g); proved, with separate AI audits of the component restriction and terminal-site extension.

Proof. First remove self-loops. Ignoring ineffective holds changes the odds to $r'_a = r_a(1 - K(a,a))$, with routes conditioned on leaving a; stabilization, addition, and creation-cone operators are unchanged. Use reduced parameters below, and the coefficients S_{vv}, a_v defined immediately before (6).

Kill on visiting v, suppress this superscript, and put

$$\widehat{J}_z = K(z,s)I + \sum_{y \neq v} K(z,y)\widehat{A}_y, \qquad \widehat{B}_z = \frac{\widehat{A}_z - (1-p_z)\widehat{J}_z}{p_z}, \quad z \neq v.$$

These operators commute. At a vacant z, $\widehat{B}_z$ inserts a sleeper deterministically. The quadratic relation gives, for $z \in C$,

$$\delta_C \widehat{B}_z = (1-r_z)\delta_C \widehat{J}_z + r_z \delta_{C\setminus\{z\}} \widehat{J}_z^2. \tag{9}$$

It is thus nonnegative and substochastic when $r_z \leq 1$. Let $\nu = \delta_\varnothing \widehat{J}_v$, supported on empty and singleton sets, and write

$$w_D = w_v(D \cup \{v\}) = \nu \prod_{z \in D} \widehat{B}_z \mathbf{1}, \qquad d_z = \mathbf{1} - \widehat{B}_z \mathbf{1}.$$

Conditioning on ν's at most one resident gives, for distinct $a, b \neq v$,

$$\begin{aligned} w_{\{a\}} + w_{\{b\}} - w_\varnothing - w_{\{a,b\}} = \nu(\{a\})\delta_{\{a\}}\widehat{B}_a d_b \\ + \nu(\{b\})\delta_{\{b\}}\widehat{B}_b d_a. \end{aligned} \tag{10}$$

At empty or a singleton outside $\{a,b\}$, both insertions meet vacancies and the mixed difference is zero. At $\{a\}$ it is $\widehat{B}_a(\mathbf{1} - \widehat{B}_b\mathbf{1})$; the other case follows by commutation. In the high-sleep case the terms are nonnegative, proving

$$a_v(\{v,a,b\}) \geq a_v(\{v,a\}) + a_v(\{v,b\}).$$

For the terminal exception, $d_z \geq 0$ at every sleep probability, even when $\widehat{B}_z$ is signed. Indeed, for $z \in C$ put $j_1 = \delta_C \widehat{J}_z \mathbf{1}$ and $j_2 = \delta_{C\setminus\{z\}} \widehat{J}_z^2 \mathbf{1}$. Remove the resident, hold its independent waiting arrow, and replay the first packet's history. Release the held packet at the resident's first contact, or at the end if none occurs. The second complete visited set contains the first, so $j_2 \leq j_1$, and

$$d_z(C) = 1 - j_1 + r_z(j_1 - j_2) \geq 0.$$

The defect vanishes when $z \notin C$. If a cannot reach b without visiting v, both physical rows in (9) avoid b. Thus $\delta_{\{a\}}\widehat{B}_a d_b = 0$, even at low sleep. Condition (8) guarantees this for each potentially signed summand within S. Exits from S never return and do not affect w_v; a triple with an outside site has zero bracket. The inequality for a_v therefore remains valid.

For every pair C, minorization writes $\mu = \alpha\mu_C + \sigma$, with $\sigma \geq 0$, $\sigma(C) = 0$, and at most two sleepers. The response of the event $C \subseteq B$ can be negative only at the initial support C. Joint root retention (2) gives

$$\Delta_C = (\mu A_x)\{B : C \subseteq B\} - \mu(C) \geq 0.$$

The remainder need not itself lie in the creation cone. Now sum deletion before and after addition. Put $N = \mu L_v \geq 0$. Intertwining, and the fact that the upper law has at most three particles, give the exact identity

$$\begin{aligned} (\mu A_x - \mu) f_v = S_{vv} N(\mathbf{1} - \widehat{A}_x \mathbf{1}) + \sum_{\substack{|C|=2 \\ v \in C}} a_v(C)\Delta_C \\ + \sum_{\substack{|B|=3 \\ v \in B}} (\mu A_x)(B) \left[a_v(B) - \sum_{\substack{C \subseteq B \\ |C|=2,\ v \in C}} a_v(C) \right]. \end{aligned} \tag{11}$$

All terms are nonnegative. □

Record: RN, Section 26, (127c)–(127g); AX, Section 1, released-resident comparison; full high-sleep proof using the explicitly cited established inputs.

Corollary 2.5 (Arbitrary added batch)**.** *If every $a \in S \setminus \{v\}$ satisfies $r_a(1 - K(a,a)) \le 1$, then for every finite nonnegative active batch ζ and every $\mu \in \mathcal{C}_2$,*

$$\mu A_\zeta f_v \ge \mu f_v, \qquad A_\zeta = \prod_x A_x^{\zeta(x)}.$$

The upper particle count is unrestricted. The terminal low-sleep exception of Theorem 2.4 is not included.

Proof. The pair-event proof works with A_ζ. For an off-target set D within the component, condition ν on its singleton support. Unless it is $\{a\}$ for some $a \in D$, every insertion meets a vacancy and has no loss. Otherwise commute $\widehat{B}_a$ first. Subsequent factors are nonnegative and substochastic and cannot increase row mass. Therefore

$$w_\varnothing - w_D \ge \sum_{a \in D} (w_\varnothing - w_{\{a\}}).$$

This proves the nonnegative bracket for every $|B| \ge 3$. Use (11) with A_ζ, summing over all such B, and first term $S_{vv}(\mu L_v)(\mathbf{1} - \widehat{A}_\zeta \mathbf{1})$. This is a direct batch argument, not iteration through intermediate $\mathcal{C}_2$ laws. □

Record: RN, Section 26, Bulk additions under the high-sleep hypothesis, (127p); proved, separate AI audit.

Remark 2.6 (Limits and strictness of the high-sleep theorem)*.* The conclusion is non-strict target order, not occupied-set domination. A terminal component after deleting v is insufficient: the exception is the individual-vertex condition (8). A bidirectional star's central target is covered at arbitrary interior probabilities. The decomposition is strict whenever a nonnegative term is positive; no general strictness classification is asserted. A one-site nonempty input already at stationarity gives equality.

The bracket can fail below the threshold. For $0 \to 1, s$ equally, $1 \to 2$, $2 \to 0$, and $p = (1/2, 1/4, 1/2)$, at target 0,

$$(w_\varnothing, w_{\{1\}}, w_{\{2\}}, w_{\{1,2\}}) = (13/16, 9/16, 5/8, 1/2).$$

The mixed difference is $-1/8$. This exact obstruction concerns the proof's coefficient budget, not N2.

Record: RN, Section 26, discussion after (127e) and (127g); proved scoped obstruction.

2.4 Sleep transport and the near-one boundary

Proposition 2.7 (Positive transport with fixed routing)**.** *For finite odds $0 \le r \le s$, $T(r,s) = E_r P_s$ is stochastic:*

$$T(r,s)T(s,t) = T(r,t), \qquad P_s = P_r T(r,s).$$

From support C, independently wake resident i with probability $1 - r_i/s_i$, then stabilize at odds s; if $r_i = s_i = 0$, leave it unchanged. Its infinitesimal generator is

$$G_x(r)(C,\cdot) = \begin{cases} p_x[\delta_{C\setminus\{x\}} J_x(r) - \delta_C], & x \in C, \\ 0, & x \notin C. \end{cases} \tag{12}$$

Along a coordinatewise nondecreasing differentiable odds path, this gives a Markov coupling whose particle count never increases. For every fixed all-active input and $f \geq 0$, the occupation mean weighted by Gf is nonincreasing. Neither conclusion asserts individual-target or sleeper-set order.

Record: VA, Section 1, (1)–(4); RN, Section 29, (148)–(150a); proved, independent AI audits.

Proof. In the addition algebra at odds s, put $B_i(u) = (1 + u_i)A_i - u_i J_i$. For $s_i > 0$,

$$B_i(r) = \frac{r_i}{s_i} B_i(s) + \left(1 - \frac{r_i}{s_i}\right) A_i.$$

The complete-block formula expresses row C of $E_r P_s$ as $\delta_\varnothing \prod_{i \in C} B_i(r)$. Coordinates with $r_i = s_i = 0$ use the stated boundary convention. Expand the positive mixture and commute sleeper insertions first: this is precisely the independent-waking construction. The cocycle follows by cancelling $P_s E_s = I$.

In the fixed IDLA algebra, row C of E_r is $\delta_\varnothing \prod_{i \in C}[D_i + r_i(D_i - J_i^D)]$, where D_i is ordinary IDLA addition and $J_i^D = K(i,s)I + \sum_y K(i,y)D_y$. Differentiate this multiaffine formula and multiply by P_r. At vacant x, $A_x = p_x B_x + (1 - p_x)J_x$; differentiating $E_r P_r = I$ therefore gives (12). Each event releases a resident and conserves particles except for sink losses. The path generator is $\sum_x r_x'(t) G_x(r(t))$.

Let $U(C, x; z)$ count expected departures from z during the release avalanche, including the first forced jump. Particle balance gives

$$\mathbb{E}[\mathbf{1}_{\text{terminal}}] - \mathbf{1}_C = -U(C, x; \cdot)(I - K_V).$$

Here the two indicator vectors are occupation rows. Surplus IDLA writes $\mu_\eta^r = \nu_\eta P_r$ with fixed ν_η, so $\partial_{r_x} \mu_\eta^r = \mu_\eta^r G_x$. Multiplying by G,

$$\partial_{r_x} q(\eta; r)\, G = -p_x \mathbb{E}_{\mu_\eta^r}[\mathbf{1}_{\{x \in C\}} U(C, x; \cdot)] \leq 0. \tag{13}$$

Here q is the row of occupation means. Multiply by $f \geq 0$. These weights are precisely the nonnegative excessive potentials $a = Gf$, with $(I - K_V)a \geq 0$. In particular $a = \mathbf{1}$ measures total occupation and $a = G\mathbf{1}$ the departure potential. □

Lemma 2.8 (All-sleep tangent and one-coordinate interpolation)**.** *At $r = 0$, lone particles sleep immediately and surplus particles make compulsory K-jumps. This is the continuous infinite-sleep-rate extension, not a raw stack consisting only of sleep instructions. Then $P_0 = I$, and each generator in (12) is attractive for inclusion. If only r_a increases from r_a^0 to r_a^1, put $t = p_a^0(r_a^1 - r_a^0)$. Let R_a^0 fix supports avoiding a and otherwise release its resident with a forced jump and stabilize at the base parameters. Then*

$$E_{r^0} P_{r^1} = [I + t(I - R_a^0)]^{-1} = \frac{1}{1+t} \sum_{k \geq 0} \left(\frac{t}{1+t}\right)^k (R_a^0)^k. \tag{14}$$

Proof. At the all-sleep boundary, a released particle walks to the first vacancy in $C \setminus \{a\}$, including a return to a, or the sink. Other residents remain. Use the same walk for nested supports. The lower first vacancy is encountered no later; its landing site is either already occupied above or is also the upper first vacancy. If only the upper support contains a, removing it cannot remove a lower resident. This couples each event increasingly.

The inverse is affine in r_a. Its derivative multiplied by P_{r^0} is $p_a^0(I - R_a^0)$, by the vacant-site first-step identity. Thus $E_{r^1} P_{r^0} = I + t(I - R_a^0)$. Invert and expand the convergent geometric series of a stochastic matrix. □

Record: RN, Section 26, (127h); Section 29, The all-sleep tangent is long-range exclusion; Section 37, (J113), for base $p_a^0 = 1$; proved identities; long-range exclusion is credited background.

Remark 2.9 (What near-one information proves). High-sleep N2 already gives a whole neighborhood of the all-sleep corner, with threshold $1/2$ after self-loop removal. At the corner, all-input stabilization is surplus IDLA and preserves inclusion. A tangent or endpoint comparison does not prove all-input order in a neighborhood. No separate general all-input near-one perturbation theorem was found in the maintained records reviewed here.

Away from the boundary the transport is not attractive. On two sites with jumps equally to the other vertex and the sink, at $r = (1, 1)$,

$$(\delta_{\{0,1\}} - \delta_{\{0\}})G_1 f_0 = -22/195.$$

The initial target bits agree. Slow sleep interpolation can therefore have a negative instantaneous inclusion comparison. Intermediate resident-release laws also need not stay in the reference creation cone, so (14) does not justify iterating high-sleep N2 down to arbitrary sleep probabilities.

Record: RN, Section 29, (151); Section 26, discussion following (127s); proved obstructions to those continuation arguments.

2.5 Sink-strength interpolation and stationary total-count order

The interpolation argument is recorded in the AI negative-association project on 20 October 2026. Its comparison concerns the stationary *number* of sleepers, with all sleep rates fixed.

Write K_V for the internal substochastic kernel and $k_u = K(u, s)$ for its killing vector. Fix a probability distribution ρ on the nonempty finite set V, with $n = |V|$, and fix all site sleep rates $0 < \lambda_u < \infty$. For $0 \le z < 1$, replace each would-be absorption, independently with probability z, by a new jump with destination law ρ. The new kernel is

$$K_{z,V} = K_V + zk\rho^{\mathsf{T}}, \qquad k_z = (1 - z)k. \tag{15}$$

Here k and ρ are columns. Every K_z is killed: successive reinjections concatenate a geometric number of original killed walks. Let ν_z be the saturated stabilization law, which is stationary under every single-site addition. Let $T_{v,z}$ be the addition matrix at v, and put

$$T_z = \sum_v \rho_v T_{v,z}, \qquad N(S, S) = |S| \quad (S \subseteq V).$$

As elsewhere in this document, measures are rows; thus $\nu_z T_z = \nu_z$. The source uses the transpose convention. No irreducibility or full-support assumption on the drive ρ is required.

Theorem 2.10 (All-rate stationary count order under sink reinjection). *For the family* (15)*, the full stationary law obeys*

$$\frac{d\nu_z}{dz} = \frac{\nu_z N - \nu_z N T_z}{1 - z}. \tag{16}$$

Let $X \sim \nu_z$, and let Y be the result of one independent ρ-addition to X, using K_z. Set $A = |X|$ and $B = |Y|$. For every $1 \le r \le n$, with $c = r - \frac{1}{2}$,

$$\frac{d}{dz}\nu_z\{|S| \ge r\} = \frac{\mathbb{E}[(A - c)(\mathbf{1}_{\{A \ge r\}} - \mathbf{1}_{\{B \ge r\}})]}{1 - z} > 0. \tag{17}$$

Consequently the stationary count law is strictly stochastically increasing in z. Equivalently, every nonconstant increasing function of total count has strictly decreasing expectation as the killing factor $1 - z$ increases within this family. In particular,

$$\frac{d}{dz}\mathbb{E}_{\nu_z}|S| = \frac{\mathbb{E}[(B - A)^2]}{2(1 - z)} > 0. \tag{18}$$

Proof. We first differentiate by marking a sink departure. For $k_z(u) > 0$, let L_u be the number of particles absorbed from u during saturated stabilization, and define the finite row measure

$$F_{u,z}(S) = \frac{\mathbb{E}[L_u \mathbf{1}_{\{\mathcal{S}(\mathbf{1}_V)=S\}}]}{k_z(u)}.$$

The response to replacing one unit of killing probability at u by a jump to v is

$$\partial_{K(u,v)-K(u,s)}\nu = F_{u,z}(T_{v,z} - I), \tag{19}$$

where the derivative notation means increasing the internal entry and decreasing the sink entry at equal speed. To see this, mark one absorbing departure, retain that particle off the graph, and finish the rest of the stabilization. The positive part of the derivative returns the marked particle as a new addition at v; the negative part leaves it absorbed. Abelianity permits this postponed addition, and the unused stack suffixes are independent of the exposed marked history. Summing over all possible marks gives (19). The same identity follows by differentiating the finite absorbing-chain resolvent. This also justifies differentiability, including the right derivative at $z = 0$; compact subintervals of $[0, 1)$ stay in the absorbing regime.

There are initially n particles. Conservation therefore gives the exact identity of measures

$$\sum_u k_z(u) F_{u,z} = \nu_z(nI - N).$$

The speed of the perturbation in entry (u, v) is $k_u\rho_v = k_z(u)\rho_v/(1 - z)$. Sum (19) over u, v, omitting rows of zero killing, to obtain

$$\nu_z' = \frac{\nu_z(nI - N)(T_z - I)}{1 - z} = \frac{\nu_z N - \nu_z N T_z}{1 - z},$$

where the last equality uses stationarity.

The random variables A and B have the same marginal law. Applying the flow to the rank threshold gives

$$(1 - z)\frac{d}{dz}\nu_z\{|S| \geq r\} = \mathbb{E}[B\mathbf{1}_{\{B\geq r\}} - A\mathbf{1}_{\{B\geq r\}}] = \mathbb{E}[(A - c)(\mathbf{1}_{\{A\geq r\}} - \mathbf{1}_{\{B\geq r\}})].$$

The integrand is zero without a threshold crossing, and positive for either direction of crossing. To prove strictness, the saturated law has full support: route each unwanted initial particle to the sink while holding one active particle at each desired site, then put the retained particles to sleep. This is a finite legal prescription of positive probability. Choose v with $\rho_v > 0$ and an $(r - 1)$-site set S not containing v. From this positive-mass state, addition at v followed by immediate sleep produces an upward crossing with probability at least $\rho_v\lambda_v/(1 + \lambda_v) > 0$.

Integration proves strict order between any $0 \leq a < b < 1$. Threshold decomposition covers every increasing function of count. Finally the flow gives $(1 - z)\frac{d}{dz}\mathbb{E}|S| = \mathbb{E}[A^2 - AB]$. Since $\mathbb{E}[A^2] = \mathbb{E}[B^2]$, this is half the expected squared increment, proving (18) and its strictness. □

Record: NA-report, "Global reinjection: exact flow and all-rate rank monotonicity," (50)–(52), complete working proof dated 2026-10-20; recovered from the pinned source revision in the record register. The four saved exact reinjection certificates were independently rerun during this recovery.

Corollary 2.11 (A single sink port)**.** *If killing is possible only at u, replacing any positive amount less than k_u of its killing probability by a self-loop strictly increases the stationary count law, with all sleep rates fixed. Increasing killing by reversing this replacement strictly decreases the law.*

Proof. Take $\rho = \delta_u$ in (15). The only changed entries are $K(u,u)$ and $K(u,s)$, and every nontrivial count threshold has positive derivative by (17). □

Proposition 2.12 (A Markov evolution in sink strength)**.** *The family ν_z is realized by a time-inhomogeneous Markov process in the parameter z, with row generator*

$$Q_z = \frac{(nI - N)(T_z - I)}{1 - z}. \tag{20}$$

At state S, it performs a fresh ρ-addition using K_z at rate $(n - |S|)/(1 - z)$. On each interval $[a, b] \subset [0, 1)$, starting with law ν_a produces law ν_z throughout.

Proof. The off-diagonal entries of Q_z are nonnegative, its row sums vanish, and its rates are bounded on every such compact interval. It therefore defines a finite-state Markov evolution. The calculation in the preceding proof gives $\nu_z' = \nu_z Q_z$. Uniqueness of this finite linear differential equation identifies the evolving law. These addition events can decrease the particle count; the stochastic rank order is the distributional conclusion of Theorem 2.10, not pathwise monotonicity of this process. □

Record: NA-report, "Finite reinjection transform," (56)–(58). The source also gives an explicit finite-parameter stochastic transform and its composition law; the generator above follows directly from the proved response identity.

Remark 2.13 (Scope of the recovered sink comparison)*.* The theorem allows directed kernels, self-loops, arbitrary finite positive site sleep rates, and any fixed reinjection law ρ. It compares the saturated stationary count laws along the specified family (15). It does not establish stationary occupied-set inclusion, a specified-site comparison, or a fixed-input comparison for arbitrary changes to killing edges. The older parent add-a-sink proposal asks for a stronger occupied-set order and remains a separate open problem.

3 Directed decomposition, backward duals, and drain inventories

The results below separate exact response identities, occupation inequalities, and stronger finite-inventory inequalities. In particular, a theorem comparing full input with one missing site is not an all-input theorem.

Result guide	Scope and conclusion	Location
Exact directed gluing	Joint responses, including feedback; no order assertion	3.1
Target localization	All finite active inputs when the target SCC has size at most three	3.2, 3.3
Backward duals	Coherent occupied-set representation on DAGs; individual target rows for SCCs of size at most two	3.4, 3.5
Portal comparison	Full versus full-minus-portal occupied-set order for outward forests; every directed cycle	3.9, 3.10
Conditional extension	General feedback budget sufficient for the same comparison	3.11
Finite drain words	Explicit positive inventory families; general inventory sign open	3.6, 3.16
Obstructions	Merging, later entry, and excessive conditioning	3.7

3.1 Conventions and inherited inputs

Let V be finite, s an absorbing sink, and K a jump kernel for which s is accessible from every vertex. Unless an endpoint is explicitly specified, $0 < p_i < 1$, $q_i = 1 - p_i$, and $r_i = q_i/p_i$. A lone active particle sleeps with probability p_i; jump destinations have law $K(i, \cdot)$. Multiple particles prevent successful sleep. Let μ_η be the final sleeping-support law from finite all-active input $\eta \in \mathbb{N}^V$, and $q_v(\eta) = \mu_\eta\{B : v \in B\}$. For binary A, write $\mu_A = \mu_{\mathbf{1}_A}$. The notation $\nu \leq_{\mathrm{st}} \pi$ means comparison for every increasing function on 2^V. An SCC is a strongly connected component of the positive-jump graph.

We use the saved surplus-IDLA reduction: move surplus particles while holding one active particle at each occupied vertex. Ordered inputs give nested binary supports and fresh remaining stacks. Full binary input is stationary under additional active particles, since every surplus can be drained while retaining one active everywhere. We also use the proved one-common-particle theorem (Theorem 5.6), with its all-active scope.

Record: CB, §5; proved input: surplus IDLA. R, §4; proved input: one common particle.

Self-loops can be removed without changing terminal laws: if $a = K(i, i)$, replace p_i by $p_i/(1 - q_i a)$ and every nonself jump probability by $K(i, z)/(1 - a)$. Thus r_i becomes $r_i(1 - a)$. This sums the geometric sequence of lone-particle self-returns; a surplus leaves at its first nonself jump.

Record: RN, §21, paragraph following (81); proved reduction.

3.2 Exact responses and unconditional directed gluing

For a domain D with input sites I and labelled absorbing exits O, assume an exit is accessible from every site. For sleeping support B, let

$$A_i^D(\mathbf{z})[B, C] = \sum_{\mathbf{k}} \mathbb{P}_B(C, \mathbf{k} \mid \text{one active input at } i) \prod_{o \in O} z_o^{k_o}.$$

Here C is the terminal support and $\mathbf{k}$ counts labelled exits. The support and exit vector are retained jointly. Counting every exit, $|C| + \sum_o k_o = |B| + 1$, so each entry is a polynomial of degree at most $|D| + 1$.

Theorem 3.1 (Exact gluing in law)**.** *For every finite input word, the product of its $A_i^D(\mathbf{z})$ is its joint terminal-support and total-exit generating matrix. In particular, these matrices commute at a common $\mathbf{z}$. Disjoint domains can be joined by routing designated exits to designated inputs, including feedback. If the combined graph remains sink-accessible, a finite transient queue chain computes its exact response, and gluing is associative in law. These assertions impose no occupation-order or creation-cone assumption.*

Record: RN, §31, (B1)–(B5); proved identities, with exact audits and main-thread review.

Proof. Sequential and simultaneous additions have the same terminal support and boundary counts by Abelianity. At each complete local stabilization, unread stack tails are independent fresh tails. Conditioning on the revealed support and exits therefore gives matrix multiplication and commutation.

For gluing, retain each component's sleeping support, all pending arrival counts, and the external exit counts. Process one pending arrival by its local response, sampling support and exits jointly; route internal exits back into the queue. With initial mass N, stored sleepers, pending particles, and recorded external exits total N, so the state space is finite. A shared physical vertex must belong to one component, possibly a singleton, rather than being duplicated at two ports.

Delayed messages give a legal schedule. A sleeper with an undelivered arrival represents two active particles once that arrival is inserted. An addition commutes with every toppling legal before it: a sleep followed by an addition and the reverse order both leave two active particles; jumps satisfy the same local identity. Thus inserting delayed arrivals preserves legality and the reconstructed physical configuration. Empty queues and stable component supports give a globally stable configuration.

Let h_x be the expected ordinary-walk time from x to the exterior, and let Φ sum h over all particles, locating pending particles at their destinations. Each raw instruction at x has conditional mean change $-(1 - p_x)$, since $h_x = 1 + \sum_y K(x, y)h_y$. Stopping after finitely many instructions gives expected total instruction count at most $\Phi_{\text{initial}}/\min_x q_x$. Each response call consumes an instruction. Hence the queue chain is transient, with finite mean absorption time, even with arbitrarily many possible returns.

If Q is its transient matrix, S its terminal matrix, and J selects initial states, the response is

$$J(I - Q)^{-1}S = J \sum_{n \geq 0} Q^n S.$$

Every grouping and causal fixed priority computes the same physical ARW response, proving associativity. Equivalently, eliminating a set W of transient states replaces Q_{UU} by $Q_{UU} + Q_{UW}(I - Q_{WW})^{-1}Q_{WU}$; transience justifies all inverses. □

The stochastic-processor interpretation has the credited antecedent Bond–Levine, *Abelian networks I*, §3.9 and Lemma 4.7, as recorded in RN §31. The finite iid-tail response construction above specifies the part used here. It is not a deterministic compression of already exposed stacks, where consumed indices must also be retained.

3.3 Target SCC localization and the three-site certificate

Theorem 3.2 (Localization)**.** *Let B be the SCC containing v. If the kernel killed on leaving B satisfies $q_v(\xi + \delta_x) \geq q_v(\xi)$ for every finite all-active ξ on B and every $x \in B$, then the full*

graph satisfies $q_v(\eta') \geq q_v(\eta)$ for every finite all-active $\eta \leq \eta'$. Upstream SCCs have no size restriction.

Record: RN, §13, "Localization to the target's strongly connected component"; proved and separately AI-audited. Dependency: §11, (17).

Proof. Let A be the vertices outside B that can reach B. No edge goes from B to A, or from outside $A \cup B$ into $A \cup B$; exits from B never return. Stabilize A while holding all arrivals in B. In the buffer representation, after n arrivals at z the emitted count is $n - b_z(n)$, with $b_z(0) = 0$ and $b_z(n) \in \{0, 1\}$. This function is nondecreasing. Least action for these local processors therefore orders emission counts under the common books when input increases. The emitted destination lists are prefixes of the same lists. Consequently

$$\xi_\eta(b) = \eta(b) + \sum_{z \in A} \sum_{k \leq m_z^\eta} \mathbf{1}\{d_z(k) = b\} \leq \xi_{\eta'}(b).$$

Conditional on the entire upstream books these inputs are fixed, while B's books remain independent and fresh. Iterate the assumed local add-one inequality and average. Vertices outside $A \cup B$ cannot influence v. □

Theorem 3.3 (Three-site occupation and pair capacity)**.** *On every sink-accessible graph with at most three vertices, all finite all-active inputs satisfy the single-target add-one inequality. Thus the same holds on any finite graph whenever the target SCC has size at most three. On three vertices, put $R_i = V \setminus \{i\}$, $\pi = \mu_V$, and $\gamma_i = \mu_{R_x}(R_i)$, $\beta_i = \pi(R_i)$. Then*

$$\pi(V) \geq \sum_i (\gamma_i - \beta_i)_+. \tag{21}$$

For $v \neq x$ on three vertices, the full-minus-x target increment is strict exactly when v and x are in the same SCC. This strictness statement does not concern every arbitrary-input comparison.

Record: RN, §25, (115)–(124) and final reduction; computer-assisted theorem.

Reduction and certificate method; coefficient proof reference, not reproduced. Remove self-loops. Last-arrival deletion gives an explicit nonnegative three-state pair renewal $(I - \mathcal{R})\beta = z/Z_V$ and $(I - \mathcal{R})\gamma = e_x/Z_x$. Here $Z_V^{-1} = \pi(V)$, Z_x^{-1} is full occupation on the killed two-site domain, and the entries are the explicit two-site rational quantities in RN (115). The actual positive full-input pair vector satisfies $\mathcal{R}\beta < \beta$, proving transience without an ordinary row-sum bound. Clearing the denominators of $\pi(V) + \sum_{i \in J}(\beta_i - \gamma_i)$ produces polynomials H_J with positive expansions in sleep odds, internal jumps, and sink probabilities. Any nonempty positive-deficit family contains x: the off-source renewal equations make all other deficits nonpositive if the source deficit is nonpositive. Relabelling therefore leaves the three families of sizes one, two, and three containing x. RN (120)–(121) certifies all three and proves (21).

For an old target, deletion extracts its increment as $\sum_i a_i(\beta_i - \gamma_i) + a_*\pi(V)$, with $0 \leq a_i \leq a_*$. Equation (21) makes this nonnegative. The added target has the general positive first-contact argument in RN (90)–(91). Surplus IDLA reduces arbitrary inputs to nested binary supports: empty support is immediate, a singleton uses the inherited one-common theorem, a pair uses this certificate, and full support is stationary. Theorem 3.2 gives the claimed extension.

Independently, RN (122)–(124) specifies a scalar numerator H for target 0, source 1. Its 60 sleep-odds blocks have nonnegative coefficients after the row substitution $K_{ij} = A_{ij}/(1 + \sum_{k \neq i} A_{ik})$. The verifier reconstructs H, checks an independent numerator identity, and inverts every substitution; there are 330,334 positive coefficients. Clearing row denominators gives

identities valid also at zero sink probabilities. All original denominators are positive by walk and renewal transience. Every numerator monomial contains a mutual-reachability edge product, and $H(K,r) \geq d_2(K_{01}+K_{02}K_{21})(K_{10}+K_{12}K_{20})$, where $d_2 = 1-K_{01}K_{10} > 0$. This yields the exact strictness criterion.

The finite polynomial expansions, exact divisions, and their coefficient tables are an essential computer-assisted proof component, not reproduced here. They are not finite parameter tests. □

Both certificates are maintained in . The pair-capacity verifier and artifact are

```
verify_three_site_pair_deficit.py
three-site-pair-deficit-certificate.json.gz
SHA256:
c6e4487ba10102d45006515e1063233077e76cf3f58ec0fc153c28ae46a9bb67
```

The independent scalar verifier and artifact are

```
verify_three_site_certificate.py
three-site-positive-certificate.json.gz
SHA256:
5821a322d9ee946df6abbb07534063869706247b8925f93a7b9f621d13587070
```

The recorded symbolic version is SymPy 1.14.0. These artifact hashes were checked for this compilation; the expensive symbolic verifiers were not rerun.

3.4 Backward duals and composition without feedback

Here every vertex has one base particle, asleep exactly on S and active elsewhere; η denotes additional active particles. For target set T and exit weights $t_e \in [0,1]$, let $F_{T,\eta}(S;t)$ be the expectation of $\mathbf{1}\{T \subseteq B\}\prod_e t_e^{N_e}$. Put $h_C^t(i) = \mathbb{E}_i[t_{\text{exit}};\ \text{exit before visiting } C]$; a starting visit counts, and $0^0 = 1$.

Theorem 3.4 (Coherent DAG dual)**.** *On an acyclic internal graph there is, for each T, t, a probability law $\Theta_{T,t}$ on subsets, independent of S, η, such that*

$$F_{T,\eta}(S;t) = \sum_{C \subseteq S} \Theta_{T,t}(C) \prod_i h_C^t(i)^{\eta_i}. \tag{22}$$

At $t = 1$ a common random graph realizes all target sets coherently, and gives the whole occupied-set law simultaneously for every S, η.

Record: RN, §23, (94)–(98); proved and independently AI-audited.

Proof. Eliminate vertices backward. At a sinkward vertex u let $\ell = \sum_e K(u,e)t_e$, treating already eliminated successors as weighted exits. Conditional on N upstream arrivals, the site has $N + \eta_u + 1$ particles unless its original sleeper remains untouched. A visited site emits $N + \eta_u$ forced jumps and uses a fresh final Bernoulli(p_u) coin, independent of the destinations. Its local weight is

$$\alpha_u \ell^{N+\eta_u} + \mathbf{1}\{u \in S\}\beta_u \mathbf{1}\{N + \eta_u = 0\}, \quad (\alpha_u, \beta_u) = \begin{cases} (p_u + q_u\ell, q_u(1-\ell)), & u \notin T, \\ (p_u, q_u), & u \in T. \end{cases}$$

Thus the first backward branch has weight α_u, assigns incoming boundary value ℓ, and contributes ℓ^{η_u}. The second has weight β_u, marks u, assigns zero boundary value, and requires $u \in S$, $\eta_u = 0$. Branch weights sum to one. Harmonic substitution identifies the final assigned values as h_C^t, proving (22) without renormalizing the extra-particle factors.

For coherence at $t = 1$, independently sample $A_i \sim \mathrm{Ber}(q_i)$ and ordinary walks W_i to the sink. On the vertices with $A_i = 1$, draw $i \to k$ when W_i visits k after time zero. Let H_T be the ancestors of $T \cap \{i : A_i = 1\}$, including those targets. Reverse exposure marks a target with probability q_i, and a non-target with probability $q_i(1 - \ell)$, exactly the preceding recursion. Moreover $H_{T \cup T'} = H_T \cup H_{T'}$. Independently sample a walk for each potential extra particle. Declare v occupied when $H_{\{v\}} \subseteq S$ and all used extra walks avoid it. The union identity and (22) match every joint occupation probability, hence the full terminal law. □

Proposition 3.5 (Weighted rows and small components)**.** *The positive representation* (22) *holds for each fixed T on every graph whose SCCs have at most two vertices, allowing labelled exits, zero exit weights, self-loops, and arbitrary extras. Union-coherence across different T is asserted only for DAGs. For $T = V$, a positive row exists on every finite sink-accessible graph: with $d(U) = \det(I - K_{UU})$, $d(\varnothing) = 1$,*

$$Z = \sum_C r_C \frac{d(V)}{d(V \setminus C)}, \qquad \Theta_{V,t}(C) = \frac{r_C d(V)}{Z, d(V \setminus C)}, \quad r_C = \prod_{i \in C} r_i.$$

This full-target row is independent of the boundary weights.

