First recording times and rare clusters in adaptive stabilizer learning

Scott Schubr
Luminary Oddities

November 2026

Abstract

We determine the full distribution of recording times for the adaptive single-copy stabilizer-learning process under global depolarizing noise at a fixed prime. With fixed retained weights, the process has Gaussian central and local laws above its noise threshold. Below threshold, recording times are exponential with explicit rank mixtures, and records arrive in compound-Poisson clusters with geometric multiplicities, both with Wasserstein rates. In the bounded critical window the scaled process converges to a reflected diffusion, and the recording counts follow its boundary local time.

Contents

1 Introduction **1**

2 First recording times and rare clusters in adaptive stabilizer learning **2**
2.1 The rank chain and the recording rule 2
2.2 The critical window 2
2.2.1 The Green operator and the extra recording row 3
2.2.2 Laplace transforms and discrete survival 4
2.3 Fixed retained weight above the threshold 6
2.4 Fixed retained weight below the threshold 7
2.5 Approaching the threshold outside the critical window 9
2.6 Continued recordings in the critical window 12
2.7 Scope of the distributional results 14

1 Introduction

Bittel, Eisert, Gong, Mele and Schatzki [2] showed that adaptive single-copy measurements learn stabilizer states optimally. The sharp noise transition for the mean behaviour of the adaptive process was determined in [1]. This article determines the distributions behind those means. Other learners under global depolarizing noise use two-copy Bell measurements [3]; the process studied here uses single copies only. Exact single-copy recovery under global white noise and under correlated diagonal noise, from $O(n^{1+\varepsilon})$ copies, is given in [5, 4]; the recording-time laws below describe how long the adaptive process takes to reach each record along the way.

2 First recording times and rare clusters in adaptive stabilizer learning

We find the recording-time distributions of the exact rank process for an adaptive single-copy stabilizer learner with global depolarization. A bounded critical window gives a reflected diffusion hitting law, with uniform $O(n^{-1})$ errors on positive diffusive time intervals and an explicit entrance coefficient at each fixed rank. The recording rule changes this entrance coefficient even though the macroscopic diffusion is the same. At a fixed retained weight below the threshold, the first recording time from maximal missing rank is asymptotically exponential with a Wasserstein error, while continued recordings converge to a compound Poisson process with explicit geometric multiplicities. Above the threshold, separate local central limit estimates describe rank exposure and first exact recording. These laws extend to vanishing drifts outside the window. At criticality the continued count converges to an explicit multiple of boundary local time. Independent thinning yields a quantitative Robin first-arrival law.

2.1 The rank chain and the recording rule

An adaptive single-copy stabilizer learner can sample in a changing commuting basis and combine two outcomes to revise its candidate group. Bittel et al. [6, Algorithm 1 and Section III] use drift bounds on an auxiliary rank chain to analyze sample complexity. For global depolarization, [7] derives the exact rank chain and its mean-time transition. We use its joint rank and recording law to find first-candidate confidence times and the distribution of continued recordings. The latent rank r runs from 0 to n, with 0 corresponding to a fully exposed target. Fix a prime p and a retained weight $0 < c \leq 1$. The input is $c|\psi\rangle\langle\psi| + (1-c)I/p^n$ for a target stabilizer state $|\psi\rangle$. Put $s = c^2$, and write a_i, b_i for the downward and upward probabilities. The full chain has

$$a_i = (1-p^{-i})[s + (1-s)p^{i-n}]/p, \qquad b_i = (1-s)(1-p^{i-n}). \tag{1}$$

All remaining probability is a hold. At rank zero, set $h_n = (1-c)p^{-n}$. The probability of an exact pre-update candidate is $F_n = c + h_n$, also the finite-dimensional fidelity of the input with the target. To see the joint transition rule, at rank zero the two outcomes are independent with probability F_n on the target and probability h_n on each other outcome. Equality of those outcomes leaves the rank at zero. Conditional on a recording the next rank is therefore zero with probability F_n, and conditional on failure it is zero with probability h_n. The next rank is one in the respective complementary events. This gives $b_0 = (1-F_n)(1-h_n) + F_n(1-F_n) = (1-c^2)(1-p^{-n})$, consistent with (1). The joint law of the mark and the successor rank, conditional on the past, depends only on the rank; the process may be continued with exactly the same sampling rule after a recording.

Let τ be the first time at rank zero. Let T be the number of rounds through the first exact recording, including the recording round. Thus $T = \sigma + 1$ in the notation of [7]. We denote the starting rank by a subscript on $\mathbb{P}$ and $\mathbb{E}$. The analysis concerns these latent times and the unobserved exact-candidate sequence. No oracle recognizing an exact candidate is required for the distributional statements.

2.2 The critical window

Set $R = p(1-s)/s$ and $c_* = \sqrt{p/(p+1)}$. We let the retained weight depend on n and write $\lambda_n = n \log R_n$. Let $\mathcal{H}_\lambda$ be the hitting time of zero for the diffusion on $[0,1]$ with generator $(\partial_x^2 + \lambda\partial_x)/(p+1)$, reflection at one, and killing at zero. Write $S_\lambda(t,x) = \mathbb{P}_x(\mathcal{H}_\lambda > t)$ and $f_{\lambda,z}(x) =$

$\mathbb{E}_x e^{-z\mathcal{H}_\lambda}$. Define

$$P_j = \prod_{i=1}^{j}(1-p^{-i}), \quad P_0 = 1, \quad C_p = P_\infty, \qquad B_r = C_p^{-1}\sum_{i=1}^{r} P_{i-1}, \qquad B_{\mathrm{rec}} = \frac{1-c_*}{c_* C_p}.$$

Theorem 2.1 (Diffusion and entrance estimates)**.** *Suppose $|\lambda_n| \le \Lambda$. Uniformly over the starting rank r and t in any fixed compact subinterval of $(0,\infty)$,*

$$\mathbb{P}_r(\tau > n^2 t) = S_{\lambda_n}(t, r/n) + O(n^{-1}), \qquad \mathbb{P}_r(T > n^2 t) = S_{\lambda_n}(t, r/n) + O(n^{-1}).$$

For fixed $r \ge 0$ and the same time intervals,

$$n\mathbb{P}_r(\tau > n^2 t) = B_r\, \partial_x S_{\lambda_n}(t,0) + O_r(n^{-1}),$$
$$n\mathbb{P}_r(T > n^2 t) = (B_r + B_{\mathrm{rec}})\, \partial_x S_{\lambda_n}(t,0) + O_r(n^{-1}).$$

For $0 \le z \le A < \infty$, both Laplace transforms at scale n^2 equal $f_{\lambda_n,z}(r/n) + O_A(n^{-1})$ uniformly over r. For fixed r, the corresponding quantities $n(1-\mathbb{E}_r e^{-z\tau/n^2})$ and $n(1-\mathbb{E}_r e^{-zT/n^2})$ equal $B_r[-f'_{\lambda_n,z}(0)]$ and $(B_r + B_{\mathrm{rec}})[-f'_{\lambda_n,z}(0)]$, respectively, up to $O_{r,A}(n^{-1})$. In these formulas, with $\kappa = (\lambda^2/4 + (p+1)z)^{1/2}$ and $H(v) = \cosh(\kappa v) - \lambda \sinh(\kappa v)/(2\kappa)$,

$$f_{\lambda,z}(x) = e^{-\lambda x/2}\frac{H(1-x)}{H(1)}, \qquad -f'_{\lambda,z}(0) = \frac{(p+1)z\,\sinh\kappa}{\kappa H(1)}. \tag{2}$$

The removable values are interpreted continuously.

