Purity frontiers and the optimal copy exponent for qutrit stabilizer measurements

Scott Schubr
Luminary Oddities

October 29, 2026

Abstract

We determine exact output frontiers for purity estimation with stabilizer operations and uniform expectation bias. For three copies of one qutrit, a four-outcome measurement attains both nonnegative and bounded-real frontiers. For outputs capped at one minus the bias tolerance, both classes have inverse-bias copy power one half in every fixed qutrit-register dimension; square-root lower bounds and upper bounds with one additional logarithmic factor establish this exponent. We also determine the exact nonnegative unbiased copy threshold of dimension squared minus one and classify three-copy single-qutrit pure-null instruments. The proofs combine phase-space positivity, moment identities, inverse-binomial calibration and one-sided likelihood curvature.

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. We study the exact tradeoff between uniform expectation bias, the upper output endpoint, and a deterministic copy cap when the allowed quantum operations are stabilizer operations.

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

Theorem 1.1 (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 3.1. The exact minimum bias for outputs in $[0, 1]$ is $2/39$.

For $0 < \varepsilon \le 1$, let $\mathcal{C}_D(\varepsilon)$ denote the smallest deterministic input-copy cap among allowed protocols with $0 \le Y \le 1 - \varepsilon$ 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 $\mathcal{C}_D^{\mathrm{R}} \le \mathcal{C}_D$.

Theorem 1.2 (Optimal copy power exponent)**.** *For each fixed $D = 3^n$, positive constants c_D, K_D satisfy, for all sufficiently small $\varepsilon > 0$,*

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

In particular,

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

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.

At the physical boundary $Y \le 1-\varepsilon$, pure inputs force $Y = 1-\varepsilon$ almost surely. The measurement support at pure states is therefore decisive. The exact frontiers follow from a common four-outcome construction and dual witnesses. The optimal copy power combines a constructive upper bound with a lower bound for every allowed adaptive protocol. The zero-bias threshold addresses a separate exact regime.

Relation to previous work. Cyclic trace identities and spectral-moment measurements are established. Fan and Luan [4, 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/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 3.1. Its marked-MUB logical test 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 [13, 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 [3], and the stabilizer normal forms are classical [8].

The author's preceding two-copy purity paper [11] 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 constructions attain that endpoint with a finite cap and determine its power dependence. MUB frames [12], classical Bernoulli factories [10], 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].

Sections 2–3 establish the measurement and exact three-copy frontiers. Section 4 develops a common correction ratio; Sections 5 and 6 prove the matching copy power. Section 7 treats the unbiased endpoint. The appendices provide an explicit numerical budget, the structure and sampling limits of pure-null instruments, a geometric moving witness, a coherent-block separation, and the larger-prime spectral obstruction.

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

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

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 B.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 B.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 [11, 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 Proof of the three-copy frontiers

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

3.1 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,

$$\mathrm{Tr}[(A - T)G] \geq 0. \tag{8}$$

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

$$E_1 = A/2, \qquad E_2 = A/2 + T, \qquad E_3 = K, \qquad E_4 = J. \tag{9}$$

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 (9). Every branch receives an outcome. In particular no antisymmetric projective measurement is assumed. Outputs (x, y, z, w) on these outcomes have expectation operator

$$F = \frac{x + y}{2} A + yT + zK + wJ. \tag{10}$$

3.2 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 [13, Proposition 3 and Table IV]. In the present notation the exact moment is

$$M_3 = \frac{1}{12} \sum_s (|s\rangle\langle s|)^{\otimes 3} = \frac{A}{6} + \frac{K}{12}. \tag{11}$$

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 (11).

3.3 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 [11]. 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 (8) applied to $UI - F$ gives $U \geq 2a - c$. Averaging the lower pure-state bias bound using (11), 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

$$U \geq 2a - c \geq 4a + 8b + 16j - 9 - 27\varepsilon \geq 3 - 39\varepsilon. \tag{12}$$

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

$$(x, y, z, w) = (3 - 39\varepsilon,\ 37\varepsilon - 1,\ 1 - \varepsilon,\ 0) \tag{13}$$

to (9). They obey $x \geq z \geq y \geq 0$, and (10) 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 $s \leq 1$, $d \geq 0$, $d \leq 1/27$ and $s \geq 1/3$. Pure inputs and $I/3$ attain the opposite endpoints. This proves attainment of (12), 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 $1/19 \leq \varepsilon \leq 1$, 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 $36d \leq 4/3 \leq 1 + s$, using $d \leq 1/27$ and $s \geq 1/3$. Therefore $2(1-\varepsilon)d \leq \varepsilon(1+s)$ whenever $\varepsilon \geq 1/19$, proving the upper error bound. Pure inputs show its bias is exactly ε. This proves the nonnegative part of Theorem 1.1.

Finally (12) forces $\varepsilon \geq 2/39$ 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.

3.4 The full bounded-real three-copy frontier

We now prove the bounded-real part of Theorem 1.1.

Proof. Retain the sectors $S, T, A, K = S - A, J$, their ranks $(10, 1, 2, 8, 16)$, and the four-outcome measurement (9). 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 (8) gives $U \geq 2a - c$. Also $F \leq UI$, hence $b \leq U$.

