

Single-copy learning of noisy stabilizers by certified filtering

Scott Schubr

October 22, 2026

Abstract

We give a single-copy recovery framework for an unknown qubit stabilizer state mixed with global white noise. For every fixed $\varepsilon > 0$, $0 < \gamma < 1$ and fixed confidence, the exact target can be recovered from $O(n^{1+\varepsilon})$ original copies when the retained signal has a known lower bound $c_0 \geq 2^{-(\log_2 n)^\gamma}$. The total number of ideal quantum operations is polynomial, while classical processing takes $2^{O(n/\log \log n)}$ time. A random Clifford frame and a known partial-derivative measurement reduce discovery of one complete target Pauli to classical noisy-parity search. Fresh verification and a short certified seed then amplify the signal enough for known high-fidelity single-copy learning. This avoids recovering every graph row through a noisy solver. We also obtain polynomial classical time for every fixed positive signal under a supplied fractional diagonal-subspace promise, including dense CSS states and a CSS-defect class. The latter uses finite-field subspace recovery, with an explicit proof of the required fixed-weight extension. Supporting bounds distinguish partial code measurements from full-basis readout. The measurements, classical solvers and high-fidelity endpoint are established ingredients; the contribution is their operational combination and original-copy resource accounting.

1 Introduction and main results

The distinction between quantum sample access and classical processing is important in stabilizer learning. Efficient pure-state learning does not automatically remain efficient under a global mixture with the maximally mixed state. Bittel, Eisert, Gong, Mele and Schatzki [3] give adaptive single-copy algorithms for pure stabilizers and for sufficiently high stabilizer fidelity. For the literal update in that work, the global-noise analysis of [9] exhibits a linear, quadratic or exponential target-recording time according to the signal regime. That is an algorithm-specific barrier, not a lower bound on alternative single-copy procedures. The present paper changes the discovery mechanism.

Throughout, the input consists of independent copies of

$$\rho = c|s\rangle\langle s| + (1 - c)I/2^n, \quad 0 < c_0 \leq c \leq 1, \quad (1)$$

where $|s\rangle$ is an unknown pure n -qubit stabilizer and c_0 is supplied. The actual c need not be known. Operations and readout are ideal. One unknown copy is operated on at a time, with classical memory and adaptation but no coherent memory coupling different input copies. Clifford gates may entangle qubits *within* one copy.

Theorem 1.1 (General recovery with nearly linear copies). *Fix $\varepsilon > 0$, $0 < \gamma < 1$ and $\delta \in (0, 1/3)$. For inputs (1) with a supplied rational floor satisfying*

$$c_0 \geq 2^{-(\log_2 n)^\gamma}, \quad (2)$$

there is a randomized single-copy algorithm returning the exact signed tableau of s with probability at least $1 - \delta$. As $n \rightarrow \infty$, it uses

$$O_{\varepsilon, \gamma, \delta}(n^{1+\varepsilon}) \quad \text{original input copies,}$$

polynomially many total ideal quantum operations, and classical time $2^{O_{\varepsilon, \gamma, \delta}(n/\log \log n)}$. Finite smaller dimensions have bounded fallback branches. Input-encoding costs are additional to the displayed bounds.

The common mechanism is a short genuinely certified target seed. Fresh checks establish its signed constraints, conditioning on verified constraints raises the signal, and the established high-fidelity learner completes recovery. Reusing discovery data across the noise grid changes the accounting of classical calls, while the copy ledger counts original preparation trials, including rejected ones.

Every fixed positive signal is included asymptotically. The confidence and the two exponents are fixed parameters of this statement. Nearly linear describes copies, not gates or classical time.

The discovery step can instead become polynomial under a structural promise. Let $S_s \leq \mathbb{F}_2^{2n}$ be the target's unsigned stabilizer Lagrangian, and let $\mathcal{Z} = \{(0, z) : z \in \mathbb{F}_2^n\}$.

Theorem 1.2 (Polynomial time with a visible diagonal subspace). *Fix rational constants $c_0 > 0$ and $\beta \in (0, 1]$. Supply a finite list $\mathcal{C}$ of known Clifford circuits, with polynomial list length and gate size, and promise that some $C \in \mathcal{C}$ satisfies*

$$\dim(CS_s \cap \mathcal{Z}) \geq \beta n. \tag{3}$$

For (1), exact signed-tableau recovery with failure at most $\delta \in (0, 1/3)$ uses time and original copies polynomial in n , the list and circuit sizes, and $\log(1/\delta)$, for these fixed c_0, β . The polynomial degree and finite-dimensional constants may depend on c_0, β .

