Sharp Onset and Unbounded Growth of Norm-Optimal Self-Commutator Rank

Lluis Eriksson

30 August 2026

Abstract

For a traceless Hermitian matrix $F \in M_d(\mathbb{C})$, define

$$\kappa_d(F) = \frac{1}{2}\min\{\|C\|_{\mathrm{HS}}^2 : CC^* - C^*C = 2F\}$$

and let $r_*(F)$ be the least rank among its minimizers. Inertia gives the universal lower bound $r_0(F) = \max\{n_+(F), n_-(F)\}$. We prove the exact low-dimensional classification

$$r_*(F) = r_0(F) \qquad (d \le 7),$$

including every singular spectrum. The bound fails sharply in dimension eight. For the entire explicit two-ray cone

$$\lambda(b, c) = (4c - 3b, b, b, b, -c, -c, -c, -c), \qquad 0 \le b \le c,$$

we determine the unrestricted and rank-four costs: their difference is $\min\{b, c - b\}$, and every interior point has $r_* = 5 > 4 = r_0$. In particular, the integer spectrum $(5, 1, 1, 1, -2, -2, -2, -2)$ has unrestricted cost 13 and rank-four cost 14. Thus eight is the least dimension in which Hilbert–Schmidt norm optimality can force factor rank strictly above inertia. The proof is an exact computer-assisted Horn-polyhedral argument. Thirty-three sign–zero strata through dimension seven yield 813 exact face rays and 272 projected facets; every facet has a rational nonnegative Farkas certificate on the unrestricted cone. A separate exact replay verifies the dimension-eight phase and a two-valued dimension-nine witness. No floating point output is used as a proof object. Finally, for the target G_t with spectrum $(5^{12t}, -4^{15t})$ in dimension $27t$, we prove

$$\kappa_{27t}(G_t) = 87t, \qquad 17t + 1 \le r_*(G_t) \le 18t.$$

Since $r_0(G_t) = 15t$, the additive excess is at least $2t + 1$ and is therefore unbounded. The amplification is symbolic: a coarse-grid restriction of any fine hive transfers a frozen order-27 rank-17 lower certificate to every t by telescoping elementary rhombi.

1 Introduction

Finite-dimensional trace-zero matrices are commutators [2], but prescribing a Hermitian target as one self-commutator is a more structured inverse problem. We ask not only for

$$CC^* - C^*C = 2F,$$

but for a factor C of minimum Hilbert–Schmidt norm and, within that optimal face, minimum rank. This differs from forward Frobenius commutator bounds and from infinite-dimensional compact-operator questions [11, 22, 4, 5].

The closest adjacent results address different optimization variables. Norm estimates, approximation by self-commutators, and Hilbert–Schmidt commutator representations control existence or the size of a forward commutator [13, 21, 12, 3]; recent work on Aluthge transforms and finite-rank Toeplitz self-commutators is likewise forward or structural [23, 1]. None of these sources asks for the least factor rank inside the minimum-norm face of a prescribed finite Hermitian self-commutator. This comparison is based on a bounded search and is not an absolute priority claim.

The elementary inertia obstruction and norm optimality pull in different directions. If $P = CC^*$ and $Q = C^*C$, then P and Q are isospectral positive matrices, while $P - Q = 2F$. Inertia only sees the dimensions of their ranges; norm optimality sees the full Horn polytope for a difference of two isospectral positive matrices. Earlier exact calculations in dimensions three through five found compatibility between the two requirements [9, 7, 6], but finite examples cannot establish a dimension threshold.

This paper gives that threshold. Its first result is universal: every traceless Hermitian target through dimension seven has a norm-optimal factor at the inertia rank. Singular spectra are not inferred by padding a smaller problem. Extra ambient zero directions can in principle lower the norm, so every triple (n_+, n_-, n_0) is certified in its original dimension. Its second result is sharpness: the entire explicit two-ray dimension-eight cone has a positive rank penalty at every interior point. We include the entire displayed-cone theorem and a two-valued dimension-nine check here so that the threshold statement is standalone and does not split one result across companion papers. The third result is qualitative growth: a coarse-grid theorem for hives amplifies one exact order-27 face certificate into targets whose minimum optimal rank exceeds inertia by an arbitrarily large amount.

