Mixed-state escape bounds and proper agnostic learning of stabilizer states in odd-prime dimensions

Scott Schubr

Updated edition — October 19, 2026

Abstract

We strengthen proper agnostic learning of stabilizer states for arbitrary mixed inputs in every odd-prime local dimension. Given an overlap promise $F \geq \tau$, the learner returns a pure stabilizer of overlap at least $F - \varepsilon$, with confidence $1 - \delta$, using $\exp(O(\log^2(2/\tau)))\,\mathrm{poly}(n, p, \varepsilon^{-1}, \log(2/\delta))$ resources with absolute exponent constants and fixed polynomial degree. The previous dimension-dependent exponent is removed by testing Weyl spectral concentration and selecting an observed eigenvalue. An even quotient-convolution argument also improves escape to $(p-1)F^r/(4p)$, where $r = 6$ for $p \equiv 3 \pmod 4$ and $r = 4$ for $p \equiv 1 \pmod 4$, using the same known measurements. Explicit rational counters cover all rejection attempts and final validation. The model still uses exact Clifford operations and coherent eigenspace measurements. We attribute the measurement primitives, uncertainty ingredients and bootstrapping framework to prior work; the strengthening concerns their uniform quantitative combination.

1 Introduction and results

An agnostic stabilizer learner seeks a classical description of a pure stabilizer whose overlap with an unknown input is nearly the largest possible. The input need not itself be a stabilizer or follow a specified noise model. Chen, Gong, Ye and Zhang [3, Corollary 6.3] established a proper mixed-qubit learner using stabilizer bootstrapping. Qi, Xu, Feng and Li [4] developed agnostic learning for pure inputs in odd-prime local dimension; their Section VI identifies extension to mixed inputs as a further direction.

The earlier author-published edition of this work [1] supplied a mixed-input escape estimate and a complete learner with a dimension-dependent exponent. This revision strengthens the resource guarantee to an absolute quasipolynomial exponent, leaving local dimension only in a fixed polynomial. The purity-sensitive mixed local inequality is retained. The new steps are an even quotient-convolution estimate, a constant-threshold test of the largest Weyl spectral outcome probability, and selection of an actual measured eigenvalue.

The learning architecture is stabilizer bootstrapping of Chen et al. and its qudit adaptation by Qi et al. The weighted/skewed measurements and their residue-class arithmetic are already present in Allcock, Doriguello, Ivanyos and Santha [5]; related tensor-Weyl measurements occur in Nöller, Tran, Gachechiladze and Kueng [7]. Hinsche, Eisert and Carrasco [6] explicitly anticipate a qudit agnostic extension using their measurement tools.

Related tasks require care in comparison. Holt and Subramanian [9, Theorem 9.27] learn the exact stabilizer group of a mixed qudit state, rather than its best pure stabilizer comparator. Bittel et al. [10, Theorems III.1 and A.1] give single-copy algorithms for mixed-qubit inputs with a known stabilizer-overlap floor greater than $\sqrt{2/3}$ and for exact pure stabilizer learning in prime dimension, respectively.

Recent algorithmic-PFR progress gives polynomial-time improper agnostic tomography for pure qubit inputs [2, Section 8.2.2]. The output may be a sum of stabilizers.

For a density operator on $(\mathbb{C}^p)^{\otimes n}$ define

$$F = F_{\mathrm{stab}}(\rho) = \max_{s \in \mathcal{S}_{p,n}} \langle s|\rho|s\rangle,$$

where $\mathcal{S}_{p,n}$ denotes pure stabilizer states up to global phase. This is squared fidelity with a pure comparator. All logarithms are natural.

Theorem 1.1 (Uniform mixed-input escape)**.** *Let $p \geq 3$ be prime, $n \geq 1$, and ρ any density operator. For the physically sampleable law B_ρ defined in Section 4, every maximizing stabilizer Lagrangian S and proper linear $T < S$ satisfy*

$$B_\rho(S \setminus T) \geq c_p F^{r_p} \geq \frac{1}{6} F^{r_p}, \qquad c_p = \frac{p-1}{4p}, \quad r_p = \begin{cases} 6, & p \equiv 3 \pmod 4, \\ 4, & p \equiv 1 \pmod 4. \end{cases} \tag{1}$$

A draw uses r_p independent input copies and a known stabilizer measurement with polynomial overhead in n, p. These are sufficient, not minimal, copy counts.

