A sharp stability transition for low-rank tomography with four missing Weyl coefficients

Lluis Eriksson
Independent researcher, Sweden

24 October 2026
Working draft 2; documentary clarifications

Abstract

Let $d = 9n$ with arbitrary integer $n \geq 1$, and omit the four cyclic Weyl coefficients associated with $X_d^{d/3}$, its adjoint, $Z_d^{d/3}$ and its adjoint. We determine the exact inverse Frobenius constant on differences of density matrices of rank at most r. Writing $x = r/n$, the squared fraction of energy that can be hidden in the missing coefficients is $2x/3$ for $0 < x \leq 1$, and $(x + \sqrt{x^2 - 3x + 3})/3$ for $1 \leq x < 2$. The resulting inverse constant is attained by explicit orthogonally supported states and diverges as $\sqrt{2n/(2n-r)}$ at the exact nonidentifiability threshold $r = 2n$. The same threshold holds when a rank-r state must be distinguished from every positive state of unrestricted rank. For this stronger task we give a separate, conservative stability bound and a convex least-squares recovery consequence. The proof reduces the missing-coefficient space to additive 3 by 3 spectral grids, then solves two four-dimensional symmetric matrix pencils. These are results for one fixed erasure pattern, not for arbitrary four erasures or a measurement scheme with subquadratically many coefficients.

1 Question, normalization and antecedents

A measurement may fail to determine arbitrary quantum states while still determining every state below a rank threshold. In that situation a quantitative question remains: how much can the omitted data amplify errors even before injectivity fails? We solve that question exactly for a four-dimensional missing-coefficient space.

Put $\omega_d = e^{2\pi i/d}$ and

$$X_d e_j = e_{j+1 \bmod d}, \qquad Z_d e_j = \omega_d^j e_j, \qquad W_d(a,b) = X_d^a Z_d^b.$$

The d^2 Weyl matrices are orthogonal, with squared Frobenius norm d. Throughout, $d = 9n$ and n is a positive integer. Define

$$U = X_d^{d/3}, \qquad V = Z_d^{d/3}, \qquad E = \{U, U^\dagger, V, V^\dagger\}.$$

These are four distinct nonidentity Weyl matrices. Let P be the Frobenius-orthogonal projection onto their complex span, and put $Q = I - P$. On Hermitian matrices P is the projection onto the real space

$$\mathcal{K} = \{aU + \bar{a}U^\dagger + bV + \bar{b}V^\dagger : a, b \in \mathbb{C}\}. \tag{1}$$

It consists of traceless Hermitian matrices. Thus $Q\rho$ means all Weyl coefficients except these four, including the known trace coefficient. For $C = \sum_w c_w W_w$ our data norm is $\|QC\|_F^2 = d \sum_{w \notin E} |c_w|^2$. All constants below refer to this normalization, not directly to a statistical probability norm.

This coefficient space can be realized by a POVM, although outcome norms then require a change of coordinates. Choose a real Hermitian basis $F_1, \dots, F_m$ of the traceless part of $\mathcal{K}^\perp$, where $m = d^2 - 5$, and rescale each element so that $\|F_j\|_{\mathrm{op}} \le 1$. The $2m$ effects $(I + F_j)/(2m), (I - F_j)/(2m)$ are positive, sum to I, and span $\mathcal{K}^\perp$. Their probability record determines $Q\rho$. The constants below do not include the conditioning of this coordinate conversion or the sampling error of this particular realization.

Let $\mathcal{D}_r = \{\rho \ge 0 : \operatorname{Tr}\rho = 1, \operatorname{rank}\rho \le r\}$ and define

$$C(n,r) = \sup_{\substack{\rho,\sigma\in\mathcal{D}_r \\ \rho\neq\sigma}} \frac{\|\rho-\sigma\|_F}{\|Q(\rho-\sigma)\|_F}, \tag{2}$$

where a zero denominator gives infinity. Ranks are integers.