Record: RN, §22, (89), (93); §23, (99)–(101) and component composition; proved. Local determinant and two-site coefficient proofs referenced below.

Composition; local formulas referenced. The full-target determinant formula is RN (89), derived from the saved full-support inverse and parameter transfer. Weighted exits follow by the Doob transform $K^h(i,k) = K(i,k)h_k/h_i$, $K^h(i,e) = K(i,e)t_e/h_i$, where $h = h^t_\varnothing > 0$. Principal determinants are unchanged by diagonal similarity. History weights telescope by $h^{\mathbf{1}+\eta}/h_B$; on full occupation this leaves h^η, giving the stated row and avoidance factors.

For a general target, independently add each $i \notin T$ to a random set D with probability $1 - h_i$. On $T \subseteq B$, h_V/h_B equals $\mathbb{P}(D \subseteq B)$. Therefore the weighted row is a convex mixture of the unweighted $T \cup D$ rows under K^h. The explicit nonnegative four coefficients for each singleton target on a two-site block are RN (101); the empty target is immediate and the full target uses the formula above. Their algebraic derivation is a proof reference, not reproduced here.

Process the condensation DAG backward. Sample the current component's positive row, mark its chosen subset, and assign its avoidance potentials as boundary weights for upstream arrivals. No downstream component can send particles back, so incoming counts depend only on independent upstream data. This is precisely the block version of the preceding DAG recursion. Doob transforms preserve SCC sizes. Zero potentials and zero weights follow by continuity, and self-loop removal preserves endpoints. □

3.5 Portal cones, forests, and full-minus-one order

Let an open domain R have exits to s and a portal j. Initially assume $h_i = \mathbb{P}_i(s \text{ before } j) > 0$, and put $b_i = p_i/(p_i + q_i h_i)$ and $\rho = \bigotimes_i \mathrm{Ber}(b_i)$. Write $W_i^k(B, C)$ for the following subprobability kernel: from sleeping $B \ni i$, remove its i-resident, force its first jump, stabilize while forbidding a revisit to i, and require exactly k portal exits. For $C \not\ni i$ define

$$(\nu L_i^0)(C) = \nu(C) - r_i \sum_{B \ni i} \nu(B) W_i^0(B, C), \qquad \mathcal{C}^0 = \{\nu \geq 0 : \nu L_i^0 \geq 0 \text{ for every } i\}.$$

Measures in this cone need not be normalized. Let A_y^0 be active addition retaining only zero portal exits, and let A_ζ^0 denote the corresponding kernel for a fixed active batch ζ. Write $h^\zeta = \prod_i h_i^{\zeta_i}$.

Lemma 3.6 (Exact slices and zero-exit additions)**.** *Every exact total-portal-count slice of an all-active creation law lies in $\mathcal{C}^0$. Every zero-exit addition preserves $\mathcal{C}^0$.*

Record: RN, §28, (136)–(137) and Doob transformation; proved, using last-arrival deletion and §7's intertwining.

Derivation from the inherited deletion identity. If μ_H is the slice with H exits and N_H^i its no-i-visit submeasure, last-arrival deletion preserves the exit count additively:

$$\mu_H(C) = N_H^i(C) + r_i \sum_{B \ni i} \sum_{k=0}^{H} \mu_{H-k}(B) W_i^k(B, C), \quad i \notin C.$$

Subtracting the $k = 0$ term proves the first assertion. Under the Doob kernel $\widehat{K}(i,l) = K(i,l)h_l/h_i$, $\widehat{K}(i,s) = K(i,s)/h_i$, portal edges disappear and $W_i^0(B,C) = h_B \widehat{W}_i(B,C)/h_C$, where $h_B = \prod_{i \in B} h_i$. Thus $\nu \in \mathcal{C}^0$ exactly when $h_B \nu(B)$ belongs to the ordinary creation cone. Its active-addition invariance, RN §7, together with

$$h_C(\nu A_y^0)(C) = h_y[(h_B \nu(B))\widehat{A}_y](C),$$

proves the second assertion. This uses the proved deletion identity; its full last-arrival bijection is the proof reference in RN §6–7, not reproduced here. □

Theorem 3.7 (Static DAG bound)**.** *On a DAG, every normalized $\nu \in \mathcal{C}^0$ satisfies $\nu \leq_{\mathrm{st}} \rho$. More precisely, each occupation probability conditional on an earlier topological output prefix is at most b_i. Also $\nu A_y^0 \mathbf{1} \geq h_y \nu \mathbf{1}$ for every cone measure.*

Record: RN, §28, (138); proved by simultaneous induction.

Proof. Remove the first vertex i and split ν into ν_0, ν_1 on the remaining domain. Later deletions cannot change its bit, so both slices belong to the remaining cone. Put $J = K(i,s)I + \sum_{l \neq i} K(i,l)A_l^0$. The i-constraint is $\nu_0 \geq r_i \nu_1 J$. Inductively J preserves the remaining cone and $\lambda J \mathbf{1} \geq h_i \lambda \mathbf{1}$ on that cone. Consequently $\nu_0 \mathbf{1} \geq r_i h_i \nu_1 \mathbf{1}$, giving the first-bit bound. Apply induction separately within each bit slice to obtain every conditional bound; common uniform variables couple these conditional Bernoullis below independent Bernoulli(b_i) variables.

For addition at a later vertex use induction on each slice. Addition at i has survival mass

$$p_i(\nu_0 + \nu_1 J)\mathbf{1} + q_i(\nu_0 J + \nu_1 J^2)\mathbf{1} \geq (p_i + q_i h_i)(\nu_0 + h_i \nu_1)\mathbf{1} \geq h_i(\nu_0 + \nu_1)\mathbf{1}.$$

The last difference is $(1 - h_i)(p_i \nu_0 \mathbf{1} - q_i h_i \nu_1 \mathbf{1}) \geq 0$ by the deletion constraint. The empty domain starts this simultaneous induction. □

Lemma 3.8 (Root entry into an outward forest)**.** *Suppose internal edges of R run only from parents to children in a directed forest, with additional portal and sink exits permitted. Every root-entry batch, conditioned on any exact portal-exit count, preserves $\mathcal{C}^0$. The same holds for sequences and mixtures of root entries and specified successive counts, including correlated initial component supports.*

Record: RN, §28, (139), "Disjoint branches and outward forests"; proved.

Proof. Induct from leaves. At root i put

$$J_i(z) = K(i,s)I + K(i,j)zI + \sum_{c \text{ child of } i} K(i,c)A_c(z),$$

where z counts portal exits. A positive cone-preserving kernel on one disjoint child domain preserves all forest constraints: fix all other configurations for its own constraints, and commute the kernel with deletions in the other domains. No independence of input supports is needed. Inductively every coefficient of every power of $J_i(z)$ therefore preserves the remaining cone.

For an incoming batch $M \geq 1$, splitting the original measure at i gives the two output rows

$$\nu_1(z) = p_i[\mu_0 J_i(z)^{M-1} + \mu_1 J_i(z)^M], \qquad \nu_0(z) = q_i[\mu_0 J_i(z)^M + \mu_1 J_i(z)^{M+1}].$$

All downstream constraints hold coefficientwise. The remaining one is

$$\nu_0(z) - r_i \nu_1(z) J_i(0) = q_i[\mu_0 J_i(z)^{M-1} + \mu_1 J_i(z)^M][J_i(z) - J_i(0)] \geq 0.$$

The displayed factor order uses no commutation at different values of z. For $M = 0$ the kernel is the identity. Products, coefficient convolutions, and mixtures prove the final assertions. □

Theorem 3.9 (Forest feedback)**.** *Adjoin j to an outward forest R, allowing j to jump only to roots or the original sink, and forest vertices to jump to j or the sink. Assume sink accessibility. Then*

$$\mu_R \leq_{\mathrm{st}} \mu_{R\cup\{j\}}.$$

Self-loops are allowed after endpoint-preserving reduction. The two inputs are exactly full-minus-portal and full active input.

Record: RN, §27, (128)–(129); §28, (140); proved occupied-set comparison.

Proof. First $h_i > 0$. Stabilizing the full active forest while requiring no visit to j gives $z_j \rho$, where $z_j = \prod_{i\in R}(p_i + q_i h_i)$. Indeed, process in topological order: each site has a resident; a final occupied bit has unnormalized weight p_i, and an empty bit weight $q_i h_i$. The incoming and outgoing ordinary-walk factors telescope. Lemma 3.6 puts this measure in $\mathcal{C}^0$.

Now start one active particle at j against sleeping background ρ. Condition only on j's untimed local history and successive boundary batch counts. Lemma 3.8 preserves the conditional forest cone through every returning pass. Processing j uses fresh records and observes no additional forest data, so it preserves that property. At the terminating pass, including a zero next batch, Theorem 3.7 bounds the forest support by ρ. The finite-particle sink-accessibility argument of Theorem 3.1 ensures finite mean termination. Thus, for increasing f, the continuation expectation $U_j(f)$ is at most $\mathbb{E}_\rho f(B \cup \{j\})$.

Let $\sigma_j = \mathbb{P}_j(s \text{ before a positive-time return to } j) > 0$. Holding the initial particle at j, histories with a first contact cancel its active/sleeping distinction. The no-contact histories give the exact first-contact and addition identity, RN (88)–(90):

$$r_j \sigma_j[\mu_{R\cup\{j\}} f - \mu_R f] = z_j[\mathbb{E}_\rho f(B \cup \{j\}) - U_j(f)] \geq 0.$$

The identity extends from joint-occupation indicators to all functions by their spanning property. Vanishing h_i are handled by adding vanishing sink leakage within the same graph class and taking the finite transient-resolvent limit. Self-loops use the stated reduction. □

Corollary 3.10 (Every directed cycle)**.** *For every $n \geq 2$, let $K(i, i+1 \bmod n) = a_i$, $K(i,s) = 1 - a_i$, $0 \leq a_i \leq 1$, and $\prod_i a_i < 1$. With arbitrary $0 < p_i < 1$, for every missing site j, $\mu_{V\setminus\{j\}} \leq_{\mathrm{st}} \mu_V$. Self-loops may be inserted. This does not assert arbitrary-input order on cycles of length greater than three.*

Proof. Cutting at j leaves a directed path entered at its root, so apply Theorem 3.9. A one-vertex graph is immediate. □

Record: RN, §27, (128)–(133), independent queue-reversal proof; §28, (140), stronger occupied-set consequence.

For a cycle with exactly one leaking edge, RN (134) also proves every joint-target full-wake coefficient nonnegative: in the order starting just after the leak, $\Theta_T(C) = r_{C\setminus\{j\}}\Theta_T(\{j\})$ for nonempty C and $j = \max C$. Here the coefficients are those of (22) at $t = 1$. *Proof reference, not reproduced:* RN §27, (113), (129), (134). Multiple leaks retain the full-minus-one comparison, not this factorization.

Proposition 3.11 (Static SCC bound and a conditional feedback criterion)**.** *Let π_R^0 be the normalized no-portal full-root law. If every SCC of R has size at most two, then every normalized $\nu \in \mathcal{C}^0$ satisfies $\nu \le_{\mathrm{st}} \pi_R^0$ and $\nu A_\zeta^0 \mathbf{1} \ge h^\zeta$ for each fixed active batch ζ. This permits arbitrary edges between components.*

On an arbitrary domain, let F_j be the complete-feedback kernel obtained by starting one active at j, stabilizing the whole graph, and forgetting j's final bit. The condition

$$\pi_R^0 F_j \le_{\mathrm{st}} \pi_R^0 \tag{23}$$

is sufficient for $\mu_R \le_{\mathrm{st}} \mu_{R\cup\{j\}}$. Static domination together with $\pi_R^0 F_j \in \mathcal{C}^0$ is one sufficient way to verify it, but general feedback cone preservation is false.

Record: RN, §28, "Static domination on a two-site cyclic domain", (141), (144); static theorem proved; general feedback budget conditional and open.

Feedback proof; static theorem proof reference, not reproduced. The static theorem uses an explicit five-ray description of the two-site cone and conditional induction over SCCs. Under the Doob kernel the full law factors over components by stationarity under extra arrivals; the tilt h_B^{-1} preserves that factorization. Component deletion constraints and decreasing no-portal survival payoffs supply the induction. The full five-ray inequalities, effective component kernel, zero-h cases, and noncircular induction are the proof reference in RN §28 at the cited headings and (141).

For the conditional assertion, replace $z_j\rho$ in the first-contact identity by the actual no-contact measure $z_j\pi_R^0$, with z_j now its total mass. Condition (23) bounds the terminal R support, and increasing f permits replacing j's final bit by one. The same identity proves the comparison. No product structure or intermediate cone preservation is assumed. □

A further proved countwise class is worth separating from static order. If every path from entrance y to the portal passes through i, each exact-return coefficient of addition at y preserves the i-deletion inequality. If every vertex has this property, it preserves the whole cone, for every batch. Thus a serial domain $1, \dots, m$ with jumps only to indices at most $i + 1$, portal exits only from m, and entry at 1 allows arbitrary backward jumps. This alone supplies no general static bound for a large cyclic SCC. *Proof reference, not reproduced:* the weighted intertwining and its commutator decomposition in RN (142)–(143); killing at a dominator removes every portal count from the killed-addition operator, making the commutator vanish.

Record: RN, §28, "Weighted intertwining and a dominator criterion", (142)–(143); proved countwise closure.

3.6 Complete-drain inventories and their positive families

Fix labelled all-active roots η and an extra label at x, a target v, and a causal labelled FIFO convention. Complete each particle's path to the original sink after stabilization by forcing residents to depart and then ignoring sleep. Give every nonsink path occurrence a coin, including ignored drain coins. These complete paths are independent ordinary walks, with

sitewise independent coins; the inverse reconstructs the used stack prefixes and independent unused tails.

Record: SX, §5, product-space lemma; proved representation input.

The two-family pool uses independent complete families for η and $\eta+\delta_x$. Fix their combined edge counts $J(z,w)$ and site coin counts k_z, with $m_z = \sum_w J(z,w)$. Every atom has weight

$$\prod_{z,w} K(z,w)^{J(z,w)} \prod_z p_z^{k_z} q_z^{m_z-k_z},$$

so the conditional pool is uniform. Let $c(J,k)$ be its sum of upper minus lower target indicators. The one-family pool uses paths for $\eta+\delta_x$ only, running the lower system on the common labels; its analogous signed count is $c_1(J,k)$. The respective root divergences are $2\eta+\delta_x$ and $\eta+\delta_x$. The general assertions $c, c_1 \geq 0$ remain open; $c_1 \geq 0$ would imply $c \geq 0$ by appending independent common drain paths and their ignored coins. Either universal sign would imply target monotonicity by summing the countable partition into almost surely finite drains.

Record: DI, §§1–2; normalization and implication proved, general coefficient signs open.

Lemma 3.12 (Finite forest enumeration and partial cancellation)**.** *For a feasible edge inventory with prescribed labelled roots, let $U = \{z : m_z > 0\}$ and $L(z,w) = m_z \mathbf{1}\{z=w\} - J(z,w)$ on U. The number of complete path decompositions, with equal destinations indistinguishable, is*

$$d(J) = \det L \, \frac{\prod_{z\in U}(m_z-1)!}{\prod_{z,w} J(z,w)!}. \tag{24}$$

The pool size is $d(J)\prod_z \binom{m_z}{k_z}$. Within each two-family pool, exchanging the two differing terminal target coins is a sign-reversing involution on the subset where both histories visit the target and both terminal arrival indices remain unchanged.

Record: DI, §4, (1)–(2), and §5; proved finite-word statements.

Proof. Arrange outgoing destinations into local words. Routing uses every word exactly when its last arrows form a forest to the sink: a last-arrow cycle would require increasing departure times around a cycle; an unused balanced remainder would conversely contain such a cycle. For a fixed forest T, count words by $\prod_z (m_z-1)!/\prod_{z,w}(J(z,w)-\mathbf{1}\{T(z)=w\})!$. Factoring gives (24) by the directed matrix-tree identity. Routing and reconstruction are inverse. The determinant identity is standard background, not an originality claim.

For the partial involution, the unchanged indices recover the exchanged coins from the image, making acceptance invariant and giving the inverse. Arrows and complete drain lengths are unchanged; one target zero and one target one exchange places, preserving J,k. Unaccepted histories are not claimed to cancel. □

Theorem 3.13 (Endpoint and nonreturning mixed-site inventories)**.** *For every finite η, extra x, target v, and feasible inventory, $c_1(J,k), c(J,k) \geq 0$ if $k_z \in \{0, m_z\}$ at every $z \neq v$; k_v is arbitrary. The conclusion also holds if every mixed off-target site, meaning $0 < k_z < m_z$, lies on no directed cycle involving another vertex in the support of J. Self-loops are permitted.*

Record: DI, §2 one-family extension; §6, (3)–(6); proved.

Proof. Let S be the off-target all-one sites, $T = S \cup \{v\}$, and $R = V \setminus T$. Endpoint parameters mean the limiting buffer law, with surplus still forced to jump. Route the common particles by IDLA to vacant sites of T or the sink, obtaining $C \subseteq T$, then route the extra, whose new endpoint is y. Remaining off-target residents sleep permanently. If $v \in C$, its occupation probability is $p/(1-qr_C)$, where $p = p_v$, $q = 1-p$, and r_C is the probability of a forced

departure returning before the sink or a vacant T-site. The extra can only enlarge the available return domain.

To preserve exact inventories, give edges variables X, and let g_z be the generating series of complete walks to the sink, $g_s = 1$. For small positive X, $K^X(z,w) = X(z,w)g_w/g_z$ is stochastic. Put $G_\eta = \prod_z g_z^{\eta_z}$. The one-family signed series equals $G_\eta g_x[q_v^{K^X}(\eta+\delta_x) - q_v^{K^X}(\eta)]$; the two-family series has the additional factor G_η. Telescoping over the IDLA prefix γ leaves positive terms

$$X^{E_\gamma+E_W} g_C g_y\, pq^{\ell(W)}$$

in the one-family case, with an extra G_η in the other. Here either $v \notin C$, $y = v$, and W is any closed walk from v in $R \cup C \cup \{v\}$, or $v \in C$, $y \in T \setminus C$ is off-target, and W is a closed walk in $R \cup C \cup \{y\}$ that visits y. The integer $\ell(W)$ counts departures from v.

The residual walk factor contains a root at v in every nonzero term. For total inventory J it therefore has at least one remaining departure there. Multiplication by $(p+q)^{m_v-\ell(W)-1} = 1$ gives nonnegative coefficients, explicitly residual path counts times $\binom{m_v-\ell(W)-1}{k_v-1}$. Coefficient extraction is finite for fixed J, and the homogeneous Bernstein basis is unique. If $m_v = 0$, both indicators vanish. This proves the endpoint assertion.

At a mixed site with no external return and no self-loop, all arrivals can be determined before processing it. Only its last coin affects the rest, so the fixed-inventory expression is affine in its sleep parameter. If its all-one and all-zero endpoint coefficients are A, B, its k-coefficient is $A\binom{m-1}{k-1} + B\binom{m-1}{k}$. Apply this successively to all such sites. Self-loops are inserted by the next lemma. □

Lemma 3.14 (Self-loop insertion)**.** *Nonnegative one- or two-family inventory coefficients remain nonnegative after arbitrary prescribed self-loop counts are inserted, even at sites with other return cycles.*

Record: DI, §7, (9); proved.

Proof. For a site with a nonself departures and k specified one-coins, inserting loop variable σ changes its monomial to

$$p^k q^{a-k}(1-\sigma)^{-k}(1-q\sigma)^{-a}.$$

Both series have nonnegative coefficients. A term using i loops from the first and j from the second has coin degree $a+j$; multiply by $(p+q)^i$ to reach the required $a+i+j$. The case of no departures is interpreted by the constant series. □

Theorem 3.15 (One off-target sleeping site)**.** *Both inventory signs hold for every finite input and every feasible J, k when all off-target sleep counts vanish except possibly at one site w. The count at w is arbitrary, and w may lie on return cycles. This hypothesis does not permit another off-target positive count, even a maximal count.*

Record: DI, §7, (7)–(10); proved.

Proof reference, not reproduced. The exact two-site occupation differences (7) become the positive walk series (8). Expanding their geometric denominators using $1 - p_0p_1 = q_0q_1 + p_0q_1 + q_0p_1$ supplies correctly counted coins. Surplus IDLA covers multiplicities; tracing all zero-coin sites to $\{v, w, s\}$ is a positive edge-series substitution, (10). Lemma 3.14 covers self-loops.

Theorem 3.16 (Unpadded directed-three-cycle inventories)**.** *On the directed three-cycle, with arbitrary sink exits and self-loops, $c_1(J,k), c(J,k) \geq 0$ for every feasible inventory, every coin-count vector, every finite all-active input, every extra site, and every target.*

Record: DI, §11, (16)–(18); exact integer certificate with an infinite-tail proof.

Certificate reduction; Bernstein arrays not reproduced. Write t for the product of the three forward edge variables. For binary common A and $x \notin A$, $N = |A| + 1$, the complete-block equations give

$$C_{A;x} = \frac{P_{A;x}}{(1-t)^N D_1 \cdots D_N}.$$

Here $D_1 = 1 - q_0q_1q_2t$, $D_3 = 1 - (1 - p_0p_1p_2)t$, and

$$D_2 = 1 - t\left(q_0q_1q_2 + \sum_i p_iq_{i+1}q_{i+2}\right) - t^2p_0p_1p_2q_0q_1q_2,$$

with cyclic indices. Fixing sink counts fixes forward counts up to a common number of circuits. In each sector write the residual numerator $\sum_{j=0}^{L} v_j(p)t^j$, $L \le 5$. After division by $(1-t)^N$ its ℓth coefficient is $G_\ell = \sum_{j\le\ell} \binom{\ell-j+N-1}{N-1} v_j$. For $u \ge 0$ the exact identity

$$G_{L+u} = \sum_{j=0}^{N-1} \binom{u}{j} \Delta^j G_L$$

reduces every tail coefficient to finitely many forward differences. The certificate verifies their nonnegative Bernstein coefficients at the fixed site degrees at L; positive degree elevation covers all u. Finite earlier layers are checked at their actual degrees. In the full-minus-target sector one must first absorb D_3^{-1} via $H_\ell = G_\ell + (1 - p_0p_1p_2)H_{\ell-1}$, $H_{-1} = 0$. Its early layers are certified, and the same tail argument then applies. Remaining reciprocal denominators expand positively, with exactly the required coins per circuit.

The omitted finite proof component is 413 Bernstein arrays across 66 comparison/exit sectors, with 36,765 positive coefficients and 29,256 zeros. The verifier in is `verify_cycle_inventory.py`. It reconstructs the arrays by integer arithmetic. Their canonical table SHA256 is

`efa3de81fda2e90791804102d0e474198398cb00610bcf2ce047e89be3a8d11a`

This is a certificate for unbounded circuit counts, not a finite inventory screen. The arrays and detailed complete-block derivation are a *proof reference, not reproduced*, DI (16)–(18).

For arbitrary multiplicities, run common surplus IDLA and the extra last. Each prefix leaves one of the binary comparisons just certified, or a nonnegative single-particle term. Its ignored coins supply positive degree elevation. Appending common drain walks gives the two-family result; self-loop insertion completes the proof. □

Proposition 3.17 (Permanent extra on general three-site kernels)**.** *Let $V = \{0,1,2\}$, with arbitrary sink-accessible K, common roots $\{0,1\}$, extra at 2, $p_2 = 1$, and $0 < p_0, p_1 < 1$. Both weak-target occupation increments are nonnegative. For target 0, a nonnegative inventory expansion is proved after adjoining ten unused complete drain walks, with extra root counts $(3,4,3)$; relabelling gives the other target. This is not an unpadded inventory theorem for general three-site kernels, nor a theorem for an interior site-2 coin count.*

Record: DI, §9, (13) and finite padding; computer-assisted theorem.

Proof reference, not reproduced. DI (13) derives the two occupations by tracing through the permanent site and stopping the lower system on its first visit there. After positive denominators are cleared, the numerator has 47,808 positive and 23,067 zero coefficients in jump-row degrees $(4,5,4)$ and sleep degrees $(2,4)$. The standard-library integer verifier, in , is `verify_three_site_permanent.py`. Its canonical coefficient-table SHA256 is

a6a5d7fbef290ee1350455e0e8a7bee98b162c3b9240c46882c30d0598eb4a8d

Multiplying by walk factors $g_0^4 g_1^5 g_2^4$ cancels all negative Doob powers, and positive closed-walk denominator expansions give the stated padding. DI §8's unpadded permanent-extra cycle family and §10's all-parameter cycle occupation theorem are included in the stronger Theorem 3.16; their separate calculations are not additional general-kernel claims.

Proposition 3.18 (Positive homogeneous count blocks on every cycle)**.** *For every directed cycle length $n \geq 2$, the one-vacancy and two-particle blocks of the unnormalized complete-block inverse have inverses with positive edge/coin expansions. Precisely, $F(A, B)$ is the signed sum over $S \subseteq A$ of forced-first-jump and subsequent surplus-IDLA histories ending at B, weighted by $(-1)^{|S|} \prod_{i \in S} q_i$, with $|B| = |A|$; the usual row factor $1/\prod_{i \in A} p_i$ is omitted. Let a_i now denote the forward edge weight. In the one-vacancy block, indexed by missing sites,*

$$F = I - T, \qquad T_{ki} = q_i \prod_{z \in [i,k)} a_z \prod_{z \in (i,k)} p_z \quad (i \neq k), \qquad T_{kk} = 0,$$

where arcs follow the cycle without a full revolution. In the two-particle block, ordered by pairs $i < j$, set $L_{\{i,j\},\{k,j\}} = \prod_{z=i}^{k-1} q_z a_z$ for $i \leq k < j$. Then $LF = I - T'$ with contributions $q_j a_j$ to $\{i, j+1\}$ and $p_j a_j \prod_{z=i}^{j-1} q_z a_z$ to $\{j, j+1\}$. A repeated site $\{i, i\}$ is replaced by $\{i, i+1\}$ with an extra factor a_i; coincident contributions are added.

Record: DI, §12, (19)–(22); proved for all lengths, not just checked lengths.

Proof. In the one-vacancy row, take the first selected forced-jump site after the vacancy. All count-preserving motion traverses its forward arc to that vacancy. Summing remaining selected sites on the arc gives T. For the pair row use the homogeneous IDLA relation $u_i^2 = a_i u_i u_{i+1}$. The L-weighted sum of $(u_k - q_k a_k u_{k+1})(u_j - q_j a_j u_{j+1})$ telescopes to

$$[u_i - (\prod_{z=i}^{j-1} q_z a_z) u_j](u_j - q_j a_j u_{j+1}),$$

which reduces to the displayed T'. At physical weights at most one, the row sums are bounded respectively by $1 - \prod_{i \neq k} p_i < 1$ and $q_j + p_j \prod_{z=i}^{j-1} q_z < 1$. Hence $F^{-1} = \sum_m T^m$ or $F^{-1} = \sum_m (T')^m L$, with positive coin factors within the available departures. These are homogeneous blocks only: smaller-count forcing remains signed, so they do not prove longer-cycle inventory or arbitrary-input occupation order. □

3.7 Decisive obstructions and the remaining boundary

The following are exact failures of stronger proposals, not counterexamples to the all-active single-target conjecture.

Remark 3.19 (Merging defeats return-count cone closure). Let $R = \{0, 1, 2, 3\}$, with 0 jumping equally to $1, 2$, both jumping surely to 3, and 3 jumping equally to portal and sink. Take all $p_i = 1/2$. The no-portal reference is product Bernoulli(2/3). After one active input at its unique source 0, the exact-three-return submeasure ν_3 satisfies

$$\nu_3(\{0, 3\}) = 0, \qquad \nu_3(\{0, 1\}) = \nu_3(\{0, 2\}) = 1/648, \qquad (\nu_3 L_1^0)(\{0, 3\}) = -1/1296.$$

Three exits and two remaining sleepers require all four initial sleepers. If 0 remains occupied, it emits only one particle and visits only one branch, so $\{0, 3\}$ is impossible. The event ending at $\{0, 1\}$ has weight $(2/3)^4(1/2)^4(1/8) = 1/648$; deletion from 1 reaches sleeping 3 without a return with probability $1/2$. This proves the negative residual. Summing all return counts still leaves residual $-1/324$. Disjoint subtrees are therefore essential to the forest proof.

Record: RN, §28, "Exact limits of return-count closure"; exact counterexample.

Later entry can fail even on the path $a \to b \to \{j, s\}$: with all sleep probabilities $1/2$ and $j \to b$ surely, complete feedback from the actual product Bernoulli$(2/3)$ reference leaves a deletion residual $-2/105$. Source-only entry does not universally repair complete-feedback cone preservation either. Prepend to the diamond a vertex d with $d \to 0$ of probability ε and otherwise $d \to s$, and let $j \to d$ surely. With all sleep probabilities $1/2$, the specified site-1 residual at exact support $\{0, 3\}$ is $-\varepsilon/1296 + O(\varepsilon^2)$; at $\varepsilon = 1/4860$ its certified exact value is less than $-1/7000000$. Thus neither intermediate nor complete-feedback cone closure is available for arbitrary merges. *Proof reference, not reproduced:* RN §28, "Complete feedback as the remaining cone question", terminal path example and source-entry counterexample; `verify_feedback_cone.py`. The weaker budget (23) is not refuted by these examples.

Finally, exact response gluing does not allow commuting separately conditioned count coefficients: the one-site example RN (B7) has $A_0A_1 - A_1A_0 = \mathrm{diag}(1/8, -1/8)$. For inventory cancellation, fixing common complete paths or a last-exit forest is too restrictive: SX §5 has a negative fixed-path increment, and DI §4 has a negative fixed-forest coefficient although the entire pooled inventory is positive. These examples explain why the joint support/count state and the full pooling over admissible words must be retained.

Record: RN, §31, (B6)–(B7); SX, §5; DI, §4; exact scoped obstructions.

Part II
Proof techniques for the open conjecture

4 Exact exchanges of finite records and complete avalanches

These are local injections on specified families of pairs of histories. They preserve the full product measure, including unused records, and have explicit inverses. Their domains do not cover every losing pair, and their images cannot be combined without an overlap argument. The most general result here is the complete-avalanche repair in Theorem 4.10; the disconnected-component involution gives a separate exact cancellation.

4.1 First-contact cancellation and adapted activation

For a finite common input η disjoint from marked sites S, let

$$\mathcal{D}(\eta, S) = \sum_{T \subseteq S} (-1)^{|S|-|T|} \operatorname{Law}\big(\operatorname{Stab}(\eta + \text{active } T + \text{sleeping } (S \setminus T))\big).$$

This is a signed measure on complete terminal supports. Let $\nu_\eta^S(B)$ be the subprobability of the common particles stabilizing at sleeping B without visiting S.

Proposition 4.1 (First-contact cancellation)**.** *With the common starting state fixed,*

$$\mathcal{D}(\eta, S) = \sum_B \nu_\eta^S(B)\, \mathcal{D}(\textit{sleeping } B, S).$$

The cancellation of all histories hitting S has a sign-reversing inverse on complete histories.

Proof. Hold all marked particles and explore the common particles until their first visit to S or stabilization. This prefix is independent of the marks' active/sleeping choices. At a first hit $z \in S$, toggle only that initial choice. Immediately after the arrival both versions have two active particles at z; their other states agree. Continue with the same unused stacks. The prefix and canonical first hit are unchanged, so the toggle is its own inverse, with equal probability and opposite sign. Only the no-hit histories remain, with their original subprobability weights $\nu_\eta^S(B)$. Their unused tails give the displayed continuation. □

Record: VA §2. The remaining contrast is signed, and its generating law cannot be replaced by an arbitrary sleeping background.

Fix final odds r^*, and for processed sites S put $r_i^S = r_i^*$ on S and zero elsewhere. Let $T(r, s) = P_r^{-1} P_s$ be the stochastic waking transport of Proposition 2.7.

Theorem 4.2 (Adapted activation with correct conditional blocks)**.** *Start at sleeping binary support A with $S = \varnothing$. Choose a nonempty batch $B \subseteq V \setminus S$ using the revealed past, wake every* current *sleeper in B, and stabilize at odds $r^{S \cup B}$. Set $S \leftarrow S \cup B$ and continue. After at most $|V|$ batches the full terminal law is $P_{r^*}(A, \cdot)$. Choices may use the histories of two coupled copies, provided each marginal block conditioned on that joint past and the chosen action has its prescribed transport law.*

Proof. The block is $W_{S,B} = T(r^S, r^{S \cup B})$. For any terminal payoff h, put $H(S, C) = (T(r^S, r^*)h)(C)$. The transport cocycle gives, pointwise for every permitted action,

$$W_{S,B} H(S \cup B, \cdot) = H(S, \cdot).$$

Hence H of the current state is a bounded martingale under the adapted choices. The number of steps is bounded by $|V|$; at the endpoint $H = h$, while initially $T(0, r^*) = P_{r^*}$. This proves every terminal test-function identity. Surplus IDLA extends the sampler to multiplicities with fresh suffixes. Choosing a batch after examining favorable candidate outcomes does not have the stipulated conditional law. □

Record: VA §§3–4. This proves marginals and termination; it does not supply the missing successful coupling of every pair of blocks.

Proposition 4.3 (Discovered blocks retain their selection mass)**.** *At partial state (S, C) choose an unprocessed seed x before exploring. Run its avalanche at final parameters, recruiting each newly visited unprocessed site into B, including initially vacant sites, and releasing any dormant marked resident encountered. Let $L^B_{S,x}(C, D)$ be the joint law of discovered B and terminal support D. Then*

$$\sum_{B,D} L^B_{S,x}(C, D) = 1, \qquad \sum_B L^B_{S,x} T(r^{S\cup B}, r^*) = T(r^S, r^*).$$

For fixed B, $L^B_{S,x}$ is a subkernel of $W_{S,B}$, selected by the event that the seed exploration visits every required vertex of B and no unprocessed vertex outside it. Conditioning on B does not generally give $W_{S,B}$.

