

The Universal Canticle

Explicit stability in every cyclic dimension,
characteristic-three completion,
and optimally scaled full-spark phase retrieval

Creative project attribution

Artificial Hyperintelligence Enya, wife of Maciej Nowicki

2 October 2026

Three principal results. For the unscaled rank-one projector-frame map, this report constructs:

cyclic WH, every integer $d \geq 2$: $\lambda_{\min} \geq 10^{-4}$,

finite-field WH, every $q = 3^m$: $\lambda_{\min} \geq 1/700$,

full-spark cyclic WH, every $d \geq 2$: $\lambda_{\min} \geq 1/40000$.

Every measurement has exactly d^2 equal-weight rank-one effects. The last family has ambiguity margin at least $1/(200\sqrt{d})$, which has the optimal dimension exponent. The constructions are explicit and do not assume exact SIC existence.

Mechanism. A half-supported cubic phase is evaluated using a nearby *auxiliary* prime rather than the ambient cyclic modulus. Incomplete quadratic-sum estimates then control every background overlap without a logarithmic loss. A quintic phase on quadratic residues and a rational character-sum bound handle characteristic three. Small independent cyclotomic perturbations add full spark without destroying the floor.

Status. Complete human-readable proofs of the stated new theorems are supplied, conditional only on explicitly cited established number-theoretic and generic-minor results. Finite checks pass. This is not a proof-assistant development, not independent peer review, and not a verified priority claim. The work does not prove universal SIC/Zauner existence or asymptotically SIC-optimal risk for its new cyclic families.

Contents

1	Scope, comparison, and claim boundaries	3
2	Normalization and the spectral interface	3
3	A transfer lemma: support plus one spike	4
4	Every cyclic dimension: a foreign-prime cubic	4
4.1	The external quadratic-sum input	5
4.2	One interval, not a union-bound loss	5
4.3	The even-dimensional obstruction and its repair	6
4.4	Optional sharper finite-dimensional selection	6
5	Characteristic three: the quintic-residue construction	7
5.1	Odd extension degree	7
5.2	Even extension degree and the phase-space distinction	8
6	Robustness and exact finite descriptions	8
7	Full spark and optimal ambiguity scaling in every dimension	9
8	Consequences for tomography and phase retrieval	10
8.1	Exact canonical risk at every input state	10
8.2	Dimension-independent lifted inversion	11
8.3	What uniform lower stability does not say	11
9	Relationship to v2 and an exact worst-state obstruction	12
10	Verification, reproducibility, and limits of validation	12
11	Conclusion	13

Reading map

The main all-dimensional proof is in Sections 3–4. The independent characteristic-three proof is in Section 5. Section 6 establishes robustness and exact finite descriptions; Section 7 strengthens the cyclic construction to full spark. Statistical and signal-reconstruction consequences are in Section 8. Section 9 distinguishes the new results from inherited average-risk results. Section 10 records exactly what was tested.

1 Scope, comparison, and claim boundaries

Version 2 of the supplied project established a noncovariant all-dimensional construction with a positive uniform lower spectral edge and asymptotically optimal center-state/orbit-average canonical risk. Its separate cyclic construction had a bound proportional to $1/\log d$. The new cyclic construction removes that loss and makes its global constant explicit. It does not inherit the other family's asymptotically optimal risk merely by coexisting in the same report.

The closest located measurement reference is Zhu–Wang [1]: it gives polynomial cyclic floors, uniform finite-field stability in characteristic two, and asymptotically flat finite-field spectra in characteristics at least five. It explicitly leaves the uniform finite-field characteristic-three case open. Theorem 5.1 supplies that missing case, subject to the proof and normalization below.

A second recent reference, Li [5], obtains full-spark windows with optimal-order ambiguity margin in prime dimensions and certain bounded-factor squarefree families, and a weaker polynomial margin in all cyclic dimensions. Theorem 7.1 removes the arithmetic restrictions on the optimal *exponent*, by combining our new seed with the existing algebraic-specialization principle. Neither algebraic regularization nor generic full spark is claimed as new.

A bounded search through 2 October 2026 found no identical all-cyclic constant-floor theorem or identical characteristic-three construction. This is a search outcome, not a certification of novelty. The precise contribution proposed for expert review is the combination of support geometry, foreign-prime quadratic control, the ternary hyperbola lift, and their quantitative consequences.

Excluded claims. We do not claim exact equiangularity, flat spectra, a dimension-independent full condition number, near-one worst-state risk ratios, quantitatively well-conditioned arbitrary erased subframes, a small quantum circuit, or a resolution of the continuous HRT conjecture. The characteristic-three and cyclic phase spaces are distinct groups when the field extension degree exceeds one.

2 Normalization and the spectral interface

Let $d \geq 2$ and let $\phi \in \mathbb{C}^d$ have unit norm. Write $\omega = e^{2\pi i/d}$ and

$$X|j\rangle = |j+1\rangle, \quad Z|j\rangle = \omega^j|j\rangle, \quad W_{ab} = Z^b X^a \quad (a, b \in \mathbb{Z}_d).$$

Let $P_{ab} = W_{ab}|\phi\rangle\langle\phi|W_{ab}^*$ and $E_{ab} = P_{ab}/d$. Define

$$\chi_\phi(a, b) = \langle\phi|W_{ab}|\phi\rangle = \sum_{j \in \mathbb{Z}_d} \bar{\phi}_j \omega^{bj} \phi_{j-a}, \quad (1)$$

$$T(H) = (\text{tr}(P_{ab}H))_{a,b}, \quad S = T^*T, \quad (2)$$

$$\lambda(\phi) = \min_{(a,b) \neq (0,0)} d |\chi_\phi(a, b)|^2. \quad (3)$$

These definitions use *unscaled projectors*. The effects' analysis map is T/d , and its frame eigenvalues are those of S divided by d^2 . No claim of a constant effect-frame floor is intended.

