Sharp noise transitions in adaptive single-copy stabilizer learning

Scott Schubr

October 22, 2026

Abstract

We determine the noise-dependent behavior of the adaptive single-copy stabilizer-learning update of Bittel, Eisert, Gong, Mele and Schatzki on globally depolarized stabilizer states. For every fixed prime local dimension p, the number of diagonal stabilizers still missing follows an exact birth–death chain. Writing c for the retained pure-state weight, the expected time to record the optimizing stabilizer is linear in the number n of qudits when $c^2 > p/(p+1)$, quadratic at equality, and exponential when $c^2 < p/(p+1)$. We obtain the exact supercritical speed, subcritical exponential prefactors, and the complete $1/n$ critical-window mean profile. An exact regeneration identity distinguishes first exposure of the computational-basis target from its actual recording. A finite-time bound shows that below threshold the unchanged algorithm has exponentially small success probability in any polynomial round budget, even on this simple noise family. The obstruction concerns this update and candidate generation rule.

1 Introduction

Stabilizer states admit compact classical descriptions and are central to quantum error correction, simulation and learning. The distinction between single-copy measurements and measurements coupling several unknown copies is important both algorithmically and experimentally. Bittel et al. [1] recently gave an adaptive single-copy learner for exact stabilizer states, extended it to prime-dimensional qudits, and proved a tolerant mixed-qubit guarantee when stabilizer fidelity exceeds $\sqrt{2/3}$. Their algorithm changes the measurement frame using the difference of two separately measured computational-basis samples. On exact stabilizers, this monotonically exposes diagonal stabilizers.

Noise can destroy that monotonicity. It is natural to ask whether the high-fidelity restriction is only a feature of a sufficient proof or reflects actual behavior of the update. We answer this question for global depolarization. The algorithm itself is unchanged. Its latent progress variable becomes an exactly computable one-dimensional Markov chain, whose transition from linear to exponential cost occurs at precisely the asymptotic qubit threshold already suggested by the positive theorem.

The result has three components. First, we identify the exact transition law, rather than bound progress and regression separately. Second, we solve its mean asymptotics, including the critical window and exponential prefactors. Third, we relate latent exposure to actual candidate recording, accounting for the dependence between a failed recording and the following frame update. The last step matters at leading order below threshold.

1.1 Relationship to earlier work

The update and its noiseless prime-dimensional analysis belong to Bittel et al. [1, Algorithm 1 and Appendix A]. Their Theorem III.1 supplies the positive high-fidelity qubit learner. Our contribution

is an exact noise-dependent analysis and an algorithm-specific failure theorem for its unchanged candidate generator. The birth–death recurrences, gambler's-ruin calculation and convergence arguments used below are classical probability methods; see, for example, Norris [7] and Levin and Peres [8].

Stabilizer bootstrapping gives proper agnostic learning with multi-copy measurements for qubits [2]; Qi et al. [3] give a pure-input odd-prime extension. The author's mixed-input odd-prime learner [4] uses coherent four- or six-copy sampling. Those results supply related contrasts using their coherent escape primitive. Tang and Steckmann [5], in their second version, learn noise-robust stabilizer structure using Bell sampling and a finder–verifier architecture. Their target and measurement access differ from the fixed single-copy candidate process studied here.

Failure of one update, even on a simple family, is not failure of every adaptive protocol. Conversely, final fidelity validation cannot compensate for a target that the candidate generator almost never records.

1.2 The model and main conclusions

Fix a prime p and an n-qudit pure stabilizer $|s\rangle$, and consider

$$\rho_n = c|s\rangle\langle s| + (1-c)\frac{I}{p^n}, \qquad 0 < c \le 1. \tag{1}$$

The parameter c is the retained pure-state weight. It is not the finite-dimensional stabilizer fidelity, which is

$$F_n = c + (1-c)p^{-n}. \tag{2}$$

Let r be the codimension of the target's diagonal stabilizer space in the current measurement frame. Initially we take $|s\rangle = |+_p\rangle^{\otimes n}$, so $r = n$. The same formulas apply to any initial target with a specified initial r.

