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

Scott Schubr
Luminary Oddities

November 26, 2026 (edition 5)

Abstract

Qi, Xu, Feng and Li (2026) gave an agnostic learner for pure qudit stabilizer states in odd prime dimension and named its extension to mixed states as a direction for further investigation, noting that this requires an anticoncentration bound for skewed Bell difference sampling on mixed inputs. We carry out this extension in every odd prime dimension p. First, for every density operator the output law of skewed Bell difference sampling is the pure-input formula of Qi et al. with the characteristic distribution replaced by a normalized Weyl-overlap law. Second, for every maximizing stabilizer Lagrangian S and proper subspace T of S, the sampled label lies in $S \setminus T$ with probability at least $(p-1)F^r/(4p)$, where F is the stabilizer fidelity, $r = 6$ in general and $r = 4$ when $p \equiv 1 \pmod 4$ and the skew matrix has a zero entry; evenness of the Weyl-overlap law replaces the Hermiticity used for qubits. Third, a proper agnostic learner for arbitrary mixed inputs with promise $F \geq \tau$ returns a pure stabilizer of overlap at least $F - \varepsilon$ with confidence $1 - \delta$, and the stabilizer fidelity is estimated to within ε. At fixed τ and δ the learner uses $O(n \log^2(n+1) + \varepsilon^{-2})$ copies, independently of p; gates and classical time are $\exp(O(\log^2(2/\tau)))\,\mathrm{poly}(n, p, \varepsilon^{-1}, \log(2/\delta))$. It uses Clifford circuits and Weyl eigenspace measurements, and combines the escape bound with a constant-threshold spectral test, isotropic rank growth and empirical spectral-mass estimates.

Keywords. Stabilizer states; agnostic quantum learning; mixed states; qudits; skewed Bell difference sampling; anticoncentration; Clifford measurements.
MSC 2020. Primary 81P45; Secondary 68Q12, 68Q32, 81P15.

1 Introduction and results

1.1 The question

An agnostic stabilizer learner is given copies of an unknown state and must output a pure stabilizer state whose overlap with it is within ε of the largest possible. For qubits, Chen, Gong, Ye and Zhang [4, Corollary 6.3] solved this task for arbitrary mixed inputs by stabilizer bootstrapping. Qi, Xu, Feng and Li [11, Theorem 13] extended stabilizer bootstrapping to pure inputs on n qudits of odd prime dimension p (Qi et al. write d for p). They replaced Bell difference sampling by the skewed Bell difference sampling of Allcock, Doriguello, Ivanyos and Santha [1], and proved an anticoncentration bound for its first-string marginal on pure inputs [11, Theorem 7].

In Section VI of [11], item 2 of the directions listed for further investigation asks for the qudit counterpart of the mixed-state extension of Chen et al. Qi et al. note that this "requires establishing an anticoncentration bound in this broader setting". They name two points to overcome: the output

distribution of skewed Bell difference sampling must be determined for mixed-state inputs, and the qubit anticoncentration proof of [4] uses the Hermiticity of Pauli operators.

This article answers item 2 of [11, Section VI] for every odd prime local dimension, the setting of [11], in three parts.

1. *Output law.* Proposition 4.2 gives the law of skewed Bell difference sampling for every density operator ρ: it is the pure-input formula of [11, Definition 7], with the characteristic distribution replaced by the normalized Weyl-overlap law A_ρ of (4). Its first-string marginal is the law B_ρ of Section 4.

2. *Anticoncentration.* Theorem 1.2 gives $B_\rho(S \setminus T) \geq \frac{p-1}{4p} F^{r_p}$ for every density operator, every maximizing stabilizer Lagrangian S and every proper subspace $T < S$. The evenness $A_\rho(-x) = A_\rho(x)$, valid for every density operator, takes the place of Hermiticity. By Remark 4.3, the first string of every skew matrix admitted in [11, Definition 7] satisfies the same bound with exponent 6, and with exponent 4 when $p \equiv 1 \pmod 4$ and the skew matrix has a zero entry. On pure inputs these bounds improve the bound $(p-1)F^6/(4p^3)$ of [11, Theorem 7] by the factor p^2, with the power of F lowered to four in the zero-entry case.