The general kernel and inertia characterization of restricted quantum state determination is due to Carmeli, Heinosaari, Schultz and Toigo [1]. Their Section 6.2, equations (12)–(13) and Proposition 13, also establish the relation between symmetric Weyl zero sets and realizable covariant POVM kernels. The elementary realization above is not claimed as a new measurement-design principle. Baldwin, Deutsch and Kalev [2], especially Sections II–III, distinguish rank-r completeness from rank-r strict-completeness and explain the latter's relation to convex reconstruction and noise robustness. We use this established distinction. We supply direct proofs of the specialized inertia and stability claims rather than assert a new general criterion.

Recent work of Zhu and Wang [3] studies spectral floors and conditioning for minimal rank-one Weyl–Heisenberg covariant measurements. Their full-space Gram-spectrum problem differs from the restricted state-difference constant of the singular coefficient projection considered here. The author's earlier Weyl rank-uncertainty reductions [4] concern support versus arbitrary complex matrix rank, rather than the exact positive-state conditioning law proved below. Spectral rearrangement, orthogonal projection duality and symmetric matrix pencils are classical tools; the contribution claimed here is their explicit solution for the kernel (1). This comparison is not a certification of worldwide priority.

2 Exact stability and identifiability

Theorem 1 (Sharp balanced-rank stability)**.** *For $1 \le r < 2n$ let $x = r/n$ and*

$$L(x) = \begin{cases} 2x/3, & 0 < x \le 1, \\ \bigl(x + \sqrt{x^2 - 3x + 3}\bigr)/3, & 1 \le x < 2. \end{cases}$$

Then

$$C(n,r) = (1 - L(r/n))^{-1/2}. \tag{3}$$

The supremum is attained by differences of states of rank exactly r with mutually orthogonal supports. For $r \ge 2n$ it is infinite. In particular $C(n,n) = \sqrt{3}$, and along integer sequences with $(2n-r)/n \to 0$ from above,

$$C(n,r) \sim \sqrt{\frac{2n}{2n-r}}.$$

Theorem 2 (Unrestricted positive competitors)**.** *Every nonzero $H \in \mathcal{K}$ has at least $2n$ strictly positive and at least $2n$ strictly negative eigenvalues. Consequently Q distinguishes every state*

of rank at most r from every other density matrix, of any rank, if and only if $r < 2n$. In this range, for $\rho \in \mathcal{D}_r$ and any density matrix σ,

$$\|\rho - \sigma\|_F \le B(n,r)\,\|Q(\rho-\sigma)\|_F\,, \qquad B(n,r) = \sqrt{1 + \frac{288n}{2n-r}}. \tag{4}$$

The bound B is conservative; it is not asserted to be the optimal constant for this larger competitor class.

The finite constants C and B have different domains. In particular, (3) cannot be substituted into (4) merely because strict-completeness holds. For $n = 1$ the finite range contains only $r = 1$; the intermediate branch becomes nontrivial for $n \ge 2$.

3 Spectral grids and the state-difference cone

The unitaries U, V commute, have order three, and generate nine distinct Weyl labels. Every nonidentity U^iV^j, $0 \le i, j < 3$, has trace zero. Character orthogonality therefore shows that each of the nine joint eigenspaces has dimension n.

By the real three-point Fourier transform, every $H \in \mathcal{K}$ has joint eigenvalues

$$h_{ij} = f_i + g_j, \qquad \sum_{i=1}^{3} f_i = \sum_{j=1}^{3} g_j = 0, \tag{5}$$

each repeated n times, where f, g are arbitrary real zero-mean triples. Conversely every such pair occurs in (1). In particular

$$\|H\|_F^2 = 3n\left(\sum_i f_i^2 + \sum_j g_j^2\right). \tag{6}$$

No restriction to dimensions that are powers of three is needed.

