

Sparse Weyl operators: a uniform binary rank theorem and sharp erasure thresholds

Lluis Eriksson

Independent researcher, Sweden

15 October 2026 Working manuscript

Abstract

For a complex matrix expanded in the cyclic Weyl basis, support size alone gives the rank bound $\text{rank}(C)s \geq d$. We refine this estimate by the largest nonzero translate overlap μ of its support:

$$\text{rank}(C) \geq \left\lceil \frac{d}{1 + \mu(1 - 1/s)} \right\rceil.$$

The estimate is an elementary fourth-moment argument. Its complement is structural: large overlaps force few partial cyclic runs. Combining cyclic compression with acyclic block selection, we prove that every nonzero combination of at most seven Weyl operators in dimension $d = 2^k$ has rank at least $d/4$. For $d \geq 4$, this is sharp under every support budget from four through seven, uniformly in k . It completes the conjectured rank envelope in dimension 16 and closes additional gaps in dimensions 27 and 32. Two exact universal erasure budgets follow for arbitrary complex low-rank matrices: seven in a binary family of ranks, and eight for rank at most two in dimension 27. Explicit examples fail at the next erasure. A quantitative fourth-moment bound also controls distance from low rank. The proofs use classical scalar uncertainty and graph lemmas; no general prime-power conjecture or Erdős problem is claimed solved.

1 The question and the results

Let $d \geq 2$, $\omega = e^{2\pi i/d}$, and let $X, Z \in M_d(\mathbb{C})$ be given by

$$Xe_j = e_{j+1 \bmod d}, \quad Ze_j = \omega^j e_j.$$

Write $W_{(a,b)} = X^a Z^b$ for $(a,b) \in G = (\mathbb{Z}/d)^2$. Thus $ZX = \omega XZ$ and the d^2 matrices W_g are orthogonal for the Hilbert–Schmidt inner product. Throughout,

$$C = \sum_{g \in S} c_g W_g \neq 0, \quad c_g \neq 0, \quad s = |S|.$$

Coefficients are arbitrary complex numbers; Hermiticity, positivity and algebraicity are not assumed. Rank always means exact complex rank. All statements concern this cyclic Weyl basis; a tensor product of qubit Pauli bases is a different setting.

For $d = p^k$ and $0 \leq t < d$, define

$$w_p(t) = \prod_{j=0}^{k-1} (t_j + 1), \quad t = \sum_{j=0}^{k-1} t_j p^j, \quad \Omega_k(m) = \max\{t < d : w_p(t) \leq m\}.$$

The scalar Fourier uncertainty law of Delvaux and Van Barel identifies $d - \Omega_k(m)$ as the minimum rank of a nonzero polynomial in Z with at most m terms [2]. The operator conjecture in [1] asks whether

$$\text{rank}(C) \geq d - \Omega_k(s) \quad (1)$$

holds for arbitrary Weyl supports. That paper proves a low-sparsity range and identifies $d = 16$, $s = 6, 7$ as its first unresolved cases. These are questions within that specific research programme, not classical named open problems.

Theorem 1 (Uniform binary seven-term bound). *Let $d = 2^k \geq 4$. Every nonzero combination of at most seven cyclic Weyl operators has rank at least $d/4$. Consequently, for every support budget $4 \leq m \leq 7$,*

$$\min\{\text{rank}(C) : C \neq 0, |\text{supp}_W(C)| \leq m\} = d/4.$$

Here supp_W denotes support in the Weyl basis.

Theorem 2 (Dimension sixteen). *For every nonzero $C \in M_{16}(\mathbb{C})$ with Weyl sparsity s ,*

$$\text{rank}(C) \geq \begin{cases} 16, & s = 1, \\ 8, & 2 \leq s \leq 3, \\ 4, & 4 \leq s \leq 7, \\ 2, & 8 \leq s \leq 15, \\ 1, & 16 \leq s \leq 256. \end{cases}$$

Each bound is attained under the corresponding support budget. Consequently (1) holds in dimension 16 at every sparsity.

Theorem 3 (Further gaps). *In dimension 32, sparsities 5, 6 and 7 imply rank at least 8, and sparsity 11 implies rank at least 4. In dimension 27, sparsities 6, 7 and 8 imply rank at least 6, and sparsity 14 implies rank at least 3.*

The bounds match the conjectured envelope at these parameters. Their distinction from the elementary product bound is recorded in Table 1. They do not settle all sparsities in dimensions 27 or 32.

d	s	$\lceil d/s \rceil$	New lower bound	Proof component
16	6	3	4	uniform binary theorem
16	7	3	4	uniform binary theorem
27	6	5	6	moment bound and periodic blocks
27	7	4	6	moment bound and single-chain blocks
27	8	4	6	moment bound and two-chain blocks
27	14	2	3	coprime support size
32	5	7	8	uniform binary theorem
32	6	6	8	uniform binary theorem
32	7	5	8	uniform binary theorem
32	11	3	4	coprime support size

Table 1: Previously uncovered parameters listed in [1] and closed here. “New” is relative to that stated baseline; the literature search does not certify worldwide priority.

