Sharp output ranges and bias frontiers for stabilizer purity estimation

Scott Schubr
Luminary Oddities

October 27, 2026

Abstract

We determine the exact tradeoff between uniform expectation bias and the upper output endpoint for projected-purity estimation by odd-prime stabilizer protocols using at most two identical input copies. If an isotropic Weyl subgroup leaves logical dimension d, the least upper endpoint at bias tolerance ε is $\max\{d-(2d+1)\varepsilon, 1-\varepsilon\}$ for $d>1$, and $1-\varepsilon$ for $d=1$. An explicit common-basis collision statistic attains every branch. In particular, the optimal bias with outputs in $[0,1]$ is $(d-1)/(2d+1)$. The proof combines the two-copy Clifford commutant with positive responses at a nonphysical Wigner phase point. The same response argument gives the sharp upper endpoint d for exactly unbiased estimation at every finite copy cap. At larger caps we derive a continuum of Chebyshev witness bounds, an improved finite-cap constant, and the exact leading asymptotic and maximizing location within this witness family. A matching polynomial extremizer on the whole density-matrix body identifies the role of operational symmetry in the exact two-copy result.

1 Introduction and principal result

Purity is a quadratic function of an unknown density operator. Randomized basis collisions estimate it using simple measurements, while their classical output may extend well above the physical purity interval. This raises a precise question: how much expectation bias is necessary when the permitted classical output range is reduced? We answer this question exactly for two-copy stabilizer protocols and projected purity in odd prime dimension.

Fix an odd prime p, $n \geq 1$, $D = p^n$, and a linear isotropic subspace $H \leq \mathbb{F}_p^{2n}$ of dimension h. With phase-consistent Weyl operators W_u, define

$$\mathcal{T}_H(X) = \frac{1}{|H|}\sum_{u\in H} W_u X W_u^*, \qquad \gamma_H(X) = \mathrm{Tr}(\mathcal{T}_H(X)^2), \qquad d = \frac{D}{|H|} = p^{n-h}. \tag{1}$$

A stabilizer protocol may append stabilizer ancillas, apply joint Clifford operations, perform coherent Weyl eigenspace measurements, discard systems, and adapt its subsequent operations to classical outcomes. Its real output is a bounded unconditional statistic, including every abort outcome. The input consists of identical copies of an arbitrary density matrix. A finite copy cap is fixed independently of that density matrix. Classical randomization and its mixtures are included.

For a protocol $\mathcal{A}$ let $P_{\mathcal{A}}(\rho)$ denote its expected output and $U_{\mathcal{A}}$ an upper bound on every output. Define the uniform bias by

$$\mathrm{bias}(\mathcal{A}) = \sup_{\rho\geq 0,\, \mathrm{Tr}\,\rho=1} |P_{\mathcal{A}}(\rho) - \gamma_H(\rho)|.$$

The optimization below has a constraint only on the upper endpoint; the output may also have negative values. Unrestricted collective measurement of the swap operator is unbiased for purity with outputs ± 1; the endpoint obstruction studied here concerns the specified stabilizer operations.

Theorem 1.1 (Exact two-copy frontier)**.** *For $\varepsilon \geq 0$, among all such protocols consuming at most two copies and having uniform bias at most ε, the infimum upper output endpoint is attained and equals*

$$U_*(\varepsilon) = \begin{cases} \max\{d - (2d+1)\varepsilon,\, 1-\varepsilon\}, & d > 1, \\ 1 - \varepsilon, & d = 1. \end{cases} \tag{2}$$

For $d > 1$ the two branches intersect at $\varepsilon_0 = (d-1)/(2d)$. The minimum uniform bias among at-most-two-copy protocols with output in $[0, 1]$ is exactly

$$\varepsilon_* = \frac{d-1}{2d+1}. \tag{3}$$

For $d = 1$ that minimum is zero.

At a logical qutrit, $d = 3$, the exact minimum in (3) is $2/7$. The attaining statistic takes values $3/7$ and 1 within a common sector, and 0 when the sectors differ. The theorem optimizes over coherent joint measurements and adaptive stabilizer circuits as well as separate basis measurements.