2.2.1 The Green operator and the extra recording row

Write Q for the transition matrix on positive ranks with killing at zero, or on all ranks with killing on a recording, as appropriate. For exposure, the inverse $G = (I-Q)^{-1}$ is given by the explicit difference equation

$$Gu(r) = \sum_{i=1}^{r}\sum_{k=i}^{n} T_{i,k}u(k), \qquad T_{i,k} = \frac{p}{s}R^{k-i}\frac{\prod_{j=i}^{k-1}(1-p^{j-n})}{\prod_{j=i}^{k}(1-p^{-j})\prod_{j=i}^{k}(1+\frac{1-s}{s}p^{j-n})}. \tag{3}$$

Indeed, if the boundary value of $v = Gu$ is zero, its increments solve $a_i(v(i)-v(i-1)) = u(i) + b_i(v(i+1)-v(i))$. Iterating from $b_n = 0$ gives (3). For recording, denote the inverse by $\widetilde{G}$, put $b_0^{\mathrm{k}} = (1-F_n)(1-h_n)$, and use the same positive-rank equations. The row at zero gives the exact identity

$$\widetilde{G}u(r) = Gu|_{\{1,\dots,n\}}(r) + \frac{b_0^{\mathrm{k}}}{F_n}\,Gu|_{\{1,\dots,n\}}(1) + \frac{u(0)}{F_n}. \tag{4}$$

In fact $v(r) - v(0) = Gu|_{\{1,\dots,n\}}(r)$ and $F_n v(0) - b_0^{\mathrm{k}}(v(1)-v(0)) = u(0)$. This calculation uses the joint recording-transition law.

Let $K_n = n^{-2}G$ or $n^{-2}\widetilde{G}$. Regard it as an integral operator on $[0,1]$ by partitioning into n or $n+1$ equal cells in rank order. It averages the input on each cell and returns the matrix output on the corresponding cell. Each positive rank lies at distance at most $2/n$ from all positions in its cell. The limiting operator satisfies

$$K_\lambda(x,y) = (p+1)\int_0^{\min(x,y)} e^{\lambda(y-v)}\,dv, \qquad D_\lambda(y) = (p+1)e^{\lambda y}. \tag{5}$$

Here D_λ is the right derivative of $K_\lambda(x,y)$ in x at zero for almost all y. Inverting $-(\partial_x^2+\lambda\partial_x)/(p+1)$ with Dirichlet boundary at zero and Neumann boundary at one gives (5).

We record the kernel bounds used below. Products in (3) and their reciprocals are bounded uniformly in the window. Their deviations from the homogeneous kernel satisfy

$$|T_{i,k} - (p+1)e^{\lambda_n(k-i)/n}| \le C(n^{-1} + p^{-i} + p^{k-n}),$$
$$|T_{i,k} - (p+1)(P_{i-1}/C_p)e^{\lambda_n(k-i)/n}| \le C(n^{-1} + p^{-k} + p^{k-n}).$$

For the first line use $p/s = p + R$ and the geometric sums of the omitted lower factors and upper factors. For the second line replace the lower reciprocal product by P_{i-1}/C_p; its missing tail has sum $O(p^{-k})$. These estimates and uniformly Lipschitz Riemann sums give

$$\|K_n - K_{\lambda_n}\|_{\infty\to\infty} \le C/n, \qquad \|K_n\|_{2\to\infty} \le C. \tag{6}$$

For clarity, the exposure step kernel on cell k at row r is $n^{-1}\sum_{i\le\min(r,k)} T_{i,k}$. The pointwise first-line error after summation is $O(n^{-1}+p^{k-n})$, which is $O(n^{-1})$ after integrating in y. The recording partition changes the normalization by $1+1/n$, and adds a rank-zero cell of width $1/(n+1)$. The extra terms in (4) have L^∞-operator norm $O(n^{-1})$. Their step kernels are uniformly bounded, including the rank-zero cell. This proves both claims of (6). For fixed r, the second-line error gives the sharper row statement

$$\|nK_n(r,\cdot) - B\,D_{\lambda_n}\|_{L^1} \le C_r/n, \qquad \|nK_n(r,\cdot)\|_{L^2} \le C_r. \tag{7}$$

Here $B = B_r$ or $B_r + B_{\mathrm{rec}}$, respectively. Indeed the omitted cells $k < r$ have total length $O_r(n^{-1})$, and replacing $e^{-\lambda_n i/n}$ by 1 for $i \le r$ has the same order. For recording, the extra positive-rank term uses the row $r = 1$; $b_0^{\mathrm{k}}/F_n = (1-c_*)/c_* + O(n^{-1})$. The rank-zero input term has bounded n-scaled kernel on its single cell. At $r = 0$ for exposure the zero row is understood throughout.

2.2.2 Laplace transforms and discrete survival

For a killed integer lifetime L with transition matrix Q, the one-step equation gives

$$f_n(z) = \big(I + n^2(e^{z/n^2} - 1)K_n\big)^{-1}\mathbf{1}.$$

The inverse has sup norm at most two for real $z \ge 0$. Indeed, on the finite rank space it equals $I - \delta(\delta + I - Q)^{-1}$, where $\delta = e^{z/n^2} - 1$; the second term is a positive contraction by the geometric series for $Q/(1+\delta)$. The bound also holds on the complementary subspace of cell-mean-zero functions, using the positive averaging projection in the same identity. For the limiting problem, $f_{\lambda,z} = 1 - zK_\lambda f_{\lambda,z}$ by the boundary equation. Using (6) gives the uniform transform claim. The finite difference $1 - f_n(z)$ equals $n^2\delta K_n f_n(z)$, so (7) proves the sharper fixed-rank assertion. Solving $f'' + \lambda f' = (p+1)zf$, $f(0) = 1$, $f'(1) = 0$ gives (2). The endpoint $r = 0$ for exposure has exactly unit transform.

For the survival estimate, the finite matrices Q are reversible. Their positive-rank detailed-balance weights satisfy $\pi_{i+1}/\pi_i = b_i/a_{i+1}$. The window product bounds in (3) give upper and lower constants for the ratio of any two positive-rank weights: use R_n^{k-i}, the geometric boundary products and the bounded ratios a_i/a_k. A recording row adds the ratio b_0^{k}/a_1, also bounded above and below. Thus the normalized weights are comparable to the uniform cell weights. In the step-function embedding, extend the weights to cellwise-constant densities. The operator K_n is then positive self-adjoint in the weighted L^2, vanishing on the orthogonal complement of the step

functions. The continuum analogue holds for K_λ with weight $e^{\lambda x}$ by integration by parts. All these L^2 norms are uniformly comparable.