Let τ_0 be the first time at which $r = 0$, and let σ be the first pre-update round index at which the recorded candidate is exactly $|s\rangle$. Indices start at zero; the actual number of measured rounds up to and including recording is $\sigma + 1$. Neither time is observable by the learner, because the target is unknown.

Put

$$s = c^2, \qquad R = \frac{p(1-s)}{s}, \qquad \mu = \frac{s}{p} - (1-s). \tag{3}$$

For fixed p, c, the mean candidate time has the following transition:

Noise regime	Condition	$\mathbb{E}_n\sigma$
Supercritical retained signal	$s > p/(p+1)$	$n/\mu + o(n)$
Critical retained signal	$s = p/(p+1)$	$(p+1)n^2/2 + o(n^2)$
Subcritical retained signal	$s < p/(p+1)$	$K_{p,c}R^n(1+o(1))$

Here $K_{p,c}$ is a strictly positive explicit convergent-series constant given in Theorem 6.2. At $c = 1$, the already known pure-state mean is $pn + O_p(1)$.

For qubits the critical weight is $\sqrt{2/3}$. The equality with the external sufficient threshold is asymptotic: $F_n \neq c$ at finite n. At critical c, F_n is slightly larger than c, but the external high-fidelity drift parameter then tends to zero exponentially. Thus its nonuniform upper bound does not contradict the quadratic behavior proved here.

2 Stabilizer geometry and the fixed update

2.1 Unsigned labels and computational support

Let $V = \mathbb{F}_p^{2n}$, with labels (a, b) representing the Pauli operator $X(a)Z(b)$ up to scalar phase, where

$$X(a)|x\rangle = |x + a\rangle, \qquad Z(b)|x\rangle = e^{2\pi i b\cdot x/p}|x\rangle.$$

The symplectic form is $[(a, b), (a', b')] = a \cdot b' - b \cdot a'$. A pure stabilizer state determines an n-dimensional isotropic subspace $L \leq V$, together with compatible eigenvalue phases. Define

$$D = \{b : (0, b) \in L\}, \qquad r = n - \dim D, \qquad A = D^{\perp} \leq \mathbb{F}_p^n,$$

where $\perp$ in this definition is the ordinary dot-product orthogonal complement.

Lemma 2.1. *The X-projection of L equals A. The computational-basis distribution of its pure stabilizer state is uniform on an affine coset $z + A$ of size p^r.*

Proof. Every $(a, b) \in L$ commutes with every $(0, d) \in L$, so $a \cdot d = 0$. The projection onto X-coordinates therefore lies in $D^{\perp}$. Its kernel has dimension $\dim D$, and hence its image has dimension r, equal to $\dim D^{\perp}$; the two spaces coincide. The diagonal stabilizer eigenvalue equations restrict nonzero amplitudes to one affine coset $z + A$. Every $a \in A$ occurs as the X-part of a stabilizer, whose action translates amplitudes by a while preserving their absolute values. Transitivity on the coset makes all these absolute values equal. Normalization proves uniformity. □

2.2 One adaptive round

The current Clifford frame is C, and the transformed target is $C|s\rangle$. In each round the algorithm does the following, exactly as in [1]:

1. Measure two fresh copies of $C\rho_n C^{\dagger}$, separately, in the computational basis, obtaining $x, y \in \mathbb{F}_p^n$.
2. Record the candidate $C^{\dagger}|x\rangle$, and set $u = x - y$.
3. If $u = 0$, leave C unchanged.
4. If $u \neq 0$, use the pivoting linear Clifford F_u to send its X-label to e_i. Then apply one uniformly chosen local Clifford H_a, $a \in \mathbb{F}_p$, such that $H_a X Z^a H_a^{\dagger}$ is proportional to Z. Replace C by $H_a F_u C$.

For $p = 2$, the two local choices are H and HS, up to the same slope convention. In general, F_u is implemented by the controlled-SUM and multiplication gates specified in [1, Appendix A]. Its action on X-coordinates is an invertible matrix M_u with $M_u u = e_i$, and its action on Z-coordinates is M_u^{-T}. All statements below are independent of the permitted choice of pivot or representative H_a.