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

$$\frac{1}{9}\sum P_{uv} = I/3, \qquad \frac{1}{9}\sum P_{uv}^{\otimes 2} = S_2/6, \qquad \frac{1}{9}\sum P_{uv}^{\otimes 3} = K/8, \tag{14}$$

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 (14) gives

$$\frac{1}{9}\sum \sigma_{uv}^{\otimes 3} = \frac{3I + \mathrm{SWAP}_{12} + \mathrm{SWAP}_{23} + \mathrm{SWAP}_{13}}{96} - \frac{K}{64} = \frac{A}{16} + \frac{3K}{64} + \frac{J}{32}. \tag{15}$$

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

$$0 \le e_2 - 9e_3 \le \frac{1}{4}, \qquad \frac{1}{2} \le s + 18e_3 \le 1. \tag{16}$$

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

$$(x, y, z, w) = (3 - 63\varepsilon, 65\varepsilon - 1, 1 + \varepsilon, -3\varepsilon). \tag{17}$$

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 (16).

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 (16). 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 $\varepsilon \geq 1/27$, 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.

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 $0 \leq h \leq 1/3$, $1 \leq R_1 \leq 3/2$, and*

$$|R_1(\rho) - R_1(\sigma)| \leq \frac{9\sqrt{D}}{4} \|\rho - \sigma\|_F, \tag{18}$$

$$|p(\rho) - p(\sigma)| \leq 2\sqrt{D} \|\rho - \sigma\|_F. \tag{19}$$

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

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

because $S^2 \geq 2A$ and $SA \geq 3B$. Therefore $h \leq S/2$, 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

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

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

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

then gives (18). 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 $|x^3 - y^3| \leq 3|x - y|$ on $[0, 1]$ proves (19). □

Choose an integer $r \geq 2$ so that

$$b^r < D^{-1}, \qquad b = \frac{D+2}{3D}. \tag{20}$$

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*

$$\frac{1}{r} \leq R_r \leq a_D := \frac{d_*}{1 - b^r} < 1. \tag{21}$$

It is Lipschitz in Frobenius norm with constant

$$L = L_{D,r} = \sqrt{D}\left(\frac{9}{4} + \frac{3}{2}r(r-1)\right). \tag{22}$$

Proof. Since $\mathrm{Tr}(\rho^3) \leq \mathrm{Tr}(\rho^2)$, we have $p \geq 2d/3$. The function $f(d) = 1 - (1 - 2d/3)^r$ is concave on $[0, d_*]$, vanishes at zero, and satisfies $q \geq f(d)$. Thus $q/d \geq f(d_*)/d_*$ for $d > 0$, giving the upper bound in (21). 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 $1 \leq S_r \leq r$, the lower bound follows. Also $|S_r'| \leq r(r-1)/2$. The quotient bound, $R_1 \leq 3/2$, and (18)–(19) give (22). □

5 The copy upper bound

We recall explicitly the complete stabilizer-MUB frame [12]. 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

$$|b, a\rangle = D^{-1/2} \sum_{x \in \mathbb{F}_D} \omega^{\tau(bx^2 + ax)} |x\rangle, \qquad a, b \in \mathbb{F}_D. \tag{23}$$

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

$$\rho = \sum_{b,a} p_{ba} P_{ba} - I, \qquad \sum_{b,a} p_{ba}^2 = 1 + \mathrm{Tr}(\rho^2), \qquad p_{ba} = \mathrm{Tr}(P_{ba}\rho). \tag{24}$$

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

$$\mathbb{E}\|T - \rho\|_F^2 = \frac{1}{m} \sum_b \left(1 - \sum_a p_{ba}^2\right) \leq \frac{D^2 - 1}{Dm}. \tag{25}$$

5.1 A square-root inverse-tolerance copy bound

The following construction improves the tolerance dependence while retaining dimension-dependent constants. All sample sizes below are deterministic maxima.

Theorem 5.1 (Square-root inverse-tolerance upper bound)**.** *For each fixed $D = 3^n$, the physical boundary $1 - \varepsilon$ is attained with nonnegative outputs using*

$$O_D\Big(\varepsilon^{-1/2}\log(1/\varepsilon)\Big)$$

copies as $\varepsilon \downarrow 0$.

Proof. Fix an integer $r \geq 2$ for which the continuous correction ratio of Section 4 satisfies $1/r \leq R \leq a < 1$. Write

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

Then

$$R = \frac{2u}{q} = A(p) + \frac{2v}{q}, \qquad q \asymp_{D,r} u, \qquad 0 \leq v \leq C_D u^2 \leq C_{D,r} q^2. \tag{26}$$

At $p = q = 0$, use the continuous values $A = R = 1/r$. The last estimate follows from the nearest-pure argument below; the comparison $q \asymp u$ also follows from $4u/3 \leq p \leq 2u$ and $p \leq q \leq rp$. Choose fixed

$$\ell = \frac{1}{2r}, \qquad \beta = \frac{1+a}{2}, \qquad \gamma = \min\left\{\frac{1}{2r}, \frac{1-a}{2}\right\} > 0.$$

Every physical R lies at least γ from the endpoints of $[\ell, \beta]$.

