AI-GENERATED RESEARCH

Negative dependence in activated random walk

Progress, obstructions, and the remaining conjecture

Generated with ChatGPT from research questions posed by Jacob Richey.
Prepared 8 November 2026.

The main question is whether the particles left sleeping by stationary activated random walk repel one another in the sense of *negative association*. The general conjecture remains open. The project has produced working proofs in several parameter regimes, exact certificates for small graphs, and inequalities on complete graphs of every size. It has also disproved several stronger statements that initially looked like possible routes to a proof.

This note summarizes that record. *Proved* below means that a complete argument or exact certificate is recorded in the project; some arguments have also received separate AI checks. It does not mean independent human verification, and no claim of originality is made. The explanations here are proof sketches.

The question

Take a finite undirected graph with an absorbing sink ∂, reachable from every other vertex. Write V for its nonsink vertices. An active particle jumps at rate one according to simple random walk. When alone, it falls asleep at rate $\lambda > 0$. An arriving particle wakes a sleeper; particles reaching the sink disappear. Add one particle, wait until every remaining particle sleeps, and repeat. With uniform additions, this gives a stationary law μ_λ on $\{0,1\}^V$, where $X_v = 1$ means that vertex v contains a sleeper.

Negative association asks that for disjoint $A, B \subseteq V$ and increasing functions f, g,

$$\mathbb{E}_{\mu_\lambda}[f(X_A)g(X_B)] \le \mathbb{E}_{\mu_\lambda}[f(X_A)]\,\mathbb{E}_{\mu_\lambda}[g(X_B)].$$

An increasing function can count sleepers, test whether a group contains any sleeper, or test whether it contains at least two. Pairwise negative correlation checks only $f = X_i$, $g = X_j$ and is substantially weaker.

The central conjecture is that this inequality holds for *every* finite undirected simple sink graph and every common $\lambda > 0$. An established theorem of Levine and Liang gives a useful starting point: stabilizing one active particle at every vertex already samples μ_λ exactly [1]. This law is invariant under addition at each vertex. Thus the investigation can work directly with stabilization, without estimating a long-time stationary distribution.

What has been established

The following are the main positive conclusions in the working record. Unless specified otherwise, the sleep rate is common to all vertices.

Setting	Conclusion and scope
Directed acyclic routing	The stationary law is a product measure; arbitrary positive site sleep rates are allowed.
Every fixed undirected sink graph	Full NA for sufficiently small sleep rate and for sufficiently large sleep rate. Thresholds depend on the graph.
Weak internal motion	Full NA for internal kernel εP, sufficiently small ε, fixed loopless bidirected support and positive site rates.
Small simple graphs	Full NA at every $\lambda > 0$ for at most four nonsink vertices. Ordinary complete graphs are covered through five nonsink vertices.
Complete graphs of any size	Strict pairwise negative correlation and a stronger block-event inequality, at all positive common rates and common positive weights on the sink edges.
Complete cores with rare killing	Full NA for sufficiently small sink-exit probability, at each fixed size and positive common sleep rate.

The acyclic case has a short explanation. Process vertices in topological order. After routing all surplus particles from a vertex, one particle remains. Its next fresh instruction leaves a sleeper with probability $\lambda_v/(1+\lambda_v)$. There can be no later return. These final decisions are independent. More generally, the stationary law factors over strongly connected components, although this does not simplify a connected undirected graph into acyclic pieces.

The small-graph result is a finite symbolic proof. At four nonsink vertices there are 44 types with connected internal graph, up to relabeling that fixes the sink. For each, exact stationary probabilities reduce all 174 disjoint increasing-event comparisons to polynomial inequalities valid for every positive rate. Disconnected cases factor. This is stronger than checking a selection of numerical rates, but does not extend the classification to arbitrary graph size.

On complete graphs, the stronger all-size statement is

$$\operatorname{Cov}\left(\mathbf{1}\left\{\sum_{i\in A} X_i > 0\right\}, g(X_B)\right) \le 0$$

for disjoint blocks and every increasing g. It controls whether a block contains *any* sleeper. Full NA also requires tests such as *at least two sleepers in this block*, whose general comparison is open.