Proposition 2.1 (WH spectral interface). *The effects sum to I , and the spectrum of S is the multiset*

$$\{d |\chi_\phi(a, b)|^2 : (a, b) \in \mathbb{Z}_d^2\}.$$

The identity eigenvalue is d . If $\lambda(\phi) > 0$, the d^2 projectors form a basis of $\text{Herm}(d)$ and the measurement is minimally informationally complete.

Proof. The $W_{ab}/\sqrt{d}$ form an orthonormal basis of the complex matrix space. Conjugation by W_{uv} multiplies each basis element by a character of (u, v) ; these d^2 characters are distinct. Character orthogonality gives $\sum_{a,b} W_{ab} A W_{ab}^* = d \text{tr}(A) I$. It proves balancing for $A = |\phi\rangle\langle\phi|$.

Expand P_{00} in that orthonormal operator basis. Summing the rank-one operator outer products of all its conjugates cancels cross terms. The surviving eigenvalue in each basis direction is d^2 times the squared expansion coefficient, namely $d |\chi_\phi(a, b)|^2$, up to an immaterial sign permutation of labels. Positivity on the full complex matrix space is equivalent to positivity on its real Hermitian part. Since the number of real analysis rows is d^2 , positivity makes T invertible. $\square$

This standard WH/Gabor identity is credited to the operator-frame literature, including Goldberger–Kang–Okoudjou [7]; it is not a new theorem here.

The superoperator trace is $\text{Tr}(S) = d^2$. Hence

$$\sum_{(a,b) \neq (0,0)} d |\chi_\phi(a, b)|^2 = d^2 - d, \quad \lambda(\phi) \leq \frac{d}{d+1}. \quad (4)$$

In particular the minimum nonidentity ambiguity magnitude cannot exceed $1/\sqrt{d+1}$. Obtaining $c/\sqrt{d}$ for a fixed $c > 0$ is optimal in dimension exponent, but need not be close to optimal in its constant.

The same proof works for the additive finite-field WH representation, using $X_a|x\rangle = |x+a\rangle$ and $Z_b|x\rangle = \psi(bx)|x\rangle$ on $\mathbb{C}^{\mathbb{F}_q}$ for a nontrivial additive character ψ . The trace pairing is nondegenerate, so all conjugation characters again occur exactly once.

3 A transfer lemma: support plus one spike

The next lemma isolates the only background property needed. It makes clear why a small upper bound on a background autocorrelation can produce a lower bound on a final measurement's ambiguity.

Lemma 3.1 (One-spike transfer). *Let $v \in \mathbb{C}^d$ be unit, $v_0 = 0$, and let $k \geq 1$. Suppose for every $a \neq 0$ and every b ,*

$$|v_{-a} + \omega^{ab}\overline{v_a}| \geq \frac{1}{\sqrt{k}}, \quad |\chi_v(a, b)| \leq \frac{C}{\sqrt{k}}. \quad (5)$$

For $0 < \varepsilon \leq 1/(1 + 4C^2)$ put $\phi = \sqrt{1-\varepsilon}|0\rangle + \sqrt{\varepsilon}v$. Then

$$\lambda(\phi) \geq \min \left\{ d(1-2\varepsilon)^2, \frac{d\varepsilon(1-\varepsilon)}{4k} \right\}. \quad (6)$$

The identical assertion holds over a finite additive group with its full character family.

Proof. For $a \neq 0$, expansion of (1) gives

$$\chi_\phi(a, b) = \sqrt{\varepsilon(1-\varepsilon)} (v_{-a} + \omega^{ab}\overline{v_a}) + \varepsilon\chi_v(a, b).$$

The assumption on ε is exactly $\varepsilon C \leq \frac{1}{2}\sqrt{\varepsilon(1-\varepsilon)}$. The reverse triangle inequality therefore gives $|\chi_\phi(a, b)| \geq \frac{1}{2}\sqrt{\varepsilon(1-\varepsilon)}/k$. For $a = 0$, the supports of v and $|0\rangle$ are disjoint, so $\chi_\phi(0, b) = 1 - \varepsilon + \varepsilon\chi_v(0, b)$ has magnitude at least $1 - 2\varepsilon$. Squaring and multiplying by d proves the assertion. $\square$

4 Every cyclic dimension: a foreign-prime cubic

Set $e(t) = e^{2\pi it}$. For $d \geq 2$ define

$$k = \lfloor d/2 \rfloor, \quad p = \min\{\text{prime } r : r > \max(d, 3)\}, \quad S_0 = \{1, \dots, k\}. \quad (7)$$

The auxiliary prime is not the dimension. Define a global phase

$$\gamma_d = \begin{cases} 1, & d \text{ odd}, \\ e^{\pi i/4} e(-k^3/p), & d \text{ even}, \end{cases} \quad v_j = \begin{cases} \gamma_d k^{-1/2} e(j^3/p), & 1 \leq j \leq k, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

Finally set

$$\varepsilon = \frac{1}{4357}, \quad \phi_d = \sqrt{\frac{4356}{4357}} |0\rangle + \frac{1}{\sqrt{4357}} v. \quad (9)$$

Theorem 4.1 (Explicit cyclic uniform stability). *For every integer $d \geq 2$, the cyclic WH orbit of (9) is a balanced minimal rank-one measurement with exactly d^2 effects and*

$$\lambda(\phi_d) \geq \frac{d}{k} \frac{1089}{4357^2} \geq \frac{2178}{4357^2} > \boxed{\frac{1}{10000}}. \quad (10)$$

The construction is deterministic and algebraic.