Development history and consolidation. Two public ARR records now provide the immediate norm-optimization background. *Sharp Rank-Adaptive Bounds for Inverse Self-Commutators* gives universal cost bounds and their equality spectra [10]; *One-Spike Inverse Self-Commutators and Exact Three-versus-Four-Kick Curvature Synthesis* gives the exact one-spike cone and optimizer singular rigidity [8]. Neither contains the rank-compatibility threshold or the unbounded-excess theorem proved here.

Four local, unpublished development packages contain the remaining chronology. *Minimum-Norm Self-Commutators Can Require Rank Above the Inertia Bound* supplies one interior dimension-eight witness, while *An Exact Rank–Norm Phase Diagram for Three-Level Self-Commutator Targets* replaces that point by the full displayed parametric family and adds the dimension-nine witness. The present manuscript absorbs and supersedes the latter two local drafts rather than treating them as separate publications, and adds the universal all-stratum theorem through dimension seven. That integrated threshold version was prepared locally as *The Sharp Dimension Threshold for Norm-Optimal Self-Commutator Rank*. The present manuscript absorbs and supersedes it by adding the coarse-grid amplification theorem and unbounded rank excess. Prepared local release or submission material is not public disclosure.

The universal part is finite but not hand-enumerative. We write the inverse problem as a rational polyhedral epigraph, project its inertia-rank face, and prove that every projected facet is already valid on the unrestricted epigraph. Exact GMP H/V conversion and exact rational Farkas identities are the proof objects. Floating linear programming is used only to propose sparse supports, which are discarded unless rational reconstruction verifies them coefficient by coefficient.

2 Horn reduction and the rank face

Let $F = F^* \in M_d(\mathbb{C})$ be traceless, with ordered eigenvalues $\lambda_1 \geq \cdots \geq \lambda_d$. Set

$$\kappa_d(F) = \frac{1}{2}\min\{\|C\|_{\mathrm{HS}}^2 : CC^* - C^*C = 2F\}, \qquad r_*(F) = \min\{\operatorname{rank} C : C \text{ attains } \kappa_d(F)\}. \tag{1}$$

Finite-dimensional self-commutator existence follows, for example, from a zero-sum ordering construction [11, 4].

Lemma 2.1 (inertia obstruction)**.** *Every feasible C satisfies*

$$\operatorname{rank} C \geq r_0(F) := \max\{n_+(F), n_-(F)\}.$$

Proof. The quadratic form of $CC^* - C^*C$ is nonpositive on $\ker(CC^*)$, hence its positive inertia is at most $\operatorname{rank} C$. Applying the same argument after interchanging C and C^* bounds the negative inertia. □

For $p_1 \geq \cdots \geq p_{d-1} \geq 0$, put

$$\alpha(p) = (p_1, \ldots, p_{d-1}, 0), \qquad \beta(p) = (0, -p_{d-1}, \ldots, -p_1), \qquad \gamma = 2\lambda.$$

Proposition 2.2 (Horn linear program)**.** *The inverse cost is*

$$\kappa_d(F) = \frac{1}{2}\min\left\{\sum_{j=1}^{d-1} p_j : (\alpha(p), \beta(p), \gamma) \text{ satisfies all Horn inequalities}\right\}. \tag{2}$$

The least cost under $\operatorname{rank} C \leq r$ *is obtained by adding* $p_{r+1} = \cdots = p_{d-1} = 0$*.*

Proof. For a feasible C, the matrices $P = CC^*$ and $Q = C^*C$ are positive, isospectral, and satisfy $P + (-Q) = 2F$. At a minimum their least common eigenvalue is zero: subtracting a positive scalar from both spectra preserves $P - Q$ and reduces the trace. Horn's theorem [17, 18, 19, 16] gives (2). Conversely, Horn feasibility supplies positive isospectral P, Q with $P - Q = 2F$. If $Q = U^*PU$, then $C = P^{1/2}U$ has the required self-commutator, norm, and rank. □

For $(I, J, K) \in T_s^d$ our convention is

$$\sum_{i \in I} \alpha_i + \sum_{j \in J} \beta_j \geq \sum_{k \in K} \gamma_k. \tag{3}$$

The recursive T-sets are those in Fulton's exposition [16].

3 The universal theorem through dimension seven

Theorem 3.1 (sharp low-dimensional compatibility)**.** *For every traceless Hermitian $F \in M_d(\mathbb{C})$ with $d \leq 7$,*

$$r_*(F) = \max\{n_+(F), n_-(F)\}. \tag{4}$$

This includes singular targets and the zero matrix.