An unbiased polarized numerator. Let $N \geq 2$ be an auxiliary accuracy parameter. Independently form three raw unbiased MUB estimates T_1, T_2, T_3 of ρ, each using m samples in each of the $D + 1$ bases. These matrices have trace one and lie in a fixed compact set; no density-matrix projection is performed. Put

$$\widehat{v} = \frac{1 - \mathrm{Tr}(T_1T_2 + T_1T_3 + T_2T_3) + \mathrm{Tr}(T_1T_2T_3 + T_1T_3T_2)}{6}. \tag{27}$$

This is real, bounded in absolute value by C_D, and exactly unbiased for v by independence. With $t = \max_i \|T_i - \rho\|_F$,

$$|\widehat{v} - v| \leq C_D(ut + t^2). \tag{28}$$

Indeed, on the trace-one Hermitian affine space,

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

Choose a rank-one projection P onto a largest-eigenvalue eigenvector of ρ. With $x = \lambda_{\max}(\rho)$ and $S = 1 - x$, one has $u \geq xS \geq S/D$ and $\|\rho - P\|_F \leq \sqrt{2}S \leq \sqrt{2}Du$. Both $v(P)$ and its first differential on trace-zero directions vanish. Bounded second derivatives therefore give $v(\rho) \leq C_D u^2$ and $\|Dv(\rho)\| \leq C_D u$. The linear term of (27) in each input is $Dv(\rho)/3$; its remaining multilinear terms are $O_D(t^2 + t^3) = O_D(t^2)$ on the compact raw-estimate set. This proves (28).

Finite-dimensional Hoeffding bounds on the bounded MUB matrix coordinates give, for every fixed $c > 0$,

$$\Pr\{t > c/\sqrt{N}\} \leq C_D \exp(-c_D c^2 m/N).$$

Thus $m = \lceil K_{D,r} N \log N \rceil$, with a sufficiently large fixed constant, makes this probability at most N^{-8} for all sufficiently large N. The small constant c will be chosen below before $K_{D,r}$.

An exact inverse-binomial identity. On independent copies collect N Bernoulli(q) outcomes, each a union of r cyclic events, and let H be the number of successes. Define

$$\widehat{N} = \frac{N+1}{H+1}, \qquad \mathbb{E}\widehat{N} = \frac{1-(1-q)^{N+1}}{q}. \tag{29}$$

The second identity follows by integrating $\mathbb{E}x^H = (1-q+qx)^N$ over $0 \le x \le 1$. At $q = 0$ it holds by continuity. Deterministically $1 \le \widehat{N} \le N+1$.

There is an event of probability at least $1 - CN^{-8}$ on which, uniformly in the input state,

$$\widehat{N}u \le C_{D,r} \log N, \qquad v\left|\widehat{N} - \frac{1}{q}\right| = o(1). \tag{30}$$

Here and below products involving v/q at $q = 0$ take their continuous value zero. To verify the claim, Hoeffding bounds $|H/N - q|$ by $C\sqrt{\log N/N}$ with failure at most N^{-8}. If $Nq > L_0 \log N$, also impose $H \ge Nq/2$; Chernoff gives the same failure order for a sufficiently large fixed L_0. In this case $\widehat{N} \le C/q$, and

$$q\widehat{N} - 1 = \frac{Nq - H + q - 1}{H+1}$$

gives

$$v|\widehat{N} - 1/q| \le C_{D,r}(|H/N - q| + 1/N) = o(1).$$

If $Nq \le L_0 \log N$, the deterministic bound $\widehat{N} \le N+1$, together with (26), instead gives

$$\widehat{N}u \le C_{D,r} \log N, \qquad v|\widehat{N} - 1/q| \le C_{D,r}((N+1)q^2 + q) \le C_{D,r}(\log N)^2/N = o(1).$$

This proves (30) without dividing by a vanishing probability in a concentration estimate.

A bounded first-success estimate of $A(p)$. One experiment performs at most N independent union trials. At the first successful union trial, return its first individual cyclic-event bit; if all trials fail, return $1/r$. Denote this $[0,1]$-valued output by X. Within a successful union, the first bit succeeds with conditional probability $p/q = A(p)$. Independence between trials gives exactly

$$\mathbb{E}X = A - (A - 1/r)(1-q)^N \le A.$$

The smooth function $A = 1/S_r(p)$ satisfies $0 \le A - 1/r \le C_r p$, hence

$$0 \le A - \mathbb{E}X \le C_r/N, \qquad 0 \le q(A - \mathbb{E}X) \le C_r/N^2. \tag{31}$$

These follow from $q \asymp p$ and the bounds on $p(1-q)^N$ and $qp(1-q)^N$. The identities include $p = q = 0$, when the output is $1/r$.

Average $L = \lceil C_{D,r} \log N \rceil$ independent such experiments to obtain $\widehat{A}$, independently of the tomography and binomial count. Hoeffding makes $\widehat{A}$ within any prescribed fixed positive distance of A with failure probability at most N^{-8}, after increasing the fixed coefficient in L and then taking N sufficiently large. The deterministic cap of these experiments is $3rNL$ copies. Their possible early stopping does not replace this maximum by an expected cost.

Clipping with an exact one-sided repair. Using the independent blocks, put

$$R_{\mathrm{raw}} = \widehat{A} + 2\widehat{v}\widehat{N}.$$

Equations (29) and (31) imply

$$\mathbb{E}R_{\mathrm{raw}} = \mathbb{E}\widehat{A} + \frac{2v}{q}(1-(1-q)^{N+1}) \le R, \tag{32}$$

$$0 \le q(R - \mathbb{E}R_{\mathrm{raw}}) = q(A - \mathbb{E}\widehat{A}) + 2v(1-q)^{N+1} \le C_{D,r}/N^2. \tag{33}$$

The last term is bounded using $v \le C_{D,r}q^2$ and $\sup_{0\le q\le 1} q^2(1-q)^{N+1} = O(N^{-2})$.