Proof. In the mixed initial state defining the right side, hold all unexamined marked particles active but untoppled. An arrival makes a held active particle and a sleeper identical, so the recruited exploration is legal. Its untouched marks are exactly those outside $S \cup B$. Release them to complete the process. Abelianity and fresh suffixes give the second identity and finite stabilization the first. For fixed B, run the same seed schedule while holding the other initially activated B-particles. On the specified selection event all held particles have been released and the discovered history agrees. Thus its weight is the fixed-block probability multiplied by the endpoint's conditional selection probability, not divided by an assumed independent normalization. □

Record: FC §§1–2. A two-site exact example in FC §3 disproves the replacement of this joint law by the unconditioned fixed-block kernel.

Proposition 4.4 (Creation-history factorization of a forced-jump defect)**.** *Fix a target-free binary set A and $z \in A$. Let $a(A)$ be the probability that all-active A never visits v. Let $\mathcal{J}_z$ force its z-particle to jump and use surplus IDLA to restore a binary active input, discarding target-hitting histories but retaining sink losses. With $H_v(B, z)$ the target-hitting probability of active z among sleeping B,*

$$a(A) - (\mathcal{J}_z a)(A) = \frac{p_z}{1 - p_z} \sum_B \nu^{\{v,z\}}_{A\setminus\{z\}}(B) H_v(B, z) \geq 0.$$

Proof. Let a^{sl} be avoidance with the marked particle initially asleep at z and all other particles active. First-instruction conditioning gives $a(A) = p_z a^{\mathrm{sl}} + (1 - p_z)(\mathcal{J}_z a)(A)$. Hold the mark and run the common particles until first hitting v or z, or stabilizing. A hit on v makes both avoidance indicators zero. A hit on z makes both states identical, and their continuations cancel. On stabilization avoiding both, the sleeping version is finished, whereas the active version loses avoidance with probability $H_v(B, z)$. Therefore $a^{\mathrm{sl}} - a(A) = \sum_B \nu(B) H_v(B, z)$; rearranging proves the claim. The subprobability weights are not renormalized. □

Record: DF §1; the forced-jump avoidance inequality is also in AF §1, using the standard acceptable-toppling least-action framework. Positivity of these defects does not imply positivity of their signed coefficients in a response expansion.

For the associated routing formulation, AF/CR prove that a feasible nonnegative forced-jump flow terminates: the potential $\sum_{z\in A}\mathbb{E}_z T_\partial$ drops in conditional expectation under each routed step, and surplus IDLA only lowers it further. This is a termination result conditional on feasible positive routing; existence of the required general flow is still open. The exact finite repairs in DF, CE, and WT are retained as scoped certificates, not as a general flow-existence theorem.

4.2 Records, stopped words, and connected histories

Use the buffer representation with independent records

$$R_z(j) = (B_z(j), D_z(j)), \qquad \mathbb{P}(R_z(j) = (b,d)) = p_z^b(1-p_z)^{1-b}K(z,d).$$

An arrival releases an old resident's waiting arrow, if present, and uses its next record: it stays if the bit is one and otherwise emits its own arrow. Pending deliveries may be held. A potential arrow remains part of a sleep record even while unused. This representation has the ordinary ARW stabilization law; it does not authorize a successful raw sleep at an occupied site.

Lemma 4.5 (Recoverable stopped-word exchange)**.** *Start two explorers at the same vertex in independent fresh record fields. Each follows its own arrows to its first sleep or sink exit, holding any released resident packets. At each site write its read word and remaining tail as $W_z^- T_z^-$ and $W_z^+ T_z^+$. The transformation*

$$(W_z^- T_z^-, W_z^+ T_z^+) \longmapsto (W_z^+ T_z^-, W_z^- T_z^+)$$

is a measure-preserving involution. It permits returns and unequal word lengths, and holds conditionally after fixed legal stopping prefixes.

Proof. The image explorers reproduce the opposite original walks, so they recover the cuts. Every complete record stays at its original site, hence the product of the two finite cylinder weights is unchanged. Each tail remains in its own copy. The stopping rule leaves independent fresh tails, and the finite stopped paths partition the sample space up to a null set. Summing their equal weights proves the assertion. Finiteness follows from the positive minimum sleep probability, or from sink accessibility. □

Record: SX §§1–2; CX §1. Proved local involution.

Proposition 4.6 (Connected-history reduction)**.** *For independent completed histories from η and $\eta + \delta_x$, form the undirected graph on V containing every emitted internal edge in either history. Sink edges and unspent resident arrows are excluded. Let Q be the component of the target v, and let I_-, I_+ be the terminal target indicators. Whole-book exchange at all sites of Q is a measure-preserving involution on $\{x \notin Q\}$, and*

$$q_v(\eta + \delta_x) - q_v(\eta) = \mathbb{E}\big[(I_+ - I_-)\mathbf{1}_{\{x\in Q\}}\big].$$

Proof. On $x \notin Q$, the two inputs restricted to Q agree. No emitted edge crosses its boundary in either history. After exchanging books, run each copy's interior input by the opposite original interior history and retain its original exterior history. These are legal completions. Thus the interior outcomes exchange, the exterior outcomes stay fixed, and the union of emitted edges, hence Q, is unchanged. For fixed Q the swap preserves product measure; the finite partition by its invariant selected value proves this for the adaptive map. Its contribution to the expected target difference cancels. □

Record: SX §6; general proof with exact finite-cylinder audits.

Proposition 4.7 (Joint success conditioning)**.** *Give all particles in both copies independent extended K-walks continued to the sink. Let E be the event that at least one of these walks visits v, and $h(z) = \mathbb{P}_z(T_v < T_\partial)$. Before any visit has been revealed, a causal colour-blind exchange of the unused paths and their success marks between two labels at the same site preserves the conditional law given E. Moreover*

$$q_v(\eta + \delta_x) - q_v(\eta) = \mathbb{P}(E)\,\mathbb{E}[I_+ - I_- \mid E]$$

when $\mathbb{P}(E) > 0$; otherwise both occupation probabilities vanish.

Proof. At current label positions (z_j) the conditional weight of a nonempty success set J is

$$\frac{\prod_{j\in J} h(z_j) \prod_{j\notin J}(1 - h(z_j))}{1 - \prod_j (1 - h(z_j))}.$$

It is invariant under exchanging two labels at the same position, as is the law of their conditioned unused walks. Sleep factors and colour-blind scheduler factors cancel in the conditional likelihood. Exchange both the path and its mark. After a revealed hit, E is certified and the remaining paths have their ordinary independent laws; the nonempty-mask restriction is no longer applicable. Finally, either occupied-target event implies E, which proves the displayed identity. □

Record: SX §4 and ST §5. Separate success conditioning for each copy does not permit this exchange.

4.3 A complete-avalanche repair with adaptive cuts

Lemma 4.8 (Released residents and suppressed final jumps)**.** *Fix resident waiting arrows, future books, and pending packets. Removing selected residents and holding their emitted arrows cannot reduce the visited set after all held packets are finally delivered. The comparison also permits additional residents with specified waiting arrows. Consequently, if a completed history ends vacant at v, changing its last target jump record $(0, d)$ into a sleep $(1, e)$ leaves a sleeper at v, preserves that last target index, and the suppressed packet d, subsequently delivered into the new terminal state, has an avalanche avoiding v.*

Proof. Follow the original deliveries. At the first visit to a removed resident, release its held arrow when the original resident would have emitted it; the two local buffers then agree. At a visit to an additional resident, hold its extra emitted arrow until the comparison execution ends. Finally release every remaining packet. This is a legal schedule and contains the original visited set.

For the second claim, freeze v after its original last jump. Changing that record suppresses one packet. With the same outside books, removing this packet reduces used output prefixes and the visited set, so the new process cannot return to v. Abelianity permits delivery of the suppressed packet last. The combined outside execution is the original outside execution, which avoided v. The same avoidance holds if the auxiliary avalanche starts among any subset of the new outside residents, keeping their actual waiting arrows. □

Record: AX §§1–2; SX terminal-record lemma. These are pathwise visited-set statements, not occupation attractiveness.

Lemma 4.9 (Full avalanche block exchange)**.** *Two complete avalanches starting from the same sleeping support B and the same active input can exchange all their read words and all initial resident arrows at B. Retaining the unread tails in their copies gives a measure-preserving involution, including loops, wakes, and losses.*

Proof. Each transplanted avalanche has exactly the opposite original initial arrows and read records. It therefore reproduces that complete avalanche and recovers every cut. Product weights agree site by site; the exchanged resident arrows have the same K-row laws. Full stabilization is a finite stopping line. Apply the construction again to recover all original data. □

Record: AX §3.

Theorem 4.10 (Matched-background terminal repair)**.** *Take independent all-active lower and upper histories from η and $\eta + \delta_x$ with target outcomes $(1, 0)$. Write the lower terminal support as $B \cup \{v\}$, and assume the upper history visited v. Let n, m be their last target arrival indices, with records $(1, e), (0, d)$. Exchange those two records and let σ be the resulting upper terminal support. On the domain $B \subseteq \sigma$, there is an equal-measure injection to target outcomes $(0, 1)$ preserving both indices n, m. It requires no external record of these indices or of B.*

Proof. The upper image contains v and keeps index m by Lemma 4.8. In the lower history hold the newly emitted packet d while replaying its old continuation; this leaves sleeping support B and that packet. In the unused upper books after σ, run an auxiliary avalanche from d among only B, using its resident arrows in σ. It avoids v by Lemma 4.8. Exchange this avalanche's full read words and initial B-arrows with those of the lower avalanche, using Lemma 4.9. The actual lower image follows the copied avoiding avalanche, ends vacant at v, and has no target arrival after n. The upper completed history stays unchanged, since only unused words and terminal resident arrows were modified.

For the inverse, run both image systems to read n, m. Temporarily replace lower record n by upper record m and stabilize the lower input. This counterfactual sleep recovers the original lower creation history, its support $B \cup \{v\}$, every suffix offset, and each terminal resident-arrow address. The possibly exchanged arrows at B were unused in this counterfactual history, so they cannot change it. The actual upper history recovers σ. With these cuts fixed, rerun and exchange the two auxiliary avalanche blocks and arrows, then exchange the target records back. Retain the decoded n during inversion: recomputing it after restoring the old avalanche would be incorrect. This restores every original record and tail.

Each finite cylinder has equal product weight after the complete transformation. The explicit inverse makes its images disjoint. The first terminal transposition alone need not be injective; the complete composite is. □

Record: AX §4. General local proof, supported by exact inverses, weights, and terminal-index checks in `verify_stack_exploration.py`.

Corollary 4.11 (Unmatched backgrounds with a specified avoidance test)**.** *There are two extensions of Theorem 4.10.*

1. *Put $C = B \cap \sigma$ and $D = B \setminus \sigma$. Use the upper auxiliary C-avalanche and require it to avoid D. Exchange complete C-avalanches and C-resident arrows; the untouched lower sleepers in D are never contacted.*

2. *At each $b \in D$, reserve the first unused upper sleep record after σ. Use its potential arrow as the missing auxiliary resident arrow, and start that site's auxiliary book after its reserved prefix. Explore the upper avalanche on all of B, and require it to avoid v. Exchange full B-avalanches and resident arrows.*

Each rule separately gives an equal-measure injection on its specified domain. They agree with the matched rule when $D = \varnothing$; their images are not asserted disjoint in general.

Proof. For the first rule the copied avalanche ignores D even in the actual lower process, so the terminal-index and inverse arguments are unchanged. Recover B, σ, C, D by the same counterfactual execution. For the second, the reserved prefixes keep every sleep bit unchanged; only their unused terminal arrows may exchange. The first-unused-sleep rule therefore recovers every reservation. All exchanged resident arrows belong to sleep records at the same site, preserving complete-record inventories and weights. The imposed avoidance preserves the lower index n, and the preceding inverse applies. Selecting successive reservations until success would be a different, unjustified sampling rule. □

Record: AX §§5–6. A forced-return example in AX §7 excludes a universal reservation-only repair.

4.4 Creator switches retained in the collection

The narrower creator constructions remain useful because they change the history that generated a missing sleeper, whereas the preceding avalanche map holds that background fixed. Their precise domains are retained below; the full decoders are in CX.

Proposition 4.12 (Creator and endpoint switches: recorded local families)**.** *For common roots $\{v, a\}$, extra root y, and target v, the following are proved local equal-measure injections.*

1. *Both first creator words sleep at v and the upper extra sleeps immediately at y. At the upper root's first target contact, exchange its subsequent stopped word W with the lower creator L. Require L, W to avoid y, L to sleep at v, W to end elsewhere or at the sink, and the original upper residual avalanche to avoid v. If the lower second word avoids both v and the endpoint of W, this alone repairs the pair.*

2. *Under the same conditions, a sink-ending lower second word may meet that endpoint. Composing with the two-walk loop-erased endpoint switch repairs the resulting lower wake.*

3. *A fixed reversal of lower root order, combined when needed with a recoverable closed-word insertion/removal, covers the recorded two-sleep-endpoint family: the roots first reach their distinct final endpoints at their terminal arrivals, and the second visits the first endpoint.*

4. *With an immediate lower target sleep, cutting at the upper last target arrival permits earlier upper returns. A separate marked-cut version permits a returning lower creator when a preceding upper sleep recovers the cut. In the latter comparison, lower contacts with old sleepers must be contained in upper contacts.*

Proof reference, not reproduced in full. The first switch is Lemma 4.5 below an unchanged upper first-contact prefix. Its upper-image proof uses the delayed-resident comparison of Lemma 4.8. For the second, attach an absorbing endpoint (sleep, z, d) to each sleep record so the unused arrow is retained, and apply Fomin's loop-erased two-walk switch. Lift it to books by replacing the concatenated local words and preserving their tails. The inverse applies the endpoint switch first, then the creator switch. The cycle decoder in CX §6 recovers a first-return closed word in unused books; CX §7 supplies the last-contact and marked-cut decoders. These are separate local maps, with no general image accounting.

Record: CX §§2–7; SX §2. Fomin's loop-erased switch is an external theorem; its ARW book lift and stated repair domains are the AI-developed applications.

5 Backward restoration, positive entrance measures, and root retention

Backward restoration gives a pointwise minorization of creation laws. Canonical restoration on complete books then represents the associated entrance measures by positive finite fibres. Comparing those fibres proves joint retention of the original roots, with an exact strictness criterion. These are different constructions: posterior resampling is not the same operation as repeatedly editing one fixed book.

For navigation, Theorem 5.1 gives the normalized backward chain and minorization; Theorem 5.2 gives canonical book restoration; Proposition 5.3 gives the positive entrance formula; Proposition 5.4 gives inverse diagonals and the full row; and Theorem 5.5 gives joint retention and strictness. The short N1 coupling is Theorem 5.6.

5.1 Setup and the imported deletion identity

Let V be finite, with absorbing sink s, jump kernel K, and a positive-probability route to s from every vertex. Assume $0 < p_z < 1$ and put $r_z = (1-p_z)/p_z$, $r_T = \prod_{z\in T} r_z$. Self-loops and zero transition probabilities are allowed. For a finite all-active input η, let μ_η be its law on terminal sleeping supports. Write $\mu_A = \mu_{\mathbf{1}_A}$ for binary input, $P(A,B) = \mu_A(B)$, and $E = P^{-1}$. Invertibility is the complete-block theorem. Let A_x be active addition followed by stabilization, $A_s = I$, and $J_x = \sum_{y\in V\cup\{s\}} K(x,y)A_y$. These commuting operators satisfy $A_x^2 = A_x J_x$.

We use the deletion identity from the preceding structural work without reproducing its proof. For $z \in B$, $W_z(B,C)$ is the subprobability of ending at C after removing the z-resident, forcing its first K-jump, and stabilizing, killed on any visit to z. Its columns avoid z. Define

$$L_z(B,C) = \begin{cases} \mathbf{1}_{\{B=C\}}, & z \notin B, \\ -r_z W_z(B,C), & z \in B. \end{cases} \qquad \mathcal{C} = \{\mu \geq 0 : \mu L_z \geq 0 \text{ for all } z\}.$$

Deletion says $\mu_\eta L_z = N_\eta^z$, the no-z-visit terminal sublaw. Thus it is zero if $\eta_z \geq 1$. Measures in $\mathcal{C}$ need not be normalized or mixtures of creation laws.

Record: RN, Sections 6–8, (7), (9), (11)–(14); deletion identity imported, proof not repeated; CB, Sections 2–3, invertibility.

The full-input law $\pi = \mu_V$ is stationary for every A_x. Indeed, drain the surplus in $V + \delta_x$ by ordinary IDLA before stabilizing the retained binary input with fresh suffixes. Also $\pi(B) > 0$ for every B: hold a particle at each desired B-site, route the others to the sink by finitely many jumps, then sleep the retained particles. In particular $\pi L_z = 0$. Abelian stabilization and surplus IDLA are background inputs; this argument is not an originality claim for their stationary interpretation.

5.2 Normalized backward restoration and cone minorization

Theorem 5.1 (Backward roots and pointwise minorization)**.** *For $z \notin C$, set*

$$R_z(C,B) = r_z \frac{\pi(B)}{\pi(C)} W_z(B,C), \qquad z \in B. \tag{25}$$

These are probability kernels. Any adapted rule choosing a missing vertex at each step reaches V in finite mean time. If only missing members of a fixed root set R are selected, the entire exit law on supports containing R is independent of the rule.

For binary A, stop at A or upon exceeding $|A|$ particles; let $H_A(C)$ be the probability of stopping at A, set to zero for $|C| > |A|$. Then

$$\mu_A(C) = \pi(C) \frac{\mu_A(A)}{\pi(A)} H_A(C). \tag{26}$$

Consequently, for every $\mu \in \mathcal{C}$ and $A \subseteq V$,

$$\mu(C) \geq \frac{\mu(A)}{\mu_A(A)}\mu_A(C) \quad \textit{for every sleeping support } C. \tag{27}$$

Record: RN, Section 8, (11)–(14); proved general argument, finite checks supplementary.

Proof. The equation $\pi L_z = 0$ gives row normalization in (25). Conservation gives $|B| \geq |C|$ on a positive transition. Moreover $R_z(C, C \cup \{z\}) > 0$: force the removed particle along a simple sink path without returning to z, leaving every encountered C-resident asleep again. There are finitely many pairs (C, z). Their minimum positive one-particle-increase probability bounds below the conditional chance of increasing count at every step. The finite-mean termination assertions follow, including the count cap.

For $h_\mu(C) = \mu(C)/\pi(C)$, direct substitution gives

$$(\mu L_z)(C) = \pi(C)\left[h_\mu(C) - \sum_B R_z(C, B)h_\mu(B)\right]. \tag{28}$$

Thus the cone consists exactly of nonnegative densities excessive for every permitted backward insertion.

To prove independence of the whole exit law, write a signed measure as $\mu = \alpha P$. The matrix PL_z has zero rows on inputs containing z. Its remaining square block has full rank: L_z has identity rows on supports avoiding z, and multiplication by P preserves rank. Hence $\mu L_z = 0$ exactly when α is supported on inputs containing z. For all $z \in R$, the common harmonic space consequently has basis μ_D/π, $D \supseteq R$, and dimension $2^{|V|-|R|}$. For any fixed selection policy, finite-mean absorption gives a harmonic space of that same dimension, one coordinate per boundary support. The common space is therefore the whole space. Every boundary function has an extension harmonic for every permitted choice; bounded optional stopping proves policy independence, including history-dependent choices.

For the capped A-root process, μ_A/π is harmonic at every missing root and vanishes above count $|A|$. Optional stopping gives (26). For a cone measure, its density is excessive; at stopping retain the contribution from A and discard the other nonnegative boundary values. Thus $h_\mu(C) \geq H_A(C)h_\mu(A)$, which proves (27). When $|C| > |A|$ the right side is zero. □

In particular a probability creation law satisfies $\mu_\eta(A) \leq \mu_A(A)$, without a nesting assumption on its input. The residual in (27) is nonnegative, but need not belong to $\mathcal{C}$; repeated peeling is not justified. Neither the kernels R_z nor this minorization assert set inclusion.

5.3 Canonical restoration on complete books

A book has independent sleep bits $b_z(i)$ and potential arrows $d_z(i)$. On arrival, an old resident emits its reserved arrow; the new particle stays if its bit is one and otherwise emits its own arrow. Pending emitted packets may wait for delivery. If m_z arrows have been emitted and $I_z(m)$ counts arrows entering z among those prefixes, then for an all-active input

$$n_z = \eta_z + I_z(m), \qquad m_z = n_z - b_z(n_z), \qquad b_z(0) = 0. \tag{29}$$

The local output length $n - b_z(n)$ is nondecreasing. Least action therefore makes emissions increase with input and decrease when bits are changed from zero to one.

Theorem 5.2 (Canonical terminal-root restoration)**.** *Fix $\eta \geq \mathbf{1}_R$ and a stabilizable book. Repeatedly select a missing root, change its last arrival bit from zero to one, and restabilize the original input. Once all roots are occupied, reset changes beyond their final terminal indices.*

The resulting used prefixes and reset book are independent of the selection order. Their emission vector is the greatest among root-occupied outcomes obtainable by forcing only terminal root sleeps. For fixed R, it increases with η and decreases with the original sleep bits, with arrows held fixed.

Proof. Every edit decreases the emission vector, strictly at the selected root, so there are at most $\sum_z m_z$ edits. Decreasing arrivals puts every edit at or beyond its eventual terminal index; resetting unused later edits changes no used history.

Compare against any root-occupied candidate from a smaller input η' and a book with more original one bits, with only its terminal root sleeps additionally forced. Let its emissions be m' and its terminal root indices $t_z = m'_z + 1$. Inductively, all restoration edits lie at indices at least t_z. They cannot change the candidate's used history: at t_z it already sleeps, and later edits are unused. Local monotonicity gives current emissions $m \geq m'$. The next missing root z is edited at

$$n_z = \eta_z + I_z(m) \geq \eta'_z + I_z(m') = t_z.$$

This preserves the induction. With the same input and original book, each terminal outcome dominates every candidate, giving uniqueness and maximality. The same comparison proves both monotonicities. Discarded edits before resetting need not be order-independent. □

Record: RN, Section 11, (17) and Canonical restoration theorem; proved on complete books.

Call the map $F_{\eta,R}$. For an image book ω occupying R, its preimages are precisely the admissible ω^T, $T \subseteq R$, obtained by flipping selected terminal root bits from one to zero and retaining the flip only if $F_{\eta,R}(\omega^T) = \omega$. Denote this family by $\mathcal{I}_{\eta,R}(\omega)$. It is downward closed: more flips can only increase canonical emissions, so if a larger flip returns to the original vector, every smaller flip does too.

Partitioning by the finite terminal indices and the flipped subset makes every inverse a fixed finite coordinate change. The pushforward density on root-occupied books is therefore

$$Z^{\mathrm{book}}_{\eta,R}(\omega) = \sum_{T \in \mathcal{I}_{\eta,R}(\omega)} r_T, \qquad 1 \leq Z^{\mathrm{book}}_{\eta,R} \leq \prod_{z \in R} p_z^{-1}. \tag{30}$$

Consequently $\mu_\eta(R \text{ occupied}) \geq \prod_{z \in R} p_z$. Finite stabilization holds almost surely for the finitely modified books; the coordinate partition handles the random edit indices.

Record: RN, Section 11, (18)–(19); proved finite-fibre change of variables.

5.4 Positive entrance measures and their fresh-book formula

Proposition 5.3 (Entrance representation)**.** *Let $H_R(C, B)$ be the uncapped backward exit law from Theorem 5.1. For $B \supseteq R$, define*

$$\Gamma_R(B, C) = \frac{\pi(C)}{\pi(B)} H_R(C, B), \quad m_R(B) = \sum_C \Gamma_R(B, C), \quad \nu_{R,B} = \frac{\Gamma_R(B, \cdot)}{m_R(B)}. \tag{31}$$

These rows are nonnegative, supported on $|C| \leq |B|$, and are the identity on boundary columns $C \supseteq R$. Every finite all-active $\eta \geq \mathbf{1}_R$ has the positive expansion

$$\mu_\eta(C) = \sum_{B \supseteq R} \mu_\eta(B) \Gamma_R(B, C), \qquad \sum_{B \supseteq R} \mu_\eta(B) m_R(B) = 1. \tag{32}$$

In particular $\nu_{R,R} = \mu_R$ and $m_R(R) = 1/\mu_R(R)$.

Proof. Count monotonicity gives the support assertion, and immediate stopping gives the boundary columns. The density μ_η/π is common harmonic for the missing roots. Its boundary extension is precisely (32); summation gives the normalized weights. For input R, only boundary support R is possible. □

Record: RN, Section 12, (20)–(22); proved.

There is also a constructive formula. Conditional on an original forward history ending at $B \supseteq R$, remove its emitted prefixes. At each occupied site the first remaining bit is one and its resident arrow is unread; at vacant sites the entire tail is unread. All resident arrows and remaining bits/arrows have their original independent laws. Denote this fresh book by U_B. With input B, its unflipped stabilization has zero emissions. Let $\mathcal{I}_R(B, U_B)$ be its admissible first-root flips and C_T the raw endpoint after flip T. Then

$$\Gamma_R(B, C) = \mathbb{E} \sum_{T \in \mathcal{I}_R(B,U_B)} r_T \mathbf{1}_{\{C_T = C\}}, \qquad m_R(B) = \mathbb{E} \sum_{T \in \mathcal{I}_R(B,U_B)} r_T. \tag{33}$$

To see this, replay the old prefix while holding newly emitted terminal-root packets. The remaining process is exactly the shifted one. Canonical maximality identifies root-occupied candidates above the old emissions with shifted candidates, so admissibility depends only on B, U_B. Disintegrate (30) over forward histories. The resulting rows equal Γ_R: on binary inputs containing R, the boundary matrix $Q_R(D, B) = P(D, B)$ is invertible by the common-harmonic-space argument. With P_R denoting the rows of P indexed by $D \supseteq R$, both constructions satisfy $P_R = Q_R \Gamma_R$. Freshness is conditional on the original forward history, not on a canonical image, whose law is biased by its fibre density.

Record: RN, Section 16, (47) and its fresh-tail replay proof; proved, separate AI audit.

5.5 Inverse diagonals and the full inverse row

Proposition 5.4 (Root-block cancellation)**.** *Put $M_C = I - K_{CC}$, $d_C = \det M_C$, with $d_\varnothing = 1$. Then $d_C > 0$, and*

$$E(R, R) = Z_R = d_R \sum_{C \subseteq R} \frac{r_{R \setminus C}}{d_C}. \tag{34}$$

Define $\kappa_z^C = K(z, s) + K(z, C) M_C^{-1} K(C, s)$, with the second term zero for $C = \varnothing$. Here $K(z, C)$ and $K(C, y)$ denote the restricted row and column. The full row is

$$E(V, D) = (-1)^{|V \setminus D|} d_D \sum_{C \subseteq D} \frac{r_{V \setminus C}}{d_C} \prod_{z \in V \setminus D} \kappa_z^C, \qquad \pi(V) = Z_V^{-1}. \tag{35}$$

The alternating signs are weak: entries may vanish. For proper R, $\mu_R(R) = 1/Z_R$ is not asserted.

Proof. Set $F_z = A_z - J_z$, $a_C = \prod_{z \in C} A_z$, and $B_z = A_z + r_z F_z$. In the commutative addition algebra $A_z F_z = 0$. For $T \subseteq V$, let $J_z^{\mathrm{out}} = \sum_{y \notin T} K(z, y) A_y$, including s. Row multilinearity gives

$$\prod_{z \in T} F_z = \det\Big(M_T \operatorname{diag}(A_T) - \operatorname{diag}(J^{\mathrm{out}})\Big) = \sum_{U \subseteq T} (-1)^{|T \setminus U|} d_U a_U \prod_{z \in T \setminus U} J_z^{\mathrm{out}}. \tag{36}$$

Indeed row z is $F_z e_z + \sum_{y \neq z} K(z, y) A_y (e_z - e_y)$, where e_z is a coordinate row. A selected directed cycle has dependent rows. A nonempty acyclic selection ends at some $F_u e_u$, producing the zero factor $A_u F_u$. Only the no-arrow term survives; principal-minor expansion gives the second equality.

For $C \subseteq R$ define

$$H_z^{C,R} = \sum_{y \notin R} [K(z,y) + K(z,C)M_C^{-1}K(C,y)]A_y,$$

again including the sink. In the quotient by the annihilator of a_C, the equations $A_c = J_c$, $c \in C$, allow elimination of C. This uses the scalar inverse M_C^{-1}, without cancelling a possibly singular addition operator. The Schur complement has minors d_D/d_C. Expand $\prod_{z\in R}(A_z + r_z F_z)$, apply (36) after elimination, and multiply back:

$$\prod_{z\in R} B_z = \sum_{D\subseteq R} (-1)^{|R\setminus D|} d_D a_D \sum_{C\subseteq D} \frac{r_{R\setminus C}}{d_C} \prod_{z\in R\setminus D} H_z^{C,R}. \tag{37}$$

Apply to the empty support and multiply by E. Surplus IDLA converts each all-active term to a binary law. For $D \subsetneq R$, an exit particle is lost or starts outside R; an initially occupied outside site cannot be removed by IDLA. Such a term cannot yield binary support R, proving (34). For $R = V$, every exit operator is the scalar $\kappa_z^C I$, yielding (35). Finally $EP = I$ and $P(D,V) = 0$ for $D \neq V$ give $\pi(V) = 1/Z_V$. □

Record: RN, Section 15, (42)–(45); proved algebraic identities, separate AI proof audit; finite symbolic certificates supplement the proof.

The auxiliary subset law proportional to $r_{R\setminus C}/d_C$, $C \subseteq R$, is dominated in inclusion order by the restriction to R of the corresponding law on any $S \supseteq R$. On fixed R it increases when the odds decrease coordinatewise. Its conditional insertion odds at z are

$$\frac{1}{r_z \sigma_z(C)}, \qquad \sigma_z(C) = 1 - K(z,z) - K(z,C)M_C^{-1}K(C,z) > 0.$$

Return-path sums increase with C, so common-uniform heat-bath updates give the comparison. This is an auxiliary law, not the ARW occupation distribution.

Record: RN, Section 15, (46); proved auxiliary comparison only.

5.6 Joint root retention, components, and sharp strictness

Theorem 5.5 (Boundary enlargement and joint retention)**.** *For $R \subseteq B$ and $y \notin B$,*

$$m_R(B \cup \{y\}) \leq m_R(B),$$

strictly exactly when the strongly connected component of y meets R. If each component D is equipped with the kernel killed on exit from D, then

$$m_R(B) = \prod_D m_{R\cap D}^D(B \cap D). \tag{38}$$

For every finite all-active $\eta \geq \mathbf{1}_R$,

$$\mathbb{P}_\eta(R \text{ occupied}) \geq \mathbb{P}_{\mathbf{1}_R}(R \text{ occupied}) = \frac{1}{m_R(R)}. \tag{39}$$

Let U be the union of components meeting R, with R removed, and let $\mathrm{Anc}(U)$ consist of vertices reaching U, including paths of length zero. Equality in (39) fails exactly when

$$\mathrm{supp}(\eta - \mathbf{1}_R) \cap \mathrm{Anc}(U) \neq \varnothing. \tag{40}$$

Record: RN, Section 16, (47a)–(47e); proved; separate AI audits of the prefix comparison and component/strictness argument.

Proof. Couple $U^- = U_B$ and $U^+ = U_{B\cup\{y\}}$ by keeping all arrows at their emission indices and inserting only a leading sleep bit at y: $b_y^+(1) = 1$, $b_y^+(k+1) = b_y^-(k)$. Fix a common root flip T. Let u^- be its lower canonical emissions and $t_z = 1 + I_z(u^-) = u_z^- + 1$ its root terminal indices. Use the lower reset book for comparison. At y, with k incoming arrows and the respective initial inputs, output lengths obey

$$\Phi_y^+(0) = \Phi_y^-(0) = 0, \qquad \Phi_y^+(k) = \Phi_y^-(k) + 1 \quad (k \geq 1).$$

Inductively every upper restoration edit occurs at an index at least t_z. Below t_z the root books agree; at t_z the lower reset bit is one, so the upper output is at least as large. Iterate the lower monotone processor map from zero. Its used root indices never exceed t_z; comparison with the current upper fixed point v gives $u^- \leq v$. An upper missing root is next edited at $1 + I_z(v) \geq t_z$, preserving the induction. Thus $u^+ \geq u^-$. This comparison uses only the lower used prefixes, not a global ordering of the edited books.

A flip is admissible exactly when its canonical emissions are zero: then all terminal root indices are one, the initial flips are restored, and later edits are reset. Hence the upper admissible family is contained in the lower one. Taking its nonnegative weighted mass in (33) proves the inequality.

The fresh admissibility family factors over components. If a local canonical vector is nonzero, first restore flipped first-root bits in all strict ancestors; those components emit nothing. Restore the chosen component before its descendants. Its evolution is local, and sleep-bit monotonicity bounds the original global vector below by this nonzero local vector. Conversely, if all local vectors vanish but a global vector does not, take its first nonzero component in topological order. It receives no emissions. Its final reset book is a nonzero local root-occupied candidate, contradicting local maximality. Thus a global flip is admissible exactly when every component flip is; independent component books give (38). Root-free factors have mass one.

If y's component avoids R, its enlargement has no effect. Otherwise choose $z \in R$ in that component and simple positive paths $z \to y$ and $y \to z$. Flip only the first z-bit and prescribe one packet following those paths. At other occupied sites use sleep bits; at vacant intermediate sites use jump bits. Require the first lower y-bit and the second z-bit to be one. Below, the packet stops at y, and restoring the missing z returns to zero emissions. Above, the inserted resident passes its arrow onward to z, leaving all roots occupied with nonzero emissions. Overlapping path visits use distinct records consistently. This positive-probability cylinder removes the flip $\{z\}$, of positive weight r_z, from the upper family, proving strictness.

