Stationary prediction from finite calibration: obstructions and sharp history rates under interval catastrophes

HUANG Man
Independent Researcher, Shanghai, China
Email: huangman@technocentrics.com
ORCID: https://orcid.org/0009-0003-3809-5105

hexagon:2611.00504v1

Abstract

Can calibration experiments confined to a fixed spatial diameter determine stationary block probabilities on larger spatial scales? We answer this question for a one-dimensional recovery process driven by independent rows of Poisson interval catastrophes with an unknown intensity law, under a common small recovery-weighted intensity budget. For every calibration diameter, two laws have exactly the same complete local damage-word distributions and deterministic-initial experimental transcripts, even under arbitrarily many adaptive repetitions, while a sufficiently large stationary block has probabilities separated by a positive constant. The pair also matches intensity moments of orders zero, one and two and the eventual power-law parameters. For the upper stationary open phase, we then prove that all-open finite-history approximations converge uniformly over a law class and all finite targets if and only if the recovery-weighted tails vanish uniformly. Under a common intensity-moment budget of order $2+\epsilon$, $0 < \epsilon \leq 1$, the worst missing-history error has matching scale $\min\{b, B[\max(m_0, T)]^{-\epsilon}\}$, with constants depending only on a fixed small recovery budget b. Here T counts retained physical rows, B is the moment budget and m_0 the minimum interval length. The lower bound holds at one fixed site; the upper bound covers every finite target and the full moment class, including zero intensity and laws without regular variation. Together, these results identify the obstruction to local calibration, the exact tail condition for uniform history approximation, and the sharp deterministic depth of the all-open truncation.

Keywords: correlated catastrophes; stationary prediction; partial identification; finite calibration; monotone recovery; history truncation.

Mathematics Subject Classification: 60K35, 60G55, 62M20.

1. Introduction

Can experiments confined to a fixed spatial diameter determine a recovery system's stationary block probabilities on larger spatial scales, even when the experiments can be repeated arbitrarily many times? If local calibration leaves an ambiguity, which tail restriction supports uniform history approximation, and how many physical rows does that approximation require? Rare long catastrophes make these questions inseparable: their frequency controls what local experiments observe, while their recovery time controls their stationary effect.

We study binary states on the integer line. At each discrete time, open sites propagate to their right neighbours and independent Poisson intervals remove the resulting open sites. Damage is spatially correlated and renewed independently across time. Starting with every site open in the remote past selects the upper stationary open phase; its complement is the stationary hole field. Our information contract supplies complete fresh damage words on windows of bounded diameter, together with recovery readouts from deterministic initial fields whose causal observations satisfy the same bound. It allows adaptive choices and arbitrarily many repetitions. Calibration diameter, retained physical history and observation count are separate resources.

The main results answer the three questions in sequence. First, exact local word and transcript equality can coexist with a constant gap in the stationary probability of a sufficiently large block, even after matching three intensity moments and the eventual power-law parameters (Theorem 2.1). Second, uniform vanishing of recovery-weighted tails is equivalent both to uniform single-site history approximation and to uniform approximation over all finite targets (Theorem 2.2). Third, a common moment budget yields matching worst-case history errors, with the lower bound at a fixed site and the upper bound for arbitrary finite target geometry (Theorem 2.4). Thus the information obstruction, the class criterion and the quantitative depth bound form one prediction theory for the selected phase.

The common tail condition controls catastrophe mass at lengths beyond the calibrated range. Its quantitative moment version converts that control into a physical-time rate. Stationary flux and a deterministic checkpoint comparison produce an auxiliary boundary distribution; a one-source restart then controls the missing past. This route applies across a common-budget class containing zero intensity, finite support and laws without regular variation.

Poisson Boolean capacity and covariogram identities describe how coverage observations encode an interval law [LP18, Theorems 17.3, 17.8 and 17.10]. Gracar, Korfhage and Mörters construct three-dimensional rotation-invariant convex-grain models with the same diameter marginals but different robustness, meaning percolation at every positive intensity [GKM26, Section 3.5]. Their ambiguity concerns dependence between grain diameters. Theorem 2.1 fixes one-dimensional interval geometry and preserves complete bounded-window damage laws and experimental transcripts before separating stationary block probabilities under recovery.

Bayraktar, Lu, Maggioni, Wu and Yang study finite-state cellular automata on a finite circular graph with conditionally independent site updates determined by a local transition matrix. They obtain trajectory-based inference under nonsingularity conditions, exhibit equal one-site stationary marginals with different joint stationary laws, and give different transition matrices with the same joint stationary law [BLMWY26, author version v4, Examples 3.3–3.4, Assumption 4.1 and Theorem 4.4]. Here the unknown object is an unrestricted sequence of interval intensities. The calibration contract supplies complete joint damage words on bounded windows; the ambiguity persists even when these word laws are known exactly and the allowed experiments are repeated arbitrarily many times.

Temporal coupling bounds derive memory loss from transition-kernel continuity [BFG99, CFF02, Theorem 4.3(iv)]. De Santis, Laxa and Löcherbach allow unknown past symbols in coupling-from-the-past algorithms for a specified transition kernel; their lower-envelope conditions yield almost-sure termination and, under stronger conditions, integrable backward times [DLL25, Section 2.4, Theorems 3–6]. Their unknown symbols represent unresolved past states. Our unknown quantity is the catastrophe law itself, and Theorems 2.2 and 2.4 characterize deterministic physical-history approximation uniformly over a class of such laws.

For spatial systems, finite ancestor clans give stationary sampling and truncation control [FFG02], high-temperature tail conditions give finite-range approximation of Gibbs measures [GLO10, Theorem 1], and a subcritical Kalikow backward sketch with a uniform positive lower bound on site clock rates gives exponential coupling from every initial configuration [GGLO13, conditions (5.1)–(5.2), Proposition 1 and Theorem 3]. The present results combine an unknown-law information contract with an exact class criterion and a matching lower bound; the time rate is derived from spatial catastrophe moments.

Further comparisons locate the model and its assumptions. In Poisson hail, spatial job extent and service duration determine workload stability at small arrival rates [FKM18, Theorem 1; MW24, Theorems 1.1–1.2]. Here spatial extent and recovery enter a binary phase's stationary prediction error. The deterministic hole map is a two-state monotone eroder. Gács and Törmä's dependent-error stability theorem assumes an exponential bound on every finite simultaneous error set [GT22, Section 2]; the polynomially likely long error blocks admitted here violate that specific hypothesis (Appendix C). Swart, Szabó and Toninelli obtain Peierls stability bounds for monotone cellular automata with intrinsic randomness using independent space-time local-rule choices [SST25, (1.7)–(1.8) and Theorems 10–11]. Our interval atoms produce spatially correlated local-rule choices already in the

first update. Recent restriction and correlation bounds treat continuous-time particle systems with bounded update diameters and unbounded dependence range, allowing phase coexistence; stationary spatial mixing is conditioned on rapid convergence from a specified initial configuration [JK26, assumption (R1) and Theorem 2.4]. We derive the physical-history rate directly for this synchronous interval model.

The positive-phase and interval-insertion arguments are shared with the author's preprint *An exact moment criterion for survival under renewed spatial catastrophes* [H26a, Version 1, Sections 3 and 5], whose principal results concern survival and stationary dilute density. The companion preprint *Finite-preparation response under heavy-tailed catastrophes: exact amplitudes and separation from Poisson coverage* [H26b, Version 1] develops compression, checkpoint comparisons and history estimates uniform over cutoffs of a fixed regularly varying law. Its Supplement S2.2–S2.3 supplies the one-source restart, finite-error contraction and remaining-time induction, including deterministic intensity strings. This paper uses those tools to establish exact local-law ambiguity, the common-tail equivalence and matching bounds over varying intensity sequences and all finite targets, with explicit budget and minimum-length dependence. The appendices reproduce the model-specific lemmas needed for a self-contained proof.

2. Model, information and principal results

2.1. Damage, the selected phase and physical history

Fix an integer $m_0 \geq 2$. For each row $t \in \mathbb{Z}$, origin $z \in \mathbb{Z}$ and length $\ell \geq m_0$, let $N_{t,z,\ell}$ be independent Poisson variables of means $\theta_\ell \geq 0$. An atom covers $[z, z+\ell-1] \cap \mathbb{Z}$; their union on row t is D_t. Assume

$$w_\ell = \frac{\ell(\ell+1)}{2}, \qquad \chi_\theta = \sum_{\ell \geq m_0} w_\ell \theta_\ell \leq b, \qquad 0 < b \leq \frac{1}{32}. \tag{2.1}$$

The point-hit rate $\lambda_\theta = \sum_\ell \ell\theta_\ell$ is finite, so every finite observed row word is well defined. Open sets and hole sets evolve by

$$O_t = (O_{t-1} \cup (O_{t-1}+1)) \setminus D_t, \qquad H_t = E(H_{t-1}) \cup D_t, \quad E(H) = H \cap (H+1). \tag{2.2}$$

Starting with $H_{-T} = \varnothing$, use the same rows for all T and let $T \to \infty$. At each site the hole indicators increase to a stationary field H_θ. Thus the selected phase is upper in open-site order and lower in hole order. Put

$$p_\theta(C) = \Pr(C \subset H_\theta), \quad p_\theta(\varnothing) = 1, \qquad h_\theta(k) = p_\theta([0,k-1]), \quad u_\theta = 1 - h_\theta(1).$$

The all-hole state is another stationary state. Our history bounds concern the selected phase.

For a finite target C, trace ancestors with the actual kill-before-expansion update

$$U_0 = C, \qquad U_{s+1} = G_{D_{s+1}}(U_s), \qquad G_D(C) = (C \setminus D) \cup ((C \setminus D) - 1). \tag{2.3}$$

Rows are relabelled in reverse chronological order. The probability of a hole target from an all-open past T rows ago is

$$p_{\theta,T}(C) = \Pr_C(U_T = \varnothing).$$

The missing-history error is

$$\Delta_\theta^C(T) = p_\theta(C) - p_{\theta,T}(C) \geq 0. \tag{2.4}$$

An empty target has zero error. Under the zero law, $p_0(C) = 0$ for nonempty C and every history error is zero.

2.2. Calibration contract

A fresh damage word is the complete binary vector $\mathbf{1}_{D\cap W}$ on a finite integer interval W. A diameter-M calibration supplies its law on every such window with $\operatorname{diam}(W) \le M$. The corresponding experimental information is formalized by an M-local observation policy. The policy has an independent random seed and the same measurable decision rules under every candidate law. Before reading an as-yet-unread physical row, it selects one containing window W from its earlier transcript and seed, then receives the complete damage word on W. It may use that word for any adaptive readouts, deterministic updates or deletion arms. All uses of the same physical row must stay within its selected window; a row shared by several arms is revealed once and reused. A new physical row has the original fresh-row law independently of the earlier transcript. The policy may stop adaptively or run countably many rounds.

Finite recovery experiments from a common deterministic initial field are included whenever their required row words obey this domain restriction. For example, reading $[0, W_0]$ through age T_0 uses row-r causal sites in $[-T_0 + r, W_0]$; $W_0 + T_0 \le M$ suffices. Once the causal words are supplied, choosing among readouts or preparation arms is measurable postprocessing. A further experiment can be selected using previous experimental transcripts and fresh rows. Independent repetitions use disjoint physical rows. Multiple observations on a shared row remain subject to one union-diameter budget, rather than one budget per arm or per separately named subwindow.

The channel supplies binary damage and deterministic-initial recovery readouts. Primitive atom labels, stationary initial observations and additional remote pieces of an already revealed row are separate information. Proposition 3.1 proves exact transcript equality for this policy class.

Define the recovery-weighted tail

$$A_\theta(N) = \sum_{\ell \ge N} \theta_\ell \frac{(\ell - N + 1)(\ell - N + 2)}{2}, \qquad N \ge 1. \tag{2.5}$$

Lengths below m_0 have zero intensity. We use either a known nonincreasing envelope

$$A_\theta(N) \le E_N, \quad E_1 = b, \quad 0 \le E_N \le b, \quad E_N \longrightarrow 0, \tag{2.6}$$

or the stronger moment class

$$\mathcal{K}_{b,B,\epsilon,m_0} = \left\{ \theta \ge 0 : \chi_\theta \le b, \sum_\ell \theta_\ell \ell^{2+\epsilon} \le B \right\}, \quad B \ge b, \quad 0 < \epsilon \le 1. \tag{2.7}$$

The budgets are supplied valid restrictions on the unknown law. A finite calibration transcript does not certify them.

2.3. Principal results

Theorem 2.1 (exact calibration and nonvanishing stationary ambiguity). Fix integers $M \ge m_0$, $K \ge 1$, a number $p > 3$ and $0 < b \le 1/32$. Put $c_b = 1 - e^{-b/48}$. For every $0 < \eta < c_b$, there exist $k \ge K$ and two fixed, strictly positive intensity sequences θ^-, θ^+ such that:

1. Their complete diameter-M fresh damage-word laws, and hence every experiment in the calibration contract, agree exactly.
2. Their total intensities, first length moments and second length moments agree; both have $\chi = b$. For a common $a > 0$, both equal $a\ell^{-p}$ for all sufficiently large ℓ.
3. Their selected stationary probabilities satisfy $h_{\theta^-}(k) \le \eta$ and $h_{\theta^+}(k) \ge c_b$.