On the intersection of the good events already established,

$$|R_{\mathrm{raw}} - R| \le |\widehat{A} - A| + 2\widehat{N}|\widehat{v} - v| + 2v|\widehat{N} - 1/q|.$$

The middle term is at most

$$C_{D,r}\left(\frac{c\log N}{\sqrt{N}} + c^2(1+1/N)\right)$$

by (28), $t \le c/\sqrt{N}$ and (30). Choose c sufficiently small depending on γ, and choose the fixed Hoeffding deviation for $\widehat{A}$ sufficiently small as well. The other terms tend uniformly to zero. For all sufficiently large N, therefore, $R_{\mathrm{raw}} \in [\ell, \beta]$ except on an event of probability at most $C_{D,r}N^{-8}$.

Let R_{clip} be its scalar projection onto $[\ell, \beta]$. Globally $|R_{\mathrm{raw}}| \le C_D(N+1)+1$, so

$$|\mathbb{E}R_{\mathrm{clip}} - \mathbb{E}R_{\mathrm{raw}}| \le C_{D,r}N^{-7}.$$

Choose a deterministic $\delta_N = C_0N^{-7}$ dominating this bound and define $R_{\mathrm{down}} = R_{\mathrm{clip}} - \delta_N$. Then, for sufficiently large N, $R_{\mathrm{down}} \in [\ell/2, \beta]$. Crucially (32) gives $\mathbb{E}R_{\mathrm{down}} \le R$, and (33) gives

$$0 \le q(R - \mathbb{E}R_{\mathrm{down}}) \le C_{D,r}/N^2. \tag{34}$$

Thus clipping is paid before the exact one-sided inequality is invoked.

Finally, on fresh independent copies measure one more Bernoulli(q) union event Q and output

$$Y = 1 - \varepsilon - R_{\mathrm{down}}Q.$$

For $\varepsilon \le (1-\beta)/2$ this lies in $[0, 1-\varepsilon]$. By independence and $qR = 1 - s(\rho)$ its error is

$$\mathbb{E}Y - s(\rho) = -\varepsilon + q(R - \mathbb{E}R_{\mathrm{down}}).$$

Choose $N \ge \max\{N_0, \sqrt{C_{D,r}/(2\varepsilon)}\}$ using (34), where N_0 is a fixed threshold for the preceding estimates. The error belongs to $[-\varepsilon, \varepsilon]$ uniformly, and pure inputs have $Q = 0$, attaining the exact upper endpoint. The deterministic copy cap is at most

$$3(D+1)m + 3rN + 3rNL + 3r = O_{D,r}(N\log N).$$

Every block is independent and every finite data table receives a bounded output; the clipping and fixed scalar correction are classical postprocessing. This proves the stated asymptotic bound. For a larger tolerance $\varepsilon \le 1$, fix a small $0 < \eta \le \varepsilon$ for which the construction above gives Y_η, and output cY_η with $c = (1-\varepsilon)/(1-\eta)$. Its upper endpoint is $1-\varepsilon$, and its bias is at most $c\eta + 1 - c = \varepsilon$, using $0 \le s(\rho) \le 1$. At $\varepsilon = 1$ this is the zero output. □

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.

Theorem 6.1 (Log-free copy lower bound)**.** *Fix $D = 3^n$. Every allowed stabilizer protocol with deterministic copy cap $M \geq 2$, bounded unconditional real output $Y \leq 1 - \varepsilon$, and uniform purity expectation bias at most $0 \leq \varepsilon \leq 1$ satisfies*

$$\varepsilon \geq \frac{c_D}{M^2}, \qquad c_D = \frac{(D-1)^3}{16384 D^6} > 0. \tag{35}$$

No prescribed lower output endpoint is required.

Proof. Use the Gross normalization $X = \sum_z W_X(z) A_z$, where the D^2 phase-point operators have trace one. For every padded stabilizer transcript effect E, the formal response

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

is nonnegative, by stabilizer-projector cone closure and Wigner positivity [6, 11]. Complete transcript effects sum to I, so their formal responses sum to one at every tuple. Write $U = 1 - \varepsilon$, let F be the expectation operator and set $G = UI - F$. Then

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

because every deficit coefficient $U - y$ is nonnegative. For a Wigner-positive state ρ,

$$\mathrm{Tr}(G\rho^{\otimes M}) = \mathbb{E}_{z_i \,\mathrm{iid}\, W_\rho} g.$$

Bounded real outputs make g finite for each fixed protocol, but no uniform upper bound on g will be used.