Relation to existing methods. Support uncertainty and additive structure have a long history, including Tao's prime-order theorem [3], Meshulam's divisor-sensitive bound [4], and the scalar prime-power law [2]. Recent work relates scalar uncertainty and recovery to additive energy [5, 6]. Our moment estimate uses the familiar principle that repeated differences control a fourth moment. We also use a classical acyclic-set bound [8]. The proposed contribution is the combination of these tools with a complete classification of the high-overlap exceptions at sparsity at most seven, yielding a uniform operator theorem rather than only isolated dimensions. The proof imports the scalar theorem only in explicitly collinear branches; it does not import an earlier operator uncertainty theorem.

2 An overlap bound and its quantitative form

For $h \in G \setminus \{0\}$ put

$$m_h = |\{g \in S : g + h \in S\}|, \quad \mu(S) = \max_{h \neq 0} m_h.$$

For a singleton support, $\mu = 0$. Let $M = \sum_g |c_g|^2$ and $Q = \sum_g |c_g|^4$.

Proposition 4 (Fourth moment). *For every dimension and every support,*

$$\frac{1}{d} \operatorname{Tr}((C^*C)^2) \leq M^2 + \mu(M^2 - Q), \quad (2)$$

$$\operatorname{rank}(C) \geq \frac{d}{1 + \mu(1 - Q/M^2)} \geq \frac{d}{1 + \mu(1 - 1/s)}. \quad (3)$$

Integer ceilings may be taken in (3).

Proof. For $g = (a, b)$ and $h = (u, v)$ the Weyl product gives

$$W_g^* W_{g+h} = \omega^{-bu} W_h.$$

Thus $C^*C = \sum_h A_h W_h$, where $A_0 = M$ and

$$A_h = \sum_{g, g+h \in S} \overline{c_g} c_{g+h} \omega^{-b_g u}.$$

Since C^*C is Hermitian, trace orthogonality gives

$$\frac{1}{d} \operatorname{Tr}((C^*C)^2) = \sum_h |A_h|^2.$$

Cauchy–Schwarz on each nonzero difference fibre yields

$$|A_h|^2 \leq m_h \sum_{g, g+h \in S} |c_g|^2 |c_{g+h}|^2.$$

Summing over $h \neq 0$ proves (2), because every ordered pair of distinct support points occurs once. If λ_i are the nonzero eigenvalues of C^*C , then

$$(dM)^2 = \left(\sum_i \lambda_i \right)^2 \leq \operatorname{rank}(C) \sum_i \lambda_i^2.$$

Finally $Q \geq M^2/s$. □

We also use the product bound

$$\text{rank}(C)s \geq d. \quad (4)$$

Indeed $\|C\|_{\text{op}} \leq \sum |c_g| \leq \sqrt{sM}$ and $\|C\|_F^2 = dM \leq \text{rank}(C)\|C\|_{\text{op}}^2$. This often outperforms (3) when $\mu = s$; the two bounds should be used together.

Corollary 5 (Aperiodic and coprime supports). *If S has no nonzero translation period, then*

$$\text{rank}(C) \geq \left\lceil \frac{ds}{s^2 - s + 1} \right\rceil.$$

In particular this holds if $\gcd(s, d) = 1$.

Proof. An overlap $m_h = s$ means $S + h = S$. Its orbits have size $\text{ord}(h) > 1$, which divides both s and d . If there is no such period, $\mu \leq s - 1$ and Proposition 4 applies. $\square$

Proposition 6 (Distance from low rank). *Set $K = 1 + \mu(1 - 1/s)$. For $1 \leq r < d$ and every matrix B of rank at most r ,*

$$\frac{\|C - B\|_F^2}{\|C\|_F^2} \geq \max \left\{ 0, 1 - \frac{r}{d} - \frac{\sqrt{r(d-r)(K-1)}}{d} \right\}. \quad (5)$$

One may replace K by $1 + \mu(1 - Q/M^2)$.

Proof. Normalize the eigenvalues of C^*C to have sum one and arrange them decreasingly. Their squared sum is at most K/d . If t is the sum of the largest r , Cauchy–Schwarz separately on the two parts gives

$$\frac{t^2}{r} + \frac{(1-t)^2}{d-r} \leq \frac{K}{d}.$$

Solving the quadratic inequality bounds t by $(r + \sqrt{r(d-r)(K-1)})/d$. The squared Frobenius error of the best rank- r approximation is the tail sum $1 - t$; this follows directly from the singular-value decomposition. $\square$

For $d = 16$, $r = 3$, the bound is strictly positive for $s = 6$, $\mu \leq 5$ and for $s = 7$, $\mu \leq 5$. For the latter it equals

$$\frac{13 - \sqrt{1170/7}}{16} > 0.$$

This certifies separation from rank three even in the presence of bounded matrix error. We do not infer a universal numerical tolerance for the exceptional supports from this formula.