3. *Mixed-input learner.* Theorem 1.1 is a proper agnostic learner for arbitrary mixed inputs, and Corollary 7.2 estimates the stabilizer fidelity of an arbitrary mixed input, the counterparts of [4, Corollaries 6.3 and 6.4] and of [11, Theorem 13 and Corollary 2]. The learner uses Clifford circuits, computational-basis measurements inside Clifford conjugation, and post-selection on Weyl eigenspaces. These are among the operations used by Qi et al.: Clifford synthesis and computational-basis measurement [11, Section II], and joint-eigenbasis measurement and post-selection [11, Algorithm 2 and Lemma 14]. By Remark 4.3, its label samples may equally be drawn by the skewed Bell difference sampling of [11].

The answer matters for the noise model that motivates agnostic learning. If a pure state is prepared through any noise channel, the resulting mixed state is a valid input, and the guarantee is measured against its own stabilizer fidelity. Moreover, at fixed promise and confidence the original-copy bound (2) carries no dependence on p. The pure-input bound of [11, Theorem 13] grows as $(p/\tau)^{O(p^2 \log(1/\tau))}$.

Edition 5 identifies the question answered, adds Proposition 4.2, Remark 4.3 and Corollary 7.2, and keeps the statements, proofs and numbering of edition 4 [13], with positive-wording edits only.

1.2 Main results

The principal algorithmic conclusion is a proper learner for arbitrary mixed inputs in every odd-prime local dimension, with an absolute quasipolynomial inverse-overlap exponent. Local dimension remains in a fixed polynomial resource factor. The learner theorem is supported by two structural improvements: an even quotient-convolution escape bound and constant-threshold amplification obtained by selecting an observed Weyl eigenvalue. The original-copy bound is uniform in the local prime and nearly linear in n at fixed promise, error and confidence. The escape, conditioning and rank-growth estimates below supply these guarantees.

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.

We use independent input copies, exact prime-dimensional Clifford operations and coherent Weyl eigenspace measurements, with finite-field arithmetic. The rational-input and accuracy conventions are specified below.

Theorem 1.1 (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

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

with absolute exponential constants and fixed polynomial degree independent of p. More precisely, with $d = \log(2/\delta)$, the number of original input copies is at most

$$\exp(O(L^2))\left[n\log^2(n+1)d + \varepsilon^{-2}d^2\right], \tag{2}$$

with absolute constants independent of p. At fixed promised τ, the bound is polynomial in inverse additive accuracy. The no-promise corollary below is quasipolynomial in inverse accuracy.

The escape estimate below is a structural input to this learner. It is proved first, before the conditioning and capped-learning construction.

Theorem 1.2 (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*

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

A draw uses r_p independent input copies and a known stabilizer measurement with polynomial overhead in n, p. These copy counts suffice for the theorem.

1.3 Relation to earlier work

Prior work and quantitative increment. 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 may be any state; it is required neither to be a stabilizer state nor to follow a specified noise model. Chen, Gong, Ye and Zhang [4, Corollary 6.3] established a proper mixed-qubit learner using stabilizer bootstrapping. Qi, Xu, Feng and Li [11] developed agnostic learning for pure inputs in odd-prime local dimension. Section 1.1 describes the mixed-input direction of their Section VI that the present results answer.

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 [1]; related tensor-Weyl measurements occur in Nöller, Tran, Gachechiladze and Kueng [10]. Hinsche, Eisert and Carrasco [5] explicitly anticipate a qudit agnostic extension using their measurement tools.

Related tasks require care in comparison. Holt and Subramanian [6, Theorem 9.27] learn the exact stabilizer group of a mixed qudit state, rather than its best pure stabilizer comparator. Bittel

et al. [3, 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.

Operation model and parameters. State access consists of independent input copies. The operation model uses exact prime-dimensional Clifford operations (including Fourier and SUM), coherent Weyl eigenspace measurements, and finite-field arithmetic, with fixed polynomial overhead in n, p. Conditioning is implemented by the charged rejection procedure below.

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.

Prior dissemination and quantitative strengthening. Edition 1 of this article [12] gave a mixed-input escape estimate and a complete learner with resource factor $\exp(O(p^2 L(L + \log p)))$, where $L = \log(2/\tau)$. The present treatment retains the mixed local inequality and replaces that factor by $\exp(O(L^2))$, with absolute exponent constants and a fixed polynomial degree independent of p. The even quotient-convolution bound, constant-threshold spectral test and observed-eigenvalue selection give this quantitative improvement. Earlier escape constants were $3(p-1)/(4p^3)$ and $(p-1)/(2p^2)$; the present bound is $(p-1)/(4p)$ with the residue-dependent powers stated above. All arguments needed for the present theorem are supplied below.

