Coherence thresholds for optimal-rate stabilizer purity testing

Scott Schubr
Luminary Oddities

October 30, 2026

Abstract

A purity test is exactly one-sided if it can certify mixedness but never does so on a pure state. We ask how many copies must be processed jointly when the available operations are restricted to stabilizer circuits. For every odd prime p and every fixed number of p-level registers, we prove that this minimum block size is p for the optimal total-copy order $\Theta(\epsilon^{-1}\log(1/\delta))$. Here ϵ is the minimum linear impurity to be detected and δ is the allowed probability of missing it. States may be entangled across registers; only classical information passes between blocks. For the upper bound, we construct a p-copy measurement on quadratic-polynomial supports and combine it with a support theorem for mutually unbiased bases. Conditioning on local measurement outcomes extends the construction to entangled inputs while retaining a click probability proportional to impurity. For the lower bound, smaller blocks cannot detect the first-order contribution of a small admixture near a pure stabilizer state. We turn this local obstruction into an $\Omega(\epsilon^{-2}\log(1/\delta))$ bound, even with classical adaptation. In dimensions five and seven, one fixed local circuit suffices, and separate support arguments exclude every shorter universal exact test. Sensitivity formulas and a global Clifford-code construction further quantify the measurement resources.

Keywords. Purity testing; stabilizer operations; coherent blocks; mutually unbiased bases; projective designs; affine codes.
2020 Mathematics Subject Classification. 81P15, 81P45, 94B05.

1 Coherence and the optimal purity-testing rate

An exact one-sided purity test certifies mixedness by a click that has probability zero on every pure state, including nonstabilizer states. With unrestricted operations, the swap test gives a two-copy signal proportional to linear impurity. When operations are restricted to stabilizer ancillas, Clifford circuits and stabilizer measurements, the minimum coherence needed to retain the optimal statistical rate is a different question.

Let $H_{p,n} = (\mathbb{C}^p)^{\otimes n}$, where p is an odd prime. A coherent block is a group of fresh copies processed jointly by a finite stabilizer circuit, with classical feedforward allowed inside the block. Between blocks only classical information is retained. The total number of copies has a deterministic finite cap, and the unknown state may be entangled across all n registers. We measure mixedness by the linear impurity $\mathcal{I}(\rho) = 1 - \operatorname{Tr}\rho^2$. The testing task is to distinguish pure states from states with $\mathcal{I}(\rho) \geq \epsilon$, never falsely declaring a pure state mixed and missing a mixed alternative with probability at most δ.

Call a measurement effect *all-pure-null* when its click probability is zero on every pure input. Such a click certifies mixedness. The theorem gives the optimal coherence threshold.

We introduce the name *Schubr's Prime-Characteristic Coherence Law* for the optimal-rate coherence threshold in the following theorem.

Theorem 1.1 (Schubr's Prime-Characteristic Coherence Law)**.** *For every odd prime p and $n \geq 1$:*

1. *A physical p-copy stabilizer test has an all-pure-null click effect with*

$$c_{p,n}\mathcal{I}(\rho) \leq q_{p,n}(\rho) \leq p\mathcal{I}(\rho)$$

for all density matrices ρ, where $c_{p,n} > 0$. A product-conditioning implementation uses at most $n(p+1)^n$ circuit settings; only one selected site requires a joint p-copy operation.

2. *For each fixed $2 \leq k < p$, every exactly one-sided protocol using blocks of at most k copies, with arbitrary classical adaptation between blocks, requires*

$$\Omega_k\big(\epsilon^{-2}\log(1/\delta)\big)$$

total copies in the small-ϵ regime to distinguish pure states from states with $\mathcal{I}(\rho) \geq \epsilon$ with miss probability at most δ. A one-copy protocol has no conclusive mixedness event.

Thus p is exactly the minimum coherent block size permitting a state-uniform linear impurity signal and the optimal fixed-register testing order.

Corollary 1.2 (Optimal total-copy order)**.** *For fixed (p, n) and $0 < \epsilon, \delta \leq 1/2$, p-copy stabilizer blocks attain*

$$\Theta_{p,n}\big(\epsilon^{-1}\log(1/\delta)\big)$$

total copies, with zero false mixedness declarations and miss probability at most δ. The lower order holds even for unrestricted collective measurements with the same exact one-sided guarantee.

The constants are uniform in the state for fixed (p, n). Theorem 1.1 determines the optimal *rate* threshold in every odd characteristic. In characteristics three, five and seven we additionally determine the universal-existence threshold by excluding every shorter exact test.

Two implementations. The general construction measures one fixed rank-(p^2-1) local projector in $p+1$ frames associated with mutually unbiased bases (MUBs), choosing the frame uniformly. In dimensions five and seven a different rank-$p(p+1)$ projector needs only one fixed local circuit, reducing the product implementation to $n(p+1)^{n-1}$ settings. Its finite geometry is proved separately. A global Clifford-code alternative trades across-register operations for an explicit polynomial dependence on the total dimension.

Why the threshold is p. Our p-copy detector responds linearly to a small admixture in every orthogonal direction; compactness makes the impurity bound uniform. Around a pure stabilizer input, however, a cancellation constraint forces every all-pure-null stabilizer effect on fewer than p copies to miss all terms containing just one orthogonal factor. Its response can therefore begin only at quadratic order. The adaptive lower bound shows that choosing later blocks from earlier classical outcomes cannot remove this loss.

Predecessors and contribution. The qutrit purity-frontier paper [9] supplies the all-pure-null resource viewpoint, a characteristic-three coherent-block separation and the generic-pure absolute-continuity method for finite classical transcripts. Its larger-prime cyclic obstruction concerns full cyclic spectral projectors, whereas our construction measures particular positive subspaces. Stabilizer-state testing [6] has a different null family. Restrictions on quantum memory also produce established purity-testing separations [1]; here we fix the register dimension and determine

a characteristic-dependent block threshold for the small-impurity rate. Unambiguous discrimination of a pure magic state and its orthogonal complement [10] requires a conclusive pure-state declaration as well, unlike the mixedness certificate here.

The general upper proof imports Li's complete-MUB support theorem [8, Theorem 4.1]. Our discriminant-phase stabilizer measurement turns its support-selection consequence into tangent detection and an exact one-sided protocol. The lower proof uses the classical affine-support normal form and cyclotomic minimal polynomial: below the characteristic, a single-excitation histogram has too few terms to cancel. Its global contact consequence and adaptive survival estimate give the matching rate threshold. MUB phase-point geometry [3], projective designs [7] and code-matroid counting [5] enter the specialized implementations and lower obstructions. The independent-pair localization has entanglement-projection ancestry; see Cai, Fu and Shao [2, Theorem 4 and discussion]. We use a finite-design version that preserves the operation model.

Organization. Section 2 gives the effect model and tangent criterion. Section 3 constructs the general local test, and Section 4 transfers it to entangled registers. Section 5 proves the contact barrier and exact adaptive lower bound. Section 6 gives the one-frame improvements in dimensions five and seven, followed by sensitivity calibration and a quantitative code lift. Appendices A–D prove the specialized support and geometry statements; Appendix E specifies their exact finite inputs. Appendix F gives the support-selection lemma in a form adapted from Li's proof.

2 Effects and a quantitative tangent criterion

Tangent completeness is the local criterion that will give a uniform impurity bound.

Fix an odd prime p, let $\omega = e^{2\pi i/p}$, and use the computational Weyls $X|x\rangle = |x+1\rangle$, $Z|x\rangle = \omega^x|x\rangle$. A stabilizer-effect cone is the set of finite nonnegative sums of normalized stabilizer rank-one projectors. Outcome effects of stabilizer protocols belong to this cone; stabilizer postselection preserves it. The lower bounds below apply to this entire cone, hence to the allowed protocols.

Definition 2.1. An effect $0 \le E \le I$ on k copies is *all-pure-null* if $\langle \psi^{\otimes k}, E\psi^{\otimes k}\rangle = 0$ for every ψ. An all-pure-null effect is *tangent-complete* if, for all nonzero independent ψ, χ,

$$\mathcal{L}_E(\psi,\chi) = \sum_{s=1}^{k} \langle \psi^{\otimes(s-1)} \otimes \chi \otimes \psi^{\otimes(k-s)}, E(\psi^{\otimes(s-1)} \otimes \chi \otimes \psi^{\otimes(k-s)})\rangle > 0.$$

A finite pure-null ray family is tangent-complete when its uniformly randomized projector effect has this property.

The sum tests a replacement of one copy by χ. Its positivity gives the term linear in a small orthogonal admixture.

Positive weighting or grouping of fine effects preserves a strictly positive response whenever the required rays remain present. Every projector on a stabilizer ray is testable by a Clifford change of basis and computational measurements. A nonnegative cone decomposition alone need not specify an efficient physical measurement; the circuits below do so explicitly.

Lemma 2.2 (Tangent criterion)**.** *Let E be tangent-complete on k copies of $\mathbb{C}^D$. There is $L_0 > 0$ such that*

$$\frac{L_0}{2D^{k-1}}\mathcal{I}(\rho) \le \operatorname{Tr}(E\rho^{\otimes k}) \le k\mathcal{I}(\rho), \qquad \mathcal{I}(\rho) = 1 - \operatorname{Tr}\rho^2.$$

Proof. For unit ψ and density matrix σ supported on $\psi^\perp$, sum the exactly-one-σ expectations. This is continuous and strictly positive, by tangent completeness and linearity in σ. Its parameter space is compact, so its minimum L_0 is positive.