Theorem 1.2 (Dimension-uniform proper learner)**.** *Let $p \geq 3$ be prime, $n \geq 1$, $0 < \tau \leq 1$ and $0 < \varepsilon, \delta < 1$. Given independent copies of any density operator ρ with $F_{\mathrm{stab}}(\rho) \geq \tau$, an algorithm returns a pure stabilizer tableau s with*

$$\langle s|\rho|s\rangle \geq F_{\mathrm{stab}}(\rho) - \varepsilon$$

with probability at least $1 - \delta$. In the model below, put $L = \log(2/\tau)$. Copies, gates and classical time are bounded by

$$\exp\big(O(L^2)\big)\, \mathrm{poly}(n, p, \varepsilon^{-1}, \log(2/\delta)), \tag{2}$$

with absolute exponential constants and fixed polynomial degree independent of p. The original-copy budget alone has polynomial dependence on $\log p$ in place of p. At fixed promised τ, the bound is polynomial in inverse additive accuracy. The no-promise corollary below is quasipolynomial in inverse accuracy.

Operation model and parameters. We retain independent input copies, exact prime-dimensional Clifford operations (including Fourier and SUM), coherent Weyl eigenspace measurements, and finite-field arithmetic, with fixed polynomial overhead in n, p. No purification, conjugated input, matrix-entry access, optimizer oracle or free postselection is assumed.

All parameters may be rational inputs. To obtain canonical short encodings, round each requested $x \in \{\tau, \varepsilon, \delta\}$ down to $2^{-\lceil \log_2(1/x) \rceil}$. This changes it by at most a factor two, preserves all requested guarantees, and keeps $\tau = 1$ unchanged. Reading arbitrary supplied encodings adds their input-processing cost. In the explicit counters we use the rounded parameters, denoted by the same symbols.

Relationship to the published edition. This is a strengthened revision of the same work [1], including the private qutrit predecessor at $p = 3$. The earlier escape constants were $3(p-1)/(4p^3)$ and $(p-1)/(2p^2)$, and its learner bound was $\exp(O(p^2 L(L + \log p)))$ times a polynomial. The new bound changes the growing-dimension guarantee; for each fixed p the quasipolynomial inverse-overlap class was already known in that edition.

2 Weyl conventions and a positive overlap law

Set $D = p^n$, $V = \mathbb{F}_p^{2n}$ and $\omega = e^{2\pi i/p}$. For $u = (a, b)$ use

$$W_u = \omega^{-a\cdot b/2} Z^a X^b, \qquad [u, v] = a \cdot b' - b \cdot a', \qquad W_u W_v = \omega^{[u,v]/2} W_{u+v}.$$

Here $X|j\rangle = |j + 1\rangle$ and $Z|j\rangle = \omega^j |j\rangle$; all exponents are in $\mathbb{F}_p$. Label subspaces are linear. A pure stabilizer is specified by a maximal isotropic subspace, or Lagrangian, together with compatible eigenvalues. We use the usual odd-prime stabilizer closure and Clifford synthesis facts [12, 4].

Write $\chi_\rho(u) = \mathrm{Tr}(\rho W_u)$ and define

$$A_\rho(x) = D^{-1} \mathrm{Tr}(\rho W_x \rho W_x^\dagger), \qquad \widehat{A}_\rho(u) = \sum_x A_\rho(x)\omega^{[u,x]}. \tag{3}$$