Let D_Z and D_X be the Z -only and X -only target direction spaces, and put $\kappa = n - \dim D_Z - \dim D_X$. The spaces are disjoint in S_s , so $\kappa \geq 0$.

Corollary 1.3 (CSS-defect class). *For any fixed $\alpha < 1$, the promise $\kappa \leq \alpha n$ implies Theorem 1.2 with the two frames $I, H^{\otimes n}$ and $\beta = (1 - \alpha)/2$. In particular every signed CSS stabilizer state is covered, with no bounded-weight assumption on its unknown generators.*

Proof. One of $\dim D_Z, \dim D_X$ is at least $(1 - \alpha)n/2$. A Hadamard on every qubit exchanges these two spaces. For CSS states their direct sum is the full stabilizer, so $\kappa = 0$. $\square$

An unknown Clifford equivalence to CSS does not supply the known frame required by this corollary.

1.1 Predecessors and the quantitative increment

Arunachalam, Bravyi, Dutt and Yoder [1] explicitly introduce single-copy random partial-derivative measurements for phase states. Their global-noise quadratic-phase result uses two-copy Bell sampling and Lyubashevsky's classical parity algorithm [2]. We use the known derivative primitive, its X/Y version for complex stabilizer phases, and the same classical solver. Random full-support Clifford frames are also known; see [8, Lemma 6.1]. Uniform Clifford generation is polynomial by [7].

Simply applying noisy derivative learning to every row in a favorable frame gives a polynomial-copy, subexponential-time construction, with roughly n^{2+a} training copies when a single row uses n^{1+a} .

We recover only $O(1 + \log(1/c_0))$ certified constraints and then use high-fidelity learning [3]. The resulting original-copy count is $n^{1+a+o(1)}$, including rejected preparations and unknown-noise handling. The weak-bias range is inherited from [2].

For Theorem 1.2, the main classical mechanism is the polynomial lift of Chen, De and Vijayaraghavan [4]. Their printed Theorem 2 has a numerical minimum-weight hypothesis. Section 5 proves the fixed-positive-weight specialization needed here rather than silently invoking that statement below its printed floor. Code conditioning and capped preparation are standard tools, also used in the mixed-input framework of [10]. Its coherent multi-copy sampler is not used here.

2 Stabilizer preliminaries and external algorithms

We use Hermitian Paulis $W_{x,z} = i^{x \cdot z} X^x Z^z$, with the exponent interpreted as the integer overlap count modulo four. An unsigned label is $(x, z) \in \mathbb{F}_2^{2n}$; two labels commute when their symplectic product vanishes. A signed tableau specifies independent commuting Hermitian operators and their consistent eigenvalues. Clifford conjugation includes its known phase. Except for logarithms in exponential tail bounds, log means $\log_2$. We write $\ln$ for the natural logarithm.

For any nonidentity Hermitian Pauli W ,

$$\text{Tr}(\rho W) \in \{0, c, -c\}. \quad (4)$$

The nonzero alternatives occur exactly on target directions. For distinct pure stabilizers s, t , the squared overlap is either zero or a negative integer power of two, and in particular

$$\langle s | \rho | s \rangle - \langle t | \rho | t \rangle = c(1 - |\langle s | t \rangle|^2) \geq c/2. \quad (5)$$

These facts follow by expanding stabilizer projectors in their commuting Pauli groups. The same expansion gives the uniform computational distribution on an affine subspace $z + A$, with $A = D^\perp$ when D is the diagonal target direction space.

We use two published algorithmic inputs, with different resource roles.

Proposition 2.1 (Classical noisy-parity input, Lyubashevsky [2]). *Fix $a > 0$ and $0 < \theta < 1$. Given $O(d^{1+a})$ independent examples $(z, a_* \cdot z + e)$ with uniform $z \in \mathbb{F}_2^d$ and independent bit errors of correctness advantage $\eta > 2^{-(\log d)^\theta}$, the unknown a_* can be recovered in time $2^{O_{a,\theta}(d/\log \log d)}$, with success tending to one. For sufficiently large d we use success at least $9/10$.*

The source uses η as correctness advantage, so the error probability is $1/2 - \eta$. Its original theorem has independent sample and bias parameters. The proof and its BKW input provide budgets from a lower advantage bound. For additional protection against any convention that supplies the exact rate, we give a finite grid in Section 4. All calls are capped at the stated resource budget, even when their samples are off-promise. The pure-label endpoint is handled by linear elimination. Uniform secret masking may be used if a solver is stated for a uniform rather than a fixed secret.