For mixed ρ take a top eigenvector and write $\rho = (1-t)|\psi\rangle\langle\psi| + t\sigma$, where $1 - t \geq 1/D$ and $\sigma\psi = 0$. Every tensor-expansion term has nonnegative expectation against E. Thus

$$\mathrm{Tr}(E\rho^{\otimes k}) \geq t(1-t)^{k-1}L_0 \geq tD^{-(k-1)}L_0.$$

Since $\mathrm{Tr}\,\rho^2 \leq \lambda_{\max}(\rho)$ and $\mathrm{Tr}\,\rho^2 \geq (1-t)^2$, one has $t \leq \mathcal{I}(\rho) \leq 2t$. For the upper bound, the all-ψ term vanishes and $E \leq I$, giving $\mathrm{Tr}(E\rho^{\otimes k}) \leq 1 - (1-t)^k \leq kt \leq k\mathcal{I}(\rho)$. Pure states follow directly. □

3 A discriminant-phase test in every odd prime dimension

The line rays below handle pure vectors with full support; the quadratic rays handle a single zero coordinate. The MUB support theorem ensures that some frame falls into one of these cases.

Fix an odd prime p, put $\omega = e^{2\pi i/p}$, and let $\mathcal{M}_p = \{I\} \cup \{Q_a F : a \in \mathbb{F}_p\}$ be the standard MUB maps, with $F|b\rangle = p^{-1/2}\sum_x \omega^{bx}|x\rangle$ and $Q_a|x\rangle = \omega^{ax^2}|x\rangle$. A ket indexed by $(y_j)_{j\in\mathbb{F}_p}$ denotes the computational word on the p copies, ordered by j. All sums and polynomial coefficients below are over $\mathbb{F}_p$. Define

$$h_\ell = p^{-1/2}\sum_t \omega^{\ell t}|(j+t)_{j\in\mathbb{F}_p}\rangle, \qquad \ell \neq 0, \tag{1}$$

$$g_{\ell,q} = p^{-1}\sum_{b,c} \omega^{\ell b + q(b^2-4c)}|(j^2+bj+c)_{j\in\mathbb{F}_p}\rangle, \qquad \ell \neq 0, \quad q \in \mathbb{F}_p. \tag{2}$$

Theorem 3.1 (A local p-copy detector)**.** *The displayed rays are orthonormal all-pure-null stabilizer rays. Their rank-(p^2-1) projector R_p is implemented by one adaptive stabilizer circuit. The physical averaged effect*

$$E_p = \frac{1}{p+1}\sum_{M\in\mathcal{M}_p} M^{\otimes p} R_p (M^*)^{\otimes p}$$

is tangent-complete.

Proof. Pure-nullness and implementation. The affine supports and quadratic phases give stabilizer rays. Translating the copy index $j \mapsto j+1$ changes b by two and preserves $b^2 - 4c$, so the plane rays have a nontrivial cyclic character when $\ell \neq 0$. The line rays likewise have a nontrivial character. Identical pure tensors are invariant under that cycle; all the rays are therefore pure-null.

Decode an input word by its unique interpolation polynomial:

$$U|y\rangle = |c_{p-1}, \ldots, c_0\rangle, \qquad y_j = \sum_{r=0}^{p-1} c_r j^r.$$

This invertible linear computational map is Clifford. Lines become $c_r = 0$ for $r \geq 2$, $c_1 = 1$ and the Fourier state $\chi_\ell = F|\ell\rangle$ at c_0. Planes become $c_r = 0$ for $r \geq 3$, $c_2 = 1$ and

$$Q_q\chi_\ell \text{ at } c_1, \qquad \chi_{-4q} \text{ at } c_0.$$

The final Fourier label identifies q, and the remaining label identifies ℓ, proving orthogonality; line and plane supports have different c_2. For $p = 3$ there are simply no coefficients above c_2.

A circuit first decodes and tests the higher coefficients. In the line branch $c_2 = 0$ it requires $c_1 = 1$ and a nonzero Fourier c_0 label. In the plane branch $c_2 = 1$ it measures the Fourier c_0 label β, sets $q = -\beta/4$, and measures c_1 in the Q_q-Fourier basis, accepting its nonzero labels. Other outcomes give no click. Each accepted fine effect is exactly a displayed ray projector. This realizes R_p by an adaptive stabilizer circuit that records the fine outcomes and groups their acceptance labels.

Detection with no zero coordinates. We next prove the tangent claim in a fixed MUB frame. If ψ has full support in a frame, single-excitation overlaps with h_ℓ are nonzero common multiples of the nonconstant Fourier coefficients of χ_x/ψ_x. They all vanish only when χ is parallel to ψ.

Detection with one zero coordinate. Suppose instead that ψ has exactly one zero. Common X shifts only the constant polynomial coefficient and changes the displayed rays by phases. We may therefore move the missing label to zero without adding a measurement setting. If $\chi_0 \neq 0$, each one-excitation line term has a unique missing-coordinate contribution, giving detection. Assume $\chi_0 = 0$, put $S = \mathbb{F}_p \setminus \{0\}$, and write $r_x = \chi_x/\psi_x$ on S, with $r_0 = 0$.

Set $d = c - b^2/4$ and let QR_0 denote zero together with the nonzero quadratic residues. The product of all pure amplitudes in a plane word is

$$P_d = \psi_d \prod_{t \in \mathrm{QR}_0 \setminus \{0\}} \psi_{d+t}^2.$$

It is independent of b. The first-slot excitation amplitude of $g_{\ell,q}$, up to a fixed nonzero normalization, is

$$\sum_d \omega^{4qd} P_d \sum_b \omega^{-\ell b} r_{d+b^2/4}.$$

This identity is valid also when $P_d = 0$: since $\chi_0 = 0$, every original term involving a missing coordinate vanishes. No division by a zero pure amplitude is used.

If these overlaps vanish for every q and $\ell \neq 0$, Fourier inversion in q gives

$$P_d \sum_b \omega^{-\ell b} r_{d+b^2/4} = 0.$$

Whenever $I_d = d + \mathrm{QR}_0$ avoids zero, $P_d \neq 0$, so inversion in the nonconstant b modes forces r to be constant on I_d. Each valid image has $(p+1)/2$ elements in the $(p-1)$-element set S, hence any two intersect. Their union is S: the allowed d are $-\mathrm{NQR}$, where NQR is the nonsquare set, so the union is $\mathrm{QR}_0 - \mathrm{NQR}$. If $x \neq 0$ were absent, $x + \mathrm{NQR}$ would be disjoint from QR_0 and, by cardinality, equal to NQR. A nonzero translation generates $\mathbb{F}_p$, contradicting that this set is nonempty and proper. Thus r is constant on S. Together with $\chi_0 = 0$, this makes χ parallel to ψ. Every independent direction is detected in the one-zero case as well.

Selecting a suitable frame. For $p \geq 5$, Li's complete-MUB support theorem [8, Theorem 4.1] states

$$\sum_{M \in \mathcal{M}_p} |\operatorname{supp}(M^* \psi)| \geq p^2 - p + 2.$$

Thus some basis has support at least $p - 1$, since $(p+1)(p-2) = p^2 - p - 2$. For $p = 3$ the same support-selection conclusion follows directly: a vector cannot be a basis ray in two mutually unbiased bases. In each case the full-support or one-zero argument supplies a positive tangent term in one frame. Averaging the finite physical family preserves positivity and all-pure-nullness, proving the theorem. □

Appendix F gives a direct derivation of the needed support-selection consequence, explicitly adapted from Li's cyclotomic-jet method. Li's support theorem strengthens the earlier complete-MUB analysis of Fiorentino and Weigert [4]. Its role here is precisely support selection. The discriminant-phase measurement, tangent response and operation-preserving transfer are the additional ingredients.

4 Entangled registers and testing rates

To extend the local detector to entangled inputs, we condition on measurements at the other sites. The following lemma ensures that some outcome preserves an independent pair at the target site. Let S_p be the union of the $p+1$ standard stabilizer MUBs. It consists of $p(p+1)$ rays and satisfies the projective two-design identity [7]

$$\sum_{a\in S_p} (|a\rangle\langle a|)^{\otimes 2} = 2P_{\mathrm{Sym}^2(\mathbb{C}^p)}.$$

In particular the second tensor powers span the symmetric square.

Lemma 4.1 (Localization of an independent pair)**.** *For independent $\psi, \chi \in (\mathbb{C}^p)^{\otimes n}$, some site j and product stabilizer bra $\langle a| = \bigotimes_{i\neq j}\langle a_i|$, $a_i \in S_p$, leave independent conditional vectors $\psi_a, \chi_a \in \mathbb{C}^p$.*

Proof. If every such pair were dependent, each coordinate of $\psi_a \wedge \chi_a$, homogeneous of degree two in each conditioning bra, would vanish on the finite Cartesian design. The quadratic spanning identity at every conditioning site promotes this to vanishing for all product bras.