Recording does not alter the state update, and no target-membership test is used. One round uses two copies measured singly, rather than a coherent two-copy measurement.

Lemma 2.2 (Exact rank change)**.** *For a current target with support direction A:*

(i) If $u = 0$, the rank variable r stays fixed.

(ii) If $u \in A \setminus \{0\}$, exactly one of the p local corrections decreases r by one, and the others leave it fixed.

(iii) If $u \notin A$, every local correction increases r by one.

Proof. For $d \in D$, the pivot component of the transformed diagonal label is

$$(M_u^{-T} d)_i = d \cdot M_u^{-1} e_i = d \cdot u.$$

If $u \in A$, every old diagonal stabilizer is therefore trivial on the pivot. Also, some stabilizer has X-part u, and after F_u it has X-part e_i. Modulo old diagonal stabilizers its pivot Pauli has one well-defined nonvertical slope: two choices with X-part e_i differ by an old diagonal label, trivial on the pivot. Exactly one H_a turns this slope diagonal. The resulting generator is independent of the old diagonal space because it acts nontrivially on the pivot. At most one generator can be added: the preimages of new diagonal labels have X-part in the one-dimensional span of e_i. The remaining corrections preserve the old space and add no generator.

If $u \notin A$, the functional $d \mapsto d \cdot u$ on D is nonzero. Its kernel has codimension one, and consists precisely of the old diagonal labels that are trivial on the pivot after F_u. Every H_a preserves these labels. It takes a nontrivial pivot Z-label to a nondiagonal label: an invertible symplectic map cannot send both the distinct lines Z and XZ^a to the Z line.

No new diagonal generator can replace the lost one. Before the local correction, such a generator must have X-part supported only at the pivot. If that part is nonzero, pulling it back through F_u puts a nonzero multiple of u in the target X-projection A, impossible. If it is zero, the generator was already diagonal and must lie in the kernel just identified. Thus exactly one diagonal generator is lost. The zero-difference case makes no update by definition. □

3 The exact chain and a finite mean formula

Theorem 3.1 (Exact transition law)**.** *For input* (1)*, conditional on any complete history, the current rank variable has a time-homogeneous birth–death law depending only on r. For $1 \le r \le n$, its downward and upward probabilities are*

$$a_r = \frac{1 - p^{-r}}{p}\big[s + (1-s)p^{r-n}\big], \tag{4}$$

$$b_r = (1-s)(1-p^{r-n}). \tag{5}$$

The remaining probability is holding, and $b_n = 0$.

Proof. Clifford conjugation preserves the form (1). By Lemma 2.1, the measured distribution is $cU_{z+A} + (1-c)U_{\mathbb{F}_p^n}$. Expanding two independent draws shows that their difference has distribution

$$sU_A + (1-s)U_{\mathbb{F}_p^n}. \tag{6}$$

The target–target difference is uniform on A; each term involving a uniform full-space draw is uniform on the full space. The probability of $u \in A \setminus \{0\}$ is $(1-p^{-r})[s + (1-s)p^{r-n}]$, and the probability of $u \notin A$ is $(1-s)(1-p^{r-n})$. Lemma 2.2 proves the formulas and their history independence. □

For hitting-time calculations we stop the chain at zero. For $0 < c < 1$, all a_r are positive, all b_r with $r < n$ are positive, and the finite chain reaches zero in finite mean. Write

$$g_n(r) = \mathbb{E}_r \tau_0, \qquad g_n(0) = 0, \qquad d_{n,r} = g_n(r) - g_n(r-1).$$

Proposition 3.2 (Exact mean identity)**.** *For $0 < c < 1$,*

$$d_{n,r} = \sum_{k=r}^{n} \frac{1}{a_k} \prod_{j=r}^{k-1} \frac{b_j}{a_j}, \tag{7}$$

$$g_n(n) = \sum_{1 \le r \le k \le n} \frac{1}{a_k} \prod_{j=r}^{k-1} \frac{b_j}{a_j}. \tag{8}$$

Empty products equal one; in particular $d_{n,r} > 0$.