Proposition 2.2 (High-fidelity endpoint, Bittel et al. [3]). *For any fixed known $f_0 > \sqrt{2/3}$, adaptive single-copy learning of an m -qubit mixed input with stabilizer fidelity at least f_0 returns a stabilizer of overlap within a prescribed additive $\xi > 0$ of the optimum, with failure at most ζ . At fixed f_0, ξ , its copy count is $O(m + \log(1/\zeta))$ and its time and gate costs are polynomial.*

We use only $f_0 = 9/10$ and $\xi = 1/10$. By (5), a successful output on a global mixture of signal at least $9/10$ is the exact target.

3 Certified constraints and paid conditioning

Lemma 3.1 (Single-copy filter). *Suppose k independent correctly signed commuting target constraints are known, and let P_k be their joint projector. Its rank is 2^{n-k} . It accepts an original copy with probability*

$$q_k = \text{Tr}(P_k \rho) = c + (1 - c)2^{-k} \geq c_0. \quad (6)$$

After a known Clifford decoder maps the constraints to coordinate Z operators, measure just those k coordinates and retain their required outcomes. The remaining logical state is exactly

$$\rho_k = c_k |s_k\rangle\langle s_k| + (1 - c_k)I/2^{n-k}, \quad c_k = \frac{c}{c + (1 - c)2^{-k}}. \quad (7)$$

Proof. The target is fixed by P_k , while the uniform component passes with probability 2^{-k} . Thus $P_k \rho P_k = c|s\rangle\langle s| + (1 - c)P_k/2^n$. Normalize, decode, and discard the fixed coordinates. $\square$

For definiteness let

$$K = \lceil \log(9/c_0) \rceil + 1. \quad (8)$$

At $k = K$ the logical signal exceeds 9/10. If $n < K$, recovering all n constraints already determines the target. Smaller sufficient seeds may be used, but are unnecessary for the bounds.

Lemma 3.2 (Validation and cumulative preparation). *For at most B adaptively proposed non-identity physical Paulis, fresh means accurate to $c_0/4$ certify membership and sign by the threshold $|\hat{\mu}| \geq c_0/2$. Accuracy holds simultaneously with failure at most ζ using*

$$v = \lceil 32c_0^{-2} \ln(2B/\zeta) \rceil \quad (9)$$

copies per observable. If a commuting candidate family shares a known measurement basis, the same v fresh basis measurements estimate all its means simultaneously. If at most Q conditional copies are requested, cap each request at

$$L = \lceil c_0^{-1} \ln(Q/\zeta) \rceil \quad (10)$$

original trials through the entire current prefix. On correctly certified prefixes, any cap is exhausted with probability at most ζ .

Proof. For a $\{-1, 1\}$ outcome Hoeffding gives $\mathbb{P}(|\hat{\mu} - \mu| > c_0/4) \leq 2e^{-vc_0^2/32}$. Equation (4) proves threshold separation and the sign. Condition on each proposed list before taking its independent validation data and union-bound over the predetermined slots. Within a commuting list the estimates may be correlated; concentration and a union bound do not require independence between observables.

Every correct full-prefix trial accepts with probability at least c_0 , so L failures have probability at most $e^{-c_0 L}$. Union-bound over Q requests. Accepted copies have the exact conditional law. Couple capped and uncapped preparations until the first cap or other charged error; no correctness is assumed after that error. Even on a bad prefix the algorithm retains the same finite caps. $\square$

An unsigned finder of just one genuine nonidentity logical target Pauli therefore suffices for full learning. Lift it with identity on the known prefix, validate on original copies, add its sign, and repeat at most $\min(n, K)$ times before Proposition 2.2. A nonidentity logical lift is independent of the prefix. Its physical expectation is still zero or $\pm c$, rather than an unobservable hidden membership bit. This reduction is uniform with polynomial overhead in $n, 1/c_0$ whenever the assumed finder has such overhead. It is not itself a construction of a polynomial-time general finder.

4 A general finder from noisy-parity examples

4.1 Random frames and one-row measurements

Lemma 4.1 (Known full-support count). *A uniform Clifford sends any fixed stabilizer to full computational support with probability*

$$p_m = \prod_{j=1}^m (1 + 2^{-j})^{-1} > e^{-1} > 1/4 \quad (11)$$

on m qubits.