3 Large overlaps force cyclic structure

Lemma 7 (Orbits and exits). *Fix a nonzero $h \in G$ of order q and partition G into cosets of $H = \langle h \rangle$. Then $s - m_h$ is the total number of occupied runs in the partially occupied directed h -cycles. Full and empty cycles contribute zero. In particular:*

- (i) *if $m_h = s$, S is a union of full H -cosets;*
- (ii) *if $m_h = s - 1$, S is a union of full H -cosets and exactly one proper consecutive chain in one other coset.*

Proof. The difference $s - m_h$ counts the points $g \in S$ with $g + h \notin S$. Each occupied run in a partial cycle has exactly one such exit. $\square$

We next give the block decomposition with its phase and injectivity details. Those details matter: distinct global Weyl labels need not remain distinct on an arbitrary subspace.

Lemma 8 (Cyclic compression). *Let $d = p^k$, let $h \in G$ have order q , and set $n = d/q$. After a unitary Weyl-coordinate change, $H = \langle h \rangle$ becomes $\langle (0, n) \rangle$. Write*

$$C = \sum_{i=1}^t W_{(a_i, b_i)} f_i(Z^n)$$

using distinct H -cosets. Let $E_j = \text{span}\{e_{j+qu} : u \in \mathbb{Z}/n\}$ for $j \in \mathbb{Z}/q$, and $\zeta_q = e^{2\pi i/q}$. Coset i contributes to the block from E_j to $E_{j+\delta_i}$, where $\delta_i = a_i \bmod q$, as a scalar $f_i(\zeta_q^j)$ times an n -dimensional Weyl unitary. Within each fixed block, distinct cosets give distinct Weyl labels, up to a common unitary and scalar phases. If every δ_i is the same, the global rank is the sum of the q block ranks.

Proof. First justify the coordinate change. Write $h = gv$, with $g = \gcd(h_1, h_2, d) = n$. Since d is a prime power, v has at least one coordinate coprime to p . Choose u_1, u_2 with $u_1 v_1 + u_2 v_2 = 1 \pmod{d}$. The matrix

$$T = \begin{pmatrix} v_2 & -v_1 \\ u_1 & u_2 \end{pmatrix} \in SL_2(\mathbb{Z}/d)$$

sends h to $(0, n)$. Such a symplectic coordinate change is implemented by a unitary, allowing phases on Weyl operators: its two column Weyls have the same commutation relation as X, Z ; rephase them to have d th powers I . An eigenvector of the second and its successive images under the first form an orthonormal basis, since their eigenvalues run through all d distinct roots. This constructs the required unitary. Phases only change the coefficients and preserve support.

Now use the normalized coordinates. Choose representatives $0 \leq \delta_i < q$ and write $a_i = qA_i + \delta_i$. Put $j' = (j + \delta_i) \bmod q$ and $\kappa = \lfloor (j + \delta_i)/q \rfloor$. On the bases indexed by $u \in \mathbb{Z}/n$,

$$W_{(a_i, b_i)}|_{E_j \rightarrow E_{j'}} = \omega^{b_i j} X_n^{A_i + \kappa} Z_n^{b_i}. \quad (6)$$

The carry κ is common to all terms with a fixed δ_i . After factoring X_n^κ , the labels are

$$(A_i \bmod n, b_i \bmod n).$$

Equal labels with equal δ_i mean equal $a_i \bmod d$ and b_i differing by a multiple of n , precisely equality of the original H -cosets. This proves injectivity. Finally, if all output shifts agree, distinct input sectors map to mutually orthogonal output sectors, so ranks add. $\square$

The binary three-term bound below is already contained in the author's three-term rank classification [9]. We give a shorter overlap-based proof of the inequality needed here; that inequality itself is not a new result of this manuscript.

Lemma 9 (Two elementary rank facts). *(i) A nonzero binomial of distinct Weyls in dimension $n = p^k$ has rank at least $n - n/p$.*

(ii) In dimension $n = 2^k$, a nonzero combination of at most three Weyls has rank at least $n/2$.

Proof. For (i), multiply by a Weyl inverse to obtain $aI + bW_u$, $u \neq 0$. If $q = \text{ord}(u)$, rephase W_u so that its q th power is I . Its eigenvalues are all q th roots, each with multiplicity n/q : indeed $\text{Tr}(W_u^j) = 0$ for $1 \leq j < q$, and Fourier inversion determines the multiplicities. Hence the nullity is at most $n/q \leq n/p$.