Write $n_+(D), n_-(D)$ for the positive and negative inertias of a Hermitian matrix. Up to a positive scalar, the nonzero differences of states in $\mathcal{D}_r$ are exactly

$$\mathcal{C}_r = \{D = D^\dagger \ne 0 : \operatorname{Tr} D = 0,\ n_+(D), n_-(D) \le r\}.$$

Indeed, a state difference has these inertia bounds by the min–max principle. Conversely the positive and negative spectral parts of such a D have a common nonzero trace τ; D/τ is the difference of their trace-one normalizations. Hence the key extremum is

$$L_* = \sup_{D \in \mathcal{C}_r} \frac{\|PD\|_F^2}{\|D\|_F^2}. \tag{7}$$

Orthogonality gives $C(n,r) = (1 - L_*)^{-1/2}$ whenever $L_* < 1$.

Lemma 3 (A common-shift dual bound)**.** *For an integer $1 \le r < 2n$ and $H = H^\dagger$ with eigenvalues $h_1 \ge \cdots \ge h_d$, put*

$$W_r(H) = \min_{\theta \in \mathbb{R}} \sum_{i=1}^{r} [(h_i - \theta)^2 + (h_{d-i+1} - \theta)^2].$$

For $D \in \mathcal{C}_r$, $|\operatorname{Tr}(HD)|^2 \le W_r(H)\,\|D\|_F^2$. Thus

$$L_* \le \sup_{0 \ne H \in \mathcal{K}} \frac{W_r(H)}{\|H\|_F^2}. \tag{8}$$

Proof. Pad the positive eigenvalues p_i of D and negative magnitudes q_i with zeros to length r, ordering both decreasingly. Their sums agree. Trace rearrangement bounds $\operatorname{Tr}(HD)$ by $\sum_{i\le r} h_i p_i - \sum_{i\le r} h_{d-i+1} q_i$. For completeness, squared overlaps of two orthonormal eigenbases give a doubly stochastic matrix; the extremal pairing of two ordered real lists bounds the trace. The same conclusion follows from weighted spectral variational inequalities. This argument does not require D to commute with H.

Subtract $\theta(\sum p_i - \sum q_i) = 0$ and apply Cauchy–Schwarz. Apply the same reasoning to $-H$ to control the absolute value, noting $W_r(-H) = W_r(H)$. Finally $\|PD\|_F = \max_{H\in\mathcal{K},\|H\|_F=1} |\operatorname{Tr}(HD)|$, which proves (8). □

4 Solving the dual extremum

4.1 Ranks at most one joint multiplicity

For $r \le n$ the extreme values of H each occur at least n times, so

$$W_r(H) = \frac{r}{2}(\max h - \min h)^2.$$

For a zero-mean triple, $\operatorname{diam}(f)^2 \le 2\sum f_i^2$. Since $\operatorname{diam}(h) = \operatorname{diam}(f) + \operatorname{diam}(g)$, (6) gives $W_r(H)/\|H\|_F^2 \le 2r/(3n)$.

Take $f = (1, 0, -1)$ and $g = -f$. The grid has extrema ± 2 on spaces of dimension n, and $\|H\|_F^2 = 12n$. Choose rank-r projectors R_+, R_- in those two spaces and set $D = 2(R_+ - R_-)$. Then $\operatorname{Tr}(HD) = \|D\|_F^2 = 8r$. Projection duality yields

$$\frac{\|PD\|_F^2}{\|D\|_F^2} \ge \frac{|\operatorname{Tr}(HD)|^2}{\|H\|_F^2\,\|D\|_F^2} = \frac{2r}{3n}.$$

This matches the upper bound, and $R_+/r, R_-/r$ are the required states.

4.2 Ranks between one and two joint multiplicities

Now $n \le r < 2n$, $x = r/n$, and $z = x - 1$. Sort f, g decreasingly and put

$$a = f_1 - f_2, \quad b = f_2 - f_3, \quad c = g_1 - g_2, \quad e = g_2 - g_3.$$

These gaps are nonnegative. Define

$$K = a^2 + ab + b^2 + c^2 + ce + e^2, \quad A = a + b + c + e, \quad u = \min(a, c), \quad v = \min(b, e).$$

Then $\|H\|_F^2 = 2nK$. The largest two cells are $h_{\max}, h_{\max} - u$ and the smallest two $h_{\min}, h_{\min} + v$, with weights $n, r - n$ at each end. Ties and zero gaps cause no problem: consistent cell ordering selects the same multiset, with partial cells of dimension $r - n$.

