Purity frontiers and optimal copy complexity for qutrit stabilizer measurements

Scott Schubr
Luminary Oddities

October 29, 2026

Abstract

We determine purity-estimation frontiers under stabilizer operations, bounded unconditional outputs and uniform expectation bias. In every fixed qutrit-register dimension, attaining the physical upper output endpoint requires and suffices to use copies of order the inverse square root of the bias tolerance. A five-block moment polarization, smooth clipping and reciprocal-weighted repair give a log-free upper bound. A likelihood-curvature argument supplies the matching lower order for bounded-real outputs. We also determine exact three-copy single-qutrit frontiers and the nonnegative unbiased copy threshold, dimension squared minus one. Pure-null instrument structure explains the role of outcome information.

Keywords. Purity estimation, stabilizer operations, qutrits, output range, uniform bias, copy complexity.
MSC 2020. 81P45, 81P68, 62F12.

1 Model and principal results

Estimators with the same expectation can have different ranges of possible outputs. Purity estimation makes this distinction particularly sharp: pure states have purity one, so a uniform bias tolerance ε forces every upper output endpoint to be at least $1-\varepsilon$. We ask how many copies are needed to attain that endpoint exactly when the measurements are restricted to stabilizer operations. This isolates an outcome-range resource question that ordinary accuracy estimates do not settle. We also determine the exact range tradeoff at three copies and the copy threshold for nonnegative unbiased estimation.

Fix $n \geq 1$, $D = 3^n$, and an arbitrary density matrix $\rho \in \mathcal{D}_D$. Independent identical input copies are supplied. The target is $s(\rho) = \mathrm{Tr}(\rho^2)$. An allowed protocol appends stabilizer ancillas, applies joint Clifford operations, measures Weyl eigenspaces, discards systems, and adapts to classical outcomes. Classical randomization and postprocessing are allowed. The input-copy cap is finite and independent of ρ. Its output is unconditional and bounded: every outcome, including any abort, is assigned a value. Write

$$\mathrm{bias}(Y) = \sup_{\rho \in \mathcal{D}_D} |\mathbb{E}_\rho Y - s(\rho)|.$$

For $0 < \varepsilon \leq 1$, let $\mathcal{C}_D(\varepsilon)$ denote the smallest deterministic input-copy cap among allowed protocols with $0 \leq Y \leq 1-\varepsilon$ and uniform purity bias at most ε. The feasible integer caps form a nonempty set by the upper construction. Pure inputs automatically force attainment of the upper endpoint. Let $\mathcal{C}_D^{\mathrm{R}}(\varepsilon)$ allow bounded real outputs with the same upper endpoint and bias, but no prescribed lower endpoint. Then $\mathcal{C}_D^{\mathrm{R}} \leq \mathcal{C}_D$.

Theorem 1.1 (Optimal copy order)**.** *For each fixed $D = 3^n$, positive constants c_D, K_D satisfy, for all sufficiently small $\varepsilon > 0$,*

$$c_D \varepsilon^{-1/2} \le \mathcal{C}_D^{\mathrm{R}}(\varepsilon) \le \mathcal{C}_D(\varepsilon) \le K_D \varepsilon^{-1/2}. \tag{1}$$

Thus both copy complexities have tight order $\Theta_D(\varepsilon^{-1/2})$. In particular,

$$\lim_{\varepsilon \downarrow 0} \frac{\log \mathcal{C}_D(\varepsilon)}{\log(1/\varepsilon)} = \lim_{\varepsilon \downarrow 0} \frac{\log \mathcal{C}_D^{\mathrm{R}}(\varepsilon)}{\log(1/\varepsilon)} = \frac{1}{2}. \tag{2}$$

The constants depend on the fixed dimension D. This theorem determines the small-tolerance asymptotic order; Corollary 6.2 gives exact integer caps in specified finite-tolerance regimes.

For a single qutrit and at most three input copies, let $U_3^+(\varepsilon)$ be the infimum of the upper endpoint U over $0 \le Y \le U$ with $\mathrm{bias}(Y) \le \varepsilon$, assigning $+\infty$ to an empty feasible class. Let $U_3(\varepsilon)$ be the corresponding infimum for bounded real outputs satisfying $Y \le U$.

Theorem 1.2 (Complete three-copy frontiers)**.** *For a single qutrit, the nonnegative three-copy frontier is*

$$U_3^+(\varepsilon) = \begin{cases} +\infty, & 0 \le \varepsilon < 1/37, \\ 3 - 39\varepsilon, & 1/37 \le \varepsilon \le 1/19, \\ 1 - \varepsilon, & 1/19 \le \varepsilon \le 1. \end{cases}$$

The bounded-real frontier, for every $\varepsilon \ge 0$, is

$$U_3(\varepsilon) = \max\left\{3 - 63\varepsilon, \frac{7 - 59\varepsilon}{5}, 1 - \varepsilon\right\}.$$

Its branches meet at $1/32$ and $1/27$. Every finite value in both frontiers is attained by postprocessing the Schubr parity refinement of Definition 7.1. The exact minimum bias for outputs in $[0, 1]$ is $2/39$.

Theorem 1.3 (Minimum nonnegative unbiased copy cap)**.** *For $D = 3^n$, the smallest deterministic copy cap of a nonnegative exactly unbiased purity estimator is $D^2 - 1$. At this cap a stabilizer protocol has output in $[0, D]$. The lower bound holds even for unrestricted quantum measurements in every dimension $D \ge 2$.*

For arbitrary fixed D, let $U_{\mathrm{fin},D}^+(\varepsilon)$ optimize the upper endpoint over all finite-copy protocols with nonnegative outputs; the cap may depend on D and ε.

Corollary 1.4 (All-finite-cap output frontier)**.** *For every $D = 3^n$,*

$$U_{\mathrm{fin},D}^+(0) = D, \qquad U_{\mathrm{fin},D}^+(\varepsilon) = 1 - \varepsilon \quad (0 < \varepsilon \le 1). \tag{3}$$

Every infimum is attained.

Relation to previous work. Cyclic trace identities and spectral-moment measurements are established. Fan and Luan [3, Sections 6.2–6.3] use controlled permutations and ancillary access for three-copy sector tests; the conditional Weyl resolution below supplies the cyclic spectral effects within the present operation model. Generalized Bell measurements and Calderbank–Shor–Steane (CSS) syndrome measurements and coherent multi-copy learning hierarchies are established [9, Section 4 and Appendix F]. The full four-outcome refinement used for the exact frontiers is specified in Definition 7.1. Its logical test in a marked mutually unbiased basis (MUB) uses the phase-point identity of Gibbons, Hoffman and Wootters [5, Section 5] and Gross's parity convention [6]. The

qutrit third-moment spectrum and CSS symmetric defect sector are established [12, Sections V.C and VI.A]; we give elementary proofs of the precise moment identities used by our witnesses. Clifford commutant structure has a broader theory [2], and the stabilizer normal forms are classical [8].

An earlier version of the present work gives the optimal power exponent with an upper logarithmic factor; the construction below supplies matching order. The author's preceding two-copy purity paper [10] supplies the transcript-cone lemma, the all-finite-cap unbiased output lower bound D, and the fixed-phase-point Chebyshev mechanism. Ordinary clipped tomography and rescaling already give the numerical positive-bias infimum in (3): a statistic $T_\delta \in [0,1]$ of bias at most $\delta < \varepsilon$ can be scaled by $(1-\varepsilon)/(1-\delta)$, with upper endpoint tending to $1-\varepsilon$. The present construction attains that endpoint with tight square-root copy order. MUB frames [11], binomial identities and concentration estimates are supporting ingredients; the quantitative result concerns the joint uniform-bias, unconditional-range and deterministic-cap contract. The exact unbiased cap uses the established Schur–Weyl purity observable [1, Section 7.1].

For even functions, smooth Gaussian approximation can improve from order $N^{-1/2}$ to N^{-1}; this improvement and its parity mechanism are established in Gaunt [4, Lemma 3.2 and Theorem 3.4]. We give a self-contained Taylor–Lindeberg version with constants uniform over the bounded sampling laws and permitting singular covariance. The operational step is to arrange a symmetric bipartite quadratic error by five-block polarization and retain an inverse-binomial weight throughout the clipping estimate. The resulting one-sided correction pays only $O_D(N^{-2})$ in the purity bias while keeping every output in the required interval.

How the two bounds meet. Write q for the probability of a fixed union of three-copy cyclic events. This probability vanishes on pure inputs. If $R = (1 - \operatorname{Tr}\rho^2)/q$, a fresh union bit Q and an independent estimate $\widehat{R}$ give

$$Y = 1 - \varepsilon - Q\widehat{R}, \qquad \mathbb{E}Y - \operatorname{Tr}\rho^2 = -\varepsilon + q(R - \mathbb{E}\widehat{R}).$$

The upper proof constructs $\widehat{R}$ in a fixed subinterval of $(0,1)$ with $0 \le q(R - \mathbb{E}\widehat{R}) \le K_D/N^2$, using $O_D(N)$ copies. Hence N of order $\varepsilon^{-1/2}$ suffices. The lower proof applies to every allowed adaptive protocol: nonnegative formal deficit coefficients on a physical Wigner-positive curve force an error of order at least the inverse square of the cap. The two arguments therefore answer the same question with the same range and bias conventions.

The finite-cap and unbiased regimes. For example, at tolerance $1/30$, three copies of one qutrit have optimal upper endpoint $17/10$ for nonnegative outputs and $151/150$ for bounded-real outputs, whereas sufficiently many copies attain $29/30$ in both classes. The complete three-copy theorem explains this finite-cap distinction. At zero bias the optimal finite-cap upper endpoint becomes D, and nonnegative estimation first becomes possible at $D^2 - 1$ copies. These endpoint phenomena complement the small-positive-tolerance copy law.

Organization. Sections 2 and 3 supply the cyclic event and stabilizer tomography frame. Section 4 gives the correction-ratio geometry. Sections 5 and 6 prove the optimal copy order. Section 7 proves the exact three-copy frontiers, and Section 8 treats zero bias. The appendices identify pure-null instrument structure and the minimum coherent block size needed by the construction.

2 Cyclic measurement primitives

Theorem 2.1 (Cyclic spectral effects)**.** *On three n-qutrit registers, let*

$$C|x,y,z\rangle = |y,z,x\rangle, \qquad J = \frac{2I - C - C^2}{3}, \qquad x,y,z \in \mathbb{F}_3^n.$$

An adaptive stabilizer measurement has the three spectral projectors of C as its coarse outcome effects. In particular its binary event B for a nontrivial cyclic eigenvalue has effect J and

$$p(\rho) := \Pr_\rho(B = 1) = \frac{2}{3}(1 - \mathrm{Tr}(\rho^3)). \tag{4}$$

It vanishes exactly on pure inputs. Three copies are the minimum for a nonzero stabilizer event vanishing on every identical pure input.

