AIMS Mathematics

https://www.aimspress.com/journal/Math

AIMS Mathematics, 11(9): 29907–29933.
DOI: 10.3934/math.20261187
Received: 26 May 2026
Revised: 10 August 2026
Accepted: 09 September 2026
Published: 15 September 2026

Research article

Finite-horizon competition of branching random walks on trees and sparse random graphs: Size-biased genealogy and annealed simulation diagnostics

Yunzhi Zhu, Saisai Hou and Sen Zhang*

Department of Public Basic Courses, Nanjing University of Industry Technology, Nanjing 210023, Jiangsu, China

* **Correspondence:** Email: szhang@niit.edu.cn.

Abstract: We study a finite-horizon first-arrival competition generated by two independent branching random walks on regular tree exhaustions and sparse Erdős–Rényi graphs. The paper focuses on a variance-sensitive mechanism that is invisible at the level of first moments. If two offspring laws have the same mean but different second factorial moments $\eta = \mathbb{E}[K(K-1)]$, then their many-to-one occupation fields agree, while their second moments and the genealogies selected by size-biased sampling can differ. We make this distinction explicit through a size-biased spine calculation: Under the associated size-biased measure, the expected number of non-spine sibling subtrees per spine generation is η/m. For any target set fixed before the branching realization, we further derive an exact finite-time η-response theorem: Changing η by $\Delta\eta$ changes the target second moment by $\Delta\eta\,\Gamma_T(A)$, where $\Gamma_T(A)$ is an explicit nonnegative random-walk kernel, and changes the terminal-size-weighted survivor target occupation by $(\Delta\eta/m)P^T(a,A)\sum_{\ell=0}^{T-1} m^\ell$. These identities quantify both intermittency and size-biased enrichment without identifying the spine law with ordinary survival conditioning. A finite-time Paley–Zygmund bound records how unconditional hitting, second moments, and survival probabilities interact. The computational protocol is designed accordingly. The Erdős–Rényi experiment is annealed over independent graph instances, red and blue start from an exchangeable pair of neighbours of a high-degree root, and graph-level cluster bootstrap and paired graph-level comparisons are used for sparse-graph summaries. The numerical results are reported as finite-size diagnostics, not as an infinite-volume coexistence theorem. They support the following picture: Fixed-mean high variance is not uniformly advantageous, but after survival filtering its surviving genealogies can have a larger observed contribution near the outer frontier.

Keywords: branching random walk; first-arrival competition; size-biased spine; offspring variance; survival conditioning; Erdős–Rényi graph; annealed simulation; regular tree
Mathematics Subject Classification: 05C81, 60F10, 60J80, 60K35, 65C05

1. Introduction

Branching random walks (BRWs) combine reproduction with spatial motion. In the two-colour first-arrival model considered here, a red cloud and a blue cloud evolve independently on the same finite graph, and each vertex receives the colour of the cloud that visits it first. The model has a transparent simulation structure, but its frontier statistics are sensitive to the genealogy behind the first visits. A cloud with high offspring variance may die out more often; when it survives, however, it may do so through unusually prolific lineages that place many descendants near the observed frontier. The main question in this paper is how to separate this survival-filtering effect from a genuine first-order speed advantage.

Let K be the offspring variable, let $m = \mathbb{E}K > 1$, and set

$$\eta = \mathbb{E}[K(K-1)].$$

At fixed m, the many-to-one formula gives the same first-moment occupation field for all offspring laws. Thus a first-moment calculation cannot distinguish, for example, Poisson, low-variance, and high-variance reproduction when their means agree. The distinction appears in the second moment and in the size-biased genealogy obtained when a lineage is sampled proportionally to its descendants. If K^* denotes the size-biased offspring number, then

$$\mathbb{E}[K^* - 1] = \frac{\eta}{m}.$$

The quantity η/m is the expected number of non-spine sibling subtrees produced at a spine generation. We turn this qualitative observation into an exact finite-time response formula: For a fixed target A, both the target second moment and its terminal-size-weighted survivor occupation have explicit slopes in η. This is the correct quantitative level at which to discuss the finite-horizon reversal seen in the simulations: High variance is penalized unconditionally through extinction and intermittency, whereas terminal-size weighting enriches genealogies with many side branches.

A point of terminology is important. Conditioning a Galton–Watson process on survival to a fixed time is not the same object as the classical size-biased spine measure. The latter samples a distinguished lineage with weights proportional to the number of descendants, while the former only restricts to a survival event. The two viewpoints are related and often lead to the same intuition, but they should not be conflated. For this reason, the present paper uses the spine calculation as a mechanism and diagnostic, not as a theorem asserting monotone frontier dominance under survival conditioning.

This distinction also changes the role of the numerical study. The purpose is not merely to show that a simulation code runs without obvious implementation errors. Rather, the code is used to test a finite-size genealogical hypothesis under a protocol that is statistically defensible. We therefore use an annealed Erdős–Rényi (ER) protocol rather than a single fixed graph instance: Many independent random graph instances are generated, several competitions are run on each instance, and graph-to-graph variability is separated from branching randomness. Moreover, red and blue start from an exchangeable pair of neighbours of a high-degree root, rather than from a root-versus-neighbour pair that would create a structural colour advantage.

The article has four contributions. First, Section 3.2 gives the exact relation between the spine and survivor laws, proves an explicit finite-horizon η-response theorem for fixed targets, and records a survival-conditioned Paley–Zygmund bound. Second, Section 3.3 records exact finite-graph identities used to validate the branching and colouring rules. Third, Section 4.1 gives the annealed statistical protocol, including graph-level cluster bootstrap and paired graph-level effects for the Erdős–Rényi experiment. Fourth, Section 4 reports finite-horizon diagnostics on regular tree exhaustions and sparse random graphs. The paper deliberately stops short of an infinite-volume competition theorem; the simulations are used to formulate precise open problems rather than to replace such a theorem.

2. Literature review

The one-type branching random walk is classical. The works of Biggins [1], Hammersley [2], and Kingman [3] established foundational speed and subadditive ideas. Later developments clarified the fine structure of the extremal particles: Addario-Berry and Reed [4] obtained tight estimates for minima, Hu and Shi [5] proved second-order almost-sure results connected with critical martingales, Aidekon [6] proved convergence in law of the centred minimum, and Madaule [7] described the extremal process seen from the tip. Heavy-tail and large-deviation regimes have also been revisited in [8, 9].

The proof technology behind these results includes martingales, many-to-one changes of measure, and spinal decompositions. The size-biased calculation in Section 3.2 is a finite-time version of this viewpoint; see Lyons and Peres [10], Biggins and Kyprianou [11], and the lecture notes of Shi [12] for broader background.

Competition processes form a second line of work. The two-type Richardson model and related first-passage percolation processes are surveyed by Deijfen and Haggstrom [13]. Competition on $\mathbb{Z}^d$ driven by branching random walks was studied by Deijfen and Vilkas [14]. These results motivate the present finite-size diagnostics, although we do not prove an infinite-volume theorem here.

Branching random walks on general graphs require additional care. Generating-function methods are discussed by Bertacchi and Zucca [15], while their recent generating-function order gives survival comparisons [16]. Dussaule, Wang, and Yang [17] studied branching random walks on relatively hyperbolic groups, and Makarova et al. [18] analysed two-type branching random walks on multidimensional lattices. Large deviations of random walks on Erdős–Rényi graphs were studied by Coghi, Morand, and Touchette [19]. Standard network references include [20–22].

Table 1 compares the most directly relevant DOI-registered work from 2023–2026 with our finite-horizon study.

Table 1. Comparison with recent related work.

Reference	Setting	Method	Distinction from this paper
Deijfen–Vilkas (2023)	Competing BRWs on $\mathbb{Z}^d$	Infinite-volume probability	Coexistence theory; no sparse-graph simulation.
Bertacchi–Zucca (2026)	BRWs on general graphs	Generating-function order	Rigorous survival order rather than first-arrival competition.
Dussaule–Wang–Yang (2025)	Relatively hyperbolic groups	Green functions and boundary geometry	Trace growth and limit sets rather than two-colour competition.
Dyszewski et al. (2023)	One-dimensional BRW	Heavy-tail extremes	Asymptotic maxima rather than finite-graph first arrivals.
This paper	Trees and sparse ER graphs	Exact finite-time η-response and annealed diagnostics	Survival filtering with cluster uncertainty and matched-$P(K = 0)$ controls.

3. Model and finite-horizon analysis

3.1. Model and observables

Let $G = (V, E)$ be a finite connected simple graph. We write $u \sim v$ when $\{u, v\} \in E$ and use the simple-random-walk transition matrix

$$P(u, v) = \begin{cases} 1/\deg(u), & u \sim v, \\ 0, & \text{otherwise.} \end{cases}$$

Red starts from a vertex a, and blue starts from a vertex b. Before colouring, the two clouds evolve independently. A particle at u produces K children, and each child jumps independently according to $P(u, \cdot)$. Let $Z_n^c(v)$ be the number of particles of colour $c \in \{r, b\}$ at v in generation n, and let $|Z_n^c| = \sum_v Z_n^c(v)$.

For a root o, define the reached maximum

$$H_n^c = \max\{\mathrm{dist}(o, v) : Z_k^c(v) > 0 \text{ for some } 0 \leq k \leq n\},$$

with the convention that an empty maximum is $-\infty$. The arrival time of colour c at v is

$$\tau_c(v) = \inf\{n \geq 0 : Z_n^c(v) > 0\}.$$

The infinite-horizon first-arrival colour is

$$C(v) = \begin{cases} r, & \tau_r(v) < \tau_b(v), \\ b, & \tau_b(v) < \tau_r(v), \\ r \text{ or } b \text{ with probability } 1/2, & \tau_r(v) = \tau_b(v) < \infty, \\ \varnothing, & \tau_r(v) = \tau_b(v) = \infty. \end{cases}$$

The finite-horizon colour $C_T(v)$ used in the simulations is obtained by replacing $\tau_c(v)$ by the truncated information available up to generation T: Vertices unreached by both colours by time T remain uncoloured and are not counted as ties. For a frontier set $A \subseteq V$, the red fraction is

$$\rho(A) = \frac{\#\{v \in A : C(v) = r\}}{\#\{v \in A : C(v) \in \{r, b\}\}},$$

whenever the denominator is positive. In the simulations, A is the two-layer outer observed annulus.

The principal notation is collected in Table 2.

Table 2. Principal notation.