Three copy regimes. Theorem 1.1 determines the full bias–upper-endpoint tradeoff with at most two copies. Theorem 4.1 identifies the exactly unbiased endpoint at every cap of at least two copies. For larger caps, Theorem 4.3 and Corollary 4.4 give approximate-bias witness bounds. The generic polynomial evaluation norm in the final mathematical section explains the additional restriction supplied by the two-copy operational symmetry. Subsequent work for unprojected qutrit registers [8] treats exact three-copy frontiers and the physical-endpoint copy order; the present projected, all-odd-prime result remains the two-copy frontier.

The proof has three ingredients. First, every complete stabilizer transcript has a nonnegative response on tensor products of Gross phase points. Second, simultaneous two-copy Clifford averaging reduces the expectation to an affine function of logical purity. Third, the affine error at formal purity d has the signed evaluation identity

$$e(d) = -d\, e(1/d) + (d+1)e(1). \tag{4}$$

The total variation $2d+1$ of these two weights supplies the lower bound. An affine combination of the usual collision estimator and the sector-match indicator attains it on the full input space.

Relation to established methods. Gross's phase-space formalism supplies the odd-dimensional phase points and stabilizer positivity [1]. Positive-Wigner simulation by Veitch, Ferrie, Gross and Emerson supplies the operational positivity principle [2]; we spell out its effect formulation below. Clifford designs and randomized collisions are established estimation tools [3, 4]. Intentional classical rescaling also has an established role in bias–variance tradeoffs for shadow estimators [5]. Related bias–spread inequalities for quantum error mitigation use distinguishability of noisy physical inputs [6]. Our objective is the exact uniform expectation-bias versus upper-output frontier for the indicated operation class and fixed copy cap. Polynomial extrapolation and signed evaluation representations are classical approximation and optimal-prediction tools [7]. They enter here with an explicit physical interval and operational witness cone. The author's mixed-state stabilizer learner [9] supplies related Weyl conventions and an overlap-learning task; the present question concerns estimation of the quadratic functional (1).

2 Phase points, transcript effects and a collision identity

2.1 Conventions and positive formal responses

For $u = (a, b)$ and $v = (a', b')$ in $\mathbb{F}_p^{2n}$ put $[u, v] = a \cdot b' - b \cdot a'$. Choose Weyls satisfying

$$W_u W_v = \omega^{[u,v]/2} W_{u+v}, \quad W_u^* = W_{-u}, \quad \mathrm{Tr}(W_u) = D\,\mathbf{1}_{u=0}, \qquad \omega = e^{2\pi i/p}.$$

The Gross phase points can be written

$$A_z = D^{-1} \sum_u \omega^{[z,u]} W_u.$$

Changing the sign convention for z relabels this family. The identities used below are

$$A_z = A_z^*, \quad A_z^2 = I, \quad \mathrm{Tr}\, A_z = 1, \quad \mathrm{Tr}(A_z A_w) = D\mathbf{1}_{z=w}. \tag{5}$$

For example A_0 is coordinate parity, $|x\rangle \mapsto |-x\rangle$; Weyl translations give all other phase points. The Wigner function of a trace-one Hermitian operator is

$$\mathcal{W}_X(z) = D^{-1} \mathrm{Tr}(A_z X), \qquad X = \sum_z \mathcal{W}_X(z) A_z.$$

It sums to one. Tensor products of phase points are phase points of the joint system in the product convention.

A pure stabilizer projector is the normalized character sum over a Lagrangian Weyl subgroup. Pairing this character sum with the Fourier sum for A_z shows that $\mathrm{Tr}(A_z|s\rangle\langle s|)$ is either zero or one. Equivalently, its Wigner function is D^{-1} on an affine Lagrangian and zero elsewhere. Thus every nonnegative combination of stabilizer projectors has a nonnegative pairing with every phase point. These familiar identities are the precise part of the formalism used here [1].