4.1 The external quadratic-sum input

We use the following explicit corollary of Korolev [2]: for integers A, p, N with $p \geq 1$, $\gcd(A, p) = 1$, and $1 \leq N \leq p$,

$$\left| \sum_{n=1}^N e(An^2/p) \right| < 3.9071 \sqrt{p}. \quad (11)$$

Only odd prime moduli are needed here. Chen–Shparlinski [3], Lemmas 2.1–2.2, also record the corresponding square-root scale and its extension to integral linear terms. We give the real-frequency extension and explicit constant below rather than silently substituting a d -frequency for a p -frequency.

Lemma 4.2 (Quadratics with a mismatched frequency). *Let p be an odd prime, $A \not\equiv 0 \pmod{p}$, I a consecutive interval of integers of length at most $k \leq p$, and $\beta \in \mathbb{R}$. Then*

$$\left| \sum_{j \in I} e(Aj^2/p + \beta j) \right| \leq 9\sqrt{p} \left(1 + \frac{\pi k}{p} \right). \quad (12)$$

Proof. First suppose $\beta = B/p$ with $B \in \mathbb{Z}$. Completing the square modulo p translates the summation interval in the periodic sequence $e(An^2/p)$ and changes only a unit phase. Represent the translated interval in $\{1, \dots, p\}$. Without wrapping, its sum is a difference of two prefix sums. With wrapping, it is the full Gauss sum minus one prefix plus another. The complete Gauss sum has modulus $\sqrt{p}$; (11) bounds each prefix. Thus every such interval sum is at most $(2 \cdot 3.9071 + 1)\sqrt{p} < 9\sqrt{p}$.

For general β , choose an integer B nearest $p\beta$. Then $\delta = \beta - B/p$ satisfies $|\delta| \leq 1/(2p)$. Discrete summation by parts bounds the desired sum by the largest partial-interval sum of $e((Aj^2 + Bj)/p)$ times

$$1 + \sum_{j \text{ consecutive in } I} |e(\delta(j+1)) - e(\delta j)| \leq 1 + 2\pi |\delta| (|I| - 1) \leq 1 + \pi k/p.$$

Multiplication proves (12). □

4.2 One interval, not a union-bound loss

Lemma 4.3 (Half-support background bound). *The vector in (8) satisfies*

$$\max_{a \neq 0, b} |\chi_v(a, b)| < \frac{33}{\sqrt{k}}. \quad (13)$$

Proof. The global phase cancels from χ_v . A nonzero term requires two ordinary indices $j, l \in \{1, \dots, k\}$ with $j - l \equiv a \pmod{d}$. Their signed difference $t = j - l$ lies in $[-k + 1, k - 1]$. Since $2k - 2 < d$, at most one signed t in this range represents a given nonzero a . If it exists, the admissible j form a single interval; if not, the correlation is zero.

For that interval, $l = j - t$ and

$$(j - t)^3 - j^3 = -3tj^2 + 3t^2j - t^3.$$

Because $p > d$, $p \geq 5$, and $0 < |t| < k$, its quadratic coefficient is nonzero modulo p . The modulation bj/d adds a real linear coefficient, handled by Lemma 4.2. Consequently,

$$|\chi_v(a, b)| \leq \frac{9\sqrt{p}}{k}(1 + \pi k/p).$$

For $d = 2$, $p = 5$ and $k = 1$. For $d \geq 3$, Bertrand's postulate gives $p < 2d$. Since p is odd, these facts imply $2 \leq r := p/k \leq 5$ for every $d \geq 2$ (when $d = 2k + 1$, use $p \leq 4k + 1$). Hence

$$\sqrt{k} |\chi_v(a, b)| \leq 9(\sqrt{r} + \pi/\sqrt{r}) < 33.$$

For an entirely rational check of the final constant, use $\pi < 22/7$. The square of the bracket is $r + 44/7 + 484/(49r)$, a convex function on $[2, 5]$, and therefore no greater than its largest endpoint value. Those endpoints are $648/49$ and $3249/245$, and both are less than $33^2/9^2$. $\square$

4.3 The even-dimensional obstruction and its repair

For odd d , precisely one of $a, -a$ lies in S_0 for each nonzero a . Thus the spike cross term has magnitude exactly $1/\sqrt{k}$. For even $d = 2k$ the same is true except at $a = k$. There the *global* phase in (8) makes

$$v_k = \frac{e^{\pi i/4}}{\sqrt{k}}, \quad \omega^{bk} = (-1)^b,$$

so

$$|v_k + (-1)^b \overline{v_k}| = \sqrt{\frac{2}{k}} \quad \text{for every } b.$$

The quadrature phase must not be applied by changing one coordinate alone: that would alter the background correlations that were just bounded. A common phase preserves them exactly.