Positivity of the trace product gives $A_\rho \geq 0$. Weyl twirling, $\sum_x W_x \rho W_x^\dagger = DI$, gives $\sum_x A_\rho(x) = 1$. Weyl orthogonality and conjugation give

$$\widehat{A}_\rho(u) = |\chi_\rho(u)|^2. \tag{4}$$

For example, expanding $\rho = D^{-1} \sum_v \chi_\rho(-v) W_v$ in (3) yields $A_\rho(x) = D^{-2} \sum_v |\chi_\rho(v)|^2 \omega^{[x,v]}$, which is the inversion of (4). The auxiliary law is normalized for mixed inputs. In contrast, the characteristic weight $p_\rho(u) = D^{-1}|\chi_\rho(u)|^2$ has total mass $\mathrm{Tr}(\rho^2)$. They agree for pure inputs, but must not be interchanged for mixed inputs. We do not assume a physical two-copy sampler for A_ρ. Positive Fourier-dual overlap distributions for mixed qubits already appear in Iyer and Liang [8, Lemma 2.12, Remark 2.13 and Facts 2.14–2.15]. Equation (4) is the standard Weyl-overlap identity in our odd-prime conventions. General Fourier-difference sampling is also treated by Holt and Subramanian [9, Lemma 8.4].

3 A mixed local inequality and the all-system lift

Lemma 3.1 (Local mixed estimate)**.** *Let τ be a density matrix on $\mathbb{C}^p$ such that $|0\rangle$ maximizes overlap with single-system stabilizers, and put $f = \tau_{00}$. Then*

$$\sum_j \tau_{jj}^2 - \frac{1}{p} \mathrm{Tr}(\tau^2) \geq c_p f^2, \qquad c_p = \frac{p-1}{4p}. \tag{5}$$

Proof. The $p + 1$ stabilizer mutually unbiased bases correspond to the lines in $\mathbb{F}_p^2$ [11, 4]. If their outcome probabilities are z_v, Weyl orthogonality gives

$$\sum_v z_v = p + 1, \qquad \sum_v z_v^2 = 1 + P, \qquad P = \mathrm{Tr}(\tau^2).$$

Since every $z_v \leq f$, we have $f \geq 1/p$ and $P \leq \min(1, (p + 1)f - 1)$. Cauchy–Schwarz on the other $p-1$ computational outcomes gives $\sum_j \tau_{jj}^2 \geq f^2 + (1-f)^2/(p-1)$. Let g denote the left side of (5) and put

$$h_0 = \frac{2}{p+1}, \qquad K_0 = \frac{(3p-1)(p+1)}{4p(p-1)}.$$

Subtraction and factorization yield

$$g - c_p f^2 \geq K_0(f - h_0)\left(f - \frac{2(2p-1)}{3p-1}\right) \geq 0, \qquad 1/p \leq f \leq h_0,$$

$$g - c_p f^2 \geq K_0(f - h_0)\left(f - \frac{2}{3p-1}\right) \geq 0, \qquad h_0 \leq f \leq 1.$$

In the first line use $P \le (p+1)f - 1$; its other root exceeds one. In the second use $P \le 1$; its other root is at most h_0. This includes both endpoints. □

Lemma 3.2 (Hyperplane masses)**.** *Let S belong to a global optimizing stabilizer of overlap F. For every hyperplane $H < S$,*

$$A_\rho(H) \ge F^2/p, \qquad A_\rho(S \setminus H) \ge c_p F^2. \tag{6}$$

Proof. Clifford covariance and a computational linear change put the optimizer at $|0^n\rangle$, with $S = \mathrm{span}(Z_1, \dots, Z_n)$ and $H = \mathrm{span}(Z_2, \dots, Z_n)$. For $y \in \mathbb{F}_p^{n-1}$ let σ_y be the positive principal block with entries $(\sigma_y)_{jk} = \rho_{(j,y),(k,y)}$. Character sums give

$$A_\rho(H) = p^{-1} \sum_y \mathrm{Tr}(\sigma_y^2), \tag{7}$$

$$A_\rho(S \setminus H) = \sum_y \left(\sum_j (\sigma_y)_{jj}^2 - p^{-1} \mathrm{Tr}(\sigma_y^2) \right). \tag{8}$$

Every summand in (8) is nonnegative: positivity implies $\mathrm{Tr}(\sigma^2) \le (\mathrm{Tr}\,\sigma)^2 \le p \sum_j \sigma_{jj}^2$. The target block has $(\sigma_0)_{00} = F$, giving the first bound. Set $t = \mathrm{Tr}\,\sigma_0 > 0$. The normalized block $\tau = \sigma_0/t$ has $f = F/t$ and $|0\rangle$ is its stabilizer optimizer: otherwise $|v\rangle \otimes |0^{n-1}\rangle$ would improve the original global optimum. Lemma 3.1 makes the target summand at least $t^2 c_p (F/t)^2 = c_p F^2$. For $n = 1$ there is a single block and $H = \{0\}$, so the same proof applies. □