Lemma 2.1 (Positive transcript responses)**.** *Let E_y be a complete-transcript POVM effect of an allowed protocol on m input systems, after padding every branch to the copy cap. Then*

$$\mathrm{Tr}\big(E_y(A_{z_1} \otimes \cdots \otimes A_{z_m})\big) \geq 0 \tag{6}$$

for every tuple, including unequal phase-point labels. These quantities sum to one over y. Consequently a trace-one Hermitian X with nonnegative Wigner function gives a probability distribution of formal transcript responses on $X^{\otimes m}$.

Proof. Feed each input port with half of a normalized maximally entangled stabilizer pair, leaving the other halves untouched. Condition on the complete transcript y and trace out the protocol's remaining systems. The reference operator is E_y^{T}/D^m. Stabilizer ancillas, Clifford gates, conditioned Weyl measurements and partial traces preserve the cone generated by stabilizer projectors. This follows branch by branch from the stabilizer projector character formula: a compatible measured character adds its linear constraint, an incompatible one gives zero, and forgetting measured characters forms a positive sum. Transposition conjugates the stabilizer character and again preserves that cone. Hence E_y belongs to it. Applying the nonnegative phase-point pairing on the joint system proves (6). Completeness gives $\sum_y \mathrm{Tr}(E_y A_{z_1} \otimes \cdots \otimes A_{z_m}) = 1$.

A branch using fewer copies is padded by discarding the unused inputs. A discarded formal input has trace one, so the padding preserves its response. For general X, expand each factor in its nonnegative Wigner weights and use (6) for every tuple. Classical mixtures integrate the same nonnegative responses. The argument also applies to measurable transcript effects and bounded-output integrals. □

For a bounded statistic f the operator $F = \sum_y f(y)E_y$ gives

$$P(X) = \mathrm{Tr}(FX^{\otimes m}). \tag{7}$$

It is a real polynomial of degree at most m on the affine space of trace-one Hermitian matrices. Lemma 2.1 gives $P(X) \le U$ on every Wigner-nonnegative formal input when $f \le U$.

2.2 The two-copy Clifford average

Lemma 2.2 (Two-copy commutant)**.** *In logical dimension $d = p^\ell$, averaging any operator F by $C \otimes C$ over the Clifford group yields $aI + b\mathsf{S}$, where S exchanges the two logical systems. Thus the averaged expectation on $\sigma^{\otimes 2}$ is $a + b\,\mathrm{Tr}(\sigma^2)$ for every trace-one σ.*

Proof. Expand F in the basis $W_u \otimes W_v$. Simultaneous Weyl conjugation multiplies a term by a character of $u+v$, so its average vanishes unless $v = -u$. The symplectic Clifford action is transitive on nonzero Weyl labels, and conjugates each phase-consistent W_u to W_{Su} up to a phase that cancels against the phase of W_{-u}. Therefore all surviving nonzero coefficients have one common value. The identity

$$\mathsf{S} = d^{-1} \sum_u W_u \otimes W_u^*$$

then proves the assertion. Finite Clifford averaging suffices; global scalar phases are immaterial. □

Lemma 2.3 (Sector collision estimator)**.** *There is a two-copy allowed experiment with indicators J and K such that*

$$f_0 = (d+1)K - J, \qquad \mathbb{E}f_0 = \gamma_H(\rho), \qquad \mathbb{E}J = s := \sum_z w_z^2 \le 1. \tag{8}$$

Here w_z are the masses of the classical sectors after dephasing by H. For $d > 1$, the possible values of f_0 are $-1, d$ within a matching sector and 0 across different sectors. For $d = 1$ it is the sector collision indicator.

Proof. An isotropic basis can be extended to a symplectic basis. A Clifford then maps H to the Z labels on h coordinates. The map $\mathcal{T}_H$ dephases those coordinates and gives subnormalized logical blocks $\sigma_z = w_z\rho_z$ of dimension d. In particular

$$\gamma_H(\rho) = \sum_z \mathrm{Tr}(\sigma_z^2), \qquad s/d \le \gamma_H(\rho) \le s. \tag{9}$$

Measure the sector on both copies, setting $J = \mathbf{1}_{z=z'}$. Apply the same uniformly random logical Clifford to the two residual systems and measure in the computational basis. Set $K = \mathbf{1}_{z=z',b=b'}$.