The target and pair may depend on M, K, η. They are fixed before any experimental repetitions. In particular, the worst expected absolute prediction error over this pair is at least $(c_b - \eta)/2$, even when the common moments and eventual tail parameters are supplied exactly.

For the following results put

$$u_0 = \frac{1+\sqrt{1-8b}}{2}, \qquad a_0 = \frac{b}{u_0}, \qquad q = 4a_0 < 1. \tag{2.8}$$

Theorem 2.2 (uniform-tail characterization). For any nonempty class $\mathcal{F}$ of laws satisfying (2.1) with the same m_0, b, put

$$e_{\mathcal{F}}(N) = \sup_{\theta\in\mathcal{F}} A_\theta(N), \quad F_{\mathcal{F}}(T) = \sup_{\theta\in\mathcal{F},\ C \text{ finite}} \Delta_\theta^C(T), \quad F_{\mathcal{F}}^{\mathrm{site}}(T) = \sup_{\theta\in\mathcal{F}} \Delta_\theta^{\{0\}}(T).$$

Then the following conditions are equivalent:

$$\begin{aligned} e_{\mathcal{F}}(N) \longrightarrow 0 \quad &\Longleftrightarrow \quad \sup_{\theta\in\mathcal{F}} \sum_{\ell\ge N} w_\ell\theta_\ell \longrightarrow 0 \\ &\Longleftrightarrow \quad F_{\mathcal{F}}^{\mathrm{site}}(T) \longrightarrow 0 \quad \Longleftrightarrow \quad F_{\mathcal{F}}(T) \longrightarrow 0. \end{aligned} \tag{2.9}$$

Section 4.2 proves the equivalence. An evaluable common envelope turns this qualitative criterion into a numerical history cap.

Theorem 2.3 (prediction intervals under a common envelope). Under (2.1) and (2.6), there is an explicit function $V_E(N) \le a_0$, tending to zero on positive integers, for which $p_\theta(C) \le V_E(|C|)$ for every nonempty finite C.

For a nonempty finite target C, if a supplied law Q on complete damage words of $W = [\min C - T, \max C]$ satisfies $\|P_{\theta,W} - Q\|_{\mathrm{TV}} \le \delta$ for every compatible θ, define, using T independent Q words,

$$z_Q = \Pr_Q(U_T = \varnothing), \qquad r_Q = \mathbb{E}_Q[\mathbf{1}_{\{U_T\neq\varnothing\}} V_E(|U_T|)], \qquad e_T = \min(1, T\delta).$$

Then every compatible law satisfies

$$p_\theta(C) \in [\max(0, z_Q - e_T),\ \min(1, z_Q + r_Q + e_T)]. \tag{2.10}$$

Total variation throughout is the supremum of event-probability differences. Under (2.7), one may use

$$V(N) = \min(a_0, C_p N^{-\epsilon}), \qquad C_p = c_p(b)B, \quad c_p(b) = \frac{q(1+q)}{b(1-q)^3}. \tag{2.11}$$

Theorem 2.4 (sharp class-uniform physical-history scale). For fixed $b \in (0, 1/32]$, there are explicit finite constants $0 < c_-(b) \le C_+(b) < \infty$ such that, for all $m_0 \ge 2$, $B \ge b$, $0 < \epsilon \le 1$ and integers $T \ge 1$,

$$\begin{aligned} c_-(b)S(T) \le \sup_{\theta\in\mathcal{K}} \Delta_\theta^{\{0\}}(T) \le \sup_{\theta\in\mathcal{K},\ C \text{ finite}} \Delta_\theta^C(T) \le C_+(b)S(T), \\ S(T) = \min\{b, B[\max(m_0, T)]^{-\epsilon}\}. \end{aligned} \tag{2.12}$$

One may take $c_-(b) = (1-a_0)/(128(1+b))$. For fixed class parameters this is $\Theta(T^{-\epsilon})$. For

$$0 < \tau < c_-(b)\min\{b, Bm_0^{-\epsilon}\}, \tag{2.13}$$

any deterministic all-open history cap with uniform error at most τ satisfies $T \ge (c_-(b)B/\tau)^{1/\epsilon}$. A sufficient cap is

$$T = \max\{1, \lceil (C_t(b)B/\tau)^{1/\epsilon} \rceil\}, \tag{2.14}$$

where $C_t(b)$ is given in Section 6. The lower-bound law may vary with T: the matching scale in (2.12) describes the worst error over the class. Define the deterministic all-open truncation depth

$$\mathcal{T}_{\mathcal{K}}(\tau) = \inf\{T \geq 1 : T \in \mathbb{N},\ F_{\mathcal{K}}(T) \leq \tau\}.$$

Equations (2.13)–(2.14) give matching orders for this deterministic all-open physical-history depth. Section 8 records separate sufficient sample and site costs for a specified observation protocol.

3. Exact local calibration and its stationary obstruction

3.1. What a complete local word determines

For a finite observed set A, the origins of length-ℓ intervals hitting it form $\bigcup_{x\in A}[x-\ell+1, x]$. Thus

$$\Pr(A \cap D = \varnothing) = \exp\left[-\sum_{\ell} \theta_\ell U_\ell(A)\right], \qquad U_\ell(A) = \left|\bigcup_{x\in A}[x-\ell+1, x]\right|. \tag{3.1}$$

If $A \neq \varnothing$, $\operatorname{diam}(A) \leq M$ and $\ell > M$, those origin intervals overlap and

$$U_\ell(A) = \ell + \operatorname{diam}(A). \tag{3.2}$$

Consequently a complete diameter-M word depends on the long intensities only through $\sum_{\ell>M}\theta_\ell$ and $\sum_{\ell>M}\ell\theta_\ell$, once the intensities at lengths at most M are specified. Void probabilities for all subsets determine every binary pattern by inclusion–exclusion. This is the lattice counterpart of the Boolean capacity calculation in [LP18, Theorem 17.3]. Matching the second length moment adds information without changing this local equality.

There is no conflict with identification from all lags. Set $\lambda = \sum_\ell \ell\theta_\ell$, $\Gamma_d = \sum_{\ell>d}(\ell-d)\theta_\ell$, and $Y_x = \mathbf{1}_{\{x\in D\}}$. Directly from (3.1),

$$\operatorname{Cov}(Y_0, Y_d) = e^{-2\lambda}(e^{\Gamma_d} - 1), \qquad \theta_\ell = \Gamma_{\ell-1} - 2\Gamma_\ell + \Gamma_{\ell+1}. \tag{3.3}$$

Thus exact covariances at every lag identify the intensity sequence inside this model class. The covariance identity is the lattice counterpart of [LP18, Theorem 17.8]. Theorem 17.10 there identifies a continuum Boolean model's intensity and centred grain law from its full capacity functional; here the second-difference formula proves identification within the declared lattice-interval class. The obstruction concerns a bounded window.

Proposition 3.1 (equality of adaptive local transcripts). Suppose two laws have the same complete damage-word distribution on every window of diameter at most M. Every M-local observation policy defined in Section 2.2 has the same transcript law under the two models. The conclusion includes all measurable forecasts based on the transcript and the policy's independent random seed.

Proof. Include the random seed in the transcript and couple it identically. Conditional on any matched prefix, the next decision rule chooses the same window and continuation or stopping decision in both models. If it requests a fresh row, the conditional word law is the same supplied marginal, since the row is independent of the prefix. Couple that word identically. Every recovery readout and every arm using it is then the same measurable function of the matched history. This inductively couples every finite prefix identically, including a stopped transcript extended by a fixed terminal symbol. Equality of all finite-prefix laws gives equality of countable transcript laws, hence of any measurable output. Reusing a row within its already supplied window is postprocessing and introduces no new kernel. Repeating experiments on fresh rows repeats this argument. □

3.2. A three-moment construction

Let $D = M+1$, $K_D = (5D^2+2D)/2$, $s = b/(2K_D)$, and

$$S_p = \sum_{\ell\geq m_0} w_\ell \ell^{-p}, \qquad a = b/(2S_p).$$

Take $q^- = (\delta_D + \delta_{3D})/2$. For an integer $L \geq \max(2k, 3D)$, put $d = L - 2D$ and

$$q^+ = \frac{D}{d+D}\delta_D + \left(1 - \frac{D}{d}\right)\delta_{2D} + \frac{D^2}{d(d+D)}\delta_L. \tag{3.4}$$

Both are probability distributions with first moment $2D$ and second moment $5D^2$. For q^+, the centred first moment about $2D$ is zero and the centred second moment is D^2: the two noncentral weights satisfy $w_D = (d/D)w_L$. Set

$$\theta_\ell^\pm = a\ell^{-p} + sq^\pm(\ell). \tag{3.5}$$

All added atoms lie above M. The zeroth, first and second moments match, and $\chi = aS_p + sK_D = b$. Equations (3.1)–(3.2) give exact complete-word equality. After their finitely many additional lengths, both laws are exactly $a\ell^{-p}$.

An interval $[z, z+L-1]$ born j rows earlier leaves, under erosion alone, the remnant $[z+j, z+L-1]$. Since (2.2) is monotone, this remnant remains contained in the actual holes, irrespective of other damage. For a target $[0, k-1]$, counting origins and ages gives

$$h_\theta(k) \geq 1 - \exp\left[-\theta_L \frac{(L-k+1)(L-k+2)}{2}\right]. \tag{3.6}$$

For the added length-L atom of θ^+, $L \geq 2k$ and $d(d+D) \leq L^2$ make the exponent at least $sD^2/8 \geq b/48$. Hence $h_{\theta^+}(k) \geq c_b$.

The low law θ^- is fixed before choosing k. Section 4 proves that its stationary block probabilities tend to zero whenever $\chi < \infty$ and the small-budget bound holds. Choose $k \geq K$ with $h_{\theta^-}(k) \leq \eta$, then choose L. This proves Theorem 2.1 without a stationary amplitude theorem or a regular-variation limit.

If a randomized forecast Z has the same output law under two targets $h_- < h_+$, the triangle inequality gives

$$\max_\sigma \mathbb{E}_\sigma|Z - h_\sigma| \geq \frac{h_+ - h_-}{2}, \qquad \max_\sigma \mathbb{E}_\sigma(Z - h_\sigma)^2 \geq \frac{(h_+ - h_-)^2}{4}. \tag{3.7}$$

For positive h_-, minimizing the maximum relative square error gives $[(h_+ - h_-)/(h_+ + h_-)]^2$. These are elementary identical-output bounds. The common eventual exponent and coefficient supply no common onset scale for stationary extrapolation, paralleling the distinction between pointwise tail assumptions and quantified uniform classes in [CK15].

4. Spatial certificates for compatible laws

The proof begins with conditional correlation (Appendix A.1), then constructs the positive phase before passing to stationary flux and the spatial tree bound (A.2–A.4). Target compression (B.1) transfers block bounds to arbitrary finite sets. The checkpoint comparison (B.2) and flux then control boundary marginals in Section 5, which feed the physical-time restart in Section 6. None of these spatial or checkpoint arguments uses a history-rate conclusion. We first record their spatial consequences.

Lemma 4.1 (phase, compression and tree bound). Under (2.1), $u_\theta \geq u_0$. For every nonempty finite C,

$$p_\theta(C) \leq h_\theta(|C|) \leq \Pr(X_\theta \geq |C|). \tag{4.1}$$

For a nonzero law, X_θ is the total cost of a binary tree with independent identically distributed node variables. A node is active with probability $a_\theta = \chi_\theta/u_\theta$ and then has a positive integer cost W_θ with

$$\Pr(W_\theta \geq N) = A_\theta(N)/\chi_\theta.$$

An active node has two children; an inactive node has cost zero and no children. For the zero law set $X_0 = 0$; no conditional positive cost is required. This is a spatial comparison tree, whose cost is a spatial length.

For a nonzero law satisfying (2.6), combine activity and cost in one node variable Y_θ, zero at an inactive node. Its tails obey

$$\Pr(Y_\theta \geq N) = A_\theta(N)/u_\theta \leq E_N/u_0. \tag{4.2}$$

The right side is a valid integer-variable tail, with mass $1 - a_0$ at zero. Independent quantile couplings at all potential vertices embed the actual comparison tree in this common tree. This couples activation and cost together; a bound on the conditional positive cost alone would divide by the unknown χ_θ.

The common tree has J active nodes with

$$\Pr(J = j) = \mathrm{Cat}(j) a_0^j (1 - a_0)^{j+1} \leq q^j.$$

Given its shape, the positive costs are independent with tails E_N/b. For integers $m \geq 2, N \geq 1$, define

$$\begin{aligned} U_{E,m}(N) &= \min\left\{ a_0, \sum_{j=1}^{m-1} \mathrm{Cat}(j) a_0^j (1 - a_0)^{j+1} \min\left(1, \frac{j E_{\lceil N/j \rceil}}{b}\right) + \frac{q^m}{1-q} \right\}, \\ V_E(N) &= U_{E,N+1}(N). \end{aligned} \tag{4.3}$$

A sum of j positive costs at least N has one cost at least $\lceil N/j \rceil$. The union bound and the Catalan bound prove (4.3). Every fixed-j term tends to zero and the summands are dominated by q^j, so $V_E(N) \to 0$. The zero law satisfies the same certificate directly.