Write $\omega = e^{2\pi i/3}$. On a register labelled by $\mathbb{F}_3^n$, set

$$X^a|u\rangle = |u + a\rangle, \qquad Z^b|u\rangle = \omega^{b\cdot u}|u\rangle.$$

First make the invertible linear coordinate change

$$t = x, \qquad u = y - x, \qquad c = x + y + z. \tag{5}$$

Its inverse is $x = t$, $y = t + u$, $z = t - u + c$. An invertible linear computational permutation is Clifford: translations transform by the linear map, and diagonal Weyls by its inverse transpose. In the new coordinates,

$$C : (t,u,c) \longmapsto (t + u, u + c, c). \tag{6}$$

Measure c in the computational basis, and t in the Fourier basis

$$|\chi_l\rangle = D^{-1/2} \sum_{t \in \mathbb{F}_3^n} \omega^{l\cdot t}|t\rangle.$$

Both measurements commute with C: c is preserved, and translation of t preserves its Fourier label. In a branch (c,l), substituting $t' = t + u$ shows that the remaining action on u is

$$V_{c,l}|u\rangle = \omega^{-l\cdot u}|u + c\rangle, \qquad V_{c,l} = X^c Z^{-l}. \tag{7}$$

This is a known Weyl operator with $V_{c,l}^3 = I$. Measure its eigenvalue, adjusting the classical label for any chosen Weyl phase convention. When $c = l = 0$, its eigenvalue is already one.

The first two measurements split the space into reducing subspaces for C. In each subspace the final Weyl measurement diagonalizes its restriction. The complete transcripts are therefore a projective refinement of the cyclic spectral decomposition. Summing transcripts with eigenvalue ω or ω^2 gives exactly J. Finally

$$\mathrm{Tr}(C\rho^{\otimes 3}) = \mathrm{Tr}(C^2\rho^{\otimes 3}) = \mathrm{Tr}(\rho^3),$$

which proves (4). A density matrix has third moment one precisely when it is pure.

Effects and conditional outputs. The construction specifies complete measurement effects. It does not implement a coarse Lüders map. For three single-qutrit inputs, Theorem A.1 proves the stronger statement that every pure-null stabilizer subchannel has a fixed sixteen-outcome measure-and-prepare form. This distinction prevents extracting a quantum conditional state merely from implementability of the cyclic event. Proposition A.2 also shows that retaining only the sixteen rare-event labels and an undifferentiated failure has the slower inverse-bias copy cost; the general upper construction uses additional tomography data.

2.1 The minimum-copy obstruction

The obstruction works in every odd local prime dimension. We give the argument for equal-size registers. A pure odd-prime stabilizer vector has, up to a common nonzero scalar, computational amplitudes that are either zero or powers of a primitive pth root, on one affine support. This is the computational normal form in [8, Appendix C]. For clarity, the part needed here follows directly from its stabilizer labels. Their translation projection is a linear space V. Diagonal labels impose an affine coset of V as the permitted support. The stabilizer translations connect all its positions, and their matrix entries and eigenvalues are pth roots. Hence every ratio of nonzero amplitudes is a pth root.

An antisymmetric vector on two equal-size registers has opposite amplitudes on every exchanged nonzero pair and zero diagonal amplitudes. Since -1 is not a pth root for odd p, no nonzero pure stabilizer vector is antisymmetric.

By the positive transcript-cone lemma [10, Lemma 2.1], every complete allowed transcript effect, padded to its copy cap, is a nonnegative combination of pure stabilizer projectors. One way to see the cone statement is to feed the input ports with halves of maximally entangled stabilizer pairs. Conditioning and tracing the circuit leaves the transpose of its effect on the reference halves. Stabilizer operations, partial traces and transposition preserve this cone. The finitely generated cone is closed, so classical mixtures are included.

Suppose an effect E on two copies vanishes on $|\phi\rangle^{\otimes 2}$ for every ϕ. Positivity implies that E annihilates their span Sym^2. In a nonnegative stabilizer decomposition each summand must annihilate that span individually, forcing its vector to be antisymmetric. The preceding paragraph excludes every nonzero summand. Thus $E = 0$. A one-copy effect vanishing on all pure states is also zero. Together with the constructed three-copy event this proves the last assertion of Theorem 2.1.

3 A stabilizer tomography frame

We recall explicitly the complete stabilizer-MUB frame [11]. In $\mathbb{F}_D$, let $\tau : \mathbb{F}_D \to \mathbb{F}_3$ be the field trace. Take the computational basis and the D bases with vectors

$$|b, a\rangle = D^{-1/2} \sum_{x \in \mathbb{F}_D} \omega^{\tau(bx^2+ax)} |x\rangle, \qquad a, b \in \mathbb{F}_D. \tag{8}$$

Fourier transforms and these quadratic phases are Clifford in n qutrit coordinates, as direct conjugation of translations and modulations shows. Within a basis, character orthogonality gives orthogonality. Between distinct bases, the relevant quadratic Gauss sum has modulus $\sqrt{D}$: square its modulus, put $x = y + u$, and sum over y; the nondegenerate trace pairing leaves only $u = 0$.

Write the projectors as P_{ba}, now using b for all $D + 1$ bases. Their traceless diagonal spaces are mutually Hilbert–Schmidt orthogonal and have total dimension $(D + 1)(D - 1) = D^2 - 1$. Consequently

$$\rho = \sum_{b,a} p_{ba} P_{ba} - I, \qquad \sum_{b,a} p_{ba}^2 = 1 + \mathrm{Tr}(\rho^2), \qquad p_{ba} = \mathrm{Tr}(P_{ba}\rho). \tag{9}$$

Use m independent copies in each basis, with outcome frequencies f_{ba}, and form $T = \sum f_{ba} P_{ba} - I$. It is unbiased for ρ. Independence and the orthogonal traceless spaces give

$$\mathbb{E}\|T - \rho\|_F^2 = \frac{1}{m} \sum_b \left(1 - \sum_a p_{ba}^2\right) \le \frac{D^2 - 1}{Dm}. \tag{10}$$

4 A correction ratio below one

Let e_j be the elementary symmetric functions of the eigenvalues of ρ, with $e_1 = 1$. Write

$$d = 1 - s(\rho) = 2e_2, \qquad p = 2(e_2 - e_3), \qquad d_* = 1 - D^{-1}.$$

For $e_2 > 0$ put $h = e_3/e_2$ and $R_1 = (1-h)^{-1} = d/p$.

Lemma 4.1. *Extending $h = 0$ and $R_1 = 1$ at pure states gives continuous functions on $\mathcal{D}_D$. They satisfy $0 \le h \le 1/3$, $1 \le R_1 \le 3/2$, and*

$$|R_1(\rho) - R_1(\sigma)| \le \frac{9\sqrt{D}}{4}\|\rho - \sigma\|_F, \tag{11}$$

$$|p(\rho) - p(\sigma)| \le 2\sqrt{D}\|\rho - \sigma\|_F. \tag{12}$$

Proof. Expanding $e_1 e_2$ gives $e_1 e_2 \ge 3e_3$, so $h \le 1/3$. Choose a largest eigenvalue x, and let S, A, B be respectively the first three elementary symmetric functions of the remaining eigenvalues. Then

$$Se_2 - 2e_3 = x(S^2 - 2A) + SA - 2B \ge 0,$$

because $S^2 \ge 2A$ and $SA \ge 3B$. Therefore $h \le S/2$, proving continuity at pure states.

With one coordinate x_i varying and the others fixed,

$$\partial_i h = \frac{A^2 - SB}{(x_i S + A)^2} \in [0, 1].$$

Indeed, for the other coordinates y, the numerator expands as

$$\sum_{i<j} y_i^2 y_j^2 + \sum_{i<j<k} y_i y_j y_k (y_i + y_j + y_k) + 2e_4(y) \ge 0,$$

and is at most A^2. Integrate along line segments and pass to boundary limits to obtain $|h(\lambda) - h(\mu)| \le \|\lambda - \mu\|_1$. The derivative of $(1-h)^{-1}$ is at most $9/4$. The Hermitian eigenvalue inequality

$$\|\lambda(\rho) - \lambda(\sigma)\|_1 \le \sqrt{D}\|\rho - \sigma\|_F$$

then gives (11). For completeness, its squared ℓ_2 version follows by expanding $\mathrm{Tr}(\rho - \sigma)^2$ in eigenbases and applying rearrangement to the doubly stochastic squared-overlap matrix; Cauchy–Schwarz gives the displayed form. Finally $|x^3 - y^3| \le 3|x - y|$ on $[0, 1]$ proves (12). □

Choose an integer $r \ge 2$ so that

$$b^r < D^{-1}, \qquad b = \frac{D+2}{3D}. \tag{13}$$

For example $r = 1 + \lceil \log D / \log(1/b) \rceil$ works. On r independent triples measure the cyclic event, and let Q be the indicator that at least one occurs. Its probability is $q = 1 - (1-p)^r$.

Lemma 4.2. *There is a continuous factorization $d = qR_r$, with*

$$\frac{1}{r} \le R_r \le a_D := \frac{d_*}{1 - b^r} < 1. \tag{14}$$

It is Lipschitz in Frobenius norm with constant

$$L = L_{D,r} = \sqrt{D}\left(\frac{9}{4} + \frac{3}{2}r(r-1)\right). \tag{15}$$

Proof. Since $\mathrm{Tr}(\rho^3) \le \mathrm{Tr}(\rho^2)$, we have $p \ge 2d/3$. The function $f(d) = 1 - (1 - 2d/3)^r$ is concave on $[0, d_*]$, vanishes at zero, and satisfies $q \ge f(d)$. Thus $q/d \ge f(d_*)/d_*$ for $d > 0$, giving the upper bound in (14). Set

$$S_r(p) = \sum_{j=0}^{r-1} (1-p)^j, \qquad R_r = R_1/S_r(p).$$

This defines the continuous pure-state value $1/r$ and proves $d = qR_r$ everywhere. Since $1 \le S_r \le r$, the lower bound follows. Also $|S_r'| \le r(r-1)/2$. The quotient bound, $R_1 \le 3/2$, and (11)–(12) give (15). □

Lemma 4.3 (Distance and third moment near a pure state)**.** *Let $u = e_2(\rho)$ and $v = e_3(\rho)$. If P projects onto a largest-eigenvalue eigenvector of ρ, then*

$$\|\rho - P\|_F \le \sqrt{2} D u, \qquad 0 \le v \le C_D u^2, \qquad \|Dv|_\rho\| \le C_D u,$$

where the differential is on trace-zero Hermitian directions.