Proof. There are $\prod_{j=1}^m (2^j + 1)$ binary Lagrangians. Those disjoint from $\mathcal{Z}$ are precisely graphs $\{(x, Ax)\}$ of symmetric binary matrices, of which there are $2^{m(m+1)/2}$. Divide these counts and use $\ln(1 + u) < u$. This standard fact is also stated as folklore in [8, Lemma 6.1]. $\square$

Lemma 4.2 (Exact noisy star law). *In a full-support frame, fix the first qubit. Some target generator is*

$$W = P_1 \otimes Z^{a_*}, \quad P \in \{X, Y\}, \quad W|s\rangle = (-1)^b |s\rangle.$$

Measuring P on that qubit and Z on the others gives features $z \in \mathbb{F}_2^{m-1}$ and outcome bit y with

$$\mathbb{P}(z, y) = 2^{-m} [1 + c(-1)^{y+a_* \cdot z+b}]. \quad (12)$$

Thus the features are uniform and the affine parity is corrupted by independent binary noise of rate $(1 - c)/2$.

Proof. The X -projection of the target Lagrangian is bijective, so the unique direction with X -label e_1 has the stated form, including a possible Y at the first qubit. Full-support amplitudes make the other Z outcomes uniform. On the pure component their parity fixes the pivot outcome by the signed generator equation. On the uniform component all m bits are uniform. Combining the two laws proves (12). $\square$

The correlation bias is c and the correctness advantage is $c/2$. This uses an ordinary single-copy measurement, not a phase-evaluation oracle. For real phase states the X/Z primitive is exactly the random partial-derivative measurement of [1]. Trying Y also covers complex stabilizer phases.

4.2 Unknown intercept, unknown rate and bad frames

For each frame try both pivot axes, using separate fresh training datasets of size $N_m = O((m - 1)^{1+a})$. For each dataset try the two label shifts $t \in \{0, 1\}$. On the correct axis and $t = b$, the labels $y + t$ are standard linear LPN labels. An independent uniform r and the replacement $y + t \mapsto y + t + r \cdot z$ make the solver's secret $a_* + r$ uniform when necessary; unmask its proposed output.

Let $\eta = c_k/2$ be the unknown correctness advantage at a logical stage. Place a grid on $[c_0/2, 1/2]$, including endpoints, with mesh at most $1/(100N_m)$ and at most $100N_m + 2$ points. Reuse the *same* training data in all its classical calls. A grid point η' satisfies $|\eta - \eta'| \leq 1/(100N_m)$. Coupling the independent error bits gives

$$\text{TV}(\mathcal{D}_\eta^{N_m}, \mathcal{D}_{\eta'}^{N_m}) \leq N_m |\eta - \eta'| \leq 1/100. \quad (13)$$

Here the feature and secret laws are identical. Consequently a solver with success 9/10 at η' retains success at least 89/100 on the actual data. The endpoint $\eta' = 1/2$ uses Gaussian elimination. All grid calls are bounded, and their dependence through data reuse is harmless: only one suitable call is needed. No quantum copy is reused.

Convert every proposed tail to its complete Pauli in the original coordinates. In a logical stage lift it with identity on the prefix. For a fixed frame and pivot axis all these candidates commute. A held-out measurement in that known decoded frame gives the eigenvalue of every tail as $(-1)^{y+a \cdot z}$, including the known Clifford phase. For validation, measure a fresh *original* input through the full decoder and logical frame and ignore prefix outcomes; no prefix postselection is needed. Equation (4) and Lemma 3.2 then certify any accepted tail and sign, even in a bad frame. Invalid output encodings are discarded.

For sufficiently large m in the admissible bias range, a trial has success probability at least $(1/4)(89/100) > 1/8$: a favorable frame and its correct call suffice. If an earlier call gives a different genuine constraint, it is equally useful. An off-promise call may do anything within its cap; its output is still tested. Repeating

$$R = \lceil 8 \ln(16K/\delta) \rceil \quad (14)$$

independent frame/dataset trials per seed stage leaves failure at most $\delta/(16K)$ per stage, apart from the separately budgeted validation and sampling errors.

4.3 Original-copy ledger and proof of Theorem 1.1

Set $a = \min(\varepsilon/4, 1/4)$ and $\theta = (1+\gamma)/2$. Before K constraints are found, every noisy-parity feature dimension is

$$d = n - O(1 + (\log n)^\gamma).$$

Uniformly over these stages and all grid points, for sufficiently large n ,

$$\eta \geq c_0/2 \geq 2^{-(\log n)^\gamma - 1} > 2^{-(\log d)^\theta}. \quad (15)$$

Indeed $(\log d)^\theta - (\log n)^\gamma \rightarrow \infty$. This is precisely the two-parameter range of Proposition 2.1; the supplied original floor suffices at every logical stage.