Table 1: Exact certificate census for Proposition 3.2.

d	strata up to reflection	projected facets	face rays	maximum full rows
3	2	6	9	17
4	4	18	28	39
5	6	37	69	103
6	9	77	189	310
7	12	134	518	1063
total	33	272	813	–

We describe the finite exact certificate behind the theorem. Use adjacent spectral gaps $g_i = \lambda_i - \lambda_{i+1} \geq 0$. Trace zero gives

$$\lambda_k = \sum_{i=1}^{d-1} \left(\mathbf{1}_{i \geq k} - \frac{i}{d} \right) g_i. \tag{5}$$

Fix a sign–zero stratum (a, b, z) with $a, b \geq 1$, $z \geq 0$, and $a + b + z = d$. Here $a = n_+$, $b = n_-$, and $z = n_0$. In gap coordinates the prescribed zero eigenvalues are imposed as rational linear equalities, each written as two opposite inequalities. The two sign edges are imposed by $\lambda_a \geq 0$ and $-\lambda_{a+z+1} \geq 0$.

Let $r = \max\{a, b\}$. Define $\mathcal{E}_{d,a,b}$ to be the homogeneous Horn epigraph in variables (g, p, t): it contains the gap inequalities, every row (3), the order and positivity rows for p, the stratum rows, and

$$t - \frac{1}{2} \sum_{j=1}^{d-1} p_j \geq 0.$$

Let $\mathcal{E}^{(r)}_{d,a,b}$ be its coordinate face $p_{r+1} = \cdots = p_{d-1} = 0$. Write $\pi(g, p, t) = (g, t)$.

Proposition 3.2 (exact epigraph equality)**.** *For every $3 \leq d \leq 7$ and every sign–zero stratum,*

$$\pi(\mathcal{E}_{d,a,b}) = \pi(\mathcal{E}^{(r)}_{d,a,b}), \qquad r = \max\{a, b\}. \tag{6}$$

Finite exact proof. The inclusion from right to left is immediate because the rank face is a coordinate face of the unrestricted cone. Sign reflection $F \mapsto -F$, implemented by $C \mapsto C^*$, preserves cost and rank, so it suffices to enumerate $a \leq b$.

For every remaining stratum, exact rational double-description conversion produces the generator representation of $\mathcal{E}^{(r)}_{d,a,b}$ and, after deleting the p coordinates, the facet representation of its projection. For each projected facet $h(g, t) \geq 0$, the certificate contains nonnegative rational weights y such that

$$h(g, t) = y^T A_{d,a,b}(g, p, t), \tag{7}$$

where $A_{d,a,b}(g, p, t) \geq 0$ is the complete unrestricted Horn system. Thus every unrestricted point satisfies every exact facet of the projected rank face, proving the reverse inclusion. All H/V rows and all identities (7) are checked over rational arithmetic. Table 1 records the complete enumeration; the certificate file contains the facets, row labels, multipliers, and hashes. □

Proof of Theorem 3.1. The zero matrix is immediate, and $d \leq 2$ is elementary. Otherwise fix the actual stratum of F. Equality (6) says that the minimum epigraph height for its spectrum is attained on the rank-$r_0(F)$ face. Proposition 2.2 gives a self-commutator factor of rank at most $r_0(F)$. Lemma 2.1 gives the reverse inequality. □

Remark 3.3 (why singular strata are explicit)**.** *Solving the restriction of F to its support and padding the factor by zero only gives an upper bound in the original dimension. A larger ambient space may have a lower unrestricted Horn optimum. The $z > 0$ rows in the certificate therefore close a genuine logical gap; they are not bookkeeping duplicates. Indeed, the dimension-seven spectrum*

$$\tfrac{1}{7}(25, 18, 18, -10, -17, -17, -17)$$

has cost $74/7$, while padding by one zero and reordering lowers the dimension-eight unrestricted cost to $73/7$. The rank-four face of the padded problem remains at $74/7$, so ambient zero padding changes both the optimum and the optimal rank.

4 Sharpness on an explicit dimension-eight cone

For $c > 0$ and $0 \le b \le c$, let $F_{b,c}$ have spectrum

$$\lambda(b, c) = (4c - 3b, b, b, b, -c, -c, -c, -c). \tag{8}$$

This family is two-dimensional before positive scaling and one-dimensional projectively. No classification of all dimension-eight spectra is claimed.