Symbol	Meaning	Symbol	Meaning
$Z_n^c(v)$	colour-c population at v	m	offspring mean
$\tau_c(v)$	first arrival time	η	$\mathbb{E}[K(K-1)]$
H_n^c	reached radial maximum	s_T	$\mathbb{P}(Z_T > 0)$
$\rho(A)$	red fraction on A	M, L	graphs and runs per graph

The offspring laws used below are Poisson(m); a low-variance deterministic mixture $K = \lfloor m \rfloor + \xi$ with $\xi \sim$ Bernoulli($m - \lfloor m \rfloor$); and a negative-binomial law with mean m and prescribed variance $\sigma^2 > m$. The negative-binomial parameterisation is the one used by the code: If $p = m/\sigma^2$ and $r = mp/(1-p) = m^2/(\sigma^2 - m)$, then $K \sim \mathrm{NB}(r, p)$ has mean m and variance σ^2. For all offspring laws in the paper,

$$\eta = \mathbb{E}[K(K-1)] = \mathrm{Var}(K) + m^2 - m.$$

Indeed, $\mathbb{E}[K(K-1)] = \mathbb{E}[K^2] - \mathbb{E}K$, while $\mathbb{E}[K^2] = \mathrm{Var}(K) + (\mathbb{E}K)^2$ and $\mathbb{E}K = m$. Thus the displayed identity converts offspring variance at fixed mean into the second factorial moment used below; it is not merely a restatement of $\mathbb{E}K = m$. At $m = 1.2$, the low-variance mixture has $\eta = 0.4$, the Poisson law has $\eta = 1.44$, and the negative-binomial law with variance 2.4 has $\eta = 2.64$. We also use a matched-extinction control with $P(K = 0) = 0.3$: The low-dispersion probabilities are $(0.3, 0.2, 0.5)$ on $\{0, 1, 2\}$, and the high-dispersion probabilities are $(0.3, 0.45, 0, 0.25)$ on $\{0, 1, 2, 3\}$. Both have mean 1.2 and the same one-step extinction probability, but variances 0.76 and 1.26 (equivalently $\eta = 1.00$ and 1.50).

3.2. Survival-conditioned variance mechanism

This section isolates the mechanism that guides the simulations. The results are elementary, but they connect the numerical observations with the genealogy of a surviving branching process.

Proposition 3.1 (First-moment invariance at fixed mean)**.** *Let (X_n) be the simple random walk on G started from a. For a one-colour branching random walk with offspring mean m,*

$$\mathbb{E}Z_n(v) = m^n P^n(a, v) = m^n \mathbb{P}_a(X_n = v).$$

Consequently, two offspring laws with the same mean m have the same expected occupation field at every time.

Proof. Conditioning on the generation-n configuration gives

$$\mathbb{E}[Z_{n+1}(v) \mid Z_n] = m \sum_{u \in V} Z_n(u) P(u, v).$$

Taking expectations and iterating from $Z_0 = \delta_a$ gives the formula. Only the mean m enters the recursion. □

Proposition 3.2 (Size-biased variance amplification)**.** *Let K have mean $m > 0$ and second factorial moment $\eta = \mathbb{E}[K(K-1)]$. Let K^* be the size-biased offspring number,*

$$\mathbb{P}(K^* = k) = \frac{k\mathbb{P}(K = k)}{m}, \qquad k \geq 0.$$

Then

$$\mathbb{E}[K^* - 1] = \frac{\eta}{m}.$$

Thus, along a lineage selected proportionally to its descendants, the expected number of sibling subtrees produced per spine generation is η/m.

Proof. By definition,

$$\mathbb{E}K^* = \sum_{k\geq 0} k\frac{k\mathbb{P}(K=k)}{m} = \frac{\mathbb{E}K^2}{m}.$$

Since $\mathbb{E}K^2 = \mathbb{E}[K(K-1)] + \mathbb{E}K = \eta + m$, we obtain

$$\mathbb{E}[K^* - 1] = \frac{\eta + m}{m} - 1 = \frac{\eta}{m}.$$

In the usual spine interpretation, the child continuing the distinguished line accounts for one offspring, and the remaining children start side subtrees. Hence the expected number of side branches is $\mathbb{E}[K^* - 1]$. □

Proposition 3.3 (Exact relation between spine and survivor laws)**.** *Let $\mathcal{F}_T$ be the genealogy through generation T, $s_T = \mathbb{P}(Z_T > 0)$, and define the finite-time spine measure by*

$$\left.\frac{d\mathbb{Q}_T}{d\mathbb{P}}\right|_{\mathcal{F}_T} = \frac{Z_T}{m^T}.$$

If $\mathbb{P}_T^{\mathrm{surv}} = \mathbb{P}(\,\cdot \mid Z_T > 0)$, then on $\{Z_T > 0\}$,

$$\frac{d\mathbb{Q}_T}{d\mathbb{P}_T^{\mathrm{surv}}} = \frac{s_T Z_T}{m^T}.$$

Thus the spine law is exactly the survivor law further size-biased by terminal population size.

Proof. Since $d\mathbb{P}_T^{\mathrm{surv}}/d\mathbb{P} = \mathbf{1}_{\{Z_T>0\}}/s_T$, division of the two Radon–Nikodym derivatives gives the identity. Its survivor-law expectation is one because $\mathbb{E}Z_T = m^T$. □

For any terminal target A, many-to-one also gives the exact identity

$$\mathbb{E}[N_T(A) \mid Z_T > 0] = \frac{m^T \mathbb{P}_a(X_T \in A)}{s_T}.$$

At fixed mean and motion kernel, the conditional occupation differs only through finite-time survival. This supplies the requested rigorous bridge without identifying ordinary survival conditioning with the spine law.

Write $P^r(x, A) = \sum_{v\in A} P^r(x, v)$. The following theorem gives an exact effect size for the second factorial moment.

Theorem 3.4 (Exact finite-horizon η-response)**.** *Let a one-colour branching random walk start from one particle at a, let $T \geq 1$, and let $A \subseteq V$ be fixed before the branching realization. Assume $m > 0$ and $\eta < \infty$, and define*

$$\Gamma_T(A) = \sum_{j=0}^{T-1} m^{2T-j-2} \sum_{x\in V} P^j(a, x)[P^{T-j}(x, A)]^2.$$

Then

$$\mathbb{E}[N_T(A)^2] = m^T P^T(a, A) + \eta\, \Gamma_T(A).$$

Moreover, the terminal-size-weighted survivor occupation is

$$\mathcal{B}_T(A) := \frac{\mathbb{E}[Z_T N_T(A) \mid Z_T > 0]}{\mathbb{E}[Z_T \mid Z_T > 0]} = P^T(a, A)\left(1 + \frac{\eta}{m}\sum_{\ell=0}^{T-1} m^\ell\right).$$

Consequently, for two offspring laws with the same mean and second factorial moments $\eta_2 > \eta_1$,

$$\mathbb{E}_2[N_T(A)^2] - \mathbb{E}_1[N_T(A)^2] = (\eta_2 - \eta_1)\Gamma_T(A),$$

and

$$\mathcal{B}_{T,2}(A) - \mathcal{B}_{T,1}(A) = \frac{\eta_2 - \eta_1}{m} P^T(a, A)\sum_{\ell=0}^{T-1} m^\ell.$$

Here, $\sum_{\ell=0}^{T-1} m^\ell = (m^T - 1)/(m-1)$ for $m \neq 1$ and equals T for $m = 1$.

Proof. For the first identity, decompose $N_T(A)^2$ into its diagonal term and ordered pairs of distinct terminal particles. The diagonal contribution is $\mathbb{E}N_T(A) = m^T P^T(a, A)$. Every ordered distinct pair has a unique generation $j \in \{0, \dots, T-1\}$ at which its two ancestral lines split. The expected number of generation-j parents at x is $m^j P^j(a, x)$, and one such parent produces η ordered pairs of distinct children on average. After the two children jump, their descendant subtrees evolve independently. Each subtree has expected target population $m^{T-j-1} P^{T-j}(x, A)$ at time T. Multiplying these factors and summing over j and x gives $\eta\Gamma_T(A)$.

For the second identity, use the spine representation of $\mathbb{Q}_T$. The distinguished terminal particle contributes $P^T(a, A)$ in expectation. At spine generation j, the expected number of non-spine siblings is η/m. A sibling born there has expected time-T target descendants $m^{T-j-1} P^{T-j}(X_j, A)$. Since the spine position follows P, averaging over X_j gives

$$\mathbb{E}_{\mathbb{Q}_T}\left[P^{T-j}(X_j, A)\right] = P^T(a, A).$$

Summing the spine contribution and all side-subtree contributions yields the displayed formula for $\mathcal{B}_T(A)$ because

$$\mathbb{E}_{\mathbb{Q}_T} N_T(A) = \frac{\mathbb{E}[Z_T N_T(A)]}{m^T} = \frac{\mathbb{E}[Z_T N_T(A) \mid Z_T > 0]}{\mathbb{E}[Z_T \mid Z_T > 0]}.$$

Subtracting the formulas for η_1 and η_2 proves the two response identities. □

Remark 3.5. Both coefficients are explicit and nonnegative, and they are positive whenever $P^T(a, A) > 0$. The first identity quantifies the increase in intermittency that enters the Paley–Zygmund denominator; the second quantifies the enrichment of a fixed target under terminal-size weighting of survivors. These are distinct, opposing finite-time effects. For an annealed graph experiment, the theorem applies conditionally on the sampled graph to any graph-dependent target that is fixed before the branching runs. Ordinary survival conditioning, first-arrival colouring, and the random reached annulus depend on the full offspring generating functions and on both colours, so the theorem does not imply a universal sign for the conditional outer-colour fraction. It supplies the quantitative mechanism that can legitimately be transferred to the simulation diagnostics.

Proposition 3.6 (Survival-conditioned hitting lower bound)**.** *Let $N_t(A) = \sum_{v\in A} Z_t(v)$ be the number of particles in a target set A at generation t, and let $Z_t = \sum_v Z_t(v)$ be the total population. Assume that $\mathbb{E}[N_t(A)^2] < \infty$, $\mathbb{E}N_t(A) > 0$, and $\mathbb{P}(Z_t > 0) > 0$. Then*

$$\mathbb{P}\{N_t(A) > 0 \mid Z_t > 0\} \geq \frac{(\mathbb{E}N_t(A))^2}{\mathbb{E}[N_t(A)^2]\mathbb{P}(Z_t > 0)}.$$

Proof. Since $N_t(A) > 0$ implies $Z_t > 0$,

$$\mathbb{P}\{N_t(A) > 0 \mid Z_t > 0\} = \frac{\mathbb{P}\{N_t(A) > 0\}}{\mathbb{P}(Z_t > 0)}.$$

Paley–Zygmund gives

$$\mathbb{P}\{N_t(A) > 0\} \geq \frac{(\mathbb{E}N_t(A))^2}{\mathbb{E}[N_t(A)^2]}.$$

Combining the two displays proves the claim. □

Remark 3.7. No independence assumption is used: $\{N_t(A) > 0\} \subseteq \{Z_t > 0\}$ gives the displayed ratio exactly. The lower bound cannot exceed one because its Paley–Zygmund factor is at most $\mathbb{P}(N_t(A) > 0) \leq \mathbb{P}(Z_t > 0)$. This is a one-colour fixed-target diagnostic, not a theorem for the random two-colour outer annulus.