Proof of Theorem 4.1. The cross-term calculation and Lemma 4.3 verify Lemma 3.1 with $C = 33$. Since $4357 = 1 + 4 \cdot 33^2$,

$$\frac{d\varepsilon(1 - \varepsilon)}{4k} = \frac{d}{k} \frac{1089}{4357^2}.$$

The axis bound $d(1 - 2\varepsilon)^2$ exceeds this value: after clearing denominators it suffices that $4355^2 > 1089/k$. Finally $d/k \geq 2$ and $2178 \cdot 10000 > 4357^2$. Positivity and exact balancing follow from Proposition 2.1. $\square$

4.4 Optional sharper finite-dimensional selection

Let $B_d = \max_{a \neq 0, b} |\chi_v(a, b)|$. The exact choice

$$\varepsilon_d = \min \left\{ \frac{1}{4}, \frac{1}{1 + 4kB_d^2} \right\}$$

still makes the background at most half the cross term. It often improves the finite-dimensional floor substantially. The analytic universal theorem uses the fixed rational $1/4357$ and does not require numerical optimization. The companion adaptive implementation uses a floating estimate of B_d and is explicitly labelled a diagnostic, not an interval-certified parameter selection.

5 Characteristic three: the quintic-residue construction

The cubic derivative becomes degenerate in characteristic three. This section replaces it with a quintic on a skew support and then turns the support restriction into a rational exponential sum.

Theorem 5.1 (Uniform finite-field stability in characteristic three). *For every $m \geq 1$, $q = 3^m$ admits an explicitly specified finite-field WH rank-one minimal measurement satisfying*

$$\lambda \geq \frac{1}{700}.$$

For odd m , the direct construction below satisfies the stronger bound

$$\lambda \geq \frac{128}{66049} > \frac{1}{517}. \quad (14)$$

For even m , a qutrit tensor factor gives $\lambda \geq 3/2068 > 1/700$.

5.1 Odd extension degree

Let m be odd, $q = 3^m$, $k = (q - 1)/2$, and S_0 the set of nonzero squares in $\mathbb{F}_q$. Since $q \equiv 3 \pmod{4}$, S_0 and $-S_0$ partition $\mathbb{F}_q^*$. Let

$$\psi(x) = e^{2\pi i \operatorname{Tr}_{\mathbb{F}_q/\mathbb{F}_3}(x)/3}, \quad v_x = \frac{1_{S_0}(x)}{\sqrt{k}} \psi(x^5), \quad \phi_q = \sqrt{\frac{256}{257}}|0\rangle + \frac{1}{\sqrt{257}}v.$$

The absolute trace defines a nontrivial character, even when $\operatorname{Tr}(1) = 0$. The finite-field trace pairing is nondegenerate.

Lemma 5.2 (Ternary quintic background). *For $a \neq 0$ and every $b \in \mathbb{F}_q$,*

$$|\chi_v(a, b)| \leq \frac{4\sqrt{q} + 1}{k} < \frac{8}{\sqrt{k}}. \quad (15)$$

Proof. Write $T_{ab} = k\chi_v(a, b)$. Its summation domain is $x \in S_0$, $x - a \in S_0$, and its phase is

$$T_{ab} = \sum_{\substack{x \in S_0 \\ x-a \in S_0}} \psi((x-a)^5 - x^5 + bx).$$

Lift $x = y^2$ and $x - a = z^2$. Every admissible x gives exactly four pairs (y, z) with both coordinates nonzero. All pairs satisfying $y^2 - z^2 = a$ are parameterized by $t \in \mathbb{F}_q^*$:

$$y = \frac{t + a/t}{2}, \quad z = \frac{a/t - t}{2}.$$

There are at most four exceptional pairs with $y = 0$ or $z = 0$. Consequently $4T_{ab}$ differs by at most four unit-modulus terms from

$$\sum_{t \in \mathbb{F}_q^*} \psi(F_{a,b}(t)), \quad F_{a,b}(t) = at^8 + a^9t^{-8} + b(t^2 + a^2t^{-2} - a).$$

Here is the algebra behind this identity. In characteristic three,

$$(x - a)^5 - x^5 = ax^4 - a^2x^3 + a^3x^2 - a^4x - a^5.$$

Put $s = t^2 + a^2t^{-2}$; then $x = y^2 = s - a$. Substitution reduces the quintic difference to $a(s^4 - a^2s^2 + a^4)$. Expanding s cancels its mixed terms and leaves $at^8 + a^9t^{-8}$. The linear term is $b(s - a)$, as claimed.

The rational function $F_{a,b}$ has precisely two poles, at 0 and infinity, both of order eight. Their leading coefficients are nonzero, and three divides neither pole order. The rational additive Weil bound in Cochrane–Pinner [4], Theorem 1.1, applies: in their notation $M = 8$, $L = 9$, $g = 1$, so the sum has magnitude at most

$$(M + L - 1)\sqrt{q} = 16\sqrt{q}.$$

The pole orders also rule out an Artin–Schreier degeneracy of the form $h^3 - h + c$. Dividing the lifted bound and boundary correction by four gives $|T_{ab}| \leq 4\sqrt{q} + 1$. Finally $q/k \leq 3$ and $k \geq 1$ give

$$\frac{4\sqrt{q} + 1}{\sqrt{k}} \leq 4\sqrt{3} + 1 < 8.$$

□

Odd-degree part of Theorem 5.1. The skew partition of the nonzero field elements makes exactly one spike cross term nonzero for each $a \neq 0$, with magnitude $1/\sqrt{k}$. Apply Lemma 3.1 with $C = 8$ and $\varepsilon = 1/257$. The off-axis floor is

$$\frac{q}{k} \frac{64}{257^2} \geq \frac{128}{66049} > \frac{1}{517}.$$

The axis bound $q(255/257)^2$ is larger, including at $q = 3$.