Theorem 4.1 (exact rank–norm phase)**.** *For every $0 \le b \le c$,*

$$\kappa_8(F_{b,c}) = \begin{cases} 10c - 7b, & 0 \le b \le c/2, \\ 9c - 5b, & c/2 \le b \le c, \end{cases} \qquad \kappa_8^{(\le 4)}(F_{b,c}) = 10c - 6b. \tag{9}$$

Consequently the rank-four penalty is $\min\{b, c-b\}$. For $0 < b < c$, $r_*(F_{b,c}) = 5$ although $r_0(F_{b,c}) = 4$; at the two boundary rays a rank-four minimizer exists.

The three exact primal branches are

$$p^L = (8c - 6b, 6c - 6b, 4c - 4b, 2c, 2b, 0, 0), \qquad 0 \le b \le c/2, \tag{10}$$

$$p^H = (8c - 6b, 4c - 2b, 2c, 2c, 2c - 2b, 0, 0), \qquad c/2 \le b \le c, \tag{11}$$

$$p^4 = (8c - 6b, 6c - 4b, 4c - 2b, 2c, 0, 0, 0), \qquad 0 \le b \le c. \tag{12}$$

Their half-traces give the three values in (9). Exact endpoint evaluation of every affine Horn residual proves feasibility throughout each chamber. Sparse nonnegative Horn combinations produce respectively the objective rows $\frac{1}{2}(1, \ldots, 1)$ and $\frac{1}{2}(1, 1, 1, 1)$ with right sides $10c - 7b$, $9c - 5b$, and $10c - 6b$. The replay records every named triple and verifies the identities exactly. This proves optimality and the rank statements.

Corollary 4.2 (integer sharpness witness)**.** *If F has spectrum*

$$(5, 1, 1, 1, -2, -2, -2, -2),$$

then $r_0(F) = 4$, $\kappa_8(F) = 13$, and $\kappa_8^{(\le 4)}(F) = 14$. The minimum optimal rank is five.

Proof. Use $(b, c) = (1, 2)$ in Theorem 4.1. Exact feasible spectra are $(10, 6, 4, 4, 2, 0, 0)$ at rank five and $(10, 8, 6, 4, 0, 0, 0)$ at rank four. □

Corollary 4.3 (sharp dimension threshold)**.** *Dimension eight is the least dimension in which Hilbert–Schmidt norm-optimality can force the rank of a self-commutator factor strictly above the inertia lower bound.*

Proof. Theorem 3.1 excludes every target through dimension seven; Corollary 4.2 supplies a dimension-eight target. □

Within this explicit cone, the relative rank-four tax is sharp:

$$\frac{\kappa_8^{(\leq 4)}}{\kappa_8} \leq \frac{14}{13},$$

with equality precisely at $b = c/2$. Thus the threshold is not inferred from an isolated point: it persists along the entire displayed projective segment.

5 A two-valued dimension-nine check

Three spectral levels are not necessary for the separation. This proposition is auxiliary to the amplification argument; it is not used to prove the sharp dimension-eight threshold.

Proposition 5.1 (two-valued target)**.** *Let $G \in M_9(\mathbb{C})$ have spectrum $(5^4, -4^5)$. Then*

$$\kappa_9(G) = 29, \qquad \kappa_9^{(\leq 5)}(G) = 30, \qquad r_0(G) = 5, \qquad r_*(G) = 6.$$

Exact feasible squared-singular-value spectra are

$$(16, 12, 10, 10, 8, 2, 0, 0) \quad \text{and} \quad (16, 14, 12, 10, 8, 0, 0, 0),$$

with costs 29 and 30. Sparse duals supported on six and five named Horn or elementary rows give the matching lower bounds. The canonical replay checks all 39,716 Horn rows and the eight order/positivity rows.

6 Hive coarse-graining

We use triangular hive vertices $h(i, \ell)$, $0 \leq i \leq \ell \leq N$. For $1 \leq i \leq \ell < N$, write the three elementary rhombus slacks as

$$\begin{aligned}
R_1(i, \ell) &= h(i, \ell+1) + h(i-1, \ell) - h(i-1, \ell+1) - h(i, \ell), \\
R_2(i, \ell) &= h(i, \ell+1) + h(i, \ell) - h(i+1, \ell+1) - h(i-1, \ell), \\
R_3(i, \ell) &= h(i, \ell) + h(i-1, \ell) - h(i, \ell+1) - h(i-1, \ell-1).
\end{aligned}$$

A hive has all these quantities nonnegative [19, 20].