Introduce the multilinear amplitude polynomials

$$f(x) = \sum_I \psi_I \prod_j x_{j,I_j}, \qquad g(x) = \sum_I \chi_I \prod_j x_{j,I_j}.$$

For each site their coefficient vectors are the conditional vectors. Wedge vanishing gives

$$f\,\partial_{j,b} g - g\,\partial_{j,b} f = 0 \quad \text{for every } j, b.$$

All derivatives of g/f vanish in the rational function field over $\mathbb{C}$. In characteristic zero such a rational function is constant, by treating one variable at a time over the field of the others. This contradicts independence. The $n = 1$ case is immediate. □

Theorem 4.2 (Consistent-label product transfer)**.** *A physical local test with L circuit settings and a tangent-complete k-copy effect transfers to a tangent-complete global effect with at most $nL(p+1)^{n-1}$ product-conditioning settings. Only one site needs a joint k-copy operation.*

Proof. Choose a target j uniformly and, independently, a MUB uniformly at each other site. Measure that basis on all k copies and continue precisely when the k labels agree separately at every conditioning site. Retain every such label tuple; no label is prescribed beforehand. At the target perform the fixed local test.

For fixed bases and target, the conclusive effect is

$$E_{j,\mathbf{M}} = \sum_{\mathbf{b}\in\mathbb{F}_p^{n-1}} E \otimes \bigotimes_{i\neq j} (M_i|b_i\rangle\langle b_i|M_i^*)^{\otimes k},$$

with tensor factors regrouped by site. Distinct label tuples occupy orthogonal subspaces, so $0 \le E_{j,\mathbf{M}} \le I$. On any global pure tensor each summand contracts to an identical local pure tensor; it is null even when the global input is entangled.

For an independent global pair, Lemma 4.1 supplies one conditioning outcome preserving independence. Its local excitation amplitudes equal the global ones for that branch. Local tangent-completeness gives a positive term; the other terms are nonnegative. Averaging the finite setting family preserves positivity. There are n targets, L target-circuit choices and $(p+1)^{n-1}$ conditioning-basis tuples; matching labels are outcomes, not extra settings. Their unnormalized probabilities are included throughout. □

Apply this theorem to the averaged tangent-complete effect E_p, realized by $L = p+1$ target-frame choices. This gives at most $n(p+1)^n$ settings and tangent completeness on every entangled register space. Lemma 2.2 proves the upper signal in Theorem 1.1.

Testing order. Independent trials miss an impurity-ϵ state with probability at most $(1-c_{p,n}\epsilon)^T \le e^{-c_{p,n}\epsilon T}$. Conversely, take orthogonal unit vectors and $\rho_\epsilon = (1-\epsilon)|\psi\rangle\langle\psi| + \epsilon|\chi\rangle\langle\chi|$. Its impurity is at least ϵ for $\epsilon \le 1/2$. Any exactly one-sided conclusive effect on N copies misses the all-ψ component of weight $(1-\epsilon)^N$. Thus $N \ge \log(1/\delta)/[-\log(1-\epsilon)]$, proving the unrestricted lower order in Corollary 1.2.

5 Quadratic contact below the characteristic

We turn to necessity. A single-excitation vector is a tensor in which one copy of a pure state is replaced by another vector. We prove that every conclusive effect on fewer than p copies annihilates these vectors at a pure stabilizer input. Consequently, a small admixture produces a click only through terms with at least two excitations. This is the quadratic-contact obstruction. We then show that it persists under arbitrary classical adaptation between blocks.

Theorem 5.1 (Single-excitation annihilation)**.** *Let $1 \le m < p$ and let $0 \le E \le I$ be an all-pure-null stabilizer-cone effect on $H_{p,n}^{\otimes m}$. For every pure stabilizer vector s, every vector χ and every slot j,*

$$E\big(s^{\otimes(j-1)} \otimes \chi \otimes s^{\otimes(m-j)}\big) = 0.$$

Consequently, if σ is a density matrix with $\sigma s = 0$, $0 \le t \le 1$, and $\rho_t = (1-t)|s\rangle\langle s| + t\sigma$, then

$$\mathrm{Tr}(E\rho_t^{\otimes m}) \le h_m(t) := 1 - (1-t)^m - mt(1-t)^{m-1} \le \binom{m}{2} t^2. \tag{3}$$

Proof. First take a normalized stabilizer-ray projector and s the global computational zero vector. Stabilizer computational amplitudes are either zero or one common scalar times pth roots of unity. All-pure-nullness says its overlap polynomial on identical pure tensors is zero. The constant histogram has zero coefficient, excluding the all-zero word. For every nonzero global label $y \in \mathbb{F}_p^n$, the histogram with $m-1$ zeros and one y has a coefficient which is a sum of at most $m < p$ such roots. A vanishing sum of pth roots with nonnegative integer multiplicities has equal multiplicities, by the minimal polynomial $1 + z + \cdots + z^{p-1}$. Its number of terms is divisible by p, so here it has no terms. All ordered single-y amplitudes are absent. Linearity covers arbitrary χ.

A common global Clifford maps any pure stabilizer s to zero and preserves stabilizer status and pure-nullness. For a positive cone sum, pure-nullness forces each nonzero ray summand to be pure-null. Their annihilation statements imply the same for E. Finally the positive tensor expansion of

$\rho_t^{\otimes m}$ has zero contribution at weights zero and one. Since $E \le I$, the remaining contribution is at most its binomial weight $h_m(t)$; the pair union bound gives the last inequality. □

In particular no such shorter-block effect admits a positive state-uniform linear impurity lower bound: take σ pure and let $t \downarrow 0$.

For adaptive protocols, we bound the probability of avoiding all mixedness-certifying block outcomes throughout the copy budget.

Theorem 5.2 (Exact adaptive survival bound)**.** *Fix $2 \le k < p$. An exactly one-sided protocol with blocks of at most k copies and deterministic total-copy cap N, with only classical memory between blocks, has miss probability satisfying*

$$\begin{aligned} \Pr(\text{miss}) \ge a_k(t)^q a_b(t) \ge a_k(t)^{N/k}, \qquad N = qk + b, \quad 0 \le b < k, \\ a_m(t) = (1-t)^{m-1}[1 + (m-1)t], \quad a_0(t) = a_1(t) = 1. \end{aligned} \tag{4}$$

on $\rho_t = (1-t)|s\rangle\langle s| + t|r\rangle\langle r|$, where s is any pure stabilizer and $r \perp s$ is unit, $0 < t < 1$. Hence miss probability at most δ requires

$$N \ge \frac{k \log(1/\delta)}{-(k-1)\log(1-t) - \log(1 + (k-1)t)}. \tag{5}$$

Proof. A conclusive transcript must contain a certificate. There are finitely many normalized stabilizer rays on at most k copies of the fixed register space. For every ray that is not all-pure-null its pure-overlap polynomial is nonzero. Choose a unit ϕ outside the finite union of these proper zero sets. Every block cone effect that is not all-pure-null then has strictly positive probability on $\phi^{\otimes m}$.

Call a conditional block outcome marked when its actual effect is all-pure-null. A marked outcome is therefore itself a mixedness certificate. Condition on a complete state-independent random tape. Any finite transcript containing no marked outcome has positive probability under ϕ whenever it has positive probability under another input. Exact one-sidedness under ϕ therefore makes conclusive mixedness transcripts without a mark have zero probability. Equivalently, the transcript law restricted to unmarked outcomes is absolutely continuous with respect to the law under ϕ. Integrating over the random tape preserves this assertion, because the conditional acceptance probabilities are nonnegative. The argument uses strict positivity of the unmarked conditional factors and nonnegative integration over the random tape.

A blockwise survival bound. At any history, the sum of all marked effects is an all-pure-null cone effect at most I. For a block of size $m \le k$, Theorem 5.1 bounds its marked probability by $h_m(t)$, so unmarked survival is at least $a_m(t) = 1 - h_m(t)$. Its per-copy logarithmic loss $-\log(a_m)/m$ increases with m. To see this, set $u = m - 1 \ge 0$ and $z = (1+ut)/(1-t) \ge 1$. Differentiation gives

$$\frac{\mathrm{d}}{\mathrm{d}u} \frac{-u\log(1-t) - \log(1+ut)}{u+1} = \frac{\log z - 1 + 1/z}{(u+1)^2} \ge 0.$$

Combining blocks under a fixed budget. Thus $a_m \ge a_k^{m/k}$. For the sharper integer bound, $\log a_x(t)$ is concave for real $x \ge 0$, since its second derivative is $-t^2/[1 + (x-1)t]^2$. For a fixed sum of two block sizes, moving them toward the endpoints $0, k$ cannot increase their product of survival factors. Repeating this operation shows that the minimum over integer partitions of N into parts at most k is

$$F(N) = a_k^q a_b, \qquad N = qk + b, \quad 0 \le b < k.$$

In particular $a_m F(N-m) \ge F(N)$ for each allowed next block. Backwards induction on the remaining copy budget therefore gives no-mark probability at least $F(N)$, including adaptive block

choices. Early termination preserves the bound because $F(N) \le 1$. A conclusive event requires a mark, so this is a miss probability lower bound. Finally $a_b \ge a_k^{b/k}$ gives the fractional-budget bound and (5). □

For fixed k, $-\log a_k(t) = k(k-1)t^2/2 + O_k(t^3)$. Calibrate the alternative to exact impurity $0 < \epsilon \le 1/2$ by choosing

$$t_\epsilon = \frac{1-\sqrt{1-2\epsilon}}{2}, \qquad \mathcal{I}(\rho_{t_\epsilon}) = \epsilon.$$

Then (5) gives

$$N \ge \left(\frac{8}{k-1} + O_k(\epsilon)\right)\epsilon^{-2}\log(1/\delta), \qquad \epsilon \downarrow 0.$$

This proves the lower order in Theorem 1.1. When $k = 1$, no nonzero effect can vanish on every pure state, so there is no conclusive event. This establishes the optimal-rate threshold for every odd prime.