Let $R = \sum_b |b,b\rangle\langle b,b|$. Its Clifford average is $\alpha I + \beta\mathsf{S}$ by Lemma 2.2. The equations $\mathrm{Tr}\,R = d$ and $\mathrm{Tr}(\mathsf{S}R) = d$ give $\alpha = \beta = 1/(d+1)$ for $d > 1$. It follows, by homogeneity also for subnormalized blocks, that

$$\mathbb{E}_C \sum_b \langle b|C\sigma_z C^*|b\rangle^2 = \frac{(\mathrm{Tr}\,\sigma_z)^2 + \mathrm{Tr}(\sigma_z^2)}{d+1}.$$

Summing sectors and subtracting $\mathbb{E}J$ proves (8). At $d = 1$ both logical outcomes agree, so $K = J$ and $f_0 = J$ directly. □

3 Proof of the exact two-copy frontier

3.1 Lower bounds

Suppose a protocol has upper output endpoint U and uniform bias at most ε. Restrict its inputs, after the coordinate Clifford above, to one fixed sector $|z\rangle\langle z| \otimes \sigma$, with arbitrary logical density σ. The target is $\operatorname{Tr}\sigma^2$. Pad to two copies and average the protocol over applying the same logical Clifford to both copies. These Cliffords preserve the sector. Averaging preserves the upper endpoint and the uniform bias bound. By Lemma 2.2, its expectation has the form

$$\overline{P}(\sigma) = a + b \operatorname{Tr}\sigma^2.$$

If $e(x) = a + (b-1)x$, the physical maximally mixed and pure logical inputs give

$$|e(1/d)| \le \varepsilon, \qquad |e(1)| \le \varepsilon.$$

At the formal input $|z\rangle\langle z| \otimes A$, where A is a logical phase point, the target is d. This formal input has a nonnegative product Wigner function. Lemma 2.1 gives $\overline{P}(A) \le U$. The affine identity (4) yields

$$U \ge d + e(d) \ge d - (2d+1)\varepsilon. \tag{10}$$

Independently, a physical pure sector input has target one, so $U \ge 1 - \varepsilon$. For $d = 1$ this physical bound is sufficient.

3.2 Attainment on all sectors

For $d > 1$ and $0 \le \varepsilon \le \varepsilon_0 = (d-1)/(2d)$ set

$$b = 1 - \frac{2d\varepsilon}{d-1}, \qquad a = \frac{(d+1)\varepsilon}{d-1}, \qquad f = bf_0 + aJ. \tag{11}$$

For a nonempty block, $\operatorname{Tr}\rho_z^2 \in [1/d, 1]$. The affine error $(b-1)\operatorname{Tr}\rho_z^2 + a$ lies in $[-\varepsilon, \varepsilon]$, taking opposite endpoint values at the two extremal purities. Hence

$$|\mathbb{E}f - \gamma_H(\rho)| = \left| \sum_z w_z^2 \big((b-1)\operatorname{Tr}\rho_z^2 + a \big) \right| \le \varepsilon s \le \varepsilon.$$

Here the w_z^2 weighting pays the full projected problem rather than only a normalized logical input. Since $b \ge 0$, the largest of the values $0, a-b, a+bd$ is

$$a + bd = d - (2d+1)\varepsilon \ge 1 - \varepsilon > 0.$$

For $\varepsilon_0 \le \varepsilon \le 1$ use $f = (1-\varepsilon)J$. By (9), each normalized-sector error is between $-\varepsilon$ and ε, because $1 - \varepsilon - 1/d \le \varepsilon$ is equivalent to $\varepsilon \ge \varepsilon_0$. Thus the global error is at most εs, and the upper endpoint is $1 - \varepsilon$. For $\varepsilon \ge 1$ the constant output $1 - \varepsilon$ has error at most ε for every target in $[0, 1]$. This pays the negative-upper-endpoint regime as well.