Let N bound all training-dataset sizes, so $N = O(n^{1+a})$, and let $J = 100N + 2$ bound the grid size. There are at most

$$B = 4KRJ$$

candidate slots. Let $Q_H = O(n + \ln(1/\delta))$ be the declared endpoint copy budget from Proposition 2.2, with allocated failure $\delta/4$, and put $Q = 2KRN + Q_H$. Choose validation size v by (9) with $\zeta = \delta/16$, and preparation cap L by (10) with the same ζ . Shared validation needs at most $2KR$ datasets, one for each attempted frame/axis. A deterministic original-copy upper bound is

$$QL + 2KRv. \quad (16)$$

The grid increases the classical calls and the logarithm in v ; it does not multiply quantum training copies by J . Even separate verification of every candidate would give $QL + Bv$ and the same claimed exponent.

Uniform Clifford sampling may be implemented by the integer indexing and factorization of [7]; cap classical rejection in uniform integer generation with total failure $\delta/16$. Allocate $\delta/16$ each to absent successful finder stages, inaccurate validation, and exhausted conditional preparations, and $\delta/4$ to the final learner. The total is below δ . On the complementary event every prefix is correct, so the exact logical-state law and all finder guarantees apply successively. At K constraints the

signal exceeds 9/10; Proposition 2.2 and (5) recover the exact residual tableau. If all constraints have been found, the signed prefix already specifies the target. Lift through the known decoder to finish.

For fixed $\varepsilon, \gamma, \delta$,

$$K, R, c_0^{-1}, \ln(B/\delta), \ln(Q/\delta) = n^{o(1)}.$$

Equation (16) is therefore $n^{1+a+o(1)} = O(n^{1+\varepsilon})$. There are $O(KRJ)$ classical solver calls, a polynomial number in n ; their polynomial factor preserves $2^{O(n/\log \log n)}$ time. Clifford synthesis, decoding and measurement use polynomial gates per original copy. This proves the asymptotic resource statement.

Choose a sufficiently large fixed threshold at which the stated source solver success and (15) hold uniformly. Below it, enumerate all nonidentity Paulis, estimate means to $c_0/4$ with a confidence union bound, and recover n independent correctly signed constraints. The lower bound (2) is positive on these finitely many dimensions, so this adds only parameter-dependent constants. All failure branches return a designated tableau within the same declared caps. The case $n = 1$ is direct Pauli estimation. This completes the proof of Theorem 1.1.

5 Polynomial-time discovery from a visible subspace

We prove the classical specialization used in Theorem 1.2. This makes the minimum-weight dependence explicit. Write $\binom{d}{\leq \ell} = \sum_{j=0}^{\min(d, \ell)} \binom{d}{j}$ and let $\phi_\ell(x)$ be the vector of all squarefree monomials of degree at most ℓ , including the constant.

Lemma 5.1 (Known extremal rank ingredient). *If $T \subseteq \mathbb{F}_2^n$ has at least 2^r distinct points, then $\dim \text{span } \phi_\ell(T) \geq \binom{r}{\leq \ell}$.*

This is the extremal polynomial-rank result used as Lemmas 18–19 of [4], attributed there to [5] and equivalently to [6]. Also $\dim \text{span } \phi_\ell(A) = \binom{d}{\leq \ell}$ for a d -dimensional linear A : an invertible linear coordinate change preserves degree, and restriction to A gives the multilinear polynomial space on d variables.

Lemma 5.2 (Recovery at every fixed positive mixture weight). *Fix rational $w_0 \in (0, 1/2]$ and $\beta \in (0, 1]$. Given independent samples from*

$$w \text{Unif}(A) + (1 - w) \text{Unif}(\mathbb{F}_2^n), \quad w_0 \leq w \leq 1 - w_0, \quad \dim A \leq (1 - \beta)n, \quad (17)$$

there is a polynomial-time and polynomial-sample algorithm recovering A with constant probability at least 9/10, for fixed w_0, β . The algorithm does not know w or $\dim A$.

Proof. Choose a fixed integer $\ell \geq 2$ with $(1 - \beta/2)^\ell \leq w_0/100$, and put $N = \binom{n}{\leq \ell}$, $\varrho = 49w_0/100$, $\eta = \varrho/(4\ell)$. For sufficiently large n , require

$$n \geq 2\ell/\beta, \quad n \geq 8\ell^2/\varrho, \quad (18)$$

and

$$e^{-w_0 N/8} + N 2^{-\eta n} + 2^n e^{-4w_0 N/25} \leq 1/10. \quad (19)$$

These are effective conditions and hold eventually for fixed parameters.