For (ii), only exact sparsity three remains after (i). Periodicity is impossible since a nontrivial power of two cannot divide three. If $\mu \leq 1$, Proposition 4 gives rank at least $3n/5 \geq n/2$. If $\mu = 2$, choose a direction of overlap two and use Lemma 7. For its order $q \geq 4$, the support is a consecutive chain of three, so after left multiplication by a Weyl, C is a degree-two polynomial in an order- q Weyl. At most two eigenvalues vanish and the rank is at least $(n/q)(q-2) \geq n/2$.

For $q = 2$, the support is a full pair and a singleton in another coset. Apply Lemma 8. If their output shifts differ, a nonzero singleton block is a scalar unitary of size $n/2$. If they agree, the singleton is nonzero in both input sectors. Both blocks are nonzero combinations of at most two distinct Weyls in dimension $n/2$, so (i) gives rank at least $n/4$ each. If $n/2 = 1$, this last case is impossible by injectivity of block labels. This completes the proof. $\square$

4 A uniform theorem in binary dimension

Lemma 10 (Acyclic block selection). *Suppose that a square matrix is partitioned into equal-sized blocks indexed by a finite set V . On a subset $A \subseteq V$ of N indices, every diagonal block has rank at least ρ , and each block column has at most $t-1$ off-diagonal nonzero blocks with row index in A . Then the matrix has rank at least $\rho \lceil N/t \rceil$.*

Proof. Draw an arc $j \rightarrow i$ for each off-diagonal nonzero block on A . Choose a uniformly random total order and retain every vertex preceding all its out-neighbours. Each vertex is retained with probability at least $1/t$. Every arc between retained vertices increases the order, so the retained graph is acyclic. Some choice retains at least $\lceil N/t \rceil$ vertices. This is a special case of the directed acyclic-set bound in [8, Theorem 4].

Topologically ordering those vertices makes the selected block matrix triangular. Its rank is at least the sum of the diagonal ranks, even when the diagonal blocks are singular: choose an invertible minor of full rank in each diagonal block. Their union is a triangular minor whose determinant is the product of those nonzero determinants. $\square$

In Lemma 8, grouping the cosets by output shift and permuting the output sectors puts any chosen group on the diagonal. If there are t distinct shifts, Lemma 10 applies with that t . In particular, a single coset active on N sectors and one other shift give rank at least $n \lceil N/2 \rceil$. If all shifts agree, ranks add without any sector loss. We call a sector active for a coset when its scalar polynomial does not vanish there. A chain of length ℓ is active on at least $q - \ell + 1$ sectors, by its degree; a singleton is active everywhere.

Proof of Theorem 1. The cases $d = 4, 8$ follow from (4). We may take $d \geq 16$. Sparsities at most three follow from Lemma 9; sparsity four follows from (4). We first prove the five-term assertion in every binary dimension, before using it for seven terms. In each high-overlap case choose h with $m_h = \mu$, put $q = \text{ord}(h)$ and $n = d/q$, and use Lemmas 7 and 8.

Five terms, in all binary dimensions. If $\mu \leq 3$, the moment bound gives $5d/17 > d/4$. Periodicity is impossible, so the only remaining overlap is four, with one partial run. For $q = 2$, the support is two full pairs and a singleton. If all shifts agree, the singleton keeps both blocks nonzero; the three-term bound gives rank at least $n/2$ in each, hence $d/2$. Otherwise the partition by shift is $2 + 1$, so a block containing just one coset is a scalar unitary somewhere, of rank $n = d/2$. For $q = 4$, there is one full cycle and a singleton. Distinct shifts give a unitary block of rank $n = d/4$; a common shift gives four nonzero binomial blocks and total rank at least $2n$. For $q \geq 8$, a single chain of length five has rank at least $n(q-4) \geq d/2$.

Six terms: zero or one exit. If $\mu \leq 3$, the moment bound gives $2d/7 > d/4$. At $\mu = 6$ a period has order two; the three full cosets give nonzero blocks with at most three terms in dimension $n = d/2$, hence a block of rank at least $d/4$. At $\mu = 5$ there is one exit. An involution has even overlap, so $q \neq 2$. At $q = 4$ the support is a full cycle and a two-point chain. Distinct shifts give a unitary block of rank $n = d/4$; a common shift gives at least three nonzero binomial blocks, of total rank at least $3n/2$. For $q \geq 8$, a chain of length six has rank at least $n(q - 5) \geq 3d/8$.