Table 1: Selected proper stabilizer-learning guarantees. Input classes, overlap promises and measurement access are part of each result.

Source	Input / local dimension	Overlap regime	Measurement access
Chen et al. [4]	Mixed / $p = 2$	Arbitrary positive promised overlap	Single- and two-copy measurements
Qi et al. [11]	Pure / odd-prime p	Arbitrary positive promised overlap	Skewed Bell sampling and SWAP-test correlation estimates
Bittel et al. [3]	Mixed / $p = 2$	Known $f_0 > \sqrt{2/3}$	Adaptive single-copy measurements
This work	Mixed / every odd-prime p	Arbitrary positive promised overlap	Exact Clifford and coherent Weyl measurements; four- or six-copy escape sampler, or skewed Bell difference sampling (Proposition 4.2)

The comparison in Table 1 separates input and access assumptions. The proofs below give the mixed local inequality and escape estimate first, then establish coherent conditioning, robust capture and the capped learner with its complete original-copy accounting.

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 [7, 11].

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

$$A_\rho(x) = D^{-1} \mathrm{Tr}(\rho W_x \rho W_x^\dagger), \qquad \widehat{A}_\rho(u) = \sum_x A_\rho(x)\omega^{[u,x]}. \tag{4}$$

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

$$\widehat{A}_\rho(u) = |\chi_\rho(u)|^2. \tag{5}$$

For example, expanding $\rho = D^{-1} \sum_v \chi_\rho(-v) W_v$ in (4) yields $A_\rho(x) = D^{-2} \sum_v |\chi_\rho(v)|^2 \omega^{[x,v]}$, which is the inversion of (5). 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; for mixed inputs the normalized law A_ρ is the one that enters below. The analysis uses A_ρ through the convolution identity (10) of Section 4. 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 (5) is the standard Weyl-overlap identity in our odd-prime conventions. General Fourier-difference sampling is also treated by Holt and Subramanian [6, 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*

$$\sum_j \tau_{jj}^2 - \frac{1}{p} \mathrm{Tr}(\tau^2) \geq c_p f^2, \qquad c_p = \frac{p-1}{4p}. \tag{6}$$

Proof. The $p + 1$ stabilizer mutually unbiased bases correspond to the lines in $\mathbb{F}_p^2$ [14, 11]. 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 (6) 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 $H < S$,*

$$A_\rho(H) \ge F^2/p, \qquad A_\rho(S \setminus H) \ge c_p F^2. \tag{7}$$

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

$$A_\rho(H) = p^{-1} \sum_y \mathrm{Tr}(\sigma_y^2), \tag{8}$$

$$A_\rho(S \setminus H) = \sum_y \left(\sum_j (\sigma_y)_{jj}^2 - p^{-1} \mathrm{Tr}(\sigma_y^2) \right). \tag{9}$$

Every summand in (9) 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 [1, 10]. 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

$$\widehat{B}_\rho(u) = \prod_i |\chi_\rho(t_i u)|^2, \qquad B_\rho = A_{t_1} * \cdots * A_{t_k}, \quad A_t(x) = A_\rho(t^{-1}x). \tag{10}$$

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 is a joint measurement on the k registers and retains their joint eigenspaces.

Proof of Theorem 1.2. 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 [1].

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

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

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 $\mu(0) = 1 - \Delta/M$. Nonzero quotient atoms occur in distinct pairs $u, -u$, hence

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

Each scaled factor A_t in (10) 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

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

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. [11, 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.

The next proposition determines the law of the skewed Bell difference sampling of Allcock et al. [1, Definition 1], in the four-copy form of Qi et al. [11, Definition 7], on an arbitrary density operator. Write $\overline{O}$ for the entrywise complex conjugate of an operator O in the computational basis and $J(a, b) = (-a, b)$. Since X is real and $\overline{Z} = Z^{-1}$, the definition of W_u gives $\overline{W_u} = W_{Ju}$. For two n-qudit registers put $|\Phi\rangle = D^{-1/2} \sum_{q \in \mathbb{F}_p^n} |q\rangle|q\rangle$ and $|\Phi_y\rangle = (W_y \otimes I)|\Phi\rangle$ for $y \in V$; these vectors form the Bell basis.