Draw N samples x_i and mark precisely the indices for which $\phi_\ell(x_i)$ lies in the span of all other sampled columns. Return the span of the marked original vectors. All marks can be computed

by one nullspace calculation: an index is marked exactly when some kernel vector has a nonzero coordinate there. Thus the procedure is bounded and polynomial in N, n , even on off-promise data.

Let $d = \dim A$ and $D_A = \binom{d}{\leq \ell}$. By (18),

$$\frac{D_A}{N} \leq \frac{\binom{d+\ell}{\ell}}{\binom{n}{\ell}} \leq \left(\frac{d+\ell}{n} \right)^\ell \leq (1 - \beta/2)^\ell \leq w_0/100. \quad (20)$$

Label mixture components only for analysis. The number N_1 of inlier draws from A is at least $w_0 N/2$, except with probability $e^{-w_0 N/8}$.

Condition on all component labels satisfying that count and on every sample other than a full-space draw x_i . The span H_i of the other lifted columns together with the entire $\phi_\ell(A)$ has dimension

$$\dim H_i \leq N - 1 - N_1 + D_A \leq (1 - \varrho)N.$$

It is independent of the uniform x_i . Set $r = \lfloor (1 - \eta)n \rfloor$. Equation (18) gives

$$\begin{aligned} \frac{\binom{r}{\leq \ell}}{N} &\geq \frac{\binom{r}{\ell}}{\binom{n+\ell}{\ell}} \geq \left(\frac{r - \ell + 1}{n + \ell} \right)^\ell \\ &\geq (1 - \varrho/(2\ell))^\ell \geq 1 - \varrho/2 > 1 - \varrho. \end{aligned}$$

By Lemma 5.1, fewer than 2^r original points can have their lift in H_i . A full-space draw therefore lies there with probability at most $2^{-\eta n}$. Union-bound over N indices. Except with probability $N2^{-\eta n}$, no outlier column is marked. This also accounts for an outlier that accidentally lands in A .

Unmarked inlier columns are coloops of the lifted matrix and hence linearly independent among themselves. There are at most D_A of them, so the marked set contains at least $.98N_1$ inliers and no outliers. For any proper hyperplane of A , Hoeffding bounds the chance it contains at least $.9N_1$ uniform inliers by $e^{-8N_1/25}$. There are fewer than $2^d \leq 2^n$ such hyperplanes. Except with probability $2^n e^{-4w_0 N/25}$, every subset containing $.9N_1$ inliers spans A . For $d = 0$ spanning is automatic. Equation (19) proves the recovery guarantee. Below the finite size threshold one can learn the finite classical distribution and enumerate subspaces; for fixed w_0, β this changes only constants. In the quantum application we instead use the finite-dimensional Pauli fallback below. $\square$

This is a specialization of the polynomial-lift mechanism of [4] to a known full-space component. It does not assume their printed minimum weight $1/100$. The degree here is $\ell = O(1 + \beta^{-1} \log(1/w_0))$; its dependence is precisely why a joint polynomial-in- $1/c_0$ claim would be unjustified.

Proof of Theorem 1.2. In each supplied frame the pure computational law is uniform on $z + A$. Two fresh separate input measurements have difference law

$$c^2 \text{Unif}(A) + (1 - c^2) \text{Unif}(\mathbb{F}_2^n). \quad (21)$$

The two signal draws differ uniformly in A ; any term containing a uniform draw is uniform on the full space. Replace each difference by an independently uniform bit string with probability one half. The law becomes (17) with $w = c^2/2$ and the known $w_0 = c_0^2/2$. This handles unknown affine phases and $c = 1$.

In a frame satisfying (3), $\dim A \leq (1 - \beta)n$. Lemma 5.2 recovers A with constant probability. Compute a basis of $D = A^\perp$, conjugate its complete Z operators back to physical coordinates,

and validate their means and signs on fresh original copies. A correct recovery gives at least βn independent target constraints. A failed recovery or a frame without the promise has no discovery guarantee, but any accepted constraint is still correct outside the allocated validation error.