Lemma 6.1 (coarse-grid restriction)**.** *Let n, t be positive integers, put $N = nt$, and let h be an order-N hive. Then*

$$H(I, L) := h(tI, tL), \qquad 0 \leq I \leq L \leq n,$$

is an order-n hive. Its boundary increments are the sums of consecutive blocks of t boundary increments of h.

Proof. Every coarse rhombus is a telescoping nonnegative sum of fine rhombi. With $1 \le I \le L < n$, the identities are

$$R_1^H(I,L) = \sum_{a=t(I-1)+1}^{tI} \sum_{q=0}^{t-1} R_1^h(a, tL+q),$$
$$R_2^H(I,L) = \sum_{a=t(I-1)+1}^{tI} \sum_{q=0}^{t-1} R_2^h(a+q, tL+q),$$
$$R_3^H(I,L) = \sum_{a=1}^{t} \sum_{q=0}^{t-1} R_3^h(t(I-1)+a, t(L-1)+a+q).$$

They follow by telescoping directional edge differences across the relevant t-by-t parallelogram. All summands lie in the triangular index range, including when $I = L$. Sampling each boundary at multiples of t gives the asserted block sums. □

For a fine squared-singular-value spectrum $p_1 \ge \cdots \ge p_N = 0$, set

$$P_s = \sum_{q=(s-1)t+1}^{st} p_q, \qquad 1 \le s \le n. \tag{13}$$

Then $P_1 \ge \cdots \ge P_n \ge 0$, the two coarse isospectral boundaries are P and $-P$ in reverse order, and

$$\frac{1}{2}\sum_{s=1}^{n} P_s = \frac{1}{2}\sum_{i=1}^{N} p_i.$$

If $\operatorname{rank} p \le rt$, then $P_{r+1} = \cdots = P_n = 0$, so the coarse point lies on the literal rank-r face. Without a rank cap, P_n may be positive. Put $Q_s = P_s - P_n$ and replace the coarse hive by

$$H'(I,L) = H(I,L) + P_n(I-L).$$

The added function is affine, so every rhombus slack is unchanged; its two isospectral boundaries become Q and $-Q$ in reverse order, while the target boundary is unchanged. Now $Q_n = 0$, and the normalized coarse cost is lower than the fine cost by $nP_n/2$. This normalization is essential in the unrestricted lower bound.

Uniform puzzle inflation validates individually block-expanded Horn rows [20, Lemma 8], but does not by itself synthesize a global rank-face objective. Lemma 6.1 instead transfers an already certified whole face bound from a coarse problem to every refined problem.

7 Unbounded rank excess

For $k \ge 1$, let $F_k \in M_{9k}(\mathbb{C})$ have spectrum

$$(5^{4k}, -4^{5k}), \tag{14}$$

where exponents denote multiplicities.

Proposition 7.1 (exact endpoint costs)**.** *For every $k \ge 1$,*

$$\kappa_{9k}(F_k) = 29k, \qquad \kappa_{9k}^{(\le 5k)}(F_k) = 30k.$$

Proof. Direct sums of the two primals in Proposition 5.1 give the upper bounds. For the reverse bounds, coarse-grain any order-$9k$ feasible hive in blocks of length k. Lemma 6.1 gives an order-nine hive with target boundary $k(10^4, -8^5)$. On the rank-$5k$ face it lies literally on the rank-five face. Without the rank cap, put $Q_s = P_s - P_9$ as above; the normalized coarse cost is the fine cost minus $9P_9/2$. Homogeneity of the two exact dimension-nine lower certificates gives $30k$ and $29k$, respectively. □

The finite seed needed for rank amplification is stronger than the dimension-nine check.

Proposition 7.2 (exact order-27 seed)**.** *For $F_3 \in M_{27}(\mathbb{C})$,*

$$\kappa_{27}^{(\le 15)}(F_3) = 90, \qquad \kappa_{27}^{(\le 16)}(F_3) = 89, \qquad \kappa_{27}^{(\le 17)}(F_3) = 88, \qquad \kappa_{27}(F_3) = 87.$$

The upper values are sorted direct sums of the two dimension-nine primals. The lower rank-16 and rank-17 values are certified by frozen rational order-27 hive primals and duals. The rank-15 and unrestricted endpoints also follow from Proposition 7.1. A dependency-free verifier reconstructs every boundary, all three rhombus families, order rows, fixed rank coordinates, stationarity, and objective identity over `Fraction`.