The same cancellation argument can eliminate several excitation orders when p is large relative to the block size. The following refinement gives stronger lower bounds in that regime.

Proposition 5.3 (Higher contact below the characteristic)**.** *Let $1 \le r < k < p$ and suppose $(k)_r = k!/(k-r)! < p$. Every all-pure-null stabilizer-cone effect on $m \le k$ copies annihilates every excitation tensor with at most r factors orthogonal to any pure stabilizer anchor. For an orthogonal density matrix σ and $\rho_t = (1-t)|s\rangle\langle s| + t\sigma$, its click probability is at most*

$$\Pr\{\mathrm{Bin}(m,t) \ge r+1\} \le \binom{m}{r+1} t^{r+1}.$$

An exact one-sided adaptive protocol with block cap k and total-copy cap N consequently has

$$\Pr(\text{miss}) \ge \left(1 - \binom{k}{r+1} t^{r+1}\right)^{N/k}$$

whenever $\binom{k}{r+1}t^{r+1} < 1$. It requires $\Omega_{k,r}(\epsilon^{-(r+1)}\log(1/\delta))$ copies on the binary hard family.

Proof. A histogram with $j \le r$ nonzero excitation labels has at most $(m)_j \le (k)_r < p$ ordered words. The same cyclotomic argument as in Theorem 5.1 forces every such amplitude to vanish. Clifford covariance and positivity give the whole-register assertion and the binomial-tail bound for arbitrary orthogonal σ.

Put $x = \binom{k}{r+1}t^{r+1}$. Since $\binom{m}{r+1}/m \le \binom{k}{r+1}/k$, unmarked survival in an m-copy block is at least $1-(m/k)x \ge (1-x)^{m/k}$, by concavity of the logarithm. The marked-outcome argument and remaining-budget induction give the displayed survival bound. Taking a pure orthogonal σ and $t = t_\epsilon$ proves the stated copy order. For example, $p = 101$ and $k = 10$ give a cubic lower order because $(10)_2 = 90 < 101$. □

6 One-frame implementations in dimensions five and seven

The general construction uses $p+1$ local MUB frames. In the two dimensions below, a smaller family removes that frame randomization.

Theorem 6.1 (A fixed local circuit)**.** *For $p = 5$ or 7, a fixed rank-$p(p+1)$ all-pure-null projector P_p is tangent-complete and has one adaptive stabilizer implementation. Its product lift uses at most $n(p+1)^{n-1}$ circuit settings and has a linear impurity signal. Moreover no universal exact one-sided protocol with blocks below p exists, even with finite-cap classical adaptation.*

The projector is the polynomial-support family proved in Appendix B; Appendices C and D prove tangent completeness. Appendix A proves the phase-independent support exclusions. The register lift is Theorem 4.2. For the stronger impossibility, compress global effects against fixed stabilizer ancillas to the one-register setting and apply the finite-ray domination argument given at the end of Appendix A. In characteristic three the no-shorter assertion also follows directly from Theorem 5.1: every computational two-copy word lies in a constant or single-excitation space, so every two-copy pure-null cone effect is zero.

The setting improvement is a factor of $p+1$ relative to the general product implementation. The two statements measure different resources: p is the coherence needed for the optimal rate in all odd characteristics, while these fixed local circuits reduce the number of measurement settings in dimensions five and seven.

7 An exact sensitivity calibration

The following family quantifies how the sensitivity of the one-frame implementations depends on the number of registers.

Proposition 7.1 (GHZ calibration)**.** *For the one-frame product implementations of Theorem 6.1, write their probabilities as $q^P_{p,n}$ and put $|\Psi_\pm\rangle = (|0\rangle^{\otimes n} \pm |1\rangle^{\otimes n})/\sqrt{2}$ and, for $0 \le t \le 1$,*

$$\rho_{n,t} = (1-t)|\Psi_+\rangle\langle\Psi_+| + t|\Psi_-\rangle\langle\Psi_-|.$$

Write $d = 2t(1-t)$. Then

$$q^P_{5,n}(\rho_{n,t}) = \frac{d(1+d)}{100\,750^{n-1}}, \qquad q^P_{7,n}(\rho_{n,t}) = \frac{d(3+10d+4d^2)}{32\cdot 7^4\,134456^{n-1}}.$$

Consequently any uniform lower comparison constants for these implementations satisfy

$$c^P_{5,n} \le \frac{1}{100\,750^{n-1}}, \qquad c^P_{7,n} \le \frac{3}{32\cdot 7^4\,134456^{n-1}}.$$

These formulas quantify the sensitivity of the specified one-frame architecture.

Proof. For the one-register response, consider P_p for any prime $p \ge 5$ and a binary-supported state $\tau = \left[\begin{smallmatrix} a & z \\ \bar{z} & b \end{smallmatrix}\right]$, $a+b=1$. All plane rays are common-Z invariant. A binary word with k ones has eigenvalue ω^k; only the two constant words can have neutral character, and their coefficients vanish by pure-nullness. Thus the plane contribution is zero, as is the computational line contribution.

In each noncomputational MUB the transformed line coefficient has modulus $p^{-(p-1)/2}$ on its prescribed total-label hyperplane. Restriction to binary words with the same weight cancels the quadratic and coherence phases. Summing the p noncomputational line bases gives

$$q^P_{p,1}(\tau) = p^{2-p}\{D_p(1,|z|^2) - a^p - b^p\},$$

where $D_p(A,c) = r_+^p + r_-^p$ for roots of $r^2 - Ar + c = 0$. Indeed the equal-weight sum is the constant Laurent coefficient of

$$\prod_{j\in\mathbb{F}_p} (A + Bu\omega^j + Cu^{-1}\omega^{-j}) = D_p(A, BC) + B^p u^p + C^p u^{-p}.$$

The identity follows by factoring each quadratic after multiplication by $u\omega^j$; it also holds when B or C is zero by polynomial continuity. Remove the two constant binary words to obtain the displayed probability.

For the local state left after conditioning the GHZ family, $a = b = 1/2$, $|z|^2 = (1-2d)/4$, and $r_\pm = (1 \pm \sqrt{2d})/2$. Hence

$$q_{p,1}^P = p^{2-p} 2^{1-p} \sum_{j=1}^{(p-1)/2} \binom{p}{2j} (2d)^j.$$

If any conditioning basis is computational, its residual state is pure or zero. Otherwise each consistent label tuple has unnormalized local state $p^{-(n-1)}\tau$, with a label-dependent phase in z that does not affect the probability. Summing its p^{n-1} labels and averaging the probability $(p/(p+1))^{n-1}$ of noncomputational bases gives the factor $[(p+1)p^{p-2}]^{-(n-1)}$. Substitution of $p = 5, 7$ proves both formulas. Divide by d and let $t \downarrow 0$ for the constant bounds. □

8 A quantitative global Clifford-code alternative

Allowing Clifford operations across registers gives an alternative: compress to a random stabilizer code, then apply the local test. This yields an explicit dimension bound with a gain for low-rank inputs. Put $D = p^n$ and

$$\ell_p = \min_{\substack{\|\psi\|=\|\chi\|=1 \\ \langle\psi,\chi\rangle=0}} \mathcal{L}_{E_p}(\psi,\chi) > 0, \qquad p \text{ odd prime.}$$

This positive local constant measures the weakest tangent response over orthonormal pairs.

Theorem 8.1 (Quantitative code lift)**.** *There is a p-copy global-Clifford stabilizer test such that every state ρ of rank s satisfies*

$$\ell_p \frac{p(p-1)^{p-1}}{D^{p-1}(D-1)^{p-1}} \left(\frac{D}{s}\right)^{p-2} \mathcal{I}(\rho) \le q_{p,n}^{\mathrm{code}}(\rho) \le p\mathcal{I}(\rho).$$

The explicit dimension factor has order $D^{-2(p-1)}$ for full-rank states and order D^{-p} for bounded-rank states at fixed p, relative to the positive local constant ℓ_p. The one-frame projectors P_5, P_7 may be used in place of E_p with their corresponding local constants.

Proof. Fix a rank-p stabilizer code and choose a uniform global Clifford image $P = VV^*$. Here V is an isometric encoding of one p-level register into the full register space. Project every copy onto the same code, decode and run E_p. A failed code projection gives no click. This is a physical stabilizer test. Its pure-nullness follows because a compressed pure vector remains unnormalized pure.

The standard Clifford second moment [6] gives $\mathbb{E}(P \otimes P) = AI + BF_D$, with

$$A = \frac{p^2 - p/D}{D^2 - 1}, \quad B = \frac{p - p^2/D}{D^2 - 1}, \quad \alpha = A - B = \frac{p(p-1)}{D(D-1)}.$$

Here F_D exchanges the two copies of the D-dimensional space. These coefficients follow from traces p^2 and p once the two-design form is known. For orthogonal unit ψ, χ, set $a = \|V^*\psi\|^2$, $b = \|V^*\chi\|^2$ and $w = ab - |\langle V^*\psi, V^*\chi\rangle|^2$. Then $\mathbb{E}w = \alpha$ and $\mathbb{E}b = p/D$.