Under (2.7), $A_\theta(N) \leq \min(b, BN^{-\epsilon})$. Conditional on $J = j$, the cost-sum tail is at most $(B/b) j^{1+\epsilon} N^{-\epsilon} \leq (B/b) j^2 N^{-\epsilon}$. Summing $j^2 q^j$ proves (2.11). For each individual law with $\chi < \infty$, use its own $A_\theta(N) \to 0$ in the same argument to obtain $h_\theta(N) \to 0$.

Stationary harmonicity and the finite-history ancestor identity give

$$p_\theta(C) = \mathbb{E}_\theta p_\theta(U_T) = p_{\theta,T}(C) + \mathbb{E}_\theta[\mathbf{1}_{\{U_T \neq \varnothing\}} p_\theta(U_T)]. \tag{4.4}$$

For T independent complete-word maximal couplings, the chance of any mismatch is at most e_T. Under Q, the bounded function $p_\theta(U_T)$ lies between $\mathbf{1}_{\{U_T = \varnothing\}}$ and that indicator plus $\mathbf{1}_{\{U_T \neq \varnothing\}} V_E(|U_T|)$. This proves (2.10).

Define $\tau_\varnothing = \inf\{s \geq 0 : U_s = \varnothing\}$. By finite-history reversal, $p_{\theta,T}(C) = \Pr_C(\tau_\varnothing \leq T)$. The same monotone prehistory construction defining $p_\theta(C)$ therefore gives directly

$$p_\theta(C) = \Pr_C(\tau_\varnothing < \infty), \qquad \Delta_\theta^C(T) = \Pr_C(T < \tau_\varnothing < \infty).$$

The Markov property and compression now yield

$$\Delta_\theta^C(s + t) = \mathbb{E}_C \Delta_\theta^{U_s}(t), \qquad 0 \leq \Delta_\theta^C(t) \leq h_\theta(|C|) \quad (C \neq \varnothing). \tag{4.5}$$

These identities hold for each law and each seed before a class supremum is taken. They include the empty target and the zero law with zero defect.

4.1. A common finite horizon for a general envelope

The envelope alone gives a uniform deterministic history horizon, although the following sufficient cap can be very large. First, $V_E(N)$ in (4.3) is nonincreasing for $N \geq 1$. For fixed m, every summand decreases with N. Increasing m to $m + 1$ adds at most q^m and decreases the geometric remainder by exactly q^m. Thus both changes decrease the bound, including after clipping at a_0.

Fix a desired error $\tau > 0$. Choose an integer $K_0 \geq 1$ with $V_E(K_0 + 1) \leq \tau/3$, put $\gamma = \tau/[3a_0(K_0 + 1)^2]$, and choose

$$N_0 \ge \lceil (m_0+1)/2 \rceil, \quad E_{N_0} \le \gamma/8, \quad R_0 = 2N_0 - 1, \quad \beta_0 = \frac{\gamma}{2\sum_{\ell=m_0}^{R_0} \ell}, \quad \eta_* = \left(\frac{\beta_0}{1+\beta_0}\right)^{K_0} > 0. \tag{4.6}$$

These choices exist because the envelope tends to zero. For every finite input set, the expected average occupation of sizes $1, \dots, K_0$ over $s = 0, \dots, R-1$ is bounded by

$$\max\left\{ \frac{K_0+1}{R} + \gamma(K_0+1)^2, \frac{1}{R\eta_*} \right\}. \tag{4.7}$$

Here is a direct proof uniform in the law. Let $R_D(C) = |C \cap D|$ count hit sites in a fresh row. For $j = |C|$, $\mathbb{E}R_D(C) \le j\lambda_\theta$, regardless of dependence among the hit indicators. If $\lambda_\theta \le \gamma$, use the bounded potential $F(C) = \min(|C|, K_0+1)$. For $1 \le j \le K_0$, no hit has probability at least $1 - j\lambda_\theta$ and raises the potential by at least one; on a hit its loss is at most j. Thus its expected drift is at least $1 - j(j+1)\lambda_\theta \ge 1 - \lambda_\theta(K_0+1)^2$.

For $j \ge K_0 + 1$, a potential decrease requires at least $j - K_0$ hit sites, and its magnitude is at most $K_0 + 1$. Markov's inequality therefore gives

$$\begin{aligned} \mathbb{E}\big[(F(C) - F(G_D(C)))_+\big] &\le (K_0+1)\Pr(R_D(C) \ge j - K_0) \\ &\le (K_0+1)\frac{j\lambda_\theta}{j-K_0} \le \lambda_\theta(K_0+1)^2. \end{aligned}$$

The empty state has zero drift. Combining the cases yields drift at least $\mathbf{1}_{\{1 \le |C| \le K_0\}} - \lambda_\theta(K_0+1)^2$. Telescoping and $0 \le F \le K_0 + 1$ give

$$\frac{1}{R}\sum_{s=0}^{R-1} \Pr_C(1 \le |U_s| \le K_0) \le \frac{K_0+1}{R} + \lambda_\theta(K_0+1)^2,$$

which proves the first term in (4.7).

If $\lambda_\theta \ge \gamma$, the inequality

$$\sum_{\ell \ge 2N_0} w_\ell \theta_\ell \le 4A_\theta(N_0) \le 4E_{N_0} \le \gamma/2 \tag{4.8}$$

follows from $(\ell - N_0 + 1)(\ell - N_0 + 2) \ge \ell(\ell+1)/4$. Since $w_\ell \ge \ell$, some supported length at most R_0 has intensity at least β_0. At every visit to a nonempty set of at most K_0 sites, require that length at each site-origin. These independent atoms kill the set in one row with probability at least $(1 - e^{-\beta_0})^{K_0} \ge \eta_*$. Absorption events at distinct times are disjoint, so their total probabilities sum to at most one. This proves the second term in (4.7).

The absorption formula makes the true remainder nonincreasing in time. For each law and seed, (4.4) gives

$$\Delta_\theta^C(R) \le V_E(K_0+1) + \frac{a_0}{R}\sum_{s=0}^{R-1} \Pr_C(1 \le |U_s| \le K_0).$$

Using (4.7) therefore gives $\sup_{\theta,C} \Delta_\theta^C(R) \le \tau$ whenever

$$R \ge \max\{1, 3a_0(K_0+1)/\tau, 3a_0/(\tau\eta_*)\}. \tag{4.9}$$

The zero law belongs to the low-intensity branch. This proof establishes a last-cut guarantee over all finite targets without assuming a monotone realized ancestor size or a monotone expected cardinality certificate. Evaluating the cap requires numerical access to the supplied envelope and conservative arithmetic. Section 6 derives the polynomial rate under the moment restriction without using this cap.

4.2. The exact uniform-tail criterion

We prove (2.9). Set $Q_{\mathcal{F}}(N) = \sup_{\theta\in\mathcal{F}} \sum_{\ell\geq N} w_\ell\theta_\ell$. The coefficient bound used in (4.8) and the definition of A_θ give

$$Q_{\mathcal{F}}(2N) \leq 4e_{\mathcal{F}}(N), \qquad e_{\mathcal{F}}(N) \leq Q_{\mathcal{F}}(N). \tag{4.10}$$

Hence either uniform tail condition implies the other. If $e_{\mathcal{F}}(N) \to 0$, take $E_1 = b$ and $E_N = e_{\mathcal{F}}(N)$ for $N \geq 2$; Section 4.1 gives $F_{\mathcal{F}}(T) \to 0$. This implies the corresponding single-site convergence.

Conversely, length-ℓ atoms at ages at least T whose individual erosion remnant covers zero have total intensity $A_\theta(T+1)$. These old atoms are independent of the latest T rows, and their remnants are contained in the actual hole field. Thus

$$\Delta_\theta^{\{0\}}(T) \geq (1 - p_{\theta,T}(\{0\}))(1 - e^{-A_\theta(T+1)}) \geq \frac{1-a_0}{1+b} A_\theta(T+1). \tag{4.11}$$

Here $p_{\theta,T}(\{0\}) \leq a_0$, $A_\theta(T+1) \leq b$, and $1 - e^{-x} \geq x/(1+b)$ for $0 \leq x \leq b$. Taking the supremum gives

$$F_{\mathcal{F}}^{\text{site}}(T) \geq \frac{1-a_0}{1+b} e_{\mathcal{F}}(T+1).$$

Thus single-site uniform forgetting forces $e_{\mathcal{F}}(N) \to 0$, closing all implications. No maximizing law is required. This proof concerns a tail restriction on the law class itself; it does not assume that bounded calibration identifies that restriction.

5. A boundary comparison valid for the entire moment class

5.1. Flux and a comparison distribution

Write $h_k = h_\theta(k)$, $u = u_\theta$, and

$$d_k = h_k - h_{k+1} = \Pr([-k+1, 0] \subset H_\theta,\ 1 \notin H_\theta). \tag{5.1}$$

The d_k are nonnegative and nonincreasing, and $\sum_{k\geq 1} d_k = h_1$. Define a nonnegative integer variable Z_θ by

$$\Pr(Z_\theta \geq k) = d_k/u, \quad k \geq 1. \tag{5.2}$$

These are valid tails: $d_1/u \leq a_0/u_0 < 1$ and $d_k \to 0$. In particular,

$$\mathbb{E}Z_\theta = h_1/u \leq \mu_b := b/u_0^2 < 1/8. \tag{5.3}$$

For $k \geq 2$, monotonicity gives $\lfloor k/2 \rfloor d_k \leq h_{\lceil k/2 \rceil}$. Therefore, under (2.7),

$$\Pr(Z_\theta \geq k) \leq Lk^{-1-\epsilon}, \qquad L = 3\,2^\epsilon C_p/u_0. \tag{5.4}$$

The bound also holds at $k = 1$, since $C_p \geq 4B/u_0$ and $L \geq 12\mu_b$. Thus $\Pr(Z_\theta > y) \leq Ly^{-1-\epsilon}$ for all real $y > 0$.

5.2. Conditional boundary kernels and single-time domination

Insert the right hole ray $[0, \infty)$ into the stationary baseline. Let Λ_t be the left endpoint of its right-infinite hole component after t fresh rows. The ray-constrained open paths are exactly those whose coordinate at row zero is at most -1. Thus the deterministic checkpoint in Lemma B.1 is row zero with cap $c = -1$. At row t every site in $[t, \infty)$ is deterministically unreachable. Put $X_t = t - \Lambda_t$ and $J_t = \Lambda_t + 1 - \Lambda_{t+1} \geq 0$.

Corollary 5.1 (ray specialization of the checkpoint comparison). At every fixed time $t \geq 0$ and every boundary value x of positive probability,

$$\Pr(J_t \geq k \mid \Lambda_t = x) \leq d_k/u, \qquad k \geq 1. \tag{5.5}$$

Proof. To display the prehistory explicitly, start fully open at row $-n$. Write $H_t^{0,n}$ for the unrestricted holes and $H_t^{c,n}$ for the holes after requiring coordinate at most -1 at row zero. This latter restriction inserts the ray into the finite-prehistory baseline at that row. For $z = x - 1$ define

$$A_z^n = \{z \notin H_t^{0,n}\}, \quad O_z^n = \{z \notin H_t^{c,n}\}, \quad R_x^n = \{[x, t-1] \subset H_t^{c,n}\}.$$

For $x \leq t$, the event $\{\Lambda_t^n = x\}$ is precisely $O_z^n \cap R_x^n$; R_t^n is certain. Both events read finitely many terminal coordinates. On O_z^n, Lemma B.1(i) identifies the old fields strictly left of z. On $O_z^n \cap R_x^n$, the eroded constrained ray starts at $x + 1$, while z and x are open in the eroded field.

For any hole field H, let

$$f_{k,x}(H) = \Pr_D\left([x - k + 1, x] \subset (E(H) \setminus \{z, x\}) \cup D\right),$$

where D is an independent fresh row. Clamping z, x open makes $f_{k,x}$ a bounded increasing function only of old hole coordinates strictly left of z: the other queried eroded sites are at most $z - 1$, and erosion there reads that site and its left neighbour. On the boundary event, this clamp changes nothing in the constrained eroded exterior. Consequently

$$\Pr(J_t^n \geq k \mid \Lambda_t^n = x) = \mathbb{E}[f_{k,x}(H_t^{0,n}) \mid O_z^n, R_x^n].$$

Part (ii) of Lemma B.1 first removes the right increasing hole event R_x^n. Part (iii) then replaces constrained openness by unrestricted openness. They give, in this order,

$$\mathbb{E}[f_{k,x}(H_t^{0,n}) \mid O_z^n, R_x^n] \leq \mathbb{E}[f_{k,x}(H_t^{0,n}) \mid O_z^n] \leq \mathbb{E}[f_{k,x}(H_t^{0,n}) \mid A_z^n]. \tag{5.6}$$

Now let $n \to \infty$ under the common increasing-prehistory construction. For fixed t, x, k, all fields and observables just used depend on finitely many terminal sites. Their indicators and expectations converge. If $\Pr(\Lambda_t = x) > 0$, the probabilities of $O_z^n \cap R_x^n$ converge to a positive number and are eventually positive, as are those of O_z^n; $\Pr(A_z^n) \to u > 0$. Passing the conditional ratios to the limit is therefore valid. The unrestricted terminal field has the stationary law of H_θ. Finiteness of the ray boundary follows as in Appendix B.2: sufficiently far to the permitted side it agrees with a positive-density stationary open field, while the other side contains a deterministic unreachable ray.