Proof. Write $x = \lambda_{\max}(\rho)$ and $S = 1 - x$. Since $x \ge 1/D$, one has $u \ge xS \ge S/D$ and $\|\rho - P\|_F \le \sqrt{2} S$. On trace-one Hermitian matrices,

$$v(X) = \frac{1 - 3\,\mathrm{Tr}\, X^2 + 2\,\mathrm{Tr}\, X^3}{6}.$$

Both $v(P)$ and its first differential on trace-zero directions vanish. Taylor expansion and bounded second derivatives on the compact density body give the remaining claims. □

5 A log-free physical-boundary upper bound

We retain the stabilizer operation model, complete stabilizer MUB frame, and cyclic event of Sections 2 and 4. All sample sizes below are deterministic maxima. Constants may depend on the fixed dimension $D = 3^n$, $n \ge 1$, and on fixed smooth functions, but never on the input state or sample size.

Theorem 5.1. *For each fixed $D = 3^n$, there are constants $K_D, \varepsilon_D > 0$ such that, for $0 < \varepsilon \le \varepsilon_D$, an allowed stabilizer protocol uses at most $K_D \varepsilon^{-1/2}$ input copies and has an unconditional output Y satisfying*

$$0 \le Y \le 1 - \varepsilon, \qquad \sup_{\rho \in \mathcal{D}_D} \left| \mathbb{E}_\rho Y - \mathrm{Tr}(\rho^2) \right| \le \varepsilon.$$

Pure inputs attain the upper output endpoint almost surely. Together with Theorem 6.1, this gives $\mathcal{C}_D^{\mathrm{R}}(\varepsilon), \mathcal{C}_D(\varepsilon) = \Theta_D(\varepsilon^{-1/2})$ as $\varepsilon \downarrow 0$.

We first give the smooth comparison estimate used in the construction. The parity improvement is classical [4]; the following finite Taylor–Lindeberg argument gives the uniform version needed here and its formulation permits singular covariance matrices.

Lemma 5.2 (Even smooth comparison)**.** *Let $\Xi_1, \dots, \Xi_N$ be independent identically distributed centred vectors in a fixed-dimensional real vector space, uniformly bounded in norm, and let $Z = N^{-1/2} \sum_i \Xi_i$. Let G be a centred Gaussian vector with covariance $\mathrm{Cov}(\Xi_1)$. If the derivatives of F through order three are bounded by a fixed polynomial in the norm, then*

$$|\mathbb{E} F(Z) - \mathbb{E} F(G)| \le C N^{-1/2}.$$

If F is even and has such derivative bounds through order six, then

$$|\mathbb{E}F(Z) - \mathbb{E}F(G)| \le CN^{-1}.$$

The constants are uniform over the bounded input distributions. If all derivative envelopes used in the comparison have a common factor α, the corresponding bounds have the factor α.

Proof. Replace the N increments successively by independent centred Gaussian increments of the same covariance. All fixed-order moments of the normalized hybrid partial sums are uniformly bounded. Taylor expansion through degree two cancels the constant, linear and quadratic terms between the two increments. The integral third-order remainder is $O(N^{-3/2})$ in expectation at each replacement, proving the first claim.

For an even F, expand instead through degree three. The fourth-order remainders sum to $O(N^{-1})$. If U_i denotes the hybrid sum excluding the increment being replaced, the remaining third-order contribution is the third moment tensor of Ξ_i, contracted with $\mathbb{E}D^3F(U_i)$, times $N^{-3/2}/6$. Every component of D^3F is odd. Its expectation at a centred Gaussian of covariance $\mathrm{Cov}(U_i)$ is zero. Applying the preceding degree-two replacement argument to that component, with the $N-1$ increments of U_i, bounds its expectation by $O(N^{-1/2})$. This uses derivatives of F through order six and is uniform over which increments have already been replaced. The sum of the third-order terms is therefore $O(N \cdot N^{-3/2} \cdot N^{-1/2}) = O(N^{-1})$. The argument uses only Taylor's formula and moment bounds, so neither invertibility of covariance nor a Gaussian density is required. Every estimate is linear in its derivative envelope, proving the last assertion. □

Put

$$u = e_2(\rho), \qquad v = e_3(\rho), \qquad p = 2(u-v), \qquad q = 1-(1-p)^r,$$
$$A(p) = \frac{p}{q} = \frac{1}{S_r(p)}, \qquad B = \frac{2v}{q}, \qquad S_r(p) = \sum_{j=0}^{r-1}(1-p)^j.$$

At pure states set $A = 1/r$ and $B = 0$. Then $A+B = (1-\mathrm{Tr}\,\rho^2)/q$, with its continuous interpretation at $q = 0$. The preceding ratio estimates give

$$q \asymp_{D,r} u, \qquad 0 \le v \le C_{D,r}q^2. \tag{LF1}$$

If P projects onto a largest-eigenvalue eigenvector of ρ, the nearest-pure estimate also gives $\|\rho - P\|_F \le C_{D,r}q$.

We choose r large enough to separate the two components in output range. Indeed $0 \le p \le 2/3$, A is increasing, and $A_r(2/3) \longrightarrow 2/3$. Also $2v/q_r$ decreases uniformly on the compact state space to $2v$, with value zero at pure states. For completeness, continuity at a pure state follows from $v = O(p^2)$ and $q_r \ge p$; monotonicity and the continuous limiting function give uniform convergence. Maclaurin's inequality gives $\max_\rho 2v = (D-1)(D-2)/(3D^2)$. Thus the limiting sum of the two separate suprema is

$$\frac{2}{3} + \frac{(D-1)(D-2)}{3D^2} = 1 - \frac{1}{D} + \frac{2}{3D^2} < 1.$$

Fix r and strict margins

$$0 < \ell_A < 1/r, \quad A_r(2/3) < b_A, \quad \sup_\rho B < b_B, \quad b_A + b_B < 1.$$

Choose $\xi > 0$ with $\xi < \ell_A/4$ and $\beta = b_A + b_B + \xi < 1$.

A bounded estimate of the analytic component. Choose a fixed C^4 function F on $[0,1]$ agreeing with A on $[0,2/3]$ and taking values strictly inside (ℓ_A, b_A). Such a function is obtained by extending the analytic function A slightly past this interval and smoothly flattening its argument before its values reach the margins. From N independent cyclic-event bits let $\widehat{p}$ be their mean. Define

$$A_0 = F(\widehat{p}) - \frac{\widehat{p}(1-\widehat{p})F''(\widehat{p})}{2N}.$$

The exact second and third binomial central moments and the fourth central-moment bound give

$$\mathbb{E}F(\widehat{p}) = F(p) + \frac{p(1-p)F''(p)}{2N} + O(N^{-2})$$

uniformly in p. The function $x(1-x)F''(x)$ has two bounded derivatives, so its expectation at $\widehat{p}$ differs from its value at p by $O(N^{-1})$. Hence $|\mathbb{E}A_0 - A(p)| \le c_A/N^2$ for a fixed c_A. Set $\widehat{A} = A_0 - c_A/N^2$. For all sufficiently large N,

$$\ell_A \le \widehat{A} \le b_A, \qquad 0 \le A - \mathbb{E}\widehat{A} \le 2c_A/N^2. \tag{LF2}$$

The cost is $3N$ copies. This is the usual smooth plug-in bias correction, applied to the analytic component of the ratio.

Five-block polarization. Take five independent raw MUB estimates A_1, A_2, B_1, B_2, B_3, each with N independent samples in each of the $D+1$ bases. They are unbiased, trace one, and uniformly bounded. For trace-one Hermitian matrices put

$$V(X,Y,Z) = \frac{1 - \mathrm{Tr}(XY + XZ + YZ) + \mathrm{Tr}(XYZ + XZY)}{6}$$

and define

$$V_{\mathrm{bip}} = \frac{1}{4} \sum_{\substack{i=1,2 \\ 1\le j<k\le 3}} V(A_i, B_j, B_k) - \frac{1}{2} V(B_1, B_2, B_3). \tag{LF3}$$

Separate affinity and independence give $\mathbb{E}V_{\mathrm{bip}} = v$. To identify its quadratic term, write $H_\rho = D^2 e_3|_\rho$ on trace-zero Hermitian directions. A single polarized triple has pair coefficient $H_\rho/6$. In (LF3), each BB edge has weight $2/4 - 1/2 = 0$, each AB edge has weight $2/4 = 1/2$, and there are no AA edges. At a pure P, $De_3|_P = 0$ and

$$H_P(E,F) = 2K_P(E,F), \qquad K_P(E,F) = \frac{\mathrm{Tr}\, P(EF+FE) - \mathrm{Tr}(EF)}{2}.$$

Consequently its pure-state quadratic term is

$$K_P\left(\frac{E_{A_1} + E_{A_2}}{2}, \frac{E_{B_1} + E_{B_2} + E_{B_3}}{3}\right). \tag{LF4}$$

In particular it is bilinear between independent groups.

Independently take $H \sim \mathrm{Bin}(N,q)$, using N union trials of r cyclic events, and put

$$a = (H+1)^{-1}, \qquad W = 2(N+1)aV_{\mathrm{bip}}.$$

The inverse-binomial identity gives, for $q > 0$,

$$\mathbb{E}W = B[1-(1-q)^{N+1}], \qquad 0 \le q(B - \mathbb{E}W) = 2v(1-q)^{N+1} \le C/N^2. \tag{LF5}$$

The corresponding weighted identities hold at $q = 0$ as well. Uniform boundedness of the raw matrices gives $|W| \le C(N+1)$.

Lemma 5.3 (Weighted smooth clipping)**.** *Let ψ_0 be a fixed odd smooth bounded saturation with range $[-\xi,\xi]$, equal to the identity on $[-\eta,\eta]$ for some $\eta>0$. There are $q_0>0$ and C such that*

$$|\mathbb{E}[\psi_0(W)-W]| \le \frac{C}{N}\,\mathbb{E}\frac{1}{H+1} \tag{LF6}$$

for every state with $q\le q_0$ and every $N\ge 2$.