Theorem 7.3 (unbounded additive excess)**.** *For every integer $t \ge 1$, let $G_t = F_{3t}$. Then*

$$\begin{aligned} \kappa_{27t}(G_t) &= 87t, & \kappa_{27t}^{(\le 15t)}(G_t) &= 90t, \\ \kappa_{27t}^{(\le 16t)}(G_t) &= 89t, & \kappa_{27t}^{(\le 17t)}(G_t) &= 88t. \end{aligned} \tag{15}$$

Consequently

$$17t + 1 \le r_*(G_t) \le 18t, \qquad r_0(G_t) = 15t, \qquad r_*(G_t) - r_0(G_t) \ge 2t + 1.$$

In particular, the excess of minimum optimal rank above inertia is unbounded.

Proof. Coarse-grain an order-$27t$ hive in blocks of length t. If its fine rank is at most $(15+s)t$, $s \in \{0,1,2\}$, then (13) lies on the order-27 rank-$(15+s)$ face, its target is t times that of F_3, and its half-trace cost is unchanged. Proposition 7.2 and homogeneity give the three restricted lower bounds in (15); the unrestricted lower bound is Proposition 7.1 with $k = 3t$. Direct sums with, respectively, $0, t, 2t, 3t$ rank-six branches among $3t$ dimension-nine blocks attain all four values. Thus every norm minimizer has rank at least $17t + 1$, while the all-full direct sum has rank $18t$. The inertia statement follows from the multiplicities $12t$ and $15t$. □

The theorem deliberately does not claim $r_*(G_t) = 18t$, nor the conjectural complete curve $\kappa_{9k}^{(\le 5k+j)}(F_k) = 30k - j$ for arbitrary k, j.

8 Reproducibility and proof boundary

The proof package separates discovery from certification.

Every route regenerates its Horn or hive rows from the stated convention, checks boundary and ordering equations, verifies primal feasibility, and checks the dual or Farkas identity coefficient by coefficient. The canonical projected certificate and the equality audit intentionally use different representations of linearity; the latter retains `cdd.gmp`'s linearity set instead of trusting the serialized orientation. The order-27 checker similarly validates the frozen seed, whereas the coarse-grid programs validate the parameter-uniform symbolic transfer. Thus the infinite theorem does not follow from agreement at a finite list of dimensions.

Table 2: Claim-to-certificate audit map. "Independent" means a separately implemented route or data model, not peer review.

Claim	Canonical proof object	Independent check
$d \leq 7$ compatibility	33 exact projected epigraph certificates; 272 H rows; rational Farkas multipliers	linearity-preserving audit of both orientations for 46 equalities and all 813 face generators
dimension-eight onset	complete affine Horn primals and sparse duals on both chambers	dependency-free 8,752-row integer-witness replay and separate direct-LR family replay
dimension-nine seed	exact rank-five/rank-six primals and matching Horn duals	independent recursive-Horn enumeration of all 39,716 Horn rows
unbounded excess	order-27 rank-face hive dual plus symbolic coarse-grid lemma	exact Fraction replay of all rhombi and a separate boundary/objective dictionary implementation

1. `verify_d_le_7_epigraph.py` is the canonical threshold verifier. It recursively regenerates Horn triples, enumerates all 33 sign–zero strata, uses exact `cdd.gmp` H/V conversion, and checks all discarded-row incidences and all reconstructed Farkas identities over rational arithmetic for the serialized orientation of each row. Its frozen JSON contains 272 projected H rows.

2. `audit_d_le_7_equalities.py` independently preserves `cdd.gmp`'s linearity set and verifies both exact Farkas orientations for all 46 projected equalities, alongside the 813 face generators.

3. `verify_low_dimension_minimality_exact.py` is an independent full-rank route with a different spectral parameterization and the standalone GMP `lcdd` executable. It is deliberately labelled nonsingular and is not used to infer the zero strata.

4. `verify_rank_gap_frontier.py` verifies the entire displayed dimension-eight phase and Proposition 5.1 using a dependency-free recursive Horn implementation and sparse rational duals.

5. `verify_d8_sharpness.py` is a smaller dependency-free replay of Corollary 4.2; it checks 8,752 Horn rows, seven order/positivity rows, both primals, and both dual coefficient identities.