Choose an outer radius larger than twice the operator norm and fixed rays at angles $\pm\vartheta$ with $0 < \vartheta < \pi/2$. For $K = K_n$ or K_{λ_n} and w on this contour, the L^2 spectral bound and (6) give

$$\|(w - K)^{-1}\|_{\infty\to\infty} \le C(|w|^{-1} + |w|^{-2}). \tag{8}$$

Use $w^{-1}(I + K(w-K)^{-1})$ to see the smoothing term. Put $\psi_t(w) = e^{-t/w}$, with value zero assigned to the spectral point zero. The same convention is used for $\psi_t(w)/w$. The spectral theorem identifies these functions of K with the integral of $\psi_t(w)(w-K)^{-1}/(2\pi i)$, or the corresponding quotient, on the two rays and outer arc. For a positive scalar spectral value this follows by closing the contour at a smaller radius and taking that radius to zero. At spectral value zero the integral is zero by the same closing argument. The exponential decay on the sector makes the integrals norm convergent. The resolvent identity and (8) now give

$$\|\psi_t(K_n) - \psi_t(K_{\lambda_n})\|_{\infty\to\infty} + \|(\psi_t/w)(K_n) - (\psi_t/w)(K_{\lambda_n})\|_{\infty\to\infty} \le C/n. \tag{9}$$

The bound is uniform on the stated time intervals since $|e^{-t/w}| \le e^{-t\cos\vartheta/|w|}$. The diffusion survival is $\psi_t(K_{\lambda_n})\mathbf{1}$ by its Dirichlet–Neumann heat equation. Factoring it as $K_{\lambda_n}(\psi_t/w)(K_{\lambda_n})\mathbf{1}$ shows uniform Lipschitz continuity in x. Its right derivative at zero is obtained by applying the row D_{λ_n} from (5). Together with (7), (9) therefore proves both continuum comparisons for $e^{-tn^2(I-Q)}\mathbf{1}$.

It remains to replace the exponential by the actual discrete lifetime. In the window, every positive-rank hold probability is bounded below by $s - 1/p$, hence by a positive constant for large n. At a recording row the off-diagonal sum b_0^{k} is bounded away from one. Gershgorin's theorem and reversibility give $\mathrm{spec}(I - Q) \subset (0, 2-\eta]$ for a fixed $\eta > 0$. For $k = \lfloor n^2 t \rfloor$,

$$\sup_{0<x\le 2-\eta} (1 + n^2 x)|(1-x)^k - e^{-n^2 t x}| = O(n^{-2}). \tag{10}$$

For $x \le 1/2$, the unweighted difference is bounded by $C(kx^2 + x)e^{-kx}$, by $-\log(1-x) - x = O(x^2)$ and the rounding error. On $[1/2, 2-\eta]$ it decays exponentially in k. For the full rows, converting weighted L^2 to step-function sup norm costs only $O(\sqrt{n})$. For fixed rows multiplied by n, factor K_n out of the matrix difference and use the second bound in (7) and the multiplier $n^2 x$ in (10). Since $\mathbb{P}_r(L > n^2 t) = (Q^k\mathbf{1})(r)$, these are errors smaller than $O(n^{-1})$. This completes the proof of Theorem 2.1.

Corollary 2.2 (Critical confidence and moments)**.** *If $\lambda_n \to \lambda$ and $r_n/n \to x > 0$, either normalized lifetime converges in law and in every fixed positive integer moment to $\mathcal{H}_\lambda$ from x. Separately, at fixed rank r, $n\mathbb{E}_r[(L/n^2)^j]$ converges to $B\,\partial_x \mathbb{E}_x \mathcal{H}_\lambda^j|_{x=0}$, for every integer $j \ge 1$, with error $O_{r,j}(n^{-1})$ when λ_n is retained in the comparison. For a fixed positive target failure probability $\epsilon < 1$ and $x > 0$, the first-recording $1 - \epsilon$ quantile divided by n^2 tends to the unique $t > 0$ such that $S_\lambda(t, x) = \epsilon$.*

Proof. The distributional convergence follows from Theorem 2.1. The transforms and their comparison extend uniformly to a small complex disk by the Neumann series for $I + n^2(e^{z/n^2} - 1)K_n$ and $I + zK_{\lambda_n}$. These are the analytic continuations of the actual transforms: the operator norm bound on K_n and the spectral interval for Q ensure a uniform positive radius of exponential integrability on the n^2 clock. By the same resolvent identity, the errors in the disk are $O(n^{-1})$ in sup norm and $O_r(n^{-1})$ for the n-scaled differences in (7). Cauchy's formula gives the moment claims. The diffusion survival is continuous and strictly decreasing from one to zero for $x > 0$. Strict decrease follows from positive probability of survival to any fixed time and of being absorbed within any subsequent positive interval from each interior starting point. Its endpoint limits follow from path continuity and finite mean. The quantile assertion follows by sandwiching at any two nearby times with distinct survival levels. □

2.3 Fixed retained weight above the threshold

Write $a = s/p$, $b = 1 - s$, and $\mu = a - b > 0$ in this section. Put

$$V = \frac{a + b - \mu^2}{\mu^3}.$$

For $r \geq 1$ let $\varphi_r(k) = (2\pi r V)^{-1/2} \exp(-(k - r/\mu)^2/(2rV))$.

Theorem 2.3 (Central and local central limit laws)**.** *Fix p, c with $c_* < c \leq 1$. For $L = \tau$ or T, uniformly in $1 \leq r \leq n$,*

$$\sup_{k \in \mathbb{Z}} |\mathbb{P}_r(L = k) - \varphi_r(k)| = O(r^{-1}).$$

The distribution of $(L - r/\mu)/\sqrt{rV}$ differs from a standard normal distribution by $O(r^{-1/2})$ in Kolmogorov distance.

Proof. Let $D_{n,i}$ be the time from rank i to the first visit to $i - 1$. The strong Markov property expresses τ from r as a sum of independent variables with laws $D_{n,i}$, $1 \leq i \leq r$. These laws have uniformly bounded exponential moments. Here are details of the boundary argument. For $k \geq J$ with J fixed and sufficiently large, $a_k - b_k \geq \mu/2$ by (1). For suitably small fixed $\theta > 0$, the expected multiplier of $e^{\theta k}$ in one step is at most $\rho < 1$ in this range. Killing on a descent below a specified level therefore bounds the tail of a downward passage of one, or of a bounded number of, levels exponentially. For $i < J$, from any level $k \in [i, J]$ try J consecutive downward moves, stopping upon reaching $i-1$. The chance of success is bounded below by a constant because $a_k \geq s(1 - p^{-1})/p$ for $k \geq 1$. A failed trial ends at most $2J$, so the wait to return to the strip has a uniformly bounded exponential moment by the drift estimate. The exponential moment of the failed trial and return, weighted by its failure indicator, is less than one for some common sufficiently small positive exponent. Summing over repeated trials proves the assertion also in the strip. The argument covers $n < J$ by making the trial within the finite state space.

In the homogeneous lazy walk with step probabilities $a, b, 1 - a - b$ for $-1, 1, 0$, let D denote the first downward passage by one. The same drift argument gives an exponential moment. The stopped position martingale and its square with compensator $(a + b - \mu^2)t$ give $\mathbb{E}D = 1/\mu$ and $\mathbb{E}(\mu D - 1)^2 = (a + b - \mu^2)\mathbb{E}D$. In particular $\operatorname{Var} D = V$. To compare moments, couple the walks started at i for $\lfloor \min(i - 1, n - i)/2 \rfloor$ steps by coupling their one-step increments. On that interval in time any possible location stays at least half the initial distance away from the boundaries. The difference in the transition probabilities is exponentially small in that distance, by (1). A union bound and the uniform exponential tails bound the failure to see an identical first passage by $Ce^{-\gamma \min(i-1, n-i)}$, after changing constants for the indices near the endpoints. Cauchy–Schwarz and bounded higher moments give the same form of bound with a possibly smaller γ for differences of first and second moments. Summing the geometric boundary errors gives total mean $r/\mu + O(1)$ and variance $rV + O(1)$.