Rare killing and weak internal motion are different limits. The former uses an internal kernel $(1-\varepsilon)P$ with closed P and small ε; the latter uses εP, so most jumps leave immediately. Neither result settles all intermediate parameters.

Where the proofs get their signs

At small sleep rate, the system is usually empty after stabilization. Let K be the killed walk's internal transition matrix and $G = (I - K)^{-1}$ its Green matrix: $G(i, j)$ is the expected number of visits to j before absorption, starting at i and including the visit at time zero. For a fixed set A,

$$\mathbb{P}_{\mu_\lambda}(X_v = 1 \text{ for every } v \in A) = \lambda^{|A|} \det G_A + O(\lambda^{|A|+1}),$$

where G_A is the principal submatrix indexed by A. Green-matrix inequalities therefore control the first terms in the dependence inequalities. The formula is asymptotic; it does not identify the full stationary law as a determinantal process.

At large sleep rate, it is more useful to study the holes, $Y_v = 1 - X_v$. Their leading probabilities are governed by escape probabilities of the underlying walk. For example, $e_A(v)$ is the probability, starting from $v \in A$, of reaching the sink before the next visit to A. The leading vacancy coefficients can be represented by volumes of regions built from these escape probabilities. Region inclusion supplies the desired inequalities.

This argument has a real degeneracy: if every route from one vertex to the sink passes through another, a leading gap can vanish. The general large-rate proof needs the next-order correction at such separating vertices. Merely being close to the fully occupied configuration would not prove NA. Both the small-rate and large-rate results are statements about each fixed graph, with no uniform threshold in graph size.

For complete graphs, Junge, Kaufman and Meisel give a representation using a binomial-refresh process and a descending boundary [2]. Building on that representation, an auxiliary process with three site states retains enough information to compare occupied sets. The project's couplings and stationarity identities yield the all-size pair and block-event inequalities. A comparison with a process traced onto fewer sites is useful, but its stationary law is not the original law's marginal.

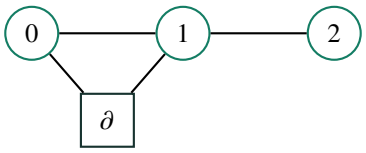
The rare-killing argument expands probabilities of small vacancy sets. Fix an irreducible stochastic kernel P invariant under every vertex permutation and a positive common sleep rate. For $(1 - \varepsilon)P$, the leading gaps can be signed for all disjoint increasing events, proving full NA at small ε. Strict inequalities also give a relative neighborhood of the reference kernel, rate vector and uniform killing profile on which one sufficiently small killing interval works.

There is also a gluing theorem for a generally directed transition-kernel extension. Attach a weakly killed complete core to finitely many recipient sites with fixed positive sleep rates and no transitions between recipients. Each recipient jump returns to the core with the same small probability b, using a common entrance distribution; otherwise it reaches the sink. Core-to-recipient routing fractions are independent of the core exit vertex, nonnegative, and sum to at most one. Full NA holds when b is sufficiently small. For each fixed sufficiently small killing rate, one bound works over all these routing fractions and common entrance distributions. It depends on the core, killing rate, recipient number and sleep rates; it is not uniform as killing tends to zero. Separate return probabilities or entrance distributions lie outside the audited theorem.

What stronger statements get wrong

A natural strategy is to place the stationary law in a familiar class of negatively dependent measures. The *strong Rayleigh* property, defined by stability of the generating polynomial, would imply NA and additional closure properties [3]. The weaker Rayleigh property requires pairwise negative correlation to survive every positive external field, which reweights a configuration by $\prod_v z_v^{X_v}$.

Both stronger properties fail for stationary ARW. The obstruction can already be seen on three nonsink vertices:



The sink is adjacent to vertices 0 and 1. All three nonsink vertices have the same sleep rate.

On this graph, an exact expansion gives

$$\mathrm{Cov}(X_1, X_2 \mid X_0 = 0) = \frac{4}{1215\lambda^2} + O(\lambda^{-3}) > 0 \qquad (\lambda \text{ sufficiently large}).$$

The unconditioned law is nevertheless NA at every positive rate. At $\lambda = 64$, the positive field $z_0 = 1/100$, $z_1 = z_2 = 1$ also gives a positive tilted covariance. Consequently, a proof of the general conjecture cannot establish Rayleigh or conditional NA along the way. Conditional negative correlation can fail on trees as well.