Proposition 4.2 (Mixed-input law of skewed Bell difference sampling)**.** *Let p be an odd prime and let $a_1, a_2 \in \mathbb{F}_p$ satisfy $a_1^2 + a_2^2 = -1$. Let B_R act on two n-qudit registers by $|q, q'\rangle \mapsto |a_1 q + a_2 q', a_2 q - a_1 q'\rangle$, coordinatewise; this is $|Q\rangle \mapsto |QR\rangle$ with $R = \left(\begin{smallmatrix} a_1 & a_2 \\ a_2 & -a_1 \end{smallmatrix}\right)$. One round takes four copies of a density operator ρ, applies B_R to registers 2 and 4, and measures the pairs $(1, 2)$ and $(3, 4)$ in the Bell basis, with outcomes $(Y_1, Y_2) \in V^2$. Two independent rounds with outcomes (Y_1', Y_2') and (Y_1'', Y_2'') give $X_i = Y_i'' - Y_i'$. Then for all $x_1, x_2 \in V$,*

$$\Pr(X_1 = x_1,\, X_2 = x_2) = \sum_{y_1, y_2 \in V} A_\rho(x_1 + y_1) A_\rho(x_2 + y_2) A_\rho(a_1 y_1 + a_2 y_2) A_\rho(a_2 y_1 - a_1 y_2). \quad (*)$$

The first string X_1 has law $A_1 * A_{a_1} * A_{a_2}$, in the notation of (10), with each factor of weight zero omitted. With $(a_1, a_2) = (a, b)$ for $p \equiv 3 \pmod 4$, and $(a_1, a_2) = (a, 0)$ for $p \equiv 1 \pmod 4$, this law is B_ρ.

For pure ρ one has $A_\rho = p_\rho$, and $(*)$ is the formula of [11, Definition 7]. Three conventions of [11, Definition 7] match the proposition. First, Qi et al. write Weyl operators with the X-exponent first; this exchanges the two halves of every label, a linear bijection of V, and leaves the form of $(*)$ unchanged. Second, their measured state is $(U_R^\dagger)^{\otimes 2}\psi^{\otimes 8}U_R^{\otimes 2}$, that is, B_{-R} in place of B_R; since $-R$ is again admissible and A_ρ is even, $(*)$ is unchanged. Third, their projector Π_{x_i} pairs the i-th Bell measurement of the two rounds, as in [1, Algorithm 1 and the paragraph preceding Definition 1]; this is the reading under which their displayed formula holds. For mixed ρ the normalized law A_ρ is the right replacement for p_ψ: with the characteristic weight p_ρ in its place, the right side of $(*)$ has total mass $(\operatorname{Tr}\rho^2)^4$, which is less than one for every mixed ρ.

Proof. Bell measurements. The operators $W_v \otimes W_w$ form an orthogonal basis of operators on two registers, and $\langle\Phi|(W_v^\dagger \otimes W_w^\dagger)|\Phi\rangle = D^{-1}\operatorname{Tr}(W_v^\dagger \overline{W_w})$. Since $\overline{W_w} = W_{Jw}$, expansion in this basis gives $|\Phi\rangle\langle\Phi| = D^{-2}\sum_v W_v \otimes \overline{W_v}$. The product rule of Section 2 gives $W_y W_v W_y^\dagger = \omega^{[y,v]}W_v$. Hence

$$|\Phi_y\rangle\langle\Phi_y| = D^{-2}\sum_v \omega^{[y,v]} W_v \otimes \overline{W_v}, \qquad \sum_y \omega^{[u,y]}|\Phi_y\rangle\langle\Phi_y| = W_u \otimes \overline{W_u},$$

the second by character orthogonality, because $[u,y]+[y,v] = [y, v-u]$. So if a two-register state σ is measured in the Bell basis with outcome Y, then $\mathbb{E}\,\omega^{[u,Y]} = \operatorname{Tr}\big(\sigma(W_u \otimes \overline{W_u})\big)$. For two pairs measured jointly, the generating function factors in the same way.