The delay $T - \tau$ is independent of τ and has a uniformly bounded exponential moment. At rank zero, a trial records with probability $F_n \geq c$. After a failure the chain holds or makes one excursion from rank one back to zero. The duration of that excursion is $D_{n,1}$; repeating the same exponential-moment trial argument gives the bound. For $c = 1$ this delay is exactly one. These facts and the uniformly bounded third absolute moments give the Kolmogorov estimate by the independent-summand Berry–Esseen inequality, allowing for the bounded errors in mean and variance.

For the local estimate, each $D_{n,i}$ has probabilities bounded below at one and two, using a single downward step or a hold followed by such a step. The hold probability is at least $s - 1/p > 0$. Thus

the modulus of its characteristic function is bounded by $\exp(-\gamma \sin^2(u/2))$ for $|u| \le \pi$. The log characteristic functions on a fixed small neighborhood of zero have a uniform third-order remainder, by the moment bounds. Taylor expansion and the summed mean and variance comparisons bound the difference of the product from $\exp(iru/\mu - rVu^2/2)$ by $Ce^{-\gamma' ru^2}(|u| + u^2 + r|u|^3)$ in that neighborhood. For instance one can interpolate the two logarithms inside the exponential, using the quadratic decay of their real parts. The recording delay multiplies the product by $1 + O(|u|)$ there and by a factor of modulus at most one elsewhere. Integrating the difference gives $O(r^{-1})$; the rest of the lattice Fourier interval and the continuum Gaussian tail are exponentially small. Fourier inversion proves the estimate for both lifetimes. □

2.4 Fixed retained weight below the threshold

Assume $0 < c < c_*$, so $R = b/a > 1$, with $a = s/p$ and $b = 1 - s$. For $j \ge 0$ put $Q_j^+ = \prod_{i=0}^{j}(1 + \frac{1-s}{s}p^{-i})$ and $Q_{-1}^+ = 1$. Define

$$U(r) = \frac{\sum_{j=0}^{r-1} R^{-j} P_j}{\sum_{j\ge 0} R^{-j} P_j}, \qquad w = \frac{c}{c + (1-c)U(1)}.$$

As usual $U(0) = 0$. The exact means from maximal missing rank give the two constants

$$D = \frac{p}{sQ_\infty^+}\left(\sum_{j\ge 1} R^{-j} P_{j-1}\right)\left(\sum_{\ell\ge 0} R^{-\ell}\frac{Q_{\ell-1}^+}{P_\ell}\right), \qquad K = D/w. \tag{11}$$

Here $D = \lim R^{-n}\mathbb{E}_n\tau$ and $K = \lim R^{-n}\mathbb{E}_n T$. Both statements follow from the mean prefactors in [7]. They also follow directly by setting $k = n - \ell$ in (3), dividing by R^n, and applying dominated convergence with summable majorant $CR^{-i-\ell}$. Equation (4) then gives K.

Theorem 2.4 (Exponential and mixture laws)**.** *For each fixed $q \ge 1$, with W_q the order-q Wasserstein distance on $\mathbb{R}$, $L = \tau, T$ from starting rank n satisfies*

$$W_q(\mathrm{Law}_n(L/\mathbb{E}_n L), \mathrm{Exp}(1)) = O_q(R^{-n/2}),$$

where $\mathrm{Exp}(m)$ *denotes the exponential distribution of mean m. Uniformly over starting ranks $0 \le r \le n$, the following bounded-Lipschitz distances tend to zero:*

$$\begin{aligned} \mathrm{Law}_r(\tau/R^n) \quad &\textit{to} \quad (1 - U(r))\delta_0 + U(r)\,\mathrm{Exp}(D), \\ \mathrm{Law}_r(T/R^n) \quad &\textit{to} \quad (1 - U(r))w\delta_0 + \big[1 - (1 - U(r))w\big]\,\mathrm{Exp}(K). \end{aligned}$$

Proof. Use the full chain, continued through recordings. Its positive return cycles to n, started from n, are independent and identically distributed, including their recording marks. Their lengths have a uniformly bounded exponential moment. To check this, select a fixed J such that $b_k - a_k \ge (b - a)/2 > 0$ when $k \le n - J$. This follows from (1), since the lower-end factor $1 - p^{-k}$ can only decrease a_k. For small $\theta > 0$, $e^{\theta(n-k)}$ has conditional one-step contraction ratio at most $\rho < 1$ below the top strip. From within a fixed distance below the strip this gives an exponential tail for the time to enter it. In the strip the chance of enough consecutive upward moves to reach n is bounded below by a positive constant, using $b_k \ge (1 - s)(1 - p^{-1})$ for $k < n$. Each failed trial ends within bounded distance of the strip. As in Section 2.3, repeated trials with a small exponential moment bound give a uniform exponential bound for return to n. From an arbitrary rank, the first entrance into this strip and then the first hit on n have an exponential tail bounded by $Ce^{\theta n - \gamma t}$ at time t. All constants depend only on p, c.

For one of the two killing events, let ϵ_n be the probability a cycle is successful and stop that cycle at its first killing event. Starting at n, the expected lifetime is the mean length of this stopped cycle divided by ϵ_n, by one-cycle renewal. The stopped length is between one and the full cycle length. Equation (11) therefore implies $\epsilon_n \asymp R^{-n}$ for both stopping rules. Write $L = \sum_{i=1}^{G} X_i + Z$ in law, with G the geometric number of failures before success, X_i the iid lengths conditional on failure, and Z the stopped length conditional on success. All variables in this representation are independent. The X_i have uniformly bounded exponential moments. The success length has $\|Z\|_{L^q} = O_q(1 + \log(1/\epsilon_n))$: divide the exponential upper bound for the full cycle length by ϵ_n, truncate the result at one and integrate the tail. For $m_f = \mathbb{E}X_1$, the L^q norm of the centered sum is $O_q(\epsilon_n^{-1/2})$, by conditioning on G and the independent-sum moment inequality (use L^2 for $q \le 2$). Also

$$\frac{m_f}{\mathbb{E}_n L} = \epsilon_n\big[1 + O(\epsilon_n \log(1/\epsilon_n))\big]$$

by the exact mean for the geometric representation and $m_f \ge 1$. With E unit exponential one can couple $G = \lfloor E/[-\log(1-\epsilon_n)]\rfloor$. Then $\|\epsilon_n G - E\|_{L^q} = O_q(\epsilon_n)$. These estimates prove the Wasserstein bound for $L/\mathbb{E}_n L$.