For large enough n that $\beta n \geq K$, repeat a bounded $O(\log(1/\delta))$ number of independent runs per supplied frame and accept the first fully certified independent seed of size K . Its absence probability is at most an allocated fraction of δ . All candidate slots and validation errors are union-bounded as in Lemma 3.2. The seed filter gives signal at least 9/10; pay its original-copy cap, apply Proposition 2.2, and lift. For fixed c_0, β all costs are polynomial in the stated variables. The finitely many n not meeting (18), (19), or $\beta n \geq K$ use exhaustive Pauli estimation as in the proof of Theorem 1.1. Their dimensions are bounded by a constant depending on c_0, β , so this preserves the fixed-parameter polynomial statement. All computations are capped off-promise, and correctness is only claimed under the supplied promise. $\square$

6 Supporting resource bounds

6.1 A standard copy lower bound

For comparison, let the target be uniform among the $M_n = 2^n \prod_{j=1}^n (2^j + 1)$ signed stabilizer states, with fixed actual signal c . Their average noisy state is $I/2^n$. Entropy concavity gives $S(\rho_s) \geq (1 - c)n$, so the Holevo information per copy is at most cn . Subadditivity bounds the information in T copies by Tcn , even allowing collective measurements. Fano's inequality yields

$$T \geq \frac{(1 - \delta) \log M_n - h_2(\delta)}{cn} = \Omega(n/c) \quad (22)$$

at fixed error. This uses standard information theory. In the weak-signal regime its scale is $n^{1+o(1)}$.

6.2 Partial code effects versus full-basis readout

The next bound explains one restricted architectural obstruction. It does not apply to unrestricted single-copy learning.

Proposition 6.1 (Adaptive code-effect bound). *Choose a uniform signed target stabilizer on n qubits. Suppose an adaptive protocol uses at most T fresh copies and that every fully resolved effect on a copy is a nonnegative scalar times a signed stabilizer-code projector of codimension at most m . There is no quantum memory between copies; arbitrary classical adaptation is allowed. The probability of outputting any genuine nonidentity unsigned target Pauli is at most*

$$\frac{1 + Tc(2^m - 1)}{2^n + 1}. \quad (23)$$

More generally replace $T(2^m - 1)$ by $\mathbb{E}_0 \sum_t (2^{m_t} - 1)$, where m_t is the realized effect's codimension under maximally mixed inputs.

Proof. Let S be the target Lagrangian and let an effect be $E_y = \alpha_y \Pi(M_y, \chi_y)$, with $\dim M_y = m_y$. Put $q_y = \text{Tr}(E_y)/2^n$. Expanding commuting Pauli projectors shows

$$\frac{\langle s|E_y|s \rangle}{q_y} = \begin{cases} 2^j, & \text{consistent eigenvalues on } S \cap M_y, \\ 0, & \text{otherwise,} \end{cases} \quad j = \dim(S \cap M_y).$$

In particular the ratio equals one whenever $j = 0$.

Introduce independent latent Bernoulli(c) tags in the true experiment and in a null experiment. A true tagged input is pure when the tag is one and maximally mixed otherwise; every null input is maximally mixed. The tags are not given to the protocol. Under the null, the full transcript and tags are independent of S . On every joint target/tag/transcript path for which each tag is zero or $S \cap M_y = \{0\}$, true and null path probabilities agree pointwise. Hence their bad-path masses are equal, and their total variation is at most the null bad-path probability.

A fixed nonzero label belongs to a uniform Lagrangian with probability $1/(2^n + 1)$. Thus $\mathbb{P}(S \cap M_y \neq \{0\}) \leq (2^{m_y} - 1)/(2^n + 1)$. Union-bound along null paths, including the factor c for the tags. Under the null an output nonzero label is independent of S , and is genuine with probability at most $1/(2^n + 1)$. Total variation transfers this event bound to the true experiment and proves the claim. Capped stopping is handled by padding with identity effects. $\square$

The code-effect property holds for resolved Clifford/Pauli instruments with pure stabilizer ancillas: apply a branch to half an EPR state and trace out its output. The reduced reference of that postselected stabilizer state is proportional to a code projector, and equals the transpose input effect up to normalization. Arbitrary coarse-graining of outcomes need not preserve this property; the codimension restriction must concern the resolved effects, not only a coarse effect's rank. Full rank-one basis readout has codimension n , so (23) is weak there. The constructive star measurement uses that full readout.

For the same projector, a compatible target intersection of dimension j gives acceptance $2^{-m}(1 - c + c2^j)$ and conditional signal $c2^j/(1 - c + c2^j)$. When $j = 0$, no signal is amplified. Artificial postmeasurement constraints therefore cannot simply be treated as original target generators.