Put $B = E(H_\theta)$ and $\Gamma = \{z, x \notin B\}$. The event $A_z = \{z \notin H_\theta\}$ is contained in Γ. On Γ clamping is immaterial, so the last conditional expectation in the stationary limit of (5.6) is at most

$$\frac{\Pr(\Gamma,\ [x - k + 1, x] \subset B \cup D)}{u} \leq \frac{h_k - h_{k+1}}{u}. \tag{5.7}$$

Indeed $B \subset B \cup D$, whose law is H_θ, and $B(x) = 0$ on Γ excludes the event that the block was already contained in B. The probabilities of the new and old block events are h_k and h_{k+1}. This proves (5.5). $\square$

We transfer the conditional increment bounds to fixed-time positions through an auxiliary chain with the same one-time marginals. Define

$$K_t(x, \cdot) = \mathcal{L}(J_t \mid X_t = x)$$

at states of positive marginal probability, and set $K_t(x, \{0\}) = 1$ at the other nonnegative integer states. On this countable state space the kernel and its quantile are measurable. For $0 < v < 1$, use the quantile convention

$$K_t^{-1}(x, v) = \min\{j \in \mathbb{N}_0 : K_t(x, \{0, \ldots, j\}) \geq v\},$$

and write $F_{Z_\theta}^{-1}(v)$ for the corresponding quantile of Z_θ. Take independent Uniform$(0, 1)$ variables $V_1, V_2, \ldots$, independent also of X_0°, whose law is that of X_0. Construct

$$\widehat{X}_0 = X_0^\circ, \qquad \widehat{X}_{t+1} = \widehat{X}_t + K_t^{-1}(\widehat{X}_t, V_{t+1}), \qquad Z_{\theta,i} = F_{Z_\theta}^{-1}(V_i). \tag{5.7a}$$

The identity $X_{t+1} = X_t + J_t$ gives the induction step for equal one-time marginals. For each $y \in \mathbb{N}_0$, using the preceding marginal equality,

$$\Pr(\widehat{X}_{t+1} = y) = \sum_{x=0}^{y} \Pr(X_t = x) K_t(x, \{y - x\}) = \Pr(X_{t+1} = y).$$

Thus $\widehat{X}_t \overset{d}{=} X_t$ at every fixed time; the argument uses only the conditional increment law given X_t. By (5.5), every projected increment in (5.7a) is at most $Z_{\theta,t+1}$ in this coupling; the zero-increment kernels satisfy the same bound. The $Z_{\theta,i}$ are independent and also independent of X_0°. Thus for each fixed integer T,

$$X_T \le_{\mathrm{st}} X_0^\circ + \sum_{i=1}^{T} Z_{\theta,i}, \qquad \Pr(X_0^\circ \ge k) = h_\theta(k), \tag{5.8}$$

where the comparison variables are independent. At time zero $X_0 \ge k$ is exactly the event that the k sites immediately left of the inserted ray are baseline holes. The construction uses the actual conditional increment kernels and permits history-dependent physical increments. It reproduces one-time marginals; the fixed-time bound (5.8) is the input used in Section 6. An initially empty exterior is coupled below the stationary hole exterior and has no larger retreat. Reflection and the shear $x \mapsto t - x$ give the other endpoint for ordinary propagation with increments $0, -1$. Under zero damage $X_0 = Z = 0$ and the edges are deterministic.

6. Physical-time restart and the matching lower bound

The one-source restart and remaining-time induction are also used in the companion response paper [H26b, Supplement S2.2–S2.3]. The input here is the class-uniform control of the boundary distribution in (5.2)–(5.8). Truncation of its comparison variables gives a fixed failure probability throughout the recovery-budget regime (2.1), with the threshold depending explicitly on B, ϵ. This yields the rate over the entire moment class (2.7), with the dependence on the budget, exponent and minimum length displayed below.

6.1. A fixed-error boundary estimate

Use the following constants, depending on b except where displayed:

$$\begin{gathered} g = 1/8 - \mu_b, \quad d_0 = g/16, \quad d_1 = 15g/16, \quad \eta_0 = 1/4, \quad \delta_0 = (\eta_0 - 3a_0)/2, \\ r = \min\left(1/2, \frac{2\delta_0 d_1^2}{3\mu_b}\right), \quad J_\epsilon = C_p d_0^{-\epsilon} + L r^{-1-\epsilon}, \quad n_0 = \max\{1, \lceil (3J_\epsilon/\delta_0)^{1/\epsilon} \rceil\}. \end{gathered} \tag{6.1}$$

All are positive. For example $u_0 \ge 3/4$, $a_0 \le 1/24$ and $\mu_b \le 1/18$, so $g > 0$ and $\delta_0 \ge 1/16$. The threshold n_0 can grow rapidly as ϵ decreases; the final multiplier is controlled through n_0^ϵ, as shown in Section 6.3.

Truncate the independent comparison variables at rT: $Y_i = Z_i \mathbf{1}_{\{Z_i \le rT\}}$. The event of an omitted large summand costs at most $L r^{-1-\epsilon} T^{-\epsilon}$. Since $Y_i^2 \le rTY_i$ and $\mathbb{E}Y_i \le \mu_b$, Chebyshev yields

$$\Pr\left(\sum_{i=1}^{T} Z_i > (\mu_b + d_1)T\right) \le L r^{-1-\epsilon} T^{-\epsilon} + \frac{\mu_b r}{d_1^2} \le L r^{-1-\epsilon} T^{-\epsilon} + 2\delta_0/3. \tag{6.2}$$

Splitting X_0° at $d_0 T$ and using (5.8) bounds either half-line's retreat beyond $T/8$ by

$$\Phi(T) = J_\epsilon T^{-\epsilon} + 2\delta_0/3. \tag{6.3}$$

Only a fixed small failure probability is needed. At $T = 8n - 1$ and $n \ge n_0$, $\Phi(T) \le \delta_0$. No endpoint logarithm or moment above order $2 + \epsilon$ enters this step.

6.2. One selected point and one output interval

For ordinary propagation $F_D(C) = (C \cup (C-1)) \setminus D$, the following deterministic path statement supplies the geometric input to the restart.

Lemma 6.1 (prefix exchange for a surviving source). Fix an integer $T \geq 0$ and any common disaster rows $D_1, \ldots, D_T$. Let A_T^S be the sites reachable from a source set S at row zero by open paths with increments $0, -1$. If $A_T^{\{y\}} \neq \varnothing$, write $L = \min A_T^{\{y\}}$ and $R = \max A_T^{\{y\}}$. Then

$$L = \min A_T^{[y,\infty)}, \qquad R = \max A_T^{(-\infty,y]}, \qquad A_T^{\{y\}} = A_T^{\mathbb{Z}} \cap [L, R]. \tag{6.3a}$$

Proof. Two allowed paths whose endpoint order reverses their starting order meet at an integer vertex: the difference of their coordinates changes by at most one at each step. At a meeting vertex, the prefix of either open path can be concatenated with the suffix of the other.

If a path from $[y, \infty)$ ended strictly left of L, it would meet a path from y to L. Swapping prefixes would give a path from y to that smaller endpoint, a contradiction. Inclusion of the point source in the half-line gives the reverse inequality and proves the first equality. A path from $(-\infty, y]$ ending strictly right of R meets a path from y to R and gives the symmetric contradiction, proving the second equality.

Now take $w \in A_T^{\mathbb{Z}} \cap [L, R]$ and an open path to w starting at q. If $q \leq y$, compare it with a path from y to L: their starting and ending orders force a meeting. If $q \geq y$, compare it with a path from y to R. In either case the prefix from y and the suffix to w form an open path, so $w \in A_T^{\{y\}}$. The reverse inclusion is immediate. All statements are pathwise and require only nonemptiness at the declared finite time. □

For an interval J of n sites, the full-open finite-history baseline has expected hole count at most na_0. Markov's inequality gives

$$\Pr(|O^{\text{full}} \cap J| < \lceil n/2 \rceil) \leq 2a_0. \tag{6.4}$$

There is no conditioning on survival in (6.4).

For actual ancestors, the exact order over $8n$ rows is

$$G_{D_{8n}} \cdots G_{D_1}(\{y\}) = B_- F_{D_{8n}} \cdots F_{D_2}(\{y\} \setminus D_1), \qquad B_- A = A \cup (A-1). \tag{6.5}$$

If the first kill leaves y alive, the $8n-1$ ordinary rows are fresh and independent of that kill. Choose a with $y \in [a, a+n-1]$ and let $J = [a-3n, a-2n-1]$. On both half-line events with retreat at most $(8n-1)/8$, the ordinary left endpoint is at most

$$y - 7(8n-1)/8 < a - 5n,$$

and the right endpoint is at least $y - (8n-1)/8 > a - n$. They enclose J. The final expansion only adds sites. Therefore

$$\Pr_y(0 < |U_{8n}| < \lceil n/2 \rceil) \leq 2\Phi(8n-1) + 2a_0. \tag{6.6}$$

For an arbitrary nonempty finite C, choose one point $y(C)$ before revealing the next rows. The update preserves unions. If C has a small nonempty descendant, the chosen point either died or also has a small nonempty descendant. Its extinction probability by that time is at most its eventual extinction probability $h_\theta(1) \leq a_0$. Hence, uniformly in C and its geometry,

$$\Pr_C(0 < |U_{8n}| < \lceil n/2 \rceil) \leq 3a_0 + 2\Phi(8n-1) \leq \eta_0, \qquad n \geq n_0. \tag{6.7}$$

The selected family may die while another family survives; the extra a_0 pays exactly that event. Conditional on any predictable seed, the future rows still have their original independent law.

6.3. Remaining-time induction

Let $F(T) = \sup_{\theta\in\mathcal{K},C \text{ finite}} \Delta_\theta^C(T)$. Apply (4.5), split at $|U_{8n}| = \lceil n/2\rceil$, use compression on the large branch and (6.7) on the small nonempty branch. For $t \geq 1, n \geq n_0$,

$$F(t+8n) \leq C_p\lceil n/2\rceil^{-\epsilon} + \eta_0 F(t). \tag{6.8}$$

Put

$$L_0 = 16(n_0+1), \qquad K = \max\{2C_p 64^\epsilon,\ a_0 L_0^\epsilon\}. \tag{6.9}$$

For an integer $Q \geq L_0$, take $n = \lfloor Q/16\rfloor$ and $t = Q - 8n$. Then $n \geq n_0$, $\lceil n/2\rceil \geq Q/64$, and $Q/2 \leq t < Q$. The base $F(j) \leq a_0$ for $1 \leq j < L_0$ and strong induction give

$$F(Q) \leq (C_p 64^\epsilon + K/2)Q^{-\epsilon} \leq KQ^{-\epsilon}, \tag{6.10}$$

because $\eta_0 2^\epsilon \leq 1/2$. The proof uses actual physical row counts throughout. It uses neither a prior time-rate theorem nor a spatial tree cost as a time variable.

The constant is linear in the moment budget, uniformly for $0 < \epsilon \leq 1$. Indeed set

$$E_b = c_p(b)[d_0^{-1} + 6/(u_0 r^2)], \qquad C_t(b) = \max\{128c_p(b), 48a_0E_b/\delta_0 + 32a_0/b\}. \tag{6.11}$$

Then $J_\epsilon \leq E_b B$. If $x = (3J_\epsilon/\delta_0)^{1/\epsilon}$, $n_0 \leq x+1$ and concavity gives $L_0^\epsilon \leq 16(3J_\epsilon/\delta_0 + 2)$. Since $B \geq b$, (6.9) implies $K \leq C_t(b)B$. Also

$$F(T) \leq \sup_\theta h_\theta(1) \leq \min(a_0, Bm_0^{-\epsilon}/u_0), \tag{6.12}$$

by $w_\ell \leq \ell^2$. Combining (6.10)–(6.12) proves the upper side of (2.12), with $C_+(b) = \max(C_t(b), 1/u_0)$.

6.4. Old remnants give the lower bound

For a target interval of length k, atoms at ages at least T whose individual eroded remnants cover it have total mean $A_\theta(k+T)$. They are independent of the most recent T rows. If the finite-history target is not already all holes and an old remnant covers it, stationary and finite-history indicators differ. Therefore

$$\Delta_\theta^{[0,k-1]}(T) \geq (1 - p_{\theta,T}([0,k-1]))(1 - e^{-A_\theta(k+T)}). \tag{6.13}$$

The age convention is that a row-0 mark has age zero; the omitted past begins at age T.

Choose a fixed positive background $\theta_\ell^0 = c\ell^{-6}$ with $\sum \theta_\ell^0 \ell^4 = H_0 = \min(b/8, B/8)$. For each T, let

$$\ell_T = \max(m_0, 2(T+1)), \quad \beta_T = \min(B/2, (b/4)\ell_T^\epsilon), \quad \theta^{(T)} = \theta^0 + \beta_T \ell_T^{-2-\epsilon}\delta_{\ell_T}. \tag{6.14}$$

The added moment is at most $B/2$ and its recovery budget at most $b/4$, so the law belongs to $\mathcal{K}$. Its eventual exponent remains six. Moreover

$$A_{\theta^{(T)}}(T+1) \geq \frac{\beta_T}{8\ell_T^\epsilon} \geq \frac{1}{32}\min(b, B\ell_T^{-\epsilon}). \tag{6.15}$$

Use $1 - e^{-A} \geq A/(1+b)$ for $A \leq b$, $p_{\theta,T}(\{0\}) \leq a_0$, and $\ell_T \leq 4\max(m_0, T)$. Since $4^\epsilon \leq 4$, (6.13) proves the lower side of (2.12). The cap consequence follows by first excluding the saturation and $T \leq m_0$ branches using (2.13). This completes Theorem 2.4.