6. `verify_zero_padding_fixture.py` verifies the three primal spectra and matching sparse rational duals behind the singular-padding warning above.

7. `verify_exact_hive_duals.py` replays the frozen order-18 rank-11 and order-27 rank-$16, 17$ hive certificates using only Python's exact `Fraction` arithmetic. It checks boundary data, all rhombus families, rank-face equalities, primal and dual feasibility, stationarity, and exact objective identities.

8. `verify_hive_coarse_graining.py` exhaustively expands the three symbolic telescoping identities of Lemma 6.1 for a broad finite index range, while `verify_block_inflation_endpoints.py` checks the endpoint scaling, block ranks, and direct-sum costs in Propositions 7.1 and 7.2.

9. `verify_coarse_graining_theorem.py` is an independent dependency-free dictionary replay of the boundary normalization, order rows, objective comparison, rank-$17t$ corollary, and the finite inputs used in Theorem 7.3.

The exact threshold certificate has SHA-256

`0103a643644977400052a68738102a7633d6371db8717383ddefa687da8a18a0`.

The canonical script has SHA-256

`e5009d579933af66283de296ac8aa46b8f0b35bae1e43dd300768ef625129bd6`.

The equality-audit script has SHA-256

`0362a3dae16a027b87b3e39c360dfa657f42fd686401c2340d46013bd7bfb95e`.

The full phase, integer-witness, and zero-padding replay scripts have, respectively, SHA-256 hashes

```
00bdbbfd8adbf3aff8f847ae3574241fc7bed887573c8fa6142bb242cf874be7
979605969b0c6df628d619b8255ba9f8ab2bef5f4f63d790d2ca4b48e1a2fa9c
6c8a804a3b7b8d77acac9d7e000dd48857858b7c616c11565b06efc38075fa2d
```

The frozen hive-dual JSON, its exact verifier, and the symbolic coarse-grid verifier have, respectively, SHA-256 hashes

```
c6b3588a2415db067f7ff34f7e23592ac9d85f3e10399dd0f8838fc244352b69
e353f0f65d7edc2b3274ba2842263747fe5a5728f65c0e1ff4cf05407fde09e2
e3bfc7df692c021b600bd85fa1b8604b5f12b7bc4995e1dee4e0b1ea8c6fef89
```

The all-k endpoint checker and the independent standalone theorem replay have SHA-256 hashes

```
f7c66a937e100e293ee68794606036afc09e6eb4e603e1a275bfb9f8d3978d80
b1ea0d5dad40b56d217cbbe32054110a63267ec7be0e862af964b17cf0cda91f
```

The GMP double-description method is implemented by `cddlib` [15, 14]. Numerical redundancy detection and HiGHS supports are hints only: a proposed reduction is retained only after every original row is nonnegative on every exact generator, and a proposed Farkas support is retained only after exact nonnegative reconstruction.

The accompanying release archive contains the manuscript source, bibliography, all canonical and independent verifiers, frozen exact JSON objects, environment pins, and a sorted SHA-256 manifest. A single command, `python run_scientific_replay.py`, executes the complete default audit. The permanent record is https://arr-research.github.io/papers/ARR-2026-5QQF95VHTC9GABH8/; its immutable release is the canonical artifact locator. No private development tree is required.

The low-dimensional threshold and the order-27 seed are computer-assisted finite proofs; the amplification step is symbolic. The result does not classify all optimizer matrices, determine the exact value of $r_*(G_t)$, give the complete rank–cost curve for arbitrary F_k, or claim that the displayed families are extremal among all higher-dimensional targets. Independent peer review has not been performed. A refreshed bounded search, including the adjacent 1986–2026 literature cited above, found no exact antecedent for the universal $d \leq 7$ statement, the sharp onset, or the unbounded optimal-rank excess. Absence from that search is not proof of novelty.

9 Conclusion

The inertia lower bound is exactly compatible with Hilbert–Schmidt optimality through dimension seven and fails in dimension eight. The result requires both halves: an exhaustive exact treatment of all low-dimensional sign–zero strata and a strict dimension-eight phase. The Horn epigraph turns the universal side into finitely checkable projection equalities, while the explicit dimension-eight primals and sparse duals show that the first failure persists along the entire displayed projective segment. The exact order-27 seed and coarse-grid theorem then show that the additive excess $r_*(F) - r_0(F)$ is unbounded. Determining the exact optimal rank along this amplified family, and the complete intermediate rank–cost curve, remains open.