6.3 Fixed menus do not guarantee a large seed

For a fixed measurement Lagrangian Z and uniform target S , choosing a k -space in Z and then a transverse quotient Lagrangian gives

$$\mathbb{P}(\dim(S \cap Z) = k) = \frac{\begin{bmatrix} n \\ k \end{bmatrix}_2 2^{(n-k)(n-k+1)/2}}{\prod_{j=1}^n (2^j + 1)}.$$

Here $\begin{bmatrix} n \\ k \end{bmatrix}_2$ is a Gaussian binomial coefficient. More simply, union-bound over k -spaces contained in Z to obtain

$$\mathbb{P}(\dim(S \cap Z) \geq k) \leq C_2^{-1} 2^{-k(k+1)/2}, \quad C_2 = \prod_{j \geq 1} (1 - 2^{-j}) > 0.$$

Indeed $\begin{bmatrix} n \\ k \end{bmatrix}_2 \leq C_2^{-1} 2^{k(n-k)}$ and a fixed k -space is in S with probability at most $2^{-kn+k(k-1)/2}$. A target-independent menu covering all targets with a k -dimensional intersection therefore needs at least $C_2 2^{k(k+1)/2}$ frames. This standard counting consequence prevents replacing the linear-gap promise by a polynomial fixed menu. It does not restrict data-dependent adaptive frames or the general one-row construction.

7 Scope, attribution and reproducibility

The main guarantee separates original copies, quantum operations and classical time. The exact global mixture and ideal gates are essential to the stated analysis. The classical search remains subexponential. The star-law representation is an *upper* reduction to ordinary classical LPN.

The prior phase-state derivatives and Bell/LPN construction [1], random-frame count [8], uniform Clifford generation [7], parity solver [2], high-fidelity learner [3], and finite-field polynomial

lift [4] are explicit dependencies. The quantitative advance is short certified filtering that avoids learning every noisy graph row, together with a polynomial-time visible-subspace regime. The supporting information, incidence and coupling bounds are applications of standard methods. Existing noisy structure learning by Bell sampling [11] uses a different inter-copy resource. The earlier noise-transition paper [9] is the predecessor motivating this change of algorithm.

The public verifier supplies exact finite controls of the star law, complex phases, parity masks, noise-grid total variation, conditional signal identity, finite symplectic counts and the polynomial-rank ingredient. These are mathematical identity checks, not an execution of Lyubashevsky’s solver or the entire quantum learner.

AI assistance

AI tools assisted with derivation, programming, literature discovery and proof checking. Analytic proofs and finite computations have distinct roles.

References

- [1] S. Arunachalam, S. Bravyi, A. Dutt and T. J. Yoder, *Optimal algorithms for learning quantum phase states*, TQC 2023, LIPIcs **266**, 3:1–3:24. doi:10.4230/LIPIcs.TQC.2023.3. Full version: arXiv:2208.07851v2.
- [2] V. Lyubashevsky, *The parity problem in the presence of noise, decoding random linear codes, and the subset sum problem*, APPROX/RANDOM 2005, LNCS **3624**, 378–389. doi:10.1007/11538462_32. Author manuscript.
- [3] 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).
- [4] A. Chen, A. De and A. Vijayaraghavan, *Learning a mixture of two subspaces over finite fields*, ALT 2021, PMLR **132**, 481–504. Proceedings paper. Full version: arXiv:2010.02841v2.
- [5] I. Ben-Eliezer, R. Hod and S. Lovett, *Random low-degree polynomials are hard to approximate*, Computational Complexity **21** (2012), 63–81.
- [6] P. Keevash and B. Sudakov, *Set systems with restricted cross-intersections and the minimum rank of inclusion matrices*, SIAM Journal on Discrete Mathematics **18** (2005), 713–727.
- [7] R. Koenig and J. A. Smolin, *How to efficiently select an arbitrary Clifford group element*, Journal of Mathematical Physics **55** (2014), 122202. doi:10.1063/1.4903507. arXiv:1406.2170.
- [8] S. Arunachalam and L. Schatzki, *Optimal stabilizer testing and learning with limited quantum memory*, arXiv:2607.02444v2 (2026).
- [9] S. Schubr, *Sharp noise transitions in adaptive single-copy stabilizer learning*, Luminary Research Archive, author-published edition 1 (2026). Public edition.
- [10] S. Schubr, *Mixed-state escape bounds and proper agnostic learning of stabilizer states in odd-prime dimensions*, Luminary Research Archive, author-published edition 2 (2026). Public edition.
- [11] L. Tang and T. Steckmann, *Learning noise-robust stabilizer structure via Bell sampling*, arXiv:2610.03307v2 (2026).