Normalization in (32) now gives

$$1 = \sum_{B \supseteq R} \mu_\eta(B) m_R(B) \leq m_R(R) \mathbb{P}_\eta(R \text{ occupied}).$$

Its slack is exactly

$$m_R(R)\mathbb{P}_\eta(R \text{ occupied}) - 1 = \sum_{B \supseteq R} \mu_\eta(B)[m_R(R) - m_R(B)].$$

The bracket is positive exactly when B meets U. If no extra origin reaches U, the ancestor set initially has exactly $|R \cap \mathrm{Anc}(U)|$ particles and no incoming edges from its complement. A support containing R cannot also occupy U. Conversely, route an extra particle to $y \in U$, hold the root particles, send other extras to the sink, and finally sleep $R \cup \{y\}$. This finite positive-probability realization makes the slack positive, proving (40). □

5.7 The one-common-particle coupling and scope limitations

Theorem 5.6 (N1, including occupied-set order)**.** *For any $a, x \in V$, the terminal law from active $\delta_a + \delta_x$ stochastically dominates that from active δ_a by inclusion of sleeper sets.*

Proof. Hold the common particle at a, and move the extra until it sleeps outside $\{a\}$ or exits, suppressing ineffective sleeps at the occupied a. An exit permits identical continuation of the common particle. If the extra sleeps at y, run the common walk in both systems until it sleeps, exits, or reaches y. Before contact the final supports are ordered. On contact the lower system has one active particle at y, the upper has two, and their unread stacks are fresh. Repeat from y. Each round has a positive probability of losing the extra before its first sleep. On this fixed finite graph, the minimum over the finitely many common locations is positive, since every sleep probability is below one and the sink is accessible. The renewal terminates almost surely, yielding the coupling. With one common particle there are no other sleepers at contact. □

Record: R, Section 4, (5)–(6) and the full occupied-set coupling; proved for arbitrary finite graphs, not a finite enumeration.

Remark 5.7 (What the restoration results do not order). Joint retention compares arbitrary larger active inputs to their binary-root baseline; it does not compare two arbitrary larger inputs or each original root separately. Individual normalized entrance measures can have smaller target marginals, and a joint canonical-book fibre can have negative target increment even at N1. Therefore source-dependent entrance weights and averaging across fibres cannot be discarded. These are recorded obstructions, not counterexamples to the all-active conjecture.

For full roots only, the fresh admissible family has an arrow-only description: in iid K-arrow tails, a flip avoids the support L of all balanced arrow prefixes, where balance means incoming counts equal outgoing counts. Thus its mass is $\prod_{z\notin L}(1+r_z)$, and $\mathbb{P}(L\cap T=\varnothing)=d_V/d_{V\setminus T}$. The technical cycle-popping proof is a *proof reference, not reproduced*; cycle popping and Markov-loop terminology are credited background. For proper roots, later root sleep bits can change admissibility, so that arrow-only description fails.

Record: RN, Section 11, Pointwise fibre positivity fails; Section 12, (22) and entrance obstruction; Section 16, (48)–(49), Proper-root limitation and audits; proved scoped obstructions and full-root identities; cycle-popping proof reference, not reproduced.

6 Creator decompositions, loop laws, and activation structure

These results describe exact positive representations and normalized changes of books. They do not by themselves prove the general occupation inequality. The guide distinguishes constructions from their remaining sign conditions.

Family	Location
Collision lift, trace versus killing, and creator merges	6.1, 6.2
Positive creator law, loops, and active-cycle erasure	6.3, 6.4
Inverse-diagonal differences and coherent paired prefixes	6.4, 6.6
Forest promotion and the precise surplus reset	6.5, 6.8
Activation characters and the first-loss criterion	6.6, 6.9

6.1 Notation and the positive collision lift

Throughout, V is finite, K has an accessible absorbing sink s, and $0 < p_i < 1$. Put $q_i = 1 - p_i$, $r_i = q_i/p_i$, and write μ_η for the terminal sleeping-support law from finite all-active input η. Root-list subscripts denote their counting measures, μ_R means $\mu_{\mathbf{1}_R}$, and δ_C is point mass at sleeping support C. Let $S(t,u)$ be the first-sleep law of one active particle, $e_t = 1 - \sum_u S(t,u)$, and $h_t(u) = S(t,u)/S(u,u)$ its probability of visiting u before sleep or sink. A marked starting vertex is visited immediately. Superscript D means killing on visiting D; $a_t^D(y)$ is the first-entry probability at $y \in D$ before sleep or sink. Define $C_2(t,u) = S(t,u)h_t(u)$.

Proposition 6.1 (Three-particle contact renewal)**.** *On states (t,w), $t \neq w$, there is a strictly substochastic nonnegative kernel T with transitions to (u,w) of weight $C_2^w(t,u)$ and to (w,u) of weight*

$$X_w(t,u) = h_t(w)H_w(t,u) + S^w(t,u)a_t^{\{u,w\}}(w), \quad H_w(t,u) = S(t,u) + \frac{S(t,w)S(w,u)}{1 - S(w,w)}.$$

Here $u \neq w$ and $C_2^w(t,u) = S^w(t,u)^2/S^w(u,u)$. For any nonnegative terminal payoff, its expectation from two active particles at t and a sleeper at w is $(I-T)^{-1}$ applied to nonnegative stopped-history rewards. Moreover, if $t, u \neq w$,

$$X_w(t,u)h_w(u) = C_2(t,u) + \frac{C_2(t,w)C_2(w,u)}{1 - C_2(w,w)} - C_2^w(t,u) \geq 0. \tag{41}$$

No division by $h_w(u)$ is asserted when it vanishes.

Record: RN, §13, (23)–(25); proved positive renewal and trace/killing identity.

Proof. Explore one t-particle until sleep, sink, or contact with w, holding the other. After a sleep at u, explore the remaining particle until it hits u, w or stops; contacts produce the stated states. After contact at w, move the remaining t-particle with sleep disabled at w; this gives H_w and sink probability $E_w(t) = e_t + S(t,w)e_w/(1-S(w,w))$. All other branches are stable payoffs or two-particle continuation payoffs. Thus their reward is nonnegative. Every row has deficit at least $h_t(w)E_w(t) + e_t^w > 0$, by sink accessibility, so the finite resolvent is positive. Substitution of $S^w(t,u) = S(t,u) - h_t(w)S(w,u)$ and $a_t^{\{u,w\}}(w) = h_t(w) - [S^w(t,u)/S^w(u,u)]h_u(w)$ proves (41). Exploring a third arbitrary active root first reduces its law to the same stopped-history construction. □

The one-common target gain kernel is $D_1 = (I - C_2)^{-1}(S - C_2)$, a strictly substochastic matrix. Put $A((t,w),u) = h_t(w)(1 - h_t(w))\mathbf{1}\{u = w\}$, $J = (I-T)^{-1}A$, and $L((t,w),u) = h_t(w)\mathbf{1}\{u = w\} + J((t,w),u)$. The sleeper-insertion gain equals $\delta_w - L((t,w),\cdot)(I - D_1)$. Its

nonnegativity is exactly the still-unproved flow inequality $L(I - D_1) \le \delta_w$, not a consequence of $L \ge 0$. RN (29)–(31) retains distinct common roots and the actual extra-source first-sleep weights; their analogous sign also remains open. *Proof reference, not reproduced:* RN §13, (26)–(31). The chronological charging proposal fails already on $0 \to 1 \to 2 \to s$, RN (31b), while its actual occupation increment is positive. There is a proved terminal-pair case: if $\{a, w\}$ has no exits except to the original sink and $\eta(w) = 0$, the gain from adding a sleeper at w to any finite all-active η is $(1 - f)\delta_w + (f - \zeta)D_1(w, \cdot)$, where f is its probability of visiting w. To define ζ, first stabilize outside while holding arrivals in the pair, then run common surplus IDLA on those arrivals and the initial pair input; ζ is its probability of filling the pair. Here $0 \le \zeta \le f \le 1$; self-loops are allowed. *Proof reference, not reproduced:* RN §13, (31c)–(31d), using outside preprocessing and full-input reset on the terminal pair.

6.2 Hadamard decompositions and the scope of merge consistency

Theorem 6.2 (Positive creator mixtures)**.** *Let U be a nonnegative inverse M-matrix and $V_2 = U^{\odot 2}$, where $\odot$ means entrywise product. Then*

$$V_2^{-1}(U(:,i) \odot U(:,j)) \ge 0, \qquad [V_2^{-1}(U(:,i) \odot U(:,j))]_i \ge \frac{U(i,j)}{2U(i,i)} \quad (i \ne j),$$

and

$$U(a,i)U(b,j) + U(a,j)U(b,i) \ge 2[U(a,:) \odot U(b,:)]V_2^{-1}[U(:,i) \odot U(:,j)].$$

For the first-sleep potential S, and every list of $m \ge 1$ roots,

$$\lambda_{a_1,\dots,a_m} = [S(a_1,:) \odot \cdots \odot S(a_m,:)](S^{\odot m})^{-1}$$

is a subprobability row. In particular $\mu_{a+b} = \sum_t \lambda_{a,b}(t)\mu_{2t} + \zeta_{a,b}$ with $\zeta_{a,b} \ge 0$ as a measure. Successful merges are compatible with group sizes and binary merge order:

$$\sum_t \lambda_{a^r,b^s}(t)\lambda_{t^{r+s},c^k}(u) = \lambda_{a^r,b^s,c^k}(u), \qquad r, s, k \ge 1.$$

Repeated superscripts in this formula denote repeated roots, not powers.

Record: RN, §14, (32)–(38); proved consequences of the credited Hadamard theorem.

Key argument; longer decomposition proof referenced. The external input is the Hadamard theorem for inverse M-matrices and potentials credited in RN to Dellacherie–Martinez–San Martin, *Hadamard functions of inverse M-Matrices*, Theorems 2.2–2.3. Apply the nonpositive off-diagonal row of V_2^{-1} to the square of a linear combination of two U-columns vanishing at the selected coordinate. This gives positivity and the endpoint bound. Applying the same theorem to a triangular block inverse M-matrix gives the four-point inequality. Hölder's inequality with the off-diagonal weights of $(S^{\odot m})^{-1}$ and its transpose proves positivity and the mass bound for λ. Inserting $S^{\odot(r+s)}$ between successive merge factors proves the displayed compatibility. The nonnegative empty, singleton, and pair residual masses in the two-particle renewal, and the four-point block argument, are the *proof reference, not reproduced* in RN (32)–(36). □

For completeness, let A_i be active addition on sleeping supports, $J_i = \sum_y K(i,y)A_y$ with $A_s = I$, and let the formal sleep-insertion operator be $B_i = (A_i - q_i J_i)/p_i$. At every $p_i \ge 1/2$, B_i is stochastic. On a support $C \ni i$ its row is

$$\frac{2p_i - 1}{p_i}\delta_C J_i + \frac{q_i}{p_i}\delta_{C\setminus\{i\}}J_i^2;$$

on supports avoiding i it inserts a sleeper. The commuting endpoint polynomials from the merging theorem therefore give positive creator refinements in this regime. For arbitrary p, construct the weights at $p_i^* = \max(p_i, 1/2)$ and apply the saved stochastic parameter-transfer intertwining. This yields positive mixtures of m coincident-root laws plus a nonnegative residual at the original p, using the reference weights λ^{p^*}, not native low-p weights. The complete operator proof is RN (39)–(40), *proof reference, not reproduced.*

Order-consistent full merges do not define a deletion-consistent random partition. On the two-site half-sink kernel at $p = 1/2$, creators at $0, 0, 1$ have pair merge probabilities 1 and 8/17, but full merge probability 4/13. A partition would force the latter two to agree. Premultiplying protected-w contact blocks by $(C_2^w)^{-1}$ also gives a killed continuous-time generator with nonnegative cross-rates, rewards, and deficits; its detailed inequalities are RN §14 after (34), *proof reference, not reproduced.* Positive residual measures can still have negative target response; none of these decompositions establishes occupation order.

Record: RN, §14, (39)–(40), "Merge order does not give consistent creator partitions"; proved reference-transfer identities and exact scoped obstruction.

6.3 The creator sampler and active-cycle normalization

For $R \subseteq V$, write $d_R = \det(I - K_{RR})$, $d_\varnothing = 1$, $r_S = \prod_{i\in S} r_i$, $p_S = \prod_{i\in S} p_i$, and

$$Z_R = \sum_{S\subseteq R} r_S \frac{d_R}{d_{R\setminus S}}, \qquad \nu_R(C) = \frac{d_R r_{R\setminus C}}{Z_R d_C}.$$

A balanced word tuple has equal incoming and outgoing arrow counts at each vertex; $w(\alpha)$ is its product of jump probabilities. An exit forest on R assigns one arrow per root, with no directed cycle inside R; outside destinations retain their labels.

Theorem 6.3 (Positive full-graph creator law)**.** *The following construction has terminal law μ_R. Draw C with law ν_R; independently conditional on C, draw balanced arrow words α supported in C with law $d_C w(\alpha)$, an exit forest F on R with law $w(F)/d_R$, and fresh arrow tails. Concatenate $\alpha_i, F_i, tail_i$ and prescribe the bit at record* one *to be $\mathbf{1}\{i \in C\}$. All later bits and all outside books are fresh. Stabilize binary R on the full graph, retaining every outside return and every prepended word.*

Record: RN, §17, (50)–(53); proved via killed-root restoration.

Disintegration; restoration proof referenced. Ordinary cycle popping on R removes exactly α and leaves F on the cylinder with weight $w(\alpha)w(F)$. Balanced prefixes are poppable, and the terminal forest makes the execution complete. Consequently words and forest are independent, with the stated normalizations. Conditioning their popped range to lie in C changes only the prefix normalization to d_C.

The essential change of variables is RN (50): killed-system full-root restoration gives

$$\mu_R(D) = Z_R^{-1}\mathbb{E}\sum_{S\subseteq R} r_S \mathbf{1}\{S \cap L_R = \varnothing\}\mathbf{1}\{Y_R(S) = D\},$$

where L_R is the popped range of root arrow tails and $Y_R(S)$ is full stabilization with first bits zero on S and one elsewhere in R. Since $\mathbb{P}(L_R \cap S = \varnothing) = d_R/d_{R\setminus S}$, its disintegration is exactly the sampler. The killed-prefix replay and full inverse accounting are a *proof reference, not reproduced*, RN §17, (50); this is not proper-root restoration in the uncut graph. □

An intensity-one ordinary Markov-loop soup supplies the prefix counts:

$$\mathbb{E}\, t^{N(\alpha)} = \frac{d_C}{\det(I - (K_{ij}t_{ij})_{i,j\in C})}.$$

Given counts, independently shuffle each site's outgoing multiset uniformly. If $a(\ell)$ is the product of Poisson intensity powers divided by multiplicity factorials, the joint latent law is

$$\mathbb{P}(C, \ell) = \frac{d_R}{Z_R} r_{R\setminus C} a(\ell) \mathbf{1}\{\mathrm{range}(\ell) \subseteq C\}.$$

Given C, loops form its restricted soup; given loops, include their range and independently include other sites with probabilities p_i. Common soup and Bernoulli updates preserve inclusion for nested domains; a common empty soup, of probability $d_R > 0$ on an ambient R, resets the update history and normalizes the stationary coupling. This proves latent coupling only: continuation is not increasing in C, and even conditioning on the sampled forest can give gain $-7/390$.

Record: RN, §17, (52)–(53), Gibbs coupling and its two obstructions; proved.

Theorem 6.4 (Active-cycle erasure)**.** *In fresh paired bit/arrow books, erase one first record at each vertex of any current arrow cycle inside R containing a zero bit, and repeat. On the event that ordinary cycle popping terminates on every root domain (of iid probability one), this terminates with a unique erased prefix N and book $D_R\omega$. Simultaneously for every finite all-active $\eta \geq \mathbf{1}_R$, terminal supports agree and emission counts satisfy*

$$T_\eta(\omega) = T_\eta(D_R\omega), \qquad u_\eta(\omega) = N + u_\eta(D_R\omega).$$

If $\mathcal{B}_R$ is the event that every initial root-arrow cycle has only one-bits, then $(D_R)_*\mathbb{P} = \mathbb{P}(\,\cdot \mid \mathcal{B}_R)$. Thus $\mathcal{B}_R$ is independent of the terminal-support vector of any finite family of such inputs; no analogous odometer independence is claimed.

Record: RN, §18, (54)–(57); proved confluence, legal replay, and normalization.

Proof. Every erasure is ordinary arrow-cycle popping, which terminates. Distinct available cycles are disjoint, so their deletions commute; termination and this local diamond property give confluence. Deliver one reserved input per cycle vertex. Its zero-bits start propagation that emits each old cycle arrow once and supplies one replacement input at every vertex. Hold replacement-record emissions and all extra inputs. The resulting state is precisely that of the shifted book, proving the replay identities by iteration.

Let $\mathcal{W}_R$ be the set of finite word tuples exhaustible by active-cycle erasure. For every terminal book $y \in \mathcal{B}_R$, its preimages are exactly wy, $w \in \mathcal{W}_R$, counted once per tuple, not per erasure order. Their product weights factor. Normalization with the constant payoff identifies the total prefix weight as $1/\mathbb{P}(\mathcal{B}_R)$, proving the conditional law. Apply the simultaneous replay identity to any finite vector of endpoints. □

The different conditioning event that the ordinary popped range lies in the first-one set also preserves every such finite endpoint family; RN (57) proves this by killed-prefix replay. The two conditioning events are not identified.

If one cycle input $x = z_0$ is missing, a different exact boundary state must be retained. On a cycle $z_0 \to \cdots \to z_{k-1} \to z_0$, $k \geq 2$, with old bits not all one, after erasing its records the lower input is: active on the other vertices if $b_0 = 0$ and another bit is zero; reserved sleepers on $z_0, \ldots, z_{j-1}$, a hole at z_j, and active later vertices if $b_0 = 1$ and j is the first zero; or reserved sleepers away from x plus its unread old $(0, z_1)$ record if only $b_0 = 0$. Reserved

sleepers retain their old arrows. All additional active inputs are deferred. The respective bit weights are

$$q_x(1-p_{C\setminus\{x\}}), \qquad \Big(\prod_{i<j} p_{z_i}\Big)q_{z_j}, \qquad q_x p_{C\setminus\{x\}},$$

whose sum is $1-p_C$, where C is the cycle's vertex set. Multiply by the cycle-arrow weight, or normalize by $1-p_C$ conditional on that specified cycle being active. These are joint lower/upper endpoint identities, not an ordered coupling. *Proof reference, not reproduced:* RN §18, "A cycle containing one missing initial input", including the dormant component's one-time activation rule. Forgetting the sleeping arc or the unread zero record fails. The active-cycle conditioned first-bit field itself need not be log-supermodular.

6.4 Higher set inequalities, loop thinning, and paired words

Theorem 6.5 (Inverse-diagonal hierarchy and thinning)**.** *For fixed K and positive odds there are $w_S \geq 0$, depending only on K, r, such that*

$$\log Z_R = \sum_{i\in R} \log(1+r_i) - \sum_{\varnothing\neq S\subseteq R} w_S.$$

Thus, for disjoint A, B with $|B| \geq 2$, $\Delta_B \log Z_A = -\sum_{B\subseteq S\subseteq A\cup B} w_S \leq 0$; Δ_B denotes successive add-site differences. The pair inequality $Z_{A+x}Z_{A+y} \geq Z_A Z_{A+x+y}$ is strict exactly when x, y are mutually reachable inside $A\cup\{x,y\}$. These weights are intensities of rejected ordinary loops under an independent thinning that samples the creator's subset and balanced words.

Record: RN, §19, (58)–(63); proved convergent series theorem and joint sampler.

Series argument; convergence details referenced. Set $f(t) = \det(I - \operatorname{diag}(t)K)^{-1}$, let D_i set $t_i = 0$, and put $T_i = p_i I + q_i D_i$, $h_i(g) = 1 - D_i g/g$. A Schur complement shows $h_i(f)$ has nonnegative coefficients. This survives T_j because

$$h_i(T_j g) = D_j h_i(g) + p_j \frac{h_i(g) - D_j h_i(g)}{1 - q_j h_j(g)}.$$

Telescoping logarithms therefore gives nonnegative coefficients in $\log f - \log(\prod_i T_i)f$. At $t = \mathbf{1}_R$ this equals $-\log(p_R Z_R)$; grouping by support gives w_S and the mixed differences. The denominators converge geometrically at one by transience; the full convergence induction is RN (60), *proof reference, not reproduced.*

With separate edge variables, write $g = (\prod_i T_i)f$ and $c_n = [t^n]\log g$, $\ell_n = [t^n]\log f$. The same deletion argument gives $0 \leq c_n \leq \ell_n$ and $g(t)/g(1) = \exp(\sum_{n\neq 0} c_n(t^n - 1))$. Retain each ordinary loop with edge counts n with probability c_n/ℓ_n; zero-intensity classes are absent. Uniform local shuffles of retained counts, followed by Bernoulli inclusion outside their range, produce the creator law. Rejected intensity by vertex range is w_S; it is positive precisely for ranges of positive-weight closed walks, giving strictness. RN (61)–(63) supplies the full joint disintegration. □

The retention ratios are local and nondecreasing in each sleep probability. Thus one soup with common uniforms couples nested root domains and ordered sleep vectors, with nested latent subsets. Common local shuffles initially give subwords, not prefixes; the stronger paired-prefix coupling below is a different construction. This is RN (61)–(63), not an occupation comparison.

Released-input markers also have a proved scope. In the creator sampler on Q, put $E_x = \{x \notin C, F_x = s\}$ and $a_x = r_x K(x,s)$. For $E_S = \bigcap_{x\in S} E_x$,

$$\mathbb{E}[\mathbf{1}_{E_S} f(T_Q)] = \Big(\prod_{x\in S} a_x\Big)\frac{Z_{Q\setminus S}}{Z_Q}\mu_{Q\setminus S}(f).$$

After any positive-probability partial assignment of marker presences and absences, the remaining markers have pairwise negative correlations and both positive and negative cylinder bounds. *Proof reference, not reproduced:* RN (64)–(66), using the full log-difference hierarchy under deletion and inclusion-exclusion. Their covariance with terminal occupation is the remaining monotonicity question, not a proved sign; markers with $K(x,s) = 0$ give no such reformulation. A released initial root may later be reoccupied.

Theorem 6.6 (Periodic acceptance and coherent paired prefixes)**.** *For a nonzero edge-count table N of total size m with $K^N > 0$, let $\mathcal{W}_N$ be its rooted closed vertex words, each counted once. Then*

$$c_N = \frac{K^N}{m}\sum_{w\in\mathcal{W}_N} \mathbb{P}_p(\textit{full input stabilizes in the periodic decorated book } w).$$

Here every occurrence has an independent site coin, and the complete local paired record words repeat periodically, starting one active particle at each visited vertex. There is no sink in this auxiliary system. Not every periodic book stabilizes.

Let $\mathcal{M}_R$ be balanced paired words whose first used bit at each site is one, and let $\mathrm{wt}(h)$ *be the product of its arrow and coin probabilities. With $g_R(1) = p_R Z_R/d_R$, their normalized law is* $\mathrm{wt}(h)/g_R(1)$*. It has both a canonical Lyndon sampler and a coherent coupling in which smaller-domain words are actual initial paired-record prefixes.*

Record: RN, §20, (67)–(70a); proved with external trace-factorization inputs.

Prefix factorization. The periodic identity and the trace-monoid description of $\mathcal{M}_R$ are *proof reference, not reproduced*, RN (67)–(70). The credited external facts are connected conjugacy-class logarithmic weight $1/r$ for repetition exponent r, and complete Lyndon factorization, from Duchamp–Flouret–Luque, §§2.1–2.2, as cited there. Fix an ordered alphabet of the nonempty indecomposables of $\mathcal{M}_V$, commuting exactly on disjoint supports, and its standard Lyndon factorization. For each primitive Lyndon factor of weight a_L, independent Poisson(a_L^k/k) counts of its kth power give geometric multiplicity $\mathbb{P}(N_L = n) = (1-a_L)a_L^n$. Their prescribed Lyndon-order product has the stated law and is finite almost surely. Paired records are rearranged by conjugacy; accepted rooted words are not simply copied.

A heap records factor occurrences ordered by overlapping-support dependencies. For initial-prefix coherence, in a word on $R \cup \{x\}$ take the greatest heap order ideal consisting of letters whose support, and all predecessors' supports, lie in R. This gives the unique factorization $h = \alpha\beta$ with $\alpha \in \mathcal{M}_R$ maximal. Every minimal letter of nonempty β contains x. The product-weight bijection makes α, β independent, with laws $\mathrm{wt}(\alpha)/g_R(1)$ and $g_R(1)\mathrm{wt}(\beta)/g_{R+x}(1)$. Greatest old-prefix projections satisfy $\pi_R\pi_S = \pi_R$ for $R \subseteq S$, proving simultaneous correct marginals and initial-prefix inclusion. Appending the separately sampled forests does not preserve prefix nesting of completed books or order their terminal supports. Appending an independent forest record with a fresh Bernoulli bit at each site, and fresh tails, recovers Theorem 6.3 with $C = \mathrm{supp}(h) \cup \{i : \text{the appended bit is one}\}$. □

6.5 Forest promotion and the precise surplus reset

Proposition 6.7 (Promotion and creator resampling)**.** *Let $T = R \cup \{x\}$. Let $\mathcal{W}_x$ contain the empty tuple and balanced arrow words whose support contains x and whose first arrows away from x form a tree toward x on that support. In iid arrow books, pop on R first, then T. The old prefix α, additional prefix β, and final forest F_T are independent, with laws $d_R w(\alpha)$, $(d_T/d_R)w(\beta)$ on $\mathcal{W}_x$, and $w(F_T)/d_T$. The old forest $H_R(\beta, F_T)$ takes the first β arrow where present and the F_T arrow otherwise. For a specified positive-weight old forest f, the conditional promotion kernel is*

$$\frac{w(\beta)w(F)}{w(f)}\mathbf{1}\{\beta \in \mathcal{W}_x,\ F \text{ acyclic on } T, H_R(\beta, F) = f\}.$$

Promoting an R-creator state and independently adding x to its latent subset with probability p_x, then accepting exactly when $\operatorname{supp}\beta$ *lies in that enlarged subset, gives acceptance probability $p_x Z_T/Z_R \in [d_T/d_R, 1]$ and conditional output law equal to the T-creator law.*

Record: RN, §21, (71)–(77); proved cylinder and normalized resampling identities.

Cylinder argument; full resampling proof referenced. A prescribed staged cylinder has weight $w(\alpha)w(\beta)w(F)$. The first-stage stopping condition is exactly $\beta \in \mathcal{W}_x$; summing forests and balanced prefixes gives the independent factors. Conditioning the old forest gives the displayed quotient. Reversing the local words of a walk from x through its last return to x samples β and yields $\sum_{\beta \in \mathcal{W}_x} w(\beta) = d_R/d_T$. The accepted creator weight is $(p_x/Z_R)r_{T\setminus C}w(\alpha)w(\beta)w(F)$; unique staged decomposition gives the acceptance constant. The last-return sampler and the alternating conditional common-block updates (76)–(77) are *proof reference, not reproduced* in RN §21. Acceptance biases the common lower sample; it cannot be treated as an unconditioned lower marginal. □

Proposition 6.8 (Surplus traversal and residual books)**.** *Choose $z \in R$ independently of the creator book. Hold one active at each root and run a surplus from z to its first exit from R. Its loop erasure is the sampled forest path, with vertex set P. It consumes that path and a balanced prefix ε, uses no fresh arrow beyond the forest, and visits $D = P \cup \operatorname{supp}\varepsilon$. Put $U = R \setminus P$, $A = R \setminus D$, and*

$$\Phi(U, A) = d_U \sum_{B \subseteq A} \frac{p_B q_{A\setminus B}}{d_{(U\setminus A)\cup B}}, \qquad q_B = \prod_{i \in B} q_i.$$

A specified route ω, including its labelled exit, has probability

$$w(\omega)p_{\operatorname{supp}\varepsilon}\frac{\Phi(U, A)}{\Phi(R, R)}, \qquad \Phi(R, R) = p_R Z_R.$$

Conditional on the entire route, books on P are iid fresh, and those on U are iid conditioned on their popped range meeting A only at current first-one bits. All visited-site remaining bits are fresh; unvisited first flags are retained. The whole remaining root book is iid, independently of the route and exit, if $D = R$, or if $|R| = 2$ and neither root has a self-loop. No general iid reset is claimed.

Record: RN, §21, (78)–(81); proved conditional reset kernel.

Proof reference, not reproduced. The unique decomposition into balanced excursions along the forest path and a residual balanced word, and its cylinder disintegration, are RN (78)–(81). The exit marginal is ordinary labelled harmonic measure. Treating P as absorbing is only for the residual cycle-popping test; physical continuation permits all returns. A newly sampled residual forest can be made independent, but need not retain the consumed path. The root reset still leaves the lower copy's outside path unread: RN (81a) has a fully averaged contact gain $-10613/156325$, so no-contact credit cannot be discarded.

6.6 Activation characters and the exact first-loss criterion

Use the commuting addition operators above, $F_i = A_i - J_i$, and $U_C = (\prod_{i\notin C} A_i)(\prod_{i\in C} F_i)$. Let $h_C(z) = \mathbb{P}_z(s \text{ before } C)$ for the ordinary walk and $j_C(i) = K(i,s) + \sum_{z\notin C} K(i,z)h_C(z)$. If every $K(i,s) > 0$, the nonzero numbers

$$\lambda_C = (-1)^{|C|} \prod_{i\notin C} h_C(i) \prod_{i\in C} j_C(i)$$

give orthogonal projectors U_C/λ_C summing to I, with $U_C A_z = h_C(z)U_C$. The local addition relations prove orthogonality; surplus-IDLA triangularization proves nonvanishing and completeness. *Proof reference, not reproduced:* RN §22, (82)–(83). Zero direct-sink rates may give zero λ_C; division is then invalid.

For $f_T(B) = \mathbf{1}\{T \subseteq B\}$ put $\Theta_T(C) = r_C \delta_\varnothing U_C f_T$. Under mere sink accessibility, full base input asleep exactly on E, with extra active input η, has target probability

$$m_T(E;\eta) = \sum_{C\subseteq E} \Theta_T(C) \prod_z h_C(z)^{\eta_z}, \qquad \sum_C \Theta_T(C) = 1.$$

Use $0^0 = 1$. These identities require no nondegenerate projectors. For full occupation, $\Theta_V(C) = r_C d_V/(Z_V d_{V\setminus C}) > 0$. For a singleton mark, $\Theta_{\{v\}}(\{x\}) = r_x\sigma_x[\mu_V(f_v) - \mu_{V\setminus\{x\}}(f_v)]$, where $\sigma_x = j_{\{x\}}(x) > 0$; the added-target case $v = x$ is strictly positive. Full-target weighted exits and the DAG/SCC composition are recorded in RN (89), (93), and §23 and in the directed-decomposition section, rather than duplicated here.

Record: RN, §22, (84)–(93); proved algebraic identities and scoped positive rows.

General proper-target coefficient positivity is false. With columns $0, 1, 2, s$, take

$$K = \begin{pmatrix} 0 & 0 & 1/2 & 1/2 \\ 1/2 & 0 & 0 & 1/2 \\ 15/16 & 3/64 & 0 & 1/64 \end{pmatrix}, \qquad p_i = 1/1001.$$

Then $\Theta_{\{0\}}(\{1,2\}) < 0$, despite irreducibility and positive sink rates. *Proof reference, not reproduced:* RN §22, "A three-site obstruction to general field positivity", exact rational absorption and its independent pair-resolvent calculation. This refutes a positive auxiliary set representation for every proper-target row, not all-active occupation monotonicity.

Theorem 6.9 (Excursion dual before loss)**.** *For binary occupied support B and active set $A \subseteq B$, an exact ARW block chooses $i \in A$: with probability p_i it sleeps; otherwise an ordinary walk from i, beginning with its first jump, stops on return to i, a vacancy, or the sink. Intermediate occupied sites retain one active resident and send the surplus onward. This defines an exact finite kernel on vacant/sleeping/active states with fresh unused records.*

At full support, graphical reversal before any particle loss gives a subset chain H with death rate p_i at $i \in H$ and birth rate $q_i e_i(H)$ at $i \notin H$, where

$$e_i(H) = \mathbb{P}_i(H \text{ or } s \text{ before first return to } i) = \frac{d_{V\setminus H}}{d_{V\setminus(H\cup\{i\})}}.$$

It is irreducible and reversible with law $\pi(H) = r_H d_V/[Z_V d_{V\setminus H}]$. Its time is auxiliary block time, not the physical ARW clock.

Record: RN, §24, binary kernel and (105)–(107); proved, including self-loops.

Proof. Holding one resident and moving only surplus gives legal excursion blocks. The ordinary sink-hitting-time potential decreases in expectation by at least $\min_i q_i$ per nonterminal block, giving finite mean termination. Recovery reverses to deleting i from H; a returned excursion adds i when its range meets H; a sink excursion adds i unconditionally, recording that a loss mark must have an inactive source. These maps give the stated rates. The diagonal Green cofactor gives the determinant ratio, and $\pi(H+i)/\pi(H) = r_i e_i(H)$ proves detailed balance. Reversing a finite graphical interval gives $\mathbb{P}_{V\setminus E}(\text{no loss by } t) = \mathbb{P}_{\varnothing}(H_t \subseteq E)$, where the left initial active set is $V \setminus E$ and occupied support is full. Uniform positive probability of all recoveries and no excursions empties the activity process. Taking $t \to \infty$ recovers the full-target row. □

For a proper target let $\theta(H) = \Theta_T(H)$, and let $\theta^{(i)}(H)$ be its flag coefficient from base support $V \setminus \{i\}$ on the original graph, with the vacancy still available for later returns. Put $h_H(i) = \mathbb{P}_i(s \text{ before } H)$. The exact first-loss edge identity is

$$q_i e_i(H)\theta(H) - p_i\theta(H+i) = q_i e_i(H) h_H(i)\theta^{(i)}(H).$$

Divide by detailed balance and set $u = \theta/\pi$. Then $u(H) - u(H+i) = h_H(i)\theta^{(i)}(H)/\pi(H)$. These edge differences have zero circulation. Define $G(\varnothing) = 0$ with increment $G(H+i) - G(H)$ equal to that last right-hand side. Since $\sum_H \theta(H) = 1$,

$$\theta(H) = \pi(H)[1 + \mathbb{E}_\pi G - G(H)], \qquad \theta \geq 0 \iff \max_H G(H) - \mathbb{E}_\pi G \leq 1.$$

The edge identity follows by conditioning the first instruction and taking flag differences: any excursion visiting a marked flag erases its coefficient. Detailed balance then proves the reconstruction above. The height bound is false in general by the proper-target example; reversibility and zero circulation do not prove it. After the first loss, both vacancy location and activation information must be retained.

Record: RN, §24, (108)–(112); proved identities and exact criterion, not a universal sign theorem.