Choose an incident stabilizer projector $P = |0\rangle\langle 0|$ and phase point A_0 with $A_0 P = P$ and $A_0|x\rangle = |-x\rangle$. Set $H = A_0 - P$, so $\mathrm{Tr}\, H = \mathrm{Tr}(PH) = 0$ and $\mathrm{Tr}\, H^2 = D - 1$. For $0 < \delta \leq 1/2$ and $|s| \leq t = \delta/(2D)$, put

$$\rho_s = (1 - \delta)P + \delta I/D + sH.$$

Its eigenvalues are $1 - \delta + \delta/D$ and $\delta/D \pm s$, the latter each of multiplicity $(D-1)/2$. Thus every ρ_s is physical, and

$$\mathrm{Tr}\, \rho_s^2 = \mathrm{Tr}\, \rho_0^2 + (D-1)s^2, \qquad 1 - \mathrm{Tr}\, \rho_0^2 = (1 - 1/D)(2\delta - \delta^2) \leq 2\delta. \tag{36}$$

Let $\mathcal{Z}_D$ be the D^2-point phase space. If $\mathcal{L}$ is the D-point Wigner support of P containing the origin, its Wigner law is $\pi_\delta + sh$, where

$$\pi_\delta = (1 - \delta)\,\mathrm{unif}(\mathcal{L}) + \delta\,\mathrm{unif}(\mathcal{Z}_D), \qquad h = \mathbf{1}_{\{0\}} - \mathrm{unif}(\mathcal{L}).$$

On $\mathcal{L}$ the unperturbed mass is at least $1/(2D)$ and $|sh| \leq \delta/(2D)$; outside it the law is unchanged. It is therefore nonnegative. Under iid $z_i \sim \pi_\delta$, the centered score $a_i = h(z_i)/\pi_\delta(z_i)$ is bounded by $B = 2D$. The likelihood ratio and formal deficit are

$$L_s = \prod_{i=1}^{M} (1 + sa_i), \qquad f(s) = \mathrm{Tr}(G\rho_s^{\otimes M}) = \mathbb{E}[gL_s].$$

All factors are at least $1 - tB = 1 - \delta \geq 1/2$. Finite-alphabet differentiation is exact. With $b_i = a_i/(1 + sa_i)$, it gives the one-sided inequality

$$L_s'' = L_s \left[\left(\sum_i b_i \right)^2 - \sum_i b_i^2 \right] \geq -4MB^2 L_s.$$

Multiplication by $g \geq 0$ yields $f''(s) \geq -4MB^2 f(s)$.

Write $e(s) = \mathbb{E}_{\rho_s} Y - \operatorname{Tr} \rho_s^2$. Uniform bias gives $e(s) \geq -\varepsilon$, so throughout the physical interval

$$0 \leq f(s) = 1 - \varepsilon - \operatorname{Tr} \rho_s^2 - e(s) \leq 1 - \operatorname{Tr} \rho_s^2 \leq 2\delta.$$

Consequently $f''(s) \geq -8MB^2\delta$. Integrating its centered second difference,

$$f(0) - \frac{f(t) + f(-t)}{2} = -\frac{1}{2} \int_0^t (t - s)[f''(s) + f''(-s)]\, ds \leq 4MB^2\delta t^2.$$

Equation (36) and $|e(s)| \leq \varepsilon$ also give the opposite estimate

$$f(0) - \frac{f(t) + f(-t)}{2} \geq (D - 1)t^2 - 2\varepsilon.$$

Choose $\delta = (D - 1)/(8B^2M)$ and $t = \delta/(2D)$. These choices satisfy all preceding conditions, and $4MB^2\delta = (D - 1)/2$. Hence

$$\varepsilon \geq \frac{D - 1}{4} t^2 = \frac{(D - 1)^3}{1024D^2B^4M^2} = \frac{(D - 1)^3}{16384D^6M^2},$$

proving the claim. The proof uses negative curvature only; an absolute-curvature tail estimate and its entropy logarithm are unnecessary. □

Proof of Theorem 1.2. The upper inequality is Theorem 5.1, and every nonnegative protocol is also a bounded-real protocol. For sufficiently small ε, Theorem D.1 excludes caps zero and one in either class. Theorem 6.1 therefore gives $\mathcal{C}_D^{\mathrm{R}}(\varepsilon) \geq \sqrt{c_D/\varepsilon}$, where c_D is the explicit constant in (35). Relabeling that fixed positive constant proves (1). Taking logarithms and dividing by $\log(1/\varepsilon)$ gives both limits in (2). □

The power exponent is optimal at fixed dimension for both output classes. The necessary and sufficient copy bounds in (1) differ by one logarithmic factor. Nonnegativity of the formal deficit makes the lower bound valid throughout the bounded-real output class. Appendix C gives an explicit formal-input family and a separate geometric extrapolation argument.

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 D.1 excludes caps at most two. The three-copy frontiers in Theorem 1.1 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 The unbiased endpoint and its optimal copy cap

The general finite-cap lower bound in [11, 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, 11]. These responses sum to one. Thus $P(A) \leq U$, whereas polynomial equality gives $P(A) = \mathrm{Tr}\, A^2 = D$.

7.1 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.

A Explicit calibration budgets and a factory comparison

The MUB frame and correction-ratio constants above yield the following directly specified budget. Project T in Frobenius norm onto the compact convex set of density matrices, obtaining $\widehat{\rho}$. Metric projection gives $\|\widehat{\rho} - \rho\|_F \le \|T - \rho\|_F$.

For $0 < \varepsilon < 1$ choose

$$\eta = \min\left\{\varepsilon, \frac{1}{2r}\right\}, \qquad m = \left\lceil \frac{L^2(D^2-1)}{D\eta^2} \right\rceil. \tag{37}$$

Then $\widehat{R} = R_r(\widehat{\rho})$ satisfies

$$|\mathbb{E}\widehat{R} - R_r(\rho)| \le \mathbb{E}|\widehat{R} - R_r(\rho)| \le \eta. \tag{38}$$

On $3r$ fresh copies, independently of this tomography, measure Q. Define

$$Y_\eta = 1 - \eta - (\widehat{R} - \eta)Q. \tag{39}$$

By (21), the coefficient is at least $1/(2r) > 0$. The $Q = 0$ output is $1 - \eta$, and the $Q = 1$ output is $1 - \widehat{R} \ge 1 - a_D > 0$. Independence is essential in the calculation

$$\mathbb{E}Y_\eta - s(\rho) = -\eta + q\big(R_r(\rho) - \mathbb{E}\widehat{R} + \eta\big) \in [-\eta, \eta], \tag{40}$$

where (38) bounds the parenthesis by $[0, 2\eta]$. There are only finitely many frequency tables for fixed m. Thus the spectral projection and ratio evaluation are classical postprocessing, equivalently a finite table of fixed real outputs.