Proof. First-step conditioning gives $a_r d_{n,r} - b_r d_{n,r+1} = 1$ for $r < n$, while $a_n d_{n,n} = 1$. Backward substitution gives (7), and summing the increments gives (8). □

Put $q = 1/p$, $t = (1-s)/s$, and

$$C_p = \prod_{j=1}^{\infty} (1 - p^{-j}). \tag{9}$$

This product is positive and $C_p \ge 1/4$. Indeed its factors are at least those for $p = 2$, whose first factor is $1/2$ and whose remaining product is at least $1 - \sum_{j \ge 2} 2^{-j} = 1/2$.

Lemma 3.3 (Interval-product bounds)**.** *For $1 \le r \le k \le n$ and $0 < c < 1$,*

$$C_p e^{-t/(p-1)} R^{k-r} \le \prod_{j=r}^{k-1} \frac{b_j}{a_j} \le C_p^{-1} R^{k-r}. \tag{10}$$

Moreover, with $S_n(R) = \sum_{\ell=0}^{n-1} (n-\ell) R^{\ell}$,

$$p C_p e^{-t/(p-1)} S_n(R) \le g_n(n) \le \frac{p}{s(1-1/p) C_p} S_n(R). \tag{11}$$

Proof. For $1 \le j < n$, direct division gives

$$\frac{b_j}{a_j} = R \frac{1 - p^{j-n}}{(1 - p^{-j})(1 + t p^{j-n})}. \tag{12}$$

Any interval product of the factors $1 - p^{j-n}$ or $1 - p^{-j}$ lies in $[C_p, 1]$. Also

$$1 \le \prod_{j=r}^{k-1} (1 + t p^{j-n}) \le \exp\left(t \sum_{h=1}^{\infty} p^{-h} \right) = e^{t/(p-1)}.$$

This proves (10), including empty products. Since $s(1-1/p)/p \le a_k \le 1/p$, substitution into (8) gives (11). □

4 Actual recording and finite-confidence budgets

First exposure of a computational-basis target does not guarantee that the next measurement records it. This distinction can change the subcritical leading constant.

Theorem 4.1 (Exact recording correction)**.** *Let* $h_n = (1-c)p^{-n}$ *and* $F_n = c + h_n$*. For every initial rank* r*,*

$$\mathbb{E}_r\sigma = g_n(r) + H_n, \qquad H_n = \frac{1-F_n}{F_n}\big[1 + (1-h_n)g_n(1)\big]. \tag{13}$$

At $c = 1$ *this reads* $H_n = 0$*.*

Proof. An exact target candidate is impossible before rank zero. At a rank-zero frame the target is a computational vector z. Its first measurement is $x = z$ with probability F_n, in which case the recording index is zero. Conditional on failure $x \neq z$, the independent second measurement equals this particular x with probability h_n. Then $u = 0$, and the next rank is zero. Otherwise $u \neq 0$ is outside the zero-dimensional target X-projection, so the next rank is one by Lemma 2.2.

Let H_n be the expected recording index from rank zero and let $H_{n,1}$ be that from rank one. Strong Markov conditioning gives $H_{n,1} = g_n(1) + H_n$. First-step conditioning, with the failed recording and update kept jointly, yields

$$H_n = (1-F_n)\{1 + h_nH_n + (1-h_n)H_{n,1}\}.$$

Solving gives (13). These expectations are finite without assuming the identity: return times to zero have finite mean in the finite chain, and each visit has conditional recording probability $F_n > 0$. Successive excursions therefore give a geometric bound on their expected number. The same argument at the first rank-zero visit proves the formula for arbitrary r. □

Corollary 4.2 (Deterministic round budgets)**.** *Let*

$$M_n = g_n(n) + H_n, \qquad B_n = \lceil 2(M_n + 1)\rceil\,.$$

For every initial frame and integer $k \geq 1$*,*

$$\mathbb{P}(\sigma + 1 > kB_n) \leq 2^{-k}. \tag{14}$$

Thus running kB_n *rounds for* $k = \lceil \log_2(1/\delta)\rceil$ *includes the target candidate with probability at least* $1-\delta$*.*