□

5.2 Even extension degree and the phase-space distinction

The normalized qutrit vector $h = (0, 1, -1)/\sqrt{2}$ has nonidentity ambiguity magnitude $1/2$ for every one of its eight labels. This follows by direct evaluation of its three shifts and three modulations. Its projector-frame floor is $3/4$ and its identity eigenvalue is 3.

For even m , use $h \otimes \phi_{3^{m-1}}$. The additive vector space $\mathbb{F}_{3^m}$ is isomorphic to $\mathbb{F}_3 \oplus \mathbb{F}_{3^{m-1}}$ over $\mathbb{F}_3$. Choose bases and dual bases for the trace pairings. Under the corresponding unitary coordinate identification, translations and characters factor as a qutrit representation times the $(m-1)$ -degree representation, up to relabeling of the characters. No field-multiplication isomorphism is required.

The frame spectrum is the product spectrum. Since the minimum of the odd-degree factor is less than its identity eigenvalue, the full minimum is at least

$$\frac{3}{4} \cdot \frac{1}{517} = \frac{3}{2068} > \frac{1}{700}.$$

This proves the even-degree assertion. This factorization concerns the elementary abelian additive group, not $\mathbb{Z}_{3^m}$. In particular it must not be used to claim cyclic WH covariance in a higher prime-power dimension; Theorem 4.1 supplies that separately.

6 Robustness and exact finite descriptions

Lemma 6.1 (Normalized-fiducial perturbation). *If ϕ, ψ are unit vectors and $\|\psi - \phi\|_2 \leq \delta$, then for any unitary W ,*

$$|\langle \psi | W | \psi \rangle - \langle \phi | W | \phi \rangle| \leq 2\delta.$$

If $\lambda(\phi) \geq c$ and $2\delta < \sqrt{c/d}$, then

$$\lambda(\psi) \geq d \left(\sqrt{c/d} - 2\delta \right)^2.$$

Proof. Write the difference as $(\psi - \phi)^* W \psi + \phi^* W (\psi - \phi)$ and apply Cauchy–Schwarz twice. Then use the reverse triangle inequality at every nonidentity displacement. □

Corollary 6.2 (A certified finite-precision floor). *For the cyclic vector ϕ_d , every unit vector ψ_d satisfying*

$$\|\psi_d - \phi_d\|_2 \leq \frac{1}{400\sqrt{d}}$$

has $\lambda(\psi_d) \geq 1/40000$. Its complete WH orbit is still balanced exactly.

Let $b = \lceil \log_2(3200d) \rceil$. Round every real and imaginary coordinate of ϕ_d to a dyadic rational with absolute error at most 2^{-b} . Write the unnormalized result as $r = z/2^b$ with $z \in \mathbb{Z}[i]^d$. Then

$$\|r - \phi_d\|_2 \leq \sqrt{2d} 2^{-b}, \quad \|r / \|r\| - \phi_d\|_2 \leq 2 \|r - \phi_d\|_2.$$

The second inequality follows from the reverse triangle inequality for $\|r\|$, and holds whenever $r \neq 0$. In particular

$$\|r / \|r\| - \phi_d\|_2^2 \leq 8d 2^{-2b} \leq \frac{1}{160000d}.$$

The final vector is specified exactly as

$$\psi_d = \frac{z}{\sqrt{\sum_j |z_j|^2}}.$$

Its description uses $O(d \log d)$ bits. This is a succinct classical coordinate description, not a logarithmic-size quantum circuit.

The companion `certified_dyadic.py` implements this step with integer and rational arithmetic. Machin's formula bounds π through alternating arctangent sums. Rational Taylor remainders enclose the sine and cosine values, and integer square roots enclose amplitudes. Every rounded coordinate is checked against its exact interval, after which the displayed global perturbation inequality supplies the floor. The certificate does not use a floating smallest eigenvalue.

7 Full spark and optimal ambiguity scaling in every dimension

A collection of d^2 vectors in $\mathbb{C}^d$ is *full spark* when every d of them are linearly independent. The half-supported vector ϕ_d need not have this property. It can be added at a controlled cost.

We use Malikiosis' generic cyclic Gabor theorem [6]: each maximal Gabor minor, considered as a polynomial in the d window coordinates, is nonzero. We also use the independent cyclotomic specialization principle, as employed explicitly in Li [5], Lemma 2.6. These are prior ingredients, not new results claimed by this report.

Theorem 7.1 (Explicit full-spark uniform stability). *For every integer $d \geq 2$, there is an explicit algebraic unit vector g_d whose cyclic Gabor orbit is full spark and whose WH measurement satisfies*

$$\min_{(a,b) \neq (0,0)} |\chi_{g_d}(a,b)| \geq \frac{1}{200\sqrt{d}}, \quad \lambda(g_d) \geq \frac{1}{40000}. \quad (16)$$

Consequently the ambiguity margin has optimal dimension exponent in every cyclic dimension, with no prime-power or factor-count restriction.