Proof. A C^∞ saturation may be used, so all derivatives needed below exist and are bounded. Write the joint centred errors of the five MUB blocks as $E_i = Z_i/\sqrt{N}$. Grouping the independent observations by their sample index represents $Z = N^{-1/2}\sum_{j=1}^{N}\Xi_j$ with iid centred bounded vectors. Their group covariances are block diagonal between the five blocks; within each block no coordinate independence is required. All fixed moments of Z, and of the Gaussian hybrid sums in Lemma 5.2, are uniformly bounded.

The exact cubic expansion of (LF3), together with $De_3|_P = 0$ and $\|\rho - P\| = O(q)$, gives

$$V_{\mathrm{bip}} = v + \frac{q}{\sqrt{N}}L(Z) + \frac{Q_P(Z)+qQ_\Delta(Z)}{N} + \frac{C_3(Z)}{N^{3/2}},$$

where L is linear, Q_Δ quadratic and C_3 cubic, with uniformly bounded coefficients. At $q=0$ the terms multiplied by q are assigned arbitrary bounded coefficients. By (LF4), Q_P is bipartite. Its value at the covariance-matched Gaussian is symmetric about zero: change the sign of the entire A group. This remains valid for singular covariance. Under simultaneous inversion of all coordinates, Q_P, Q_Δ are even and L, C_3 are odd.

Condition on H and set $\lambda = Nq$. Absorbing the bounded factor $2(N+1)/N$ into polynomial coefficients gives the exact expression

$$W = b + aQ(Z) + T, \qquad b = 2(N+1)av = O(a\lambda^2/N), \quad T = \frac{a}{\sqrt{N}}(\lambda L + C_3) + \frac{a\lambda}{N}Q_\Delta. \tag{LF7}$$

All coefficient and moment bounds are independent of H, N, ρ. Let $g = \psi_0 - \mathrm{id}$; it is odd, and g and all its positive-order derivatives vanish in $[-\eta,\eta]$.

First suppose $b \le \eta/8$. The function $g(aQ(Z))$ is even in Z. Its derivatives of orders three through six have bounds $Ca^2(1+\|Z\|^m)$ for a fixed m: since Q is quadratic, every term of derivative order at least three contains at least two differentiated factors carrying a. Gaussian bipartite symmetry makes its Gaussian expectation zero. The proof of Lemma 5.2 therefore gives

$$|\mathbb{E}g(aQ)| \le Ca^2/N.$$

By the mean-value formula and support of g',

$$|g(aQ+b) - g(aQ)| \le Cba^2Q^2.$$

Indeed a nonzero derivative along that segment requires $|aQ| \ge 7\eta/8$. The expectation of this shift is at most $Ca^3\lambda^2/N$.

For either $R = L$ or $R = C_3$, the function $g'(aQ+b)R$ is odd in Z. Its derivatives through order three are bounded by Ca^2 times a fixed polynomial envelope. For the terms with zero or one differentiated factor, this follows from localization:

$$|g'(aQ+b)| \le Ca^2Q^2, \qquad |ag''(aQ+b)| \le Ca^2(1+Q^2);$$

terms with two or more such factors already have a^2. Its Gaussian expectation is zero, so the ordinary comparison gives

$$|\mathbb{E}[g'(aQ+b)R]| \le Ca^2/\sqrt{N}.$$

The linear and cubic part of the first Taylor term in T is thus bounded by $Ca^3(\lambda+1)/N$. Localization bounds the quadratic part by $Ca^3\lambda/N$.

Finally, for $0 \le \theta \le 1$,

$$|g''(aQ+b+\theta T)| \le C(a^2Q^2+T^2).$$

If this derivative is nonzero, $|aQ|+|T| \ge 7\eta/8$; otherwise the inequality is immediate. The uniform polynomial moments and (LF7) yield

$$\mathbb{E}(T^2Q^2) \le Ca^2(\lambda+1)^2/N, \qquad \mathbb{E}T^4 \le Ca^4(\lambda+1)^4/N^2.$$

Taylor's integral second remainder is therefore bounded by

$$C\left(\frac{a^4(\lambda+1)^2}{N}+\frac{a^4(\lambda+1)^4}{N^2}\right).$$

Combining these estimates proves, conditionally on a good value of H,

$$|\mathbb{E}[g(W) \mid H]| \le C\left[\frac{a^2}{N}+\frac{a^3(\lambda+1)^2}{N}+\frac{a^4(\lambda+1)^2}{N}+\frac{a^4(\lambda+1)^4}{N^2}\right]. \tag{LF8}$$

We average without restricting λ to a bounded interval. For each fixed integer $j \ge 1$,

$$\lambda^j\mathbb{E}a^{j+1} \le C_j\mathbb{E}a, \qquad \lambda^j\mathbb{E}a^j \le C_j, \qquad \mathbb{E}a \ge \frac{1}{N+1}. \tag{LF9}$$

For $\lambda < 1$ these follow from $a \le 1$ and Jensen's inequality. For $\lambda \ge 1$, split at $H \ge \lambda/2$, where $\lambda a \le 2$. The other event has probability at most $e^{-c\lambda}$, and Jensen gives $\mathbb{E}a \ge 1/(\lambda+1)$, which absorbs its fixed polynomial factors. Thus every term of (LF8) averages to at most $C\mathbb{E}a/N$. For its last term use $\lambda^4\mathbb{E}a^4 \le C$ and $N^{-2} \le C\mathbb{E}a/N$.

It remains to control $b > \eta/8$. Since $b \le Ca\lambda^2/N$, this entails $\lambda \ge c\sqrt{N}$ and $H+1 \le C\lambda^2/N = Cq\lambda$. Choose q_0 sufficiently small that $Cq_0 \le 1/2$. The exceptional event is then contained in $H < \lambda/2$ and has probability at most $e^{-c'\sqrt{N}}$. Since $|g(W)| \le C(N+1)$, its contribution is $O(Ne^{-c'\sqrt{N}}) \le C/N^2 \le C\mathbb{E}a/N$. Enlarging constants covers the finitely many small N. This proves (LF6). □

A switch valid on the whole state space. Choose a smooth bounded ψ_1 with range $[-\xi, b_B+\xi]$ and equal to the identity on $[-\eta, b_B]$, reducing η if necessary. By (LF1), $B \le Cq$. Choose a fixed $\tau > 0$ so small that $\tau/2 \le q_0$ and $B \le \eta/4$ whenever $q \le 2\tau$. Set

$$\psi_H = \begin{cases} \psi_0, & H \le \tau N, \\ \psi_1, & H > \tau N. \end{cases}$$

We claim, uniformly on all states,

$$|\mathbb{E}[\psi_H(W)-W]| \le C\mathbb{E}[(H+1)^{-1}]/N. \tag{LF10}$$

For $q \le \tau/2$, the switch differs from ψ_0 only on an event of probability $O(e^{-cN})$ by the binomial upper-tail bound. The two saturations have bounded difference, and $\mathbb{E}a \ge 1/(N+1)$, so Lemma 5.3 proves the claim.

For $q \ge \tau/2$, bounded MUB-coordinate concentration and binomial concentration give a common event of probability $1-Ce^{-cN}$ on which W is within any prescribed fixed small distance of B,

for all sufficiently large N. To see uniformity, H/N is then bounded away from zero; $(N+1)/(H+1)$ is uniformly close to $1/q$, and the polynomial V_{bip} is uniformly close to v. If $q \le 2\tau$, both saturations are the identity throughout a fixed neighborhood of $B \in [0, \eta/4]$. If $q > 2\tau$, the probability of $H \le \tau N$ is exponentially small, and the strict margin $\sup B < b_B$ makes ψ_1 the identity throughout a common neighborhood of every such B. Hence $\psi_H(W) = W$ except on an event of probability Ce^{-cN}. The global bound $|\psi_H(W) - W| \le C(N+1)$ proves (LF10); again adjust the constant for finitely many small N.

Proof of Theorem 5.1. Let C be a constant valid in (LF10), and define

$$\delta(H) = \frac{C}{N(H+1)}, \qquad \widehat{R} = \widehat{A} + \psi_H(W) - \delta(H).$$

Equation (LF10) gives $\mathbb{E}[\psi_H(W) - \delta(H)] \le \mathbb{E}W$. Together with (LF2) and (LF5) this implies $\mathbb{E}\widehat{R} \le R = A + B$. Moreover the exact inverse-binomial identity gives

$$q\mathbb{E}\delta(H) = \frac{C[1-(1-q)^{N+1}]}{N(N+1)} \le C/N^2.$$

Using (LF10) once more, with both signs, therefore yields

$$0 \le q(R - \mathbb{E}\widehat{R}) \le K/N^2 \tag{LF11}$$

for a fixed K. The component margins give, deterministically for every data table and all sufficiently large N,

$$\ell_A/2 \le \widehat{R} \le b_A + b_B + \xi = \beta < 1,$$

because $\delta(H) \le C/N$ and $\xi < \ell_A/4$.

On independent fresh copies take one further union-event bit $Q \sim \mathrm{Bernoulli}(q)$ and output $Y = 1 - \varepsilon - Q\widehat{R}$. For $\varepsilon \le (1-\beta)/2$, this is in $[0, 1-\varepsilon]$. Independence and $qR = 1 - \operatorname{Tr}\rho^2$ give

$$\mathbb{E}Y - \operatorname{Tr}\rho^2 = -\varepsilon + q(R - \mathbb{E}\widehat{R}).$$

Choose the integer $N \ge \max\{N_0, \sqrt{K/(2\varepsilon)}\}$. Equation (LF11) puts this error in $[-\varepsilon, \varepsilon]$ uniformly. At a pure state $q = 0$, so $Y = 1 - \varepsilon$ almost surely.

The five tomography blocks cost $5(D+1)N$ copies, the estimator of A costs $3N$, the N union trials cost $3rN$, and the final fresh union trial costs $3r$. Thus the deterministic cap is

$$5(D+1)N + 3N + 3rN + 3r = O_D(\varepsilon^{-1/2}).$$

All measurements are the stated stabilizer MUB or cyclic-event measurements, and all subsequent operations are classical. The unknown projector P occurs only in the analysis. The lower bound of Theorem 6.1 applies to bounded real outputs and hence to the nonnegative subclass. It gives the two asserted tight asymptotic orders. □

Attainment at every positive tolerance. For any $0 < \varepsilon \le 1$, choose $0 < \eta < \min\{\varepsilon, \varepsilon_D\}$ and an estimator Y_η given by Theorem 5.1. Put $a = (1-\varepsilon)/(1-\eta)$ and $Y_\varepsilon = aY_\eta$. Then $0 \le Y_\varepsilon \le 1-\varepsilon$, and, writing $s = \operatorname{Tr}\rho^2$,

$$\mathbb{E}Y_\varepsilon - s = a(\mathbb{E}Y_\eta - s) - (1-a)s \in [-a\eta - (1-a), a\eta] \subseteq [-\varepsilon, \varepsilon].$$

The left endpoint is $-\varepsilon$ because $a(1-\eta) = 1-\varepsilon$. The copy cap remains finite, and pure inputs attain the endpoint exactly. This proves feasibility at every positive tolerance as used in Corollary 1.4.