Proof. By Proposition 3.2, $g_n(r) \leq g_n(n)$. Conditional on any complete history of unsuccessful rounds, the current input is again a depolarized stabilizer in some frame and the expected additional inclusive round count is at most $M_n + 1$. Markov's inequality bounds the probability of needing more than B_n further rounds by $1/2$. Iterating this conditional bound over k blocks proves (14). □

This is an implementable fixed-budget statement; it does not require an oracle announcing τ_0 or σ. It costs twice the stated number of rounds in unknown-state copies for candidate generation. For this parameterized budget, p, n, c are supplied; alternatively any certified upper bound on M_n may be used. The finite sums above determine the budget from these parameters without knowledge of the target. Final selection from the list requires additional fresh copies and computation. For example, direct independent measurements of each listed stabilizer projector give additive accuracy ε for all N candidates using $O(N\varepsilon^{-2}\log(N/\delta))$ copies by Hoeffding's inequality. This elementary validation cost is separate from our exposure analysis; the sharper shadow-based validation in [1] may instead be used when its hypotheses and implementation are supplied.

5 Three regimes and the critical window

Theorem 5.1 (Mean phase transition)**.** *Fix p and $0 < c < 1$. If $R > 1$, then $g_n(n)$ and $\mathbb{E}_n\sigma$ are $\Theta_{p,c}(R^n)$. If $R = 1$, both are $\Theta_p(n^2)$. If $R < 1$, both are $\Theta_{p,c}(n)$. At the pure endpoint $c = 1$, $\mathbb{E}_n\sigma = pn + O_p(1)$.*

Proof. For $R < 1$, $n \le S_n(R) \le n/(1-R)$. For $R = 1$, $S_n(R) = n(n+1)/2$. For $R > 1$, setting $j = n - 1 - \ell$ shows

$$R^{n-1} \le S_n(R) = R^{n-1}\sum_{j=0}^{n-1}(j+1)R^{-j} \le \frac{R^{n-1}}{(1-R^{-1})^2}.$$

Lemma 3.3 gives the asserted orders for $g_n(n)$. Its same estimates applied to (7) at $r = 1$ give $g_n(1) = O_{p,c}(R^n)$ when $R > 1$, $O_p(n)$ when $R = 1$, and $O_{p,c}(1)$ when $R < 1$. Since $F_n \ge c > 0$, the correction (13) has these same upper orders. The lower orders transfer from $\sigma \ge \tau_0$. If $c = 1$, the chain is pure death with $a_r = (1 - p^{-r})/p$, so

$$\mathbb{E}_n\sigma = \sum_{r=1}^{n}\frac{p}{1-p^{-r}} = pn + p\sum_{r=1}^{n}\frac{1}{p^r - 1} = pn + O_p(1).$$

This endpoint is already the noiseless analysis of [1]. □

Theorem 5.2 (Critical-window profile)**.** *Fix p. Suppose $c_n \in (0,1)$ and*

$$R_n = \frac{p(1-c_n^2)}{c_n^2}, \qquad n\log R_n \longrightarrow \lambda \in \mathbb{R}.$$

Then

$$\frac{g_n(n)}{n^2},\ \frac{\mathbb{E}_n\sigma}{n^2} \longrightarrow (p+1)\int_0^1 (1-x)e^{\lambda x}\,dx = \begin{cases}(p+1)\dfrac{e^\lambda - 1 - \lambda}{\lambda^2}, & \lambda \ne 0,\\ (p+1)/2, & \lambda = 0.\end{cases} \tag{15}$$

In particular, at $c^2 = p/(p+1)$ the exact leading mean is $(p+1)n^2/2$.

Proof. The hypothesis gives $c_n^2 \to p/(p+1)$ and $t_n \to 1/p$. For fixed $\eta > 0$, uniformly on $\eta n \le r \le k \le (1-\eta)n$, the product in (10) divided by R_n^{k-r} tends to one. To see this, use (12): the sums of all deviations from one in its three correction products are $O_p(p^{-\eta n})$. On the same region, $1/a_k \to p+1$ uniformly.