For an explicit sufficient constant, take $b = B = 31/2048$ and $\epsilon = 1$. Then $u_0 = 31/32$, $a_0 = 1/64$, $C_p = 272/3375$, $\mu_b = 1/62$, $n_0 = 8060$, $L_0 = 128976$ and $K = 8061/4$. The bound $F(T) \leq KT^{-1}$ certifies one-percent absolute hole error with $T = 201525$. This sufficient analytic cap counts physical rows. Section 8 separately accounts for row widths and observations used to learn a prediction.

7. A decaying extrapolation obstruction inside the moment class

Proposition 7.1. Fix b, B, ϵ, m_0 as in (2.7). There are an explicit D_0 and $\zeta = \min(B/176, b/76) > 0$ such that for every integer $D \geq D_0$, two laws in $\mathcal{K}$ have identical diameter-$(D-1)$ calibration contracts, equal primitive moments of orders zero, one and two, and the same eventual power-law parameters, but at $k_D = \lceil 7D/2 \rceil$ satisfy

$$h_+(k_D) - h_-(k_D) \geq \frac{\zeta}{16(1+b)} D^{-\epsilon}. \tag{7.1}$$

Proof. Use the background θ^0 in Section 6.4, put $s_D = \zeta D^{-2-\epsilon}$ and take

$$\theta^- = \theta^0 + s_D\delta_D + 3s_D\delta_{3D}, \qquad \theta^+ = \theta^0 + 3s_D\delta_{2D} + s_D\delta_{4D}. \tag{7.2}$$

Both add-ons have mass $4s_D$, first moment $10s_DD$, second moment $28s_DD^2$ and recovery budget $c_D = s_D(14D^2 + 5D) \leq 19\zeta D^{-\epsilon}$. Their higher-moment coefficients are at most $88\zeta \leq B/2$, and $c_D \leq b/4$. For $D \geq m_0$, the laws belong to $\mathcal{K}$. Equations (3.1)–(3.2) prove the calibration equality.

The high law's length-$4D$ atom gives $A_+(k_D) \geq s_DD^2/8$, and hence

$$h_+(k_D) \geq \zeta D^{-\epsilon}/[8(1+b)]. \tag{7.3}$$

For the low law, dominate its tree by one with activity exactly a_0. Conditional on a positive node, give it a background label with probability χ_0/b and background cost tail $A_0(N)/\chi_0$; an add-on label with probability c_D/b and cost tail $A_{\mathrm{add}}(N)/c_D$, supported on $[1, 3D]$; or a dummy label with the remaining probability and cost one. For $N \geq 2$ its joint active-cost tail is $A_-(N)/u_0 \geq A_-(N)/u_-$; at $N = 1$ activity a_0 dominates χ_-/u_-. Thus Lemma 4.1's node coupling applies.

Let J be its number of positive nodes, $\mu = \mathbb{E}J = a_0/(1-2a_0)$, and

$$s_2 = \mathbb{E}J^2 = \frac{a_0(1+4\mu+2\mu^2)}{1-2a_0}, \qquad f_2 = s_2 - \mu.$$

Put $V_i = 0$ at an add-on label and equal to its cost otherwise. Tail summation, or $W_0 \leq \ell$ in the length-biased background mixture, gives $\chi_0\mathbb{E}W_0^2 \leq \sum_\ell \theta^0_\ell w_\ell \ell^2 \leq H_0$. Hence $\mathbb{E}V_i^2 \leq m_2 = 1 + H_0/b$. If total cost is at least k_D and at most one add-on node occurs, the non-add-on sum is at least $D/2$. Conditional on $J = j$, its second moment is at most j^2m_2. Two or more add-on labels cost at most $f_2(c_D/b)^2/2$. Therefore

$$h_-(k_D) \leq C_{\mathrm{bg}}D^{-2} + C_{\mathrm{pair}}\zeta^2D^{-2\epsilon}, \quad C_{\mathrm{bg}} = 4m_2s_2, \quad C_{\mathrm{pair}} = \tfrac{1}{2}f_2(19/b)^2. \tag{7.4}$$

Take D_0 to be the ceiling of the maximum of $m_0, 1$ and

$$[32(1+b)C_{\mathrm{bg}}/\zeta]^{1/(2-\epsilon)}, \qquad [32(1+b)C_{\mathrm{pair}}\zeta]^{1/\epsilon}.$$

Each term in (7.4) is then at most $\zeta D^{-\epsilon}/[32(1+b)]$. Subtract from (7.3) to prove (7.1). □

Equation (3.7) gives a worst expected absolute error of at least half of (7.1), even with the exact local word law. If an interval has coverage at least $1-\alpha$ under each law, $\alpha < 1/2$, its common output distribution satisfies

$$\Pr(|I| \geq h_+(k_D) - h_-(k_D)) \geq 1 - 2\alpha. \tag{7.5}$$

This is a one-sided spatial extrapolation bound at a declared growing target. The common moment budget makes its gap decay, while Theorem 2.1 retains a constant gap over a broader tail class.

8. Prediction procedures and distinct resource accounts

8.1. Complete-word calibration

For a nonempty finite target C under (2.7), choose T with $KT^{-\epsilon} \le \tau$, using (6.9). If Q approximates the true complete row word on $[\min C - T, \max C]$ within δ, uniformly over compatible laws, then

$$p_\theta(C) \in [\max(0, z_Q - e_T),\ \min(1, z_Q + \tau + e_T)]. \tag{8.1}$$

Row coupling transfers the emptiness probability, and the true stationary remainder lies in $[0, \tau]$. The bound covers every finite target whose causal window is supplied. Its accuracy input is total variation of the complete joint word on that window; low-order moments or separate site marginals do not determine this input.

A corresponding result applies to a shared-history preparation. At row zero delete disjoint finite root sets I_1, I_3 from the stationary open field. Let X_A be the terminal all-hole indicator for a finite target C after deleting roots indexed by $A \subset \{1, 3\}$ and evolving for age t. The four-arm contrast

$$D = X_{13} - X_1 - X_3 + X_\varnothing \in \{0, 1\} \tag{8.2}$$

uses the same old field and future rows. An empty target has contrast zero and needs no prediction. For a nonempty target, reverse the t future rows first, remove the roots, then reverse the same T older rows for every arm. The cut sets satisfy $C_{13} \subset C_1 \cap C_3$ and $C_\varnothing = C_1 \cup C_3$. With an all-open cut field the contrast is

$$d = \mathbf{1}_{\{C_{13}=\varnothing, C_1\neq\varnothing, C_3\neq\varnothing\}}.$$

For any recent history, discrepancy from the true contrast has conditional probability at most

$$\pi(C_{13}) + \pi(C_1) + \pi(C_3), \qquad \pi(S) = \mathbf{1}_{\{S\neq\varnothing\}} p_\theta(S). \tag{8.3}$$

If C_{13} is nonempty, a positive true response requires it all holes. If it is empty, disagreement requires one of the two nonempty single-arm sources to be all holes. Equality of a single-arm source with C_{13} gives exact response zero. These cases prove (8.3).

Condition on future rows and preparation. Older rows are independent of the resulting seeds, so each of the three expectations in (8.3) is a true remainder at time T, at most τ. They need not be independent of one another. If the complete word on $[\min C - (t + T), \max C]$ is within δ of Q, and $d_Q = \mathbb{E}_Q d$, then

$$\mathbb{E}_\theta D \in [\max(0, d_Q - 3\tau - e_{t+T}),\ \min(1, d_Q + 3\tau + e_{t+T})]. \tag{8.4}$$

8.2. Direct fresh-row observations

The inputs to this protocol are the declared class parameters and observed joint row words; replay uses the known deterministic update. In contrast, the kernel-based constructions in [DLL25, Section 2] specify lower-envelope transition probabilities over admissible pasts. Fixed-diameter calibration does not supply those kernel inputs.

One can avoid estimating an entire word distribution by observing fresh words and replaying them. Fix in advance $J = [a, b_0]$ of d sites and a maximum age $t_{\max}$. Consider all hole targets $C \subset J$ and all response queries with $C \subset J$, disjoint $I_1, I_3 \subset J$ and $0 \le t \le t_{\max}$. There are at most

$$m = 2^d + (t_{\max} + 1)6^d \le (t_{\max} + 2)6^d \tag{8.5}$$

indexed queries. The 6^d counts two target memberships and three root labels per site.

For $0 < \eta, \alpha < 1$, put $\rho = \eta/6$, $s = \eta/2$,

$$T = \max(1, \lceil (K/\rho)^{1/\epsilon} \rceil), \quad H = T + t_{\max}, \quad W = [a - H, b_0],$$
$$N = \left\lceil 2\eta^{-2} \left[d \log 6 + \log \frac{2(t_{\max} + 2)}{\alpha} \right] \right\rceil. \tag{8.6}$$

Read NH independent fresh row words on W and divide them into N disjoint chronological blocks. Each block's first T rows supplies old history and its remaining rows supply future history. Reverse, delete roots and reverse the old rows in the order specified above. All queries may reuse a block; distinct blocks use disjoint rows.

Hole outputs are Bernoulli emptiness indicators. Response outputs are also Bernoulli. To see this directly, split the row-zero ancestor set into $A_1 = S \cap I_1$, $A_3 = S \cap I_3$ and $E = S \setminus (I_1 \cup I_3)$. For a fixed finite-history hole field put $Q_Z = \mathbf{1}_{\{Z \subset H_T\}}$. The response is

$$Q_E(1 - Q_{A_1})(1 - Q_{A_3}) \in \{0, 1\}. \tag{8.7}$$

The hole bias lies in $[0, \rho]$; response absolute bias is at most 3ρ. For a Bernoulli variable Y, the log moment-generating function of $Y - \mathbb{E}Y$ has second derivative at most $1/4$, so it is bounded by $\lambda^2/8$. Chernoff and independent blocks give $\Pr(|\bar{Y} - \mathbb{E}Y| > s) \le 2e^{-2Ns^2}$. A union bound over (8.5) gives one probability-$1 - \alpha$ event on which every hole interval

$$[\max(0, \bar{Y} - s),\ \min(1, \bar{Y} + s + \rho)] \tag{8.8}$$

and every response interval

$$[\max(0, \bar{Y} - s - 3\rho),\ \min(1, \bar{Y} + s + 3\rho)] \tag{8.9}$$

covers its target. The statement holds under whichever admissible law generates the observations. Queries may be chosen after seeing the data within the declared domain. Enlarging that domain or stopping early requires additional confidence control.

The procedure reads NH row words and $NH(d + H)$ site values. A one-step all-open experiment realizes each damage-word observation, provided W and its left predecessor are prepared open. Such an implementation uses up to NH resets. Replaying one requested query requires at most order $4NH(d + H)$ site operations. At fixed class parameters, domain, age cap and confidence, the sufficient row and site counts are respectively $O(\eta^{-2-1/\epsilon})$ and $O(\eta^{-2-2/\epsilon})$. These are sufficient costs for the stated fresh-row protocol. The matching optimization in Theorem 2.4 concerns the deterministic all-open history depth.

8.3. Continuity for all-hole events

Let θ, θ' belong to the same fixed class $\mathcal{K}$, and set $d_w(\theta, \theta') = \sum_\ell w_\ell |\theta_\ell - \theta'_\ell|$. Couple common Poisson atoms and independent difference atoms. At an explored set U_s, difference atoms hitting it have mean at most $\frac{2}{3}|U_s| d_w$, because $\ell / w_\ell \le 2/3$. Before the first discrepancy $|U_s| \le |C|(s + 1)$. Thus the probability of a first discrepancy by row T is at most $|C|T(T + 1)d_w/3$, regardless of target diameter.

For $|C| \le T$, the two remainders lie in $[0, KT^{-\epsilon}]$, so their difference costs at most $KT^{-\epsilon}$. For $|C| > T$, both stationary probabilities are at most $C_p T^{-\epsilon}$. Consequently, with $K_m = \max(K, C_p)$,

$$\sup_{C \text{ finite}} |p_\theta(C) - p_{\theta'}(C)| \le \min\{1, K_m T^{-\epsilon} + (2/3)T^3 d_w\}. \tag{8.10}$$

For $0 < d_w \le 1$, choosing $T = \lceil d_w^{-1/(3+\epsilon)} \rceil$ gives

$$\sup_{C \text{ finite}} |p_\theta(C) - p_{\theta'}(C)| \le \min\{1, (K_m + 16/3) d_w^{\epsilon/(3+\epsilon)}\}. \tag{8.11}$$

Equation (8.11) gives a Hölder continuity bound in the all-hole event metric. Passing to arbitrary Boolean combinations can introduce growing inclusion–exclusion factors, so this metric is weaker than total variation of the entire infinite spatial field.

9. Discussion

Finite local calibration, uniform tail control and retained physical history play complementary roles in stationary prediction. Exact agreement of local data leaves a stationary ambiguity; uniform recovery-weighted tails characterize vanishing history truncation error; a common moment budget gives its matching quantitative scale. Together, Theorems 2.1, 2.2 and 2.4 identify what local observations determine, what class information must supplement them, and how deep the all-open history approximation must be.