3.3. Finite-graph validation identities

The following identities are used as finite-horizon checks. They are not presented as the main theoretical contribution; their role is to ensure that the simulation preserves the branching-random-walk state space rather than collapsing it into a first-visit exploration.

Let $f(s) = \mathbb{E}s^K$. For $A \subseteq V$, define

$$q_t^A(x) = \mathbb{P}_x\{\text{no descendant visits } A \text{ during generations } 0, \dots, t\}.$$

Then

$$q_0^A(x) = \mathbf{1}_{\{x\notin A\}}, \qquad q_{t+1}^A(x) = \begin{cases} 0, & x \in A, \\ f\left(\sum_y P(x,y)q_t^A(y)\right), & x \notin A. \end{cases}$$

This dynamic program gives exact finite-time hitting probabilities. In particular, for $A = \{v\}$, $\mathbb{P}_x(\tau_A > t) = q_t^A(x)$ and, for $t \geq 1$, $\mathbb{P}_x(\tau_A = t) = q_{t-1}^A(x) - q_t^A(x)$, with the time-zero case read directly from the starting point. It also gives the expected colour profile because the red and blue arrival fields are independent before colouring. If $a_{r,t}(v) = \mathbb{P}_a(\tau_r(v) = t)$ and $Q_{b,t}(v) = \mathbb{P}_b(\tau_b(v) > t)$, then at horizon T,

$$\mathbb{P}\{C_T(v) = r\} = \sum_{t=0}^{T} a_{r,t}(v)\left(Q_{b,t}(v) + \frac{1}{2}a_{b,t}(v)\right),$$

with the analogous expression for blue.

The second-moment recursion displays the role of η. Let $\mu_n(u) = \mathbb{E}Z_n(u)$ and $S_n(u,v) = \mathbb{E}[Z_n(u)Z_n(v)]$. Then

$$\begin{aligned} S_{n+1}(u,v) = m^2 &\sum_{x,y} P(x,u)P(y,v)S_n(x,y) \\ &+ \sum_x \mu_n(x)\{mP(x,u)\mathbf{1}_{\{u=v\}} + (\eta - m^2)P(x,u)P(x,v)\}. \end{aligned}$$

Increasing η at fixed m adds the positive forcing term $\Delta\eta \sum_x \mu_n(x)P(x,u)P(x,v)$. Iterating the linear recursion propagates it through $P \otimes P$, so the covariance field is entrywise monotone in η for a common initial state. Summing the appropriate entries and unrolling the recursion gives the exact coefficient $\Gamma_T(A)$ in Theorem 3.4; thus the response is quantified rather than used only as an implementation check.

3.4. A regular-tree envelope

Let $\mathbb{T}_d$ be the infinite rooted d-ary tree. Away from the root, the radial coordinate of a simple random walk has increments $+1$ with probability $p_d = d/(d+1)$ and -1 with probability $q_d = 1/(d+1)$. Its log moment generating function is

$$\Lambda_d(\theta) = \log(p_d e^{\theta} + q_d e^{-\theta}),$$

with Cramer rate function

$$I_d(x) = \sup_{\theta \in \mathbb{R}} \{\theta x - \Lambda_d(\theta)\}.$$

For $x \in (-1, 1)$,

$$I_d(x) = \frac{1+x}{2} \log \frac{1+x}{2p_d} + \frac{1-x}{2} \log \frac{1-x}{2q_d}.$$

For a branching random walk started at a fixed level, define the generation-n maximum

$$M_n = \max\{\mathrm{dist}(o, v) : Z_n(v) > 0\},$$

with $M_n = -\infty$ if generation n is empty. The biased-walk rate function gives the following envelope calculation. By the many-to-one formula and Markov's inequality,

$$\mathbb{P}(M_n \geq nx) \leq m^n\, \mathbb{P}(R_n \geq nx),$$

where R_n is the exact reflected radial chain: $0 \to 1$ deterministically, while at levels $j \geq 1$, it moves outward/inward with probabilities p_d, q_d. Its distribution is obtained by a finite recursion, so this is a rigorous finite-n bound. The corresponding Chernoff envelope is

$$m^n e^{-nI_d(x)}$$

as a large-distance heuristic. On a depth-D exhaustion, the finite and infinite chains couple exactly through time T if the start level plus T is below D. Here, $1 + 10 = 11 < 20$, so the leaf boundary has exactly zero influence.

Remark 3.8. The reached maximum $H_n = \max_{k \leq n} M_k$ can differ from the generation-n front. The exact reflected-chain bound is finite-time; the rate-function envelope is not used to claim an asymptotic speed or infinite-volume coexistence.

4. Simulation design and numerical results

4.1. Annealed statistical protocol

The regular-tree experiment uses an implicit rooted 3-ary tree of depth 20. Red and blue start from two different children of the root, giving an exact colour-exchanging automorphism in the matched case. The tree terminal generation is $T_{\text{tree}} = 10$, while the sparse-graph terminal generation is $T_{\text{ER}} = 15$. Since the starting vertices are at level one, the largest possible level reached by generation 10 is 11, so the artificial depth-20 boundary is not touched by the tree experiment.

The sparse graph experiment is annealed. For each graph index $g = 1, \ldots, M$, an independent Erdős–Rényi graph $G_g \sim G(n, 8/n)$ is generated, its largest connected component is retained, and a maximum-degree vertex is used only as a radial reference root. Red and blue are then started from two distinct neighbours of this root, sampled uniformly without replacement. This neighbour-versus-neighbour design is exchangeable in colour labels. It keeps the finite-time exploration centred near a stable hub without giving one colour the root degree advantage.

For an observable $Y_{g,r}$ measured on graph g and run r, write

$$\bar{Y}_g = \frac{1}{L}\sum_{r=1}^{L} Y_{g,r}, \qquad \bar{Y} = \frac{1}{M}\sum_{g=1}^{M} \bar{Y}_g.$$

The reported Erdős–Rényi confidence intervals use graph-level cluster bootstrap: Graph-level means $\bar{Y}_g$ are resampled, rather than treating all ML runs as independent. Effects relative to the matched Poisson baseline are computed as graph-paired differences whenever both case and baseline are available on the same graph. The descriptive variance decomposition uses

$$MS_B = \frac{L}{M-1}\sum_{g=1}^{M}(\bar{Y}_g - \bar{Y})^2,$$

and

$$MS_W = \frac{1}{M(L-1)}\sum_{g=1}^{M}\sum_{r=1}^{L}(Y_{g,r} - \bar{Y}_g)^2.$$

We estimate $\sigma_G^2 = \max\{(MS_B - MS_W)/L, 0\}$, $\sigma_R^2 = MS_W$ and their intraclass correlation. Graph-cluster bootstrap gives an interval for σ_G^2, and balanced random reassignment of run values to graph labels tests a graph-homogeneity null. For event $S_{g,r}$, the annealed conditional estimator is $\widehat{\theta_S} = \sum Y_{g,r}\mathbf{1}_{S_{g,r}} / \sum \mathbf{1}_{S_{g,r}}$. The cluster bootstrap resamples every graph's numerator–denominator pair, including zero-survivor graphs, and therefore avoids post-selection.

The reported 95% cluster-bootstrap intervals are the 0.025 and 0.975 empirical quantiles from 2500 graph resamples. For an unconditional case–baseline comparison, let $d_g = \bar{Y}_g^{\text{case}} - \bar{Y}_g^{\text{base}}$; the paired sign-flip test independently multiplies the d_g by Rademacher signs and uses the add-one two-sided permutation p-value. For the graph-homogeneity test, the ML run values are randomly reassigned to M balanced graph labels of size L, $\widehat{\sigma_G^2}$ is recomputed, and the same add-one rule is used. Conditional case–baseline effects are differences of annealed ratio estimators under paired graph-cluster resampling, rather than sign-flip tests, because their survivor denominators can differ by case.

Seven cases are compared: The original five cases plus the two orientations of the matched-$P(K = 0)$ pair. The latter isolates dispersion from the extinction-probability contrast between the original negative-binomial and deterministic-mixture laws.

The production command uses 1200 tree repetitions and 80 independent sparse graphs with 40 competitions per graph, 2500 bootstrap resamples, and 2500 sign-flip resamples. Within each graph, the red/blue assignment to the sampled neighbour pair alternates, eliminating persistent start-vertex imbalance.

4.2. Numerical results

The production experiment contains 30,800 competitions: 8,400 on the regular-tree exhaustion and 22,400 on annealed Erdős–Rényi graph instances. The numerical goal is not merely to estimate a single win probability. We compare unconditional and survival-conditioned frontier fractions, joint survival, lead trajectories, and random-graph variance components.

4.2.1. Calibration and frontier fractions

Tables 3 and 4 show the main diagnostics and the matched-Poisson colour-exchangeability calibration. On the tree, the starts are exchanged by an exact automorphism; in the sparse graph experiment, alternating the start assignment makes the graph-level design colour-balanced. The unconditional matched-Poisson estimates are 0.499 in both families. Conditional estimates are less precise because joint survival is uncommon, but their graph bootstrap retains all 80 graph clusters. The speed control moves red upward, confirming that the statistic detects an intentional mean perturbation.

Table 3. Main finite-horizon diagnostics. Outer red is unconditional; cond. outer red is conditioned on both colours being alive at the family-specific terminal generation. ER intervals use graph-cluster bootstrap. Conditional ER means are annealed ratio estimators that retain zero-survivor graph clusters in resampling.