Six terms: two exits. It remains to take $\mu = 4$. At $q = 2$ the support consists of two full pairs and two singletons. A common shift gives two nonzero blocks with at most four terms, of rank at least $n/4$ each by (4). Otherwise the shift partition is $3 + 1$ or $2 + 2$; a scalar-unitary block or a nonzero binomial block has rank at least $n/2 = d/4$.

At $q = 4$ the complete list is two three-point chains; a full cycle and two singleton cosets; or a full cycle and an opposite pair in another coset. In the first case at least two sectors are active for a chain; a common shift therefore gives two binomial blocks of rank at least $n/2$ each, and distinct shifts give a unitary block of rank n . The second case has a singleton shift group unless all shifts agree; in the latter case four nonzero blocks have at most three terms and total rank at least $2n$. In the third case the opposite pair is active on at least two sectors, and the first argument applies. Thus every case gives at least $n = d/4$.

For $q \geq 8$ there are no full cycles. If both runs lie in one coset, the scalar binary uncertainty law [2, Theorem 10, equation (23), TW470] gives rank at least $n(q/4) = d/4$, since the scalar support has six terms. If they lie in two cosets, the shorter chain has length at most three and is active on $N \geq q - 2$ sectors. A common shift gives rank at least $Nn/2$ by the binomial bound; distinct shifts give rank at least $n \lceil N/2 \rceil$ by Lemma 10. Both exceed $d/4$.

Seven terms: one exit. If $\mu \leq 3$, the moment bound gives $7d/25 > d/4$. Periodicity is impossible. At $\mu = 6$ there is one exit. For $q = 2$, there are three full pairs and a singleton. A common shift gives two active blocks with at most four terms and total rank at least $n/2 = d/4$. If shifts split, a nonzero block contains at most three terms, hence has rank at least $n/2$. For $q = 4$, there is a full cycle and a three-point chain. A common shift gives at least two binomial blocks of rank at least $n/2$ each; distinct shifts give a unitary block of rank $n = d/4$. For $q \geq 8$, the chain of length seven gives rank at least $n(q - 6) \geq d/4$.

Seven terms: two exits. At $\mu = 5$, parity excludes $q = 2$. For $q = 4$, the support is a full cycle, a two-point chain and a singleton, in three cosets. If the shifts split, a singleton shift group gives a unitary block of rank $n = d/4$; if all shifts agree, the singleton keeps four blocks active, each of rank at least $n/2$. For $q \geq 8$, one coset is handled by the scalar theorem as above. Otherwise the shorter of two chains has length at most three, and the same $N \geq q - 2$ argument used for six terms applies.

Seven terms: three exits, small orders. At $\mu = 4$, $q = 2$ gives two full pairs and three singletons. If all shifts agree, both blocks are active and have at most five terms. The already proved five-term assertion in dimension $n = d/2$ gives total rank at least $2(n/4) = d/4$. Otherwise the partition by the two possible shifts is $4 + 1$ or $3 + 2$. The first gives a unitary block; in the second a nonzero block in the two-coset group is a binomial, with rank at least $n/2 = d/4$.

For $q = 4$, all possibilities, up to cycle rotation, are

$$(3, 2, 2), \quad (3, 3, 1), \quad (4, 2_{\text{opp}}, 1), \quad (4, 1, 1, 1),$$

where a part below four denotes a consecutive chain unless marked as an opposite pair. This list follows by distributing three exits among partial four-cycles, while full cycles contribute zero. In each three-coset case, a split of shifts has a singleton group and gives a unitary block of rank $n = d/4$. With a common shift, at least two sectors are active: this follows from a singleton in the middle two cases, and a length-two chain in the first. Each active block has at most three terms and rank at least $n/2$. For $(4, 1, 1, 1)$, a common shift gives four nonzero four-term blocks, of total rank at least n . A split with a singleton group again gives a unitary block. The only remaining partition is $2 + 2$. One of the two groups contains a singleton, so its diagonal blocks, after permuting output sectors, are active in all four sectors and have rank at least $n/2$. There is just one other shift. Lemma 10 retains at least two sectors, giving rank at least n .

Seven terms: three exits, orders at least eight. There are no full cycles, and at most three occupied cosets. One coset is again scalar. With two cosets, one contains exactly one run, of length at most five because the other contains two nonempty runs. This chain is active on $N \geq q - 4$ sectors. A common shift gives N binomial blocks of rank at least $n/2$; distinct shifts give $\lceil N/2 \rceil$ invertible diagonal blocks. Both give rank at least $d/4$ for $q \geq 8$.