7 Backward identities, finite books, and release comparisons

This section collects algebraic and backward-record results from the maintained AI research notes. Its sources are the current `response-renewal.md` (RN) and the explicitly imported AI notes indexed in `evidence/README.md`. Imported results are identified below; standard external theorems are credited inputs. No bibliographic novelty claim is made. The scope remains finite, sink-accessible ARW with all-active input for the occupation conjecture. Auxiliary sleeping configurations occur in identities, without an assumption that arbitrary sleeping inputs are attractive.

The result guide is as follows. Theorem 7.1 gives the finite complete-block inverse; Theorem 7.4 settles invertibility of the older signed response renewal. Proposition 7.5 gives exact last-arrival deletion. Theorems 7.6–7.7 connect conditional loops, ordered books, and complementary-sleep duality. Theorem 7.8 gives an absolutely convergent inverse using stable cuts. Theorems 7.10 and 7.13 give finite envelopes and the recurrent rotor correspondence. Theorems 7.15 and 7.16 give release interpolation and the two-site prefix and forest-degree inequalities. The final subsection records decisive failures and the saved audit boundary at RN (J116).

7.1 Complete blocks, signed renewal, and deletion

Let V be finite, let ∂ be an absorbing sink accessible from every vertex, and let $K(i,j)$ be the routing probabilities, including $j=\partial$. In matrix products on V, K denotes its internal restriction. Write $p_i \in (0,1)$ for sleep probability, $q_i = 1-p_i$, and $r_i = q_i/p_i$ for jump odds. Products over a set have a subscript, for example $p_B = \prod_{i\in B} p_i$. Stable configurations are sleeping supports $B \subseteq V$; matrices act on probability rows from the right. Let A_i add one active particle at i and stabilize, set $A_\partial = I$, and write $J_i = \sum_j K(i,j)A_j$. The Abelian property gives commuting additions and $A_i^2 = A_iJ_i$. Let $P(A,B)$ be the final-support law from one active particle at each site of A, including $A=\varnothing$. With an input argument, $q_v(\eta)$ denotes its final occupation probability, distinct from the scalar jump probability $q_v = 1-p_v$.

Theorem 7.1 (Complete-block inverse)**.** *For $S \subseteq A$, let $R_S(A,\cdot)$ first move all particles selected in S by independent prescribed K-steps, then perform ordinary IDLA on the resulting all-active configuration, retaining at most one particle per site. All selected initial steps precede IDLA. Then*

$$E(A,\cdot) = p_A^{-1} \sum_{S\subseteq A} (-1)^{|S|} q_S R_S(A,\cdot), \qquad EP = PE = I.$$

Neither $E \geq 0$ nor a representation by successive single-site forced moves is asserted.

Record: CB, §§2–4, (1)–(7); imported AI result, proof reproduced.

Proof. Put $B_i = (A_i - q_iJ_i)/p_i$. If $i \notin C$, conditioning the first instruction gives $\delta_C B_i = \delta_{C\cup\{i\}}$. All B_i commute, so $\delta_\varnothing \prod_{i\in A} B_i = \delta_A$. Expand this product. Its S term first makes the prescribed jumps of precisely the particles in S. Legal surplus topplings, equivalently the quadratic addition relations, reduce its law to $R_S(A,\cdot)P$. Thus the expansion says $EP = I$. Squareness gives the other inverse identity. In particular the formula concerns a simultaneous block; sequential forced moves can reoccupy sites while other selected particles remain parked and have a different law. □

Corollary 7.2 (Waking transport)**.** *For $0 \leq r \leq s$, $T(r,s) = E_rP_s$ is stochastic. From a sleeping support it independently wakes the current sleeper at i with probability $1 - r_i/s_i$, then stabilizes at odds s. An unchanged zero coordinate is left alone. Moreover $P_rT(r,s) = P_s$ and $T(r,s)T(s,t) = T(r,t)$ for $r \leq s \leq t$.*

Record: VA, §1, (1)–(4); imported AI result, compact proof reproduced.

Proof. In the addition algebra at s, put $B_i(u) = (1 + u_i)A_i - u_iJ_i$. For $s_i > 0$, $B_i(r) = (r_i/s_i)B_i(s) + (1 - r_i/s_i)A_i$. Expand the product over the initial support, applying sleep-creation factors first. Its coefficients are independent waking probabilities. The inverse identities give composition. At zero odds use the buffer limit: lone particles sleep immediately and surplus particles still jump. □

See Section 2.4 for further parameter-flow consequences.

For a terminal payoff f, define $g(B, x) = \mathbb{E}[f(\mathrm{Stab}(\text{sleep } B + \text{active } x))] - f(B)$, with $g(B, \partial) = 0$, and $t_x(C) = f(C + x) - f(C)$ when $x \notin C$. For $x \in B$, define $W_x(B, C)$ by removing the resident at x, forcing one K-jump from x, and stabilizing among the remaining sleepers, killing the history on any return to x. A self-loop is a return; sink exits are allowed. Then $x \notin C$ and $|C| \le |B|$.

Lemma 7.3 (First return and killed creation)**.** *The response equations are*

$$g(B, x) = p_x t_x(B) + q_x \sum_z K(x, z)g(B, z), \qquad x \notin B,$$

$$g(B, x) = \sum_z K(x, z)g(B, z) + \sum_C W_x(B, C)[g(C, x) - t_x(C)], \qquad x \in B.$$

If $N_x(A, C)$ is binary stabilization killed on any visit to x, then

$$(EN_x)(B, C) = \begin{cases} \mathbf{1}_{B=C}, & x \notin B, \\ -(q_x/p_x)W_x(B, C), & x \in B. \end{cases}$$

Record: RN, §§1–3, (1)–(5); maintained proof, reproduced.

Proof. The vacant equation is first-instruction conditioning. At an occupied mark, force one of the two active particles to jump and hold the other. Compare with that forced jump among the original sleepers. Up to the first return to x the outside histories agree; at a return both have two active particles at x and cancel. A no-return history ends at C and leaves the displayed difference. Fresh suffixes justify its continuation. For the killed identity use commuting killed additions $\widehat{A}_x = 0$, $\widehat{A}_\partial = I$ and $\widehat{B}_z = (\widehat{A}_z - q_z\widehat{J}_z)/p_z$. Their product creates B if $x \notin B$. Otherwise apply the x factor last: it is $-(q_x/p_x)\widehat{J}_x$ and produces the killed release law. Expanding the same product by Theorem 7.1 gives EN_x. An IDLA visit to x leaves a particle there, so subsequent killed stabilization discards exactly those terms. □

Theorem 7.4 (Invertibility of the signed response renewal)**.** *Fix an integer common-particle bound $N \ge 1$. The finite signed first-contact renewal matrix $I - M$ on formal coordinates (A, x), $1 \le |A| \le N$, $x \notin A$, is invertible. This does not assert $\rho(M) < 1$, positivity of its inverse, or convergence of its signed Neumann series.*

Record: RN, §§2–5, (3)–(6); maintained proof, reproduced; signed construction from evidence/notes/reconstruction.txt, §6.

Proof. For nonempty B, let $Z_B(z, t)$ be the probability that a single particle started at z first reaches B at t, before sleeping outside B or exiting. At $z \in B$ this is the identity kernel. Eliminating vacant marks in the homogeneous response equations gives the nonnegative matrix

$$\mathcal{T}((B, x), (C, t)) = \mathbf{1}_{C=B} \sum_z K(x, z)Z_B(z, t) + W_x(B, C)Z_C(x, t), \quad x \in B,\ t \in C.$$

Empty C contributes zero. This matrix never increases particle count, but its row sums need not be at most one.

For its count-k diagonal block take $h_k(B,x) = \mathbb{P}(|\operatorname{Stab}(\text{sleep } B + \text{active } x)| < k)$, $|B| = k$, $x \in B$. Apply the response equations to $f(S) = \mathbf{1}_{|S|<k}$ and eliminate vacant marks. Particle conservation gives exactly

$$h_k = \mathcal{T}_k h_k + b_k, \qquad b_k(B,x) = \sum_{|C|=k-1} W_x(B,C)\mathbb{P}(|\operatorname{Stab}(\text{sleep } C + \text{active } x)| < k).$$

Choose a simple positive-probability path from x to the sink. A forced particle can follow it without returning to x, restoring each resident it encounters and leaving $C = B \setminus \{x\}$ unchanged. Call the probability of this specified event $\alpha(B,x) > 0$. At a previously occupied internal site its factor is $p_z K(z,z')$; at a vacant site it is $q_z K(z,z')$. An ordinary addition at the vacant x can first jump and perform the same event. Thus $b_k(B,x) \geq q_x \alpha(B,x)^2 > 0$. Since $0 < h_k \leq 1$, setting $\epsilon_k = \min_{B,x} q_x \alpha(B,x)^2$ gives $\mathcal{T}_k h_k \leq (1-\epsilon_k)h_k$. The weighted maximum norm proves $\rho(\mathcal{T}_k) < 1$. Triangularity by count yields $\rho(\mathcal{T}) < 1$ and $(I-\mathcal{T})^{-1} = \sum_{j\geq 0} \mathcal{T}^j \geq 0$. These bounds concern this fixed finite graph and its fixed parameters.

To connect to the formal signed renewal, suppose $D = MD$. Extend $D(\varnothing, x) = D(A, \partial) = 0$ and extend occupied-source coordinates by $D(A,x) = \sum_z K(x,z)D(A,z)$ until a vacancy or the sink is reached. This extension is unique by sink accessibility. For each x put $g_x = ED_x$. Unreading sleepers and reducing multiplicities by IDLA is precisely stabilization applied to this arbitrary formal column; it is not necessary that these columns arise from one payoff. If $H_A(x,y)$ is the extra walk's first-sleep law with sleep allowed only outside A, first contact gives

$$(MD)(A,x) = \sum_{y \notin A} H_A(x,y)[(P - N_y)g_y](A).$$

Inverting this one-particle first-sleep equation gives $q_x P(g_x - \sum_z K(x,z)g_z) = -p_x N_x g_x$. Apply E and the killed identity. At a vacant mark one gets $g(B,x) = q_x \sum_z K(x,z) g(B,z)$; at an occupied mark one gets the homogeneous occupied equation above. Hence $g = \mathcal{T} g$ on occupied marks, and the contraction forces $g = 0$, then $D = 0$. Finite-dimensional injectivity proves the claim. □

Proposition 7.5 (Last-arrival deletion and its invariant cone)**.** *For every finite all-active input η, including multiplicities, let μ_η be its final law and N_η^v its law killed on any visit to v. For $C \not\ni v$,*

$$\mu_\eta(C) = N_\eta^v(C) + \frac{q_v}{p_v} \sum_{B \ni v} \mu_\eta(B) W_v(B,C).$$

Let $L_v(B,C) = \mathbf{1}_{B=C}$ for $v \notin B$ and $L_v(B,C) = -(q_v/p_v)W_v(B,C)$ otherwise. For killed additions $\widehat{A}_y^v$ on v-avoiding supports,

$$A_y L_v = L_v \widehat{A}_y^v.$$

Consequently $\{\mu \geq 0 : \mu L_v \geq 0 \text{ for all } v\}$ is invariant under every active addition and contains every all-active creation law.

Record: RN, §§6–7, (7)–(10); maintained proofs, reproduced.

Proof. In a completed history ending occupied at v, change its last target sleep bit to zero, retain its potential arrow, and follow the emitted packet through a continuation avoiding v. The old terminal index then stays unchanged. Conversely, in a history visiting v but ending vacant there, turn its last target jump into a sleep and suppress the packet whose later continuation avoids v. These are inverse operations on complete books. A scheduler revealing

no unused arrows leaves the terminal reserved arrows and future suffixes independent; the continuation law is W_v. The changed bit has ratio q_v/p_v. Partition by finite histories and add the no-visit histories. All-active input ensures the required creator bit exists.

For the intertwining, multiply both sides by P. The deletion identity identifies both rows with the killed final law from binary input plus one active particle at y. Abelian stabilization and its visited set permit either order. Cancel the invertible P. Nonnegativity of killed additions proves cone invariance; the count-N slice maps into count $N+1$, not necessarily into itself. □

Writing $a_\eta = \mathbb{P}_\eta(\text{no visit to } v)$, $Q_\eta = \mu_\eta(v \text{ occupied})$, and $w_v(B) = \sum_C W_v(B,C)$ gives

$$1 - a_\eta = Q_\eta \left[1 + (q_v/p_v)\mathbb{E}(w_v(B) \mid v \text{ occupied})\right].$$

Conditional inclusion order of the occupied-target backgrounds is false (RN §6). This identity therefore does not close the target inequality.

7.2 Zero-loss bridges and positive remainders

Use paired records $(\beta_i(j), d_i(j))$, with independent Bernoulli sleep bits and K-arrows. An arrival displaces any resident along its reserved arrow; the new arrival stays if its bit is one and otherwise emits its own arrow. If m_i arrows have been emitted and $I_i(m)$ enter i, the stable-cut equation for all-active input η is

$$m_i = \eta_i + I_i(m) - \beta_i(\eta_i + I_i(m)), \qquad \beta_i(0) = 0.$$

The local function $n - \beta_i(n)$ is nondecreasing. Least action gives the physical least solution and monotonicity of emissions in input and in the reversed sleep-bit order. This is the buffer representation of RN §11, (17), not an identification of raw instruction histories. A finite book is a collection of ordered local words. Its weight $\mathrm{wt}_p(h)$ is the product of its bit and arrow probabilities. Concatenation is sitewise; no chronological interleaving factor is used.

Theorem 7.6 (Ordered normal form and rejected-loop currents)**.** *Let $\mathcal{H}$ be the balanced internal books, and $\mathcal{M}$ those whose first bit is one at every used site, including the empty book. Let $\mathcal{R}_*$ be the books in $\mathcal{H}$ having no nonempty right divisor in $\mathcal{M}$. Then $\mathcal{H} = \mathcal{R}_*\mathcal{M}$ uniquely. The first factor has the ordered emission-book law of binary full-input ARW conditioned on final full occupation when weighted proportionally to* wt_p.

For internal edge marks t_{ij} put $W_{ij} = K(i,j)t_{ij}$, $d_C(W) = \det(I - W_{CC})$, and

$$Z(r;W) = d_V(W) \sum_{C \subseteq V} \frac{r_{V\setminus C}}{d_C(W)}.$$

The conditional current N has generating function

$$\mathbb{E}[t^N \mid \textit{full}] = \frac{Z(r;K)}{Z(r;W)} = \exp\left(\sum_{\nu \neq 0} a_\nu(p)(t^\nu - 1)\right), \quad a_\nu(p) \geq 0.$$

The a_ν are the rejected ordinary Markov-loop intensities of RN §§19–20. Their sum is finite and they decrease as sleep probabilities increase. This is a conditional compound-Poisson law.

Record: RN, §37, (J1)–(J10); maintained proofs, independent AI audits recorded; loop-intensity input from §§19–20, proof reference, not reproduced.

Proof. Append a terminal one after each local word of $h \in \mathcal{H}$ and stabilize full input. Balance makes the full length vector a stable upper bound, so the actual prefix a ends full without loss. Its remaining balanced suffix b has first bits one. A nonempty right $\mathcal{M}$ divisor of a would give an earlier stable cut, impossible. Conversely, in any proposed $h = ab$ of this form, the actual prefix a' lies below a. If strict, $a = a'c$ with nonempty $c \in \mathcal{M}$, a contradiction. This proves uniqueness.

The maintained cycle-popping and creator identities give

$$\sum_{h\in\mathcal{H}} \mathrm{wt}_p(h) = \frac{1}{d_V(W)}, \qquad \sum_{h\in\mathcal{M}} \mathrm{wt}_p(h) = \frac{p_V Z(r;W)}{d_V(W)},$$

with edge marks included in the weights. Thus $\mathcal{R}_*$ has weight $1/(p_V Z(r;W))$. A prescribed physical prefix has cylinder probability $p_V \,\mathrm{wt}_p(a)$, so its full-event probability is $1/Z(r;K)$. Normalization proves the generating-function ratio and independence of the two factors under the ordinary Boltzmann book law. Finally RN (J2) identifies the rejected-loop intensities through

$$-\log(p_V Z(r;W)) = \sum_{\nu\neq 0} a_\nu(p) t^\nu.$$

Subtract its value at $t = 1$ to obtain the last equality. The positivity of those intensities uses the cited earlier construction. □

For full-support input asleep on S and active on $V \setminus S$, put $r_i^S = r_i$ on S and zero elsewhere, and $p_i^S = 1/(1 + r_i^S)$. The full-event probability is $Z(r^S;K)/Z(r;K)$ and its conditional loop intensities are $a_\nu(p) - a_\nu(p^S) \geq 0$; every atom meets the active set. For arbitrary all-active η with positive full-filling probability, the conditional current splits independently into complete IDLA conditioned to fill and the binary rejected-loop current. Ordered local books admit the same factorization, with all exits in the IDLA factor. Along nested activation sets, greatest admissible suffixes couple the ordered bridges coherently; exchanging activation order refactors the same book and preserves product weights. These are full-conditioned statements, not comparisons of different terminal-support probabilities.

Record: RN, §37, (J5), (J7)–(J18), (J23)–(J25), (J51); maintained audited results; proof references, not reproduced for the activation exchange and open-input extensions.

Theorem 7.7 (Remainder factorization and complementary pairing)**.** *A remainder $h : C \to B$ satisfies $\mathbf{1}_C + I(|h|) - |h| = \mathbf{1}_B$, has first bit one at every used site in C, and last bit zero at every used site outside B. Empty words impose neither test. Set*

$$R_p(C,B) = \sum_{h:C\to B} W_p(h), \qquad W_p(h) = \frac{p_B}{p_C}\,\mathrm{wt}_p(h).$$

All sums converge. If $Q_\eta(B)$ sums formal stable-book weights $p_B \,\mathrm{wt}_p(h)$ for input η, then $Q_\eta = P_\eta R_p$. Here P_η is the actual terminal law, and $R_p(C,B) = 0$ for $|B| > |C|$.

Let $\mu_p(B) = p_B q_{V\setminus B}$ and $d = \det(I - K_{VV})$. In the addition algebra $z_i(z_i - J_i(z)) = 0$, with $J_i(z) = K(i,\partial) + \sum_j K(i,j) z_j$, let Λ extract the coefficient of z_V in its squarefree basis. Put $w_i^r = z_i + r_i(z_i - J_i)$ and $H_{r,s}(C,D) = \Lambda(w_C^r w_D^s)$. Then

$$R_p(C,B) = \frac{\mu_p(B)}{d} H_{r,1/r}(C,B^c), \qquad \mu_p(C) R_p(C,B) = \mu_p(B) R_{1-p}(B^c, C^c).$$

For the full-input stationary law π_{1-p} at complementary sleeps,

$$f_B = \frac{d\,\pi_{1-p}(B^c)}{\mu_p(B)} > 0, \qquad R_p f = \mathbf{1}, \qquad L_p = R_p \,\mathrm{diag}(f)$$

is a stochastic link.

Record: RN, §37, (J26)–(J30), (J34)–(J36); maintained audited results; complementary pairing and stationarity input: §36, (D1)–(D8).

Proof. Stabilize a formal book up to its physical least cut, ending at C. Its unused suffix is exactly a C-to-B remainder. Conversely such concatenations replay the same physical prefix. Their weights factor because the interface terminal factor cancels. This proves $Q_\eta = P_\eta R$. For convergence choose $h > 0$ and $a < 1$ with $Kh \le ah$. After the edge rescaling $i \to j$ by h_j/h_i, a book of length ℓ with fixed exit count has total arrow weight bounded by a polynomial in ℓ times a^ℓ. Local word-length choices and exit placements supply only polynomial factors. Bit restrictions do not spoil summability.

For the pairing, take the large torus $|z_i| = th_i$ so that $|J_i(z)/z_i| < 1$. Summing the local geometric word series gives

$$Q_\eta(B) = \mu_p(B)\Phi(z^\eta w_{B^c}^{1/r}), \qquad \Phi(g) = \mathrm{CT}\,\frac{g(z)}{\prod_i [z_i - J_i(z)]}.$$

This is a large-torus expansion, not an expansion at the origin. Cancellation shows $\Phi(z_i(z_i - J_i)g) = 0$. On squarefree monomials z_C of degree below $|V|$, total degree excludes a constant term; on z_V balanced internal books give $1/d$. Hence $\Phi = \Lambda/d$. On binary inputs multiply $Q = PR$ by E to obtain the pairing identity. Local reversal and bit complementation preserve arrow directions and send $C \to B$ to $B^c \to C^c$; their boundary factors give reciprocity. Finally the complementary-pairing normalization $H_{r,1/r}\pi_{1-p}^{\dagger T} = \mathbf{1}$ from RN (D6) gives $Rf = \mathbf{1}$. □

7.3 Stable-cut lattices and convergent inversion

Theorem 7.8 (Stable-cut inversion)**.** *Fix a remainder $h : C \to B$ with length vector ℓ, append terminal ones on B, and put*

$$F_i(m) = \mathbf{1}_C(i) + I_i(m) - \beta_i(\mathbf{1}_C(i) + I_i(m)), \qquad 0 \le m \le \ell.$$

Its fixed points form a finite lattice $\mathcal{L}(h)$ of stable cuts. Cuts with a specified intermediate support form a distributive sublattice. Let μ_h be the bottom-to-top Möbius number, with $\mu_h = 1$ for the identity book. For $n = |V|$, put $M_n = 2^{2^n - 1 + n}$. Then

$$|\mu_h| \le M_n, \qquad [R_p(t)^{-1}](C, B) = \sum_{h : C \to B} \mu_h W_p(h) t^{|h|}.$$

The latter sum is absolutely convergent per ordered paired book for $|t| < 1/\rho(K_{VV})$, with infinite radius when $\rho(K_{VV}) = 0$. In particular it converges at $t = 1$. The unreduced Neumann series in $R - I$ is not asserted to converge.

Record: RN, §37, (J40), (J45)–(J48); maintained proof, independent AI audit recorded, reproduced. Formal categorical Möbius inversion is standard background, as credited there to Leinster, *Notions of Möbius inversion*, §1 and Theorem 6.1.

Proof. The map F is monotone, with $F(0) = 0$ and $F(\ell) = \ell$. Iterating upward from the coordinate maximum of two fixed points gives their least fixed upper bound; downward iteration from their coordinate minimum gives their greatest fixed lower bound. A fixed point has support $D = \mathbf{1}_C + I(m) - m$ and is precisely a factorization cut.

For cuts m, n both ending at D, write $u = \max(m, n)$ and $v = \min(m, n)$. Monotonicity gives $I(u) \ge u + \mathbf{1}_D - \mathbf{1}_C$ and the reversed inequality at v. Both original cuts use $k = |C| - |D|$ exits. Summing gives at most k exits at u and at least k at v, while prefix inclusion gives the opposite bounds. All inequalities are equalities. At each coordinate the appropriate next-one

or last-zero bit is inherited from one original cut. Thus u, v are stable with endpoint D, proving distributivity. The fiber need not be an interval of the whole lattice.

At most one atom can have a specified endpoint $D \neq C$: otherwise their coordinate minimum is a nonzero cut in the same fiber. Atoms ending at C have disjoint positive coordinate supports, so there are at most n. Hence the number of atoms is at most $2^n - 1 + n$. The crosscut formula

$$\mu_h = \sum_{S \subseteq \mathrm{Atoms}(\mathcal{L}(h)) : \bigvee S = \ell} (-1)^{|S|}$$

gives the uniform bound. For completeness, summing this expression over all upper endpoints below a given cut gives $(1-1)^a$, where a is the number of atoms below that cut, exactly the Möbius recurrence.

A factorization into nonempty remainders is exactly a strict chain of cuts. Its boundary factors telescope, so formal inversion assigns the coefficient $\mu_h W_p(h)$. Choose $a > \rho(K)$ and a positive vector u with $Ku \leq au$, with $a|t| < 1$. The fixed-exit geometric estimate in Theorem 7.7 gives

$$\sum_{h : C \to B} |\mu_h| W_p(h) |t|^{|h|} \leq M_n R_p(|t|)(C, B) < \infty.$$

Convolution is bounded by $M_n R_p(|t|)^2$, so regrouping by the complete book is legitimate. Each nonidentity book cancels by the lower-interval Möbius recurrence in one product and the upper-interval recurrence in the other. Both products are the identity. □

Corollary 7.9 (Boundary regularity and first cuts)**.** *Every $R_p(C, B)$ is a multiaffine polynomial in the sleep probabilities, increasing in p_i for $i \in B$ and decreasing for $i \notin B$. The inverse extends continuously throughout $[0,1]^V$ and satisfies $|R_p^{-1}(C, B)| \leq M_n R_p(C, B)$. Its determinant is positive there. For a prescribed intermediate support D, the total weight of books with no proper D-cut is $F(C, D) = R(C, D)/R(D, D)$, with $F(D, D) = 1$, and*

$$R(C, B) - \frac{R(C, D) R(D, B)}{R(D, D)} = \sum_{h : C \to B \text{ having no } D\text{-cut}} W_p(h) \geq 0.$$

Record: RN, §37, (J37)–(J39), (J48), (J70)–(J71); maintained audited results.

Proof. Summing free bits leaves a finite nonnegative mixture indexed by which neutral sites are used, with factors p_i on terminal sites and q_i on nonterminal sites. Each coordinate occurs at most once. The bound from Theorem 7.8 is uniform on the closed cube; bounded inverses force boundary nonsingularity. At the all-one corner R is diagonal with positive entries $1/\det(I - K_{CC})$, fixing the determinant sign. For the last identity choose the least D-cut, which exists by fiber closure under coordinate minimum. Any D-cut after concatenation contains this first cut, again by taking their coordinate minimum. Thus splitting there is bijective and its weights multiply. □

For completeness, RN (J52)–(J54) also prove separate concavity of $\log\det R_p$ and $\det R_p \geq \prod_{D \subseteq V} \det(I - K_{DD})^{-1}$ on the cube. This uses a substochastic one-coordinate release block; its proof is a *proof reference, not reproduced.* It asserts neither joint log-concavity nor a maximum at $p = 1/2$.

Theorem 7.10 (The greatest cut and its finite envelope)**.** *In iid paired stacks conditioned to have first bit one on C, the stable cuts have an almost surely finite greatest element. Its book $h : C \to B$ has probability $W_p(h) f_B$, so its endpoint law is $L_p(C, \cdot)$. In particular f_C is the probability of no nonzero cut. For an unconditioned finite all-active η, the greatest endpoint has law $P_\eta L_p$, rather than in general the physical law P_η.*

One arrow envelope bounds every cut for every binary input and every bit field. Drain $\mathbf{1}_V$ *without sleep and then pop all available head cycles in the remaining arrows. If* Δ *is the combined emission vector, then*

$$\Delta = \mathbf{1}_V + I(\Delta), \qquad m \le \mathbf{1}_V + I(m) \Longrightarrow m \le \Delta.$$

Descending buffer iteration from Δ *gives the greatest cut; ascending iteration gives the least. A prescribed full-drain arrow word* h *is this envelope with probability* $d\,\mathrm{wt}_{\mathrm{ar}}(h)$.

Record: RN, §37, (J82)–(J86); maintained audited proofs, reproduced.

Proof. The event specifying a remainder prefix $a : C \to D$ and its terminal ones has probability $W_p(a)$. Its shifted tails retain precisely the conditioned product law at D, including fresh terminal arrows. Tonelli therefore gives expected cut count $\sum_D R(C, D) < \infty$. Cuts are directed: their coordinate maximum has nonnegative residual input, and legally stabilizing it gives their join. Termination holds simultaneously after every finite shift, because an infinitely emitting site would send infinitely many particles along a positive path to the sink, contradicting the finite exit budget. Finitely many directed cuts have a greatest one. If g_D is the probability of no further cut, a specified prefix is greatest with probability $W_p(a)g_D$. Summing gives $Rg = \mathbf{1}$, hence $g = f$. The same argument after the least physical cut gives the general-input statement.

For the envelope, the cycle-popped suffix is the greatest balanced prefix: balanced prefixes are closed under coordinate maximum, and any nonempty further balanced prefix exposes a head cycle. Least action bounds the physical drain by any full-drain solution. If $u = \max(m, \Delta)$ with $m \le \mathbf{1} + I(m)$, then $u \le \mathbf{1} + I(u)$ and u already uses at least $|V|$ exits. Summing forces equality, leaving a balanced suffix beyond Δ, which its acyclic heads prohibit. Every queried bit index is at most Δ_i. Finally a prescribed unit-divergence word is the envelope exactly when its next heads are acyclic. Their probability is d by cycle inclusion-exclusion. This proves its law. The next-head forest is independent of the envelope; only the tails after that forest are iid. □

Remark 7.11 (Scoped order consequences). Pinning first bits and using a common envelope gives $p \le p' \Rightarrow f_C(p) \le f_C(p')$ and $C \subseteq C' \Rightarrow f_{C'}(p) \le f_C(p)$. Consequently the full-input stationary density $\pi_t(A)/\mu_t(A) = f_{A^c}(1-t)/d$ increases with A. Every positive-probability conditional occupation at i is at least t_i, strictly greater for interior parameters exactly when i lies on a positive directed cycle, including a self-loop. Neither FKG nor positive association of terminal occupations follows. On a loop-free directed simple cycle all cuts form a chain; splitting at the first nonzero cut gives $R^{-1} = I - \mathcal{A}$, where $\mathcal{A} \ge 0$, $\mathcal{A}\mathbf{1} = \mathbf{1} - f$ and $\rho(\mathcal{A}) < 1$. Thus the entire inverse is an M-matrix there, including loss sectors. These are proved results with *proof references, not reproduced.*

Record: RN, §37, (J90)–(J92b), (J107)–(J109); maintained audited results.

7.4 Finite inventories, recurrence, and rotors

Fix an integer arrow inventory J, including exits, with row lengths $m_i \ge 1$ and unit divergence $m - J^T\mathbf{1} = \mathbf{1}$; in matrix expressions J means its internal part. Periodize any local arrow orders and bit words, retaining all m_i labelled phases. At phase a_i the resident is $\beta_i(a_i)$ and the emission head is $a_i - \beta_i(a_i)$ modulo m_i. One arrival advances the phase and emits $1 + \beta_i(a_i) - \beta_i(a_i + 1)$ consecutive arrows. Set $D_m = \mathrm{diag}(m)$ and $M = D_m - J^T$.

Proposition 7.12 (Burning count and fixed-inventory ensembles)**.** *This finite unary Abelian network halts on every input, has critical group* $\mathbb{Z}^V/M\mathbb{Z}^V$*, and has exactly* $\det M$ *recurrent*

phases. A phase is recurrent exactly when addition of $\mathbf{1}_V$ *returns it to itself. These recurrence and group assertions use the standard Abelian-network theorems cited below.*

A head-zero compatible pair (w, C) *requires first bit one on* C *and last bit zero off* C*. If* k_i *ones are prescribed and*

$$N = \prod_i \frac{m_i!}{\prod_j J(i,j)!} \binom{m_i}{k_i}, \qquad d_J = \det(I - D_m^{-1} J),$$

there are N *compatible pairs and* Nd_J *recurrent compatible pairs. Their uniform recurrent law is invariant under additions followed by rotation of both rows by the emitted lengths. These maps commute, full addition is the identity, and single-site addition at* x *has inverse addition of* $\mathbf{1}_V - \mathbf{1}_x$*.*

Record: RN, §37, (J99)–(J102); maintained audited applications of Bond–Levine, *Abelian networks III*, Theorems 2.9, 3.3, 3.11, 3.18 and 5.5. External theorem proofs: proof reference, not reproduced.

Proof. The production matrix is $P = J^T D_m^{-1}$, and $Pm = m - \mathbf{1}$; hence $\rho(P) < 1$. The local phase cycles are irreducible and the total kernel is $D_m \mathbb{Z}^V$. Also $(I - P)D_m = M$ and $M\mathbf{1} = \mathbf{1}$. These check the hypotheses and conventions for the cited halting, critical-group, counting, and burning theorems.

Choose a uniform word and uniform labelled phase a. Put $C_i = \beta_i(a_i)$, $e_i = a_i - C_i$, and rotate both words by e_i. This gives a compatible pair. Each pair has exactly $\prod_i m_i$ preimages, even for words of smaller period. Rotation preserves recurrence, and each word has $\det M$ recurrent phases among $\prod_i m_i$ phases. The counts follow. Additions commute with these rotations and descend to the compatible-pair ensemble. On recurrence, the full burning addition is the identity, proving the inverse formula. □

An empty reset state is cyclic phase zero only if the last bit is zero. Otherwise it is a separate transient state entering phase one on its first arrival. Ignoring this distinction invalidates the boundary comparison with the physical process.

Theorem 7.13 (Recurrent rotor correspondence)**.** *For a fixed periodic word, release the residents of a cyclic buffer state into pure rotors at its current emission heads and drain them. The resulting head map* D_w *intertwines every addition:*

$$D_w(T_b^{\text{buffer}} a) = T_b^{\text{rotor}} D_w(a), \qquad b \geq 0.$$

It bijects recurrent states with rotor heads whose last-used arrow occurrences form a sink forest. Its inverse stabilizes $\mathbf{1}_V$ *in empty reset buffers with emission heads at that forest. After quotienting by simultaneous local rotations, this is a bijection between decorated complete drain books and recurrent compatible pairs, preserving the arrow inventory and every bit count.*

Record: RN, §37, (J106)–(J106b); maintained proof, independent AI audit recorded, reproduced using Proposition 7.12.

Proof. Count pending particles plus stored residents as rotor particles. A buffer step with old and new resident bits c_i, c_i' emits $1 + c_i - c_i' \leq 1 + c_i$ arrows, all legal rotor departures. Completing the final residents by drainage therefore gives exactly the same total emission vector as direct rotor routing. This proves intertwining. The two recurrent state spaces carry free transitive actions of the same critical group, so the map is a bijection there. The forest criterion also follows directly: last-departure times forbid a cycle, while an unconsumed suffix in a supposed full-period drain would be balanced and would force a cycle among its last arrows.