The shift minimizing W_r is

$$\theta = \frac{h_{\max} + h_{\min}}{2} + \frac{z(v - u)}{2x}.$$

Expanding the squares gives

$$\frac{W_r(H)}{n} = \frac{xA^2}{2} - zA(u + v) + z(u^2 + v^2) - \frac{z^2(v - u)^2}{2x}. \tag{9}$$

The single common shift is essential because the positive and negative parts of a state difference must have the same trace.

There are four choices for the two minima; exchange of f, g reduces them to $(u, v) = (a, b)$ and (a, e). Write $w = (a, b, c, e)^T$ and

$$M = \begin{pmatrix} 2 & 1 & 0 & 0 \\ 1 & 2 & 0 & 0 \\ 0 & 0 & 2 & 1 \\ 0 & 0 & 1 & 2 \end{pmatrix} > 0.$$

Then $2K = w^T M w$, and the numerator in (9) is $w^T N w$, with

$$N_{\text{same}} = \begin{pmatrix} 1 - \frac{1}{2x} & \frac{1}{2x} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2x} & 1 - \frac{1}{2x} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{x}{2} & \frac{x}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{x}{2} & \frac{x}{2} \end{pmatrix},$$

$$N_{\text{cross}} = \begin{pmatrix} 1 - \frac{1}{2x} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2x} \\ \frac{1}{2} & \frac{x}{2} & \frac{x}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{x}{2} & \frac{x}{2} & \frac{1}{2} \\ \frac{1}{2x} & \frac{1}{2} & \frac{1}{2} & 1 - \frac{1}{2x} \end{pmatrix}.$$

Their exact pencil determinants are

$$\det(N_{\text{same}} - \lambda M) = -\frac{\lambda}{x}(\lambda x - x + 1)[-9\lambda^2 + 3(x+1)\lambda - x + 1], \tag{10}$$

$$\det(N_{\text{cross}} - \lambda M) = -\frac{\lambda}{x}(3\lambda x - 2x + 2)[-3\lambda^2 + 2x\lambda - x + 1]. \tag{11}$$

These identities can be checked by direct four-by-four expansion. As M is positive definite, the largest generalized eigenvalue bounds the quotient even after relaxing the gap and branch inequalities.

Put $\ell = (x + \sqrt{x^2 - 3x + 3})/3$. It is the larger root of $3\lambda^2 - 2x\lambda + x - 1 = 0$ and is at least $2/3$ for $1 \le x \le 2$. The other roots of (11) are 0, $2(x-1)/(3x) \le 1/3$, and the smaller quadratic root. Thus the cross branch has maximum generalized eigenvalue ℓ.

For the same branch the additional roots are 0 and $(x-1)/x \le 1/2$. With $q(\lambda) = 9\lambda^2 - 3(x+1)\lambda + x - 1$,

$$q(\ell) = (x-1)(3\ell - 2) \ge 0.$$

Moreover ℓ lies to the right of the vertex $(x+1)/6 \le 1/2$. It follows that the larger root of q is at most ℓ. This proves $L_* \le \ell$ in all branches.

5 Extremizing states and the singular boundary

For $n \le r < 2n$ put

$$t = 1 - x + \sqrt{x^2 - 3x + 3}, \qquad f = \left(\frac{2+t}{3}, \frac{t-1}{3}, -\frac{1+2t}{3}\right), \quad g = -f.$$

Here $0 < t \le 1$, and

$$t^2 + 2(x-1)t + x - 2 = 0. \tag{12}$$

The grid of H has eigenvalues 0 on $3n$ dimensions and $\pm(1+t), \pm 1, \pm t$, each listed sign repeated n times. At $t = 1$ distinct cells may share an eigenvalue; the indicated cell decomposition is still valid.