With three cosets, each has a single chain, and their lengths sum to seven. If all shifts agree, a shortest chain has length at most two and keeps at least $q - 1$ blocks active; the three-term bound gives total rank at least $(q - 1)n/2$. If exactly two shifts agree, the isolated coset has length at most five. Put it on the diagonal: at least $q - 4$ diagonal blocks are invertible and there is one other shift, giving rank at least $n \lceil (q - 4)/2 \rceil \geq d/4$. If the three shifts are distinct, put a shortest chain on the diagonal. At least $q - 1$ diagonal blocks are invertible and there are two other shifts, so Lemma 10 gives

$$\text{rank}(C) \geq n \lceil (q - 1)/3 \rceil \geq nq/4 = d/4 \quad (q \geq 8).$$

This exhausts all overlap values and all orders. When a compressed dimension is one, a nonzero scalar has rank one; any branch requiring distinct labels in the same scalar block is impossible by Lemma 8. Thus no small-block exception has been used implicitly.

Finally the four-term subgroup sum $\sum_{j=0}^3 Z^{jd/4}$ has rank $d/4$, proving sharpness under every budget from four through seven. $\square$

Proof of Theorem 2. One term is unitary. Two and three terms are covered by Lemma 9, four through seven by Theorem 1, and the remaining ranges by (4). The sums $\sum_{j=0}^{2^t-1} Z^{j16/2^t}$, $0 \leq t \leq 4$, attain each bound under the corresponding support budgets. $\square$

5 Dimensions twenty-seven and thirty-two

Proof of Theorem 3. The binary cases with five through seven terms follow from Theorem 1. Two remaining coprime cases follow from Corollary 5:

$$(32, 11) : \lceil 352/111 \rceil = 4, \quad (27, 14) : \lceil 378/183 \rceil = 3.$$

In dimension 27 with six terms, $\mu \leq 5$ gives $\lceil 162/31 \rceil = 6$. If $\mu = 6$, a period has order three, and the support consists of two full cosets. Compression gives blocks of size nine. If the output shifts agree, some nonzero block is a binomial and has rank at least six by Lemma 9(i). If the shifts differ, a nonzero block is a scalar unitary and has rank nine.

With seven terms, $\mu \leq 5$ gives $\lceil 189/37 \rceil = 6$; overlap seven is impossible. At overlap six, orders $q = 9, 27$ give a single seven-point chain. The degree-six argument gives rank at least $(27/q)(q - 6) \geq 9$. For $q = 3$, the support is two full three-cycles and one singleton. If the output shifts are not all equal, each block contains at most two nine-dimensional Weyls, and

a nonzero block has rank at least six. If the shifts are all equal, the singleton makes all three blocks nonzero. Each has at most three terms, so (4) gives rank at least three in each block, hence at least nine in total.

Eight terms. If $\mu \leq 5$, the moment bound gives $\lceil 216/43 \rceil = 6$. Overlap eight is impossible because no nontrivial power of three divides eight.

If $m_h = 7$, there is a single partial chain. For $q = 9, 27$ the chain has length eight and the degree-seven bound gives rank at least $(27/q)(q - 7) \geq 6$. For $q = 3$, there are two full cycles and a two-point chain in a third coset. The chain polynomial is nonzero on at least two of the three sectors. If all output shifts agree, these two nonzero nine-dimensional blocks have at most three terms each and total rank at least six. If the shifts do not all agree, every block has at most two terms and a nonzero block already has rank at least six.

It remains to consider $m_h = 6$, which gives exactly two exits. If $q = 3$, each partial cycle has a single run, and the size count forces two full cycles and two singleton cosets. If all four output shifts agree, all three input blocks are nonzero, have at most four terms in dimension nine, and have rank at least three each. Otherwise partition the four cosets by output shift. A singleton part gives a nonzero scalar-unitary block of rank nine somewhere, since its polynomial is not identically zero on the three roots. If there is no singleton part, the only nontrivial partition is $2 + 2$, and a nonzero binomial block has rank at least six.

For $q = 9$ or 27 , there are no full cycles. If the two runs occupy just one coset, the problem is collinear. The classical scalar law [2, Theorem 10, equation (23), in the TW470 report] gives ranks at least $(27/9)H_9(8) = 6$ or $H_{27}(8) = 6$, respectively, where $H_q(m) = q - \Omega(m)$.

If the runs occupy two different cosets, their lengths sum to eight. The shorter has length at most four, so its polynomial f_1 is nonzero in at least $N \geq q - 3$ input sectors. If the output shifts agree, then $q = 9$ (the case $q = 27$, block size one, has only one coset label per block). The $N \geq 6$ nonzero blocks are binomials in dimension three, each of rank at least two.