Set $c = (1-\varepsilon)/(1-\eta)$ and output $Z = cY_\eta$. Since $0 \le c \le 1$,

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

Here $c\eta + 1 - c = \varepsilon$. The output lies in $[0, 1-\varepsilon]$ and equals its upper endpoint on every pure input. The copy cap is

$$(D+1)m + 3r. \tag{41}$$

For $\varepsilon = 1$, output zero. Conversely every protocol with bias at most ε has $U \ge 1 - \varepsilon$ on pure states. This proves the positive-bias part of Corollary 1.4.

Classical factory comparison. The qualitative calibration has a classical interpretation. Paes Leme and Schneider [10, Theorem 3.2 and its sufficiency proof] construct multiparameter Bernoulli factories for continuous functions bounded away from zero and one. Their positive series construction supplies finite truncations with a uniform one-sided expectation error. An informationally complete vector of stabilizer measurement probabilities identifies ρ continuously on its compact image. Extend R_r continuously to the whole probability cube, preserving its interval $[1/r, a_D]$, and choose $a_D < A_0 < 1$. Applying that theorem to R_r/A_0 supplies a finite Bernoulli output B_k with expectation $g_k \le R_r/A_0$ and arbitrarily small uniform error. With independent Q, take

$$Y_\eta = 1 - \eta - A_0 Q B_k, \qquad 0 < \eta \le \min\{\varepsilon, (1 - A_0)/2\}, \qquad \sup(R_r - A_0 g_k) \le 2\eta.$$

Its range is nonnegative with upper endpoint $1 - \eta$, and its bias is $-\eta + q(R_r - A_0 g_k) \in [-\eta, \eta]$. The same scaling attains $1 - \varepsilon$. Thus general continuous-ratio attainment follows qualitatively from classical factory machinery. The explicit tomography proof above provides a concrete copy bound by counting the input copies used in each measurement block.

A.1 An explicit single-qutrit budget

For one qutrit, the arbitrary-positive-bias construction has an explicit budget with six event copies. Here $h \le 1/9$, since for positive eigenvalues $e_2/e_3 = 1/x + 1/y + 1/z \ge 9$, with boundary cases by continuity. Thus $R_1 \le 9/8$, and $p \le 16/27$. With two independent events,

$$q = 2p - p^2, \qquad R_2 = \frac{R_1}{2 - p} \in \left[\frac{1}{2}, \frac{243}{304}\right].$$

Indeed $R_1 \le 9/8$ gives $p \ge 8d/9$. Monotonicity of $2p - p^2$ and $d \le 2/3$ yield $q \ge 16d/9 - 64d^2/81 \ge (304/243)d$. Both endpoint bounds are sharp, by continuity at pure states and evaluation at $I/3$. In this dimension R_1 has Frobenius Lipschitz constant $81\sqrt{3}/64$. Since $2 - p \ge 4/3$, the quotient estimate gives the constant

$$\frac{567\sqrt{3}}{256} < 4$$

for R_2. For $0 < \varepsilon < 1/3$, take $m = \lceil 128/(3\varepsilon^2) \rceil$ copies in each of the four MUBs. The output

$$1 - \varepsilon - (R_2(\widehat{\rho}) - \varepsilon)Q$$

has bias at most ε, upper endpoint $1 - \varepsilon$, and lower endpoint at least $61/304$. Its cap is $4m + 6$. For $1/3 \le \varepsilon \le 1$, the constant $1 - \varepsilon$ suffices.

A.2 The inherited fixed-phase-point witness

The all-finite-cap upper-endpoint formula (3) also holds when bounded negative outputs are permitted, with $1 - \varepsilon$ attained for $\varepsilon > 1$ by a constant. For a fixed cap $M \ge 2$, the preceding paper's Chebyshev witness [11, Theorem 4.3] gives $U \ge D - \varepsilon T_M(D)$. Hence any such cap attaining $U = 1 - \varepsilon$, $0 < \varepsilon < 1$, obeys

$$\varepsilon(T_M(D) - 1) \ge D - 1, \qquad M \ge \frac{\operatorname{arcosh}(1 + (D - 1)/\varepsilon)}{\operatorname{arcosh} D}. \tag{42}$$

The construction (37)–(41) supplies an explicit polynomial upper bound, and Theorem 5.1 gives the stronger asymptotic upper bound in the main argument. Equation (42) supplies a logarithmic necessary bound.

B 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 B.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

$$\sigma(\rho) = \frac{\rho - \rho^3}{1 - \operatorname{Tr}\rho^3} \tag{43}$$

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, 11]. 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 $E \le I$ gives $\operatorname{Tr} \tau_i \le 1$. Arbitrary stabilizer ancillas, adaptation, discard and state-independent mixtures preserve this finite-dimensional closed cone argument.