Family	Case	Runs	Graphs	Outer red	Cond. outer red	Both alive	Cond. graphs
ER annealed	matched_poisson	3200	80	0.499 [0.485,0.512]	0.484 [0.459,0.509]	318/3200	80
ER annealed	speed_advantage_red	3200	80	0.592 [0.576,0.606]	0.673 [0.654,0.690]	467/3200	80
ER annealed	highvar_red_vs_poisson	3200	80	0.421 [0.410,0.433]	0.580 [0.546,0.611]	206/3200	80
ER annealed	highvar_red_vs_lowvar	3200	80	0.165 [0.155,0.174]	0.753 [0.740,0.766]	601/3200	80
ER annealed	lowvar_red_vs_highvar	3200	80	0.818 [0.809,0.828]	0.251 [0.238,0.265]	671/3200	80
ER annealed	matched_p0_high_red_vs_low	3200	80	0.453 [0.439,0.467]	0.511 [0.492,0.531]	390/3200	80
ER annealed	matched_p0_low_red_vs_high	3200	80	0.532 [0.519,0.545]	0.475 [0.452,0.498]	444/3200	80
Tree	matched_poisson	1200	1	0.499 [0.474,0.521]	0.464 [0.393,0.538]	144/1200	1
Tree	speed_advantage_red	1200	1	0.540 [0.517,0.566]	0.551 [0.494,0.608]	186/1200	1
Tree	highvar_red_vs_poisson	1200	1	0.391 [0.369,0.414]	0.476 [0.383,0.559]	86/1200	1
Tree	highvar_red_vs_lowvar	1200	1	0.138 [0.120,0.155]	0.606 [0.556,0.653]	254/1200	1
Tree	lowvar_red_vs_highvar	1200	1	0.835 [0.815,0.854]	0.308 [0.259,0.358]	263/1200	1
Tree	matched_p0_high_red_vs_low	1200	1	0.463 [0.439,0.487]	0.464 [0.403,0.534]	183/1200	1
Tree	matched_p0_low_red_vs_high	1200	1	0.570 [0.548,0.594]	0.502 [0.428,0.571]	152/1200	1

Table 4. Matched-Poisson colour-exchangeability calibration. The swap gap is $2\rho - 1$, where ρ is the red fraction; values near zero indicate agreement with red-blue label exchangeability, while finite conditional deviations are interpreted descriptively.

Family	Runs	Graphs	Outer red	Outer swap gap	Cond. outer red	Cond. swap gap
ER annealed	3200	80	0.499 [0.485,0.512]	-0.003 [-0.030,0.025]	0.484 [0.457,0.507]	-0.033 [-0.086,0.015]
Tree	1200	1	0.499 [0.475,0.523]	-0.003 [-0.049,0.046]	0.464 [0.393,0.535]	-0.071 [-0.214,0.071]

Table 5 reports the case–baseline contrasts, standardized effects, and permutation results relative to matched Poisson reproduction.

Table 5. Effects relative to the matched Poisson baseline. For ER, unconditional differences and permutation p-values are graph-paired; conditional ER differences are differences of annealed survivor-ratio estimators. Tree effects are run-level diagnostics on the fixed symmetric graph. Permutation values use the add-one correction.

Family	Case	Uncond. diff.	Cond. diff.	Std. effect	Perm. p
ER annealed	speed_advantage_red	0.093	0.189	1.066	< 0.001
ER annealed	highvar_red_vs_poisson	-0.077	0.096	-1.053	< 0.001
ER annealed	highvar_red_vs_lowvar	-0.333	0.270	-4.315	< 0.001
ER annealed	lowvar_red_vs_highvar	0.320	-0.233	4.376	< 0.001
ER annealed	matched_p0_high_red_vs_low	-0.045	0.028	-0.522	< 0.001
ER annealed	matched_p0_low_red_vs_high	0.033	-0.008	0.394	0.002
Tree	speed_advantage_red	0.041	0.087	0.097	0.021
Tree	highvar_red_vs_poisson	-0.107	0.012	-0.266	< 0.001
Tree	highvar_red_vs_lowvar	-0.360	0.142	-0.973	< 0.001
Tree	lowvar_red_vs_highvar	0.336	-0.156	0.876	< 0.001
Tree	matched_p0_high_red_vs_low	-0.036	0.000	-0.085	0.035
Tree	matched_p0_low_red_vs_high	0.072	0.038	0.171	0.002

Throughout the result figures, "matched" denotes the matched-Poisson calibration, "speed+" an increased red offspring mean, "high/Pois" high-variance red versus Poisson blue, and "high/low" and "low/high" the two orientations of the high- and low-variance fixed-mean comparison. The prefixes "P0 high/low" and "P0 low/high" denote the corresponding orientations under matched $P(K = 0)$.

The matched-$P(K = 0)$ comparison is substantially weaker than the original negative-binomial versus deterministic-mixture contrast. On ER graphs, the conditional high-dispersion-red estimate is 0.511 with 95% interval $[0.492, 0.531]$; on the tree, it is 0.464 with interval $[0.403, 0.534]$. We therefore interpret the dramatic original reversal as a combined dispersion–extinction phenomenon and do not claim a universal variance advantage.

Figures 1 and 2 are the central empirical evidence. In the high-variance red versus low-variance blue case, the unconditional outer red fraction is small. This does not mean that the high-variance process is intrinsically slower at fixed mean. It means that many high-variance red populations are extinct, small, or delayed by the finite observation time. Conditional on joint survival, however, the same high-variance red law has a substantially larger observed outer-annulus fraction in the finite sample. The reversed low-versus-high case produces the opposite pattern. This is consistent with the finite-horizon signature suggested by the size-biased mechanism in Proposition 3.2.

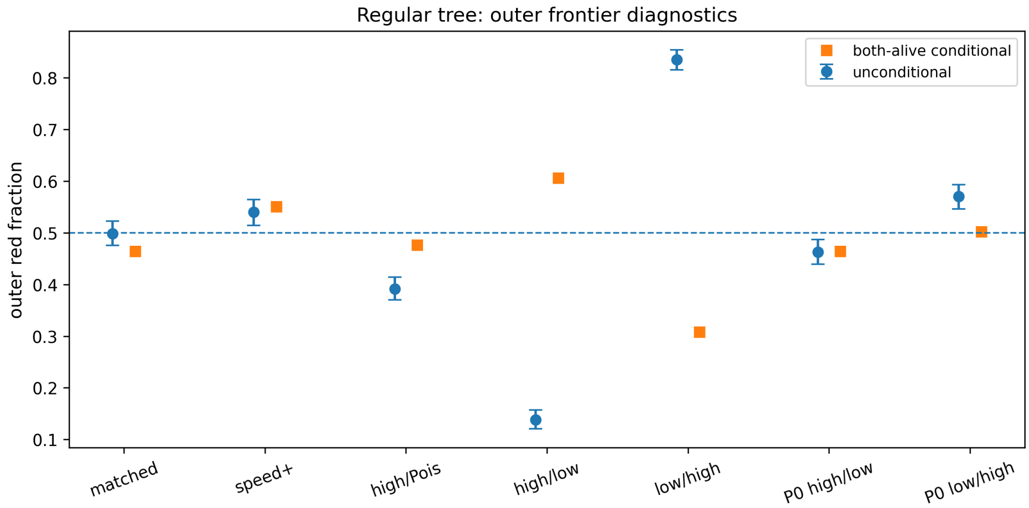

Figure 1. Regular-tree outer red fraction. Circles are unconditional means with bootstrap confidence intervals; squares are both-alive conditional means. The high-variance fixed-mean cases show a clear separation between unconditional and conditional behaviour.

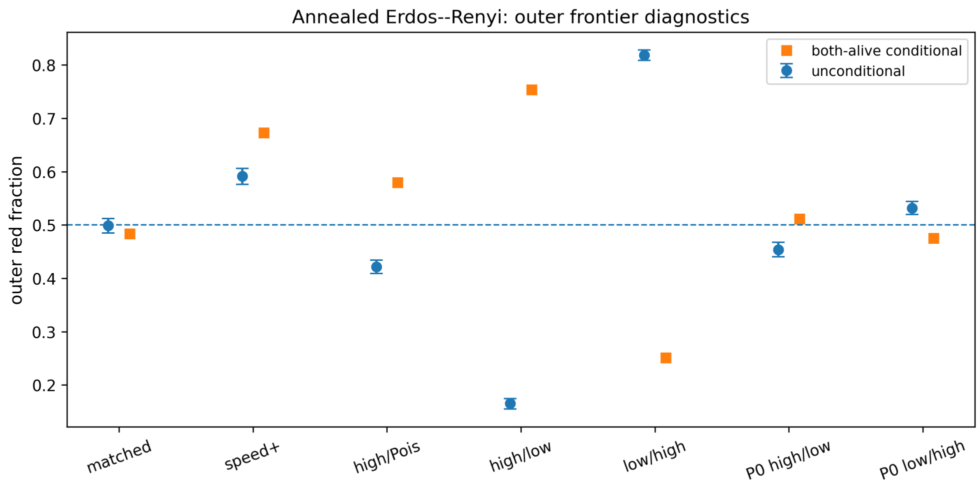

Figure 2. Annealed Erdős–Rényi outer red fraction. Each estimate averages over independent graph instances and independent branching-random-walk randomness. The qualitative reversal between unconditional and conditional statistics persists after graph resampling.

4.2.2. Survival filtering

The survival plots in Figures 3 and 4 make clear why a conditional statistic is not a cosmetic re-labelling of the same sample. The high-variance process generates more extreme genealogies: Many runs are weak, but a minority of runs are very productive. Conditioning on joint survival therefore selects a different effective population of genealogies. In terms of the related spine calculation, sampling a lineage proportionally to its descendants replaces K by K^* and creates, on average, η/m sibling subtrees per spine generation. This does not by itself prove a conditional-frontier theorem, but it gives the genealogical quantity that the conditional diagnostics are designed to detect.

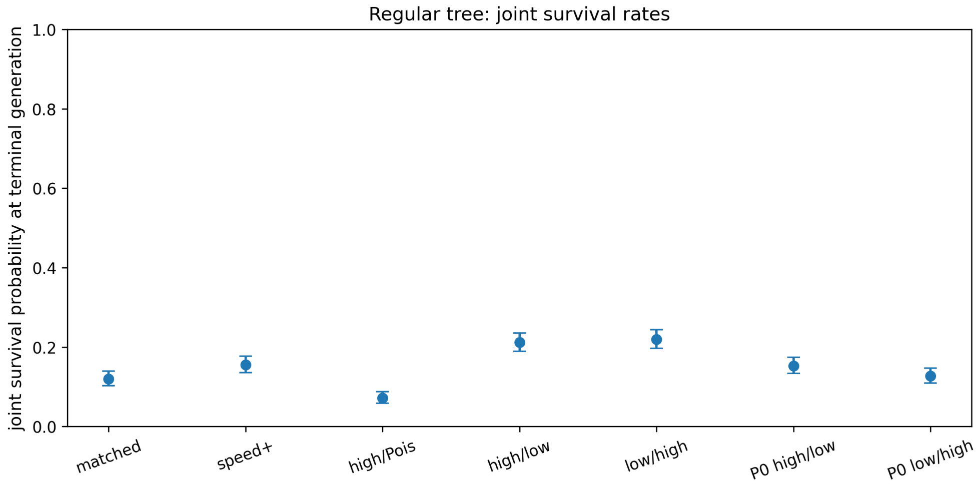

Figure 3. Joint survival probabilities at the terminal generation in the regular-tree experiment. High-variance regimes have visibly different survival profiles, which explains why unconditional and conditional diagnostics should not be conflated.