Homogeneity and removal of the component parallel to $V^*\psi$ give local response at least $\ell_p a^{p-2} w$. At $a = 0$ set $w/a = 0$. Hölder applied to

$$w = (a^{p-2}w)^{1/(p-1)} (w/a)^{(p-2)/(p-1)}$$

and $w/a \le b$ yields $\mathbb{E}(a^{p-2}w) \ge \alpha^{p-1}(D/p)^{p-2}$. Thus the averaged global tangent response is at least $\ell_D = \ell_p \alpha^{p-1}(D/p)^{p-2}$ on every orthonormal pair.

For a mixed state, write $\rho = \sum_{i=1}^{s} \lambda_i |i\rangle\langle i|$ using its positive eigenvalues. Retain all ordered tensor words with $p-1$ occurrences of i and one of $j \neq i$. Since $p > 2$, these patterns do not overlap for different repeated indices. Positivity gives

$$q_{p,n}^{\mathrm{code}}(\rho) \ge \ell_D \sum_i \lambda_i^{p-1}(1-\lambda_i).$$

If $d = 1 - \sum_i \lambda_i^2 > 0$, use weights $w_i = \lambda_i(1-\lambda_i)/d$. Their mean of λ_i is at least $1/s$. Writing e_j for the elementary symmetric polynomial of degree j in the positive eigenvalues, this is $(s-2)e_2 \ge 3se_3$, which follows from normalized Newton inequalities and $e_2 \le \binom{s}{2}/s^2$. Indeed $e_3/e_2 \le (s-2)/(3s)$; the $s = 2$ case has $e_3 = 0$.

Jensen therefore gives $\sum_i \lambda_i^{p-1}(1-\lambda_i) \ge s^{-(p-2)}d$. Combining the constants yields the lower bound. The pure case follows directly. The upper bound is the pure-null effect bound of Lemma 2.2. □

The displayed bound gives an explicit dimension dependence in terms of the positive local constants ℓ_p. Numerical sensitivity estimates require evaluating those constants.

Methods disclosure: use of AI tools

OpenAI Codex, configured with the recorded model identifier `gpt-6.1-sol` and the low reasoning setting, generated mathematical arguments and proof drafts, verification programs, and manuscript text, and assisted with literature comparison and internal analytic checks under the author's research and editorial direction during October 2026. The tools were used to explore candidate arguments, check finite constructions, compare source material, and organize the exposition. Session-creation records report Codex client/runtime versions `0.165.0-alpha.1.2` and `0.166.0-alpha.5.1` (OpenAI).

Conclusion

The characteristic p is the exact coherence threshold for optimal-rate exact one-sided stabilizer purity testing in every odd prime dimension and every fixed register count. In dimensions five and seven, one-frame implementations reduce the measurement settings, while exact support obstructions also establish the universal-existence threshold.

Acknowledgments

This research received no external funding.

Author declarations

Conflict of interest. The author declares no financial or nonfinancial competing interests.

Author contributions (CRediT). Scott Schubr: Conceptualization; Resources; Project administration; Supervision; Writing—review & editing.

Data availability

The verification programs and preprint source are available in the source package accompanying the Luminary Oddities Research Archive preprint at the article page and source download.

A Short-block affine-support obstructions

This appendix proves the stronger short-block exclusions used for dimensions five and seven. The first step removes the quadratic phase from the obstruction, leaving a finite question about affine supports. An odd-prime stabilizer vector has the computational normal form [6]

$$|A, Q\rangle = |A|^{-1/2} \sum_{x \in A} \omega^{Q(x)} |x\rangle,$$

up to a global phase, with affine support $A \subseteq \mathbb{F}_p^m$ and quadratic phase Q. A coordinate histogram records how many times each alphabet symbol occurs in a word, without recording its position. Grouping the overlap with $\psi^{\otimes m}$ by these histograms gives a phase-independent obstruction.

Lemma A.1 (Histogram divisibility)**.** *If a stabilizer ray is all-pure-null, every coordinate-histogram fiber of its affine support has cardinality divisible by p.*

Proof. Each histogram coefficient is a sum $\sum_{r=0}^{p-1} n_r \omega^r$. Its vanishing and the minimal polynomial $1 + z + \cdots + z^{p-1}$ force all integer phase multiplicities n_r to agree. Their sum is divisible by p. □

For $A = b + L$ define its complete weight enumerator $W_A(X) = \sum_{x \in A} \prod_i X_{x_i}$. Let $d = \dim L$ and let $B(L)$ count the d-coordinate projections invertible on L.

Proposition A.2 (Basis-count criterion)**.** *In $\mathbb{F}_p[t]$,*

$$W_A(1 + t, 1, \ldots, 1) = B(L) t^d + O(t^{d+1}).$$

Consequently $p \nmid B(L)$ excludes all-pure-nullness for every stabilizer phase on A.

Proof. The coefficient of t^r sums the cardinalities of the intersections $A \cap \{x_I = 0\}$ over r-coordinate sets I. Each nonempty intersection has size $p^{d - \operatorname{rank}(L \to \mathbb{F}_p^I)}$. For $r < d$ it is divisible by p. At $r = d$, precisely the invertible projections survive, each with one point. □

This is an elementary code-matroid rank-counting mechanism [5]. In particular $1 \le B(L) \le \binom{m}{d}$, so the condition $\binom{m}{\lfloor m/2 \rfloor} < p$ excludes every support of that length.

Theorem A.3 (The required short-block exclusions)**.** *No nonzero stabilizer-cone effect is all-pure-null on at most four copies of $\mathbb{C}^5$, or on at most six copies of $\mathbb{C}^7$.*

Proof. We use reduced row echelon form (RREF) to enumerate each linear direction once, together with one representative of each affine coset. For characteristic five, the exact RREF support certificate enumerates every affine subspace of $\mathbb{F}_5^4$. Its counts by dimension are $625, 19500, 20150, 780, 1$, totaling 41056. Each support has a histogram fiber whose size is nonzero modulo five. Lemma A.1 excludes every phase simultaneously.

For characteristic seven it suffices to work at length six. Centering the coordinates reduces the ambient dimension before enumeration. Put $H = \{x : \sum_i x_i = 0\} \cong \mathbb{F}_7^5$ and $Px = x - (\sum_i x_i/6)\mathbf{1}$. If $\mathbf{1} \notin L$, centering is injective on A and each centered histogram merges whole original histogram fibers. If $\mathbf{1} \in L$, the mean-zero slice lowers the dimension by one; the mean of a histogram specifies

its unique alphabet translation. A support with all fibers divisible by seven would therefore give one in H.

Centered dimensions zero and five are excluded by a singleton histogram and the zero word, respectively. At dimension one, seven points would have to share a histogram, impossible because a varying coordinate takes all seven symbols while a word has only six entries. A dimension-four support in H has equation $c\cdot x=\beta$, with $c\neq 0$ and $\sum c_i=0$. For $\beta=0$ it contains zero. Otherwise choose j with $c_j\neq 0$ and $a=-\beta/(6c_j)$. The word with five entries a and entry j equal to $-5a$ lies in the support. Its nonempty histogram fiber has at most six elements.

The preceding arguments settle all but centered dimensions two and three. These two cases are the finite computational part of the proof. The exact certificate enumerates 140050 RREF directions in each. It represents 48037150 and 6862450 affine cosets, respectively. Directions with $7\nmid B(L)$ are discharged by Proposition A.2; for the remaining 61560 directions it explicitly tests all 17587080 cosets and exhibits a nondivisible histogram fiber in each. Appendix E specifies the complete finite input. Thus no centered support remains.

Padding lower lengths with computational zeros preserves pure-nullness. Finally, a nonnegative sum of stabilizer expectations can vanish on every pure tensor only if each nonzero ray summand does. This gives the effect-cone assertions. □

Embedded inputs and classical adaptation. Throughout this paragraph let $p\in\{5,7\}$. Compress any global block effect against the same computational zero state on every nonselected register and every copy. Stabilizer postselection gives a local cone effect. Theorem A.3 therefore makes an all-pure-null block effect with fewer than p copies vanish on every embedded state $\sigma\otimes|0\rangle\langle 0|^{\otimes(n-1)}$.

For adaptive blocks, use the finite-ray absolute-continuity principle from [9]. On one to $p-1$ local copies there are finitely many normalized stabilizer rays, and none has identically zero pure-overlap polynomial. Choose a unit ϕ outside the finite union of their proper zero sets. The minimum squared overlap is positive, say $a>0$. Every compressed cone effect F then satisfies

$$\operatorname{Tr}(F\sigma^{\otimes m})\le\operatorname{Tr}F\le a^{-1}\langle\phi^{\otimes m},F\phi^{\otimes m}\rangle,\qquad m<p.$$

This uniform domination holds at every history and setting. Finite-cap classical-adaptive transcript laws, including state-independent randomization, are absolutely continuous under the embedded mixed input with respect to this embedded pure input. A conclusive event null on every pure input is therefore null on the embedded mixed family. No quantum memory across blocks is allowed. This proves the stronger impossibility assertion in Theorem 6.1.

B One polynomial-support projector

For $p\ge 5$ set $F|b\rangle=p^{-1/2}\sum_x\omega^{bx}|x\rangle$, $Q_s|x\rangle=\omega^{sx^2}|x\rangle$, and $R|x\rangle=|-x\rangle$. Write $|\chi_\ell\rangle=F|\ell\rangle$ and $\mathcal{M}_p=\{I\}\cup\{Q_sF:s\in\mathbb{F}_p\}$. These Clifford maps commute with parity and give the standard $p+1$ stabilizer MUBs. We show that their line and plane rays can be combined into one physically realizable projector.