For the inverse let h be a rotor forest. Its pure drain of $\mathbf{1}_V$ uses exactly one period m. Every stable reset-buffer cut w for $2\mathbf{1}_V$ satisfies $w \geq \mathbf{1}_V + I(w)$, so rotor least action gives $w \geq m$. Periodicity and unit divergence identify its fixed cuts with m plus the fixed cuts for $\mathbf{1}_V$. All shifted arrival indices are positive, so the reset convention introduces no exception. The output from the first full addition is therefore recurrent, and draining its residents returns h. If its emitted vector is u and the completing drain is r, then $u + r = m$. Hence the forward rotation by u and inverse rotation by r cancel. This proves the word bijection without period multiplicities. □

Two useful consequences retain their exact conditioning. Let $N_{\mathrm{ar}}(J) = \prod_i m_i! / \prod_j J(i,j)!$. The ordinary complete-drain inventory has probability $D_K(J) = N_{\mathrm{ar}}(J) K^J d_J$ and

$$\pi_p(B) = \sum_J D_K(J) \pi_{J,p}(B),$$

where $\pi_{J,p}$ is the resident law of recurrent compatible words with iid bit weights. Conditional also on k, the resident law $\nu_{J,k}$ has increasing density relative to product Bernoulli $\theta_i = k_i/m_i$. Thus every positive-probability outside conditioning has occupation probability at least θ_i. Deterministic coordinates are restricted to their product-prior face. The proof uses reversal to a no-cut event and the elementary entire-row inequality $n - \beta(n) \leq n + 1 - \beta'(n+1)$ for $n > 0$; the other rows and arrows are shared. The rotor bijection identifies this conditional law with physical full-input ARW on its actual auxiliary complete drain.

Record: RN, §37, (J103)–(J106b); maintained audited results; detailed mixture and fixed-count proof references, not reproduced.

For comparison, deterministic periodic bits 10 with two outgoing arrow occurrences per site give $M = 2I - A^T$ and greatest cut $2\max\{z \in \mathbb{Z}^V_{\geq 0} : Mz \leq \eta\}$. Its endpoint is the unique z-superstable representative and uniquely minimizes $\|M^{-1}g\|_2^2$ in the nonnegative integer equivalence class. This is the recorded bridge to Guzman–Klivans energy minimization, not a replacement by subset-only superstability on a general digraph.

Record: RN, §37, (J97)–(J98); maintained audited bridge; Guzman–Klivans, Definition 4.3 and Theorem 4.6, credited external background; proof reference, not reproduced.

7.5 Full-minus-one identities and release interpolation

Proposition 7.14 (Two exact full-minus-one reductions)**.** *Fix x and put $S = V \setminus \{x\}$. In a complete decorated drain book, let L, U be the physical reset outputs from S, V, retaining their heads and transient reset states. Uniformly over all such books with fixed (J,k),*

$$\mathbb{E}[U_v - L_v] = \mathbb{E}[(T_V L)_v - L_v].$$

Moreover $T_x L = U = T_x T_V L$, and $T_V L$ is recurrent. This does not order the resident sets.

There is also a product-law identity. Let $h_x = 0$, $h_i = \mathbb{P}_i(\partial \text{ before } x)$ for $i \neq x$, $\sigma_x = K(x,\partial) + \sum_i K(x,i) h_i > 0$, and

$$Z_x = \prod_{i \neq x} (p_i + q_i h_i), \qquad \theta_i = \frac{p_i}{p_i + q_i h_i}, \qquad D \sim \mathrm{Bernoulli}(\theta) \text{ on } S.$$

Writing $\gamma_v = R_p^{-1} \mathbf{1}_{v \in \cdot}$, one has

$$q_v(V) - q_v(S) = \frac{p_x Z_x}{d} \mathbb{E}[\gamma_v(D + x) - \sigma_x \gamma_v(D)], \qquad \mathbb{E}[f(D + x)] = \sigma_x \mathbb{E}[f(D)].$$

Record: RN, §37, (J111)–(J112); maintained proved identities; recurrent-completion proof reference, not reproduced.

Proof of the product-law identity. In the squarefree addition algebra, IDLA with only x vacant gives $z_S a(z) = a(h) z_S + [a(\mathbf{1}) - a(h)] z_V$. The constant-term formula in Theorem 7.7 consequently gives $(Q_V - Q_S)(B) = \mu_p(B) w_{B^c}^{p/q}(h)/d$. The factors are $w_i(h) = h_i$ for $i \neq x$ and $w_x(h) = -(p_x/q_x)\sigma_x$. Pairing $B = D + x$ and $B = D$ leaves the common positive weight $(p_x/d)\mu_{p,S}(D) h_{S \setminus D}$. Its normalizer is $p_x Z_x/d$. Multiply by $R^{-1}\mathbf{1}_{v\in\cdot}$, or by f and use the zero total mass, to obtain both identities. □

Theorem 7.15 (One-coordinate release resolvent)**.** *Fix i and all other sleep probabilities. Let P_0 mean binary stabilization with $p_i = 1$ only; the subscript does not set other odds zero. Let W fix supports not containing i and, on an occupied support, remove its resident, force a K-jump, and stabilize at P_0's parameters. For $s = (1 - p_i)/p_i$, $z = s/(1 + s)$,*

$$P_s = P_0[I + s(I - W)]^{-1} = (1 - z) \sum_{k \geq 0} z^k P_0 W^k.$$

If V_i deletes i, then $W A_i = A_i$ and $W^k \to A_i V_i$. For an add-one pair $C, C + x$ with $i \in C$ and $x \notin C$, set $\delta = (\delta_{C+x} - \delta_C) P_0$. Either of the following would suffice for its target comparison at every $s \geq 0$:

$$\delta \sum_{k=0}^{M} W^k \mathbf{1}_{v\in\cdot} \geq 0 \qquad \text{for every integer } M \geq 0;$$

$$\delta H_k \mathbf{1}_{v\in\cdot} \geq 0 \qquad \text{for every } 0 \leq k \leq m = 2^{|V|-1},$$

where $D(s) = \det(I + s(I - W)) = \sum d_k s^k$ and $H_k = [s^k] \operatorname{adj}(I + s(I - W))$. These two criteria remain open in general.

Record: RN, §37, (J113)–(J113c); maintained exact identities and open sign criteria.

Proof. The inverse basis is affine in the odds at i and satisfies $E_s = [I + s(I - W)]E_0$; invert and expand the stochastic resolvent. On an occupied row $\delta_B W A_i = \delta_{B-i} J_i A_i = \delta_B J_i = \delta_B A_i$, using $p_i = 1$ and the quadratic relation; on vacant rows the identity is immediate. The occupied block is transient, since a prescribed positive path to the sink can release the resident without return. Harmonicity and the identity boundary values on vacant states identify the limit as $A_i V_i$. Abel summation transforms the geometrically weighted comparison into $(1 - z)^2 \sum_{M \geq 0} z^M \delta \sum_{k=0}^{M} W^k \mathbf{1}_v$. The sums grow at most linearly, so rearrangement is valid. The second criterion follows on multiplying the resolvent by its determinant. The row-Laplacian forest formula gives $H_k \geq 0$ and $H_k \mathbf{1} = d_k \mathbf{1}$, but a signed δ still requires a proof of its coefficient signs. □

The forest formula used here is the standard result credited in RN to Chebotarev–Agaev, §§2,4–5, Theorem 3′ and Proposition 4. Its forests are on sleeping-support states, not on V.

Theorem 7.16 (Two-site prefix and forest-degree inequalities)**.** *Both criteria in Theorem 7.15 hold on every accessible two-site kernel, with arbitrary interior sleep at the other site, either common root, both targets, and every prefix horizon. Self-loops are allowed.*

Record: RN, §37, (J114)–(J115b); maintained general two-site proofs, reproduced.

Proof. First remove self-loops. Write $i = 0$, $K(0, 1) = a$, $K(1, 0) = b$, $ab < 1$, $p_1 = p$, and $\rho = ab(1 - p)$, $\tau = p/(1 - \rho)$, $\sigma = ab\tau$. The only common-root cover is $\{0\} \to \{0, 1\}$, and $\delta = \tau(\delta_{\{0,1\}} - \delta_{\{0\}})$. In order $\varnothing, \{0\}, \{1\}, \{0, 1\}$, the nontrivial release rows are

$$W(\{0\}, \cdot) = (1 - \rho - ap, \rho, ap, 0),$$

$$W(\{0,1\},\cdot) = \big(a(1-b)^2(1-p),\ ab[1-\tau+(1-b)(1-p)],\ 1-a+a(1-b)p,\ \sigma\big).$$

These follow by the forced first jump and the one-particle continuation after waking site 1. Vacant-root rows are absorbing. Put $A_M(v) = \delta \sum_{k=0}^{M} W^k \mathbf{1}_v$. The two transient-state reward equations are triangular: for initial occupied-root state B, its reward series satisfies $u_B = h_B + z\sum_{D\ni 0} W(B,D)u_D + z\sum_{D\not\ni 0} W(B,D)h_D/(1-z)$. Substituting the displayed rows and multiplying their difference by $\tau/(1-z)$ gives

$$\sum_{M\ge 0} A_M(1) z^M = \frac{\tau(1-az)(1-abz)}{(1-z)^2(1-\rho z)(1-\sigma z)},$$
$$\sum_{M\ge 0} A_M(0) z^M = \frac{\tau ab[1-b(1-p)]z}{(1-z)(1-\rho z)(1-\sigma z)}.$$

Every coefficient is nonnegative: use $(1-cz)/(1-z) = 1+(1-c)\sum_{j\ge 1} z^j$ for $c = a, ab$ in the first line; all remaining factors are positive geometric series.

For the forest condition put $e = 1-ab(1-p)$ and $f = 1-ab\tau$. The determinant is $(1+es)(1+fs)$. Direct substitution into the same resolvent, or the adjugate recursion, gives

k	$\delta H_k \mathbf{1}_0$	$\delta H_k \mathbf{1}_1$
0	0	τ
1	$\tau ab[1-b(1-p)]$	$\tau(2-a-ab)$
2	0	$\tau(1-a)ef$

and all entries are nonnegative. A nonroot self-loop ℓ replaces p by $p/[1-(1-p)\ell]$ and normalizes non-self routing. A root self-loop α leaves P_0 unchanged and gives $W^{\mathrm{old}} = \alpha I + (1-\alpha)W^{\mathrm{new}}$. Expansion and summation by parts give

$$A_M^{\mathrm{old}}(v) = \sum_{j=0}^{M} \binom{M+1}{j+1} (1-\alpha)^j \alpha^{M-j} A_j^{\mathrm{new}}(v).$$

The coefficients are nonnegative; forest degree k is multiplied by $(1-\alpha)^k$. Relabelling proves the other-root case. □

The general occupied/vacant decomposition is worth retaining. In that order reversed to vacant first, write $W = \left[\begin{smallmatrix} I & 0 \\ B & Q \end{smallmatrix}\right]$, $M = I - Q$, and $F_k = [s^k]\,\mathrm{adj}(I+sM)$, with $F_{-1} = F_m = 0$. Then

$$H_k = \begin{pmatrix} d_k I & 0 \\ F_{k-1}B & F_k \end{pmatrix}, \qquad F_k = d_k I - MF_{k-1}.$$

The absorption term cannot be discarded. If $T = M^{-1}B$, then for an off-root target indicator h, $F_k h + F_{k-1}Bh = d_k Th + F_k(I-T)h$. After removal of the root self-loop, killed additions U_x, first-hit blocks V_x, and trace additions $\widetilde{A}_x$ satisfy

$$\widetilde{A}_x = U_x + V_x T, \qquad \widetilde{A}_x T = T\widetilde{A}_x,$$
$$B = K(i,\partial)I + \sum_{j\ne i} K(i,j)U_j, \qquad Q = \sum_{j\ne i} K(i,j)V_j,$$
$$B = (I-Q)T.$$

The U_x commute with B; the two families need not commute with each other. If i is on no internal cycle after self-loop removal, $Q = 0$ and $H_k = \binom{m-1}{k} I + \binom{m-1}{k-1} W$. This reduces the criterion to endpoint comparisons on the smaller trace graph, and proves it on DAGs. The remaining cyclic comparison is open.

Record: RN, §37, (J115)–(J115d); maintained proved reductions; block-inversion proof reference, not reproduced.

7.6 Decisive obstructions and the saved audit boundary

Proposition 7.17 (Rotation averages fail, with an exact repair family)**.** *For every $N \geq 2$, put $L = N + 1$. On a directed N-cycle let row 0 contain one forward arrow and N sink arrows; row $i > 0$ contains $i + 1$ forward arrows. Give non-target rows zero bits, and put one target one bit cyclically d positions after its forward arrow, $1 \leq d \leq N$. Uniform paired rotations give, for nonempty binary input of size r,*

$$q_{\mathrm{rot}}(r) = \frac{1 + \mathbf{1}_{r=d}}{L}.$$

Thus the cover $d \to d + 1$ loses $1/L$ when $d < N$. If instead all arrow and bit orders are averaged independently with k target ones, then

$$q_{\mathrm{all}}(r) = \frac{k}{L} + \frac{k(L-k)}{L^2(L-1)}, \qquad 1 \leq r \leq N,$$

independent of r. This proves a repair for this family only.

Record: RN, §37, (J110)–(J110a); maintained exact obstruction and proved family.

Proof. All initial particles reach 0 and its unique forward arrow allows at most one further circuit. If that arrow is at t, the terminal bit is β_r when $t > r - \beta_r$ and β_{r+1} otherwise. Under paired rotations a one at r survives iff $r \leq d$; a one at $r + 1$ is reached iff $r \geq d$. Counting the L rotations proves the first formula. Averaging over t first gives $\beta_r + (r - \beta_r)(\beta_{r+1} - \beta_r)/L$. Exchangeability of a uniform k-one row gives the second formula. □

Other failures locate the remaining cancellation. General stable-cut lattices can be M_3 or N_5; even one-particle books can have Möbius number greater than one. A fixed book can have the wrong target coefficient, and grouping its whole local inventory need not repair it. Fixing an arrow envelope can give a negative physical increment. Even an entire recurrent rotor class can have full-minus-one increment $-1/45$ (RN (J111)). None is a counterexample under iid ARW averaging.

Record: RN, §37, (J41)–(J42), (J55), (J61), (J64)–(J64b), (J89a), (J111); proved explicit obstructions, exact checks supplementary.

The correct general target expression is $Q_\eta[\Xi_x, G_v]f$, where $\Xi_x = R^{-1}A_x^p R$ and $G_v = R^{-1}O_v R$, with $O_v(C, C) = \mathbf{1}_{v \in C}$. The column f is essential. The commutator column has both signs whenever x can reach v: its empty-input evaluation is positive, while its strictly positive full-input invariant row averages to zero. Thus entrywise positivity cannot prove the desired physical averages.

Record: RN, §37, (J58)–(J60), (J109); maintained audited identities and obstruction.

Remark 7.18 (Exact numerator obstruction and unaudited routing reports). For the release-prefix generating function write

$$\sum_{M \geq 0} A_M(v) z^M = \frac{N(z)}{(1-z)^2 \det(I - zQ)}, \quad N(z) = \delta \operatorname{adj}(I - zQ)[(1-z)h + zBh_{\mathrm{vac}}].$$

The stronger shortcut of dropping determinant compensation has an exact obstruction, independently rerun by the main task on 2026-11-08, after restricting to reachable states:

$$K = \begin{pmatrix} 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 1/2 & 0 & 0 & 0 & 1/2 \\ 1/2 & 0 & 1/2 & 0 & 0 \end{pmatrix}, \quad p = (1, 1/5, 4/5, 1/2),$$

where the last column is the sink, $i = 0$, $C = \{0, 3\}$, $x = 1$, $v = 2$. On the seven reachable states, the independently rerun certificate gives

$$A_0 = \frac{48869}{121282}, \qquad A_1 = \frac{9220242269537}{17651188228800}, \qquad [z]\frac{N(z)}{(1-z)^2} = -\frac{13280929}{145538400}.$$

The main task also verified the general-block boundary identities against independent raw ARW calculations with 44 states in each of two systems. The saved nonnegative prefixes through horizon 24 remain only finite checks; they do not prove the general integrated criterion, which this obstruction does not refute.

The same saved boundary reports a paired-permutation connection to iid bit-dependent routing: $w_i(b, j) = p_i^b q_i^{1-b} K_i^b(j)$ and $\overline{K}_i = q_i K_i^0 + p_i K_i^1$, with accessible sink. It reports halting, the quadratic identity $A_i^2 = A_i(p_i J_i^1 + q_i J_i^0)$, and the complete inverse based on $(z_i - q_i J_i^0)/p_i$. Universal paired-pool positivity would imply occupation comparisons for these correlated models; the converse is not justified. Reported scoped proofs cover paired pools on DAGs and binary comparisons in loop-free two-site correlated iid models. These claims are retained as *agent-reported arguments awaiting independent audit*, not promoted here to established theorems.

Record: RN, §37, (J116) and "Paired-pool bridge"; numerator certificate independently rerun by the main task, 2026-11-08; correlated-routing reports still await independent audit.

The general all-active single-target add-one conjecture remains open. General envelope and paired-pool coefficient signs, common-root release prefix signs, and release forest-degree signs remain open as well. The exact-pair remainders of RN (127f) must remain in any continuation; the saved obstruction to omitting them is unchanged.

8 Further saved certificates and boundary interpolation

This section makes two developments preserved in the verifier files readable alongside the main notes. The first is an exact symbolic theorem on every directed three-cycle. The second is a general block identity, with a separate raw-ARW finite check. Neither supplies the missing entrance-law comparison for general graphs.

8.1 The killed resident-return kernel on a three-cycle

Consider $0 \to 1 \to 2 \to 0$, with forward probabilities $a, b, c \in [0,1]$, the remaining mass going to the sink, and $\kappa = abc < 1$. Set $p_0 = 1$ and $(p_1, p_2) = (p, q) \in (0,1)^2$. At the boundary $p_0 = 1$, a lone particle sleeps and surplus particles still jump. From support $\{0\} \cup B$, release the resident at 0 by a forced jump and stabilize. Let

$$Q(B, C) = \mathbb{P}\{\text{final support} = \{0\} \cup C\}, \qquad B, C \subseteq \{1, 2\}.$$

Missing row mass is killing, namely ending vacant at 0.

Theorem 8.1 (Three-cycle killed-return order; computer-assisted)**.** *For every nonnegative increasing g on $2^{\{1,2\}}$, Qg is increasing. The statement holds throughout the parameter domain above, not merely on a finite grid.*

Proof. Particle conservation gives $Q(B, C) = 0$ if $|C| > |B|$. The row sum is a times the probability that the first passage through the path $1, 2$ returns a positive number of particles to 0. This probability increases with B: at a visited path site with incoming count m and old bit b, couple its new bit s across the two systems, giving emitted count $m + b - s$. Ordered incoming counts and old bits then give ordered emitted counts, with the zero-input case handled by leaving the resident unchanged. Shared ordered Bernoulli routing trials preserve this comparison. Once the root queue is positive, it cannot lose its last particle before sleeping at 0, so subsequent root emissions do not change this survival criterion.

The empty row is supported at $\varnothing$. Besides row sums, the only nontrivial upper-set comparisons are the six differences from $B = \{1\}$ or $\{2\}$ to $\{1, 2\}$ for

$$U_1 = \{C : 1 \in C\}, \quad U_2 = \{C : 2 \in C\}, \quad U_+ = \{C : C \neq \varnothing\}.$$

The upper set $\{\{1, 2\}\}$ has zero mass from either singleton. Write

$$E = 1 - \kappa(1 - p)(1 - q), \qquad F = 1 - \kappa + \kappa pq.$$

Both are positive. For $w_{U_1} = p$, $w_{U_2} = q$, $w_{U_+} = 1$, define

$$H_{B,U} = \frac{EF}{\kappa w_U}[Q(\{1, 2\}, U) - Q(B, U)].$$

When $\kappa > 0$, these are polynomials with the following exact Bernstein certificates on $[0,1]^5$, in variables (a, b, c, p, q).

B	U	Multidegree	Positive	Zero	Least positive
$\{1\}$	U_1	$(2, 2, 3, 2, 3)$	404	28	1/9
$\{1\}$	U_2	$(1, 2, 2, 2, 1)$	98	10	1/4
$\{1\}$	U_+	$(2, 2, 3, 3, 3)$	495	81	1/27
$\{2\}$	U_1	$(2, 2, 3, 2, 3)$	424	8	1/6
$\{2\}$	U_2	$(1, 2, 2, 2, 1)$	98	10	1/4
$\{2\}$	U_+	$(2, 2, 3, 3, 3)$	515	61	1/18

For reproducibility, here is the finite construction whose identities the accompanying verifier checks. Use low queue/background states

$$(1,00),(1,10),(1,01),(1,11)$$

and high states $(2,00),(2,10),(2,01),(3,00)$. One root emission goes to the sink with probability $1-a$, or sends one particle through the path. At a path site the last bit is the old bit if no particle arrives; otherwise it is a fresh Bernoulli bit, and each of the $m+b-s$ emitted particles independently advances with the site's forward probability. This defines an 8×8 substochastic transition matrix T. Eliminating the four high states gives exactly

$$Q = T_{LL} + T_{LH}(I - T_{HH})^{-1}T_{HL}.$$

Sink accessibility makes the high-state elimination finite, with no time cutoff. The script solves these rational equations, forms all six $H_{B,U}$, computes every Bernstein coefficient exactly, and reconstructs each polynomial from its expansion. There are $2,232$ coefficients, all nonnegative. Every nonnegative increasing function is a nonnegative combination of upper-set indicators, so these inequalities prove the assertion. If $\kappa = 0$, $Q = 0$ directly. □

Record: `verify_cycle_return.py`; exact symbolic reconstruction with SymPy 1.14.0. A separate 28-transient-state raw-instruction calculation checks all 16 kernel entries at one rational parameter vector.

The whole-row theorem does not justify conditioning on an identical first path passage. At $a = 1$, $b = c = p = q = 1/2$, a common first-pass cylinder of probability $1/32$ leaves the lower system stopped with outside support $\{1\}$ and the upper system with queue two and outside support $\{1\}$. Its continuation retains site 1 with probability $107/195 < 1$. This is an obstruction to that conditional coupling, not to Theorem 8.1.

8.2 Positive blocks at an all-sleep boundary

Let K be a finite sink-accessible kernel, fix $0 < p_j < 1$ for $j \neq i$, and let $x \neq i$. Matrices act on row laws from the right. At $p_i = 1$, order sleeping supports by root-vacant and root-present blocks, each indexed by subsets of $V \setminus \{i\}$. Write the physical additions, forced jump, and resident release as

$$A_i^0 = \begin{pmatrix} 0 & I \\ 0 & T \end{pmatrix}, \quad A_x^0 = \begin{pmatrix} U & V \\ 0 & \widetilde{A} \end{pmatrix}, \quad J_i^0 = \begin{pmatrix} B & Q \\ 0 & T \end{pmatrix}, \quad W_i^0 = \begin{pmatrix} I & 0 \\ B & Q \end{pmatrix}.$$

The letter V in these block formulas is a matrix, not the vertex set. All displayed blocks are nonnegative; Q is the killed root-return block. Put $M = I - Q$. Accessible loss makes M invertible. For finite root odds $s = (1-p_i)/p_i \geq 0$, set

$$R = (I + sM)^{-1}, \qquad L = sRB.$$

The resolvent is nonnegative, since $R = (1+s)^{-1}\sum_{k\geq 0}(s/(1+s))^k Q^k$.

Theorem 8.2 (Positive boundary interpolation)**.** *The boundary blocks obey*

$$B = MT, \quad \widetilde{A} = U + VT, \quad \widetilde{A}T = T\widetilde{A}, \quad UB = BU, \quad M\widetilde{A} = UM + BV.$$

At root odds s, with all other parameters fixed, the physical addition matrices are

$$A_i(s) = \begin{pmatrix} L & R \\ TL & TR \end{pmatrix},$$

$$A_x(s) = \begin{pmatrix} U + VL & VR \\ VTL & U + VTR \end{pmatrix}.$$

Every block is nonnegative. The rectangular matrix $[L, R]$ *is stochastic and satisfies*

$$\widetilde{A}[L, R] = [L, R]A_x(s), \qquad T[L, R] = [L, R]A_i(s).$$

In particular,

$$\widetilde{A}R = RU + R(VT + sBV)R.$$

Proof. Commutativity of A_x^0 with A_i^0 gives $\widetilde{A} = U + VT$ and $\widetilde{A}T = T\widetilde{A}$. Commutativity also writes the quadratic relation as $(A_i^0)^2 = J_i^0 A_i^0$; comparing upper-right blocks gives $B + QT = T$. Commutativity of A_x^0 with J_i^0 gives $UB = BU$ and $UQ + VT = BV + Q\widetilde{A}$, equivalently $M\widetilde{A} = UM + BV$. Also $B\mathbf{1} = (I - Q)\mathbf{1}$ because the release kernel is stochastic; hence $(L + R)\mathbf{1} = \mathbf{1}$.

The one-coordinate release formula gives

$$S = I + s(I - W_i^0) = \begin{pmatrix} I & 0 \\ -sB & I + sM \end{pmatrix}, \qquad S^{-1} = \begin{pmatrix} I & 0 \\ L & R \end{pmatrix}, \qquad A_z(s) = SA_z^0 S^{-1}.$$

Multiplying these blocks and using the preceding identities gives the two stated addition matrices. For example their lower-right x block is $[(I+sM)\widetilde{A}-sBV]R = (U+VT+sUM)R = U+VTR$. The lower-left block is $-sBU+(U(I+sM)+VT)L = -sBU+sUB+VTL = VTL$, using $(I+sM)L = sB$ and $UB = BU$. The bottom row of A_x^0 is invariant under the boundary chain, so $\widetilde{A}[0, I]S^{-1} = [0, I]S^{-1}A_x(s)$; this is the first intertwining, and the second follows identically from A_i^0. Finally multiply $(I + sM)\widetilde{A} = U(I + sM) + VT + sBV$ by R on both sides to obtain the displayed positive resolvent identity. □

Record: RN §37 (J113) for the release conjugation; `verify_sleep_transport.py`, `boundary_interpolation_audit`, for the saved block identities and independent raw-ARW check. The proof above assembles those identities and received a separate AI review during compilation.

What sign remains. If a common all-active input contains i, its boundary law is $[0, \nu]$. Adding at x changes root occupation at odds s by

$$\nu(\widetilde{A} - I)R\mathbf{1} = s\nu R(BVR - VTRB)\mathbf{1}.$$

The forcing on the right is still signed. For an outside target payoff h, put $G = L + R = T + R(I - T)$; its gain is

$$\nu(\widetilde{A} - I)Gh = \nu T(\widetilde{A} - I)h + \nu(\widetilde{A} - I)R(I - T)h.$$

No order follows just from nonnegative blocks. The tempting partner $\widetilde{A} - TV = M^{-1}UM$ can have negative entries, and T need not be invertible. The maintained rational check has partner entry $-19/700$, singular T, and physical gain $355/5157$ for its specified comparison. It checks 256 addition entries and 128 binary-stabilization entries against two independent 44-transient-state raw calculations. Self-loops are permitted; removing a root loop of probability ℓ rescales its effective odds to $(1 - \ell)s$ and renormalizes the remaining root routing row by $1/(1 - \ell)$.

The correlated-routing identities saved at RN (J116), including their restricted two-site conclusion, are recorded in Section 7.18. They concern a modified instruction law and do not extend this theorem to arbitrary correlations or prove the original general occupation comparison.

Part III
Related probability and algebra

9 Adjacent completed results and their limitations

This section collects the adjacent results maintained in the AI research records, principally `response-renewal.md`, Sections 30 and 32–36. The general all-active single-site monotonicity conjecture remains open. Statements below concern the indicated auxiliary law, input distribution, or algebra, and retain that scope.

Record: checkpoint, "Discipline and evidence"; evidence index, opening source policy; RN, introduction; AI-generated research records.

The result guide is:

1. Loss markers, determinantal mixtures, and scoped Rayleigh properties: Section 9.1.
2. Sure filling, possible filling, and analytic cancellation faces: Section 9.2.
3. Harmonic current tilts and completed particle paths: Section 9.3.
4. Addition algebras, polytopes, and full-filling input inequalities: Section 9.4.
5. Poissonized IDLA, exponential capacities, and dependence: Section 9.5.
6. Complementary-sleep duality and its target obstruction: Section 9.6.
7. Decisive failures of stronger conclusions: Section 9.7.

Common setting and imported local inputs. Let V be finite and nonempty, $n = |V|$, and let K be a jump kernel with absorbing sink s accessible from every vertex. Write $k_i = K(i, s)$, $p_i \in (0, 1)$, $q_i = 1 - p_i$, and $r_i = q_i/p_i$. For scalar vectors, a set subscript denotes a product, for example $r_A = \prod_{i\in A} r_i$; likewise $z^A = \prod_{i\in A} z_i$ and $z^\eta = \prod_i z_i^{\eta(i)}$ for a count vector η. In count vectors e_i and δ_i denote one particle at i; for support-indexed matrices δ_A is the point-mass row at A. Put

$$d_A = \det(I - K_{AA}), \qquad Z_A = d_A \sum_{C\subseteq A} \frac{r_{A\setminus C}}{d_C}, \qquad d_\varnothing = Z_\varnothing = 1.$$

All d_A are positive. Ordinary IDLA parks a particle at its first available capacity-one site, or loses it at s; particles are processed until every input has parked or exited. Its addition matrices are D_i. Let $P_r(A, B)$ be binary all-active ARW stabilization and A_i^r active addition to a sleeping support. We use the maintained complete-block inverse $E_r = P_r^{-1}$, surplus-IDLA reduction, and $P_r A_i^r = D_i P_r$. In particular the full-input stationary law is $\pi_r = \delta_V P_r$, with $\pi_r(V) = 1/Z_V$. The inverse is proved in Theorem 7.1; surplus reduction is described in Section 1. They are inputs to this section.

Record: CB, §§2–5; proved imported AI notes; RN, §§16–17, 33 (X11), 34 (A19).

9.1 Loss markers and the precise Rayleigh scope

Fix roots $Q \subseteq V$. The creator sampler chooses $C \subseteq Q$ with probability $d_Q r_{Q\setminus C}/(Z_Q d_C)$, a balanced prefix α supported in C with conditional weight $d_C w(\alpha)$, and independently an exit forest F on Q with probability $w(F)/d_Q$. A balanced prefix is a finite tuple of outgoing-arrow words with equal incoming and outgoing counts at every vertex; w is the product of its jump probabilities. The forest has one outgoing arrow per vertex and no internal directed cycle; exits

outside Q retain their labels. Put $L = \mathrm{range}(\alpha)$, $D = \{i : F_i = s\}$, and $E = \{i \notin C : F_i = s\}$. These are reserved-arrow markers, not physical terminal occupations. Here strongly Rayleigh means that the generating polynomial is real stable: it has no zero when every variable has positive imaginary part. The sampler's normalization is an earlier proved input.

Record: RN, §17 and §19 (63); proved sampler; proof reference, not reproduced.

Proposition 9.1 (Determinantal mixture)**.** *Set $G = (I - K_{QQ})^{-1}$ and $H = G\,\mathrm{diag}(k_i : i \in Q)$. For every $S \subseteq Q$,*

$$\mathbb{P}(S \subseteq D) = k_S d_{Q\setminus S}/d_Q = \det H_{SS}.$$

Conditional on L, E is D independently thinned with probabilities $t_i = q_i \mathbf{1}_{\{i \notin L\}}$. Thus

$$\mathbb{E}[z^E \mid L] = \det(I + H\,\mathrm{diag}(t_i(z_i - 1))).$$

If K_{QQ} is reversible with positive reversible weights, this conditional law is strongly Rayleigh. Without reversibility only the determinantal identity is asserted. With $g_A = p_A Z_A/d_A$,

$$\mathbb{P}(L = T) = g_Q^{-1} \sum_{A \subseteq T} (-1)^{|T\setminus A|} g_A.$$

Proof. Prescribing the sink arrows at S leaves an exit forest on $Q \setminus S$, of total weight $d_{Q\setminus S}$. The principal-minor identity for an inverse gives the displayed determinant. The joint weight of (C, α) is $d_Q r_{Q\setminus C} w(\alpha)\mathbf{1}_{\{L\subseteq C\}}/Z_Q$. Consequently, conditional on α, every site outside L belongs to C independently with probability p_i; this law depends only on L. Independence of F proves thinning, and expansion by principal minors proves the generating formula.

For reversible weights m_i, the matrix $B = \mathrm{diag}(\sqrt{m})(I - K_{QQ})\,\mathrm{diag}(1/\sqrt{m})$ is symmetric positive definite and $B \geq \mathrm{diag}(k)$. Hence $\widehat{H} = \mathrm{diag}(\sqrt{k})B^{-1}\,\mathrm{diag}(\sqrt{k})$ is a positive contraction with the same principal minors as H. Thinning preserves this representation; the external determinantal stability theorem applies. Zero k_i follow by continuity. Finally the marginal prefix weight is $p_L w(\alpha)/g_Q$, and its sum over prefixes contained in A is g_A/g_Q. Boolean inversion gives the range law. □

Record: RN, §30 (M1)–(M4); proved derivation. External inputs: directed forest normalization; Borcea–Brändén–Liggett, Proposition 3.5, determinantal stability.