The obstruction is generated by the recovery footprint of a long catastrophe. An interval of length L with intensity of order L^{-2} occupies space-time area of order L^2 as it erodes. Shorter atoms can compensate its contribution to every calibrated local word and to the matched moments, while its stationary effect stays finite. A common tail envelope controls precisely the recovery-weighted mass that can escape to such lengths. Single-site error already detects escaping old remnants; sufficiency then extends to every finite target geometry.

The upper bound converts stationary spatial control into a physical-time guarantee. Flux produces an integrable auxiliary boundary law, and checkpoint comparison bounds the actual conditional increment kernels. The auxiliary chain copies the one-time boundary marginals and couples them to sums of independent comparison increments. A fixed-error estimate feeds the shared one-source restart and remaining-time induction, yielding constants uniform over the moment class. This is the mechanism behind the matching saturation, minimum-length and time regimes in (2.12).

The result connects established approximation ideas to an unknown-law prediction problem. FFG's finite-clan construction controls perfect-simulation truncation [FFG02, Sections 3–6]; here the remainder assigns stationary all-hole probability to a cut source that can survive forever. Complete-connection coupling uses temporal kernel continuity [BFG99, CFF02], whereas our rate follows from spatial moment data. GLO compares Gibbs interactions with finite-range versions [GLO10, Theorem 1], and GGLO obtains coupling from arbitrary initial configurations under a subcritical sketch condition and a uniform positive lower bound on site clock rates [GGLO13, conditions (5.1)–(5.2) and Theorem 3]. The present class criterion and matching single-site lower bound quantify history approximation for a selected stationary phase.

Sharpness concerns the deterministic depth of the all-open physical-history truncation. The class supremum in (2.12) allows the lower-bound law to change with the depth. Proposition 7.1 separately quantifies spatial extrapolation to growing targets, and the fresh-row protocol supplies explicit sample and site costs. Establishing valid tail budgets from richer observations and finding matching total prediction costs are concrete further questions. They concern what an expanded observation contract can certify about unobserved recovery-weighted mass and how efficiently that information can be used.

Statements and declarations

Competing interests. The author is employed by a public fund management institution in China. This research was conducted independently and does not represent the views of the author's employer, any fund, or any client. The study did not use non-public employer or client information and does not constitute investment advice.

Data and code. The results are mathematical. Finite diagnostic source and recorded output are packaged separately in `reproducibility.zip` (directory `repro/`); these checks do not replace the proofs. No empirical dataset or simulation of the infinite stationary field is used.

Generative AI and AI-assisted technologies. During the preparation of this work, the author used OpenAI GPT-6.1 Sol and xAI Grok 4.7 to assist with literature synthesis, development of argumentative and theoretical structure, expansion of the author's propositions into manuscript drafts, and language revision. AI-assisted mathematical checking and finite diagnostic code were also used. The author retains responsibility for the research question, interpretation, conclusions and final content. No AI tool or service is credited as an author.

References

[BFG99] X. Bressaud, R. Fernández and A. Galves (1999). Decay of correlations for non-Hölderian dynamics. A coupling approach. *Electronic Journal of Probability* **4**, paper 3, 1–19. DOI; author text.

[BHK06] J. van den Berg, O. Häggström and J. Kahn (2006). Proof of a conjecture of N. Konno for the 1D contact process. *IMS Lecture Notes–Monograph Series* **48**, 16–23. DOI; author text.

[BLMWY26] E. Bayraktar, F. Lu, M. Maggioni, R. Wu and S. Yang (2026). Probabilistic cellular automata with local transition matrices: Synchronization, ergodicity, and inference. *Bernoulli* **32**(2), 1073–1097. DOI; author version, arXiv:2405.02928v4, 22 May 2025.

[CK15] A. Carpentier and A. K. H. Kim (2015). Adaptive and minimax optimal estimation of the tail coefficient. *Statistica Sinica* **25**(3), 1133–1144. DOI; published full text.

[CFF02] F. Comets, R. Fernández and P. A. Ferrari (2002). Processes with long memory: regenerative construction and perfect simulation. *Annals of Applied Probability* **12**(3), 921–943. DOI; author text.

[DLL25] E. De Santis, K. Laxa and E. Löcherbach (2025). A new look at perfect simulation for chains with infinite memory. Preprint, arXiv:2510.24996v1, 28 October 2025. DOI; full text.

[FFG02] P. A. Ferrari, R. Fernández and N. L. Garcia (2002). Perfect simulation for interacting point processes, loss networks and Ising models. *Stochastic Processes and their Applications* **102**(1), 63–88. DOI; author text.

[FKM18] S. Foss, T. Konstantopoulos and T. Mountford (2018). Power Law Condition for Stability of Poisson Hail. *Journal of Theoretical Probability* **31**(2), 684–704. Published online 24 November 2016. DOI; author text.

[GLO10] A. Galves, E. Löcherbach and E. Orlandi (2010). Perfect simulation of infinite range Gibbs measures and coupling with their finite range approximations. *Journal of Statistical Physics* **138**(1–3), 476–495. Published online 1 December 2009. DOI; author text.

[GGLO13] A. Galves, N. L. Garcia, E. Löcherbach and E. Orlandi (2013). Kalikow-type decomposition for multicolor infinite range particle systems. *Annals of Applied Probability* **23**(4), 1629–1659. DOI; published reprint.

[GKM26] P. Gracar, M. Korfhage and P. Mörters (2026). Robustness in the Poisson Boolean model with convex grains. Preprint, arXiv:2410.13366v2, 12 October 2026; first posted 17 October 2024. DOI; full text.

[GT22] P. Gács and I. Törmä (2022). Stable Multi-Level Monotonic Eroders. *Theory of Computing Systems* **66**(1), 322–353. Published online 26 October 2021. DOI and full text; author text.

[H26a] HUANG Man (2026). An exact moment criterion for survival under renewed spatial catastrophes. *Zenodo preprint*, Version 1, 3 November 2026. Version DOI.

[H26b] HUANG Man (2026). Finite-preparation response under heavy-tailed catastrophes: exact amplitudes and separation from Poisson coverage. *Zenodo preprint*, Version 1, 3 November 2026, with Online Resource 1: Effective threshold supplement. Version DOI.

[JK26] B. Jahnel and J. Köppl (2026). Restriction and mixing properties of interacting particle systems with unbounded range. Preprint, arXiv:2603.21817v1; full text.

[LP18] G. Last and M. Penrose (2018). *Lectures on the Poisson Process*. Cambridge University Press, Institute of Mathematical Statistics Textbooks 7. DOI; author version, 21 August 2017; publisher copyright page.

[MW24] T. Mountford and Z. Wang (2024). Stable systems with power law conditions for Poisson hail. *Advances in Applied Probability* **56**(1), 315–353. Published online 22 August 2023. DOI; author text.

[SST25] J. M. Swart, R. Szabó and C. Toninelli (2025). Peierls bounds from random Toom contours. Preprint, arXiv:2501.09660v1, 16 January 2025. DOI; full text.

Appendix A. Conditional correlation, positive phase and spatial tree comparison

For the nonzero intensity sequence in the main text, write $\varrho = \sum_\ell \theta_\ell$, $\nu_\ell = \theta_\ell/\varrho$, $A_* = \sum_\ell \nu_\ell w_\ell$, and $c = \varrho A_* = \chi_\theta$. Thus $S_1 = \sum_\ell \ell\nu_\ell < \infty$. The point-hit rate defined in Section 2.1 is $\lambda_\theta = \varrho S_1$. In last-row interpolation formulas, α is a scalar intensity ranging from zero to ϱ; it is distinct from the fixed primitive intensity sequence. We suppress the subscript of the selected stationary law, write $H^* = H_\theta$, $u = 1 - \beta$, and set $V(r) = w_r$. The zero law is handled directly in Section 2.

The conditional path-resampling construction is adapted from [BHK06]. The interval-specific insertion, phase and tree arguments shared with the companion manuscripts are reproduced below.

A.1. Conditional negative association

Lemma A.1. At every finite time, conditional on the origin being open, bounded increasing functions depending on finitely many hole coordinates strictly to its left and strictly to its right are negatively correlated. In a finite model, the statement allows arbitrary independent type-dependent probabilities. In the infinite model used here, the type means in a finite history are bounded by $\varrho\nu_r$, including a final row of intensity $\alpha\nu_r$ with $0 \le \alpha \le \varrho$.

Proof. First restrict to finitely many interval types in a finite slab containing the relevant backward cones. Set $Y_e = 1$ when disaster type e is absent. The variables are independent with parameters strictly between zero and one. The bottom row is fully occupied. Let Γ be the event that a path reaches terminal site 0, and let μ be the product law conditioned on Γ.

The pointwise minimum and maximum of two allowed paths are still allowed: their increments remain in $\{0, 1\}$. For open paths these operations choose an already open site in each row. Every configuration in Γ therefore has a leftmost path L and a rightmost path R to 0. For a specified path π, a connected disaster interval in its row is strictly left of π, intersects π, or is strictly right of π. Denote these disjoint type sets by $\mathcal{L}_\pi, \mathcal{C}_\pi, \mathcal{R}_\pi$.

Define a right update by resampling the variables in $\mathcal{R}_L$ from their original independent laws. This leaves L unchanged. Indeed, if a new path p appeared partly to the left of L, then $\min(p, L)$ would have been open before the update, because strictly right types cannot affect it. This contradicts leftmostness. Moreover, the event $L = \pi$ depends only on $\mathcal{L}_\pi \cup \mathcal{C}_\pi$: existence of a path further left can be tested among paths everywhere at or left of π. Conditional on the remaining variables and L, the right variables retain their product distribution. The right update thus preserves μ. Resampling $\mathcal{L}_R$ analogously preserves R and μ.

Alternating these kernels gives a finite irreducible, aperiodic chain. To verify irreducibility, set all resampled variables to 1 in a complete right-then-left step, an event of positive probability. After the right step the vertical path at 0 is open: any interval hitting it either hits the old leftmost path and was already absent, or is strictly to its right and has just been removed. The vertical path is the largest allowed path to 0. The left step then removes every remaining disaster strictly to its left; all other types are already absent. Thus every state can reach the all-1 state. Since μ is invariant and has full support on Γ, there can be no transient states or another closed class. Self-loops have positive probability.

Define a comparison between two configurations in Γ by $\widehat{\omega} \preceq \omega$ if (i) $\widehat{L} \le L$ and $\widehat{R} \le R$ rowwise, (ii) $\widehat{Y}_e \ge Y_e$ for types strictly left of $\widehat{L}$, and (iii) $\widehat{Y}_e \le Y_e$ for types strictly right of R. No ordering is imposed on the remaining variables. This is the interval-type version of the comparison in [BHK06, Section 3.1]. For a fixed absence vector y, write the updates explicitly as

$$(T_R^y Y)_e = \begin{cases} y_e, & e \in \mathcal{R}_{L(Y)}, \\ Y_e, & e \notin \mathcal{R}_{L(Y)}, \end{cases}$$

$$(T_L^y Y)_e = \begin{cases} y_e, & e \in \mathcal{L}_{R(Y)}, \\ Y_e, & e \notin \mathcal{L}_{R(Y)}. \end{cases}$$

The same vector y is used for the two configurations; each configuration selects its own resampling set.

A common right update preserves the comparison. Write $\alpha = \widehat{L} \le L = \gamma$ and let p be open after the update in $\widehat{\omega}$. We check $q = \max(p, \gamma)$ row by row in the updated ω. At a site of γ, every covering type is absent because γ remains open. At a site $p_t > \gamma_t$, a covering interval either intersects γ_t, in which case it is absent in ω, or has its entire interval strictly right of γ_t. In the latter case it is selected in both configurations and receives the same entry y_e; openness of p forces that entry to be 1. Connectedness of the interval exhausts these cases. Thus q is open, proving the order of the new rightmost paths. The leftmost paths are unchanged. Types strictly left of α are unchanged in both configurations. A type strictly right of the new rightmost path in ω is also strictly right of γ and α, so it was selected in both configurations and now has equal values. These observations verify all three comparison conditions after the right update.

For a common left update, set $\alpha = \widehat{R} \le R = \gamma$. Given a path p open after the update in ω, check $q = \min(p, \alpha)$ in the updated $\widehat{\omega}$. At a site of α, all covering types are absent. At a site $p_t < \alpha_t$, a covering interval either meets α_t and is absent in $\widehat{\omega}$, or is strictly left of both α_t and γ_t and receives the same entry y_e. This proves that q is open and orders the new leftmost paths. The rightmost paths are unchanged. Every type strictly left of the new $\widehat{L}$ is strictly left of both old rightmost paths, so it was identically resampled. Types strictly right of γ are unchanged in both configurations. Hence this update also preserves all three conditions, including when the two selected type sets differ.

For a single right update from identical input configurations, decreasing the resampled absence variables adds only right-side disasters. The old leftmost path stays open, and the new rightmost path can only move left; all types beyond the larger rightmost path have the required absence order, while types left of the common leftmost path are unchanged. This establishes the comparison after that update. For a single left update, increasing the resampled absences adds only left-side open paths; the rightmost path stays fixed and the leftmost path can only move left, with the required order on the two outer type sets. The common-update preservation just proved propagates either comparison through subsequent updates. Finally, if $\widehat{\omega} \preceq \omega$, every path to a terminal $x < 0$ open in ω yields an open path in $\widehat{\omega}$ by taking its minimum with $\widehat{L}$. An interval covering a strictly left point either intersects $\widehat{L}$ or belongs to its strictly left type set, where condition (ii) applies. Terminal reachability on the left therefore increases, while reachability on the right decreases by the corresponding maximum construction and condition (iii).