Stationarity is another essential qualification. On a two-site path with both sites adjacent to the sink, start from one prescribed sleeper. At rate $\lambda = 1$, one uniformly located addition followed by stabilization gives covariance $31/3042 > 0$. Thus the dynamics do not preserve NA from arbitrary NA initial laws. A stationary-law proof needs more than an induction asserting preservation by each addition.

Even abelianity has a precise boundary. It permits different legal toppling orders for fixed *site* stacks [4]. It does not allow a revealed sleep instruction to be carried away and tried again. On the triangle with vertices x, y, ∂, start one active particle at each nonsink vertex and use stack prefixes

$$x : (\text{sleep}, x \to y, x \to \partial), \qquad y : (y \to x, y \to \partial).$$

Topple y first. The sleep at x fails because two particles are present; both particles subsequently leave. Moving that sleep to y, immediately after the jump $x \to y$, instead leaves a sleeper there. Fresh, unexposed randomness may be assigned to particles with the correct law, but fixed particle tapes do not inherit site-stack abelianity.

The remaining work

The largest gap is still the original one: full NA on arbitrary undirected sink graphs at intermediate sleep rates. The endpoint arguments do not give a continuation principle between their parameter ranges. The counterexamples above explain why several standard stronger properties cannot supply one.

Complete graphs provide a focused intermediate target. Full NA at every size and rate remains open, despite all-size pair and block-OR inequalities. Exact reductions express the missing threshold comparisons through stationary queue quantities. Their signs require averaging with the actual source weights: some individual histories, entrances, or terms contribute with the wrong sign. Proving each term positive is therefore unavailable.

A related two-copy approach compares how the events of a count process are allocated between two systems. A preserved ordering has been proved for one restricted class of event sequences. Extending it requires a property of the profiles actually generated by the process; several plausible general shape conditions have exact counterexamples. These failures eliminate particular induction schemes, rather than disproving the desired comparison.

For an irreducible reversible closed kernel P and a fixed common sleep rate, replace P by $(1-\varepsilon)P$. Its stationary law has an expansion $\nu_\varepsilon(V \setminus I) = \varepsilon^{|I|} c_I + O(\varepsilon^{|I|+1})$ for the probability of the exact sleeper set $V \setminus I$. Already the universal leading pair inequality $c_{\{i\}} c_{\{j\}} \geq c_{\{i,j\}}$, for distinct i, j, remains open; equality would leave higher-order terms unresolved. Some natural inequalities imposed before averaging are false even on a three-site path, although the actual pair gap there is positive.

Avalanches offer another possible route to gluing. If X and Y are successive stationary configurations and N counts particles, the loss is $L = 1 + N(X) - N(Y)$. On complete graphs, one target is that this preceding loss decreases stochastically as $N(Y)$ increases. Ordinary monotonicity of the reversed rank chain is too weak: the needed comparison also bounds its change by one. The desired upper bound remains open.

How to read the evidence. The project combines analytic arguments with exact local calculations. Rational stationary equations, alternative toppling orders, and symbolic inequalities provide checks with explicit scopes. Numerical agreement at a few sizes is kept separate from an all-size argument. Several expensive coefficient calculations did not produce a general proof; current work emphasizes structural questions and small tests capable of rejecting a specific claim. Failed implications are retained in a compact working record and supplied to subsequent agents. Separate AI checks reduce some errors, but do not replace human mathematical review.

References

[1] L. Levine and F. Liang. *Exact sampling and fast mixing of Activated Random Walk*. arXiv:2110.14008. The saturated sampler is an input from this work.

[2] M. Junge, H. Kaufman and J. Meisel. *Scaling limit and density conjecture for activated random walk on the complete graph*. arXiv:2604.04747. The complete-graph representation is an input from this work.

[3] J. Borcea, P. Brändén and T. M. Liggett. *Negative dependence and the geometry of polynomials*. arXiv:0707.2340.

[4] L. T. Rolla and L. Tournier. *Non-fixation for Biased Activated Random Walks*. arXiv:1507.04732. Sections 2 and 4 distinguish site-wise and particle-wise constructions.