The bounds in Lemma 3.3 are uniform for all large n, and R_n^n remains bounded above and below. Thus every summand in (8) is bounded by a constant independent of its indices. The at most $O(\eta n^2)$ pairs outside the indicated interior contribute $O(\eta)$ after division by n^2. Inside it, the sum converges to the corresponding Riemann integral. Sending η to zero proves

$$n^{-2}g_n(n) \to (p+1)\int_{0\le x\le y\le 1} e^{\lambda(y-x)}\,dx\,dy,$$

which equals (15). The same uniform estimates give $g_n(1) = O_{p,\lambda}(n)$ along this window. Equation (13) therefore adds $O_{p,\lambda}(n)$ and leaves the limit unchanged. □

The regime statements are for fixed p. The constants and exact formulas remain available at every specified finite prime.

6 Exact fixed-regime constants

Theorem 6.1 (Supercritical speed)**.** *For fixed p, c with $c^2 > p/(p+1)$,*

$$\frac{g_n(n)}{n},\ \frac{\mathbb{E}_n\sigma}{n} \longrightarrow \frac{1}{c^2/p - (1-c^2)}. \tag{16}$$

Proof. Here $R < 1$. The summands in (7), indexed by $h = k - r$, are bounded uniformly by a constant times R^h. For each fixed h, when $r \to \infty$ and $n-r \to \infty$, their factors satisfy $1/a_{r+h} \to p/s$ and $b_j/a_j \to R$. Dominated convergence against the geometric majorant gives, uniformly away from the two boundaries,

$$d_{n,r} \to \frac{p}{s}\sum_{h=0}^{\infty} R^h = \frac{p}{s(1-R)} = \frac{1}{\mu}.$$

Explicitly, first truncate h so its geometric tail is uniformly small, then remove a fixed number of indices at each boundary to control all remaining coefficients. Every $d_{n,r}$ is uniformly bounded, so those boundary indices have vanishing contribution to $n^{-1}\sum_r d_{n,r}$. This proves the first limit. Here $g_n(1) = O_{p,c}(1)$, and the bounded recording correction proves the second. □

For the subcritical constants let

$$P_m = \prod_{j=1}^{m}(1-q^j), \quad P_0 = 1, \qquad Q_m = \prod_{j=0}^{m}(1+tq^j), \quad Q_{-1} = 1, \tag{17}$$

and $Q_\infty = \prod_{j=0}^{\infty}(1+tq^j)$.

Theorem 6.2 (Subcritical prefactors)**.** *Fix p, c with $0 < c^2 < p/(p+1)$, and put $z = 1/R \in (0,1)$. Define the positive finite constants*

$$A = \sum_{r=1}^{\infty} z^r P_{r-1}, \qquad B = \sum_{\ell=0}^{\infty} z^\ell \frac{Q_{\ell-1}}{P_\ell}. \tag{18}$$

Then

$$\frac{g_n(n)}{R^n} \longrightarrow \frac{p}{s}\frac{AB}{Q_\infty}, \tag{19}$$

$$\frac{g_n(1)}{R^n} \longrightarrow \frac{p}{s}\frac{zB}{Q_\infty}, \tag{20}$$

$$\frac{\mathbb{E}_n\sigma}{R^n} \longrightarrow K_{p,c} := \frac{p}{s}\frac{B}{Q_\infty}\left[A + \frac{1-c}{c}z\right]. \tag{21}$$

Proof. Let $T_{n,r,k}$ denote the summand in (8). Combining $1/a_k$ with the product in (12) gives exactly

$$T_{n,r,k} = \frac{p}{s}R^{k-r}\frac{\prod_{j=r}^{k-1}(1-q^{n-j})}{\prod_{j=r}^{k}(1-q^j)\prod_{j=r}^{k}(1+tq^{n-j})}. \tag{22}$$

Set $k = n - \ell$. For each fixed $r \geq 1$ and $\ell \geq 0$, the numerator product tends to C_p/P_ℓ, and the two denominator products tend to C_p/P_{r-1} and $Q_\infty/Q_{\ell-1}$, respectively. Thus

$$\frac{T_{n,r,n-\ell}}{R^n} \longrightarrow \frac{p}{s}z^{r+\ell}\frac{P_{r-1}Q_{\ell-1}}{P_\ell Q_\infty}.$$