4 Known weighted measurements and mixed-input escape

We describe the measurement explicitly to fix signs and mixed-input normalization. It is an instance of the known weighted/skewed stabilizer measurements [5, 7]. Choose nonzero $t_1, \dots, t_k \in \mathbb{F}_p$ with $\sum_i t_i^2 = 0$ and put

$$R_u = \bigotimes_{i=1}^{k} W_{t_i u}, \qquad P_y = D^{-2} \sum_{u \in V} \omega^{[y,u]} R_u.$$

The cancellation makes $R_u R_v = R_{u+v}$ and $R_u^\dagger = R_{-u}$. Character orthogonality shows that the P_y are orthogonal Hermitian projectors summing to I, with rank D^{k-2}. Indeed, all nonzero R_u have trace zero. If Y is the outcome on $\rho^{\otimes k}$, then

$$\mathbb{E}\omega^{[u,Y]} = \prod_i \chi_\rho(t_i u).$$

For independent outcomes Y, Y', define B_ρ to be the law of $Y - Y'$. Consequently

$$\widehat{B}_\rho(u) = \prod_i |\chi_\rho(t_i u)|^2, \qquad B_\rho = A_{t_1} * \cdots * A_{t_k}, \quad A_t(x) = A_\rho(t^{-1}x). \tag{9}$$

The scalars are nonzero, so every A_t is normalized. For each physical position the two generating Paulis on the k registers commute; standard symplectic Clifford synthesis implements their joint measurement. The construction factors across the n positions. It retains the joint eigenspaces and does not separately measure the tensor factors.

Proof of Theorem 1.1. If $p \equiv 3 \pmod 4$, choose a, b with $a^2 + b^2 = -1$: the squares and minus one minus the squares each have size $(p+1)/2$, so intersect. Both coefficients are nonzero because -1 is nonsquare. Use $k = 3$ and weights $(1, a, b)$. If $p \equiv 1 \pmod 4$, choose $a^2 = -1$, use $k = 2$ and weights $(1, a)$. These are the known arithmetic specializations [5].

Fix a hyperplane $H < S$. The auxiliary law is even: cyclicity of trace gives $A_\rho(-x) = A_\rho(x)$. In the optimizing stabilizer basis, Weyl phase sums give

$$M = A_\rho(S) = \sum_z \rho_{zz}^2 \geq F^2, \qquad d = A_\rho(S \setminus H) \geq c_p F^2$$

by Lemma 3.2. Restrict A_ρ to S, normalize by $M > 0$ and push to $S/H = \mathbb{F}_p$. The resulting probability law μ is even and $\mu(0) = 1 - d/M$. Nonzero quotient atoms occur in distinct pairs $u, -u$, hence

$$\|\mu\|_\infty \leq \max\big(\mu(0), (1-\mu(0))/2\big), \qquad 1 - \|\mu\|_\infty \geq \min(d/M, 1/2).$$

Each scaled factor A_t in (9) has mass M on S and a quotient conditional law that is a scalar permutation of μ. For probability laws ν, λ, the elementary contraction $\|\nu * \lambda\|_\infty \leq \|\nu\|_\infty$ therefore bounds the zero atom of the quotient sum. The event that all k factors lie in S and sum outside H has mass at least

$$M^k \min(d/M, 1/2) = \min(M^{k-1} d, M^k/2) \geq c_p F^{2k}.$$

Other factor configurations contribute nonnegative mass. Enlarge any proper $T < S$ to a hyperplane H and use $S \setminus H \subset S \setminus T$. Since $r_p = 2k$ and $c_p \geq 1/6$, the theorem follows. The probability argument is an application of standard symmetry and convolution contraction. □

Remark 4.1. For pure input $A_\rho = p_\rho$, so the convolution agrees with the pure-input marginal used in Qi et al. [4, Eq. (27)], after its skew change of variables. Zero-weight registers are omitted, never inverted. In particular the two-register branch handles $p = 5$, where $a^2 + b^2 = -1$ has no solution with both a, b nonzero. The pure-input algorithm of Qi et al. uses its own measurement primitives.