Realize the resampling chain on a product space by sampling arrays $Y^R_{k,e}$ and $Y^L_{k,e}$ in advance for every update index k and every retained interval type e. Use an entry only when its type belongs to the currently selected resampling set; otherwise ignore it. All these entries are independent with the original type-specific absence probabilities. If one entry is changed, the two chains have identical histories before its update. At that update either the entry is ignored in both chains or the single-update comparison established above applies. Every subsequent common update preserves that comparison, even when the two resampling sets differ. Thus, after any fixed number of updates, every left increasing hole functional is decreasing in the left-update entries and increasing in the right-update entries. Right increasing functionals have the opposite signs. Complementing the right-update entries places the first functional and the negative of the second in the same coordinatewise decreasing class. The product-measure association inequality gives their nonnegative covariance, equivalently the desired negative correlation. A direct proof for the type-dependent parameters used here is given below. Start at the all-1 configuration and let the finite conditional chain converge to μ to obtain the asserted conditional inequality.

For the infinite model, first fix the time and the finite terminal coordinates of the functions. Their backward cones are finite. Remove types of length greater than K; the total mean of omitted types intersecting the cones is at most a constant times $\varrho T \sum_{r>K} \nu_r(r+C)$, which tends to zero by $S_1 < \infty$. The conditional event has probability at least that of an open vertical path, $e^{-\varrho T S_1} > 0$. Thus the conditional inequality passes to the limit. Deterministic absence variables can be removed or obtained by parameter limits. The same argument applies to the stated type-dependent parameters and last-row interpolation. □

The product-measure correlation bound follows by induction over independent Bernoulli coordinates. For bounded increasing F, G, condition on the last coordinate Y: the total covariance is

$\mathbb{E}\,\mathrm{Cov}(F, G \mid Y) + \mathrm{Cov}(\mathbb{E}[F \mid Y], \mathbb{E}[G \mid Y])$. The first term is nonnegative by induction; the second is $p(1-p)\Delta F \Delta G \geq 0$, where $p = \Pr(Y = 1)$ and $\Delta F = \mathbb{E}[F \mid Y = 1] - \mathbb{E}[F \mid Y = 0]$, with ΔG defined similarly. Negating one function gives the opposite-monotonicity form. The argument permits different parameters at different coordinates.

The proof adapts the extremal-path resampling strategy of [BHK06] to connected interval errors. For a disconnected error closing $\{-1, 1\}$ with probability $1/2$ while leaving 0 open, the two side-hole indicators have joint probability $1/2$ and product of marginals $1/4$. The conclusion is therefore false for arbitrary spatially correlated row errors.

A.2. Finite-history phase construction

Lemma A.2 (full insertion count). In a translation-invariant hole field F, insert $[0, r-1]$. Let L_z be the closed run immediately to the left of z, and R_r the closed run beginning at r. All indicator-weighted run expressions in this appendix are interpreted by restriction: $\mathbf{1}_A Y$ equals Y on A and zero off A, including when run lengths in Y are infinite off A. In particular, each summand below is zero when $z \in F$. The count is initially allowed to be infinite; under the hypotheses stated below, (A2) establishes its integrability. The number of newly closed, nonempty finite intervals is

$$\Delta_r(F) = \sum_{z=0}^{r-1} \mathbf{1}_{\{z \notin F\}} (L_z + 1)(r - z + R_r). \tag{A1}$$

Every newly closed interval has a unique first previously open point z. Its left endpoint has $L_z + 1$ choices and its right endpoint has $r - z + R_r$ choices.

Assume $u_F = \Pr(0 \notin F) > 0$ and the conditional inequality of Lemma A.1. If $h_k(F)$ is the k-interval closure probability and $\zeta = \lim_k h_k(F)$, boundary tail sums give

$$\mathbb{E}[\mathbf{1}_{\{z \notin F\}}(L_z + 1)] = u_F + \beta_F - \zeta \leq 1.$$

For $d = r - z$, split $\{z \notin F, [r, r + b - 1] \subset F\}$ by the last open site in $[z, r-1]$, discard other restrictions, and sum over $b \geq 1$. Each of the d possible boundary sites contributes at most $\beta_F - \zeta$, so

$$\mathbb{E}[\mathbf{1}_{\{z \notin F\}} R_r] \leq d\beta_F, \qquad \mathbb{E}[\mathbf{1}_{\{z \notin F\}}(d + R_r)] \leq d.$$

Conditional negative association, followed by bounded truncation and monotone convergence, now yields

$$\mathbb{E}\Delta_r(F) \leq V(r)/u_F. \tag{A2}$$

Lemma A.3 (finite-history bound). If $c \leq 1/32$, then

$$\sup_t \beta_t \leq \frac{1 - \sqrt{1 - 8c}}{2} < \frac{1}{2}. \tag{A3}$$

Proof. Use the last-row interpolation field

$$F_{t,\alpha} = E(H_t) \cup D^{(\alpha)}, \qquad 0 \leq \alpha \leq \varrho,$$

where the first t rows have intensities $\varrho\nu_r$ and the last row has intensities $\alpha\nu_r$. It is exactly a finite-history process of Lemma A.1 with non-identical type parameters; its open density is at least $u_{t+1} > 0$.

Work first on the infinite integer line with only lengths $r \leq K$ retained. This model remains translation invariant. A fixed finite observation and its derivative involve finitely many retained types in a backward cone of this spatially translation-invariant model. Define $\phi_M(F, x) = \sum_{j=1}^{M} \mathbf{1}_{\{[x, x+j-1] \subset F\}}$. Poisson differentiation of the last row, followed by translation mass counting on this infinite-line law, gives

$$\frac{d}{d\alpha}\mathbb{E}\phi_M(F_{t,\alpha},0)=\sum_r \nu_r\,\mathbb{E}\Delta_{r,\le M}(F_{t,\alpha})\le A_*/u_{t+1}. \tag{A4}$$

Here $\Delta_{r,\le M}$ counts new intervals of length at most M. The same upper bound holds for all K since removal of long disasters only increases openness. Integrate over α and let $K\to\infty$ by bounded local convergence. Then let $M\to\infty$ in the nonnegative insertion increments. Starting from zero, induction proves that $\Phi_t:=\sum_{j\ge1}h_{j,t}<\infty$ for each finite t: erosion does not increase it, and each row adds at most c/u_{t+1}. Since erosion replaces $h_{j,t}$ by $h_{j+1,t}$,

$$\Phi_{t+1}-\Phi_t\le c/u_{t+1}-\beta_t. \tag{A5}$$

Adding past rows makes the hole marginals larger, so β_t is nondecreasing. A single new row gives $\beta_{t+1}\le\beta_t+q$ with $q=1-e^{-\varrho S_1}\le c$. Independent time-block concatenation gives $u_{a+b}\ge u_a u_b$: choose, by a fixed rule using only the last a rows, an open path to 0 and its initial site z when one exists. The first b rows reach that site with probability u_b, independently. The shared boundary row is checked by the first block only. Thus $\beta_{a+b}\le\beta_a+\beta_b$.

Summing this inequality over complementary time indices, and summing (A5) while retaining $\Phi_T\ge0$, gives

$$\frac{T-1}{2}\beta_T\le\sum_{t<T}\beta_t\le Tc/u_T,\qquad \beta_T(1-\beta_T)\le\frac{2T}{T-1}c\le 4c\quad(T\ge2). \tag{A6}$$

If T is the first crossing of $1/2$, then $1/2\le\beta_T\le1/2+c$. Its product with $1-\beta_T$ is at least $1/4-c^2\ge255/1024>1/8\ge4c$, contradicting (A6). Hence the crossing never occurs. Taking $T\to\infty$ in the sharper version of (A6) gives $\beta(1-\beta)\le2c$ on the small branch, proving (A3). The constant $1/32$ supplies a uniform positive-density regime. □

A.3. Stationary flux

The stationary field is a translation-equivariant measurable function of the independent columns of marked Poisson variables, hence is spatially ergodic. Equation (A3) implies $u>0$, so it has no all-closed component and $h_j\to0$. A stationary field cannot have a closed half-line with a finite endpoint at positive probability: the events specifying its unique endpoint are disjoint translates with equal probabilities. Thus closed runs are finite almost surely. The interpolation fields $F_\alpha=E(H^*)\cup D^{(\alpha)}$ are coupled below F_ϱ, which has law H^*; they inherit this property and have open densities $u_\alpha\ge u$.

Lemma A.4 (local stationary interpolation). Let $\Delta_{r,j}(F)$ count new closed intervals of length j after inserting $[0,r-1]$. Only the $r+j-1$ windows intersecting the insertion can change, so

$$0\le\Delta_{r,j}(F)\le r+j-1.$$

For fixed j, this bound is integrable under ν by $S_1<\infty$. Approximate H^* by finite histories and truncated lengths. Each local insertion difference converges and is bounded by the same integrable envelope. Dominated convergence passes the last-row interpolation to stationarity, giving

$$q_j:=h_j-h_{j+1}=\int_0^\varrho\sum_r\nu_r\mathbb{E}\Delta_{r,j}(F_\alpha)\,d\alpha. \tag{A7}$$

The conditional inequality also passes to these fields from Lemma A.1, since their open densities are bounded below by $u>0$.

In a stationary field with $h_j(F)\to0$, the boundary tail identity sharpens to $\mathbb{E}[\mathbf{1}_{\{z\notin F\}}L_z]=\beta_F$. Substituting it in (A1) yields

$$\mathbb{E}\Delta_r(F)=V(r)+\sum_{z=0}^{r-1}\mathbb{E}[\mathbf{1}_{\{z\notin F\}}(L_z+1)R_r]\ge V(r).$$

For the interpolation fields under consideration, $h_j(F) \to 0$. Combining the identity with (A2) and summing (A7) by Tonelli gives

$$\beta = \int_0^{\varrho} \sum_r \nu_r \mathbb{E}\Delta_r(F_\alpha)\, d\alpha, \qquad c \le \beta \le c/u. \tag{A8}$$

The inequality $\beta < 1/2$ selects the small root of $\beta(1-\beta) \le c$, giving $\beta \le (1-\sqrt{1-4c})/2$. Also $\beta \le 2c$ and $\beta - c \le \beta^2 \le 4c^2$. The phase bound (A3) determines the branch used in the stationary flux inequality.

A.4. Branching comparison and the spatial tail bound

Put $q_0 = u$ and $q_j = h_j - h_{j+1}$ for $j \ge 1$. This is a probability mass function on the nonnegative integers; let Q have this law, so $\Pr(Q \ge k) = h_k$. For an interpolation field F, define $h_0^F = 1$ and $q_a^F = h_a^F - h_{a+1}^F$ for $a \ge 0$. This is also the probability that a site is open and the closed run immediately to its left has length at least a.

Assign each new interval of length j in (A1) to its first old open point z. Let its left and right endpoint distances from z be a, b with $a + b = j - 1$, and put $d = r - z$. If $b < d$, its right requirement is automatic. If $b \ge d$, partition by the last old open site in $[z, r-1]$ and discard additional restrictions. Both cases give

$$\Pr(z \notin F,\ d + R_r \ge b + 1) \le \sum_{i=0}^{\min(d-1,b)} q_{b-i}^F.$$

Conditional negative association separates the left run event from this right requirement. Summing over endpoints and z gives

$$\mathbb{E}\Delta_{r,j}(F) \le \frac{1}{u_F} \sum_{i=0}^{r-1} (r-i)(q^F * q^F)_{j-1-i}, \tag{A9}$$

with negative subscripts zero. In the coupling $F_\alpha \subset F_\varrho$, Q_α is stochastically bounded by Q. Applying this stochastic domination to tails, sum (A7) over $j \ge k$ and use (A9) to obtain

$$h_k \le \frac{\varrho}{u} \sum_r \nu_r \sum_{i=0}^{r-1} (r-i) \Pr(Q_1 + Q_2 + i \ge k - 1). \tag{A10}$$

Here Q_1, Q_2 are independent auxiliary copies of Q. Define I by

$$\Pr(I = i) = \frac{\sum_r \nu_r (r-i)_+}{A_*}, \qquad \Pr(I \ge m) = A_{m+1}/A_*.$$

With $a = c/u$, (A10) says $Q \le_{\mathrm{st}} B(1 + I + Q_1 + Q_2)$, all variables on the right independent and B Bernoulli with parameter a. Iterate the distributional upper map on a binary tree. The chance of any active leaf at depth n is at most $(2a)^n$. On its complement the depth-n comparison no longer depends on the leaf copies of Q. Since $2a < 1$, this proves

$$Q \le_{\mathrm{st}} X = \sum_{i=1}^{J} W_i, \qquad W = 1 + I,$$

where J is the number of successful nodes, success having probability a and producing two children. The costs are independent of the tree, and

$$\Pr(J = j) = \mathrm{Catalan}(j) a^j (1-a)^{j+1} \le (4a)^j, \qquad \mathbb{E}J = \frac{a}{1-2a}. \tag{A11}$$

Here $A_j = \sum_{r \ge j} \nu_r (r - j + 1)(r - j + 2)/2$, so the positive cost $W = 1 + I$ has tail $A_\theta(j)/\chi_\theta$. The tree therefore has exactly the law in Lemma 4.1. Root inactivity gives $h_1 \le c/u$; no subexponential assumption is needed for this comparison.