Extend this expression by zero when $r+\ell > n$. Lemma 3.3 bounds it by a constant times $z^{r+\ell}$, summable on $r \geq 1, \ell \geq 0$. Dominated convergence for the double series proves (19). Restricting the same calculation to $r = 1$ proves (20). Finally $F_n \to c$, $h_n \to 0$, and $R > 1$, so the constant term in (13) divided by R^n vanishes. Substituting (20) into that identity proves (21). □

The prefactors are explicit approximable constants, not only qualitative existence statements. The elementary bounds

$$1 \leq Q_\infty \leq e^{t/(1-q)}, \qquad C_p \geq 1/4$$

give, for $M \geq 0$,

$$\begin{aligned} 0 &\leq \sum_{r=M+1}^{\infty} z^r P_{r-1} \leq \frac{z^{M+1}}{1-z}, \\ 0 &\leq \sum_{\ell=M+1}^{\infty} z^\ell \frac{Q_{\ell-1}}{P_\ell} \leq \frac{Q_\infty}{C_p} \frac{z^{M+1}}{1-z}, \\ 1 &\leq \frac{Q_\infty}{Q_M} \leq \exp\left(\frac{tq^{M+1}}{1-q}\right). \end{aligned}$$

They allow certified truncation using outward arithmetic if numerical prefactors are desired.

7 A finite-time failure theorem

Theorem 7.1 (Full-height exponential barrier)**.** *Fix p and $0 < c < 1$ with $R > 1$. For $n \geq 2$ and every integer $T \geq 0$,*

$$\mathbb{P}_n(\sigma \leq T) \leq \mathbb{P}_n(\tau_0 \leq T) \leq T C_p^{-1} e^{t/(p-1)} R^{-(n-1)}. \tag{23}$$

In particular, the unchanged candidate generator has exponentially small target-inclusion probability in any polynomial round budget.

Proof. For the chain stopped upon hitting 0 or n, let v_r be the probability of hitting zero first from r. Holding probabilities cancel from its first-step equations:

$$b_r(v_{r+1} - v_r) = a_r(v_r - v_{r-1}), \qquad v_0 = 1,\ v_n = 0.$$

With $w_0 = 1$ and $w_j = \prod_{i=1}^{j}(a_i/b_i)$ for $j < n$, solving the increment recurrence gives

$$v_{n-1} = \frac{w_{n-1}}{\sum_{j=0}^{n-1} w_j} \leq w_{n-1} \leq C_p^{-1} e^{t/(p-1)} R^{-(n-1)},$$

where the last inequality is the lower product bound (10) inverted on $1, \ldots, n-1$. Every path from n to zero must succeed during an excursion beginning with a downward step n to $n-1$. There are at most T such excursions in T transitions. Applying the strong Markov property at their entrance times and a union bound proves the second inequality in (23). Recording requires prior exposure, proving the first. The statement remains valid when its right side exceeds one. □

The mean lower bound has the same exponential rate. Indeed the factors $1 - p^{j-n}$ and $1 - p^{-j}$ cancel over the full interval, giving

$$\prod_{j=1}^{n-1} \frac{b_j}{a_j} = \frac{R^{n-1}}{\prod_{j=1}^{n-1}(1 + tp^{-j})}, \qquad g_n(n) \geq p e^{-t/(p-1)} R^{n-1}.$$

The finite-time bound is stronger than a mean lower bound alone: it rules out success hidden in a large mass of short trajectories with a small number of extremely long failures.