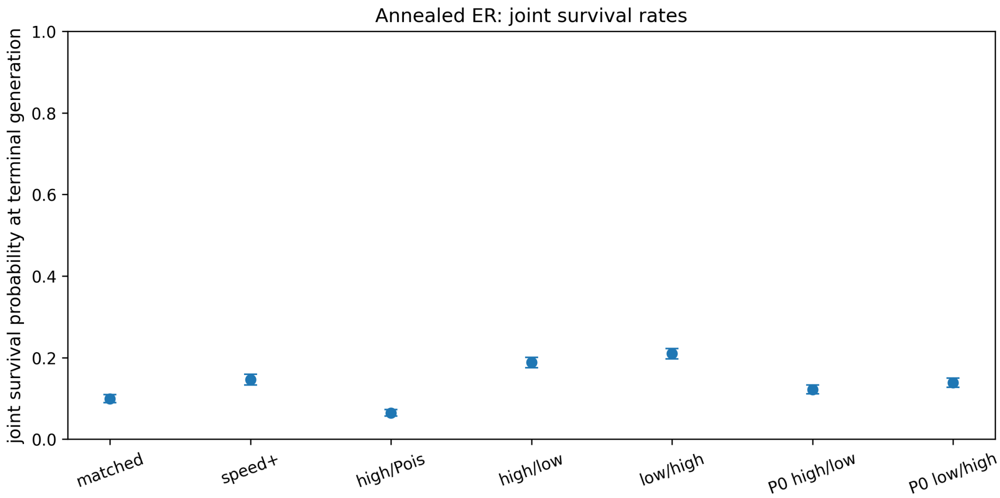

Figure 4. Joint survival probabilities at the terminal generation in the annealed Erdős–Rényi experiment. The survival rates are estimated jointly over graph and branching randomness.

Figure 5 summarises the survival-filtering effect. The sign and magnitude of the gap depend on which colour carries the high-variance law. This diagnostic is useful because it isolates the part of the observed colour bias that is caused by conditioning, rather than by the unconditional mixture of extinction, delay, and prolific growth.

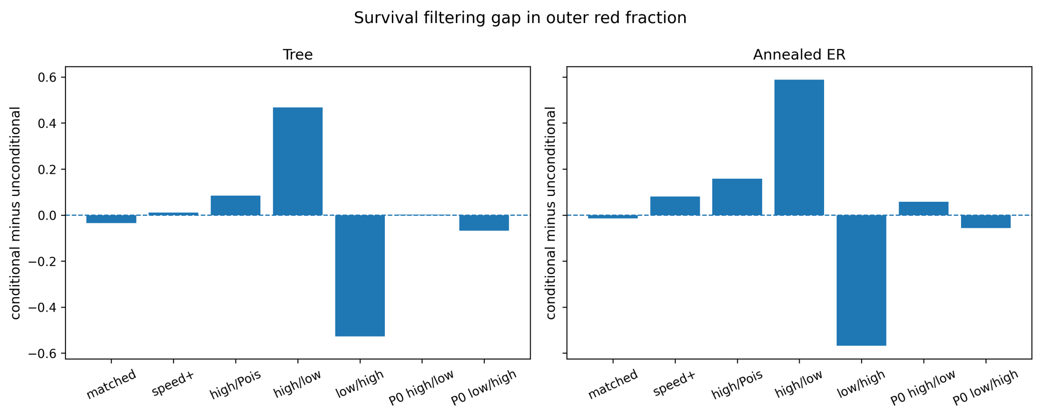

Figure 5. Difference between survival-conditioned and unconditional outer red fractions. Positive values indicate that conditioning on joint survival favours red relative to the unconditional sample; negative values indicate that conditioning favours blue.

4.2.3. Lead trajectories and population trajectories

The lead trajectories in Figures 6 and 7 complement the colour fractions. They measure the reached radial front rather than the final colour of an annulus. The high-versus-low and low-versus-high fixed-mean comparisons are almost mirror images, indicating that the variance perturbation changes the finite-time distribution of the frontier even when the reproduction mean is fixed. These lead curves are not asymptotic speeds; they are finite-horizon diagnostics designed to detect genealogical intermittency.

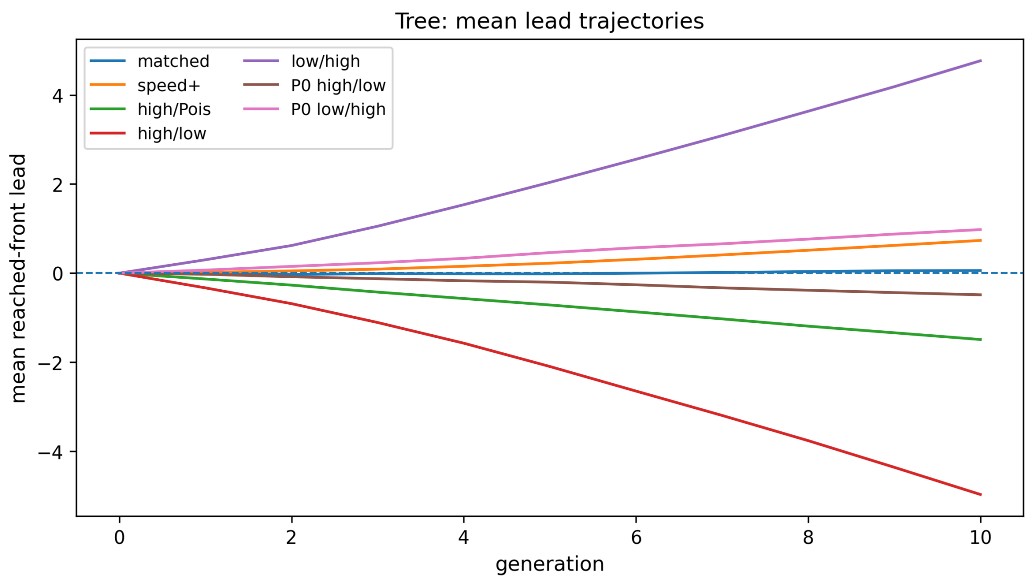

Figure 6. Mean reached-front lead on the regular tree. A positive value means that the red-reached maximum is larger than the blue-reached maximum. The fixed-mean high-versus-low cases show strong finite-time asymmetry even though the means are equal.

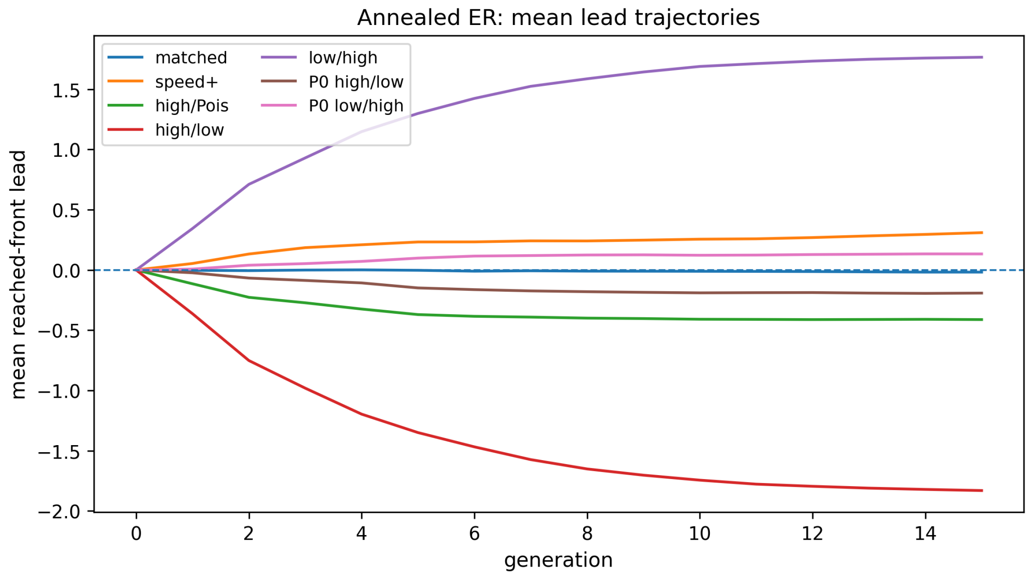

Figure 7. Mean reached-front lead in the annealed Erdős–Rényi experiment. The trajectories average over independent graph instances and repeated competitions on each graph.

Figures 8 and 9 report the corresponding mean population trajectories for selected cases on the regular tree and in the annealed ER experiment, respectively.

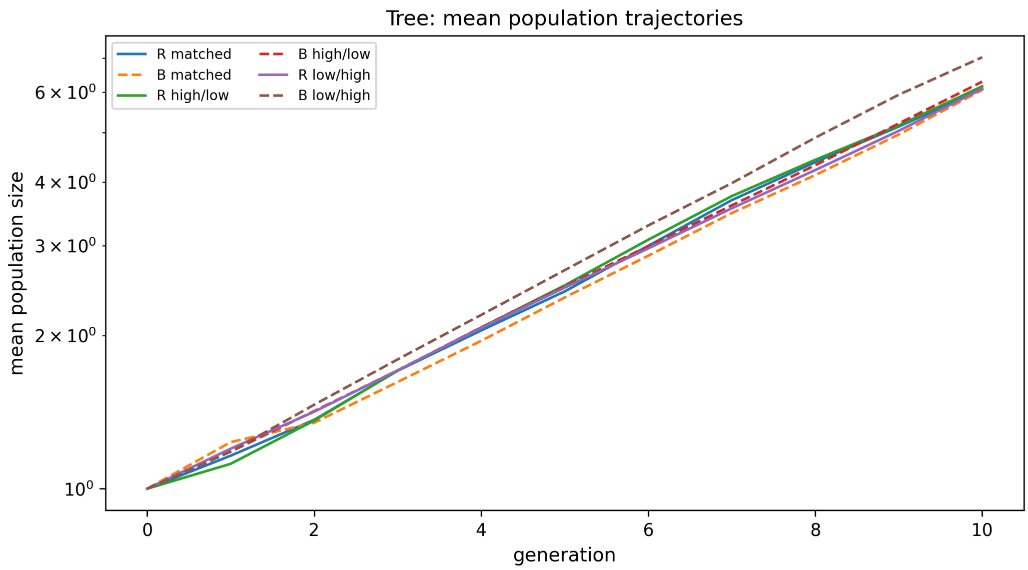

Figure 8. Mean population trajectories on the regular tree for selected cases. The logarithmic vertical scale highlights how high-variance laws produce more variable finite-time population fields even when the mean is fixed.