At $d = 1$, $\gamma_H = s$ and $(1-\varepsilon)J$ attains the physical bound for $0 \le \varepsilon \le 1$; the same constant construction applies thereafter. This proves (2).

Finally $U \le 1$ forces $\varepsilon \ge (d-1)/(2d+1)$. At that value (11) has $b = 1/(2d+1)$ and $a = (d+1)/(2d+1)$. Its values are $d/(2d+1), 1$ in a common sector and 0 across sectors. They all lie in $[0, 1]$, proving (3) and Theorem 1.1.

4 Every finite copy cap: exact endpoint and stable witnesses

4.1 Sharp exactly unbiased endpoint

Theorem 4.1. *Every allowed finite-capped protocol that is exactly unbiased for γ_H on every density matrix has upper output endpoint at least d. For every cap $m \geq 2$ this endpoint is attained.*

Proof. With cap m, the expectation polynomial (7) agrees with the quadratic polynomial γ_H on all density matrices. Positive-definite density matrices form a relatively open subset of the affine real space of trace-one Hermitian matrices. A polynomial vanishing on that open set vanishes on the whole affine space, as follows by successively restricting to coordinate lines. Therefore equality holds also at $X = |z\rangle\langle z| \otimes A$. There $\gamma_H(X) = d$, and Lemma 2.1 gives $P(X) \leq U$. Lemma 2.3 attains upper endpoint d with two copies; extra copies can be discarded. □

4.2 A continuum of approximate witnesses

Let $d > 1$ and A be a logical phase point. Its eigenvalues are $+1, -1$, with multiplicities $(d+1)/2, (d-1)/2$. Put

$$L_t = (1-t)I_d/d + tA, \qquad X_t = |z\rangle\langle z| \otimes L_t.$$

The eigenvalues of L_t are $[1 + (d-1)t]/d$ and $[1 - (d+1)t]/d$. Thus it is physical precisely on

$$J_d = [-1/(d-1), 1/(d+1)]. \tag{12}$$

For $0 \leq t \leq 1$ it is Wigner-nonnegative because it is a convex combination of I_d/d and A. Its continued target is

$$\gamma_H(X_t) = 1/d + (d - 1/d)t^2.$$

The affine coordinate $s = ((d^2-1)t+1)/d$ maps J_d to $[-1, 1]$ and $t = 1$ to $s = d$. In this coordinate the target is

$$G_d(s) = \frac{ds^2 - 2s + d}{d^2 - 1}. \tag{13}$$

Write T_m for the Chebyshev polynomial, $T_m(\cos\theta) = \cos(m\theta)$.

Lemma 4.2 (Classical extrapolation)**.** *If a real polynomial r has degree at most m and $|r| \leq 1$ on $[-1, 1]$, then $|r(s)| \leq T_m(s)$ for $s \geq 1$.*

Proof. For $m = 0$ the assertion is immediate. If $r(s) > T_m(s)$ at $s > 1$, multiply r by $c = T_m(s)/r(s) \in (0, 1)$. At the $m + 1$ successive alternating extrema of T_m, the polynomial $T_m - cr$ has strictly alternating signs. It therefore has m distinct roots inside $(-1, 1)$ and also its root at s, impossible for degree at most m. Apply the same argument to $-r$. Continuity pays $s = 1$. □

Theorem 4.3 (Finite-cap witness bound)**.** *For cap $m \geq 2$, bias at most ε, and upper endpoint U, one has, for $d > 1$,*

$$U \geq G_d(s) - \varepsilon T_m(s) \qquad (1 \leq s \leq d), \tag{14}$$

$$\varepsilon \geq \sup_{1 \leq s \leq d} \frac{(G_d(s) - U)_+}{T_m(s)}. \tag{15}$$

For $d = 1$ the physical bound is $U \geq 1 - \varepsilon$.