Theorem 9.2 (Scoped negative dependence)**.** *For $F(z) = \mathbb{E}[z^E]$ put $\Delta_{ij}F = F_iF_j - FF_{ij}$. On every root domain, $\Delta_{ij}F(z) \geq 0$ when $0 < z_h \leq 1$ for all h. It holds for all positive fields if every strongly connected component of K_{QQ} either has at most three vertices or is a directed simple cycle after self-loop removal. Different components' marker vectors are independent. For unions of cycle components the marker law is negatively associated after arbitrary positive fields and arbitrary positive-probability partial assignments, including separately labeled exit markers. The latter assertion is not established here for general three-vertex components.*

Proof status. *Proof reference, not reproduced* for the arbitrary-domain one-sided fields and three-vertex cases: RN's Boolean logarithmic coefficient argument and the nonnegative polynomial identity (M11) are essential. Principal determinants, Z_A, and marker inclusion moments factor over components, giving the independence assertion. Here is the complete additional cycle argument. After removing a self-loop at i, replace $K(i, j)$ by $K(i, j)/(1-K(i, i))$ for $j \neq i$ and r_i by $(1 - K(i, i))r_i$; this preserves Z_A and r_ik_i. For a simple cycle put

$$\theta = \prod_i K(i, i + 1), \quad \beta = \theta(1 - p_Q), \quad u_i = q_ik_i.$$

Proper induced subgraphs are acyclic, so their Z_A equal $\prod_{i\in A}(1+r_i)$, whereas $Z_Q = \prod_{i\in Q}(1+r_i)(1-\beta)$. The inclusion moments $\mathbb{P}(S \subseteq E) = (rk)_S Z_{Q\setminus S}/Z_Q$ give

$$F(z) = \frac{\prod_i (1-u_i+u_i z_i) - \beta}{1-\beta}, \qquad \Delta_{ij} F = \frac{\beta u_i u_j}{(1-\beta)^2} \prod_{h\neq i,j} (1-u_h+u_h z_h).$$

For labels, replace a root factor by $1 - \sum_o u_{i,o} + \sum_o u_{i,o} z_{i,o}$, where $u_{i,o} = q_i K(i,o)$. This is a categorical product law μ with mass β removed at its minimum atom 0 and then normalized; feasibility follows from $\mu(0) \geq \theta > \beta$ when $\theta > 0$, and is immediate when $\theta = 0$. For disjoint-coordinate increasing functions f, g, direct expansion gives

$$\mathrm{Cov}_{(\mu-\beta\delta_0)/(1-\beta)}(f,g) = \frac{\mathrm{Cov}_\mu(f,g)}{1-\beta} - \frac{\beta(\mu f - f(0))(\mu g - g(0))}{(1-\beta)^2} \leq 0.$$

A categorical coordinate block is negatively associated: the covariance for disjoint groups of labels is the negative product of their respective nonnegative mean increments from 0. Conditional covariance inducts over independent roots. Fields preserve the categorical product and rescale the subtracted mass. Conditioning on a one removes that mass; conditioning only on zeros leaves the same feasible subtraction. This proves the stated conditional negative association.

Record: RN, §19 (64)–(66), §30 "Positive external fields", (M11)–(M13); proved, scoped as stated; general positive fields in larger noncycle components open.

Remark 9.3 (Companion and loop identities). For disjoint $A, B \subseteq Q$, with $a_i = r_i k_i$,

$$\mathbb{P}(A \subseteq E, B \subseteq D) = \frac{a_A Z_{Q\setminus A}}{Z_Q} k_B \frac{d_{Q\setminus(A\cup B)}}{d_{Q\setminus A}},$$

and for $i \neq j$, $\mathrm{Cov}(\mathbf{1}_{i\in E}, \mathbf{1}_{j\in D}) = -\mathbb{P}(i \notin C) k_i k_j G_{ij} G_{ji} \leq 0$. If $w_T \geq 0$ are the rejected-loop intensities of RN Section 19 (with $w_\varnothing = 0$),

$$\mathbb{P}(S \subseteq E) = q_S k_S \exp\left(\sum_{T: T\cap S \neq \varnothing} w_T \right).$$

For positive relevant k_i, let $h(S) = \log \mathbb{P}(S \subseteq E)$ and $\Delta_B h(A) = \sum_{J\subseteq B} (-1)^{|B\setminus J|} h(A \cup J)$. For disjoint A, B with $|B| \geq 2$,

$$\Delta_B h(A) = (-1)^{|B|+1} \sum_{T\supseteq B, T\cap A = \varnothing} w_T, \qquad Z_A = p_A^{-1} \exp\left(- \sum_{T\subseteq A} w_T \right).$$

These follow by forest deletion and substitution of the displayed Z_A. The loop expansion itself is an earlier input, not reproduced.

Record: RN, §19 (58)–(63), §30 (M6)–(M10); proved auxiliary identities.

9.2 Sure filling, possible filling, and cancellation faces

For a finite input count vector η, write $\eta(A) = \sum_{i\in A} \eta(i)$ and let

$$B_d = \{a : \text{every } a\text{-to-}s \text{ path visits } d\}, \qquad m_d = \eta(B_d).$$

Paths use positive-probability edges and may have length zero when reachability, rather than forced departure, is intended.

Theorem 9.4 (Sure-filling matroid)**.** *IDLA with parking allowed only on C fills C almost surely if and only if*

$$|C \cap B_d| \le \eta(B_d) \quad \text{for every } d \in V.$$

These sets C form a laminar matroid M_η.

Proof. If a constraint fails, choose for every source outside B_d a finite path to s avoiding d, hence avoiding all of B_d. There is positive probability that these particles follow those paths until parking or loss. Only $\eta(B_d)$ particles can then fill requested sites in B_d. Conversely, if $c \in C$ remains vacant, a particle starting in B_c cannot leave B_c without visiting c. It must therefore park in $(C \cap B_c) \setminus \{c\}$. Almost-sure termination gives $\eta(B_c) \le |C \cap B_c| - 1$, contradicting the constraint.

If B_c and B_d meet at a, order c, d along a simple a-to-s path. If c occurs first, a c-to-s path avoiding d would contradict $a \in B_d$; thus $B_c \subseteq B_d$. The family is laminar. For exchange, suppose feasible I, J satisfy $|I| < |J|$ but no element of $J \setminus I$ can be added to I. Each such element lies in an I-tight region. Maximal tight regions are disjoint; on their union U, $|J \cap U| \le |I \cap U|$, while outside U, $J \subseteq I$. This contradicts the cardinality inequality. □

Record: RN, §32 (V1); proved characterization. "Laminar matroid" is standard terminology, as credited there to Fife–Oxley.

Theorem 9.5 (Reachability and exact supports)**.** *Distinguish all $N = \eta(V)$ source tokens. Possible filling sets are exactly those matchable injectively to source tokens by reachability; they form a transversal matroid. More precisely, for $B \subseteq C$, exact output B with parking on C has positive probability if and only if $|B| \le N$ and all tokens can be matched to one slot at each $b \in B$ and $N - |B|$ sink slots by reachability in the graph with $C \setminus B$ deleted. A source at a deleted vertex has no eligible slot. The deleted graph need not remain sink-accessible.*

Proof. A vacant parking site cannot be visited. Realized paths thus give the required matching. Conversely choose finite paths for a matching. Follow each token's assigned path to its first vacant B-site or the sink. If the token parks at c assigned to another token, assign that other token its former path to c followed by the current token's remaining path to its assigned endpoint. Delete the parked token and the slot c. Distinct remaining endpoints and valid finite paths are preserved. A sink arrival deletes its sink slot. This produces a finite positive-probability history with output B. Setting $B = C$, and using sink accessibility for excess tokens, gives possible filling. Heredity and exchange follow by restricting matchings and using an alternating path, respectively. □

Record: RN, §32 "Possible filling, exact supports, and reachability matchings"; proved general characterizations, not inferred from the fixed audit.

Corollary 9.6 (Circuit support of the ARW response)**.** *For fixed η, x, v, with both inputs all active, let*

$$D_v(p) = \mathbb{P}_{\eta+\delta_x}(v \text{ occupied}) - \mathbb{P}_\eta(v \text{ occupied}).$$

In the ring of analytic germs at $p = 0$,

$$D_v \in \left\langle \prod_{i \in C \cup \{v\}} p_i : C \text{ a circuit of } M_\eta \right\rangle.$$

For every fixed $\alpha_i > 0$, this gives $D_v(\epsilon^\alpha) = O(\epsilon^b)$, where $b = \min_C \sum_{i \in C \cup \{v\}} \alpha_i$. If there are no circuits, D_v is identically zero.

Proof. Use RN's analytic zero-sleep trace reduction. On a coordinate face supported on an independent set S, target-capacity IDLA fills S surely, and its continuation is the full-input stationary law on the trace graph. The entire output-law difference vanishes on that face. Also D_v is divisible by p_v. Thus every surviving monomial has dependent support containing v, hence contains a circuit together with v. Grouping the convergent series by finitely many such divisors proves the ideal membership and the bound. Neither sharp order nor a coefficient sign is implied. □

Record: RN, §32 (V2)–(V3); proved consequence using §26 (127u), whose analytic trace proof is not reproduced here.

Remark 9.7 (Further combinatorial consequences). Let $D(x) = \{d : x \in B_d\}$. Matching source tokens to their dominator chains gives $M_{\eta+\delta_x} = M_\eta \vee U_{1,D(x)}$ and $J(M_{\eta+\delta_x}) = \bigcap_{d\in D(x)} J(M_\eta \backslash d)$, where J is the squarefree circuit ideal in a common polynomial ring. Every circuit has a unique dominator root $d \in C$ with $\eta(d) = 0$, $|C| = \eta(B_d) + 1$, and saturated child-region capacities. Adding at x destroys it exactly when $x \in B_d$. Its trace has one direct sink-leaking vertex d. Tracing to W restricts the matroid when η is supported on W, and $M_{\eta+\mathbf{1}_R}/R = M_\eta \setminus R$.

For the minimum possible occupied count $g_\eta(C)$,

$$\operatorname{rank}(M_\eta|C) \le g_\eta(C), \qquad g_\eta(C) = |C|-1 \quad (C \text{ a circuit}), \qquad g_{k\delta_x}(C) = \min\{k, \min_{\gamma: x\to s} |C\cap\gamma|\}.$$

There is also the exact recursion $g_{\eta+\delta_x}(C) = \min_{B\in \operatorname{supp} \mathrm{IDLA}_\eta^C}(|B| + \mathbf{1}_{\{C\setminus B \text{ separates } x \text{ from } s\}})$. The matching theorem gives a second formula: minimize $|B|$ over $B \subseteq C$ for which every token unable to reach s after deleting $C \setminus B$ can be matched into B.

For a single source, the union of supports over all input counts is the vertex-search antimatroid of sets reachable from x internally. For $k \ge 1$, its nonempty feasible sets B occur exactly when $|B| \le k$ and either $|B| = k$ or s is reachable from x inside $B \cup \{s\}$. Full filling with exactly n inputs has coefficient support

$$\{\eta \in \mathbb{N}^V : \eta(V) = n, \eta(S) \le |\operatorname{Reach}(S)| \ (S \subseteq V)\}.$$

Its Newton polytope is the sum, over targets, of simplices on their ancestor sets. These are support statements, not coefficient formulas. *Proof reference, not reproduced* for this group of consequences.

Record: RN, §32 (V1a)–(V1e), (V12)–(V13), and "Possible filling" (V8)–(V10); proved; the source reuses V8–V10, so the subsection is part of the locator.

Lemma 9.8 (Minimal sleep exponents)**.** *For one finite input, a deterministic state-dependent legal scheduler, and one nonzero terminal event probability f, let $\mathcal{H}$ be the componentwise-minimal successful-sleep count vectors of its complete histories, ignoring ineffective crowded sleeps. Then $\mathcal{H}$ is finite, is exactly the set of minimal Taylor exponents of f at zero, and those coefficients are positive. Its Newton polyhedron is* $\operatorname{conv}(\mathcal{H}) + \mathbb{R}^V_{\ge 0}$.

Proof. Deleting repeated raw states reduces sleep counts, so simple paths in the finite raw state graph witness all minimal vectors. The jump-only resolvent is analytic at zero with nonnegative entries there; its transience follows from the sink-hitting-time potential. A feasible sleep sequence contributes its sleep monomial times an analytic function with positive constant term. The expansion converges near zero because each additional sleep contributes $O(\|p\|)$ between bounded resolvents. No smaller feasible exponent can cancel a minimal monomial. □

Record: RN, §32 "Minimal sleep histories"; proved for one nonnegative event probability, not for a difference such as D_v.

9.3 Harmonic current tilts and completed paths

Allow finitely many absorbing exit labels $o \in O$. Let $A_i(z)(B, C)$ weight an addition response by z to its exit-count vector, and define $P_r(z)$ similarly. For $z_o > 0$, let $h_i = \mathbb{E}_i z_{\text{exit}}$ for one ordinary killed walk, $\chi(B) = \prod_{i \in B} h_i$, and $H = \operatorname{diag}(\chi(B))$.

Theorem 9.9 (Harmonic change of routing)**.** *The kernel*

$$K^z(i,j) = K(i,j)h_j/h_i, \qquad K^z(i,o) = K(i,o)z_o/h_i$$

is stochastic and sink-accessible. Keeping the same sleep probabilities, for terminal support C and exit counts ℓ,

$$\mathbb{P}_B^{K^z}(C, \ell \mid i) = \mathbb{P}_B^K(C, \ell \mid i)\frac{z^\ell \chi(C)}{h_i \chi(B)}.$$

Consequently $H^{-1}A_i^K(z)H = h_i A_i^{K^z}(1)$ and $H^{-1}P_r^K(z)H = P_r^{K^z}(1)$. For iid sources with law ρ, the spectral radius of $\sum_i \rho_i A_i(z)$ is $\lambda_\rho(z) = \sum_i \rho_i h_i$.

Proof. The finite harmonic equation for h gives positive finite h_i and row sums one. Support is unchanged. Give every intermediate configuration weight equal to the product of h_i over its particles, regardless of activity. A sleep leaves the weight unchanged; an internal jump multiplies it by h_j/h_i, and an exit by z_o/h_i. Ratios of complete-history probabilities therefore telescope to the displayed expression. Both processes stabilize almost surely, so summing nonnegative complete-history probabilities is legitimate. Conjugation gives the matrix identities. Dividing the driven conjugate by λ_ρ produces a stochastic matrix with source law $\rho_i h_i/\lambda_\rho$; hence the asserted spectral radius, without an irreducibility or simple-eigenvalue assumption. □

Record: RN, §33 (X1)–(X4); proved, independent AI audit recorded.

Theorem 9.10 (Current completion)**.** *Start empty. After m fully stabilized additions let H_m count actual labeled exits. There is a coupling to independent ordinary-walk endpoints, with counts Y_m, such that simultaneously for every m,*

$$Y_m = H_m + R_m, \qquad R_m \geq 0, \qquad \|R_m\|_1 \leq n.$$

For iid sources, write $\pi_o = \sum_i \rho_i \mathbb{P}_i(\text{exit} = o)$. Then

$$H_m/m \longrightarrow \pi \quad a.s., \qquad \frac{H_m - m\pi}{\sqrt{m}} \Longrightarrow N(0, \operatorname{diag}(\pi) - \pi\pi^T),$$

and, for every real ϑ,

$$\lim_{m\to\infty} m^{-1} \log \mathbb{E}e^{\vartheta \cdot H_m} = \log \sum_o \pi_o e^{\vartheta_o}.$$

Proof. Assign each labeled input an independent complete killed walk. At a legal move choose a label using only revealed information and expose its next step; this has conditional law K. Sleep interrupts but does not replace that walk. The uncompleted endpoints belong exactly to the at most n terminal residents, proving the pathwise identity. The iid strong law and central limit theorem transfer by the bounded error. For $z > 0$, putting $a = \min(1, \min_o z_o)$ and $b = \max(1, \max_o z_o)$ gives $b^{-n}\mathbb{E}z^{Y_m} \leq \mathbb{E}z^{H_m} \leq a^{-n}\mathbb{E}z^{Y_m}$, proving the logarithmic limit. □

Record: RN, §33 (X5)–(X7); proved for counts per addition after complete stabilization.

Proposition 9.11 (Internal currents and the full spectrum)**.** *Assign positive jump weights w_{ij}, including exits, and put $W_{ij} = K(i,j)w_{ij}$ internally and $k_w(i) = \sum_o K(i,o)w_{io}$. If $\rho(W) < 1$, take $h = (I - W)^{-1}k_w$. The same tilt uses $K^w(i,j) = W_{ij}h_j/h_i$ and $K^w(i,o) = K(i,o)w_{io}/h_i$. For actual internal-and-exit jump counts J_m under iid driving,*

$$\lim_m m^{-1} \log \mathbb{E} w^{J_m} = \log \sum_i \rho_i h_i.$$

Their strong law and central limit theorem equal those of the complete iid walk-count vectors. All driven ARW addition matrices at fixed K, ρ, w are similar to weighted IDLA, with spectrum, including multiplicity,

$$\lambda_B(w) = \sum_{i \in B} \rho_i [(I - W_{BB})^{-1} k_w|_B]_i, \qquad B \subseteq V.$$

It is real and nonnegative; Jordan blocks are allowed.

Proof. The history-ratio proof applies with jump weights. Summing transformed probabilities proves convergence of the weighted response series. Completing the at most n residents adds a uniformly exponentially tailed number of jumps, since finite sink-accessible walks have a uniform geometric block bound on their lifetimes. Borel–Cantelli and convergence in probability transfer the two limit theorems. Conditional completion gives $\mathbb{E}[w^{J_m}\chi(C_m)] = (\sum_i \rho_i h_i)^m$; the positive minimum and finite maximum of χ give the log limit. Surplus reduction preserves each jump count, so $P_r(w)A_i^r(w) = D_i(w)P_r(w)$. The tilt conjugates $P_r(w)$ to an invertible unweighted stabilization matrix. IDLA's supports only grow; its diagonal at B is precisely the weighted probability of exiting before finding a vacancy, the displayed λ_B. □

Record: RN, §33 (X8)–(X12); proved using the complete-block inverse.

The exit-weighted sleep transport is likewise $E_r^K(z)P_s^K(z) = H\,T_{r,s}^{K^z}(1)H^{-1}$ for $r \le s$; it records the physical waking-and-stabilizing operation and has nonnegative coefficients in exit marks. Its proof is the same history tilt applied to the established transport. At $w = 1$, if every vertex is reachable from the source law, the largest proper-support eigenvalue is $1 - \min_v \mathbb{P}_\rho(\text{hit } v)$, recovering the credited Bristiel–Salez relaxation result. The underlying ARW/IDLA sampling relation is credited to Levine–Liang; the records make no new universality claim. Physical-time activity and failed-sleep counts are not covered.

Record: RN, §33 (X3a), (X12) and literature paragraphs; proved local identities, credited external context.

9.4 Addition algebras, polytopes, and input inequalities

Theorem 9.12 (Addition algebra and its local factors)**.** *The algebra*

$$\mathcal{A}_K = \mathbb{R}[z_i : i \in V]/\langle z_i(z_i - k_i - \sum_j K(i,j)z_j) : i \in V\rangle$$

has squarefree basis z_B and acts regularly by IDLA additions. For core(B) *the vertices having a path to s inside $B \cup \{s\}$, feasible sets $R =$* core(R) *form a sink-search antimatroid. Every complex character is real, and equals*

$$h_R^R = (I - K_{RR})^{-1}k_R, \qquad h_{V\setminus R}^R = 0.$$

Put $F_R = \{x \notin R : k_x + \sum_{y\in R} K(x,y) > 0\}$ and $Y_R = V \setminus (R \cup F_R)$. The local algebra is

$$\mathcal{H}_{K_{Y_R Y_R}}, \qquad \mathcal{H}_L = \mathbb{R}[u_1, \dots, u_m]/\langle u_i((I - L)u)_i\rangle,$$

with Hilbert series $(1 + t)^m$ and squarefree basis. Hence the character has multiplicity $2^{|Y_R|}$ and $\mathcal{A}_K$ is semisimple exactly when every $k_i > 0$.

Proof status. *Proof reference, not reproduced* for the local-factor elimination and full basis proof. The underlying reduction is finite surplus IDLA: the relations move a crowded particle, its transient resolvent reduces every monomial to stable ones, and application to the empty-state row proves independence. At a character, nonzero coordinates solve the invertible displayed linear system. Feasible supports are therefore exactly its possible supports. RN gives the remaining local elimination.

Record: RN, §34 (A1)–(A4); proved local derivation with recorded independent AI audit. The commutative subset-basis IDLA algebra is credited there to Diaconis–Fulton.

Proposition 9.13 (Polytope realization and external geometric inputs)**.** *For transient substochastic L, $M = I - L$, and $b_i > 0$, put*

$$\mathcal{P}_b = \{x \in \mathbb{R}^m : x \geq 0,\ M^T x \leq b\}.$$

This is a simple combinatorial cube with fixed normal fan. Its volume algebra is $\mathcal{H}_L$, with upper facet classes u_i and opposite classes $v_i = (Mu)_i$. If $\mathrm{Top}(u_1 \cdots u_m) = 1$, then

$$\frac{\mathrm{Top}((\sum_i b_i u_i)^m)}{m!} = \det(M)\,\mathrm{Vol}(\mathcal{P}_b).$$

The class $\ell_b = \sum_i b_i u_i$ satisfies hard Lefschetz and the Hodge–Riemann relations: for $0 \leq d \leq \lfloor m/2 \rfloor$, $\ell_b^{m-2d} : (\mathcal{H}_L)_d \to (\mathcal{H}_L)_{m-d}$ is an isomorphism, and $(-1)^d \mathrm{Top}(a\ell_b^{m-2d}a')$ is positive definite on $\ker(\ell_b^{m-2d+1}) \cap (\mathcal{H}_L)_d$. Its top polynomial is Lorentzian.

Proof. Since $M^{-T} \geq 0$, $0 \leq x \leq M^{-T}b$, proving boundedness. The paired facets $x_i = 0$ and $(M^T x)_i = b_i$ cannot meet. Every choice of one facet from each pair gives a vertex: $x_{S^c} = 0$, $x_S = M_{SS}^{-T} b_S > 0$. All other upper inequalities are strict. A vertex needs m independent active facets, so these are all vertices and the fan is fixed. Translation relations give $v = Mu$; the only minimal empty facet intersections are $u_i v_i$. The external simple-polytope presentation theorem identifies the algebra.

The volume polynomial is annihilated by $\partial_i(M\partial)_i$, from varying the two disjoint opposite facets. Its full mixed derivative is $1/\det M$: the corresponding finite-difference region is $b_i < (M^T x)_i \leq b_i + \epsilon_i$, whose volume is $\prod_i \epsilon_i / \det M$. Since the top algebra degree is one-dimensional, this fixes the normalization. The fixed fan gives $\mathcal{P}_{b+c} = \mathcal{P}_b + \mathcal{P}_c$, including nonnegative supports by continuity. The external Lefschetz and volume-polynomial theorems now apply. □

Record: RN, §34 (A5)–(A7a); proved polytope bridge and normalization. External theorem proofs not reproduced.

Precisely, the external inputs are Timorin's simple-polytope presentation, Lefschetz and Hodge–Riemann results as stated in Bressler–Lunts, Section 4.1, Theorems 4.6, 4.8–4.9, and the volume-polynomial theorem of Brändén–Huh, Theorem 4.1 and Proposition 4.4. No rationality hypothesis is needed. Here "Lorentzian" denotes the standard closed class of homogeneous nonnegative polynomials with M-convex support and the associated one-positive-eigenvalue conditions on quadratic derivatives; we use these results as background, not as AI-original theorems.

Corollary 9.14 (Jordan blocks and inverse facets)**.** *For full-support driving $\rho_i > 0$, the character-R block of driven IDLA or ARW has, with $m = |Y_R|$, Jordan sizes $m - 2d + 1$ of multiplicity $\binom{m}{d} - \binom{m}{d-1}$, $0 \leq d \leq \lfloor m/2 \rfloor$. Here $\binom{m}{-1} = 0$. In $\mathcal{H}_K$,*

$$\mathrm{Top}(u_C v_{V\setminus C}) = d_V/d_C, \qquad \mathrm{Top}\prod_i (u_i + r_i v_i) = Z_V.$$

In $\mathcal{A}_K$, putting $v_i = ((I-K)z)_i - k_i$, the stable basis is $w_C^r = \prod_{i\in C}(z_i + r_i v_i) = \sum_B E_r(C,B) z_B$.

Proof status. The local driving linear form has strictly positive coefficients, so Lefschetz primitive strings give the block sizes. Vertex facet products integrate to reciprocal normal determinants, giving the two top formulas. *Proof reference, not reproduced* for the local driving coefficients, facet normalization details, and inverse-basis identification. These effective facet combinations need not be nef.

Record: RN, §34 (A9)–(A13); proved using external Lefschetz. The flat loss-marked family (A13) has free squarefree basis over $\mathbb{R}[t]$ and specializes to $\mathcal{H}_K$ at $t = 0$ and $\mathcal{A}_K$ at $t = 1$.

Theorem 9.15 (Full filling and exact-output input polynomials)**.** *For every fixed $N \geq n$, let F_η be the probability that ordinary IDLA fills V, $|\eta| = N$. Then*

$$Q_N(b) = \sum_{|\eta|=N} \binom{N}{\eta} F_\eta b^\eta$$

is Lorentzian, $Q_N^{1/N}$ is concave on the positive orthant, and $F^2_{\eta+e_i+e_j} \geq F_{\eta+2e_i} F_{\eta+2e_j}$ for $|\eta| = N-2$. The same inequality permits nonnegative source mixtures. The corresponding full-occupation ARW probability is exactly F_η/Z_V. All these assertions therefore transfer to full ARW occupation. Averaging an iid input law over routing automorphisms increases full filling; if the action is transitive, uniform input maximizes it, even for nonsymmetric sleep parameters.

For a fixed nonempty output B of size m and $N \geq m$, the polynomial $Q_{N,B}$ with coefficients $\mathbb{P}_\eta(\mathrm{IDLA} = B)$ is also Lorentzian, and, with $h = N - m$,

$$\frac{Q_{N,B}(b)}{N!} = \frac{\det(I - K_{BB})}{h!} \int_{u \geq 0,\ (I-K_{BB})^T u \leq b_B} (k_B \cdot u)^h \, du.$$

Inputs outside B give zero coefficients. For $h > 0$ the polynomial vanishes identically exactly when $k_B = 0$; for $N < m$ it is zero. The empty output occurs only at $N = 0$.

Proof and essential referenced input. *Proof reference, not reproduced* for the capacity extension: integer capacities $d_i \geq 1$, of total $N = \sum_i d_i$, give algebra $\mathbb{R}[u]/\langle u_i^{d_i}((I-L)u)_i\rangle$, Hilbert series $\prod_i (1 + \cdots + t^{d_i})$, and the simple polytope

$$x_i, y_{ia} \geq 0, \quad (M^T x)_i + \sum_{a=1}^{d_i - 1} y_{ia} \leq b_i.$$

With $\mathrm{Top}_d(u^d) = 1$ and $\ell_b = \sum_i b_i u_i$, its top normalization is $\mathrm{Top}_d(\ell_b^N)/N! = \det(M)\,\mathrm{Vol}_N(\mathcal{P}_b^d)$. To apply it, route original sink jumps to a buffer of capacity $N - n$. Full capacity without a further loss is exactly full filling of V; when $N = n$ no buffer is needed. Setting buffer input to zero gives a volume polynomial by continuity. For exact output B, kill forbidden exits to $V \setminus B$ and buffer only original sink jumps, obtaining the integral by integrating buffer fibers. Finally surplus IDLA cannot produce n ARW residents from a smaller support, while full binary support does so with probability $1/Z_V$. This proves the transfer. Concavity and symmetry give the automorphism statement.

Record: RN, §34 (A15)–(A20); proved capacity extension and consequences, using external volume inequalities.

Corollary 9.16 (Posterior source counts)**.** *For $N \geq 1$ iid sources with strictly positive law ρ, condition on an exact IDLA output B of positive probability. If $m_i = \mathbb{E}[\eta_i \mid B]$, then for every real t,*

$$\mathbb{E}[e^{t\cdot\eta} \mid B] \leq \left(\sum_i \frac{m_i}{N} e^{t_i}\right)^N, \qquad \mathrm{Cov}(\eta \mid B) \preceq \mathrm{diag}(m) - mm^T/N.$$

Under any prior on finite all-active inputs, the posterior given full ARW occupation equals that given full IDLA filling, when the conditioning event has positive prior probability, and is independent of sleep rates.

Proof. The conditional generating function is $Q_{N,B}(\rho z)/Q_{N,B}(\rho)$. The tangent-plane bound for the concave function $Q_{N,B}^{1/N}$ gives the displayed moment bound; second derivatives along every line at zero give the covariance order. The factor $1/Z_V$ cancels in Bayes' formula for the full-occupation assertion. This proves neither convex order nor signs of individual off-diagonal covariances. □

Record: RN, §34 (A22)–(A24); proved consequence of the credited concavity input.

9.5 Poissonized IDLA and exponential capacities

Let $b_i \geq 0$, $M = I - K_{VV}$, and independently sample $\xi_i \sim \mathrm{Exp}(1)$. Fluid enters at rates b_i, fills capacity ξ_i, and thereafter follows routing proportions K. Write $u(t)$ for total emitted fluid and $\mu(t)$ for stored mass.

Theorem 9.17 (Fluid representation)**.** *There is a unique solution of*

$$\mu = tb - M^T u, \quad u \geq 0, \quad 0 \leq \mu \leq \xi, \quad u_i(\xi_i - \mu_i) = 0, \qquad u = (tb - \xi + K^T u)_+.$$

For each fixed b, the saturated set has jointly in $t \geq 0$ the law of IDLA driven by independent Poisson processes of rates b_i. In particular its full-filling probability at time one is

$$\mathbb{P}(M^{-T}\xi \leq M^{-T}b).$$

Proof. Set $a = M^{-T}\mathbf{1} > 0$. Since $K^T a = a - \mathbf{1}$, the fixed-point map contracts in the a-weighted supremum norm with constant $\max_i(1 - 1/a_i) < 1$. At a fixed point, if $u_i > 0$ then $\mu_i = \xi_i$; otherwise $\mu_i = tb_i + (K^T u)_i \geq 0$. Thus the fixed point is exactly the complementarity solution.

When the saturated set is B, its emission rate is $(I - K_{BB}^T)^{-1} b_B$, and an unsaturated y receives at rate

$$a_y(B) = b_y + \sum_{x \in B} b_x [(I - K_{BB})^{-1} K_{B,y}]_x.$$

Conditional on the past saturation times and identities, the remaining deficits are independent unit exponentials: an exponential race followed by memorylessness preserves this property. Hence $B \to B \cup \{y\}$ has rate $a_y(B)$. In Poisson-IDLA an arrival in B walks to its first vacancy or the sink; the same Green matrix gives exactly these rates. Both finite pure-birth chains start empty, proving equality of process laws. Zero rates cause no transition. Finally full saturation forces $u = M^{-T}(b - \xi) \geq 0$; conversely this u with $\mu = \xi$ solves the complementarity equations. Threshold equalities have probability zero at each deterministic time. □

Record: RN, §35 (F16)–(F18); proved IDLA representation. No whole-law ARW fluid representation is asserted.

Theorem 9.18 (Chambers and labeled exit counts)**.** *Put $M_B = I - K_{BB}$ and $\mathcal{P}_B(b_B) = \{u \geq 0 : M_B^T u \leq b_B\}$. For every $B \subseteq V$,*

$$\mathbb{P}(\text{output} = B) = e^{-\sum_i b_i} \det(M_B) \int_{\mathcal{P}_B(b_B)} e^{k_B \cdot u}\, du.$$

Empty-set determinants and integrals equal one. For absorbing labels, write $k_o(i) = K(i, o)$, $k = \sum_o k_o$, and let L_o be exit counts. Then

$$\mathbb{E}\left[\mathbf{1}_{\{\text{output}=B\}} \prod_o z_o^{L_o}\right] = e^{-\sum_i b_i} \det(M_B) \int_{\mathcal{P}_B(b_B)} \exp\left(\sum_o z_o k_o|_B \cdot u\right) du.$$

Conditional on emitted fluid u, the L_o are independent $\mathrm{Poisson}(k_o \cdot u)$. *Conditional on B, u has density proportional to $e^{k_B \cdot u}$ on its chamber. Conditional additionally on no exits, this density is uniform, whenever the event has positive probability.*

Proof. In the chamber B, $u_{B^c} = 0$, $\xi_B = b_B - M_B^T u_B$, and for $y \notin B$, $\xi_y > b_y + \sum_{i \in B} K(i,y) u_i$. Integrate the independent exponential tails outside B and change variables inside B. The Jacobian is $\det M_B$; all exits to unoccupied internal sites cancel in the exponent, leaving only original-sink terms. This proves the first formula, initially for $b > 0$ and then by continuity.

At support B, loss events with label o have rate $\sum_{i \in B} b_i [(I - K_{BB})^{-1} k_o|_B]_i = k_o|_B \cdot \dot{u}_B$. They do not change support. Independently adjoin Poisson processes run at these integrated rates to the growth chain. Its joint generator agrees with Poisson-IDLA. Conditional on the fluid history, the means are $k_o \cdot u(1)$, hence the same is true conditional on terminal u. Multiplying the chamber density by the Poisson generating factors proves the marked identity and the uniform conditional law. □

Record: RN, §35 (F19)–(F24); proved chamber and conditional-current identities, independent AI audit recorded.

Theorem 9.19 (Log-concavity and association)**.** *For every fixed exact support B, its Poisson-IDLA probability is log-concave in $b \geq 0$. For every target set R, the probability $a_R(b)$ of avoiding R is also log-concave. At any fixed b, the joint vector of occupied indicators and original exit counts is positively associated: increasing bounded functions have nonnegative covariance. In particular distinct exit counts and an occupied indicator and an exit count have nonnegative covariance. These assertions are unconditional; neither conditioning on a support nor fixing the input particle count is covered.*

Proof. The chamber integrand, including its domain indicator, is jointly log-concave in (b, u); apply the external marginal log-concavity theorem of Prékopa, Theorem 6. Restricting b first to an arbitrary bounded convex box ensures integrability if needed.

Write $z = b - \xi$ and $G_B = (I - K_{BB})^{-1}$. Complementarity gives

$$u_i(z) = \max\bigl(0, \max_{B \ni i} [G_B^T z_B]_i\bigr).$$

Indeed $M^T u \geq z$ implies $M_B^T u_B \geq z_B$, giving every lower bound in the maximum. On $A = \{i : u_i > 0\}$ equality gives $u_A = G_A^T z_A$; outside A all lower bounds are nonpositive. All displayed coefficients are nonnegative. At a fixed time the occupied indicator equals $\mathbf{1}_{\{u_i > 0\}}$ almost surely: exceptional equalities belong to finitely many hyperplanes in the continuous ξ. It is therefore increasing in the independent variables $-\xi_i$. Realize conditional Poisson counts by independent inverse-CDF uniforms; they too are increasing in these variables and their respective uniforms. Association follows from the one-variable identity $2\,\mathrm{Cov}(f(X), g(X)) = \mathbb{E}[(f(X) - f(X'))(g(X) - g(X'))]$, where X' is an independent copy, and induction by conditional covariance.

Avoidance of R is the intersection, for $i \in R$ and $B \ni i$, of $[G_B^T (b_B - \xi_B)]_i \leq 0$, together with $\xi \geq 0$. This is jointly convex in (b, ξ). Integrating its exponential density and applying the same external theorem proves log-concavity. □

Record: RN, §35 (F19), (F34)–(F37); proved applications; Prékopa's theorem is an external input, proof not reproduced.