6 A log-free square-root copy lower bound

A one-sided likelihood-curvature estimate strengthens the lower bound and applies even when negative outputs are permitted. For a padded transcript effect E and Gross phase-point operators A_z,

$$\mathrm{Tr}[E(A_{z_1} \otimes \cdots \otimes A_{z_M})] \geq 0,$$

by stabilizer-projector cone closure and Wigner positivity [6, 10]. Complete effects sum to I, so these formal responses sum to one at every tuple. If F is the expectation operator and $G = (1-\varepsilon)I - F$, then

$$g(z_1, \ldots, z_M) = \mathrm{Tr}[G(A_{z_1} \otimes \cdots \otimes A_{z_M})] \geq 0,$$

because every outcome deficit $1 - \varepsilon - y$ is nonnegative. For Wigner-positive ρ, the identity $\rho = \sum_z W_\rho(z) A_z$ gives $\mathrm{Tr}(G\rho^{\otimes M}) = \mathbb{E}_{z_i \,\mathrm{iid}\, W_\rho} g$. Bounded outputs make every such formal weight finite. This representation covers classical randomization and adaptive stabilizer transcripts and uses no upper bound on g.

Theorem 6.1 (A dimension-independent lower bound)**.** *Let $D = 3^n$ with $n \geq 1$. Suppose an allowed stabilizer protocol has a deterministic copy cap $M \geq 2$, a bounded real unconditional output $Y \leq 1 - \varepsilon$, and*

$$\sup_\rho |\mathbb{E}_\rho Y - \mathrm{Tr}(\rho^2)| \leq \varepsilon.$$

Then

$$\varepsilon \geq \frac{1}{162M^2}.$$

The lower endpoint of the output range may be any finite real number.

Proof. First consider one qutrit. Let A_0 be the Gross phase-point operator at the origin and let P be a pure stabilizer projector whose phase-space line $\mathcal{L}$ contains that point. Thus $A_0 P = P$. For $0 < \delta \leq 1/6$ put

$$\rho_s = (1-\delta)P + \frac{\delta}{3} I + s(A_0 - P), \qquad |s| \leq t := \frac{\delta}{3}.$$

These are density matrices: on P their eigenvalue is $1 - 2\delta/3$, and on its orthogonal complement their eigenvalues are $\delta/3 \pm s$. Their purity deficit is

$$d(s) := 1 - \mathrm{Tr}(\rho_s^2) = d_0 - 2s^2, \qquad d_0 = \frac{2}{3}(2\delta - \delta^2).$$

In Gross's normalization their Wigner laws are

$$\pi_s(0) = \frac{a + 2s}{3}, \qquad \pi_s(z) = \frac{a - s}{3} \quad (z \in \mathcal{L} \setminus \{0\}), \qquad \pi_s(z) = \frac{\delta}{9} \quad (z \notin \mathcal{L}), \quad a = 1 - \frac{2\delta}{3}.$$

All these masses are positive. Set $c = 1 - 4\delta/3$. The logarithmic score $b(z) = \pi_s'(z)/\pi_s(z)$ belongs to $[-1/c, 2/c]$ uniformly on the interval. Hence

$$b^2 \leq \frac{b}{c} + \frac{2}{c^2}.$$

For a phase-point tuple of length M, let $L_s = \prod_{i=1}^M \pi_s(z_i)/\pi_0(z_i)$ and $S = \sum_i b(z_i)$. Since the individual masses are affine in s,

$$\frac{L_s''}{L_s} = S^2 - \sum_i b(z_i)^2 \geq S^2 - \frac{S}{c} - \frac{2M}{c^2} \geq -K, \qquad K = \frac{2M + 1/4}{c^2}.$$

The positive transcript-cone representation of the protocol, padded to its cap, gives a nonnegative formal deficit weight: each outcome has coefficient $1-\varepsilon-Y\geq 0$. Therefore its expected deficit

$$f(s)=1-\varepsilon-\mathbb{E}_{\rho_s}Y$$

satisfies $f''\geq -Kf$. The output bound and bias hypothesis imply $0\leq f(s)\leq d(s)\leq d_0$. The centered second-difference identity now gives

$$f(0)-\frac{f(t)+f(-t)}{2}\leq \frac{Kd_0t^2}{2}.$$

The same difference, evaluated using the uniform bias bound and $d(s)=d_0-2s^2$, is at least $2t^2-2\varepsilon$. Consequently

$$2\varepsilon\geq\left(2-\frac{Kd_0}{2}\right)t^2.$$

Choose $\delta=1/(3M)$. Since $M\geq 2$, we have $c\geq 7/9$ and $2M+1/4\leq 17M/8$. Thus

$$\frac{Kd_0}{2}\leq\frac{2K\delta}{3}\leq\frac{153}{196}<1.$$

It follows that $\varepsilon\geq t^2/2=1/(162M^2)$.

For $D=3^n$, restrict each input to $\rho\otimes P_0^{\otimes(n-1)}$, where P_0 is a fixed pure qutrit stabilizer state. Any allowed register protocol is then simulated with the same number of unknown qutrit copies and free stabilizer ancillas. Purity, the output range and the expectation-bias contract are preserved. The one-qutrit bound therefore applies unchanged. □

Proof of Theorem 1.1. The upper inequality is Theorem 5.1, and every nonnegative protocol is also a bounded-real protocol. For sufficiently small ε, Theorem B.1 excludes caps at most two in either class. Theorem 6.1 gives the lower inequality. The logarithmic limits follow directly from the matching-order inequalities. □

Corollary 6.2 (Exact low-copy regimes for one qutrit)**.** *For positive tolerances,*

$$\mathcal{C}_3(\varepsilon)=3\quad(1/19\leq\varepsilon<1/3),\qquad \mathcal{C}_3^{\mathrm{R}}(\varepsilon)=3\quad(1/27\leq\varepsilon<1/3).$$

Both caps equal zero when $1/3\leq\varepsilon\leq 1$. *Below* $1/19$ *the nonnegative cap is at least four, and below* $1/27$ *the bounded-real cap is at least four.*

Proof. The constant $Y=1-\varepsilon$ uses no input copies and has bias at most ε exactly when $\varepsilon\geq 1/3$. Below that threshold, Theorem B.1 excludes caps at most two. The three-copy frontiers in Theorem 1.2 meet the physical endpoint at $1/19$ in the nonnegative class and at $1/27$ in the bounded-real class: the latter follows by comparing both affine branches with $1-\varepsilon$. The stated values and exclusions follow. □

7 Proof of the three-copy frontiers

We work with a single qutrit. The shared code-parity witness and measurement below serve both output classes.

7.1 The Schubr parity refinement

Let S and T be the symmetric and alternating projectors on three qutrits. Together with the existing cyclic projector J they have ranks $10, 1, 16$ and sum to I. Introduce the stabilizer encoding

$$V|u\rangle = 3^{-1/2} \sum_{t\in\mathbb{F}_3} |t, t+u, t-u\rangle, \qquad W|u\rangle = |-u\rangle.$$

The code projector $R = VV^*$ is the joint eigenvalue-one projector of XXX and ZZZ. In the coordinates of Section 2 it is exactly the branch $c = l = 0$. On this code a transposition acts as logical parity W. Consequently

$$A = V(I+W)V^*/2, \qquad T = V(I-W)V^*/2, \qquad K = S - A$$

are orthogonal projectors, and

$$I = A + K + T + J, \qquad (\operatorname{rank} A, \operatorname{rank} K, \operatorname{rank} T, \operatorname{rank} J) = (2, 8, 1, 16).$$

An orthonormal basis for A is

$$a_0 = (|000\rangle + |111\rangle + |222\rangle)/\sqrt{3}, \qquad a_1 = \sum_{\pi\in S_3} |\pi(012)\rangle/\sqrt{6}.$$

For every operator G in the cone generated by pure stabilizer projectors,

$$\operatorname{Tr}[(A-T)G] \geq 0. \tag{16}$$

Indeed, projection into R and Clifford decoding send a pure stabilizer vector to zero or an unnormalized logical stabilizer vector. The logical parity expectation is 1 or 0 on a computational state. A full-support logical stabilizer has amplitudes $3^{-1/2}\omega^{au^2+bu}$, and its parity expectation is $3^{-1}\sum_u \omega^{-2bu} \in \{0, 1\}$. Positive combinations prove the claim. This is a concrete use of stabilizer phase-point positivity.

Definition 7.1 (Schubr parity refinement)**.** *The Schubr parity refinement is the complete four-outcome measurement on three qutrits with effects*

$$E_1 = A/2, \qquad E_2 = A/2 + T, \qquad E_3 = K, \qquad E_4 = J. \tag{17}$$

It refines the binary cyclic-event measurement $\{I - J, J\}$ through an unsharp logical test in the zero CSS syndrome.

To implement it, retain the coordinate measurement implementing the cyclic effects. Outside the branch $c = l = 0$, classify trivial cyclic phase as outcome 3 and nontrivial phase as outcome 4. The former effect is K: the fixed space of the cycle is $S \oplus T$, whose intersection with R is $A \oplus T$. Inside $c = l = 0$, the remaining register is the decoded logical qutrit. Choose one of its four stabilizer MUBs uniformly. Mark $|0\rangle$ in the computational basis and

$$\phi_k = 3^{-1/2} \sum_u \omega^{ku^2} |u\rangle, \quad k = 0, 1, 2,$$

in the three quadratic bases. Return outcome 1 for the marked state and 2 otherwise. The marked-basis identity is the standard phase-point/MUB relation [5, Section 5, equation (48)], with the origin phase point identified as parity [6, Theorem 6]. Here the four marked projectors sum to $I + W = 2P_{\text{even}}$: their 00 entry is 2, their $11, 22, 12, 21$ entries are 1, and the other entries vanish. Their average is $P_{\text{even}}/2$, proving the first two effects in (17). Every branch receives an outcome. In particular no antisymmetric projective measurement is assumed. Outputs (x, y, z, w) on these outcomes have expectation operator

$$F = \frac{x+y}{2} A + yT + zK + wJ. \tag{18}$$

7.2 An exact MUB third moment

The twelve pure qutrit stabilizers comprise the three computational states and the nine vectors $3^{-1/2}(1, \omega^p, \omega^q)$, $p, q \in \mathbb{F}_3$. Their third-moment spectrum is a special case of [12, Proposition 3 and Table IV]. In the present notation the exact moment is

$$M_3 = \frac{1}{12} \sum_s (|s\rangle\langle s|)^{\otimes 3} = \frac{A}{6} + \frac{K}{12}. \tag{19}$$

For a direct verification, phase summation at degree three retains equal multi-indices and cross pairs of pure cubes. For $z \in \mathbb{C}^3$ it gives

$$\sum_s |\langle s, z\rangle|^6 = \|z\|^6 + \frac{1}{3}|z_0^3 + z_1^3 + z_2^3|^2 + 6|z_0 z_1 z_2|^2 = \|z\|^6 + \langle z^{\otimes 3}, A z^{\otimes 3}\rangle.$$

Both operators are supported on the symmetric subspace, and equality of these polynomial forms implies equality of every coefficient. This proves the qutrit stabilizer moment identity (19).