Proof. The expectation error restricted to X_t has degree at most m, since $m \geq 2$, and is bounded by ε on the physical interval (12). Apply Lemma 4.2 in the s coordinate. For $1 \leq s \leq d$ the corresponding t belongs to $[0,1]$, so Lemma 2.1 bounds its expectation above by U. Rearrangement proves both formulas. At $d = 1$ use a pure physical sector. □

The endpoint $s = d$ gives $\varepsilon \geq (d-U)_+/T_m(d)$. An interior point can be stronger. At $d = 3, U = 1, s = 2$, one has $t = 5/8$, $G_3(2) = 11/8$, $T_4(2) = 97$, and $T_8(2) = 18817$. Hence

$$m = 4: \quad \varepsilon \geq \frac{3}{776} > \frac{2}{577}, \qquad m = 8: \quad \varepsilon \geq \frac{3}{150536} > \frac{2}{665857}.$$

The fractions on the right are the respective endpoint-only bounds.

Corollary 4.4 (Optimized finite-cap constant)**.** *For fixed odd $d \geq 3$ and $U = 1$ set*

$$s_* = \frac{1 + (d-1)\sqrt{d+1}}{d}, \quad b_d = \frac{2ds_* - 2}{d^2 - 1}, \quad v_d = \sqrt{s_*^2 - 1}, \quad \alpha_d = \operatorname{arcosh} s_*.$$

For every integer $m \geq 2$ with $v_d/m \leq d - s_*$,

$$\varepsilon \geq \frac{2b_d v_d}{em(1 + e^{-2m\alpha_d})} e^{-m\alpha_d}. \tag{16}$$

For $d = 3$, this gives $\varepsilon \geq 8/[3em3^m(1 + 3^{-2m})]$ for all $m \geq 2$.

Proof. Solving $G_d(s) = 1$ gives s_*, with $1 < s_* < d$ and $b_d = G_d'(s_*) > 0$. Use $s = s_* + v_d/m$ in (15). Convexity gives $G_d(s) - 1 \geq b_d v_d/m$, and concavity of arcosh gives $m \operatorname{arcosh} s \leq m\alpha_d + 1$. The exact identity

$$T_m(s) = \tfrac{1}{2} e^{m \operatorname{arcosh} s}(1 + e^{-2m \operatorname{arcosh} s}) \leq \tfrac{1}{2} e^{m\alpha_d + 1}(1 + e^{-2m\alpha_d})$$

proves the claim. At $d = 3$, $s_* = 5/3$, $b_d = 1$, $v_d = 4/3$ and $\alpha_d = \log 3$; the displacement condition holds for every $m \geq 2$. □

Proposition 4.5 (Sharp asymptotic within the witness family)**.** *For fixed odd $d \geq 3$, let*

$$W_m = \max_{1 \leq s \leq d} \frac{(G_d(s) - 1)_+}{T_m(s)}.$$

With the notation of Corollary 4.4,

$$W_m \sim \frac{2b_d v_d}{em} e^{-m\alpha_d} \quad (m \longrightarrow \infty).$$

Every sequence of maximizers s_m satisfies $m(s_m - s_*) \to v_d$.

Proof. The numerator is positive precisely when $s > s_*$. For bounded $x \geq 0$, put $s = s_* + x/m$. Taylor expansion and the exact hyperbolic-cosine identity give, uniformly on compact x intervals,

$$me^{m\alpha_d} \frac{G_d(s) - 1}{T_m(s)} \longrightarrow 2b_d x e^{-x/v_d}.$$

To justify maximizing this limit, put $\Delta = s - s_*$. On $[s_*, d]$,

$$G_d(s) - 1 \leq C_d \Delta, \qquad \operatorname{arcosh} s - \alpha_d \geq \frac{\Delta}{\sqrt{d^2 - 1}}, \qquad T_m(s) \geq \tfrac{1}{2} e^{m \operatorname{arcosh} s},$$

where C_d is a fixed bound for G'_d on that interval. Hence the normalized witness is bounded by $2C_d x \exp[-x/\sqrt{d^2-1}]$, with $x = m\Delta$. This tends uniformly to zero as $x \to \infty$ and also as $x \downarrow 0$. Evaluation at $x = v_d$ gives a positive limiting value, so maximizing x stay in a fixed compact subinterval of $(0, \infty)$. The limiting function has its unique maximum at $x = v_d$, of value $2b_d v_d/e$. Compact uniform convergence proves both assertions. □

These are sharp constants and locations for the displayed one-dimensional witness envelope. Their proof uses elementary endpoint asymptotics; an attaining higher-copy measurement protocol is not supplied by that scalar optimization.