Lemma 7.2 (Stabilizer overlap gap)**.** *Distinct pure n-qudit stabilizers in prime dimension satisfy $|\langle s'|s\rangle|^2 \le 1/p$. For (1), the target is the unique maximizing stabilizer, and every other stabilizer's overlap is smaller by at least $c(1-1/p)$.*

Proof. The stabilizer projector is the average of its p^n phase-corrected Pauli stabilizers. Orthogonality of Pauli operators shows that the trace of two such projectors is either zero, if their characters disagree on the intersection, or p^{k-n}, where k is the intersection dimension of their unsigned spaces. If the states are distinct and the trace is nonzero, $k \le n-1$; when $k = n$ and the characters agree, the states are the same. Thus the trace is at most $1/p$. The identity component in (1) contributes the same $(1-c)p^{-n}$ to every pure-state overlap, so the gap is at least $c(1-1/p)$. □

Corollary 7.3 (Failure of the unchanged agnostic extension)**.** *For a fixed subcritical c and any $0 < \varepsilon < c(1-1/p)$, success of the unchanged candidate-and-select procedure at additive accuracy ε requires target inclusion. Consequently any polynomial candidate round budget has success probability tending to zero exponentially, even if final selection from the list is perfect.*

Proof. Every non-target listed stabilizer fails the required overlap inequality by Lemma 7.2. Apply Theorem 7.1. □

8 Interpretation and reproducibility

The limiting bulk downward probability is c^2/p, while the upward probability is $1-c^2$. Their difference is μ in (3). This explains the location of the transition. The finite-dimensional boundary factors do more than perturb this drift: they determine the critical mean profile and the subcritical prefactors. The exact recording correction further separates exposure of a suitable frame from obtaining the correct observed eigenvalue.

The critical result gives an actual finite-confidence inclusion guarantee of $O_{p,K}(n^2 \log(1/\delta))$ rounds for this promised input family, uniformly in critical windows with $|n \log R_n| \le K$ for fixed K. The generic block argument gives $O_{p,c}(n \log(1/\delta))$ rounds above threshold. Below threshold, the failure bound concerns this prescribed update and output-list construction. It neither excludes guarded updates with a genuinely different discovery mechanism nor shows that global depolarization makes stabilizer recovery information theoretically impossible.

Known classical obstacles also require precise scope. Over $\mathbb{F}_2$, recovering an unknown hyperplane from a mixture of its uniform distribution and uniform full-space noise is exactly a version of learning parity with noise; Chen, De and Vijayaraghavan [6, Section 6] give the explicit reduction. This explains why an unspecified "recover the noisy support" subroutine is not an immediate repair. It is a known classical barrier, not an additional theorem about arbitrary quantum measurements.

The analytic proofs here establish all-n statements. Finite symplectic checks of gate actions and exact transition probabilities, and direct evaluations of the finite recurrence, provide reproducible supporting tests. The supplied finite controls use exact arithmetic; the asymptotic claims rest on the analytic proofs.

AI assistance

AI tools assisted with derivation, programming, literature discovery and proof checking. Analytic arguments and finite computations have distinct roles. AI-assisted internal checking was used.

References

[1] 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). arXiv:2610.02031v1.

[2] S. Chen, W. Gong, Q. Ye and Z. Zhang, *Stabilizer bootstrapping: A recipe for efficient agnostic tomography and magic estimation*, Proceedings of STOC 2025, 429–438. doi:10.1145/3717823.3718191.

[3] W. Qi, B. Xu, S. Feng and L. Li, *Agnostic learning of qudit stabilizer states*, arXiv:2607.15559v1 (2026). arXiv:2607.15559v1.

[4] S. Schubr, *Mixed-state escape bounds and proper agnostic learning of stabilizer states in odd-prime dimensions*, updated author-published edition, October 19, 2026. Luminary Research Archive, edition 2.

[5] L. Tang and T. Steckmann, *Learning Noise-Robust Stabilizer Structure via Bell Sampling*, arXiv:2610.03307v2 (2026). arXiv:2610.03307v2.

[6] A. Chen, A. De and A. Vijayaraghavan, *Learning a mixture of two subspaces over finite fields*, arXiv:2010.02841v2 (2021). arXiv:2010.02841v2.

[7] J. R. Norris, *Markov Chains*, Cambridge University Press, 1997.

[8] D. A. Levin and Y. Peres, with contributions by E. L. Wilmer, *Markov Chains and Mixing Times*, second edition, American Mathematical Society, 2017.