Effect of B_R. Since $R^\top = R$ and $R^2 = -I$, the inverse of B_R is B_{-R}. Hence $B_R(X^b \otimes X^{b'})B_R^\dagger = X^{a_1 b + a_2 b'} \otimes X^{a_2 b - a_1 b'}$. For the Z-part, write a basis state as $|QR\rangle$ and use $Q = -(QR)R$; this gives $B_R(Z^a \otimes Z^{a'})B_R^\dagger = Z^{-(a_1 a + a_2 a')} \otimes Z^{-(a_2 a - a_1 a')}$. Replace R by $-R$ and apply the result to $\overline{W_u} \otimes \overline{W_{u'}} = W_{Ju} \otimes W_{Ju'}$. On the Z-part the sign introduced by J meets the sign in the Z-rule, and on the X-part the sign comes from $-R$. Both halves of each label are therefore transformed by the same scalar combination, and for some scalar ζ of modulus one,

$$B_R^\dagger(\overline{W_u} \otimes \overline{W_{u'}})B_R = \zeta\, W_{-(a_1 u + a_2 u')} \otimes W_{-(a_2 u - a_1 u')}.$$

Generating function. Combining the two steps, one round satisfies

$$\mathbb{E}\,\omega^{[u,Y_1]+[u',Y_2]} = \zeta\,\chi_\rho(u)\chi_\rho(u')\chi_\rho(-(a_1 u + a_2 u'))\chi_\rho(-(a_2 u - a_1 u')).$$

The two rounds are independent, and $|\chi_\rho(-w)| = |\chi_\rho(w)|$ because $W_{-w} = W_w^\dagger$. Therefore

$$\mathbb{E}\,\omega^{[u,X_1]+[u',X_2]} = |\chi_\rho(u)|^2|\chi_\rho(u')|^2|\chi_\rho(a_1 u + a_2 u')|^2|\chi_\rho(a_2 u - a_1 u')|^2.$$

The right side of $(*)$. Multiply by $\omega^{[u,x_1]+[u',x_2]}$ and sum over x_1, x_2. The substitution $z_i = x_i + y_i$ produces the factor $\widehat{A}_\rho(u)\widehat{A}_\rho(u')\,\omega^{-[u,y_1]-[u',y_2]}$. Put $(w_1, w_2) = (y_1, y_2)R$, so that $(y_1, y_2) = -(w_1, w_2)R$. Then $-[u,y_1]-[u',y_2] = [a_1 u + a_2 u', w_1] + [a_2 u - a_1 u', w_2]$, and the remaining sum is $\widehat{A}_\rho(a_1 u + a_2 u')\widehat{A}_\rho(a_2 u - a_1 u')$. By (5), the two transforms agree. The characters $x \mapsto \omega^{[u,x]}$, $u \in V$, are all the characters of V because the form is nondegenerate, so the two laws on V^2 coincide.

The first string. Setting $u' = 0$ and using $\chi_\rho(0) = 1$ gives

$$\mathbb{E}\,\omega^{[u,X_1]} = |\chi_\rho(u)|^2\,|\chi_\rho(a_1 u)|^2\,|\chi_\rho(a_2 u)|^2,$$

where a weight $a_i = 0$ contributes the factor $|\chi_\rho(0)|^2 = 1$. For nonzero t the definition of A_t and (5) give $\widehat{A_t}(u) = \widehat{A}_\rho(tu) = |\chi_\rho(tu)|^2$. Hence the right side is the transform of $A_1 * A_{a_1} * A_{a_2}$ with zero weights omitted, which by (10) is B_ρ for the stated weights. □

Remark 4.3. The proof of Theorem 1.2 given above uses only one property of the label law: it is a convolution of k dilations A_t with nonzero t. For the first string X_1 of Proposition 4.2, let $k \in \{2, 3\}$ be the number of nonzero entries of $(1, a_1, a_2)$. The same proof then gives $\Pr(X_1 \in S \setminus T) \geq c_p F^{2k}$ for every maximizing stabilizer Lagrangian S and every proper $T < S$. With the choices of Section 4, this is Theorem 1.2 itself. The learner of Section 7 uses its label samples only through their law. If they are drawn by skewed Bell difference sampling, each label costs eight conditional copies instead of r_p. In Section 7 the symbol r then plays its two roles separately: the escape exponent r in (12) becomes $2k \leq 6$, and the per-draw copy counts rm and r in the formula for Q in (15) become $8m$ and 8. This changes only the absolute constants in (1) and (2).

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 [7] 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 $\|E+G\| \leq 1+p^{-1/2} \leq 1+1/\sqrt{3} < 33/20$. 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

$$z_{\mathrm{target}}(x) \geq \langle s|\sigma|s\rangle = F_{\mathrm{stab}}(\sigma) \geq \tau. \tag{11}$$

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 $F_{\mathrm{stab}}(\sigma)$ to $F_{\mathrm{stab}}(\sigma)/z_{\mathrm{target}} \geq F_{\mathrm{stab}}(\sigma)/q$. This yields a constant amplification factor $10/9$ independent of p.