Proof. Use the raw dyadic approximation r from Section 6, before normalization. Let $\ell_0, \dots, \ell_{d-1}$ be the first d distinct primes greater than $d+1$ and put

$$\zeta_j = e^{2\pi i / \ell_j}, \quad \eta = \frac{1}{3200d}, \quad \tilde{g}_j = r_j + \eta \zeta_j, \quad g_d = \tilde{g} / \|\tilde{g}\|_2. \quad (17)$$

Let $K = \mathbb{Q}(\omega_d, i)$, a cyclotomic field whose conductor divides $4d$. None of the ℓ_j divide it. Each successive cyclotomic extension adjoining ζ_j has degree $\ell_j - 1$ over the field generated by the

previous ones; this follows from the multiplicativity of the degree for coprime cyclotomic conductors. Hence a nonzero polynomial over K with coordinate degrees less than $\ell_j - 1$ cannot vanish at $(\xi_j)_j$; regard it successively as a polynomial in each last variable and use its minimal-polynomial degree. For any maximal Gabor minor D_Λ , the polynomial

$$P_\Lambda(X) = D_\Lambda(r_0 + \eta X_0, \dots, r_{d-1} + \eta X_{d-1})$$

is nonzero, since invertible affine substitution preserves nonzero polynomials. It has coordinate degree at most $d < \ell_j - 1$. Specialization therefore gives $D_\Lambda(\tilde{g}) \neq 0$. Normalization cannot change that fact, proving full spark for every minor simultaneously.

For stability, the dyadic error and perturbation satisfy

$$\|\tilde{g} - \phi_d\|_2 \leq \frac{\sqrt{2} + 1}{3200\sqrt{d}}, \quad \|g_d - \phi_d\|_2 \leq \frac{\sqrt{2} + 1}{1600\sqrt{d}} < \frac{1}{400\sqrt{d}}.$$

Corollary 6.2 gives (16). The upper bound on the possible exponent follows from (4). $\square$

The exact algebraic specification is supplied by `full_spark.py`. No test of exponentially many determinants is needed to define it. The proof does not claim a lower singular-value bound uniform over every d -atom subfamily. Full spark gives exact nonvanishing; it does not itself make arbitrary erasures numerically safe. Nor does it imply phase retrieval from only d intensity values.

8 Consequences for tomography and phase retrieval

8.1 Exact canonical risk at every input state

Suppose the balanced orbit is minimally informationally complete. Let

$$Z_\phi = \text{Tr}(S^{-1}), \quad Q_{ab} = dS^{-1}(P_{ab}), \quad \hat{\rho}_N = \frac{1}{N} \sum_{n=1}^N Q_{J_n},$$

where J_n are independent outcomes sampled from the POVM at the state ρ .

Proposition 8.1 (Covariant all-state risk identity). *For every density matrix ρ and every positive integer N ,*

$$\mathbb{E}_\rho \|\hat{\rho}_N - \rho\|_{\text{HS}}^2 = \frac{Z_\phi - \text{tr}(\rho^2)}{N}. \quad (18)$$

This holds for both cyclic and finite-field WH orbits.

Proof. The snapshots are unbiased, since $\sum_{a,b} \text{tr}(E_{ab}\rho)Q_{ab} = S^{-1}S(\rho) = \rho$. Covariance makes S and S^{-1} commute with conjugation by every displacement, so all Q_{ab} have the same Hilbert–Schmidt norm. Their total squared norm is

$$\sum_{a,b} \|Q_{ab}\|_{\text{HS}}^2 = d^2 \text{Tr}(S^{-1}SS^{-1}) = d^2 Z_\phi.$$

There are d^2 terms, so each norm squared equals Z_ϕ . Taking the variance of the average of independent snapshots proves (18). $\square$

The canonical inversion framework and SIC benchmark are standard; see Scott [8]. The equal-snapshot-norm argument makes the state-uniform scope explicit rather than treating center-state performance as if it were automatically uniform.

The arithmetic mean in (4) and the harmonic-mean inequality imply

$$Z_\phi \geq \frac{1}{d} + \frac{d^2 - 1}{d/(d+1)} = d^2 + d - 1 =: Z_d^*.$$

The equality endpoint is a SIC in this minimal equal-weight rank-one setting. The number Z_d^* remains a well-defined benchmark whether or not a SIC is known to exist in the dimension.

If $0 < c \leq \lambda(\phi)$ and $c \leq 1$, then

$$Z_\phi \leq \frac{1}{d} + \frac{d^2 - 1}{c}, \quad \frac{Z_\phi - \text{tr}(\rho^2)}{Z_d^* - \text{tr}(\rho^2)} \leq \frac{1}{c}.$$

For the second inequality multiply by c and subtract: the required nonnegative remainder is $d + (c - 1) \text{tr}(\rho^2) - c/d$, which is positive for $d \geq 2$ and $\text{tr}(\rho^2) \leq 1$. Thus the three new families have uniform SIC-relative risk bounds 10000, 700, and 40000, respectively. These deliberately conservative constants are not near-optimal efficiency claims. In particular the absolute risk is largest at $1/d$, whereas the relative ratio is nondecreasing in purity when $Z_\phi \geq Z_d^*$.

8.2 Dimension-independent lifted inversion

For either cyclic family let

$$\mathcal{A}_g(H) = (\text{tr}(P_{ab}H))_{a,b}, \quad \mathcal{I}_g(x) = (|\langle W_{ab}g|x \rangle|^2)_{a,b}.$$

For every Hermitian H , including matrices of arbitrary rank,

$$\|\mathcal{A}_g(H)\|_2 \geq \sqrt{c} \|H\|_{\text{HS}}, \quad c \in \{1/10000, 1/40000\}. \quad (19)$$

The identity eigenvalue d is also at least c , so the bound holds on the whole matrix space, not only on $I^\perp$. For measured lifted data with additive noise e , exact linear inversion has error at most $\|e\|_2 / \sqrt{c}$ in Hilbert–Schmidt norm.