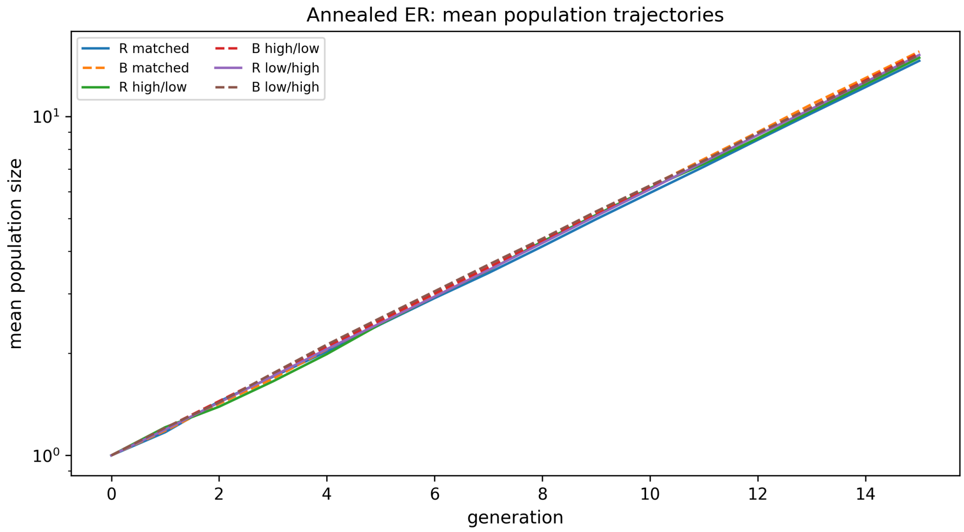

Figure 9. Mean population trajectories in the annealed Erdős–Rényi experiment. These curves are averaged over graph instances and branching randomness.

4.2.4. Random-graph variability and empirical scaling

The decomposition in Table 6 and Figure 10 directly addresses the statistical weakness of a single-graph experiment. Sparse random graphs have local degree fluctuations, short cycles, dangling trees, and small bottlenecks, any of which can influence a finite-time front. In the annealed protocol, those features are resampled rather than conditioned on a selected graph. The estimated between-graph component is small and is truncated at zero in most cases; its bootstrap intervals and the homogeneity randomization p-values therefore do not support a claim that graph-to-graph variability is dominant. The inferential conclusions remain annealed and graph-cluster based.

The scaling plot in Figure 11 condenses the fixed-mean variance comparisons into a single horizontal variable. A larger value of $\log(\eta_r/\eta_b)$ means that the red reproduction law creates more expected side branches under the size-biased spine. The observed association with the conditional red outer fraction is consistent with Proposition 3.2. Because there are only finitely many parameter pairs, the scatterplot is read only as a visual finite-size summary, not as a universal regression model.

Table 6. Random-effects analysis of variance (ANOVA) decomposition for the annealed Erdős–Rényi outer red fraction. Between-graph variance is $\max\{(MS_B - MS_W)/L, 0\}$, and within-graph variance is MS_W.

Case	Between var. (95% CI)	Within var.	Total var.	ICC	Homog. p	L
matched_poisson	0.0000 [0.0000,0.0012]	0.1576	0.1575	0.000	1.000	40
speed_advantage_red	0.0007 [0.0000,0.0021]	0.1587	0.1593	0.004	0.129	40
highvar_red_vs_poisson	0.0000 [0.0000,0.0007]	0.1427	0.1425	0.000	1.000	40
highvar_red_vs_lowvar	0.0000 [0.0000,0.0004]	0.0933	0.0931	0.000	1.000	40
lowvar_red_vs_highvar	0.0000 [0.0000,0.0002]	0.0984	0.0980	0.000	1.000	40
matched_p0_high_red_vs_low	0.0000 [0.0000,0.0014]	0.1577	0.1577	0.000	0.456	40
matched_p0_low_red_vs_high	0.0000 [0.0000,0.0005]	0.1551	0.1545	0.000	1.000	40

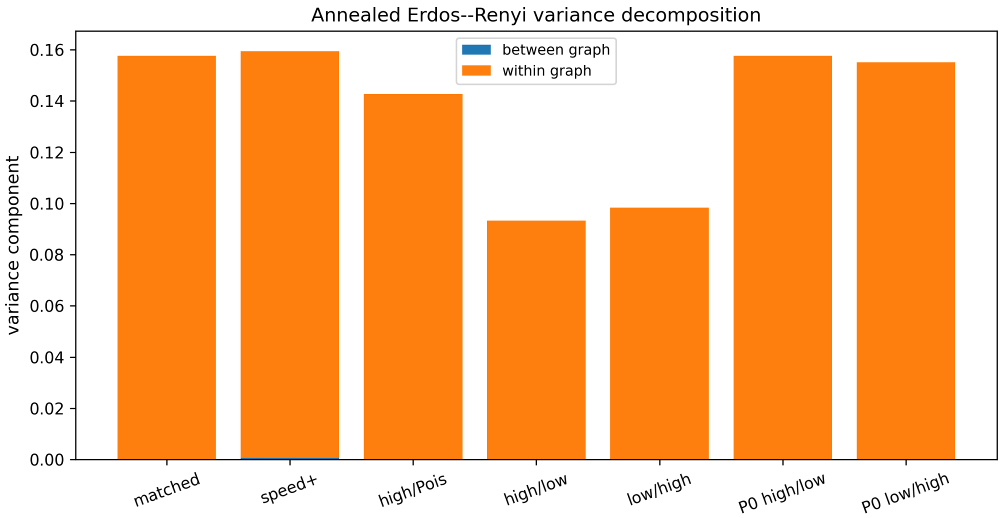

Figure 10. Variance decomposition for the annealed Erdős–Rényi experiment. The between-graph component measures the variability of graph-level means, while the within-graph component measures branching randomness on a fixed graph.

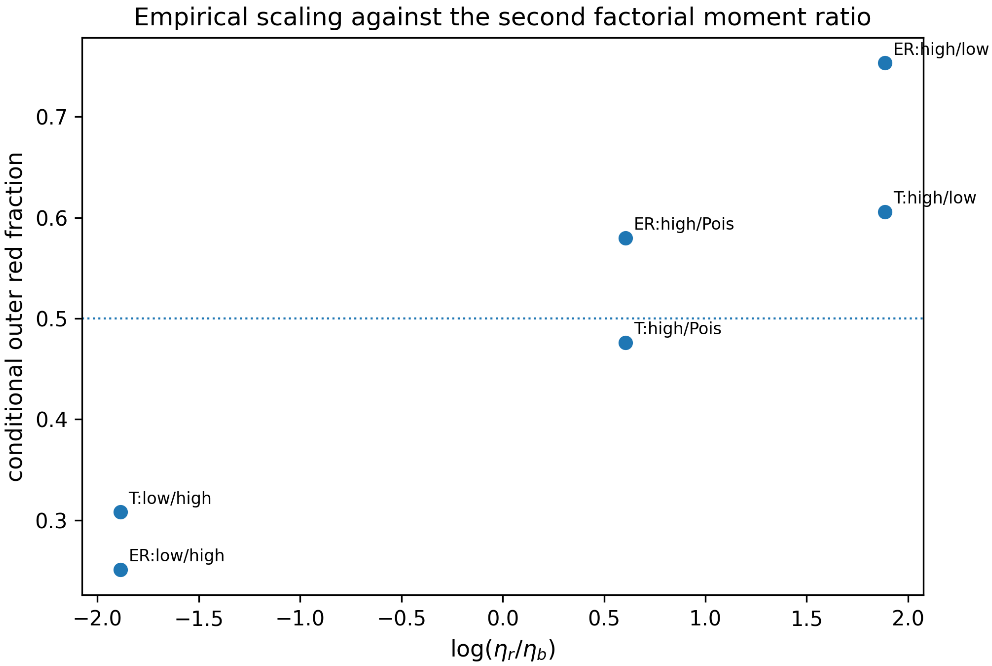

Figure 11. Empirical scaling of the survival-conditioned outer red fraction against $\log(\eta_r/\eta_b)$ for the fixed-mean variance comparisons. This is a descriptive finite-size summary based on a small finite collection of parameter-pair/family points only, not an asymptotic theorem or a systematic regression design.

Figures 12 and 13 summarise the total red fraction among all coloured vertices on the regular tree and in the annealed ER experiment. They complement the outer-annulus diagnostics by showing the balance over the whole coloured set.

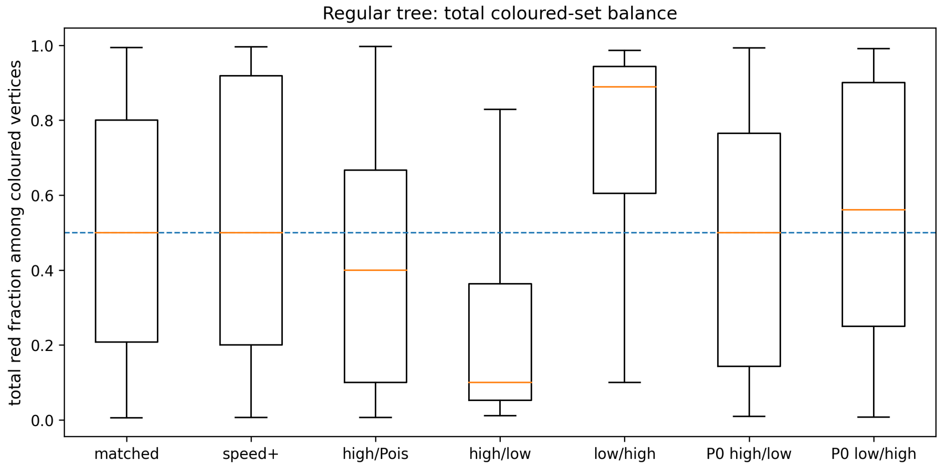

Figure 12. Total red fraction among coloured vertices on the regular tree. This global statistic is less sensitive to the frontier than the outer annulus, but it helps distinguish local frontier effects from whole-coloured-set bias.

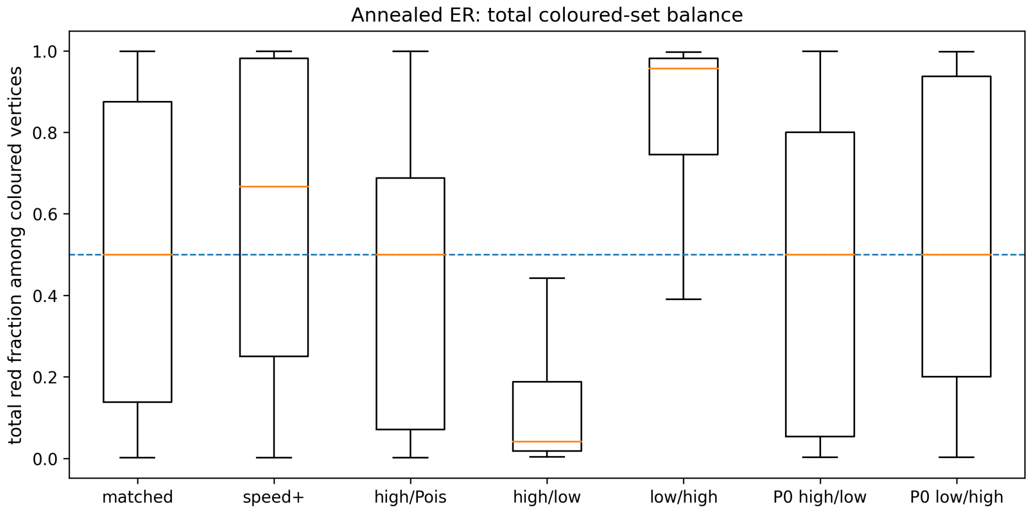

Figure 13. Total red fraction among coloured vertices in the annealed Erdős–Rényi experiment. The comparison with the outer-annulus statistics shows that frontier effects are stronger than global coloured-set effects.