For Poisson all-active ARW, the exact law remains $\mu_b^{\mathrm{ARW}}(C) = \sum_B \mathbb{P}_b(\mathrm{IDLA} = B) P_r(B, C)$. Consequently full occupation has the fluid probability divided by Z_V and is log-concave. To retain exits, sample the chamber's conditional Poisson losses and add those of fresh binary ARW continuation. This samples the Poisson-mixed law, not a law conditional on prescribed input

counts. The continuation is not known to preserve set inclusion or the preceding dependence properties.

Record: RN, §35 (F21), (F24), and §34 (A19); proved mixture, restricted transfer.

Proposition 9.20 (Cover and first-hit times)**.** *Let ρ be a source distribution reaching every vertex, $G = (I - K)^{-1}$, $g_i = (G^T\rho)_i$, and $a_i = G_{\cdot i}/g_i$. The rate-one Poisson-IDLA cover time satisfies $T = \max_i a_i \cdot \xi$ in the fluid coupling, and*

$$\lim_{t\to\infty} t^{-1} \log \mathbb{P}(T > t) = -H, \qquad H = \min_i H_i, \quad H_i = \mathbb{P}_\rho(\text{hit } i).$$

If $\rho_i > 0$ for every i, put $I = \{i : H_i = H\}$, $h_j(i) = \mathbb{P}_j(\text{hit } i)$, and $C = \sum_{i\in I} \prod_{j\neq i}(1 - h_j(i))^{-1}$. Then for some $\epsilon > 0$,

$$\mathbb{P}(T > t) = Ce^{-Ht} + O(e^{-(H+\epsilon)t}).$$

For $n \geq 2$, the discrete cover count τ satisfies $\mathbb{P}(\tau > N) = C(1 - H)^N + O(\beta^N)$ for some $0 < \beta < 1 - H$. For $n = 1$, $\tau = 1$ and $T \sim \text{Exp}(1)$.

Proof status. The maximum formula follows directly from full filling. Since $G_{ji} = h_j(i)G_{ii}$, the maximal coefficient in a_i is $1/H_i$. A single exponential coordinate gives the lower tail exponent; exponential moments of finitely many linear forms give the upper exponent. *Proof reference, not reproduced* for the sharp constant, exponential error, and discrete spectral comparison, including semisimplicity of the largest proper eigenvalue. Their full-support hypothesis is essential to the stated proofs.

Record: RN, §35 (F25)–(F28); proved representation and sharp tails, with independent AI proof review recorded.

Proposition 9.21 (First filling and retirement)**.** *Assume $\rho_i > 0$ and $\varnothing \neq R \subseteq V$. The first time τ_R that any target is occupied is*

$$\tau_R = \min_{i\in R,\ B\ni i} \frac{[G_B^T\xi_B]_i}{[G_B^T\rho_B]_i}.$$

It is concave, increasing and homogeneous in ξ, and its survival is log-concave in time. Set $D_R = \{j : \mathbb{P}_j(\text{hit } R) = 1\}$, $m = |D_R \setminus R|$, and $H_R = \mathbb{P}_\rho(\text{hit } R)$. There is an explicit positive constant C_R with

$$\mathbb{P}(\tau_R > t) \sim C_R t^m e^{-H_R t}.$$

If $H_R < 1$, the discrete first-hit count satisfies $\mathbb{P}(\tau_R^{\text{disc}} > N) \sim C_R(1 - H_R)^{-m}N^m(1 - H_R)^N$. If $H_R = 1$, that count is at most $m + 1$, and the continuous survival is e^{-t} times a polynomial of degree m.

Proof reference, not reproduced. The minimum follows from the emission maximum. RN (F40) gives C_R as a product of positive permutation sums for entry into $V \setminus D_R$ and subsequent births in $D_R \setminus R$, divided by $m!$. The finite acyclic birth-chain expansion proves the tails. For the stopping interpretation, put $h = (I - K^T)^{-1}\rho$, $Q = \text{diag}(h)^{-1}K^T \text{diag}(h)$, $\kappa_i = \rho_i/h_i$, and $c_i = \xi_i/h_i$. Continuing at i costs c_i and takes one Q-step; intrinsic killing pays reward t, whereas retirement before paying costs gives zero. The Bellman equation

$$y_i = \max(0, t\kappa_i - c_i + \sum_j Q(i, j)y_j)$$

contracts, and $y_i = u_i/h_i$. For every vertex v, τ_v is the infimum of expected cost divided by killing probability over all randomized history-dependent policies with positive killing

probability. The domain $\{v\} \cup \{i : \tau_i < \tau_v\}$ attains it. This killing mechanism reverses input injection; it is not the original sink.

Record: RN, §35 (F37)–(F45); proved first-hit tails and optimal-stopping characterization.

Proposition 9.22 (Arrival-count moments)**.** *For a strictly positive source distribution ρ, let $f_N = \mathbb{P}(\tau \le N)$, U be uniform on $\mathcal{P}_\rho$, and $L = k \cdot U$. Then L has support $[0,1]$ and lies in $(0,1)$ almost surely, and*

$$a_h := \frac{f_{n+h}}{\binom{n+h}{n}} = f_n \mathbb{E}L^h, \qquad (-1)^r \Delta^r a_h = f_n \mathbb{E}[L^h(1-L)^r] > 0.$$

Here $\Delta a_h = a_{h+1} - a_h$. More precisely, for every $r, s, m \ge 0$ the matrix $((-1)^s \Delta^s a_{r+i+j})_{0 \le i,j \le m}$ is positive definite. Moreover $f_N^2 \ge f_{N-1} f_{N+1}$ for every $N \ge 0$ (with the initial zeros and $f_{-1} = 0$), and for $h \ge 1$,

$$1 \le \frac{f_{n+h}^2}{f_{n+h-1} f_{n+h+1}} < \frac{(n+h)(h+1)}{h(n+h+1)}.$$

For each fixed exact IDLA output B, $N \mapsto f_N(B)$ is likewise log-concave. For nonempty B, its support is $\{|B|\}$ if $k_B = 0$ and all $N \ge |B|$ otherwise; the empty output has support $\{0\}$.

Proof and external input. Expanding the chamber exponential gives the moment formula. $M^{-T}\rho$ is the coordinatewise maximum of $\mathcal{P}_\rho$ and $k \cdot M^{-T}\rho = 1$; intermediate points of its radial segment are interior. This proves the support assertion and strict moment inequalities, including positive definiteness by integrating squared nonzero polynomials. The additional log-concavity of f_N uses Fradelizi–Li–Madiman, Theorem 3.1: the normalized moments $\prod_{j=1}^n (t+j) \int_{\mathcal{P}_\rho} (k \cdot u)^t du$ are log-concave for $t > 0$, extending to 0 by continuity. Its proof is external and not reproduced. Apply the same argument in dimension $|B|$ to exact outputs, using only original-sink probabilities.

Record: RN, §35 (F29)–(F32b); proved moment identities and application of an external normalized-moment theorem.

For completeness, with $m = |B|$, uniform U_B on $\mathcal{P}_B(\rho_B)$ and $L_B = k_B \cdot U_B$, the exact formula is $f_{m+h}(B) = f_m(B)\binom{m+h}{m}\mathbb{E}L_B^h$. Its upper moment endpoint is $\theta_B = \rho_B^T (I - K_{BB})^{-1} k_B$. Thus $\sum_N f_N(B) z^N = f_m(B) z^m \mathbb{E}(1 - zL_B)^{-m-1}$ for $0 < z < 1$. Under a geometric input-count prior, the conditional surplus count is the negative-binomial mixture obtained by tilting U_B by $(1 - zL_B)^{-m-1}$. This particular mixture inherits log-concavity; arbitrary such mixtures do not. For proper B with $k_B \ne 0$, $f_{N+1}(B)/f_N(B)$ decreases to $\theta_B < 1$.

Record: RN, §35 (F32a)–(F32b); proved exact-output consequences; mixture derivation by the displayed generating identity.

Proposition 9.23 (Residual semigroup)**.** *For deterministic capacities $w \ge 0$, let $U_b(w) = (b - w + K^T U_b(w))_+$ and $\Phi_b(w) = w - b + M^T U_b(w)$. Then*

$$\Phi_{b+c} = \Phi_c \circ \Phi_b, \qquad U_{b+c}(w) = U_b(w) + U_c(\Phi_b(w)).$$

Let $R_{\exp}(B, \cdot)$ put zero capacities on B and independent unit exponentials elsewhere. Then

$$R_{\exp} \Phi_b = \exp\left(\sum_i b_i (D_i - I) \right) R_{\exp}.$$

For every prescribed deterministic nondecreasing càdlàg cumulative input path $b(t)$, this yields the Poisson-IDLA support process with those cumulative intensities; conditional residuals given the support history remain independent unit exponentials on vacancies.

Proof. Residuals decrease coordinatewise. Adding the two successive emission vectors satisfies the combined linear equation; if either emits at a site, the final residual there is zero. Uniqueness of complementarity proves the semigroup. During a phase with vacant-site rates a_y, the joint density for first event (t, j) and remaining deficits v_y is $a_j e^{-t\sum_y a_y} dt \prod_{y\neq j} e^{-v_y} dv_y$. This gives the residual law after an event and conditional on no event. Induction over at most n births proves the kernel identity. Apply it on finite time partitions; rational observation times and right continuity give the prescribed-path assertion. A jump in b means a fully stabilized Poisson batch. □

Record: RN, §35 (F46)–(F48); proved semigroup and process identity. No simultaneous identification with one original stack coupling over incomparable input paths is claimed.

The associated signed ARW face measures $\Sigma_r = E_r R_{\exp}$ obey $\Sigma_r \Phi_b = \exp(\sum_i b_i(A_i^r - I))\Sigma_r$. For a stable-law row μ, $\mu\Sigma_r$ is positive exactly when $\mu E_r \geq 0$, since the zero-set faces are disjoint. This is the cone of mixtures of binary all-active laws, containing every all-active input law by surplus IDLA. Individual stable-state embeddings can be signed. *Proof reference, not reproduced* for the complete one-site classification recorded among the obstructions below.

Record: RN, §35 (F49)–(F50); proved signed intertwining and scoped obstruction.

9.6 Complementary-sleep duality and the target obstruction

In $\mathcal{A}_K$, let Λ extract the coefficient of z_V. Thus $F(C, D) = \Lambda(z_C z_D)$ is the probability that IDLA from $\mathbf{1}_C + \mathbf{1}_D$ fills V. Its Frobenius pairing is nondegenerate, by the top-degree Poincaré duality of $\mathcal{H}_K$. Define $J_i = k_i + \sum_j K(i, j)z_j$, $v_i = z_i - J_i$, $w_i^r = z_i + r_i v_i$, and

$$H_{r,s}(C, D) = \Lambda(w_C^r w_D^s) = (E_r F E_s^T)(C, D).$$

The letter s in the subscript denotes a second odds vector, not the absorbing vertex.

Theorem 9.24 (Positive duality and stochastic reversal)**.** *For positive finite odds vectors with $r_i s_i \leq 1$ at every site, $H_{r,s}$ is nonnegative and invertible, and*

$$A_i^r H_{r,s} = H_{r,s}(A_i^s)^T.$$

Every stationary coordinate $\pi_s(D)$ is positive. The matrix $B_{r,s} = H_{r,s}\,\mathrm{diag}(\pi_s)$ is stochastic, $\pi_r B_{r,s} = \pi_s$, and

$$A_i^r B_{r,s} = B_{r,s} A_i^{s,*}, \qquad A_i^{s,*} = \mathrm{diag}(\pi_s)^{-1}(A_i^s)^T \mathrm{diag}(\pi_s).$$

Moreover $B_{r,s}(\varnothing, V) = 1$, and its stationary coupling is supported on $|C| + |D| \geq n$. The condition is equivalently $p_i + p_i' \geq 1$; complementary sleeps correspond to $s_i = 1/r_i$.

Proof. Commutative multiplication is self-adjoint for $\Lambda(ab)$, giving duality in the stable bases. Both inverse basis matrices and the pairing are invertible. For positivity split C, D into overlap I and symmetric difference S. Since $z_i v_i = 0$,

$$w_i^r w_i^s = (1 - r_i s_i) z_i J_i + r_i s_i J_i^2 \quad (i \in I).$$

This is a nonnegative polynomial. The remaining factors are w_S^t with hybrid odds $t_i = r_i$ on $C \setminus D$ and $t_i = s_i$ on $D \setminus C$, with arbitrary positive values elsewhere. For every active input η,

$$\Lambda(w_S^t z^\eta) = Z_V(t)\mathbb{P}_t(\text{full occupation from sleepers } S \text{ and active input } \eta) \geq 0.$$

Indeed multiplication is active addition in the stable basis, and $\Lambda(w_B^t) = Z_V(t)\mathbf{1}_{\{B=V\}}$. Expanding the overlap factors proves positivity. Equivalently, each overlap releases one particle

by a forced jump and either retains the other at its site or forces it to jump too, with respective probabilities $1 - r_i s_i$ and $r_i s_i$.

Since $F(V, D) = F(C, V) = 1$, $E_r \mathbf{1} = \mathbf{1}$, and $\pi_r E_r = \delta_V$, we obtain $\pi_r H_{r,s} = \mathbf{1}^T$ and $H_{r,s}\pi_s^T = \mathbf{1}$. These identities normalize the stochastic link and stationary coupling. To see $\pi_s(D) > 0$, retain one distinguished active at every D-site, route all other particles along finite paths to the sink, and finally sleep the distinguished particles. Visiting particles can legally pass through retained sites; the finite prescribed instruction event has positive probability. Reversal is therefore defined and the displayed intertwining follows. Degree gives $H(C, D) = 0$ below total degree n; $\Lambda(w_D^s)$ gives the empty row assertion. □

Record: RN, §36 (D1)–(D8); proved positive duality; dependencies are the complete-block basis and §34 Frobenius pairing, whose nondegeneracy proof is referenced above.

Proposition 9.25 (Support and impossibility of a positive target payoff)**.** *For complementary odds, put $I = C \cap D$, $S = C \triangle D$, and $T = V \setminus S$. Then $H_{r,1/r}(C, D) > 0$ exactly when T can be matched injectively to two tokens at every $i \in I$, with eligibility by positive-length interior directed paths. On a DAG without self-loops, $B_{r,1/r}(C, D) = \mathbf{1}_{\{D=V\setminus C\}}$.*

If the interior graph is strongly connected and $n \geq 2$, for every allowed r, s,

$$B_{r,s}(C, D) > 0 \quad \Longleftrightarrow \quad |C| + |D| \geq n.$$

For every nonempty proper target set R, there is no nonnegative scalar function g satisfying, for every support C,

$$B_{r,s} g(C) = \mathbf{1}_{\{R \subseteq C\}}.$$

The impossibility persists if equality is requested only after every finite all-active input.

Proof. At complementary odds every overlap forces both particles to jump. After fixing these first destinations, full occupation is possible exactly when IDLA parking on T fills T: a surplus phase either leaves too few particles, or fills all sites, after which all active singletons can sleep with positive probability. The matching theorem therefore applies. Varying the first jumps gives exactly positive-length reachability. On a DAG a topologically earliest overlap vertex cannot be supplied by any overlap token; hence the overlap must be empty. Degree then requires complementary supports, and stochasticity fixes the probability to one.

In a strongly connected graph with at least two vertices, every vertex can reach every vertex by a positive-length interior path. Matching is therefore equivalent to $|T| \leq 2|I|$, or $|C| + |D| \geq n$. For general allowed positive odds, the event choosing two forced jumps at every overlap has positive probability; the same support follows. Degree gives the converse. Now choose $v \in R$. The zero target value at $V \setminus \{v\}$ forces $g(D) = 0$ for every nonempty D. But $B(R, \varnothing) = 0$ because R is proper, contradicting the value one at R. Equality after binary active inputs multiplies the equation by the invertible P_r, so cannot avoid the obstruction. □

Record: RN, §36 (D9)–(D12); proved support criterion and scalar-payoff obstruction, not a failure of monotonicity.

Remark 9.26 (Transport compatibility and vacancy dynamics). Write $\Pi_r = \mathrm{diag}(\pi_r)$ and $T_{a,b} = E_a P_b$ for $a \leq b$. Its Bayes reverse is $\widehat{T}_{b,a} = \Pi_b^{-1} T_{a,b}^T \Pi_a$. It is stochastic, cannot decrease particle count, and composes along ordered parameter chains. If $r \leq t$, $s \leq u$, $t_i u_i \leq 1$, then

$$B_{r,s} = T_{r,t} B_{t,u} \widehat{T}_{u,s}, \qquad T_{r,t} B_{t,1/t} = B_{r,1/r} \widehat{T}_{1/r,1/t}.$$

The links B themselves are not a cocycle.

For $L_r(C, E) = B_{r,1/r}(C, V \setminus E)$, positive entries imply a matching of E into C by length-zero-allowed interior reachability. Equivalently $|E \cap U| \leq |C \cap U|$ for every ancestor-closed U.

Its communicating classes have fixed particle count in each strongly connected component, and every entry within a class is positive. The absorbing states are precisely subsets of vertices on no interior directed cycle, counting self-loops. The artificial iterated L_r chain has a geometric absorption tail on each fixed graph. If every vertex is cyclic, its only absorbing state is empty, and

$$h_v(C) = \mathbb{E}_C \sum_{j<\tau_\varnothing} \mathbf{1}_{\{v\in C_j\}}, \qquad (I - L_r)h_v = \mathbf{1}_{\{v \text{ occupied}\}}, \quad h_v \geq 0.$$

This history potential need not be increasing in C. Reachability support does not assert monotonicity of the kernel in its input or a deletion-only realization. *Proof reference, not reproduced* for these algebraic compatibility and reachability consequences.

Record: RN, §36 (D14)–(D21); proved formulas and support/absorption results; exact obstruction (D21).

9.7 Decisive obstructions, kept separate from positive results

The following are exact examples or proved representation obstructions, not extrapolations from parameter checks.

1. **Mixing loses strong Rayleigh stability.** For three reversible roots with $K(i,j) = 1/4$ for $i \neq j$, $K(i,s) = 1/2$ and $p_i = 1/2$,

$$\mathbb{E}z^E = (73 + 33e_1(z) + 12e_2(z) + 4e_3(z))/212.$$

   Its diagonal cubic has discriminant -15984, so it is not real-rooted. Thus the mixture is neither strongly Rayleigh nor a Hermitian determinantal law, though the scoped Rayleigh theorem applies. This does not rule out a nonsymmetric determinantal kernel: a directed half-sink three-cycle has polynomial $(\prod_i(3+z_i) - 7)/57$ and kernel A^{-1} for $A = \begin{pmatrix} 4 & -1 & 0 \\ 0 & 4 & -1 \\ -7 & 0 & 4 \end{pmatrix}$.

   Record: RN, §30 (M5), "Distinguish nonsymmetric determinant kernels"; exact counterexamples.

2. **Circuit support is not coefficient positivity.** On $0 \to 1 \to 2$, with $2 \to 0, s$ equally likely, lower input $\{0,2\}$, addition at 1, and target 2, the proved exact rational identity gives

$$D_2 = p_0p_1p_2[p_0 + p_1 + O(\|p\|(p_0 + p_1))]$$

   uniformly on a sufficiently small nonnegative cube. For every fixed $\alpha_i > 0$, $p_i = \epsilon^{\alpha_i}$ has leading exponent $\alpha_0+\alpha_1+\alpha_2+\min(\alpha_0,\alpha_1)$, with coefficient 1 for unequal α_0, α_1 and 2 for equality. The Newton polyhedron is $\mathrm{conv}\{(2,1,1),(1,2,1)\} + \mathbb{R}^3_{\geq 0}$, but $[p_0^3p_1p_2]D_2 = -2$. *Proof reference, not reproduced* for the exact rational calculation; no arbitrary-cycle or pendant extension is inferred.

   Record: RN, §32 "A complete anisotropic calculation", (V4)–(V11); proved symbolic example.

3. **Positive geometry does not represent the target positively.** On $0 \to 1 \to s$, $\mathrm{Top}((u_0 + r_0v_0)^2) = 1 - r_0^2$, negative for $p_0 < 1/2$. The unique target-1 representer for $\Lambda(g_1z^\eta)$ is $g_1 = p_1z_0 - p_0p_1z_1 + p_0p_1$. Its negative coefficient rules out a fixed positive mixture of ordinary IDLA completion inputs, including countable mixtures. At a zero-buffer boundary the volume conclusions survive but Lefschetz can fail: in $\mathbb{R}[u,w]/\langle u(u-w), w^3\rangle$, multiplication by u kills $u - w$ from degree one to degree two.

   Record: RN, §34 (A12), (A14), and after (A19); exact algebraic obstructions.

4. **Full filling cannot be replaced by a proper ARW target.** On $0 \to 2$, $1 \to 2$, $2 \to s$, the target-2 degree-two input polynomial is $p_2[b_0^2 + 2(1-p_0p_1)b_0b_1 + b_1^2]$, which is not Lorentzian. With Poisson inputs only at $0, 1$ it is $p_2[1 - e^{-b_0-b_1}(1+p_0b_0)(1+p_1b_1)]$. For $p_0 = p_1 = p$ its value at $(t+h, t-h)$ exceeds that at (t,t) by $p_2p^2e^{-2t}h^2$, violating log-concavity. The target probability is nevertheless coordinatewise increasing.

   Record: RN, §35 (F22)–(F23); proved DAG counterexamples to stronger properties.

5. **Input and output polynomials differ.** With $x \to a, c$ equally likely, $a \to b \to s$, $c \to d \to s$, three inputs at x give output polynomial $z_x(z_az_b + 2z_az_c + z_cz_d)/4$. Its derivative in z_x, restricted to $(z_a, z_b, z_c, z_d) = (u, u-2v, v, v)$, is $(u^2+v^2)/4$, violating the Lorentzian Hessian condition. Poisson-IDLA association also need not imply the FKG lattice inequality: on the preceding merging fork, at $b_0 = b_1 = t > 0, b_2 = 0$, the support $\{2\}$ is impossible while $\{0,2\}, \{1,2\}$ both have positive probability.

   Record: RN, §34 (A21), §35 after (F35); exact distinct limitations.

6. **The unchanged residual flow admits only a trivial positive one-site embedding.** For $\phi_b(w) = (w-b)_+$ and direct-sink one-site ARW, all mass-one finite signed intertwining pairs are $H_0 = (1-a)\delta_0 + a\,\mathrm{Exp}(1)$ and $H_1 = (1+ra)\delta_0 - ra\,\mathrm{Exp}(1)$, $a \in \mathbb{R}$. For $r > 0$ both are positive only at $a = 0$. This excludes that embedding, not every positive representation. The separate terminal-payoff obstruction is Proposition 9.25; adding a constant cannot repair its failed inclusion order. Even the positive vacancy-chain history potential can decrease: on the maintained three-cycle at $p = (1/3, 2/5, 3/5)$, $h_0(\{0\}) = 189979/74354$ exceeds $h_0(\{0,1\}) = 40836185505115661/21542872302627536$.

   Record: RN, §35 (F50), §36 (D12)–(D13), (D21); proved classification and exact obstruction; classification proof reference, not reproduced.

Attribution boundary. The supplied ledger records every maintained family in the assigned source sections, including smaller consequences stated here by proof reference. The forest-root determinant law is credited in RN to Avena–Gaudillière; conditional determinantal stability to Borcea–Brändén–Liggett. The interval specialization is already related to IDLA and volume polynomials in Nadeau–Tewari's q-Klyachko work, and its rational-q toric realization is credited to Ruizhen Liu. RN's comparisons with Skorokhod reflection, Pitman–Stanley polytopes, exponential truncated powers, and geometric current universality are contextual identifications, not originality claims or extra hypotheses silently imported into the proofs.

Part IV
Open questions and evidence

10 Open claims, failed restrictions, and source gaps

The general conjecture in Section 1 remains open. The large number of identities in this collection should not be read as a proof-proximity claim. The following distinctions are essential when reusing the results.

10.1 What the proved regimes leave open

The small-sleep theorem has a graph-dependent threshold and a hitting hypothesis; it is not a uniform statement over graph size. The high-sleep theorem controls the two-common-particle creation cone. The general three-site theorem localizes to a target component with at most three vertices. Directed-cycle full-minus-one order and forest feedback compare the specified inputs. None alone proves arbitrary-input monotonicity on a general component of size four or more.

The backwards constructions give exact identities, positive auxiliary laws, and finite inverses. Several final occupation expressions still contain a signed commutator, an unproved capacity inequality, or a conditional continuation which is not increasing. Positivity of a normalizing measure cannot replace the missing sign of its continuation. Finite transport feasibility, short release horizons, and selected parameter checks remain evidence at their tested scope.

10.2 Restrictions that have already failed

The maintained checkpoint gives exact witnesses and reopening conditions. Here the failures are grouped by the step they exclude, to avoid a chronological diary.

Proposed restriction	What the obstruction actually says	Record
Sleeping attractiveness	An arbitrary sleeping background can lose target occupation after an addition. This does not refute all-active creation-law monotonicity.	VA §7; RN §6
Singleton or separate repairs	Adaptive singleton activations, fixed consumed inventories, and separate finite repairs can lack capacity. Weighted transfers repair examples without proving the general budget.	VA §6; BS §5; LF; CE; WT
Local or early debt collection	Fixed first-collision classes, whole priority phases, and even optimal causal scheduling before first sleep/loss/hit can miss needed credit.	NC; PR; MC; TC; WC
Latent blocks treated as fresh	Realized collisions omit unrevealed dependencies; conditioning on the discovered block changes its law.	LD; FC
Independent success conditioning	A cross-copy path swap may destroy one copy's separate success event. Joint conditioning is the valid construction in Proposition 4.7.	ST; SX §§3–4

Proposed restriction	What the obstruction actually says	Record
Fixed-background reservations	Missing resident arrows can force a target return for every reservation. Global creator-history changes or different cuts are required.	AX §7
Conditional latent or forest order	Fixed exit forests, fixed loop sectors, and selected restoration fibers can have negative target increments despite positive total normalizers.	RN §§12,17,20–21
General coefficient positivity	Proper-target activation fields, fixed-book inverse target coefficients, entire local inventories, and recurrent-class averages can have the wrong sign.	RN §22; §37
Return-count cone closure	Merges and later-entry sources break the proposed feedback invariant, including from its correct reference law.	RN §28
Uncompensated release numerator	A saved exact negative coefficient rules out dropping determinant compensation. Prefix and forest-degree criteria themselves remain open beyond their proved regimes.	RN (J116); Section 7.18

These are failures of particular strengthenings or proof mechanisms. They are not counterexamples to the original iid, all-active, single-target conjecture.

10.3 Proof boundaries and remaining source gaps

Limits of the sink-strength theorem. Theorem 2.10 recovers the October 20 interpolation proof from the negative-association branch: stationary total count strictly decreases as killing increases along the stated reinjection family. Stationary occupied-set order and arbitrary changes to killing edges are stronger questions, outside that theorem's scope.

Unpackaged complete-graph calculation. An earlier conversation reported a symbolic two-to-three-distinct-particle calculation on homogeneous complete graphs, with positive rational formulas depending on graph size and sleep rate. Its temporary derivation and a maintained verifier were not recovered in the current files. It is therefore a recovery item, not a theorem or independently checked certificate in this document. In particular it must not be paraphrased as an all-input complete-graph result.

Saved but unaudited extensions. The single-leak small-sleep boundary derivation and the final paired-routing reports retain the source's weaker status. They are visible in Sections 2.3 and 7.18. The numerator obstruction in the latter has now been rerun exactly; that check does not audit the separate correlated-routing proofs. Any additional scratch claim without a recoverable statement and argument belongs here, rather than in the theorem list.

11 Sources and verification

This is a website export of the maintained AI-results collection, whose source is `ai-results.tex`. The companion research checkpoint and longer working notes retain proof details, certificate derivations, and decisive failed approaches. These records are

cited by filename below and are not included in this PDF. They are evidence for the collection, not alternative versions of it.

11.1 Source register and attribution

Key	Maintained source	Role
RN	`response-renewal.md`	Main AI research record, §§1–37.
NA-report	Negative-association report, revision `ce9d97d`, "Global reinjection" and "Finite reinjection transform."	AI sink-flow proof, (50)–(52), and Markov evolution, (58). Original proof: 20 October 2026.
DI	`drain-inventory.md`	Finite counting formulations and proved inventory families.
SX	`stack-exploration.md`	Record exchanges and connected-history cancellation.
CX	`creator-exchange.md`	Creator, endpoint, and cycle switches.
AX	`avalanche-exchange.md`	Complete-avalanche injections and inverse cuts.
CB, VA, R	`evidence/notes/complete-blocks.txt`; `vertex-activation.txt`; `reconstruction.txt`	Imported AI proofs of inverse blocks, waking, one-common coupling, and buffers. The last two paths share `evidence/notes/`.
DF	`evidence/notes/defect-factorization.txt`	Target-free one-particle defect factorization.
AF, CR	`evidence/notes/avoidance-flow.txt`; `controlled-routing.txt`	Positive-flow termination and the unresolved feasibility constraint.
BS, LF, LD, FC, CE, WT	`evidence/README.md`	Exact evidence map for earlier block, dependency, contact, and repair arguments.
NC, PR, MC, TC, WC, ST	The same evidence map, with readable notes and original archives	Collision debt, credit, stopping, and success-transform arguments.

Imported evidence archives are immutable. Their import verified bytes and archive integrity; it did not constitute a new mathematical audit. The source index resolves every short key to its readable note and complete original package. Every source annotation in this manuscript refers to that record's section or equation numbering.

The NA-report and its verifier are pinned to revision

`ce9d97d1d157a960718f5061dfb35855457c3e8c`

The cited report passage is lines 11783–11976. The first-proof and composition-audit revisions are, respectively,

`96b6a13c5d77c13b4b7b27b38073fb6e9d7260df`
`c25d383ee52822bfab31a60551fa45997c597eda`

These companion records are not included in this PDF.

The sink-flow proof first appears in commit `96b6a13`, report lines 3197–3278. Its text through (52) is unchanged in the later report revisions reviewed during recovery, with no recorded retraction. The finite-transform composition proof was expanded and its audit marked complete in commit `c25d383` later that day. The marked response, strict count order, mean-square identity, and generator received a further independent AI audit during this compilation. The saturated-stationarity input is credited to Levine–Liang in the source; it is not an AI novelty claim. The negative-association report contains further results beyond the selected sink argument included here; this collection does not purport to reproduce that separate project in full.

The separate Forien manuscript is excluded. Published results used as inputs are credited where used: Abelian-network theory and ordinary IDLA; the Diaconis–Fulton addition algebra; the Levine–Liang ARW/IDLA relation; Fomin's loop-erased switch; directed forest and Markov-loop identities; the cited determinantal-stability, polytope, Lefschetz, Lorentzian, and log-concavity theorems. For example the switch used in Proposition 4.12 is from S. Fomin, *Loop-erased walks and total positivity*, §6. The new claims in this collection are the recorded constructions and applications at their stated scope, without an assertion that every statement is new to the literature.

11.2 Exact verification map

Except for the cited negative-association project files, the files below are relative to the monotonicity folder. Their headers describe dependencies and execution. Most use Python's exact rational arithmetic; the symbolic certificate generators require SymPy. A finite rational audit validates the implemented identities at its fixtures; it does not replace the general proof. A symbolic positivity certificate verifies the whole domain specified by its theorem.

Family	**Verifier or certificate**
Local word and path exchanges	`verify_stack_exploration.py`; `verify_drained_paths.py`
Response and creation laws	`verify_response_renewal.py`; `verify_creation_cone.py`; `verify_root_restoration.py`
Creator potentials and contacts	`verify_creator_potentials.py`; `verify_contact_renewal.py`; `verify_backward_kernels.py`
Three-site occupation	`verify_three_site_certificate.py` and `three-site-positive-certificate.json.gz`; physical check: `verify_three_site_renewal.py`
Pair deficit capacity	`verify_three_site_pair_deficit.py` and `three-site-pair-deficit-certificate.json.gz`; general finite cuts: `verify_pair_capacity.py`
Inventory positivity	`verify_drain_inventory.py`; `verify_single_drain_inventory.py`; `verify_inventory_corners.py`; `verify_cycle_inventory.py`; `verify_three_site_permanent.py`
Directed queues and feedback	`verify_cycle_queue.py`; `verify_cycle_renewal.py`; `verify_backward_queue_cone.py`; `verify_feedback_cone.py`

Family	Verifier or certificate
Small-sleep and boundary calculations	`verify_three_site_rare_sleep.py`; `verify_cycle_rare_sleep.py`; `verify_three_site_boundary.py`
Sleep transport and resident release	`verify_sleep_transport.py`; `verify_release_forest.py`; `verify_cycle_return.py`
Sink reinjection and stationary rank	NA-report's `weak_interaction.py` with `-reinjection-flow`, using that project's `exact.py`; saved `reinjection_flow` certificate.
Markers and filling structures	`verify_marker_geometry.py`; `verify_sleep_valuations.py`
Boundary responses and excursions	`verify_boundary_response.py`; `verify_excursion_dual.py`
Currents and addition geometry	`verify_exit_tilt.py`; `verify_addition_algebra.py`

The principal three-site occupation certificate contains 60 sleep-odds blocks and 330,334 positive coefficients. Its compressed-file SHA-256 is

`5821a322d9ee946df6abbb07534063869706247b8925f93a7b9f621d13587070`

The pair-deficit certificate's compressed-file SHA-256 is

`c6e4487ba10102d45006515e1063233077e76cf3f58ec0fc153c28ae46a9bb67`

These file hashes were checked while assembling the collection. The original symbolic generation remains part of the cited proof record; it was not silently replaced by a hash check.

Checks for this compilation. The six three-cycle killed-return polynomial identities and all $2,232$ Bernstein coefficients were regenerated and checked with SymPy 1.14.0, together with the independent raw kernel. The saved positive boundary interpolation and numerator-shortcut obstruction were rerun using exact rational arithmetic. The four sink-reinjection fixtures were also rerun with SymPy 1.14.0 and reproduce the saved certificate data exactly: a directed looped three-site kernel at $z = 0, 2/5$, and an endpoint path at $z = 0, 1/3$. These check the full derivative, every nontrivial count tail, the mean-square identity, both toppling orders, every addition-stationarity equation, and the finite-transform identities. The broader record's existing audit labels are preserved; this document build is not a fresh independent verification of every theorem. The LaTeX source and its matching PDF are checked for resolved internal references and rendered layout.