7.3 All-protocol bounds and attaining outputs

Pad an allowed protocol to three copies and let F be its expectation operator. For nonnegative output in $[0, U]$, both F and $UI - F$ belong to the stabilizer cone by the transcript-cone lemma [10]. Set

$$a = \mathrm{Tr}(AF)/2, \quad b = \mathrm{Tr}(KF)/8, \quad c = \mathrm{Tr}(TF), \quad j = \mathrm{Tr}(JF)/16.$$

These are normalized traces, not presumed sector eigenvalues. Equation (16) applied to $UI - F$ gives $U \geq 2a - c$. Averaging the lower pure-state bias bound using (19), and using the upper bias bound at $I/3$, gives

$$a + 2b \geq 3(1 - \varepsilon), \qquad 2a + 8b + c + 16j \leq 9 + 27\varepsilon.$$

Since $j \geq 0$, it follows that

$$U \geq 2a - c \geq 4a + 8b + 16j - 9 - 27\varepsilon \geq 3 - 39\varepsilon. \tag{20}$$

Pure inputs also give $U \geq 1 - \varepsilon$.

There is a stronger feasibility restriction independent of U. Haar averaging the lower pure-state bias bound gives $\mathrm{Tr}(SF) \geq 10(1 - \varepsilon)$. Positivity gives $\mathrm{Tr}\, F \geq \mathrm{Tr}(SF)$, whereas the maximally mixed state gives $\mathrm{Tr}\, F \leq 9 + 27\varepsilon$. Therefore $\varepsilon \geq 1/37$. This infeasibility statement holds even for unrestricted quantum measurements at this copy cap.

For $1/37 \leq \varepsilon \leq 1/19$, assign the outputs

$$(x, y, z, w) = (3 - 39\varepsilon,\ 37\varepsilon - 1,\ 1 - \varepsilon,\ 0) \tag{21}$$

to (17). They obey $x \geq z \geq y \geq 0$, and (18) becomes

$$F = (1 - \varepsilon)S + (37\varepsilon - 1)T.$$

Writing $s = \mathrm{Tr}\, \rho^2$ and $d = \det \rho$, the symmetric and alternating moments are $\mathrm{Tr}(S\rho^{\otimes 3}) = s + d$ and $\mathrm{Tr}(T\rho^{\otimes 3}) = d$. Thus

$$\mathbb{E}_\rho Y - s = \varepsilon(36d - s) \in [-\varepsilon, \varepsilon],$$

since $s \le 1$, $d \ge 0$, $d \le 1/27$ and $s \ge 1/3$. Pure inputs and $I/3$ attain the opposite endpoints. This proves attainment of (20), including the feasibility floor. At $\varepsilon = 1/37$ the outputs are $(72/37, 0, 36/37, 0)$; the expectation operator is $(36/37)S$, whose smaller operator norm should not be confused with the operational output cap.

For $1/19 \le \varepsilon \le 1$, use $Y = (1-\varepsilon)(1-B)$ for the binary cyclic event. Its error is

$$-\varepsilon s + 2(1-\varepsilon)d.$$

This is at least $-\varepsilon$. Moreover $36d \le 4/3 \le 1 + s$, using $d \le 1/27$ and $s \ge 1/3$. Therefore $2(1-\varepsilon)d \le \varepsilon(1+s)$ whenever $\varepsilon \ge 1/19$, proving the upper error bound. Pure inputs show its bias is exactly ε. This proves the nonnegative part of Theorem 1.2.

Finally (20) forces $\varepsilon \ge 2/39$ for outputs in $[0, 1]$. The attaining outputs at this tolerance are

$$(1, 35/39, 37/39, 0).$$

This improves the binary-event affine value $1/19$ through the coherent measurement refinement.

7.4 The full bounded-real three-copy frontier

We now prove the bounded-real part of Theorem 1.2.

Proof. Retain the sectors $S, T, A, K = S - A, J$, their ranks $(10, 1, 2, 8, 16)$, and the four-outcome measurement (17). The deficit $G = UI - F$ of a bounded-real expectation operator F remains in the stabilizer cone, although F need not be positive. With normalized sector traces

$$a = \mathrm{Tr}(AF)/2, \qquad b = \mathrm{Tr}(KF)/8, \qquad c = \mathrm{Tr}(TF), \qquad j = \mathrm{Tr}(JF)/16,$$

the code-parity witness (16) gives $U \ge 2a - c$. Also $F \le UI$, hence $b \le U$.

We need two finite ensembles. Put

$$|h_{uv}\rangle = X^u Z^v \frac{|1\rangle - |2\rangle}{\sqrt{2}}, \qquad P_{uv} = |h_{uv}\rangle\langle h_{uv}|, \qquad u, v \in \mathbb{F}_3.$$

Their distinct squared overlaps are $1/4$: a nontrivial diagonal Weyl has overlap $-1/2$ with the displayed fiducial, and a nonzero translation has one overlapping support position, of modulus $1/2$. Thus these are the usual nine Hesse SIC vectors. Their moments satisfy

$$\frac{1}{9}\sum P_{uv} = I/3, \qquad \frac{1}{9}\sum P_{uv}^{\otimes 2} = S_2/6, \qquad \frac{1}{9}\sum P_{uv}^{\otimes 3} = K/8, \tag{22}$$

where S_2 is the two-copy symmetric projector. The first identity is Weyl averaging. For the second, the support is in the six-dimensional symmetric space and the squared Hilbert–Schmidt norm is $[9 + 72(1/4)^2]/81 = 1/6$, forcing the normalized projector $S_2/6$. For the third, each $|h_{uv}\rangle^{\otimes 3}$ lies in K. For the fiducial this follows by its zero overlap with the two displayed basis vectors of A; collective Weyls preserve A. The squared norm of the third moment is $[9 + 72(1/4)^3]/81 = 1/8$, forcing $K/8$ on its eight-dimensional support.

Let $\sigma_{uv} = (I - P_{uv})/2$, of purity $1/2$. Expanding their tensor cubes with (22) gives

$$\frac{1}{9}\sum \sigma_{uv}^{\otimes 3} = \frac{3I + \mathrm{SWAP}_{12} + \mathrm{SWAP}_{23} + \mathrm{SWAP}_{13}}{96} - \frac{K}{64} = \frac{A}{16} + \frac{3K}{64} + \frac{J}{32}. \tag{23}$$

Here the transposition class sum acts by $3, -3, 0$ on S, T, J. The pure-input upper bias and rank-two lower bias now imply

$$b \le 1 + \varepsilon, \qquad 4a + 12b + 16j \ge 16 - 32\varepsilon.$$

The upper bias at $I/3$ gives $2a + 8b + c + 16j \le 9 + 27\varepsilon$. Using $c \ge 2a - U$ in this last inequality and subtracting yields

$$U + 4b \ge 7 - 59\varepsilon.$$

The bounds $b \le U$ and $b \le 1 + \varepsilon$ give, respectively, $U \ge (7 - 59\varepsilon)/5$ and $U \ge 3 - 63\varepsilon$. Pure inputs also give $U \ge 1 - \varepsilon$. These arguments use no positivity of F and no invariant-twirl assumption.

For attainment, write $s = \mathrm{Tr}(\rho^2)$ and $e_3 = \det \rho$. We use

$$0 \le e_2 - 9e_3 \le \frac{1}{4}, \qquad \frac{1}{2} \le s + 18e_3 \le 1. \tag{24}$$

The lower bound follows from $e_2/e_3 = 1/x + 1/y + 1/z \ge 9$ when the eigenvalues are positive, with boundary cases by continuity. For the upper bound use Schur's $4e_2 \le 1 + 9e_3$. One direct proof, ordering $x \ge y \ge z$, is

$$1 + 9xyz - 4e_2 = \sum_{\mathrm{cyc}} x(x-y)(x-z) = (x-y)^2(x+y-z) + z(z-x)(z-y) \ge 0.$$

Finally $s + 18e_3 = 1 - 2(e_2 - 9e_3)$.

For $0 \le \varepsilon \le 1/32$, assign to the four outcomes

$$(x, y, z, w) = (3 - 63\varepsilon, 65\varepsilon - 1, 1 + \varepsilon, -3\varepsilon). \tag{25}$$

The largest output is $3 - 63\varepsilon$. Their expectation effect is

$$F = (1 + \varepsilon)S + (65\varepsilon - 1)T - 3\varepsilon J,$$

and the error is

$$\mathbb{E}Y - s = \varepsilon(4s + 72e_3 - 3) \in [-\varepsilon, \varepsilon]$$

by (24).

For $1/32 \le \varepsilon \le 1/27$, put

$$U = \frac{7 - 59\varepsilon}{5}, \qquad k = \frac{9}{5}(1 - 12\varepsilon) \in [1, 9/8], \qquad Y = U - kB,$$

where B is the binary event of effect J. Set $L_k = s + k\,\mathrm{Pr}(B = 1) = k + (1 - k)s - 2ke_3$. Its exact minimum and maximum are

$$\min L_k = \frac{9 + 16k}{27} \quad \text{at } I/3, \qquad \max L_k = \frac{1 + k}{2} \quad \text{at spectrum } (1/2, 1/2, 0).$$

At $k = 1$, this is the interval of $1 - 2e_3$; at $k = 9/8$ it follows from $L_k = 9/8 - (s + 18e_3)/8$ and (24). The same extremizers attain both endpoint bounds, so interpolation in k proves the full interval. Its midpoint is $(45 + 59k)/108 = U$ and its half-range is $(9 - 5k)/108 = \varepsilon$. Thus the displayed output has the required bias and endpoint.