To exclude (43), 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 (43). 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.

B.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 B.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.

C The moving Wigner-positive witness for signed outputs

A formal input approaching a stabilizer pure state gives a polynomial lower bound for boundary attainment. The witness uses the discrete-Wigner positivity of stabilizer effects [6, 11]; the cap-dependent choice of the formal input is essential.

Definition C.1 (Moving Wigner-positive witness)**.** *Fix an incident stabilizer projector P and Gross phase point A_0, with $A_0 P = P$. For a cap $M \ge 2$, set*

$$t_M = \frac{1}{(D-1)M^2}, \qquad X_{t_M} = (1 - t_M)P + t_M A_0.$$

The cap-dependent formal input X_{t_M} together with its affine extrapolation from I/D is the moving Wigner-positive witness. Its positive property is nonnegativity of the discrete Wigner function in Gross's convention [6]; the phase-point operators and that positivity convention are the established ingredients.

Theorem C.2 (Moving Wigner-positive witness bound)**.** *Fix $D = 3^n$ and an integer $M \geq 2$. Suppose an allowed protocol uses at most M input copies, has a bounded real output $Y \leq 1-\varepsilon$, and satisfies*

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

Then

$$\varepsilon \geq \frac{1}{(D-1)(\cosh 2 - 1)M^4}. \tag{44}$$

The conclusion holds for bounded real outputs.

Proof. Let A_0 be the Gross phase-point operator at the origin and let $P = |0\rangle\langle 0|$ be the stabilizer projector whose phase-space support contains that point. In standard computational coordinates $A_0|x\rangle = |-x\rangle$, so

$$A_0^* = A_0, \qquad A_0^2 = I, \qquad \mathrm{Tr}\, A_0 = 1, \qquad A_0 P = P A_0 = P.$$

For $0 < t \leq 1$ define

$$X_t = (1-t)P + tA_0 = P + t(A_0 - P).$$

Its eigenvalues are 1 and $\pm t$, each sign having multiplicity $(D-1)/2$. In particular,

$$\mathrm{Tr}\, X_t = 1, \qquad \mathrm{Tr}\, X_t^2 = 1 + (D-1)t^2. \tag{45}$$

Although X_t is not a density matrix, its discrete Wigner function is a probability distribution: it is the convex combination of the uniform distribution associated with P and the point mass associated with A_0.

Pad the protocol to M inputs and write its expectation as $F(X) = \mathrm{Tr}(KX^{\otimes M})$, where K is its output-weighted expectation operator. Every padded transcript effect belongs to the stabilizer-projector cone. Its Wigner symbol is nonnegative, and $X_t^{\otimes M}$ has a normalized nonnegative Wigner function. The resulting formal transcript weights are therefore nonnegative and sum to one. The output bound consequently gives

$$F(X_t) \leq 1 - \varepsilon. \tag{46}$$

This argument also covers state-independent classical mixtures and adaptive protocols: their effects and output weights are integrated into the same expectation operator.

Consider the affine line

$$\rho_t(\alpha) = (1-\alpha)I/D + \alpha X_t.$$

Its three possible eigenvalues are

$$\frac{1 + (D-1)\alpha}{D}, \qquad \frac{1 - \alpha + Dt\alpha}{D}, \qquad \frac{1 - \alpha - Dt\alpha}{D}.$$

They are simultaneously nonnegative exactly when

$$a := -\frac{1}{D-1} \leq \alpha \leq b := \frac{1}{1+Dt}. \tag{47}$$

For $0 < t \leq 1$ the other eigenvalue constraints are weaker at these endpoints. The affine map from $[a, b]$ to $[-1, 1]$ sends $\alpha = 1$ to

$$s_t = \frac{2 - a - b}{b - a} = \frac{1 + (2D-1)t}{1+t} = 1 + \frac{2(D-1)t}{1+t} > 1. \tag{48}$$

The bias polynomial

$$e_t(\alpha) = F(\rho_t(\alpha)) - \mathrm{Tr}(\rho_t(\alpha)^2)$$

has degree at most M, because $M \geq 2$, and has absolute value at most ε on $[a, b]$. The Chebyshev extremal inequality therefore gives

$$|e_t(1)| \leq \varepsilon T_M(s_t), \qquad T_M(s) = \cosh(M \operatorname{arcosh} s), \quad s \geq 1.$$

For completeness, the extremal inequality follows from the alternating values of T_M at $\cos(k\pi/M)$: if a polynomial bounded by one on $[-1, 1]$ exceeded T_M in absolute value at an exterior point, subtracting a suitable multiple of it from T_M would produce a root between every successive pair of these extrema and one more root at the exterior point, contradicting degree at most M. The case $\varepsilon = 0$ follows directly because the bias polynomial then vanishes on an interval.

By (45) and (46),

$$e_t(1) \leq -\varepsilon - (D-1)t^2.$$

Consequently every $0 < t \leq 1$ supplies the necessary condition

$$\varepsilon(T_M(s_t) - 1) \geq (D-1)t^2. \tag{49}$$

Choose $t = 1/((D-1)M^2)$. Then $s_t \leq 1 + 2/M^2$ and

$$\operatorname{arcosh}(s_t) \leq 2/M,$$

since $\cosh x \geq 1 + x^2/2$. Hence $T_M(s_t) \leq \cosh 2$ and $(D-1)t^2 = 1/((D-1)M^4)$. Substitution into (49) proves (44). □

The geometric bound holds on the bounded-real output domain of Theorem C.2. Its explicit formal-input construction and Chebyshev extrapolation complement the square-root lower bound obtained by one-sided likelihood curvature on that same domain.