Retain the full n-dimensional cells of $\pm(1+t)$, rank-$(r-n)$ subspaces in the cells of ± 1, and zero elsewhere, obtaining D. Then $\operatorname{Tr} D = 0$, $n_\pm(D) = r$, and

$$\operatorname{Tr}(HD) = \|D\|_F^2 = 2[n(1+t)^2 + (r-n)], \qquad \|H\|_F^2 = 2n[(1+t)^2 + 1 + t^2].$$

Using (12),

$$\frac{\|D\|_F^2}{\|H\|_F^2} = \frac{(1+t)^2 + x - 1}{2(1+t+t^2)} = \frac{x + \sqrt{x^2 - 3x + 3}}{3} = \ell.$$

Projection duality forces $\|PD\|_F^2 / \|D\|_F^2 \geq \ell$, so equality holds by the previously proved upper bound. The positive and negative parts of D, divided by their common trace $n(1+t)+(r-n)$, are explicit orthogonally supported rank-r states.

At $r = 2n$, the matrix

$$H_0 = U + U^\dagger - V - V^\dagger$$

has eigenvalues $3, -3$ each on $2n$ dimensions and zero elsewhere. Its positive and negative spectral parts, each divided by $6n$, are distinct rank-$2n$ states whose difference lies in $\mathcal{K}$. This proves failure for all $r \geq 2n$. For $x < 2$, $L(x) < 1$, completing (3). Finally $L'(2^-) = 1/2$, whence $1 - L(x) \sim (2-x)/2$ and the stated asymptotic follows.

6 Strict-completeness and convex reconstruction

Proof of Theorem 2. For a nonzero grid (5), sort both triples and put $F = f_1 \geq 0$, $G = g_1 \geq 0$, $S = F + G > 0$. Zero mean and ordering give $f_2 \geq -F/2$, $g_2 \geq -G/2$. The second largest cell satisfies

$$h_2 = \max(F + g_2, f_2 + G) \geq \max(F - G/2, G - F/2) \geq S/4 > 0.$$

Thus at least two cells, each of dimension n, are strictly positive. Apply the argument to $-H$ for the negative cells.

We also have $\sum f_i^2 \leq 6F^2$: indeed $f_2 \in [-F/2, F]$ and the convex quadratic $F^2 + f_2^2 + (F + f_2)^2$ is at most $6F^2$ on that interval. Similarly $\sum g_j^2 \leq 6G^2$. Consequently

$$\|H\|_F^2 \leq 18n(F^2 + G^2) \leq 18nS^2.$$

The largest $2n$ positive eigenvalues are therefore all at least $\|H\|_F / \sqrt{288n}$, and the same bound holds for the $2n$ largest negative magnitudes.

Suppose $D = D^\dagger$, $\operatorname{Tr} D = 0$, and either $n_+(D) \leq r$ or $n_-(D) \leq r$. Set $H = PD$. If $H = 0$ there is nothing to prove. In the first case the Hermitian eigenvalue distance inequality bounds $\|D - H\|_F^2$ below by the sum of squares of the positive eigenvalues of H after its r largest ones. To see this, sort the eigenvalues of both matrices; D has at most r positive entries, so each subsequent positive eigenvalue of H contributes at least its square to their squared Euclidean distance, which is bounded above by $\|D - H\|_F^2$. This last comparison follows from trace rearrangement: if d_i, h_i are the respective decreasing eigenvalue lists, then

$$\|D - H\|_F^2 = \|D\|_F^2 + \|H\|_F^2 - 2\operatorname{Tr}(DH) \geq \sum_i (d_i - h_i)^2.$$

Hence

$$\|QD\|_F^2 \geq \frac{2n - r}{288n} \|PD\|_F^2 .$$

Apply the same argument to $-D, -H$ in the second case. Pythagoras now gives $\|D\|_F \leq B(n,r) \|QD\|_F$.