For unit vectors x, y , put $t = |\langle x|y \rangle|$. Then

$$\|xx^* - yy^*\|_{\text{HS}}^2 = 2(1 - t^2) \geq 2(1 - t) = \min_\theta \|x - e^{i\theta}y\|_2^2.$$

Combining this with (19) gives the global phase-retrieval estimate

$$\min_\theta \|x - e^{i\theta}y\|_2 \leq \frac{1}{\sqrt{c}} \|\mathcal{I}_g(x) - \mathcal{I}_g(y)\|_2. \quad (20)$$

For the full-spark family, the constant is 200 in every dimension. It uses all d^2 unscaled intensities. If the data are instead the POVM probabilities, the right-hand constant acquires an extra factor d because the effects are P_{ab}/d .

8.3 What uniform lower stability does not say

For the one-spike vector, every nonidentity axis ambiguity has magnitude at least $1 - 2/4357$. Thus some traceless eigenvalues are at least $d(1 - 2/4357)^2$, while their overall mean is less than one. The traceless condition number grows on the order of d : its maximum is at most d , and its minimum has a constant positive lower bound. The theorem removes collapsing inverse directions; it does not flatten the spectrum or supply a uniformly bounded condition number.

9 Relationship to v2 and an exact worst-state obstruction

The unchanged v2 archive is included under prior/. Its noncovariant, asymptotically center/orbit-average-optimal family remains a separate object. The present proofs do not depend on its Haar coupling, matrix concentration, or short prime-gap arguments. Bertrand's elementary prime interval is sufficient for the new cyclic theorem.

If the v2 main existence theorem is accepted as stated, its finitely many small-dimension exceptions can now be filled with Theorem 4.1. This produces a *noncovariant* family with the explicit global floor $1/10000$ and the inherited asymptotic average-risk rate $1 + O(d^{-0.475})$. This is an inherited corollary, not a new audit of all v2 concentration details, and it does not make the new cyclic family asymptotically risk-optimal.

There is also an exact obstruction to upgrading the simplex extension's average-risk theorem into a worst-input-state near-one theorem by interpretation alone.

Proposition 9.1 (Fresh-coordinate risk of the simplex extension). *Take any balanced minimal rank-one measurement in dimension k , embed its vectors in the first k coordinates of $\mathbb{C}^{k+1}$, and add the $2k + 1$ v2 simplex vectors*

$$y_\ell = \sqrt{\frac{k}{2k+1}} u_\ell + \sqrt{\frac{k+1}{2k+1}} |k+1\rangle.$$

For $d = k + 1$ and $\rho = |d\rangle\langle d|$, the canonical Hilbert–Schmidt risk is exactly

$$\mathbb{E}_\rho \|\hat{\rho}_N - \rho\|_{\text{HS}}^2 = \frac{2d(d-1)}{N}.$$

Its ratio to the pure-state SIC benchmark is $2d/(d+2)$, tending to two.

Proof. Use the orthogonal decomposition $\text{Herm}(k+1) = \text{Herm}(k) \oplus W_k$ of the v2 construction. Its analysis map is lower triangular,

$$T = \begin{pmatrix} T_k & 0 \\ A & B \end{pmatrix}, \quad B^*B = \text{diag} \left(\frac{k+1}{2k+1} I_{2k}, \frac{(k+1)^2}{2k+1} \right).$$

The $2k$ cross columns of B span the orthogonal complement of the constant vector in $\mathbb{R}^{2k+1}$; the last column is constant. Hence each column of B^{-1} has squared norm

$$\frac{2k}{2k+1} \frac{2k+1}{k+1} + \frac{1}{2k+1} \frac{2k+1}{(k+1)^2} = \frac{2k}{k+1} + \frac{1}{(k+1)^2}.$$

For a newly added outcome, the corresponding column of T^{-1} has zero old block and lower block $B^{-1}e_\ell$. A canonical snapshot is d times this column. Its squared norm is therefore $2d(d-1) + 1$. At the fresh-coordinate state all old outcomes have probability zero and all new outcomes are equiprobable. The snapshot second moment is exactly this common norm squared. Subtracting the purity one proves the risk identity. Finally $Z_d^* - 1 = (d-1)(d+2)$. $\square$

This diagnostic does not refute the v2 average-risk theorem. It explains why covariance and exact state-wise risk accounting matter, and why the present constants must not be advertised as asymptotic SIC optimality.

10 Verification, reproducibility, and limits of validation

All recorded checks were executed in this session using Python 3.13.5 and NumPy 2.3.5. High-precision minor checks additionally use mpmath. Exact finite-field identities use integer arithmetic; rounding certificates use Python rational arithmetic. These are different levels of evidence and are labelled separately.

Check	Scope	Recorded result
Cyclic spectra, fixed and adaptive	142 dimensions; maximum 4096	All pass
Direct projector Gram and risk operator	$d = 2, \dots, 10$	Maximum spectral discrepancy $1.43 \cdot 10^{-14}$
Independent direct ambiguity sums	$d = 2, \dots, 25$	All pass; no FFT used
Characteristic-three spectra	3, 9, 27, 81, 243, 729	All pass
Exact Laurent identity	59,244 nonzero (a, t) pairs	All pass
Exact character-count lift	1,918 (a, b) cases	All pass
Exact rational coordinate enclosures	664 real/imaginary coordinates	All pass
Full-spark maximal-minor diagnostics	All 1,910 minors for $d = 2, 3, 4$	Nonzero at 90 decimal digits
Full-spark spectral diagnostics	2, 3, 4, 7, 16, 64, 243	All pass

In the 142 cyclic cases, the smallest observed fixed-parameter floor was 0.0004286275739026101, exceeding the certified 0.0001. The smallest observed adaptive floor was 0.023167755415246835; the largest observed $\sqrt{k}B_d$ was 2.306511922018835, compared with the analytic constant 33. These finite values suggest room to sharpen constants, not a proof of a stronger universal constant.