5 Coherent conditioning and correlation thresholds

Lemma 5.1 (Preservation of a global optimizer)**.** *Let s maximize stabilizer overlap F and let P be a consistent stabilizer-code projector fixing s. Then $w = \mathrm{Tr}(P\rho) \geq F$, and $\rho_P = P\rho P/w$ has stabilizer optimum exactly F/w, still attained by s.*

Proof. Stabilizer measurement closure [12] gives, for any pure stabilizer t, either $P|t\rangle = 0$ or $P|t\rangle = \sqrt{c}\, e^{i\theta}|t'\rangle$ for another stabilizer t' and $0 < c \leq 1$. Thus $\langle t|P\rho P|t\rangle = c\langle t'|\rho|t'\rangle \leq F$, whereas s attains F. Since P contains $|s\rangle\langle s|$ as a subprojection, $w \geq F$. Division by w proves the claim, including tied optimizers. □

For nonzero x let $E_j(x)$, $j \in \mathbb{F}_p$, be the spectral projectors of W_x, and write

$$z_j(x) = \mathrm{Tr}(\sigma E_j(x)), \qquad M_\sigma(x) = \max_j z_j(x).$$

For $x = 0$ use its deterministic outcome, so $M_\sigma(0) = 1$.

Lemma 5.2 (Spectral concentration and constant amplification)**.** *Labels with $M_\sigma(x) > 33/40$ commute pairwise. If σ is supported on a code with span C, each such label commutes with C. Put $q = 9/10$ and $G_\sigma = \{x : M_\sigma(x) > q\}$. Outside G_σ, every spectral outcome has probability at most q.*

Proof. For spectral projectors E, G of noncommuting Weyls, character orthogonality gives $EGE = E/p$ and $GEG = G/p$. Thus $\|E+G\| \le 1+p^{-1/2} \le 1+1/\sqrt{3} < 33/20$. If both labels have a spectral outcome of probability greater than $33/40$, their probabilities sum to more than $33/20$, a contradiction. If x does not commute with the code, conjugation by a suitable definite-eigenvalue code Weyl permutes the p spectral outcomes transitively. Primality makes every outcome probability $1/p$, precluding a high label. The last assertion is the definition of M_σ. □

On a code prefix fixing a global optimizer s, Lemma 5.1 retains its optimality. For a target label $x \in S$, the target eigenspace contains $|s\rangle\langle s|$, so

$$z_{\text{target}}(x) \ge \langle s|\sigma|s\rangle = F_{\text{stab}}(\sigma) \ge \tau. \tag{10}$$

Measure W_x on one fresh conditional copy and choose its observed eigenvalue. The target-correct event therefore has probability at least τ, without an optimizer oracle. If also $x \notin G_\sigma$, that event raises target overlap from $F_{\text{stab}}(\sigma)$ to $F_{\text{stab}}(\sigma)/z_{\text{target}} \ge F_{\text{stab}}(\sigma)/q$. This yields a constant amplification factor $10/9$ independent of p.

6 Robust capture of high spectral masses

For rational $x \ge 1$ put $\ell(x) = \lceil \log_2 x \rceil$, computed by integer comparisons. Since $\ell(x) \ge \log x$, it provides conservative counters without transcendental comparisons. Write $r = r_p$ and set

$$a = \tau^r/12, \qquad m = \left\lceil \frac{4n^2\ell(p)+3}{a} \right\rceil, \qquad h = 800\ell(16pm). \tag{11}$$

Take m independent B_σ samples. For each sampled label, measure its Weyl operator on h fresh independent conditional copies, and estimate M_σ by the largest empirical outcome frequency. Retain empirical maxima greater than $17/20$. Include the current code span C and extend the retained span to a full Lagrangian H; abort on observed nonisotropy or actual preparation failure. The hidden mass guarantee below is not tested by the algorithm.

Lemma 6.1 (Capture)**.** *With uncapped conditional-copy access this gives, with probability at least $3/4$, a full Lagrangian $H \supset C$ satisfying*

$$B_\sigma(G_\sigma \setminus H) \le a. \tag{12}$$

Only integer frequency counts and rational threshold comparisons are needed.