Finally suppose the output shifts differ. After permuting output sectors, the first coset supplies invertible diagonal blocks on those N sectors; the second supplies blocks on the directed edges of translation by a nonzero $\delta \in \mathbb{Z}/q$. Let $r = \text{ord}(\delta) \geq 3$. Remove at most one active sector from every translation cycle that is entirely active. The remaining set I has

$$|I| \geq N - q/r \geq q - 3 - q/3.$$

The submatrix with input and output sectors in I has invertible diagonal blocks and an acyclic directed graph of off-diagonal blocks. A topological ordering makes it block triangular, hence its rank is $n|I|$, where $n = 27/q$. This is at least nine when $q = 9$ and at least fifteen when $q = 27$, both more than required. $\square$

The bounds are sharp as budget statements at these parameters. In dimension 32, subgroup sums of four and eight powers of Z have ranks eight and four. In dimension 27, writing $\zeta = e^{2\pi i/3}$,

$$(Z^{18} + Z^9 + I)(Z^3 - \zeta I)$$

has exactly six terms and rank six: the first factor vanishes on eighteen eigenvalues, the second on three further eigenvalues. A sum of nine subgroup powers has rank three. These witnesses also explain why one must distinguish a budget from exact sparsity.

6 Consequences, verification and limitations

Sharp erasure thresholds. For a set S of omitted Weyl labels, let $\mathcal{M}_{Sc}(A)$ be the list of coefficients $d^{-1} \text{Tr}(W_g^* A)$ for $g \notin S$.

Corollary 11 (Binary erasure budget). *Let $d = 2^k \geq 16$ and let r be an integer with $d/16 \leq r < d/8$. The map $\mathcal{M}_{S^c}$ is injective on all complex matrices of rank at most r whenever $|S| \leq 7$. Some set of eight omitted labels makes it noninjective. The universal erasure budget is therefore exactly seven.*

Proof. A difference of two rank-at-most- r matrices has rank at most $2r < d/4$. If it is supported on at most seven Weyl labels, Theorem 1 forces it to vanish. For failure at eight, take $S = \{(0, jd/8) : 0 \leq j \leq 7\}$. The sum

$$P = \sum_{j=0}^7 Z^{jd/8} = 8 \sum_{\substack{0 \leq u < d \\ 8|u}} E_{uu}$$

has rank $d/8 \leq 2r$. Partition its nonzero diagonal positions into two groups of size at most r . Let A equal P on the first group and let B equal $-P$ on the second, with zero entries elsewhere. Then $A - B = P$ and the measured coefficients agree. $\square$

The product bound alone guarantees $\lceil d/(2r) \rceil - 1$ erasures. For example, at $(d, r) = (32, 3)$ it guarantees five, whereas Corollary 11 gives the exact value seven. This is a uniform family of rank-sensitive statements, not a claim about arbitrary ranks or all sparsity ranges.

Corollary 12 (Universal rank-two erasure budget). *In dimension 27, the map $\mathcal{M}_{S^c}$ is injective on the set of all complex matrices of rank at most two for every S with $|S| \leq 8$. There exists a set of nine labels for which it is not injective. Thus the largest universal erasure budget is exactly eight.*

Proof. If A, B both have rank at most two and agree on the measured coefficients, then $D = A - B$ has rank at most four and Weyl support contained in S . A nonzero D with at most five terms has rank at least six by (4); with six, seven or eight terms it has rank at least six by Theorem 3. Hence $D = 0$.

For failure at nine, take $S = \{(0, 3j) : 0 \leq j \leq 8\}$ and

$$P = \sum_{j=0}^8 Z^{3j} = 9(E_{00} + E_{99} + E_{18,18}).$$

With $A = 9(E_{00} + E_{99})$ and $B = -9E_{18,18}$, the distinct matrices A, B have ranks two and one, and $A - B = P$ is supported in S . $\square$

The product estimate alone guarantees only six erasures in this setting, since $27 > 4 \cdot 6$ but $27 \leq 4 \cdot 7$. The counterexample above is not a pair of density matrices: B is negative, and the identity label is among the omissions. The corollary concerns arbitrary complex matrices, with no positivity or fixed-trace constraint. It proves noiseless identifiability, not an efficient reconstruction algorithm. Related Weyl-support obstructions for Hermitian traceless rank-two matrices and phase retrieval are studied by Yin [7]; we do not claim a new universal phase-retrieval threshold.

Quantitative certificates depend on the support. For fixed S satisfying a sufficiently small overlap bound, Proposition 6 gives an explicit error margin for the same low-rank obstruction. If the effective nonempty support is $T \subseteq S$, then $\mu(T) \leq \mu(S)$ and $1 - 1/|T| \leq 1 - 1/|S|$. Thus $K_S = 1 + \mu(S)(1 - 1/|S|)$ is already a valid uniform choice for all those subsets. The integer improvements in Table 1 alone do not constitute uniform, dimension-independent conditioning estimates.