References

[1] Mankunikuzhiyil Abhinand, Raul E. Curto, and Thankarajan Prasad. Hyponormal block Toeplitz operators with finite rank self-commutators. *Constructive Mathematical Analysis*, 9(1):19–30, 2026.

[2] A. Adrian Albert and Benjamin Muckenhoupt. On matrices of trace zero. *Michigan Mathematical Journal*, 4(1):1–3, 1957.

[3] Omer Angel and Gideon Schechtman. The Hilbert–Schmidt version of the commutator theorem for zero trace matrices. *Bulletin of the London Mathematical Society*, 47(4):715–719, 2015.

[4] Daniel Beltiţă, Sasmita Patnaik, and Gary Weiss. $B(H)$-commutators: A historical survey II and recent advances on commutators of compact operators. In *The Varied Landscape of Operator Theory*, volume 17 of *Theta Series in Advanced Mathematics*, pages 57–75. Theta, Bucharest, 2014.

[5] Albrecht Böttcher and David Wenzel. How big can the commutator of two matrices be and how big is it typically? *Linear Algebra and its Applications*, 403:216–228, 2005.

[6] Lluis Eriksson. The exact five-level inverse commutator cost: Twelve Horn chambers, optimal rank, and the sharp 5/2 resource tax. ai.viXra 2608.0049, 2026.

[7] Lluis Eriksson. The exact four-level inverse commutator cost: Horn–Littlewood–Richardson facets, rank transitions, and sharp loop synthesis. ai.viXra 2608.0031, 2026.

[8] Lluis Eriksson. One-spike inverse self-commutators and exact three-versus-four-kick curvature synthesis. ARR-2026-7NPRNBW4488HG90K, 2026.

[9] Lluis Eriksson. Sharp costs and exact semigroups from forgotten quantum order. ai.viXra 2608.0034, 2026.

[10] Lluis Eriksson. Sharp rank-adaptive bounds for inverse self-commutators. ARR-2026-1D2QV1RP1292JREW, 2026.

[11] Peng Fan and Che-Kao Fong. Which operators are the self-commutators of compact operators? *Proceedings of the American Mathematical Society*, 80:58–60, 1980.

[12] N. Filonov and Y. Safarov. On the relation between an operator and its self-commutator. *Journal of Functional Analysis*, 260(10):2902–2932, 2011.

[13] Che-Kao Fong. Norm estimates related to self-commutators. *Linear Algebra and its Applications*, 74:151–156, 1986.

[14] Komei Fukuda. cddlib: Polyhedral representation conversion and linear programming. Software and reference manual. Accessed 24 August 2026.

[15] Komei Fukuda and Alain Prodon. Double description method revisited. In *Combinatorics and Computer Science*, volume 1120 of *Lecture Notes in Computer Science*, pages 91–111. Springer, 1996.

[16] William Fulton. Eigenvalues, invariant factors, highest weights, and Schubert calculus. *Bulletin of the American Mathematical Society*, 37(3):209–249, 2000.

[17] Alfred Horn. Eigenvalues of sums of hermitian matrices. *Pacific Journal of Mathematics*, 12(1):225–241, 1962.

[18] Alexander A. Klyachko. Stable bundles, representation theory and hermitian operators. *Selecta Mathematica*, 4:419–445, 1998.

[19] Allen Knutson and Terence Tao. The honeycomb model of $GL_n(\mathbb{C})$ tensor products I: Proof of the saturation conjecture. *Journal of the American Mathematical Society*, 12(4):1055–1090, 1999.

[20] Allen Knutson, Terence Tao, and Christopher Woodward. The honeycomb model of $GL_n(\mathbb{C})$ tensor products II: Puzzles determine facets of the littlewood–richardson cone. *Journal of the American Mathematical Society*, 17(1):19–48, 2004.

[21] P. Maher. Self-commutator approximants. *Proceedings of the American Mathematical Society*, 134(1):157–165, 2006.

[22] Gary Weiss. Commutators of hilbert–schmidt operators. II. *Integral Equations and Operator Theory*, 3:574–600, 1980.

[23] Teng Zhang. On a conjecture of lambda-aluthge transforms and Hilbert–Schmidt self-commutators. arXiv:2603.04655, 2026.