Proof. Hoeffding and a union over the p outcome probabilities bound estimation error greater than $1/40$ by $2pe^{-h/800}$ per label. Conditional on the sampled labels, a union over all m labels fails with probability at most $1/8$, since $h \ge 800\log(16pm)$. On this event every sampled G_σ label is retained, with empirical mass greater than $7/8$; all retained labels have true mass greater than $33/40$. Lemma 5.2 makes them and C isotropic.

There are at most p^{4n^2} subspaces of V, overcounting their column-span representations. For any fixed H violating (12), the probability that all samples miss $G_\sigma \setminus H$ is at most e^{-am}. Every sampled high label is retained, so the actual output H contains those labels on the estimation event. The union bound $p^{4n^2}e^{-am} \le 1/8$ handles every possible data-dependent output simultaneously. The two failures sum to at most $1/4$; independence between them is unnecessary. □

7 A capped learner with uniform resources

Set the following counters in order:

$$q = 9/10, \qquad J = \min\{j \in \mathbb{Z}_{\geq 1} : \tau q^{-j} > 1\}, \tag{13}$$

$$b = 3a\tau/4 = \tau^{r+1}/16, \qquad s_0 = b^J,$$

$$Q = J(rm + mh + r + 1), \qquad K = \left\lceil \frac{\ell(4Q/s_0)}{\tau} \right\rceil, \tag{14}$$

$$R = \left\lceil \frac{2\ell(2/\delta)}{s_0} \right\rceil, \qquad N = R(J+1),$$

$$v = \left\lceil \frac{2\ell(4N/\delta)}{\varepsilon^2} \right\rceil. \tag{15}$$

The strict comparison defining J is rational; $\tau = 1$ gives $J = 1$, while $\tau = 9/10$ gives $J = 2$. All loops have deterministic upper bounds.

One run. Initialize the code projector $P = I$ and its unsigned span $C = \{0\}$. For at most J stages:

1. Prepare independent $\sigma = \rho_P$ copies as requested using the cap below, and construct H by Section 6. Abort on preparation failure or observed nonisotropy.
2. Measure one fresh *original* ρ copy in H's stabilizer basis and record the resulting pure stabilizer candidate.
3. Draw an independent $x \sim B_\sigma$. Abort if $x \in H$ or x fails to commute with C. Otherwise measure W_x on *one additional fresh conditional copy*, observe its eigenvalue, and extend the code by that eigenvalue constraint.
4. If code rank reaches n, record its unique pure stabilizer and stop. Otherwise continue. Preserve all recorded candidates on every abort.

Only observable code/label tests and measurements are implemented; membership in the unknown target space or a hidden favorable event is never queried. Because $C \subset H$, an accepted outside-H label gives an independent commuting constraint.

Capped conditional preparation and copy ledger. For each ρ_P request, try at most K original copies, coherently measure the joint code eigenvalues, and keep only the desired outcome. Clifford conjugation to coordinate measurements and undoing realizes P while preserving within-code coherence. Separate destructive measurements of physical tensor factors would not suffice. Abort if all K trials fail. Wrong prefixes have the same deterministic cap, with no lower bound on their acceptance weights.

Each stage requests at most rm conditional copies for capture draws, mh for histograms, r for the progress draw, and *one* for eigenvalue selection. This proves the Q formula, including its additional $+1$. Separately there is at most one original basis sample per stage. Each run lists at most $J + 1$ candidates, allowing a terminal code state. After R independent runs, estimate each candidate's original-state overlap by its rank-one projector on v fresh original copies and return a largest estimate; on an empty list return a fixed stabilizer.

Proof of Theorem 1.2. Fix an original global optimizer s and its label Lagrangian S. Analyze an auxiliary witness using uncapped conditional copies only while active. Absorb as success on first recording s; otherwise continue only after a target-fixing low-spectral-mass jump with the target-correct eigenvalue. Kill the witness at the first other continuation or abort. No normalized conditional state is needed after death or absorption. On active prefixes, P fixes s, its weight is at least $F \geq \tau$, and s remains globally optimal.