Reproducibility. The accompanying `verify_overlap.py` records exact rational checks of the moment thresholds; 66,332 exhaustive binary cycle patterns at orders 2, 3, 4, 8, 9, 16; 108,201 exact integer checks of block indices and phases; and 2,069 determinant-one normalizers in dimensions 8, 16, 27, 32. Sharp budget witnesses are certified by polynomial greatest common divisors or cyclotomic remainders over $\mathbb{Q}$. The complementary `verify_two_runs.py` checks the eight-term moment threshold, the order-three occupancy count and 90,000 acyclic submatrix selections. The file `verify_erasure.py` checks all 729 coefficients of the sharp dimension-27 erasure witness over $\mathbb{Q}(\zeta_3)$.

The reviewer’s `astra-binary-plateau-checks.py` enumerates the short-orbit profiles used in the uniform theorem. It checks the random-order selection on all 2,401 loopless four-vertex digraphs with maximum outdegree two and all 24 orders; 406 exact branch inequalities at power-of-two orders through 2^{60} ; and 131,071 pairs (d, r) for the erasure arithmetic and its comparison with the product bound. These finite checks support the written classification and proof, and do not prove the arbitrary-dimension quantifier by enumeration.

A separate set of 480 floating-point instances checks the twisted fourth-moment identity and tail inequalities; its maximum relative identity residual was below 1.8×10^{-15} . Those numerical checks are diagnostics, not proofs of the universal statements. No exhaustive search over all complex coefficient vectors is asserted.

What remains open. The argument does not settle the general prime-power envelope, nor even the complete binary envelope in every dimension. In particular it leaves $d = 32$, $s = 12, \dots, 15$, and $d = 27$, $s = 15, 16, 17$ from the previous gap list. Theorem 1 is uniform in the dimension but restricted to sparsity at most seven. More generally, overlap $s - j$ permits j partial cyclic runs. Extending the exceptional-support classification beyond this range is a concrete next question. The overlap bound can be loose because it discards phase cancellation and replaces all difference multiplicities by their maximum.

Research and AI disclosure. This manuscript arose from a catalog-level sweep of mathematical work, followed by selected full-text reading. The sweep is documented separately; it is not a claim to have verified every proof in the surveyed corpus. Codex assisted the search, derivations, writing and computational checks. A GPT-6 Astra reviewer with shared task context supplied the binary three-term argument, audited the uniform seven-term extension and identified the graph reference. This is mathematical involvement, not an independent blind review. Model reports, when supplied with the frozen artifact, are assessments rather than human peer review or formal proof verification. No new Opus review is claimed. Worldwide novelty remains subject to further literature review.

References

- [1] L. Eriksson, *Sparse cyclic Weyl operators in prime-power dimension: an operator extension of the scalar uncertainty law through sparsity $2p - 1$* , AIRR.SCIENCE, ARR-2026-1KTM07SEEV968BCB, v1 (2026). <https://airr.science/papers/ARR-2026-1KTM07SEEV968BCB/versions/v1/>.
- [2] S. Delvaux and M. Van Barel, Rank-deficient submatrices of Fourier matrices, *Linear Algebra Appl.* **429** (2008), 1587–1605. doi:10.1016/j.laa.2008.04.043. Author report TW470 (2006), whose theorem numbering is used here: <https://www.cs.kuleuven.be/publicaties/rapporten/tw/TW470.pdf>.
- [3] T. Tao, An uncertainty principle for cyclic groups of prime order, *Math. Res. Lett.* **12** (2005), 121–127. <https://arxiv.org/abs/math/0308286>.

- [4] R. Meshulam, An uncertainty inequality for finite abelian groups, *European J. Combin.* **27** (2006), 63–67. <https://arxiv.org/abs/math/0312407>.
- [5] K. Aldahleh, A. Iosevich, J. Iosevich, J. Jaimangal, A. Mayeli and S. Pack, *Additive energy, uncertainty principle and signal recovery mechanisms*, arXiv:2504.14702 (2025). <https://arxiv.org/abs/2504.14702>.
- [6] I. Bortnovskyi et al., *Refined additive uncertainty principle*, arXiv:2510.26664 (2025). <https://arxiv.org/abs/2510.26664>.
- [7] P. Yin, *A Sharp Rank-Two Uncertainty Principle for Finite Weyl Transforms*, public manuscript and accompanying formalization repository, accessed 15 October 2026. <https://github.com/holdonyb/weyl-stft-paper>. The repository’s formalization claims were not independently replayed for this paper.
- [8] H. Gruber, Bounding the feedback vertex number of digraphs in terms of vertex degrees, *Discrete Appl. Math.* **159** (2011), 872–875. <https://arxiv.org/abs/1101.1291>.
- [9] L. Eriksson, *Three-term Weyl operators: central multiplicities, rank laws, and exact list decoding*, AIRR.SCIENCE, ARR-2026-0WESHW4YMM9FG8EK, v1 (2026). <https://airr.science/papers/ARR-2026-0WESHW4YMM9FG8EK/versions/v1/>.