Theorem B.1 (A polynomial-support projector)**.** *For each prime $p\ge 5$, the p^2+p stabilizer rays*

$$\{M^{\otimes p}h_\ell:M\in\mathcal{M}_p,\ \ell\neq 0\}\ \cup\ \{M^{\otimes p}g:M\in\mathcal{M}_p\}$$

are orthonormal and all-pure-null. Their projector P_p has rank $p(p+1)$ and its binary effects are implementable by one adaptive stabilizer circuit. Here

$$h_\ell=p^{-1/2}\sum_t\omega^{\ell t}|(j+t)_{j\in\mathbb{F}_p}\rangle,\qquad g=p^{-1}\sum_{b,c}\omega^b|(j^2+bj+c)_{j\in\mathbb{F}_p}\rangle.\tag{B1}$$

Appendices C and D establish universal detection by this projector in characteristics five and seven.

Proof. Every displayed vector is a normalized affine-support stabilizer ray. Cyclic translation of the copy labels gives a nontrivial character for h_ℓ and g; common M commutes with that translation. Hence all rays are orthogonal to identical pure tensors.

The commuting Paulis $\mathsf{X} = X^{\otimes p}$ and $\mathsf{Z} = Z^{\otimes p}$ give character exponents $(-\ell, 0)$ on h_ℓ, and $(2s\ell, -\ell)$ on $(Q_sF)^{\otimes p}h_\ell$. These are the $p^2 - 1$ distinct nonzero pairs in $\mathbb{F}_p^2$. Common X fixes g by shifting c, and common Z fixes it since $\sum_j (j^2 + bj + c) = 0$ for $p \geq 5$. All $p + 1$ plane rays are therefore in the neutral sector, orthogonal to every line ray.

Decode each computational word through its interpolation polynomial,

$$U|y\rangle = |c_{p-1}, \ldots, c_0\rangle, \qquad y_j = \sum_{r=0}^{p-1} c_r j^r.$$

This invertible linear map is Clifford. The untransformed plane state becomes the product state with $c_r = 0$ for $r \geq 3$, $c_2 = 1$, and Fourier states χ_1, χ_0 at c_1, c_0.

For $g_s = (Q_sF)^{\otimes p}g$, finite Fourier orthogonality gives

$$\langle y|F^{\otimes p}g\rangle = p^{-(p-2)/2}\mathbf{1}_{\{\sum_j y_j = 0,\ 1 + \sum_j j y_j = 0\}}\omega^{\sum_j j^2 y_j}.$$

Using $\sum_j j^k = -1$ for positive multiples of $p-1$ and zero for other positive powers, the constraints become $c_{p-1} = 0$, $c_{p-2} = 1$, and the phase becomes $-c_{p-3}$. On this support,

$$\sum_j y_j^2 = -2c_1 - H(\mathbf{c}), \qquad H(\mathbf{c}) = \sum_{\substack{2 \leq r,t \leq p-3 \\ r+t=p-1}} c_r c_t, \quad \mathbf{c} = (c_{p-3}, \ldots, c_2).$$

No other multiple of $p - 1$ occurs in the square, whose degree is at most $2p - 4$. Thus

$$Ug_s = |0, 1\rangle \otimes |\eta_s\rangle \otimes |\chi_{-2s}\rangle \otimes |\chi_0\rangle,$$

where the normalized $(p - 4)$-qupit stabilizer state is

$$|\eta_s\rangle = p^{-(p-4)/2} \sum_{\mathbf{c}} \omega^{-c_{p-3} - sH(\mathbf{c})}|\mathbf{c}\rangle.$$

The p states g_s have distinct Fourier labels at c_1. They are orthogonal to g, whose c_{p-2} value is zero. This proves orthonormality and the rank.

For the measurement circuit, measure X, Z. Their joint eigenvalue labels form the syndrome. A nonzero syndrome uniquely identifies a line ray; test its stabilizer projector within that branch. In the neutral branch apply U and measure c_{p-1}, c_{p-2}. At $(0, 0)$, test the product stabilizer state describing Ug above. At $(0, 1)$, measure c_1 in the Fourier basis with label b, set $s = -b/2$, and test η_s on the middle coordinates and χ_0 at c_0. A full-support quadratic phase state such as η_s is prepared and tested by Clifford operations. Every other branch is no-click. Since each target ray lies in its measured sector, each conclusive effect is exactly its projector. The complete binary effect is P_p and its complement. The instrument resolves the fine stabilizer outcomes and classically groups them into click and no-click events. □

C Tangent completeness in dimension five

Specialize to $p = 5$, write $\mathcal{M} = \mathcal{M}_5$, and let $W = \ker(R + I)$, with orthonormal basis

$$u = (|1\rangle - |4\rangle)/\sqrt{2}, \qquad v = (|2\rangle - |3\rangle)/\sqrt{2}.$$

The fine family of P_5 consists of twenty-four line rays and six plane rays. The line rays detect most independent directions. We identify the remaining odd-parity planes and use the plane rays to detect directions within them.

C.1 Line responses and the exceptional pure locus

If ψ has full support in one basis, its slot-excitation amplitudes for these line rays are nonzero multiples of the nonconstant Fourier coefficients of χ_x/ψ_x. Vanishing for every $\ell \neq 0$ forces this ratio to be constant. Thus a nonparallel χ is detected. If instead ψ has exactly one zero coordinate, say at b, the excitation amplitude at a fixed slot has exactly one surviving term and is a nonzero multiple of

$$\chi_b \prod_{x \neq b} \psi_x. \tag{C1}$$

The only pure vectors not already covered by a full-support basis have a zero in every MUB. We locate them by summing the projectors onto chosen zero rays. For one chosen ray in each basis let S be the sum of their projectors. Fixing the computational ray at $|0\rangle$ by a Weyl displacement leaves labels $b_a \in \mathbb{F}_5$ for the five chirp rays $5^{-1/2} \sum_x \omega^{ax^2 + b_a x} |x\rangle$.

Lemma C.1 (Singular MUB sums)**.** *Among the* 3125 *label tuples, S is singular exactly in the following cases: the five constant tuples; the 100 tuples with four equal entries and one exception; and the 20 nonconstant affine tuples $b_a = ca + d$. The kernel of every singular sum is contained in a Weyl displacement of W.*

Proof. The determinant classification is the exact cyclotomic calculation specified in Appendix E. The associated spectra are part of the phase-point enumeration in [3, Table III]. Here the labels identify the kernels needed by the construction.

For constant $b_a = b$, summing the chirp projectors gives $S = I + Z^b R Z^{-b} = 2P_+$, where P_+ is the corresponding even-parity projector. For four equal labels, replace one of these canonical rays by another. The five retained canonical rays span the same even space: their projector sum has trace five and squared trace nine, so its rank is at least $25/9 > 2$, and the even space has dimension three. The new kernel is consequently contained in the odd plane.

For $b_a = ca + d$, $c \neq 0$, the five chirp rays are canonical rays of the phase point whose computational ray is $|-c/2\rangle$. They span its even space by the same trace calculation. Adding $|0\rangle\langle 0|$ therefore leaves a kernel contained in its odd plane. There are 25 Weyl-displaced parity planes in total. □

C.2 The odd-plane tangent

Lemma C.2 (In-plane tangent response)**.** *Let $\psi = \alpha u + \beta v$ and $\chi = su + tv$. If $\alpha\beta \neq 0$ and ψ, χ are independent, the first-slot excitation amplitude of g is nonzero.*

Proof. The projection of g onto $W^{\otimes 5}$ has ten nonzero binary words, each with two u's and three v's. Two representative words and their unnormalized coefficients are

$$(2, 2, 1, 1, 2) : 1, \quad (2, 1, 2, 2, 1) : -1,$$

and a left cyclic shift by r positions multiplies either coefficient by ω^{2r}. Direct substitution into the 25 terms in (B1) gives

$$\langle g, \chi \otimes \psi^{\otimes 4} \rangle = \frac{K}{5\sqrt{32}} \alpha\beta^2(\alpha t - \beta s), \qquad K = -\omega^{-1} + \omega^{-2} + \omega^{-3} - \omega^{-4} = -\sqrt{5}. \tag{C2}$$

The coefficient certificate supplies all ten words. One can first put $\chi = v$ to obtain $K\alpha^2\beta^2$; linearity in χ and pure-nullness then give the determinant factor in (C2). It is nonzero under the hypotheses. □

Proposition C.3. *The rank-thirty projector P_5 is tangent-complete.*

Proof. Every ray is stabilizer and all-pure-null. Fix independent ψ, χ, taking ψ unit. A full-support MUB handles the pair by the Fourier-ratio argument above. Otherwise choose one zero ray from each MUB. Lemma C.1 puts ψ in a displaced odd plane. The line projectors are unchanged by a common Weyl displacement: $X^{\otimes 5} h_\ell = \omega^{-\ell} h_\ell$, $Z^{\otimes 5} h_\ell = h_\ell$, and $M^* D M$ is Pauli. We may therefore undo the displacement and work in W.

An odd vector has support two or four in each MUB. It has support two in at most one. Indeed, if its supports in two unbiased bases have sizes r, s, then Cauchy–Schwarz and the entrywise overlap bound imply $1 \le \sqrt{rs/5}$. Two supports of size two would contradict this inequality. Thus at least five bases are *good*: exactly one zero coordinate and four nonzero coordinates.