6 Robust capture of high spectral masses

Write $r = r_p$. For rational $x \geq 1$ let $\ell(x) = \lceil \log_2 x \rceil$. Set

$$a = \tau^r/12, \qquad m = \left\lceil \frac{2n+24}{a} \right\rceil, \qquad h = 3200\ell(16m). \tag{12}$$

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 C and extend the retained span to a full Lagrangian H, aborting on observed nonisotropy or actual preparation failure. The hidden capture event is used in the analysis.

Lemma 6.1 (Capture with dimension-independent copy counts)**.** *With uncapped conditional-copy access the procedure outputs, with probability at least $3/4$, a full Lagrangian $H \supset C$ such that*

$$B_\sigma(G_\sigma \setminus H) \leq a. \tag{13}$$

The sample counts m, h have uniform upper bounds independent of p.

Proof. Order the spectral outcomes by their representatives $0, \ldots, p-1$. For one sampled label write F_x and $\widehat{F}_x$ for the outcome CDF and its empirical CDF. The one-sample Dvoretzky–Kiefer–Wolfowitz–Massart inequality [9] gives

$$\Pr\left(\sup_t |\widehat{F}_x(t) - F_x(t)| > 1/80 \right) \leq 2\exp(-h/3200).$$

For the first atom take the preceding CDF value to be zero. Every atom is a difference of two successive CDF values. Thus on the complementary event every empirical atom differs from its probability by at most $1/40$, and the empirical maximum differs from $M_\sigma(x)$ by at most $1/40$. This controls all p atoms simultaneously, including unobserved ones. Conditional on all sampled labels, the fresh histogram copies have their corresponding outcome laws. A union bound over the m histograms gives failure at most

$$2m\exp(-h/3200) \leq 1/8.$$

Call its complement the estimation event. On that event every sampled G_σ label is retained, and each retained label has true maximum greater than $33/40$. The spectral concentration lemma therefore makes the retained labels and C isotropic.

For high-label capture we use the standard rank-growth principle of Chen et al. [4, Lemmas 4.21 and 6.7], applied to the isotropic high-label span. Let $X_1, \ldots, X_m$ be the sampled labels and define the auxiliary span

$$U_t = \operatorname{span}(C, \{X_j : j \leq t, X_j \in G_\sigma\}), \qquad U_0 = C.$$

These spans are analysis variables. The high labels commute pairwise and with C, so all U_t are isotropic and have dimension at most n. Their uncovered mass $B_\sigma(G_\sigma \setminus U_t)$ is nonincreasing.

Until the uncovered mass becomes at most a, the conditional probability that $X_t \in G_\sigma \setminus U_{t-1}$ is greater than a; on that event the rank increases by one. Define Y_t to be the indicator of this rank increase while the preceding uncovered mass exceeds a. After coverage has been reached define Y_t instead to be a fresh independent Bernoulli variable of parameter a, used only in this proof. For the filtration including those auxiliary coins,

$$\Pr(Y_t = 1 \mid \mathcal{F}_{t-1}) \geq a.$$

For $\theta > 0$, iterated conditioning consequently gives

$$\mathbb{E}e^{-\theta\sum_{t=1}^{m} Y_t} \le (1 - a + ae^{-\theta})^m \le \exp\{-ma(1 - e^{-\theta})\}.$$

Put $\mu = ma$ and choose $\theta = \log 2$. Markov's inequality yields

$$\Pr\left(\sum_{t=1}^{m} Y_t \le \mu/2\right) \le \exp\{-\mu(1 - \log 2)/2\} \le e^{-\mu/8}.$$

If the final uncovered mass exceeds a, coverage was never reached, so every Y_t counts an actual rank increment. Moreover U_m then has dimension less than n: an isotropic space of dimension n equals its symplectic orthogonal, and every high label commutes with U_m. Hence on final failure

$$\sum_{t=1}^{m} Y_t = \dim U_m - \dim C < n \le \mu/2,$$

since $\mu \ge 2n+24$. Thus final uncovered mass exceeds a with probability at most $e^{-\mu/8} \le e^{-3} < 1/8$.