For $\varepsilon \ge 1/27$, take $Y = 1 - \varepsilon - B$. Its error is $2\det\rho - \varepsilon \in [-\varepsilon, \varepsilon]$, and its upper endpoint is $1 - \varepsilon$. The formulas agree at the two breakpoints, proving all branches. □

The nonnegative and bounded-real frontiers optimize different output classes. The former has a positive three-copy bias floor, whereas the latter includes exactly unbiased estimation at upper endpoint three. Both are attained by postprocessing the same four-outcome stabilizer measurement.

8 The unbiased endpoint and its optimal copy cap

The general finite-cap lower bound in [10, Theorem 4.1] is $U \geq D$. Its proof is useful here. With a cap M, a bounded protocol's expectation is a polynomial $P(X) = \mathrm{Tr}(FX^{\otimes M})$ on trace-one Hermitian matrices. If it equals $\mathrm{Tr}(X^2)$ on density matrices, polynomial identity on the open positive-definite set extends equality to that entire affine space. An odd-dimensional Gross phase point A has $\mathrm{Tr}\, A = 1$, $A^2 = I$ and nonnegative formal responses to every allowed stabilizer transcript [6, 10]. These responses sum to one. Thus $P(A) \leq U$, whereas polynomial equality gives $P(A) = \mathrm{Tr}\, A^2 = D$.

8.1 The exact minimum copy cap at zero bias

Proof of Theorem 1.3. For attainment, choose one of the $D + 1$ stabilizer MUBs uniformly and measure all $m = D^2 - 1$ copies in that single chosen basis. Let k_a be its D outcome counts and $C_m = \sum_a k_a(k_a - 1)$. Output

$$Y = \frac{D+1}{m(m-1)} C_m - 1.$$

Conditional multinomial expectation and the MUB second-moment identity give

$$\mathbb{E}Y = \sum_{b=1}^{D+1} \sum_{a=1}^{D} p_{ba}^2 - 1 = \mathrm{Tr}\, \rho^2.$$

The integer collision count is minimized by balanced counts: transferring one count between entries differing by at least two strictly reduces it. At $m = D^2 - 1$ these counts consist of $D - 1$ entries equal to D and one equal to $D - 1$. Their collision count is

$$(D-1)D(D-1) + (D-1)(D-2) = (D-1)(D^2-2) = \frac{m(m-1)}{D+1}.$$

Also $C_m \leq m(m-1)$. Thus every output lies in $[0, D]$. The total copy count is m, since the basis is randomized once rather than each basis being separately sampled.

For necessity, suppose a nonnegative unbiased protocol uses at most $D^2 - 2$ copies and pad it to $N = D^2 - 2$. Its expectation operator F is positive semidefinite and satisfies $\mathrm{Tr}(F\rho^{\otimes N}) = \mathrm{Tr}\, \rho^2$. Average F mathematically over tensor-factor permutations and all conjugations $U^{\otimes N}$. Positivity is preserved; physical implementation of this average is unnecessary. By Schur–Weyl duality, the result is scalar on every sector Π_λ, $\lambda \vdash N$, $\ell(\lambda) \leq D$.

The average must equal the standard normalized transposition-class observable

$$T_N = \binom{N}{2}^{-1} \sum_{i<j} \mathrm{Swap}_{ij}\,.$$

Indeed T_N has the same expectation and both invariances. For eigenvalues x of ρ, sector expectations are $\mathrm{Tr}(\Pi_\lambda \rho^{\otimes N}) = \dim(\mathrm{Sp}_\lambda) s_\lambda(x)$. The difference of the two expectations is a homogeneous degree-N polynomial vanishing on the positive simplex, hence by homogeneity on the positive orthant. Linear independence of the degree-N Schur polynomials in D variables forces equality of every sector coefficient.

The standard Schur–Weyl purity-observable formula [1, Section 7.1, Proposition 7.4 and equation (12)] gives the scalar of T_N on sector λ as

$$\binom{N}{2}^{-1} \sum_{(i,j)\in\lambda} (j - i).$$

Choose $\lambda = (D^{D-2}, D-1, D-1)$, a partition of N with D rows and therefore a nonzero sector. A D-by-D square has total content zero. This partition removes its cells $(D-1, D)$ and (D, D), of contents 1 and 0, respectively. Its total content is -1, so T_N has eigenvalue $-\binom{N}{2}^{-1} < 0$. This contradicts positivity of the averaged F. The same argument covers $D = 2$, where $N = 2$ and $\lambda = (1, 1)$. □

The normalized transposition observable and its Schur spectrum are established ingredients; the conclusion here is the exact nonnegative unbiased copy threshold together with an attaining stabilizer collision estimator. It is distinct from optimizing the output endpoint over arbitrary finite caps.

Proof of Corollary 1.4. At zero bias the phase-point argument at the start of Section 8 gives upper endpoint at least D, and the collision estimator attains D. At positive tolerance a pure input forces every upper endpoint to be at least $1-\varepsilon$. The construction of Section 5, including its scaling argument for every positive tolerance, attains this endpoint with a finite cap and nonnegative outputs. □

A Pure-null three-copy instruments are classical

The cyclic event is implementable as an effect, but that fact does not supply a coarse Lüders instrument. In the single-qutrit case the obstruction is stronger.

Theorem A.1 (Pure-null instrument structure)**.** *Let Λ be an allowed stabilizer subchannel from three single-qutrit inputs to retained qutrit registers. Suppose its effect $E = \Lambda^*(I)$ satisfies*

$$\langle\phi|^{\otimes 3} E |\phi\rangle^{\otimes 3} = 0 \qquad \text{for every pure single-qutrit } \phi.$$

There is a fixed orthonormal basis $v_1, \dots, v_{16}$ of J and positive unnormalized stabilizer mixtures τ_i, with $\operatorname{Tr}\tau_i \le 1$, such that

$$\Lambda(X) = \sum_{i=1}^{16} \langle v_i|X|v_i\rangle \tau_i.$$

Thus every conditional retained output is a stabilizer mixture. In particular, no nonzero such subchannel outputs

$$\sigma(\rho) = \frac{\rho - \rho^3}{1 - \operatorname{Tr}\rho^3} \tag{26}$$

for every identical input $\rho^{\otimes 3}$ on which its event has positive probability.

Proof. We first classify the pure stabilizer rays orthogonal to $S = \mathrm{Sym}^3$. The code-parity inequality $\langle v|(A-T)|v\rangle \ge 0$ and $A \le S$ force their T component to vanish. They therefore lie in $J = I - S - T$.

Use the affine-support quadratic-phase stabilizer normal form [6, 10]. Since $(I + C + C^2)v = 0$, the amplitudes on every nontrivial cyclic orbit sum to zero. Two nonzero amplitudes cannot cancel, because -1 is not a third root of unity. Three nonzero amplitudes cancel precisely when their phases are the three distinct third roots. Consequently the affine support is C-invariant, avoids the fixed diagonal, and its inverse-cycle phase difference never vanishes.

In the coordinates of Section 2, $C(t, u, c) = (t + u, u + c, c)$. The nilpotent map $C - I$ is a single Jordan block, whose invariant linear subspaces are 0, the t-axis, the (t, u)-plane and the whole space. For an invariant affine support $x_0 + W$ one also needs $(C - I)x_0 \in W$. Avoiding the fixed diagonal leaves exactly the lines $u = u_0 \neq 0, c = 0$ and the planes $c = c_0 \neq 0$.

On a line a quadratic phase is $f(t) = \alpha t^2 + dt$, up to a constant. Its inverse-cycle difference is $\alpha(-2tu_0 + u_0^2) - du_0$. It never vanishes exactly when $\alpha = 0$ and $d \neq 0$, giving four rays. On a plane write

$$f(t,u) = \alpha t^2 + \beta tu + \gamma u^2 + dt + eu.$$

The inverse cycle is $(t,u) \mapsto (t - u + c_0, u - c_0)$. The difference's t coefficient is $-2\alpha u + 2\alpha c_0 - \beta c_0$, so nonvanishing for every t, u forces $\alpha = \beta = 0$. The remaining difference is

$$(-d + \gamma c_0)u + dc_0 + \gamma c_0^2 - ec_0.$$

It is nowhere zero exactly when $d = \gamma c_0$ and $e \neq 2\gamma c_0$. These conditions give twelve rays. Different supports are disjoint. Within one plane, different γ give a nonzero t-linear character, and equal γ with different e give a nonzero u-linear character. The analogous character sum handles the two rays on each line. Thus these sixteen rays are mutually orthogonal and exhaust the rank-sixteen space J.

No stabilizer code of rank at least three lies in J. Otherwise, fixing all but one logical qutrit would give a rank-three stabilizer subcode. Its twelve logical stabilizer rays would all be physical stabilizer rays in J, hence members of the sixteen mutually orthogonal rays just classified. Twelve orthogonal rays cannot lie in dimension three.

We use the established stabilizer Choi and code-marginal structure [7]. Form the Choi operator of Λ by feeding maximally entangled stabilizer pairs into its inputs. Stabilizer closure makes this operator a positive combination of pure stabilizer projectors. Its input-reference marginal is proportional to E^T. Positivity and the pure-null premise give $\operatorname{supp} E \subseteq S^\perp$; decomposing E in the stabilizer cone and applying the code-parity inequality gives $\operatorname{supp} E \subseteq J$. The same is true of E^T, since J is real.

The marginal of a bipartite pure stabilizer state is proportional to a stabilizer-code projector: trace its stabilizer-group projector expansion and retain only group elements trivial on the traced side. Every pure Choi summand therefore has input marginal a stabilizer code supported in J. The preceding rank obstruction makes that marginal rank one. Hence the summand is a product of stabilizer projectors. Grouping summands by the sixteen input rays gives the displayed channel formula; conjugation merely permutes those rays. Finally $E \leq I$ gives $\operatorname{Tr} \tau_i \leq 1$. Arbitrary stabilizer ancillas, adaptation, discard and state-independent mixtures preserve this finite-dimensional closed cone argument.