Capture is good with probability at least $3/4$. If a good $H = S$, the original basis sample records s with probability $F \geq \tau$. If a good $H \neq S$, Theorem 1.1 applied to $S \cap H$ yields $B_\sigma(S \setminus H) \geq F_{\text{stab}}(\sigma)^r/6 \geq 2a$. Removing $G_\sigma \setminus H$ costs at most a, leaving low target-label mass at least a. These labels commute with $C \subset S$. Equation (10) gives target-correct observed eigenvalue probability at least τ, conditionally on the prefix, H and label, because that copy is fresh. If the preceding basis sample already recorded s the witness is absorbed; otherwise the same jump event is favorable continuation. Partitioning over good H gives probability of success or favorable continuation at least

$$\frac{3}{4}\min(\tau, a\tau) = \frac{3}{4}a\tau = b.$$

Each favorable nonabsorbing jump raises overlap by at least q^{-1}. The strict J definition makes J such jumps impossible since $\tau q^{-J} > 1$. Favorable rank-n termination already records s. Induction with absorption therefore bounds ideal witness success below by $s_0 = b^J$.

On active favorable prefixes every conditional request accepts an original copy with probability at least τ. Its cap-failure probability is at most $e^{-\tau K}$. Accepted states have exactly the ideal conditional law. Sequential coupling, stopped when the witness dies or records s, loses at most

$$Qe^{-\tau K} \leq s_0/4.$$

Only active favorable requests need this lower weight bound. All other paths remain cost-capped. A later abort cannot erase a recorded candidate, so an actual run records s with probability at least $3s_0/4 \geq s_0/2$.

Independent runs miss s with probability at most $e^{-Rs_0/2} \leq \delta/2$. Conditional on the completed list of at most N candidates, fresh original-state Bernoulli validation has simultaneous error at most $\varepsilon/2$, except with probability at most $2Ne^{-v\varepsilon^2/2} \leq \delta/2$. Selecting the largest estimate loses at most ε from a listed optimizer. This proves the confidence and proper-output guarantees, including $n = 1$, tied optimizers and $\tau = 1$.

The deterministic original-copy budget is

$$RQK + RJ + Nv. \tag{16}$$

Put $L = \log(2/\tau)$. Since $r \leq 6$ and $q = 9/10$ are absolute,

$$\begin{aligned} J &= O(L), & \log(1/s_0) &= O(L^2), \\ m &= O(\tau^{-r}(n^2 \log p + 1)), & h &= O(\log p + \log m + 1), \\ Q &= O(Jm(r + h + 1)), & K &= O(\tau^{-1}(\log Q + L^2 + 1)), \\ R &= \exp(O(L^2))O(\log(2/\delta)), & v &= O(\varepsilon^{-2}(L^2 + \log(2/\delta) + 1)). \end{aligned}$$

Every constant is independent of p. Powers of L and τ^{-1} are absorbed by $\exp(O(L^2))$. Equation (16) therefore has a fixed polynomial in $n, \log p, \varepsilon^{-1}, \log(2/\delta)$ outside that exponential. Fixed algorithms for finite-field elimination, tableaux, isotropic completion and exact Clifford measurements add polynomial overhead in n, p and the counters. Brute-force choice of sampler weights uses at most $O(p^2)$ field operations. Histogram storage may use p bins. This establishes (2).

After canonical power-of-two rounding, b^J has $O(L^2)$ representation bits and integer counters have polynomial bit lengths in $L, \log n, \log p, \log(1/\varepsilon), \log(1/\delta)$. Integer comparisons compute ℓ, J and the retention test $20\,\mathrm{max_count} > 17h$. Arbitrary long input encodings incur their stated additional processing cost. No exact comparison of transcendental numbers is used. □

Corollary 7.1 (Unknown optimum)**.** *For arbitrary mixed input and every odd prime p, the no-promise proper additive-ε learner uses*

$$\exp(O(\log^2(2/\varepsilon)))\,\mathrm{poly}(n, p, \log(2/\delta)),$$

with absolute exponent constants and fixed polynomial degree. Its inverse-accuracy dependence is quasipolynomial.