On the estimation event the actual output H contains U_m. Consequently its uncovered high-label mass is at most that of U_m. The two failure probabilities sum to less than $1/4$, without an independence assumption. This proves the lemma. □

7 A capped learner with uniform resources

Set the following counters in order:

$$q = 9/10, \qquad J = \min\{j \in \mathbb{Z}_{\ge 1} : \tau q^{-j} > 1\}, \tag{14}$$

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

$$Q = J(rm + mh + r + 1), \qquad K = \left\lceil \frac{\ell(4Q/s_0)}{\tau} \right\rceil, \tag{15}$$

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

$$v = \left\lceil \frac{2\ell(4N/\delta)}{\varepsilon^2} \right\rceil. \tag{16}$$

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.

The algorithm uses only observable code and label tests and measurements; the target space and the favorable events enter only the analysis. 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 as a joint measurement of the code generators, preserving within-code coherence. Abort if all K trials fail. Wrong prefixes have the same deterministic cap; acceptance-weight bounds are used only on active prefixes.

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 1.1. 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. Normalized conditional states are needed only while the witness is active. 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 1.2 applied to $S \cap H$ yields $B_\sigma(S \setminus H) \geq F_{\mathrm{stab}}(\sigma)^r/6 \geq 2a$. 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 (11) 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 $2Ne^{-v\varepsilon^2/2} \le \delta/2$. 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

$$RQK + RJ + Nv. \tag{17}$$

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

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

Every constant is independent of p. Write $u = \log(n+1)$ and $d = \log(2/\delta)$. Since $n \ge 1$ and $L, d \ge \log 2$, powers of L and τ^{-1} are absorbed by $\exp(O(L^2))$. The displayed counters give

$$RQK + RJ \le \exp(O(L^2))nu^2d, \qquad Nv \le \exp(O(L^2))\varepsilon^{-2}d^2.$$

Indeed $N \le \exp(O(L^2))d$ and $\log(4N/\delta) = O(L^2+d)$, which pay the validation term. Equation (17) proves (2). The two residue-class choices $r \in \{4, 6\}$ have the same uniform upper bound. 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. A histogram stores at most h observed labels and their counts, each label using $O(\log p)$ bits. This establishes (1).

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 $20\max_\text{count} > 17h$. Arbitrary long input encodings incur their stated additional processing cost. All comparisons are between integers or rationals. □

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. Original copies satisfy the same bound with p omitted from the polynomial factor.

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 1.1 applies. Costs remain bounded when the promise fails, since every run and request is capped. Polynomial inverse-accuracy factors are absorbed into the exponential. □

The learner also estimates the stabilizer fidelity, a standard measure of magic, of an arbitrary mixed input. This is the qudit, mixed-input counterpart of [4, Corollary 6.4] and [11, Corollary 2].

Corollary 7.2 (Stabilizer-fidelity estimation)**.** *Let p be an odd prime, $n \ge 1$ and $0 < \varepsilon, \delta < 1$. Given independent copies of any density operator ρ on n qudits, an algorithm outputs a number $\widehat{F}$ with $|\widehat{F} - F_{\mathrm{stab}}(\rho)| \le \varepsilon$ with probability at least $1 - \delta$. Its gates and classical time are bounded by $\exp(O(\log^2(2/\varepsilon)))\,\mathrm{poly}(n, p, \log(2/\delta))$, with absolute exponent constants and fixed polynomial degree. Its original copies satisfy the same bound with p omitted from the polynomial factor.*

Proof. Run the learner of Corollary 7.1 with accuracy $\varepsilon/2$ and confidence $\delta/2$. Except with probability at most $\delta/2$, it returns a pure stabilizer s with $\langle s|\rho|s\rangle \geq F_{\mathrm{stab}}(\rho) - \varepsilon/2$. By the definition of F_{stab}, also $\langle s|\rho|s\rangle \leq F_{\mathrm{stab}}(\rho)$. Next, measure the rank-one projector $|s\rangle\langle s|$ on $v' = \lceil 2\varepsilon^{-2}\log(4/\delta)\rceil$ fresh original copies, as in the validation step of Section 7, and let $\widehat{F}$ be the observed frequency of the outcome $|s\rangle\langle s|$. By Hoeffding's inequality, $|\widehat{F} - \langle s|\rho|s\rangle| \leq \varepsilon/2$ except with probability at most $2e^{-v'\varepsilon^2/2} \leq \delta/2$. On the intersection of the two events,

$$F_{\mathrm{stab}}(\rho) - \varepsilon \leq \widehat{F} \leq F_{\mathrm{stab}}(\rho) + \varepsilon/2.$$