Their canonical zero rays span $W^\perp$, by the trace-rank calculation in Lemma C.1. Equation (C1) therefore detects every χ with a nonzero component in $W^\perp$. If $\chi \in W$, use a good basis and its parity-preserving map M. In the odd coordinates of this basis both components α, β of ψ are nonzero. Lemma C.2 detects the independent in-plane pair. This proves tangent completeness. □

D Tangent completeness in dimension seven

Now let $p = 7$ and $W = \ker(R + I)$, with odd basis $u_j = (|j\rangle - |-j\rangle)/\sqrt{2}$, $j = 1, 2, 3$. We prove that the fixed projector P_7 is tangent-complete. The proof follows the same division as in dimension five: locate the exceptional pure vectors, detect directions transverse to their odd spaces, and then detect directions within those spaces. Exact determinant and coefficient calculations establish the two finite steps.

D.1 A finite MUB kernel reduction

Lemma D.1 (The exceptional pure locus)**.** *If a pure vector has a zero coordinate in every standard stabilizer MUB of $\mathbb{C}^7$, it lies in one of the forty-nine Weyl-displaced odd spaces DW.*

Proof. Choose a zero ray in each of the eight bases and let S be the sum of their projectors. The vector lies in $\ker S$. A Weyl displacement makes the computational choice $|0\rangle$; the remaining seven chirp labels form a tuple $(b_a)_{a \in \mathbb{F}_7}$, giving $7^7 = 823543$ possibilities.

For a phase point (u, s) the canonical computational ray is $|u\rangle$ and its chirp label in basis a is $s - 2au$. If at least five selected rays are canonical at this phase point, they span its even four-space. Indeed their projector sum has trace five and squared trace $5 + 20/7 = 55/7$, so its rank is at least $35/11 > 3$; all these rays lie in that four-space. Hence $\ker S$ lies in the corresponding odd three-space.

The exact finite certificate accepts 92071 tuples by this geometric criterion. For every other tuple, it proves $\det(7S) \neq 0$ by a nonzero modular determinant: 704816 are excluded modulo 29,

and 26656 more modulo 43. Each prime is 1 modulo seven, so a chosen nontrivial seventh root defines a homomorphism from $\mathbb{Z}[\omega]$ to the finite field. Nonzero image implies a nonzero algebraic determinant. The counts total 823543, with no unresolved tuple. The explicit loop and acceptance conditions are specified in Appendix E. This proves the kernel containment, and undoing the displacement proves the claim. □

D.2 Directions transverse to the odd space

For a basis in which ψ has full support, the line first-order amplitudes are nonzero multiples of the nonconstant Fourier coefficients of χ_x/ψ_x. They all vanish only when χ is parallel to ψ. Otherwise Lemma D.1 puts ψ in a displaced odd space. Common-Weyl character invariance of the line projectors and neutrality of the plane projectors allow us to undo that displacement.

In W, the odd MUB directions are the three coordinate rays and the twenty-one rays Q_aFu_j. Their twenty-four projective coordinate columns have no three dependent: the exact geometry certificate gives 2024 nonzero three-column determinants under $\omega \mapsto 16$ modulo 29. Li's total-support bound also gives this odd-ray independence; the displayed minors supply direct exact certificates for this implementation. Thus a nonzero odd vector is orthogonal to at most two of these rays. At least six MUBs are therefore *good*, with exactly their canonical even zero and six nonzero coordinates.

The canonical even zero rays are $|0\rangle$ and $Q_aF|0\rangle$. In an unnormalized even basis their columns are $(1,0,0,0)$ and $(1,\omega^a,\omega^{4a},\omega^{2a})$. The same certificate gives all seventy four-column determinants nonzero modulo 29. Any four good zero rays consequently span $W^\perp$. In a single-zero basis, the line excitation is a nonzero multiple of χ's zero-coordinate component times the product of the six nonzero coordinates of ψ. Some line ray therefore detects every direction with a nonzero component in $W^\perp$.

D.3 The odd-plane coefficient certificate

It remains to detect directions within the odd three-space. Write $x = (\alpha, \beta, \gamma)$ for odd coordinates, and define the row of degree-six forms $A(x)$ by

$$\langle g, \chi \otimes \psi^{\otimes 6} \rangle = A(x)\chi.$$

Expanding the forty-nine support words of g and projecting to W gives twenty-one nonzero tensor coefficients. Up to one common nonzero normalization these are signed seventh roots, so the row has coefficients in $\mathbb{Q}(\omega)$. Pure-nullness gives the identity $A(x)x = 0$.

For an actual plane ray $M^{\otimes 7}g$, its row is $A(M^{-1}x)M^{-1}$. On W, $F^2 = R = -I$, so the eight coordinate pullbacks are projectively I and FQ_s, $s \in \mathbb{F}_7$. The adjoint order is important: these are precisely the eight plane rays in P_7.

For any pair of rows $a(x), b(x)$, their cross product has the form

$$a(x) \times b(x) = x\, P_{a,b}(x)$$

with a homogeneous degree-eleven form $P_{a,b}$. To see the global polynomial divisibility, the identities $a \cdot x = b \cdot x = 0$ imply $\alpha C_2 = \beta C_1$ and $\alpha C_3 = \gamma C_1$ for $C = a \times b$. Unique factorization gives $\alpha \mid C_1$, and then $C = xP_{a,b}$ identically, including the chart $\alpha = 0$.

The next calculation rules out a common nonzero zero of these cross forms. There are twenty-eight such cross forms. Multiply each by all six quadratic monomials. Their 168 coefficient vectors span the entire 105-dimensional space of ternary degree-thirteen forms. This is the exact finite

assertion checked by the plane tangent certificate modulo 29 at $\omega = 16$; it verifies the eight pullbacks, row syzygies, cross divisions and coefficient rank 105.

A nonzero modular minor proves the corresponding characteristic-zero minor nonzero. Projective Fourier normalization removes its common square-root factor; the remaining coefficients lie in a localization of $\mathbb{Z}[\omega]$ whose used denominators have nonzero image. Thus this rank calculation is on *coefficient vectors*, not evaluations at a finite set of complex or finite-field points.

If every $P_{a,b}$ vanished at a nonzero complex x, their quadratic multiples would all vanish there. Their full span contains $\alpha^{13}, \beta^{13}, \gamma^{13}$, a contradiction. Hence at each $x \neq 0$ two rows have nonzero cross product. All rows annihilate x, so they span its two-dimensional annihilator. Every χ independent of x has nonzero amplitude for at least one plane ray.

Proposition D.2. *The rank-fifty-six projector P_7 is tangent-complete.*

Proof. A full-support MUB handles the nonexceptional case. In an odd space the line effects handle transverse directions, and the coefficient certificate handles in-plane directions. Weyl covariance covers every displaced odd space. These cases exhaust all pure vectors and independent directions. □

E Exact finite certificates and reproduction

This appendix specifies the finite computations used above, including their coverage and acceptance conditions. All arithmetic is integer or finite-field arithmetic, and a failed condition raises an exception. The Python programs require only the standard library; the C# certificates have source and portable .NET run instructions. Partial or timed-out runs are not accepted as complete proofs. The programs and geometric reductions together constitute the finite inputs; hashes identify deterministic evidence and do not replace their mathematical coverage.

Characteristic-five support and geometry. `ququint_support_certificate.py` enumerates RREF bases in $\mathbb{F}_5^4$ and the unique coset representatives zero at pivot coordinates. Counts agree separately with $5^{4-d}\begin{bmatrix}4\\d\end{bmatrix}_5$, totaling 41056. Here the Gaussian binomial coefficient counts the d-dimensional linear subspaces; the power of five counts their affine cosets. Each sorted-coordinate histogram table has a fiber nonzero modulo five. The deterministic witness-stream SHA-256 is

`a123307fde2dc3146eb052e2f297b4d0768f1321eb61f379eb3c211b7a827d72.`

`mub_singular_certificate.py` checks all 3125 chirp tuples in $\mathbb{Q}(\omega)$: five constant, one hundred four-equal and twenty nonconstant affine tuples are singular; 3000 are nonsingular. Kernel containment is proved in the text. `parity_detector_certificate.py` gives the ten projected plane coefficients and the tangent identity used in Section C.

Characteristic-seven support. `CenteredSevenCertificate.cs` handles centered dimensions two and three in $\mathbb{F}_7^5$. Each has 140050 RREF directions. The full represented coset counts are 48037150 and 6862450. Nonzero basis count modulo seven discharges the other directions analytically. The remaining direction counts are 49560 and 12000, giving 16999080 and 588000 explicitly checked cosets. The program precomputes all 16807 centered words and their 132 histograms, then exhibits a nondivisible fiber for every tested coset. Complete runs require these counts and the witness digest

`9f7d6ad88e26349bff2b964033c157fc1a827ba651ec0a833aeb0524563647e5.`

The centering reduction and remaining dimensions are proved in Theorem A.3; the code does not silently assume them.

Characteristic-seven MUB kernels. `SevenMubLocus.cs` runs every base-seven label tuple exactly once. A geometric acceptance certifies five canonical rays spanning an even four-space. Otherwise exact Gaussian elimination must give a nonzero determinant modulo one of the specified primes. The totals are 92071 geometric cases, 704816 exclusions modulo 29 and 26656 additional exclusions modulo 43. No unresolved tuple is allowed. A nonzero modular determinant is an exact sufficient test for the algebraic determinant, not a numerical rank tolerance.