For the initial rank, let $U_n(r)$ be the probability to reach n before zero, with $U_n(0) = 0$, $U_n(n) = 1$. Solving the two-boundary difference equation shows that its successive increments are proportional to $\prod_{i=1}^{j} a_i/b_i$, $j = 0, \dots, n-1$. By (1) each product is bounded by CR^{-j} and tends for fixed j to $R^{-j}P_j$; thus $\sup_r |U_n(r) - U(r)| \to 0$. The endpoint-hitting time is $O_{\mathbb{P}}(n)$ uniformly by comparison with the first hit on n in the full chain. Starting at zero, the probability of a recording before n is exactly

$$w_n = \frac{F_n}{F_n + b_0^{\mathrm{k}} U_n(1)} \longrightarrow w, \tag{12}$$

by conditioning on the next trial and the excursion if it enters rank one. Thus the probability of a recording before reaching n from r is $(1 - U_n(r))w_n$. Time until recording or first hitting n is again bounded by the time to n in the full chain, up to one round under the inclusive convention. On reaching n the remaining time has the just-proved exponential limit independently of the past. This and the analogous exposure decomposition give the two mixtures uniformly. □

Theorem 2.5 (Continued recordings)**.** *Run the chain from rank n without stopping on a recording. The point measure of recording times divided by R^n converges in law on $[0, \infty)$, for the vague topology, to a Poisson process of cluster locations of rate K^{-1}, with iid multiplicities independent of the locations. The multiplicities are geometric on $\{1, 2, \dots\}$, with success parameter*

$$\frac{(1-c^2)U(1)}{c + (1-c)U(1)}. \tag{13}$$

For fixed initial rank r, the same limit has an additional cluster at time zero with probability $(1 - U(r))w$, independent of the subsequent process and with the same multiplicity law conditional on occurring.

Proof. Call a positive return cycle to n marked if it has at least one recording, with probability ϵ_n. After any recording, the next rank is zero with probability F_n and otherwise one. Hence another recording before returning to n occurs with probability

$$d_n = w_n[F_n + (1 - F_n)(1 - U_n(1))] = w_n[1 - (1 - F_n)U_n(1)].$$

The strong Markov property applied just after that recording transition gives the positive geometric distribution with success parameter $1 - d_n$ conditional on a marked cycle. Its limit is (13).

This computation keeps the post-recording state in the renewal. An immediate fresh trial with probability w_n alone would give a different parameter.

Let M_n be the mean of a full return cycle. The mean of the cycle stopped on its first recording differs from M_n by $o(1)$, since the difference is bounded by the full cycle length times the mark indicator, whose expectation is $O(\sqrt{\epsilon_n})$ by Cauchy–Schwarz. Renewal for T and (11) imply

$$\frac{M_n}{\epsilon_n R^n} \longrightarrow K.$$

In the cycle index, iid Bernoulli marks of probability ϵ_n with iid conditional multiplicities converge, on multiplying the index by ϵ_n, to a rate-one Poisson process with the independent limiting marks. This follows by multiplying the one-cycle Laplace factors for continuous compactly supported test functions and taking their Riemann sum; the conditional geometric probability generating functions converge uniformly on $[0,1]$. For the calendar time, the maximum deviation of a partial sum of the first $\lfloor A/\epsilon_n \rfloor$ cycle lengths from its index times M_n, divided by R^n, tends to zero in probability by the independent-sum maximal second-moment inequality. The largest length among these cycles is $o_{\mathbb{P}}(R^n)$ by their uniform exponential tails. Taking A larger than any given time-window length divided by K shows that moving each cycle mark to its actual times gives the stated calendar limit. The total multiplicity on bounded cycle-index windows is tight by the bounded conditional geometric mean. No independence between the cycle length and the multiplicity within the same cycle is used.

For fixed starting rank, the time to first hit n is $O_{\mathbb{P}}(n)$, so all preliminary recordings converge to time zero. The chance of any preliminary recording tends to $(1 - U(r))w$ by (12). Once the first one occurs its continued multiplicity has the same geometric law. The full process after first hitting n is independent by the strong Markov property, proving the last assertion. □

Corollary 2.6 (First-recording confidence)**.** *From initial rank n, at fixed $0 < c < c_*$ the time for a fixed target failure probability $\epsilon \in (0,1)$ is asymptotic to $KR^n \log(1/\epsilon)$. At fixed $c > c_*$, its quantile is $n/\mu + \sqrt{nV}\,\Phi^{-1}(1-\epsilon) + O_\epsilon(1)$, with Φ the standard normal distribution function.*

Proof. The exponential law yields the first limit by continuity and strict monotonicity. The normal density is positive at the fixed quantile; the Kolmogorov $O(n^{-1/2})$ bound in Theorem 2.3 therefore gives an $O(1)$ error in the unscaled integer quantile. □

2.5 Approaching the threshold outside the critical window

The preceding fixed-parameter laws extend to the two neighboring limits where $c_n \to c_*$ and the total drift on the interval diverges. The left side has increasingly large recording clusters, and the right side is asymptotically normal once the initial rank exceeds the reciprocal drift length.

Theorem 2.7 (Vanishing drift towards zero)**.** *Let $a = s/p$, $b = 1-s$, and suppose $0 < \mu = a - b \to 0$. For $L = \tau, T$ and $r\mu \to \infty$, the distribution of*

$$\frac{L - r/\mu}{\sqrt{r(a+b-\mu^2)/\mu^3}}$$

has Kolmogorov distance $O((r\mu)^{-1/2})$ from the standard normal.

Proof. Here $R = b/a = 1 - \mu/a$, so $1 - R \asymp \mu$. Fix i and let D_i be the passage from i to $i-1$. The same Green computation as in (3), with absorbing boundary $i-1$, gives expected occupation $T_{i,k}$ at $k \geq i$ starting at i. The geometric boundary products give

$$T_{i,k} \leq CR^{k-i}, \qquad |T_{i,k} - a^{-1}R^{k-i}| \leq CR^{k-i}(p^{-i} + p^{k-n}).$$

Writing $d_i = \mathbb{E}D_i$, summing this comparison and paying the homogeneous tail beyond n yields

$$|d_i - \mu^{-1}| \le C\mu^{-1}(p^{-i} + R^{n-i}). \tag{14}$$

For instance the upper-boundary product contribution is $O(R^{n-i})$, since R stays bounded away from $1/p$.

If L_i is the passage to $i-1$ from a general $k \ge i$, its mean is $\sum_{j=i}^{k} d_j$. The one-step equation for its second moment has forcing $2\mathbb{E}_k L_i - 1$. Thus

$$\mathbb{E}D_i^2 = \sum_{k=i}^{n} T_{i,k} \left(2 \sum_{j=i}^{k} d_j - 1 \right).$$

Compare with the homogeneous version on the half-line, whose mean increment is μ^{-1} and variance is $V = (a+b-\mu^2)/\mu^3$, as in Section 2.3. The lower-boundary errors are at most $C\mu^{-3}p^{-i}$. The homogeneous tail beyond n is at most $C\mu^{-3}(1+\mu(n-i))R^{n-i}$. For the upper-boundary errors in $T_{i,k}$, sum $R^{k-i}p^{k-n}(k-i+1)/\mu$, which is bounded by $C(n-i+1)R^{n-i}/\mu$. For those from (14), interchange the two sums and use $\sum_{k=j}^{n} R^{k-i} \le CR^{j-i}/\mu$. This gives $C(n-i+1)R^{n-i}/\mu^2$. In particular

$$|\mathbb{E}D_i^2 - (\mu^{-2} + V)| \le C\mu^{-3}(p^{-i} + (1+\mu(n-i))R^{n-i}).$$

Summing over $i \le r$ and using (14) to subtract the squared means gives

$$\mathbb{E}_r \tau = r/\mu + O(\mu^{-2}), \qquad \mathrm{Var}_r(\tau) = rV + O(\mu^{-4}).$$

For the third moment, the one-step equation has forcing $3\mathbb{E}_k L_i^2 - 3\mathbb{E}_k L_i + 1$. The independent-passage representation and the preceding uniform first and second moment bounds give $\mathbb{E}_k L_i^2 \le C[(k-i+1)\mu^{-3} + (k-i+1)^2\mu^{-2}]$. Summing against $T_{i,k} \le CR^{k-i}$ yields $\mathbb{E}D_i^3 \le C\mu^{-5}$.