5 Exact generic polynomial evaluation norm

The smaller factor $2d+1$ in Theorem 1.1 uses an operational constraint beyond polynomial degree and physical bias. The following equality identifies that distinction.

Proposition 5.1. *Let $X = |z\rangle\langle z| \otimes A$ and let $\mathcal{K}$ be the full body of D-dimensional density matrices. For real polynomials of degree at most m on the trace-one Hermitian affine space,*

$$\sup\{|R(X)| : \sup_{\rho \in \mathcal{K}} |R(\rho)| \le 1\} = T_m(d).$$

For $d = 1$ this norm is one.

Proof. The physical line (12) and Lemma 4.2 give the upper bound. For the reverse bound let $B = |z\rangle\langle z| \otimes A$, extended by zero on other sectors. It has operator norm one, so $|\operatorname{Tr}(B\rho)| \le 1$ for every physical ρ. The polynomial $R_*(\rho) = T_m(\operatorname{Tr}(B\rho))$ therefore has physical norm at most one, while $\operatorname{Tr}(BX) = \operatorname{Tr} A^2 = d$. This attains $T_m(d)$. At $d = 1$ the formal input is physical, and constants attain one. □

For $d = 3$, this polynomial extremizer gives values 17 at $m = 2$ and 577 at $m = 4$. At two copies, the Clifford-invariant expectation is affine in purity and its error obeys (4); this additional structure gives the exact frontier. The general evaluation-norm viewpoint is consistent with signed minimum-total-variation prediction representations [7].

6 Reproducible demonstrations

The accompanying standard-library Python verifier checks the exact rational frontier branches, their intersection and the globally projected sector weighting. It also reconstructs the one-qutrit common-basis collision operator over $\mathbb{Q}(\omega)$, checks phase-point responses, and verifies the interior-witness fractions and Chebyshev recurrence. Acceptance gates raise exceptions and remain active under Python's optimized mode. The reader instructions give both commands and identify which finite calculations support each analytic statement. The analytic proofs above supply the all-dimension and all-protocol quantifiers.

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

References

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

[2] V. Veitch, C. Ferrie, D. Gross and J. Emerson, Negative quasi-probability as a resource for quantum computation, *New J. Phys.* **14** (2012), 113011. arXiv:1201.1256.

[3] D. Gross, K. Audenaert and J. Eisert, Evenly distributed unitaries: on the structure of unitary designs, *J. Math. Phys.* **48** (2007), 052104. arXiv:quant-ph/0611002.

[4] J. Lee and C. Wadhwa, Optimal purity estimation with incoherent measurements (2026). arXiv:2609.40248.

[5] Z. Cai, A. Chapman, H. Jnane and B. Koczor, Biased estimator channels for classical shadows, *Phys. Rev. A* **111** (2025), L030402. arXiv:2402.09511.

[6] R. Takagi, S. Endo, S. Minagawa and M. Gu, Fundamental limits of quantum error mitigation, *npj Quantum Information* **8** (2022), 114. doi:10.1038/s41534-022-00618-z.

[7] L. Bos, A characterization of optimal prediction measures via ℓ_1 minimization (2023). arXiv:2312.03091.

[8] S. Schubr, Purity frontiers and optimal copy complexity for qutrit stabilizer measurements, Luminary Research Archive, edition 2 (2026). Archive edition 2.

[9] S. Schubr, Mixed-state escape bounds and proper agnostic learning of stabilizer states in odd-prime dimensions, Luminary Research Archive, edition 3 (2026). Luminary Oddities research archive.