Characteristic-seven tangent. `seven_parity_plane_certificate.py` reconstructs the twenty-one odd projected tensor coefficients from forty-nine support words. `seven_mub_tangent_geometry.py` checks twenty-four distinct odd rays, all 2024 three-minors, and all seventy even four-minors at the valid specialization $\omega = 16$ modulo 29.

`seven_plane_tangent_certificate.py` uses precisely the eight actual coordinate pullbacks I, FQ_s (projective Fourier normalization). It checks every pure-null syzygy and cross divisibility, constructs twenty-eight degree-eleven cross forms and their 168 quadratic multiples, and requires coefficient rank 105 in degree thirteen. Its deterministic matrix-and-independent-multiple witness digest is

`a0083b19b6747d023970229976529830efeabff6d87206eaa1f80c7ee0bc43a3.`

The coefficient matrix has one column for each degree-thirteen monomial. Full rank therefore certifies the spanning statement used in the proof, and exhibits a characteristic-zero nonzero minor under a cyclotomic ring homomorphism. It does not infer complex-projective nonvanishing from tests on finite-field points. This distinction and the inverse order of the plane pullbacks are essential.

General-construction controls. `check_discriminant_detector.py` checks the quadratic-image cover and the inside-support tangent rank at exact modular specializations for $p = 3, 5, 7, 11, 13$. These finite controls illustrate the general construction; the all-prime proof and the cited support-selection input are analytic.

Availability and scope. The accompanying source package contains these exact acceptance programs, editable manuscript, build and run instructions. All analytic arguments are in the text. The specialized finite certificates are computer-assisted proof inputs.

F A direct support-selection proof following Li

Lemma F.1 (Near-full MUB support; adapted from Li)**.** *For every odd prime p and every nonzero vector in $\mathbb{C}^p$, at least one of the standard $p+1$ mutually unbiased bases has at least $p-1$ nonzero coordinates.*

Proof. We prove the support-selection statement used in Section 3 by adapting the cyclotomic-jet argument of Li [8]. The argument assumes two zeros in every basis and bounds the total number of those zeros. A jet here means a coefficient of the truncated power series introduced below. For $p = 3$, failure would make the vector a basis ray in two mutually unbiased bases, which is impossible. Suppose $p \geq 5$ and the conclusion fails. Its computational support R has size $k \leq p-2$. Writing $\omega = e^{2\pi i/p}$, the other bases give the exponential sum

$$F(a,b) = \sum_{x \in R} c_x \omega^{ax^2+bx}, \qquad (a,b) \in \mathbb{F}_p^2.$$

Changing signs or multiplying the quadratic parameter by a nonzero constant merely relabels these bases. The zero set Z of F has at least $2p$ points.

We first reduce arbitrary complex coefficients to coefficients in a cyclotomic number field. Replace the coefficients, if necessary, by a nonzero vector with minimal support in the kernel of the equations indexed by Z. Relabel its actual support as R and its size as k. The replacement preserves every zero in Z, and $k \le p-2$. A one-point support has no zeros, so $2 \le k \le p-2$. On that support the kernel is one-dimensional: two independent kernel vectors could be combined to cancel a coordinate and reduce the support. Since the equation matrix is over $K = \mathbb{Q}(\omega)$, row reduction supplies a generator over K. Thus, after this replacement and a scalar rescaling, all coefficients lie in the discrete valuation ring $\mathcal{O} = \mathbb{Z}[\omega]_{(1-\omega)}$, and at least one is a unit. This step permits arbitrary original complex coefficients and discards every coordinate that becomes zero.

There is an isomorphism

$$\mathcal{O}/((1-\omega)^{p-1}) \simeq \mathbb{F}_p[T]/(T^{p-1}), \qquad \omega \longmapsto e^T = \sum_{j=0}^{p-2} \frac{T^j}{j!}.$$

Indeed, p is a unit times $(1-\omega)^{p-1}$; the displayed substitution respects the cyclotomic relation and sends $1-\omega$ to a generator of the maximal ideal. Both rings have length $p-1$, so this surjective map is an isomorphism. Write the images of the normalized coefficients as $C_x(T)$, and expand

$$G(a,b,T) = \sum_{x \in R} C_x(T) e^{T(ax^2+bx)} = \sum_{j=0}^{p-2} H_j(a,b) T^j \pmod{T^{p-1}}.$$

Here H_j has total degree at most j. At every $(a,b) \in Z$, all $H_j(a,b)$ vanish.

Let n be the first index with $H_n \neq 0$. Then $n \le \lfloor k/2 \rfloor$. To see both existence and this bound, if $H_0, \dots, H_{\lfloor k/2 \rfloor}$ all vanished, their highest-total-degree coefficients would give $\sum_x C_x(0) x^r = 0$ for $0 \le r \le k-1$: each r occurs as $2u+v$ with $u+v \le \lfloor k/2 \rfloor$. The Vandermonde matrix on the distinct elements of R would force all $C_x(0) = 0$, contrary to the unit normalization. In particular $n \le (p-3)/2$.

The identity $\partial_b^2 G = T \partial_a G$ gives $\partial_b^2 H_j = \partial_a H_{j-1}$, with $H_{-1} = 0$. Thus $Q_0 := H_n = A(a) + bB(a)$, where $\deg A \le n$ and $\deg B \le n-1$. For $0 \le j \le n$, put $Q_j = H_{n+j}$. These jets are available since $2n \le p-3$, and induction in the recurrence gives

$$\deg_b Q_j \le 2j+1, \qquad \partial_b^{2j} Q_j = \partial_a^j Q_0.$$

The degree bound follows without a characteristic-p kernel term because $\deg Q_j \le n + j < p$.

We count zeros one vertical fiber at a time, fixing a and varying b. If $B \neq 0$, let r be its number of roots in $\mathbb{F}_p$. At a nonroot of B, the equation $Q_0 = 0$ permits at most one value of b. At a common root α of A, B, let $m_\alpha = \min\{\operatorname{ord}_\alpha A, \operatorname{ord}_\alpha B\}$, using infinite order for an identically zero polynomial. The displayed derivative identity shows that $Q_{m_\alpha}(\alpha, b)$ is nonzero: its $2m_\alpha$-th derivative is the nonzero affine polynomial $A^{(m_\alpha)}(\alpha) + bB^{(m_\alpha)}(\alpha)$. It therefore has at most $2m_\alpha + 1$ roots. Roots of B that are not roots of A contribute none. Consequently

$$|Z| \le p - r + \sum_{A(\alpha)=B(\alpha)=0} (2m_\alpha + 1) \le p + 2\deg B \le p + 2n - 2 < 2p.$$

If $B = 0$, then $A \neq 0$. A root α of A of multiplicity m similarly permits at most $2m+1$ values of b, by the nonzero polynomial $Q_m(\alpha, b)$. Hence $|Z| \le 3 \deg A \le 3n < 2p$. If A is constant there are no zeros. Both cases contradict $|Z| \ge 2p$. □

References

[1] S. Arunachalam and L. Schatzki, *Optimal stabilizer testing and learning with limited quantum memory*, arXiv:2607.02444v2 (2026). https://arxiv.org/abs/2607.02444v2.

[2] J.-Y. Cai, Z. Fu, and S. Shao, *From Holant to quantum entanglement and back*, in ICALP 2020, LIPIcs 168 (2020), 22:1–22:16. doi:10.4230/LIPIcs.ICALP.2020.22.

[3] A. Casaccino, E. F. Galvão, and S. Severini, *Extrema of discrete Wigner functions and applications*, Phys. Rev. A 78 (2008), 022310. doi:10.1103/PhysRevA.78.022310.

[4] V. Fiorentino and S. Weigert, *Uncertainty relations for the support of quantum states*, J. Phys. A: Math. Theor. 55 (2022), 495305. doi:10.1088/1751-8121/aca9a4.

[5] C. Greene, *Weight enumeration and the geometry of linear codes*, Stud. Appl. Math. 55 (1976), 119–128. doi:10.1002/sapm1976552119.

[6] D. Gross, S. Nezami, and M. Walter, *Schur–Weyl duality for the Clifford group with applications: Property testing, a robust Hudson theorem, and de Finetti representations*, Comm. Math. Phys. 385 (2021), 1325–1393. doi:10.1007/s00220-021-04118-7.

[7] A. Klappenecker and M. Rötteler, *Mutually unbiased bases are complex projective 2-designs*, in IEEE ISIT 2005 (2005), 1740–1744. doi:10.1109/ISIT.2005.1523643.

[8] D. Li, *Sharp total-support uncertainty for the standard complete set of mutually unbiased bases*, arXiv:2609.16030v2 (2026). https://arxiv.org/abs/2609.16030v2.

[9] S. Schubr, *Purity frontiers and optimal copy complexity for qutrit stabilizer measurements*, Luminary Oddities Research Archive, edition 2 (2026). https://research.luminaryoddities.com/purity-frontiers-optimal-copy-exponent-qutrit-stabilizer-measurements/v/2.

[10] C. Zhu, Z. Liu, C. Zhu, and X. Wang, *Limitations of classically simulable measurements for quantum state discrimination*, Phys. Rev. Lett. 133 (2024), 010202. doi:10.1103/PhysRevLett.133.010202.