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

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

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 $k \leq 2$ 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

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

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

Since $\rho^{\otimes k} \leq I$, 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 Sections 4–5 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.

E A prime-cycle obstruction outside characteristic three

The cyclic measurement above has a genuine characteristic-three restriction. The obstruction concerns exact effects on arbitrary joint inputs, not merely their probabilities on identical tensor powers.

Theorem E.1. *Let $p \geq 5$ be prime, and let C_p cyclically permute p registers of $n \geq 1$ p-level qudits. No stabilizer protocol has the full spectral projectors of C_p as its coarse outcome effects. Nor can such a protocol have the exact binary effects $\{I - \Pi_0(C_p), \Pi_0(C_p)\}$, where $\Pi_0(C_p)$ projects onto the eigenvalue-one subspace.*

Proof. It suffices first to take $n = 1$. On computational labels in $\mathbb{F}_p^p$, the cycle matrix S is similar to one unipotent Jordan block. Indeed, if e_0 is a coordinate vector, then the columns $b_j = (S - I)^{p-1-j} e_0$, $0 \leq j < p$, form a basis and satisfy $(S - I)b_j = b_{j-1}$, with $b_{-1} = 0$. Independence follows from the triangular binomial change of basis from $S^j e_0$. Thus an invertible linear computational change, which is Clifford, conjugates the physical cycle to

$$C_J : (t_0, \ldots, t_{p-1}) \longmapsto (t_0 + t_1, t_1 + t_2, \ldots, t_{p-2} + t_{p-1}, t_{p-1}).$$

Write $\omega = e^{2\pi i/p}$ and $|\chi_k\rangle = p^{-1/2} \sum_a \omega^{ka} |a\rangle$. Consider the stabilizer encoding of a logical qudit

$$V|z\rangle = |\chi_1\rangle \otimes |\chi_{-z}\rangle \otimes |z\rangle \otimes |1\rangle \otimes |0\rangle^{\otimes(p-4)}.$$

This is a deterministic Clifford encoding: prepare the displayed fixed stabilizer states and obtain $|\chi_{-z}\rangle$ from $|\chi_0\rangle$ by a controlled phase $\omega^{-t_1 z}$. On its support,

$$(a, b, z, 1, 0, \ldots) \longmapsto (a + b, b + z, z + 1, 1, 0, \ldots).$$

Putting $a' = a + b$, $b' = b + z$ gives $a - zb = a' - (z + 1)b' + z^2 + z$. Consequently this code is invariant and

$$C_J V = VU, \qquad U|z\rangle = \omega^{z^2 + z} |z + 1\rangle.$$

The logical unitary has the nondegenerate eigenbasis

$$|\psi_r\rangle = p^{-1/2} \sum_{z\in\mathbb{F}_p} \omega^{(z^3-z)/3-rz}|z\rangle, \qquad r \in \mathbb{F}_p, \qquad U|\psi_r\rangle = \omega^r|\psi_r\rangle.$$

Here $1/3$ denotes the field inverse. The eigenvalue equation follows by taking the first difference of the displayed cubic; orthogonality follows from the nonzero linear character in the difference between two exponents. These are Alltop-type cubic vectors, and the occurrence of nonstabilizer eigenbases for order-p Clifford unitaries is established [2, Sections 2–3]. The required nonstabilizerness is also immediate here: a one-qudit stabilizer vector of full computational support has quadratic phase [8], whereas the displayed exponent has nonzero cubic coefficient. A nonzero degree-three polynomial cannot vanish at all $p \geq 5$ field points.

Prepending the deterministic stabilizer encoding to a proposed cyclic spectral measurement would therefore give logical effects $|\psi_r\rangle\langle\psi_r|$. Every allowed outcome effect is a nonnegative combination of pure stabilizer projectors, by the transcript cone lemma [11, Lemma 2.1]. A rank-one member of this cone must itself have a stabilizer range, since positivity forces every summand to lie in that range. This contradicts the preceding nonstabilizerness. The invariant code is reducing, so its compressed spectral projectors are exactly those of U. The same argument using only the eigenvalue-one projector rules out the stated binary effects.

For $n > 1$, fix the remaining $n - 1$ qudits of each register to $|0\rangle$. The cycle acts trivially on those factors, reducing a proposed measurement to the prohibited $n = 1$ case. □

Theorem E.1 concerns the exact cyclic spectral effects on arbitrary joint inputs. Its stabilizer encoding identifies the obstruction to extending this measurement from characteristic three to larger odd primes. The measurement specification is equality of the spectral effects on every joint input state.

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. These finite checks support the displayed identities; the arbitrary-register, channel and statistical bounds are proved analytically above. 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] I. Bengtsson, K. Blanchfield, E. T. Campbell and M. Howard, *Order 3 symmetry in the Clifford hierarchy*, J. Phys. A: Math. Theor. **47** (2014), 455302. https://arxiv.org/abs/1407.2713.

[3] 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/.

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

[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] R. Paes Leme and J. Schneider, *Multiparameter Bernoulli factories*, arXiv:2202.07216v1 (2022). https://arxiv.org/abs/2202.07216v1.

[11] 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.

[12] 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.

[13] 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.