If $\rho \in \mathcal{D}_r$ and σ is any density matrix, the quadratic form of $D = \rho - \sigma$ is nonpositive on $\ker \rho$. By min–max, $n_+(D) \leq r$, so (4) follows. It implies uniqueness for $r < 2n$. The states constructed from H_0 prove failure at and above $2n$. □

Corollary 4 (A convex noisy-data estimator)**.** *Let $1 \leq r < 2n$, $\rho \in \mathcal{D}_r$, and $y = Q\rho + e$ with $\|e\|_F \leq \epsilon$. Every minimizer*

$$\widehat{\rho} \in \arg \min_{\sigma \geq 0,\ \operatorname{Tr}\sigma = 1} \|Q\sigma - y\|_F$$

satisfies $\|\widehat{\rho} - \rho\|_F \leq 2B(n,r)\epsilon$.

Proof. The density-matrix domain is nonempty, compact and convex, so a minimizer exists. Its residual is at most that of ρ, hence at most ϵ. The triangle inequality gives $\|Q(\widehat{\rho} - \rho)\|_F \leq 2\epsilon$, and Theorem 2 applies without a rank restriction on $\widehat{\rho}$. □

This is a deterministic guarantee in coefficient space. It supplies no experimental sampling bound, approximate-rank guarantee or running-time analysis. In particular it does not assert an improvement over compressed sensing protocols using many fewer measured coefficients.

7 Scope, checks and remaining questions

The argument depends only on the additive spectral kernel with nine equal-multiplicity cells. It therefore applies after arbitrary unitary conjugation and to any commuting order-three pair with all nine joint multiplicities equal to n. The cyclic Weyl construction realizes this kernel in every dimension divisible by nine.

Exact symbolic checks reproduce both pencil determinants and the attainer identities, including a direct computation of projection energy. Finite exact tests additionally compare sorted spectra with the formulas in 2,400 formula checks on 2,000 distinct instances, using $n = 1, \ldots, 5$ and 80 nonzero gap vectors. The 400 instances at $r = n$ are checked in both branches. These tests supplement the universal proofs; they do not establish unexamined dimensions by extrapolation.

The main open quantitative question here is the optimal constant in (4). Further questions concern unequal joint multiplicities, other erasure patterns, and error bounds for approximately low-rank states. The present result does not classify those problems. Only four coefficients are removed out of d^2; the theorem isolates an exact conditioning mechanism rather than a large measurement reduction.

AI and responsibility disclosure. The named author directed substantial OpenAI Codex assistance in derivation, writing and exact computational checks. GPT-6 Astra xhigh mathematical audits used shared filesystem and declared reused context. No runtime isolation, independent human review, Lean formalization, certified worldwide originality or predicted assessment score is claimed. A full assessment of draft 1 is preserved separately. This draft makes documentary clarifications; no new full assessment score is assigned to it.

References

[1] C. Carmeli, T. Heinosaari, J. Schultz and A. Toigo, *Tasks and premises in quantum state determination*, arXiv:1308.5502v2 (2014). https://arxiv.org/abs/1308.5502v2.

[2] C. H. Baldwin, I. H. Deutsch and A. Kalev, *Strictly-complete measurements for bounded-rank quantum-state tomography*, Physical Review A **93**, 052105 (2016). https://doi.org/10.1103/PhysRevA.93.052105; arXiv:1605.02109v1.

[3] X. Zhu and Y. Wang, *Uniformly Stable Minimal Weyl–Heisenberg Measurements Approaching the SIC Benchmark*, arXiv:2608.11850v1 (2026). https://arxiv.org/abs/2608.11850v1.

[4] L. Eriksson, *Removing the binary and ternary factors from cyclic Weyl rank uncertainty*, AIRR.SCIENCE, ARR-2026-1QEPAQ29TN8FZRK5, v1 (2026). https://airr.science/papers/ARR-2026-1QEPAQ29TN8FZRK5/versions/v1/.