Appendix B. Compression and deterministic checkpoint comparison

In B.1 retain Appendix A's normalization $\theta_r = \varrho\nu_r$. In B.2 all independent type parameters may vary deterministically by row and type.

B.1. Compression of a finite target

For a finite target $U = \{x_1 < \cdots < x_q\}$ and a gap $x_{i+1} - x_i \geq 2$, let U' shift the right suffix one site to the left. We prove $p_t(U') \geq p_t(U)$ and hence

$$p_t(U) \leq p_t(\{0, \ldots, q-1\}), \qquad p_\varrho(U) \leq h_\varrho(|U|). \tag{B1}$$

The one-row statement needed for the induction is an order on survivor ranks. A set function f is supermodular if $f(A \cup C) + f(A \cap C) \geq f(A) + f(C)$. On a finite rank set, the disaster row is a pure-death chain with generator

$$(\mathcal{L}_U f)(S) = \varrho \sum_r \nu_r \sum_{\ell \in \mathbb{Z}} [f(S \setminus J_{r,\ell}(U)) - f(S)], \quad J_{r,\ell}(U) = \{j : \ell \leq x_j < \ell + r\}.$$

Its total rate on every nonempty state is at most $\varrho|S|S_1 < \infty$.

Fix S and expand $F(A) = f(S \setminus A)$ in Boolean Möbius coefficients. The number of length-r origins covering every member of a nonempty rank subset T is $(r - x_{\max T} + x_{\min T})_+$. Compression changes this number by one exactly when T crosses the gap and its original diameter is at most r. Group the coefficients by their extreme ranks $j \leq i < k$. The intermediate-rank Möbius sum becomes a mixed difference on the survivor set outside $[j, k]$. Thus

$$\begin{aligned}(\mathcal{L}_{U'} - \mathcal{L}_U) f(S) = \varrho \sum_{\substack{j,k \in S \\ j \leq i < k}} \overline{\nu}_{x_k - x_j} \big[& f(T_{jk}) + f(T_{jk} \cup \{j,k\}) \\ & - f(T_{jk} \cup \{j\}) - f(T_{jk} \cup \{k\}) \big] \geq 0,\end{aligned} \tag{B2}$$

where $T_{jk} = S \setminus [j,k]$ in rank coordinates and $\overline{\nu}_d = \sum_{r \geq d} \nu_r$. Deleting a fixed set of ranks preserves supermodularity, so each row semigroup preserves it. Write $P_U(s) = \exp(s\mathcal{L}_U)$ for the semigroup on rank subsets, where s is the Poisson exposure parameter and one fresh row corresponds to $s = 1$. Duhamel's formula gives

$$(P_{U'}(1) - P_U(1)) f = \int_0^1 P_{U'}(s)(\mathcal{L}_{U'} - \mathcal{L}_U) P_U(1-s) f \, ds \geq 0.$$

The inner semigroup preserves supermodularity, (B2) gives a nonnegative function, and the outer semigroup preserves nonnegativity. Thus every supermodular function has larger expectation under the compressed one-row survivor law.

Induct on physical time, proving the single-gap inequality simultaneously for every finite target and every admissible gap, with no restriction on target cardinality. At time $t = 0$, $p_0(\varnothing) = 1$ and $p_0(A) = 0$ for nonempty A; compression preserves emptiness, so the claim holds. Assume the claim at time t. The function $p_t(A) = \Pr(A \subset H_t)$ is decreasing and supermodular. Put $\Psi(A) = A \cup (A - 1)$. The function $F_U(S) = p_t(\Psi(U_S))$ is supermodular, because Ψ preserves unions and $\Psi(U_{S \cap T}) \subset \Psi(U_S) \cap \Psi(U_T)$. For fixed surviving ranks S, put $L = U_S \cap (-\infty, x_i]$ and $R = U_S \cap [x_{i+1}, \infty)$, so $U'_S = L \cup (R - 1)$. If either part is empty, the predecessor targets are translates. Otherwise write $\ell = \max L$, $r = \min R$, $d = r - \ell \geq 2$, $A = \Psi(L)$ and $B = \Psi(R)$. Then

$$\Psi(U_S) = A \cup B, \qquad \Psi(U'_S) = A \cup (B - 1), \qquad \max A = \ell, \quad \min B = r - 1.$$

If $d \geq 3$, the separating distance between A and B is $d - 1 \geq 2$; shifting B left by one is a permitted gap compression, so the time-t induction hypothesis compares the two predecessor probabilities. If $d = 2$, the shifted pieces meet at ℓ. Delete the preimage of this duplicated requirement and set

$$C_0 = A \cup (B \setminus \{\ell + 1\}).$$

The remaining right piece is nonempty, starts at $r = \ell + 2$, and has separating distance two from A. Compressing it by one gives exactly $A \cup (B - 1)$, because the removed point would shift to $\ell \in A$. The induction hypothesis and decreasingness under target inclusion give

$$p_t(\Psi(U'_S)) \geq p_t(C_0) \geq p_t(\Psi(U_S)).$$

Thus the predecessor probabilities compare pointwise for every S, including internal gaps and deleted ranks. This is the predecessor comparison used in [H26b, Appendix B.1]. Let $\mathcal{S}_U, \mathcal{S}_{U'}$ be the surviving rank sets under the two respective fresh-row laws, and put $F_{U'}(S) = p_t(\Psi(U'_S))$. The kill-before-expansion recursion gives

$$p_{t+1}(U) = \mathbb{E}_D p_t(\Psi(U \setminus D)) = \mathbb{E}F_U(\mathcal{S}_U).$$

The pointwise predecessor comparison, followed by the one-row order applied to the supermodular function F_U, yields

$$p_{t+1}(U') = \mathbb{E}F_{U'}(\mathcal{S}_{U'}) \geq \mathbb{E}F_U(\mathcal{S}_{U'}) \geq \mathbb{E}F_U(\mathcal{S}_U) = p_{t+1}(U).$$

This proves the time-$t+1$ claim for every finite target and closes the induction. Iterated gap closure and the increasing remote-past limit prove (B1). The one-row order applied to $\mathbf{1}_{\{S=\varnothing\}}$ gives the single Boolean-row full-coverage comparison. The argument also applies to every length truncation.

B.2. Conditioning on a deterministic checkpoint

Lemma B.1 (checkpoint comparison). Consider a finite time slab with increments $0, 1$, a fully open initial row and independent absence variables for finitely many connected interval types. Fix a checkpoint row r and an integer c. A constrained path has coordinate at most c at row r. Let H^0, H^c be the terminal hole sets for unrestricted and constrained paths, coupled by the same types. At a deterministic terminal site z, put $A_z = \{z \notin H^0\}$ and $O_z = \{z \notin H^c\}$. If $\Pr(O_z) > 0$, then:

(i) On O_z, the two fields agree strictly left of z.

(ii) Given O_z, increasing functions of constrained holes strictly left and strictly right of z are negatively correlated.

(iii) If f is a bounded increasing function of unrestricted holes strictly left of z, then

$$\mathbb{E}[f(H^0) \mid O_z] \leq \mathbb{E}[f(H^0) \mid A_z]. \tag{B3}$$

If Z is the rightmost constrained open site and $\Pr(Z = z) > 0$, then

$$\mathbb{E}[f(H^0) \mid Z = z] \leq \mathbb{E}[f(H^0) \mid A_z]. \tag{B4}$$

The statements pass to the infinite interval-type model at a fixed finite time whenever $\sum_\ell \ell\theta_\ell < \infty$, and to increasing all-open prehistories when the limiting conditioning probabilities are positive.

Proof. Work in a finite spatial slab containing every backward cone used in the observations. Initially take all type parameters strictly between zero and one. If an unrestricted open path ends at $y < z$ and a constrained open path ends at z, their minimum is a constrained open path ending at y. Both increments and the checkpoint restriction are preserved. This proves (i).

For (ii), condition on existence of a constrained path to z. Resample types strictly right of its leftmost path, and then types strictly left of its new rightmost path, exactly as in Lemma A.1. Each conditional kernel preserves the conditioned product measure: the extremal path is unchanged, and specifying it reads only the types on its other side and crossing it. Minima and maxima preserve the checkpoint, so the common-update comparison in Lemma A.1 is unchanged.

We check the irreducibility step with the constraint included. If the terminal row is n, the greatest feasible all-absent path to z is

$$\pi_s = \min\{z, c + (s - r)_+\}, \qquad 0 \le s \le n. \tag{B5}$$

Feasibility gives $z \le c + n - r$. Before the checkpoint every feasible path lies below both z and c; after it, it lies below both z and $c + s - r$. Thus every constrained path is below π. In a right-then-left update, force every selected absence variable to one. After the right update π is open: an interval hitting it either crosses the old leftmost path and was absent already, or lies strictly to its right and was cleared. The rightmost path is now π. The left update clears every remaining type strictly left of it. Types crossing π are absent and those strictly right were already cleared. Hence the all-absent state is reachable. The conditional invariant measure has full support, so the finite chain is irreducible; positive self-loops give convergence.

For clarity, the common comparison orders both extremal paths leftward, the absence variables upward strictly left of the hatted leftmost path, and downward strictly right of the unhatted rightmost path. For a common right update, a hatted path p maps to $\max(p, L)$ in the other configuration. A covering type at a point right of L either intersects L, and is absent, or is strictly right of both leftmost paths, and is identically resampled. This orders the rightmost paths; the leftmost paths stay fixed. Types beyond the new rightmost path were identically resampled, and types left of the old hatted leftmost path were untouched. A common left update uses $\min(p, \widehat{R})$ and the same interval dichotomy, preserving all three orders. The checkpoint is preserved in both path operations.

Increasing a left-update absence entry moves the configuration left in this comparison; decreasing a right-update entry has the same effect. Terminal left-hole functions then decrease in left-update entries and increase in right-update entries, with opposite signs on the right. Independent pre-sampled update arrays, product association and convergence of the conditional chain prove (ii), just as in Lemma A.1.

For (iii), instead run the unrestricted conditional chain given A_z and let P be its leftmost path to z. On A_z, a constrained path exists if and only if $P_r \le c$. This event increases in left-update absence entries and decreases in right-update entries, opposite to the monotonicities of a left increasing hole function. Product association therefore yields

$$\mathbb{E}[f\mathbf{1}_{\{P_r \le c\}} \mid A_z] \le \mathbb{E}[f \mid A_z] \Pr(P_r \le c \mid A_z).$$

Division proves (B3). On O_z, part (i) identifies $f(H^0)$ with its constrained version. The edge event adds only the right increasing hole event $[z + 1, c + n - r] \subset H^c$. Apply (ii), then (B3), to obtain (B4). Empty right intervals impose no additional condition.

For infinitely many types at a fixed time and fixed terminal coordinates, truncate lengths at K. The mean of omitted atoms meeting their finite cones is bounded by a finite constant times $\sum_{\ell > K}(\ell + C)\theta_\ell \to 0$. A feasible open path gives a positive lower bound for each conditioning probability as K increases. This passes bounded conditional expectations to the limit; deterministic parameters follow by parameter limits. Finally, extending an all-open prehistory gives monotone convergence of every finite terminal cylinder. A fixed checkpoint restricts only a finite cone for each terminal coordinate, so the same conditional limit works whenever the limiting event has positive probability. This proves the last assertion without a uniform history estimate. □

To justify finite endpoints in the stationary application, fix the number of rows after the checkpoint. Sufficiently far on its permitted side, every ancestor of a target is permitted, so constrained and unrestricted fields coincide there. The stationary unrestricted field has positive open density and is spatially ergodic (Appendix A). It has infinitely many open sites on that side. On the forbidden side there is a deterministic unreachable ray. The constrained boundary is therefore finite. Its event at any fixed location is a finite terminal cylinder, as required in Section 5.2.

Appendix C. Exact error-set boundary for the eroder comparison

Gács and Törmä's ε-perturbation definition requires that, for every finite space-time set A, the probability of a violation at every member of A is at most $\varepsilon^{|A|}$ [GT22, Section 2]. It allows dependence; its restriction is on simultaneous errors.

Choose an admissible positive background $\theta_\ell = c\ell^{-6}$ with $c > 0$ small enough for the declared budgets. Start from all-open sites, so $H_0 = \varnothing$ and the actual rule-violation set at the first update is exactly D_1. A length-k atom at origin zero closes k consecutive sites, giving

$$\Pr([0, k-1] \subset D_1) \geq 1 - e^{-ck^{-6}}. \tag{C1}$$

For every fixed $0 < \varepsilon < 1$, (C1) eventually exceeds ε^k. Thus the common moment class does not satisfy that finite-error-set hypothesis. At the same time, (3.3) gives summable spatial pair covariances for this background. Summable pair dependence cannot substitute for an all-orders error-set condition. This proves the stated direct-application boundary; it does not assert failure of eroder stability under this noise.

The independent space-time local-rule choices in [SST25, (1.7)–(1.8)] give a second precise comparison. Our first-update local rules share interval atoms across sites, so their law lies outside that independent-rule set-up. The moment conditions in this paper control the recovery footprint of those correlated errors.