To exclude (26), put $|S_0\rangle = (|1\rangle - |2\rangle)/\sqrt{2}$ and

$$\rho_* = \frac{99}{100}\left(\frac{3}{4}|S_0\rangle\langle S_0| + \frac{1}{4}|0\rangle\langle 0|\right) + \frac{I}{300}.$$

Its eigenvalues are $179/240, 301/1200, 1/300$, so every nonzero positive E accepts $\rho_*^{\otimes 3}$ with positive probability. For the origin phase point $W|x\rangle = |-x\rangle$, direct rational substitution gives

$$\operatorname{Tr}[W\sigma(\rho_*)] = -\frac{262219}{1612806} < 0.$$

Thus $\sigma(\rho_*)$ has a negative Gross Wigner value, whereas every conditional output of Λ is a stabilizer mixture. This proves the final assertion. □

For a hypothetical coarse Lüders-J branch, permutation symmetry gives $\operatorname{Tr}_{23}(J\rho^{\otimes 3}J) = 2(\rho - \rho^3)/3$, whose normalization is (26). The theorem rules out that conditional marginal even on identical inputs. On arbitrary joint inputs, the simpler test $J|012\rangle = (2|012\rangle - |120\rangle - |201\rangle)/3$ already forbids the exact map $X \mapsto JXJ$: its nonzero pure output has unequal computational magnitudes and is nonstabilizer. The instrument classification and conditional-state obstruction apply to three single-qutrit inputs.

A.1 The cost of retaining only pure-null labels

The following separation concerns a fixed seventeen-symbol classical experiment. Each trial measures three fresh single-qutrit copies in the cyclic transcript measurement, retains its sixteen fine J labels, and collapses all other outcomes to one failure symbol. No quantum system or finer failure data is retained.

Proposition A.2 (Pure-null-label sampling cost)**.** *Any statistic based only on this experiment, with state-independent classical randomization and stopping under a deterministic trial cap B, output $0 \le Y \le 1-\varepsilon$, and uniform purity bias at most $0 < \varepsilon \le 1/512$, satisfies*

$$B \ge \frac{1}{1024\varepsilon}.$$

Thus this restricted architecture has copy cost $\Omega(\varepsilon^{-1})$.

Proof. For $0 < t \le 1/4$, $0 \le z \le 1/4$, set $a = 1-t$ and

$$\rho_{t,z} = \operatorname{diag}(a, t/2, t/2) + \frac{t\sqrt{z}}{2}(|1\rangle\langle 2| + |2\rangle\langle 1|).$$

Its eigenvalues are $a, t(1 \pm \sqrt{z})/2 > 0$. Substitution in the line and plane vectors classified above gives four success probabilities $\ell = at^2/4$ and twelve success probabilities

$$h = \frac{a^2 t/2 + at^2/4 + t^3(1-z)/8}{3}.$$

Their sum is $p = 4\ell + 12h = 2at + t^3(1-z)/2$, and the failure probability is $1-p$. Derivatives in z are $0, -t^3/24, t^3/2$, respectively. The single-trial Fisher information is therefore

$$I_t(z) = \frac{t^6}{48h} + \frac{t^6}{4(1-p)} \le t^5,$$

using $h \ge 3t/32$, $p \le 2t \le 1/2$ and $t \le 1/4$.

Pad a stopped procedure to B independent trials and integrate its auxiliary randomness. Its conditional deficit $g = 1 - \varepsilon - Y$ lies in $[0,1]$. Write $f(z) = \mathbb{E}_z g$. Uniform bias and the upper endpoint imply

$$0 \le f(z) \le 1 - \operatorname{Tr}\rho_{t,z}^2 = 2t - (3+z)t^2/2 \le 2t.$$

For the full-record score S_z, finite-alphabet differentiation gives $f'(z) = \mathbb{E}_z(gS_z)$ and $\mathbb{E}_z S_z^2 = BI_t(z)$. Since $g^2 \le g$, Cauchy–Schwarz yields $|f'(z)| \le \sqrt{2B}\, t^3$. Hence

$$|f(1/4) - f(0)| \le \frac{\sqrt{2B}\, t^3}{4}.$$

Purity increases by $t^2/8$ between those endpoints, so bias also requires $f(0) - f(1/4) \ge t^2/8 - 2\varepsilon$. Choose $t = \sqrt{32\varepsilon} \le 1/4$. The last lower bound is $t^2/16$, and comparison gives $B \ge 1/(32t^2) = 1/(1024\varepsilon)$. □

The inverse-bias sampling bound applies to the specified seventeen-symbol experiment. The square-root-power construction combines cyclic events with additional tomography data.

B Coherent-block resource separation

A block protocol partitions its input copies into disjoint blocks. Each block may use a joint stabilizer circuit on its copies and fresh stabilizer ancillas; only classical data, not a quantum system, may pass between blocks. The choices may adapt to all earlier classical outcomes. The total copy cap remains deterministic and finite.

Theorem B.1 (Minimum coherent block size)**.** *For $D = 3^n$ and $0 \le \varepsilon \le 1$, a block protocol whose blocks have at most two copies can have bounded real output $Y \le 1 - \varepsilon$ and bias at most ε if and only if*

$$\varepsilon \ge \frac{1 - D^{-1}}{2}.$$

For every positive tolerance, blocks of at most three copies suffice even with nonnegative outputs. Thus three is the minimum coherent block size for finite-cap boundary attainment at any $0 < \varepsilon < (1 - D^{-1})/2$.

Proof. There are finitely many pure stabilizer rays on one and two D-dimensional registers. For each such vector v on $k \le 2$ copies, $\langle v|\phi^{\otimes k}\rangle$ is a nonzero polynomial in ϕ: otherwise v would be orthogonal to Sym^k, excluded above. Choose a unit ϕ outside the finite union of their zero sets. Its expectation is strictly positive on every nonzero stabilizer-cone effect on either block size. More precisely there is a common $c > 0$ such that

$$\langle \phi^{\otimes k}|E|\phi^{\otimes k}\rangle \ge c \operatorname{Tr} E \qquad (k \le 2)$$

for every such effect, by minimizing over the finite normalized rays.

Since $\rho^{\otimes k} \le I$, every conditional block outcome has probability under ρ at most c^{-1} times its probability under ϕ. Multiplying conditional bounds along a transcript, with a finite bound on the number of input-containing blocks, shows that the full outcome law under any ρ is absolutely continuous with respect to the law under this pure input. Input-free classical randomization is shared and does not change this domination. The pure-state bias and upper bound force $Y = 1 - \varepsilon$ almost surely under ϕ, hence under every ρ. Comparing this constant with $s(I/D) = D^{-1}$ yields the necessary inequality. Conversely that constant suffices whenever the inequality holds. The construction in Section 5 uses single-copy tomography blocks and independent three-copy cyclic blocks, proving sufficiency of block size three at positive tolerance. □

The restriction on quantum memory is essential to the model: coherently retaining systems between blocks can produce a joint measurement on more copies. The theorem permits arbitrarily many two-copy blocks within the finite cap, so increasing only the amount of classical adaptation cannot replace the three-copy event.

Reproducibility and assistance. The accompanying standard-library Python program reconstructs all 27 transcript eigenvectors over $\mathbb{Q}(\omega)$ for one qutrit, checks their cyclic eigenvalues, orthogonality and completeness, and verifies that the 16 nontrivial-eigenvalue projectors sum to J. Its exception-based checks run in ordinary and optimized Python. A second standard-library exact program verifies the sector, MUB, Hesse and four-outcome identities, the frontier breakpoint outputs, and selected collision/content arithmetic. Two further exact programs check inverse-binomial, polarized-moment and likelihood-derivative identities, and the sixteen conditional-label polynomials and full-rank negative-parity witness, respectively. A fifth exact program checks the five-block edge cancellations and reciprocal-binomial controls. These finite checks support the displayed identities; the arbitrary-register, channel and statistical bounds are proved analytically above.

Statements and Declarations

Funding. This research received no external funding.

Competing interests. The author owns Luminary Oddities. Its Research Archive hosts the preprint. The author declares no additional competing interests.

Data and code availability. The results are analytic. Editable source and five exact-arithmetic verification programs accompany this manuscript. An earlier version is available as the author's preprint, edition 1, in the Luminary Research Archive.

Author contributions. Scott Schubr is the sole author and is responsible for the research and final manuscript.

AI assistance. OpenAI's Codex assisted with exploration, proof checking, code, literature comparison and preparation. The author takes responsibility for the mathematics and accompanying files.

References

[1] C. Bădescu, R. O'Donnell and J. Wright, *Quantum state certification*, arXiv:1708.06002 (2017). https://arxiv.org/abs/1708.06002.

[2] L. Bittel, J. Eisert, L. Leone, A. A. Mele and S. F. E. Oliviero, *A complete theory of the Clifford commutant*, Quantum **10** (2026), 2171. https://quantum-journal.org/papers/q-2026-07-22-2171/.

[3] X. Fan and Y. Luan, *Clifford Invariants and Three-Copy Quantum Tests*, arXiv:2610.14758v1 (2026). https://arxiv.org/abs/2610.14758v1.

[4] R. E. Gaunt, *Stein's method for functions of multivariate normal random variables*, Ann. Inst. H. Poincaré Probab. Statist. **56** (2020), 1484–1513. https://doi.org/10.1214/19-AIHP1011.

[5] K. S. Gibbons, M. J. Hoffman and W. K. Wootters, *Discrete phase space based on finite fields*, Phys. Rev. A **70** (2004), 062101. https://doi.org/10.1103/PhysRevA.70.062101.

[6] D. Gross, *Hudson's theorem for finite-dimensional quantum systems*, J. Math. Phys. **47** (2006), 122107. https://arxiv.org/abs/quant-ph/0602001.

[7] A. Heimendahl, M. Heinrich and D. Gross, *The axiomatic and the operational approaches to resource theories of magic do not coincide*, J. Math. Phys. **63** (2022), 112201. https://arxiv.org/abs/2011.11651v3.

[8] E. Hostens, J. Dehaene and B. De Moor, *Stabilizer states and Clifford operations for systems of arbitrary dimensions and modular arithmetic*, Phys. Rev. A **71** (2005), 042315. https://arxiv.org/abs/quant-ph/0408190.

[9] J. Nöller, V. T. Tran, M. Gachechiladze and R. Kueng, *An infinite hierarchy of multi-copy quantum learning tasks*, arXiv:2510.08070v2 (2026). https://arxiv.org/abs/2510.08070v2.

[10] S. Schubr, *Sharp output ranges and bias frontiers for stabilizer purity estimation*, author-published edition 1, Luminary Oddities (2026). https://research.luminaryoddities.com/sharp-output-ranges-bias-frontiers-stabilizer-purity-estimation/v/1.

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

[12] H. Zhu, C. Mao and C. Yi, *Third moments of qudit Clifford orbits and 3-designs based on magic orbits*, arXiv:2410.13575v1 (2024). https://arxiv.org/abs/2410.13575v1.