The terminal delay is independent of the exposure time. It has first, second and third moments bounded by $C\mu^{-1}, C\mu^{-3}, C\mu^{-5}$, respectively. Indeed the number of failed rank-zero trials is geometric with parameter F_n bounded away from zero, each failed trial with its potential rank-one excursion lasts no longer in law than $1 + D_1$, and conditional on the number of failures the cycles are independent. For an integer $j \le 3$, the conditional j-th moment of their sum is at most this number to the power j times a common j-th moment. Add the final round. The independent-summand Berry–Esseen ratio is bounded by a constant times $r\mu^{-5}/(rV)^{3/2}$, hence is $O((r\mu)^{-1/2})$. The change from the exact mean and variance to the displayed ones costs at most that order in Kolmogorov distance, using the bounded normal density and bounded $|x|$ times that density. □

Theorem 2.8 (Vanishing drift away from zero)**.** *Suppose $\delta = \log R \to 0$ with $n\delta \to +\infty$, and put $A_n = (p+1)e^{n\delta}/\delta^2$. For either $L = \tau, T$, $\mathbb{E}_n L \sim A_n$ and*

$$W_2(\mathrm{Law}_n(L/\mathbb{E}_n L), \mathrm{Exp}(1)) = O(e^{-n\delta/2}).$$

If $r_n\delta \to x \in [0,\infty]$, either L/A_n from rank r_n converges to $e^{-x}\delta_0 + (1-e^{-x})\,\mathrm{Exp}(1)$. For the continued recording process, put $\vartheta_n = 1 - w_n[1-(1-F_n)U_n(1)]$. On dividing times by A_n and assigning each recording mass ϑ_n, the random atomic measures from rank n converge vaguely in law to a rate-one Poisson process with independent unit exponential masses. The normalizing mass satisfies

$$\vartheta_n \sim \frac{1-c_*^2}{c_* C_p}\,\delta.$$

From ranks with $r_n\delta \to x$, there is in addition a cluster at zero with probability e^{-x}, with independent unit exponential mass conditional on occurring.

Proof. Write $\lambda = n\delta$. Taking $k = n - \ell$ in (3), the homogeneous sum divided by e^{λ} satisfies

$$\sum_{\substack{i\geq 1,\ \ell\geq 0\\ i+\ell\leq n}} e^{-\delta(i+\ell)} \sim \delta^{-2}.$$

Indeed the omitted part of the product geometric sums has relative size $O((1+\lambda)e^{-\lambda})$. The product corrections are bounded by a constant times $p^{-i} + p^{-\ell}$, whose relative weighted sum is $O(\delta)$. Since $p/s \to p+1$, $\mathbb{E}_n\tau \sim A_n$. The bound $G\mathbf{1}(1) \leq Ce^{\lambda}/\delta$ and (4) give the recording mean. In the gambler's ruin formula, the increment product at index j equals $R^{-j}P_j[1 + O(p^{j-n})]$ by (1). The denominator is asymptotic to C_p/δ: the sum of $R^{-j}(P_j - C_p)$ is bounded, the upper correction sums to $O(1)$, and the omitted geometric tail costs a fraction $O(e^{-\lambda})$. The same absolute errors in any partial sum show $U_n(r) = 1 - e^{-\delta r} + o(1)$ uniformly. In particular $U_n(1) \sim \delta/C_p$ by its exact numerator 1. Thus $w_n \to 1$ and the formula for ϑ_n follows from

$$\vartheta_n = \frac{(1-F_n)(1-h_n+F_n)U_n(1)}{F_n + b_0^{\mathrm{k}}U_n(1)}.$$

A positive return cycle to n has exposure mark probability $a_n(1-U_n(n-1)) \asymp \delta e^{-\lambda}$; its recording mark probability additionally multiplies this by w_n. Use ϵ_n for either probability.

We need a uniform second moment for full return cycles. In the reversed coordinate $y = n - r$, the downward and upward transition probabilities have ratio

$$\frac{\mathrm{up}(y)}{\mathrm{down}(y)} = R^{-1}\frac{(1-p^{y-n})(1+\frac{1-s}{s}p^{-y})}{1-p^{-y}} \qquad (1 \leq y \leq n).$$

The downward probabilities are bounded below for $y \geq 1$. Products of these ratios are bounded by $Ce^{-\delta(j-i)}$ over an interval from i through $j-1$, since the upper products have factors at most one and the p^{-y} factors have bounded products and reciprocals. The Green kernel toward $y = 0$ therefore gives the mean from y at most Cy/δ. Its first row is at most $Ce^{-\delta(k-1)}$ in column k. Applying it to twice the mean minus one gives second moment $O(\delta^{-3})$ from $y = 1$. Let H_n denote a positive full return cycle length. From $y = 0$ it either takes one step or adds the passage from $y = 1$, so $\mathbb{E}H_n^2 \leq C\delta^{-3}$.

For the stopped cycle, write $m = \epsilon_n\mathbb{E}_nL \asymp \delta^{-1}$ for its mean. Then

$$\frac{\sqrt{\epsilon_n\mathbb{E}H_n^2}}{m} = O(e^{-\lambda/2}).$$

In the geometric representation of Theorem 2.4, the centered failure sum and the final success length have L^2 norm $O(\sqrt{\mathbb{E}H_n^2/\epsilon_n})$. The ratio of the failure mean to m differs from one by $O(\epsilon_n + \sqrt{\epsilon_n\mathbb{E}H_n^2}/m)$, since the success contribution to the stopped mean is bounded by $\sqrt{\epsilon_n\mathbb{E}H_n^2}$. The same coupling of the geometric count to an exponential then gives the asserted W_2 bound. From any initial rank, the first hit on n has mean $O(n/\delta) = o(A_n)$. Use the endpoint decomposition from Theorem 2.4, with $U_n(r)$ and w_n as estimated above, for both mixtures.