For the costs, note that $\log(4/\varepsilon) \leq 2\log(2/\varepsilon)$ and $\log(4/\delta) \leq 2\log(2/\delta)$, because $\varepsilon, \delta < 1$. So Corollary 7.1 at $(\varepsilon/2, \delta/2)$ stays within the stated bound. The additional $v' = O(\varepsilon^{-2}\log(2/\delta))$ copies and the Clifford circuit for $|s\rangle\langle s|$ are absorbed in the same way. □

8 Scope, attribution and verification

The mixed local estimate of edition 1 [12] 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 copy bound (2) is uniform over all odd primes, with $O_{\tau,\delta}(n\log^2(n+1) + \varepsilon^{-2})$ copies at fixed promise and confidence. Gate and classical resources retain a fixed polynomial dependence on p. Relative to the dimension-uniform learner of earlier editions, rank-growth capture removes subspace enumeration and the empirical-CDF estimate removes the alphabet-size factor in the histogram budget. The rank-capture principle and distribution-free concentration are inherited inputs; their use with the mixed-state escape and charged conditioning gives the improved copy guarantee.

The proofs are analytic. The edition 4 package [13] provides exact-rational and integer checks of the finite identities and physical-copy request counts, together with floating-point diagnostics and negative controls. The scripts of this edition simulate skewed Bell difference sampling and the weighted measurement of Section 4 on random mixed states, and confirm Proposition 4.2 to rounding accuracy: the joint law $(*)$ for $p \leq 11$ with $n = 1$, and its first marginal together with (10) for $p \leq 13$ with $n = 1$ and for $p \in \{3, 5\}$ with $n = 2$. They also test Lemma 3.2 and Theorem 1.2 on random, noisy and tied inputs for $p \leq 13$ with $n = 1$ and $p \in \{3, 5\}$ with $n = 2$, and on locally optimized inputs for $p \leq 7$ with $n = 1$ and $p = 3$ with $n = 2$, and verify the two factorizations in the proof of Lemma 3.1 symbolically. These computations make the finite calculations reproducible; the uniform resource guarantee follows from the analytic argument.

Statements and Declarations

Funding. This research was self-funded.

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

Data and code availability. The verification scripts of this edition and their outputs are included with it. The edition 4 verification package is archived with edition 4 [13].

AI assistance. Edition 4 was developed with AI assistance in derivation, programming, literature discovery and proof checking. For edition 5, Anthropic Claude identified the answered question, derived and proved Proposition 4.2, Remark 4.3 and Corollary 7.2, wrote the verification scripts, edited this edition and checked the argument. Scott Schubr directed the research and is responsible for the content.

References

[1] 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.

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

[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, 1 October 2026.

[4] S. Chen, W. Gong, Q. Ye and Z. Zhang, *Stabilizer bootstrapping: A recipe for efficient agnostic tomography and magic estimation*. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing (STOC 2025)*, pp. 429–438, 2025. doi:10.1145/3717823.3718191. Version cited: arXiv:2408.06967v3, 5 December 2024.

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

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

[7] 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**(4) (2005), 042315. doi:10.1103/PhysRevA.71.042315.

[8] V. Iyer and D. Liang, *Tolerant testing of stabilizer states with mixed state inputs*. arXiv:2411.08765v2, 12 May 2025.

[9] P. Massart, *The tight constant in the Dvoretzky–Kiefer–Wolfowitz inequality*, Annals of Probability **18**(3) (1990), 1269–1283. doi:10.1214/aop/1176990746.

[10] 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.

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

[12] S. Schubr. *Mixed-state escape bounds and proper agnostic learning of stabilizer states in odd-prime dimensions*, edition 1. Luminary Research Archive, October 2026. `https://research.luminaryoddities.org/mixed-state-escape-bounds-proper-agnostic-learning-stabilizer-states-odd-prime-dimensions/v/1`.

[13] S. Schubr. *Mixed-state escape bounds and proper agnostic learning of stabilizer states in odd-prime dimensions*, edition 4. Luminary Research Archive, November 2026. `https://doi.org/10.5281/zenodo.24472597`.

[14] W. K. Wootters and B. D. Fields, *Optimal state-determination by mutually unbiased measurements*, Annals of Physics **191**(2) (1989), 363–381. doi:10.1016/0003-4916(89)90322-9.