Proof. Choose $\tau = \varepsilon$, or a canonical lower approximation. If $F < \tau \le \varepsilon$, every stabilizer has overlap at least $0 > F - \varepsilon$; otherwise Theorem 1.2 applies. Costs remain bounded when the promise fails, since every run and request is capped. Polynomial inverse-accuracy factors are absorbed into the exponential. □

8 Scope, attribution and verification

The mixed local estimate of the earlier edition is retained. The strengthened escape uses elementary evenness and convolution contraction. Spectral-projector uncertainty, empirical frequencies and outcome sampling are standard tools. The contribution is their quantitative fit to the known learning framework: all local-dimension dependence moves into a fixed polynomial, with an absolute quasipolynomial exponent in inverse promised overlap.

The improvement matters when p grows. At fixed promised overlap the new upper bound is polynomial in n, p and inverse additive accuracy; for each fixed p the earlier guarantee was already quasipolynomial in inverse overlap. These are improvements of upper bounds, not lower bounds proving previous algorithms inherently slow. No practical implementation advantage has been demonstrated.

The output is a pure stabilizer comparator, not a reconstruction of ρ or all Weyl amplitudes. Exact mixed-state symmetry learning [9], high-fidelity/exact-pure single-copy alternatives [10], and pure-qubit improper learning [2] have different guarantees. Multi-copy amplitude-estimation bounds [7] address another task. No composite dimension, approximate-local-optimizer, finite-universal-gate robustness or uncontrolled-noise theorem is included.

The proof is analytic. Internal proof reviews are private preparation records. The public reproducibility supplement contains targeted exact-rational/integer and floating-point diagnostics, including negative controls. The revised verification entry points enforce the expected proof identity and independently check the physical-copy request ledger; these engineering gates protect source/evidence bindings, rather than turn numerical tests into proof-assistant verification.

Development disclosure. This revision was developed with AI assistance in derivation, programming, literature discovery and proof checking. Analytic proof, finite diagnostics and independent assistant reviews are distinct.

References

[1] S. Schubr, *Mixed-state escape bounds and proper agnostic learning of stabilizer states in odd-prime dimensions*, Luminary Research Archive, edition 1, 18 October 2026. Preserved author-published edition.

[2] S. Arunachalam, A. Dutt, S. Grewal and A. Gupte, *Marton's conjecture in polynomial time*. arXiv:2609.20771v1, 17 September 2026.

[3] S. Chen, W. Gong, Q. Ye and Z. Zhang, *Stabilizer bootstrapping: A recipe for efficient agnostic tomography and magic estimation*. arXiv:2408.06967v3, 5 December 2024.

[4] W. Qi, B. Xu, S. Feng and L. Li, *Agnostic learning of qudit stabilizer states*. arXiv:2607.15559v1, 17 July 2026.

[5] J. Allcock, J. F. Doriguello, G. Ivanyos and M. Santha, *Reconquering Bell sampling on qudits: stabilizer learning and testing, quantum pseudorandomness bounds, and more*. arXiv:2510.06848v3, 25 August 2026.

[6] M. Hinsche, J. Eisert and J. Carrasco, *The abelian state hidden subgroup problem: Learning stabilizer groups and beyond*. arXiv:2505.15770v3, 27 May 2026.

[7] J. Nöller, V. T. Tran, M. Gachechiladze and R. Kueng, *An infinite hierarchy of multi-copy quantum learning tasks*. arXiv:2510.08070v2, 1 September 2026.

[8] V. Iyer and D. Liang, *Tolerant Testing of Stabilizer States with Mixed State Inputs*. arXiv:2411.08765v2, 12 May 2025.

[9] I. Holt and S. Subramanian, *Learning quantum symmetries*. arXiv:2609.38085v2, 30 September 2026.

[10] 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, 1 October 2026.

[11] W. K. Wootters and B. D. Fields, *Optimal state-determination by mutually unbiased measurements*, Annals of Physics **191** (1989), 363–381. doi:10.1016/0003-4916(89)90322-9.

[12] E. Hostens, J. Dehaene and B. De Moor, *Stabilizer states and Clifford operations for systems of arbitrary dimensions and modular arithmetic*, Physical Review A **71** (2005), 042315. doi:10.1103/PhysRevA.71.042315.