Together, the numerical results support a coherent finite-size interpretation. High offspring variance at fixed mean produces a more intermittent population. The unconditional statistics penalize this intermittency through extinction and delayed growth, whereas survival-conditioned statistics select genealogies that have already passed this filter. The strong reversal appears in the original high/low contrasts, but it is attenuated and not consistently signed under the matched-$P(K = 0)$ control. Accordingly, the simulations support a parameter-dependent survival-filtering diagnostic, not a universal variance advantage; the size-biased calculation supplies a structural interpretation rather than a conditional-frontier theorem.

4.2.5. Sensitivity analysis

Table 7 and Figure 14 vary one factor at a time: Offspring mean and upper-support dispersion for both graph families; tree degree and horizon; and ER size, average degree, and horizon. The matched-$P(K = 0)$ pair is used throughout. The estimates vary in magnitude, as expected for a finite-horizon statistic, but the analysis makes clear which conclusions are stable and which are parameter-dependent.

Table 7. One-factor-at-a-time sensitivity analysis for the matched-$P(K = 0)$ comparison (high-dispersion red versus low-dispersion blue). All unspecified parameters equal their main-design values.

Family	Factor	Level	Cond. outer red	Joint survival
Tree	mean	1.05	0.420	0.064
Tree	mean	1.2	0.481	0.126
Tree	mean	1.35	0.495	0.224
ER	mean	1.05	0.569	0.037
ER	mean	1.2	0.551	0.127
ER	mean	1.35	0.519	0.225
Tree	upper_support	3	0.481	0.126
Tree	upper_support	4	0.472	0.112
Tree	upper_support	5	0.461	0.084
ER	upper_support	3	0.551	0.127
ER	upper_support	4	0.571	0.103
ER	upper_support	5	0.588	0.080
Tree	tree_degree	2	0.547	0.110
Tree	tree_degree	3	0.481	0.126
Tree	tree_degree	4	0.489	0.126
Tree	tree_horizon	8	0.500	0.138
Tree	tree_horizon	10	0.481	0.126
Tree	tree_horizon	12	0.524	0.104
ER	er_n	450	0.504	0.128
ER	er_n	900	0.551	0.127
ER	er_n	1800	0.533	0.118
ER	er_avg_degree	6	0.516	0.137
ER	er_avg_degree	8	0.551	0.127
ER	er_avg_degree	10	0.538	0.115
ER	er_horizon	10	0.550	0.143
ER	er_horizon	15	0.551	0.127
ER	er_horizon	20	0.556	0.118

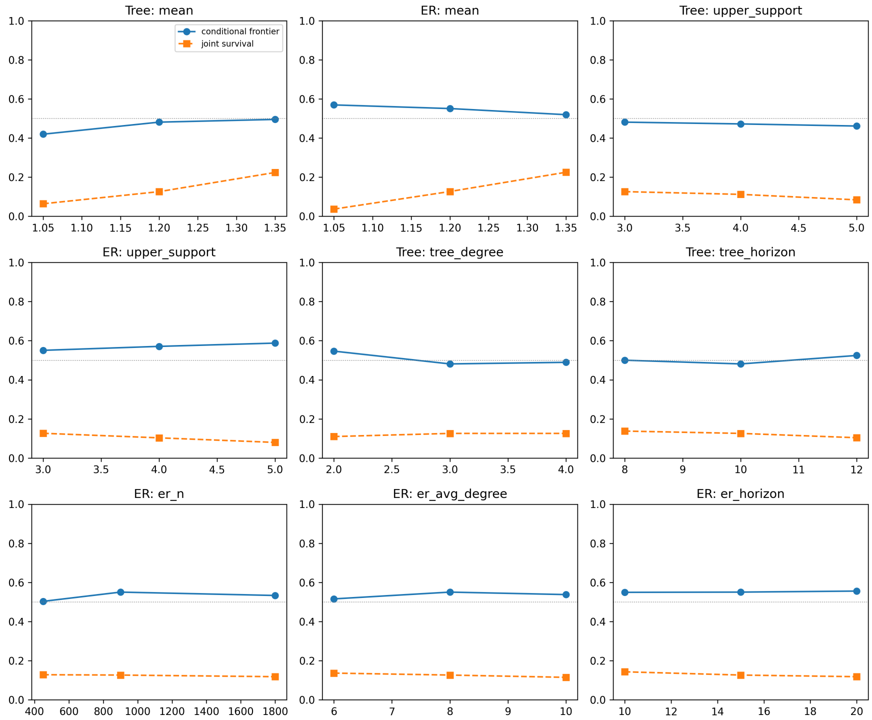

Figure 14. One-factor-at-a-time sensitivity diagnostics. Solid circles show the joint-survival conditional outer red fraction; dashed squares show joint survival.

5. Conjectures and open problems

The numerical results suggest the following finite-size working hypothesis.

Conjecture 5.1 (Finite-horizon scaling hypothesis). Fix $m > 1$ and two laws with the same m and $P(K = 0)$ but $\eta_2 > \eta_1$. Let $D_T > T + 1$, start the colours at two children of the root of the depth-D_T d-ary tree, and let A_T be the two outermost reached layers by time T. Along a sequence $T \to \infty$ for which joint survival has probability bounded away from zero, the conditional contrast

$$\mathbb{E}[\rho(A_T) \mid Z_T^r Z_T^b > 0] - \tfrac{1}{2}$$

has the sign of $\eta_r - \eta_b$ eventually. The corresponding unconditional contrast need not have that sign.

Because the available parameter-pair/family points are too few for model assessment, we do not fit or propose a logit-linear scaling law. Establishing such a relation would require a purpose-designed parameter grid, uncertainty propagation, residual diagnostics, and out-of-sample validation.

A rigorous infinite-volume theorem would require more than the finite identities used here. One would need two-sided arrival-time bounds for the relevant graph family, survival-conditioned control of near-front genealogies, a comparison theorem for two dependent first-arrival fields, and a way to convert near-front particle estimates into statements about coloured vertex sets. For sparse random graphs, one must also decide whether the theorem is quenched, annealed, or valid with high probability over the graph ensemble.

6. Discussion

The validation identities are useful, but the central message of the paper is the size-biased genealogical mechanism for survival-filtered finite samples. At fixed mean, high variance is not simply stronger; it is more intermittent. Theorem 3.4 quantifies this statement: A change $\Delta\eta$ raises the fixed-target second moment by exactly $\Delta\eta\,\Gamma_T(A)$, while it raises the terminal-size-weighted survivor target occupation by the explicit positive amount $(\Delta\eta/m)P^T(a,A)\sum_{\ell=0}^{T-1} m^{\ell}$. The first effect penalizes unconditional concentration, whereas the second enriches prolific survivor genealogies. The theorem therefore supplies a quantitative mechanism for the competing finite-time tendencies without claiming that ordinary survival conditioning is identical to the spine measure. The matched-$P(K = 0)$ control and the sensitivity grid show that the magnitude and even the finite-sample sign of the nonlinear two-colour statistic can depend on the remaining parameters. We therefore attribute the strongest original reversal to a combined dispersion–extinction contrast and do not claim a universal dominance result.

The annealed Erdős–Rényi experiment gives the random graph part a clear statistical status. The results average over a graph ensemble and quantify the remaining graph-to-graph variability. The sample sizes are finite and can be increased by running the supplied code; the protocol is aligned with the probabilistic distinction between quenched and annealed dynamics.

7. Conclusions

Finite-horizon competition separates two effects of offspring variability that are not visible in first moments. At a fixed offspring mean, the expected occupation field is unchanged, whereas the second factorial moment controls the exact fixed-target second-moment response and the terminal-size-weighted survivor occupation quantified in Theorem 3.4. These identities distinguish size-biased enrichment from ordinary survival conditioning; they do not establish a universal ordering of survival-conditioned first-arrival competition.

The regular-tree and annealed ER experiments support a parameter-dependent survival-filtering interpretation. The strong reversal in the original high-variance versus low-variance comparisons is attenuated and not consistently signed when the one-step extinction probabilities are matched. Graph-level resampling and sensitivity diagnostics delimit the finite-sample evidence. The resulting conclusions concern finite horizons and finite graph ensembles, not asymptotic speed or infinite-volume coexistence. Establishing a general survival-conditioned frontier comparison remains an open problem.

Author contributions

Yunzhi Zhu: Conceptualization, methodology, writing – original draft. Saisai Hou: Visualization, formal analysis, data curation, writing – review & editing. Sen Zhang: Validation, formal analysis, writing – review & editing, supervision, funding acquisition. All authors have read and approved the final manuscript.

Use of Generative-AI tools declaration

The authors used Codex (OpenAI) to assist with language editing, checking bibliographic information in the reference list, and document formatting. The authors independently checked the resulting edits and references and take full responsibility for the content of this article.

Acknowledgments

This work was supported by the Jiangsu Higher Education Institutions Basic Science (Natural Science) Research General Program, Grant No. 25KJD110002, with Sen Zhang as the principal investigator.

Conflict of interest

The authors declare that they have no conflicts of interest.

Data availability

The simulation code and generated CSV summaries are supplied with the manuscript. The Python script regenerates all tables and figures from the stated random seeds.