For continued recordings the marked-cycle multiplicity is geometric with parameter ϑ_n, as in Theorem 2.5. Multiplication by ϑ_n gives convergence to a unit exponential mark. The full mean return length differs from m by at most $\sqrt{\epsilon_n\mathbb{E}H_n^2} = o(m)$, so $\mathbb{E}H_n/(\epsilon_nA_n) \to 1$ and the calendar rate on the A_n clock is asymptotically one. The maximal second-moment bound and the union bound for the first $O(1/\epsilon_n)$ cycles are both controlled by $\mathbb{E}H_n^2/(\epsilon_nA_n^2) = O(e^{-\lambda})$ after rescaling. Consequently both the calendar deviations and the largest cycle length tend to zero in probability on this clock. On bounded cycle-index windows the total rescaled mark mass is tight, since its

expectation is bounded. The same iid marked-cycle Laplace product thus gives the claimed vaguely convergent weighted point measures. For initial rank r_n, the preliminary interval has length $o_{\mathbb{P}}(A_n)$; it contains a recording with probability $(1-U_n(r_n))w_n \to e^{-x}$. Its conditional multiplicity has the same geometric law, independently of the process after the first top visit. □

2.6 Continued recordings in the critical window

Now leave the full chain running when an exact candidate is recorded. Write $X_n(t)$ for the rank at round $\lfloor n^2 t \rfloor$ divided by n, and $J_n(t)$ for the cumulative number of exact recordings through that many rounds. The relevant limit of the record count is a boundary local time. We fix its normalization by taking the lower regulator L_0 in

$$X(t) = x + \sqrt{\frac{2}{p+1}}\, B(t) + \frac{\lambda}{p+1} t + L_0(t) - L_1(t). \tag{15}$$

The regulators L_0, L_1 are continuous, start at zero, are nondecreasing, and have increments supported on $X = 0$ and $X = 1$, respectively. Put $\rho_p = c_*(p+1)C_p$.

Theorem 2.9 (Boundary occupation and thinned recordings)**.** *If $\lambda_n \to \lambda$ and $r_n/n \to x \in [0,1]$, then*

$$(X_n, J_n/n) \implies (X, \rho_p L_0)$$

on each compact time interval in the Skorokhod topology. Independently retain each candidate with probability θ/n for fixed $\theta > 0$. The point process of retained exact recordings on the n^2 clock converges jointly with X_n to the process that, conditional on X and its regulators, is Poisson with intensity $\theta\rho_p\, dL_0$. Its first-arrival survival probability from x is

$$u_\lambda(t,x) = \mathbb{E}_x \exp(-\theta\rho_p L_0(t)).$$

This is the diffusion survival with lower Robin condition $\partial_x u_\lambda(t,0) = \theta\rho_p u_\lambda(t,0)$ and upper Neumann condition $\partial_x u_\lambda(t,1) = 0$.

Proof. For $1 \le i \le n$, define a corrected rank increment

$$d_i = C_p^{-1} \prod_{j=1}^{i-1} (R_n a_j / b_j), \qquad H(r) = \sum_{i=1}^{r} d_i.$$

The exact products in (1) give $|d_i - 1| \le C(p^{-i} + p^{i-n})$. Indeed the lower product is P_{i-1}/C_p and the upper correction is a product of the factors $(1 + \frac{1-s}{s} p^{j-n})/(1-p^{j-n})$. The sum of the resulting errors is bounded, so $H(r)/n = r/n + O(n^{-1})$ uniformly. Write $\widehat{X}_n$ for the corrected process and $l_n = H(n)/n$. Its one-step conditional mean increment from $1 \le r < n$ is $(R_n - 1)a_r d_r/n$, by $b_r d_{r+1} = R_n a_r d_r$. The conditional mean increments at the boundaries are $b_0 d_1/n$ and $-a_n d_n/n$. Let $L_{0,n}$ and $L_{1,n}$ be the sums of these two boundary drift magnitudes, respectively. With M_n the centered-increment martingale and A_n^{int} the sum of interior drifts,

$$\widehat{X}_n(t) = \widehat{X}_n(0) + M_n(t) + A_n^{\mathrm{int}}(t) + L_{0,n}(t) - L_{1,n}(t).$$

For a fixed t_1, let V_i count the visits to rank i at the starts of rounds through $\lfloor n^2 t_1 \rfloor$. The test function $g(u) = u(l_n - u)$ has conditional mean increment $g'(\widehat{X}_n)\mathbb{E}(\Delta\widehat{X}_n \mid r) - \mathbb{E}((\Delta\widehat{X}_n)^2 \mid r)$. The first term is bounded below by a positive constant times $1/n$ at either boundary, and is $O(n^{-2})$ in the interior. The second term is $O(n^{-2})$ uniformly. Summing, using bounded g, gives $\mathbb{E}(V_0 + V_n) \le C_{t_1} n$.

For each cut between i and $i+1$, $b_i\mathbb{E}V_i - a_{i+1}\mathbb{E}V_{i+1}$ is bounded in absolute value by one, since the net crossing count equals the change of a zero-one indicator. The detailed-balance weights with initial weight one at zero stay between positive constants by the same geometric product argument used for (8). For $i < n$, their products with b_i are bounded below. Dividing each crossing relation by this product and summing from 0 to $i-1$ gives $\sup_i \mathbb{E}V_i \le C_{t_1} n$.

It follows that the accumulated interior drift converges uniformly in probability to $\lambda t/(p+1)$. Indeed $n(R_n - 1) \to \lambda$, $a_i d_i = (p+1)^{-1} + O(n^{-1} + p^{-i} + p^{i-n})$, and the occupation estimate makes the summed boundary errors $O(n^{-1})$ in L^1 after division by n^2. The same reasoning applies to the predictable quadratic variation of M_n, whose interior one-step conditional second moment before centering is $[a_i d_i^2 + b_i d_{i+1}^2]/n^2$. It converges uniformly in probability to $2t/(p+1)$; the centering correction vanishes and the boundary visits have negligible duration. The jumps are $O(n^{-1})$, so the martingale functional central limit theorem applies.

For completeness, the boundary support persists in the limit. On an interval where the oscillation of $M_n + A_n^{\mathrm{int}}$ is smaller than $l_n - O(n^{-1})$, the two regulators cannot both increase: between the last push at one boundary and the first push at the other, the driver must traverse the interval apart from one push of size $O(n^{-1})$. When only one regulator increases, its total increment is bounded by the driver's oscillation plus $O(n^{-1})$, by evaluating the path just before the last push at that boundary. These estimates and the continuous limiting driver give tightness with continuous limits for both regulators and for $\widehat{X}_n$. In a limit, either push is constant on each compact time interval where the path is separated from the corresponding endpoint. The interval Skorokhod problem has a unique solution for a given continuous driver [12], so the full corrected process and its pushes converge to (15).

The recording count minus F_n times the number of rank-zero visits is a martingale even though a recording and its next rank can be dependent. Its quadratic variation in expectation is $O(n)$ by the visit bound. Dividing by n and applying the maximal L^2 inequality gives a uniformly vanishing error in probability. Since $F_n/(b_0 d_1) \to \rho_p$, this proves the recording-count limit. Given all trials and their recording marks, thinning multiplies the conditional Laplace factors $1 - (\theta/n)(1 - e^{-v(t_j)})$ for each recording time t_j, with v a nonnegative continuous function of compact support. The logarithm equals $-\theta \int (1 - e^{-v})\, d(J_n/n) + o_{\mathbb{P}}(1)$. The expected number of retained recordings on each compact interval is bounded by the visit estimate. Tightness and the conditional Laplace factors prove the joint limit by testing also against bounded continuous functions of the unthinned process. Taking the probability of no retained recording gives the displayed survival. For the heat equation, apply Itô's formula to $e^{-\theta\rho_p L_0(v)} u_\lambda(t-v, X(v))$ for $v < t$ along (15) and pass to $v = t$ by continuity. The lower push gives $\partial_x u_\lambda - \theta\rho_p u_\lambda$ and the upper push gives $-\partial_x u_\lambda$, yielding the stated boundary conditions with initial value one. This is the conventional local-time representation of Robin killing [13]. □