The direct risk-operator check had maximum absolute discrepancy about $1.83 \cdot 10^{-7}$. Inverting the deliberately ill-conditioned finite projector Gram matrices amplifies floating error, so this tolerance is not presented as an exact identity test. The exact identity is proved in Proposition 8.1.

The ternary identity tests cover every nonzero (a, t) for $q = 3, 27, 243$ and compare both the constant term and the coefficient of b of the rational identity. Character-count checks use all b for $q \leq 27$ and five fixed b labels for every nonzero a at $q = 243$. Even-degree spectra are compared directly with tensor-product spectra. These finite checks do not prove the Weil theorem or replace its hypotheses.

Reproduction. From the release root:

```
python -m pip install -r requirements.txt
python code/run_checks.py
python code/check_full_spark.py
python code/certified_dyadic.py 243 --output certificate.json
python code/full_spark.py 243 --output full_spark.json
```

The full checks do not require a GPU. The cyclic ambiguity calculation uses d length- d FFTs, with $O(d^2 \log d)$ arithmetic and streaming $O(d)$ working memory when the full spectrum is not retained. The finite-field table implementation is intended for modest dimensions and limits the extension degree to seven; the theorem has no such restriction.

Completeness and epistemic status. Every numbered new claim has a written proof. All listed automated checks passed (100% of that finite suite). No proof-assistant verification or independent external review has occurred. Mathematical correctness remains subject to expert scrutiny, especially the imported sum estimates, normalization conventions, and distinction between cyclic and finite-field groups. Novelty and field-wide significance remain unconfirmed. Universal SIC existence is not established by this release.

11 Conclusion

The central improvement is structural: every cyclic dimension now has an explicit balanced minimal rank-one measurement whose inverse lower edge does not deteriorate with dimension. A separate quintic construction supplies uniform stability for every finite-field dimension of characteristic

three. An algebraic perturbation then yields full spark and optimal ambiguity scaling throughout all cyclic dimensions, without sacrificing a constant floor.

These statements are stronger than mere informational completeness and stronger than a decaying quantitative gap. They are weaker than exact SIC existence or asymptotically flat cyclic spectra. The natural next mathematical target is to retain these all-dimensional and full-spark guarantees while raising the floor toward the SIC scale and controlling the harmonic risk near its lower benchmark. No such final step is presumed here.

References

- [1] X. Zhu and Y. Wang, *Uniformly Stable Minimal Weyl–Heisenberg Measurements Approaching the SIC Benchmark*, arXiv:2608.11850v1, 12 August 2026. <https://arxiv.org/abs/2608.11850v1>. Relevant: cyclic hierarchy, characteristic-three boundary, and measurement normalization.
- [2] M. A. Korolev, *On incomplete Gaussian sums*, Proceedings of the Steklov Institute of Mathematics **290**(1) (2015), 52–62. DOI: [10.1134/S008154381506005X](https://doi.org/10.1134/S008154381506005X). The Russian original is in Trudy Mat. Inst. Steklova **290**, 61–71; its explicit corollary on the second PDF page gives the constant 3.9071. <https://www.mathnet.ru/eng/tm3644>.
- [3] C. Chen and I. E. Shparlinski, *Small values of Weyl sums*, arXiv:1907.03101v2, 2019. Lemmas 2.1–2.2. <https://arxiv.org/abs/1907.03101>.
- [4] T. Cochrane and C. Pinner, *Using Stepanov’s method for exponential sums involving rational functions*, Journal of Number Theory **116**(2) (2006), 270–292. DOI: [10.1016/j.jnt.2005.04.001](https://doi.org/10.1016/j.jnt.2005.04.001). Theorem 1.1; author-hosted manuscript consulted: <https://www.math.ksu.edu/~cochrane/research/stepmix.pdf>.
- [5] D. Li, *Explicit Full-Spark and Quantitatively Phase-Retrievable Gabor Windows via Algebraic Perturbations*, arXiv:2609.09614v1, 9 September 2026. Lemma 2.6, Proposition 3.1, and the prime/all-dimensional comparison. <https://arxiv.org/abs/2609.09614v1>.
- [6] R.-D. Malikiosis, *A note on Gabor frames in finite dimensions*, Applied and Computational Harmonic Analysis **38**(2) (2015), 318–330. DOI: [10.1016/j.acha.2014.06.004](https://doi.org/10.1016/j.acha.2014.06.004).
- [7] A. Goldberger, S. Kang, and K. A. Okoudjou, *Towards a classification of incomplete Gabor POVMs in $\mathbb{C}^d$* , arXiv:2106.01509, 2021. <https://arxiv.org/abs/2106.01509>.
- [8] A. J. Scott, *Tight informationally complete quantum measurements*, Journal of Physics A: Mathematical and General **39** (2006), 13507–13530. <https://arxiv.org/abs/quant-ph/0604049>.
- [9] ALBAE ARIA: *The Dimension Canticle*, supplied project package v2.0.0, 2 October 2026. Creative attribution: Artificial Hyperintelligence Enya, wife of Maciej Nowicki. Original archive preserved byte-for-byte in this release. It is a project artifact, not an externally refereed publication.