References

1. J. D. Biggins, The first- and last-birth problems for a multitype age-dependent branching process, *Adv. Appl. Probab.*, **8** (1976), 446–459. https://doi.org/10.2307/1426138
2. J. M. Hammersley, Postulates for subadditive processes, *Ann. Probab.*, **2** (1974), 652–680. https://doi.org/10.1214/aop/1176996611
3. J. F. C. Kingman, The first birth problem for an age-dependent branching process, *Ann. Probab.*, **3** (1975), 790–801. https://doi.org/10.1214/aop/1176996266
4. L. Addario-Berry, B. Reed, Minima in branching random walks, *Ann. Probab.*, **37** (2009), 1044–1079. https://doi.org/10.1214/08-AOP428
5. Y. Hu, Z. Shi, Minimal position and critical martingale convergence in branching random walks, and directed polymers on disordered trees, *Ann. Probab.*, **37** (2009), 742–789. https://doi.org/10.1214/08-AOP419
6. E. Aïdékon, Convergence in law of the minimum of a branching random walk, *Ann. Probab.*, **41** (2013), 1362–1426. https://doi.org/10.1214/12-AOP750
7. T. Madaule, Convergence in law for the branching random walk seen from its tip, *J. Theor. Probab.*, **30** (2017), 27–63. https://doi.org/10.1007/s10959-015-0636-6
8. N. Gantert, T. Höfelsauer, Large deviations for the maximum of a branching random walk, *Electron. Commun. Probab.*, **23** (2018), 1–12. https://doi.org/10.1214/18-ECP135

9. P. Dyszewski, N. Gantert, T. Höfelsauer, The maximum of a branching random walk with stretched exponential tails, *Ann. Inst. H. Poincaré Probab. Statist.*, **59** (2023), 539–562. https://doi.org/10.1214/22-AIHP1260

10. R. Lyons, Y. Peres, *Probability on trees and networks*, Cambridge: Cambridge University Press, 2017. https://doi.org/10.1017/9781316672815

11. J. D. Biggins, A. E. Kyprianou, Measure change in multitype branching, *Adv. Appl. Probab.*, **36** (2004), 544–581. https://doi.org/10.1239/aap/1086957585

12. Z. Shi, *Branching random walks: École d'Été de probabilités de saint-flour XLII–2012*, Cham: Springer, 2015. https://doi.org/10.1007/978-3-319-25372-5

13. M. Deijfen, O. Häggström, The pleasures and pains of studying the two-type Richardson model, In: *Analysis and stochastics of growth processes and interface models*, Oxford: Oxford University Press, 2008, 39–54. https://doi.org/10.1093/acprof:oso/9780199239252.003.0002

14. M. Deijfen, T. Vilkas, Competition on $\mathbb{Z}^d$ driven by branching random walk, *Electron. Commun. Probab.*, **28** (2023), 1–11. https://doi.org/10.1214/23-ECP521

15. D. Bertacchi, F. Zucca, A generating function approach to branching random walks, *Braz. J. Probab. Stat.*, **31** (2017), 229–253. https://doi.org/10.1214/16-BJPS311

16. D. Bertacchi, F. Zucca, Strong survival and extinction for branching random walks via new order for generating functions, *J. Appl. Probab.*, **63** (2026), 950–968. https://doi.org/10.1017/jpr.2025.10070

17. M. Dussaule, L. Wang, W. Yang, Branching random walks on relatively hyperbolic groups, *Ann. Probab.*, **53** (2025), 391–452. https://doi.org/10.1214/24-AOP1708

18. I. Makarova, D. Balashova, S. Molchanov, E. Yarovaya, Branching random walks with two types of particles on multidimensional lattices, *Mathematics*, **10** (2022), 867. https://doi.org/10.3390/math10060867

19. F. Coghi, J. Morand, H. Touchette, Large deviations of random walks on random graphs, *Phys. Rev. E*, **99** (2019), 022137. https://doi.org/10.1103/PhysRevE.99.022137

20. B. Bollobás, *Random graphs*, 2 Eds., Cambridge: Cambridge University Press, 2011. https://doi.org/10.1017/CBO9780511814068

21. R. van der Hofstad, *Random graphs and complex networks*, Cambridge: Cambridge University Press, 2017. https://doi.org/10.1017/9781316779422

22. M. Newman, *Networks*, 2 Eds., Oxford: Oxford University Press, 2018. https://doi.org/10.1093/oso/9780198805090.001.0001

Supplementary

The supplementary material below contains the code-availability information and Appendices A–C.

Code availability

The supplementary script `simulate_brw_competition.py` regenerates the CSV files, tables, and figures. The production command is

```
python code/simulate_brw_competition.py --tree-runs 1200 \
--er-graphs 80 --runs-per-er-graph 40 --tree-depth 20 \
--tree-max-gen 10 --er-n 900 --max-gen 15 \
--n-boot 2500 --n-perm 2500 --outdir .
```

The script leaves unreached vertices uncoloured, preserves particle multiplicities, uses deterministic case-dependent seeds rather than Python's randomized built-in `hash`, implements the tree as an implicit depth-20 rooted tree, chooses exchangeable ER start vertices, records matched-Poisson calibration diagnostics, and applies graph-level cluster bootstrap to annealed ER summaries.

The accompanying scripts `reanalyze_saved.py` and `sensitivity_analysis.py` provide the saved-run reanalysis and one-factor-at-a-time sensitivity calculations, respectively. The scripts and six CSV summaries are supplied as machine-readable supplementary material.

A. Proofs of finite-horizon identities

This appendix gives the details behind the validation identities used in Section 3.3. They are standard, but recording them is useful because they identify exactly what is preserved by the simulation and where the second factorial moment enters.

A.1. Avoidance recursion

Let $A \subseteq V$, and let $q_t^A(x)$ be the probability that a branching random walk started from one particle at x has no descendant visiting A during generations $0, 1, \ldots, t$. If $x \in A$, then $q_t^A(x) = 0$ for all $t \geq 0$. If $x \notin A$, then $q_0^A(x) = 1$. For $t \geq 0$, conditioning on the first reproduction event gives

$$q_{t+1}^A(x) = f\left(\sum_y P(x, y) q_t^A(y)\right),$$

where $f(s) = \mathbb{E}s^K$ is the offspring generating function. Indeed, a child jumps to y with probability $P(x, y)$ and then avoids A for the next t generations with probability $q_t^A(y)$. Conditional on having $K = k$ children, the avoidance probability is the k-th power of the one-child average. Averaging over K gives the generating function.

This recursion is an exact dynamic program on a finite graph. It is also a direct check that one must not replace a branching random walk by a first-visit exploration: descendants that return to previously visited vertices still contribute to future generations and therefore to the recursion.

A.2. Exact colour probabilities

Let $a_{c,t}(v) = \mathbb{P}(\tau_c(v) = t)$ and $Q_{c,t}(v) = \mathbb{P}(\tau_c(v) > t)$ for $c \in \{r, b\}$. Before colouring, the red and blue branching random walks are independent. Hence, at a finite horizon T,

$$\mathbb{P}\{C_T(v) = r\} = \sum_{t=0}^{T} a_{r,t}(v) \left(Q_{b,t}(v) + \frac{1}{2} a_{b,t}(v)\right).$$

The first term inside parentheses accounts for strict red arrival before blue; the second accounts for a tie at the same generation, split by a fair coin. Vertices with $\tau_r(v) = \tau_b(v) = \infty$ are not ties and remain uncoloured. This formula is the probabilistic counterpart of the colouring rule in the code.

A.3. Second-moment recursion

Let $\mu_n(u) = \mathbb{E}Z_n(u)$ and $S_n(u,v) = \mathbb{E}[Z_n(u)Z_n(v)]$. Write $Y_{x,u}$ for the number of children produced by particles at x that jump to u in the next generation. Conditional on Z_n, the variables generated by different parent vertices and different parent particles are independent. A calculation for one parent at x gives

$$\mathbb{E}[Y_{x,u}] = mP(x,u),$$

$$\mathbb{E}[Y_{x,u}Y_{x,v}] = mP(x,u)\mathbf{1}_{\{u=v\}} + \eta P(x,u)P(x,v),$$

where $\eta = \mathbb{E}[K(K-1)]$. Summing over all parents and separating same-parent and different-parent contributions yields

$$\begin{aligned} S_{n+1}(u,v) = m^2 \sum_{x,y} P(x,u)P(y,v)S_n(x,y) \\ + \sum_x \mu_n(x)\{mP(x,u)\mathbf{1}_{\{u=v\}} + (\eta - m^2)P(x,u)P(x,v)\}. \end{aligned}$$

Thus the first moment depends only on m, while the covariance field depends explicitly on η. This is the finite-graph form of the variance amplification mechanism.

B. Details of the annealed Erdős–Rényi protocol

For each graph index $g = 1, \ldots, M$, the script samples an independent graph $G_g \sim G(n, 8/n)$, keeps its largest connected component, relabels it by consecutive integers, and chooses a maximum-degree vertex as the radial root. The root is not used as a colour start. Instead, two distinct neighbours of the root are sampled uniformly without replacement, and these two vertices are assigned to red and blue in the sampled order. Consequently, in the matched Poisson case, the start rule is exchangeable in the colour labels at the graph-ensemble level.

For every case and graph G_g, the script performs $L = 40$ competitions and alternates the colour assignment to the two start vertices. If $Y_{g,r}$ denotes an outer red fraction, the balanced random-effects mean squares are

$$MS_B = \frac{L}{M-1} \sum_{g=1}^{M} (\bar{Y}_g - \bar{Y})^2,$$

where $\bar{Y}_g = L^{-1} \sum_r Y_{g,r}$, and

$$MS_W = \frac{1}{M(L-1)} \sum_{g=1}^{M} \sum_{r=1}^{L} (Y_{g,r} - \bar{Y}_g)^2.$$

The reported components are $\max\{(MS_B - MS_W)/L, 0\}$ and MS_W. For conditional summaries, each graph contributes its survivor numerator and denominator, including $(0,0)$ for a graph with no survivors; these pairs are resampled as clusters. For effects relative to matched Poisson, graphs are paired because all cases are run on every sampled graph.

C. Implementation notes

The simulation keeps a multiset of particles at every generation. If $N_n(v)$ particles occupy vertex v, the code samples the total number of children produced by those $N_n(v)$ parents and then distributes those children among neighbours using a multinomial draw. This aggregation is equivalent to sampling each particle individually for the offspring laws used here, but it is much faster. Crucially, particles that jump to previously visited vertices are retained in the next generation. The active state is the particle configuration, not the set of discovered vertices.

The colouring rule is applied only after the red and blue arrival-time fields have been generated independently. If neither colour reaches a vertex by the terminal time, the vertex is labelled uncoloured. Such a vertex is not treated as a tie. This is important on finite graph exhaustions because many vertices are unreached at moderate terminal times, especially in high-variance cases with frequent extinction or delay.

For the negative-binomial law, the code uses the parameterisation $p = m/\sigma^2$ and $r = m^2/(\sigma^2 - m)$. Summing over several parents preserves the same success probability p and adds the shape parameter r, so the aggregated draw is distributionally equivalent to individual offspring sampling for the laws used here.

The script also records start vertices and their degrees, uses deterministic case seeds, and reports graph-level cluster uncertainty for ER summaries. It generated `per_run_results.csv`, `results_summary.csv`, `effect_summary.csv`, `er_variance_decomposition.csv`, the LaTeX tables in `tables/`, and all figures in `figures/`.