Corollary 2.10 (Quantitative Robin approximation)**.** *Suppose only that $|\lambda_n| \le \Lambda$, and let T_θ be the inclusive time of the first exact recording retained at probability θ/n for fixed $\theta > 0$ and sufficiently large n. For t in a compact subinterval of $(0,\infty)$ and z in a compact subinterval of $[0,\infty)$, uniformly over r,*

$$\mathbb{P}_r(T_\theta > n^2 t) = u_{\lambda_n}(t, r/n) + O_\theta(n^{-1}),$$
$$\mathbb{E}_r e^{-zT_\theta/n^2} = f_{\lambda_n, z}(r/n)\, \frac{\theta\rho_p}{\theta\rho_p - f'_{\lambda_n, z}(0)} + O_\theta(n^{-1}).$$

Proof. For the killed chain, the row at zero has killing probability $q_n = \theta F_n/n$, and upward probability $b_0^{(\theta)} = b_0 - (\theta/n)F_n(1 - F_n)$. Equation (4) therefore holds with F_n, b_0^{k} replaced there by

$q_n, b_0^{(\theta)}$. By (7) for the exposure row $r = 1$,

$$n^{-2}\widetilde{G}_\theta = K_{\lambda_n} + \frac{1}{\theta\rho_p}\,\mathbf{1} \otimes D_{\lambda_n} + O_{\infty\to\infty}(n^{-1}).$$

Indeed $b_0^{(\theta)}/(nq_n) = 1/((p+1)\theta c_*)+O(n^{-1})$, $B_1 = 1/C_p$, and $u(0)/(n^2 q_n)$ has operator norm $O(n^{-1})$. The step kernel is uniformly bounded, including the latter cell's contribution. The new continuum kernel is the inverse of the negative generator with these Robin and Neumann conditions: it adds the constant $(\theta\rho_p)^{-1}$ times the old solution's derivative at zero. Positivity and self-adjointness in the weighted space follow by integration by parts, including the positive Robin form term. Reversibility and the positive holds in the finite chain remain uniform. The resolvent and de-Poissonization proof of Theorem 2.1 now gives the survival error. For the transform the same one-step inverse gives the stated error against the continuum equation $g = 1 - zK_{\mathrm{Robin}}g$. Thus g satisfies the interior ODE of $f_{\lambda,z}$, the upper Neumann condition, and $g'(0) = \theta\rho_p(g(0) - 1)$. The upper condition makes g a multiple of $f_{\lambda,z}$; the lower condition gives the displayed factor. □

2.7 Scope of the distributional results

Passage-time representations and exponential escape for birth–death chains have classical antecedents [8, 9]. Glynn [10] proves regenerative exponential limits for rare hitting times. Erhardsson [11] analyzes scalar compound Poisson counts for extreme sojourns in stationary birth–death chains. Here the nonstationary initial rank and the full recording transition determine the cluster rate, multiplicity, and probability of an initial cluster. Fan [14] proves joint boundary local-time approximations for lattice walks. Singer et al. [15] derive Robin killing from partially absorbing Euler steps. Here the recording marks supply a Poisson intensity on the lower regulator and an explicit first-arrival approximation for the nonuniform rank chain. The exact chain (1), its finite means and their leading asymptotics are those of [7]. The distributional conclusions here use the boundary layer of that chain, including the full recording and transition law at zero. In the critical window the boundary correction survives in the fixed-rank entrance tail, while from ranks comparable with n both lifetimes have the same leading diffusion. In the rare-recording regime the same recording rule affects both the mean first-recording time and the geometric multiplicities of the continued process. These conclusions concern the exact process under global depolarization at fixed prime p.

Statements and Declarations

Funding. This research was self-funded.

Competing interests. The author owns Luminary Oddities and has no other relevant competing interests.

AI assistance. OpenAI Codex assisted with derivations, proof development, source checks and independent internal review of each part. Anthropic Claude edited, consolidated and cross-referenced the parts and checked computations. Scott Schubr directed the research and is responsible for the content.

References

[1] S. Schubr, Sharp noise transitions for an adaptive single-copy stabilizer-learning process, Luminary Research Archive, edition 4 (2026). `https://research.luminaryoddities.org/sharp-noise-transitions-adaptive-single-copy-stabilizer-learning`.

[2] L. Bittel, J. Eisert, W. Gong, A. A. Mele and L. Schatzki, Adaptivity is all you need: optimal stabilizer learning using just single-copy measurements, arXiv:2610.02031v1 (2026). `https://arxiv.org/abs/2610.02031v1`.

[3] A. Zabokritskiy, Stabilizer learning under global depolarizing noise with polynomial copies and 2^t time, arXiv:2611.03324 (2026). `https://arxiv.org/abs/2611.03324`.

[4] S. Schubr, Single-copy learning of stabilizers with correlated diagonal noise, Luminary Research Archive, edition 2 (2026). `https://research.luminaryoddities.org/single-copy-learning-stabilizers-correlated-diagonal-noise`.

[5] S. Schubr, Single-copy learning of noisy stabilizers by certified filtering, Luminary Research Archive, edition 2 (2026). `https://research.luminaryoddities.org/single-copy-learning-noisy-stabilizers-certified-filtering`.

[6] L. Bittel, J. Eisert, W. Gong, A. A. Mele, and L. Schatzki, *Adaptivity is all you need: Optimal stabilizer learning using just single-copy measurements*, arXiv:2610.02031v1, 2026.

[7] S. Schubr, *Sharp noise transitions for an adaptive single-copy stabilizer-learning process*, version 4, Luminary Research Archive, 2026. Archive edition.

[8] J. A. Fill, *The passage time distribution for a birth-and-death chain: Strong stationary duality gives a first stochastic proof*, J. Theoret. Probab. **22** (2009), 543–557. arXiv:0707.4042.

[9] J. Barrera, O. Bertoncini, and R. Fernández, *Abrupt convergence and escape behavior for birth and death chains*, J. Stat. Phys. **137** (2009), 595–623. arXiv:0903.1832.

[10] P. W. Glynn, *On exponential limit laws for hitting times of rare sets for Harris chains and processes*, J. Appl. Probab. **48A** (2011), 319–326. Author copy.

[11] T. Erhardsson, *Compound Poisson approximation for counts of rare patterns in Markov chains and extreme sojourns in birth-death chains*, Ann. Appl. Probab. **10** (2000), 573–591. DOI.

[12] Ł. Kruk, J. Lehoczky, K. Ramanan, and S. Shreve, *An explicit formula for the Skorokhod map on* $[0, a]$, Ann. Probab. **35** (2007), 1740–1768. arXiv:0710.2977.

[13] D. S. Grebenkov, *Probability distribution of the boundary local time of reflected Brownian motion in Euclidean domains*, Phys. Rev. E **100** (2019), 062110. arXiv:1909.07677.

[14] W.-T. L. Fan, *Discrete approximations to local times for reflected diffusions*, arXiv:1410.3585v4, 2016.

[15] A. Singer, Z. Schuss, A. Osipov, and D. Holcman, *Partially reflected diffusion*, SIAM J. Appl. Math. **68** (2008), 844–868. DOI.