The Exact Five-Level Inverse Commutator Cost

Twelve Horn chambers, optimal rank, and the sharp 5/2 resource tax

Lluis Eriksson

11 August 2026

Abstract

For a traceless Hermitian five-level target F, consider

$$\kappa_5(F) = \inf\{\|A\|_{\mathrm{HS}} \|B\|_{\mathrm{HS}} : A = A^*,\ B = B^*,\ -i[A, B] = F\}.$$

If $g_i = \lambda_i - \lambda_{i+1}$ for the ordered spectrum $\lambda_1 \geq \cdots \geq \lambda_5$, we prove the closed formula

$$\boxed{\kappa_5(F) = \max_{q \in \mathcal{Q}} q \cdot g,}$$

where $\mathcal{Q}$ consists of six explicit rational rows and their reversals. Thus the answer has twelve exposed linear facets, rather than the four facets present at dimension four. The proof reduces the problem to the Horn–Littlewood–Richardson cone for the spectra of CC^* and $-C^*C$. Exact enumeration gives $15 + 56 + 56 + 15 = 142$ Horn triples. Twelve sparse dual identities give the lower bounds, while twelve rational maps $p = P_j g$ attain them. Completeness is certified by exact equality between every facet-dominance cone and its full Horn-feasibility cone; no sampling is used in the proof.

The dimensional transition is structural: the qutrit problem has two adjacent-gap facets and optimal rank at most two; dimension four has four facets and first requires rank three; dimension five has twelve facets and two open chambers that force rank four. The corresponding sharp trace-norm tax rises from 3/2 through 2 to 5/2.

For every nonzero target, the least rank of an optimal factor is

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

The exact trace-norm tax is

$$1 \leq \frac{\kappa_5(F)}{\|F\|_1 / 2} \leq \frac{5}{2}.$$

The lower equality holds exactly for spectra $(x, y, 0, -y, -x)$, while the upper equality occurs only on the two rays generated by $(4, -1, -1, -1, -1)$ and $(1, 1, 1, 1, -4)$. As a control corollary, the least squared Hilbert–Schmidt perimeter of a balanced triangular loop with geometric Hamiltonian F is $12\sqrt{3}\,\kappa_5(F)$. Two exact rational certificates and a deterministic 4096-spectrum LP campaign accompany the paper.

1 The five-level inverse problem

Every complex trace-zero matrix is a commutator, and finite-dimensional trace-zero Hermitian matrices are self-commutators [1, 2]. These existence theorems do not determine the smallest factors for a prescribed target. We study

$$\kappa_d(F) := \inf\{\|A\|_{\mathrm{HS}} \|B\|_{\mathrm{HS}} : A = A^*,\ B = B^*,\ -i[A, B] = F\}, \tag{1}$$

where $F = F^* \in M_d(\mathbb{C})$, $\operatorname{Tr} F = 0$, and $\|X\|_{\mathrm{HS}}^2 = \operatorname{Tr}(X^*X)$ is unnormalized.

The nearby quantitative literature asks different questions. Böttcher–Wenzel inequalities bound a commutator from above in terms of its factors [9, 11, 12]. Dimension-dependent commutator theorems control an operator–Hilbert–Schmidt mixed cost for arbitrary trace-zero matrices [13, 14]. Work on self-commutators establishes existence or forward norm bounds [4, 10, 15]. Here the complete Hermitian target is prescribed, both factors are Hermitian, and their HS–HS product is minimized. To the best of our knowledge, no previous closed spectral formula is available for this five-dimensional inverse optimization problem.

The engine is classical: Horn's theorem and its representation-theoretic completion characterize spectra of sums of Hermitian matrices [3, 5, 6, 7, 8]. We do not claim that cone or the Littlewood–Richardson rule. The contribution is the exact elimination of the specialized five-level program: twelve value facets, twelve attaining singular spectra, their complete chamber geometry, optimal ranks, and sharp resource ratios.

Relation to the eleven source manuscripts

This is a new twelfth manuscript, not a replacement for the eleven submitted papers. The immediate predecessor [16] solved the four-level problem and stated the general Horn program used below. The present result is not obtained by substituting $d = 5$ in its four-level formula: the facet count rises from four to twelve, the sharp trace tax rises from 2 to 5/2, two rank-four optimal chambers appear, and the full proof requires exact elimination of 142 Horn–LR inequalities. We repeat the reduction so that the theorem is logically self-contained apart from the classical Horn existence theorem.

The structural transition from three to five levels

Table 1 isolates what changes in the exact low-dimensional solutions [17, 16]. The general Horn program predicts a polyhedral optimization problem, but it does not determine the exposed facets, their chamber adjacency, the ranks of attaining factors, or the sharp tax constant. Those data change at each dimension.

Table 1: Exact structural ladder for the inverse Hermitian commutator cost.

d	exposed facets	spectral structure	largest forced rank	sharp tax
3	2	adjacent gaps	2	3/2
4	4	two gap and two mixed Horn forms	3	2
5	12	twelve mixed Horn forms in six reflection pairs	4	5/2

In particular, passing from four to five levels creates eight new exposed facets and the first open chambers of inertia $(1, 4)$ and $(4, 1)$, where rank four is unavoidable. The exact elimination below is therefore not a numerical rerun of the predecessor: it determines new chamber geometry and new extremal resource laws.

2 Reduction to one Horn linear program

Lemma 2.1 (Balanced self-commutator reduction)**.** *For every traceless Hermitian $F \in M_d(\mathbb{C})$,*

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

Proof. Put $C = A + iB$. Then

$$CC^* - C^*C = -2i[A, B], \qquad \|C\|_{\mathrm{HS}}^2 = \|A\|_{\mathrm{HS}}^2 + \|B\|_{\mathrm{HS}}^2 .$$

The rescaling $(A, B) \mapsto (sA, s^{-1}B)$ preserves the commutator and minimizes the sum of squares at $2\,\|A\|_{\mathrm{HS}}\,\|B\|_{\mathrm{HS}}$. Conversely, the Hermitian real and imaginary parts of every feasible C give feasible factors. ☐

For decreasing real vectors $\alpha, \beta, \gamma \in \mathbb{R}^d$ of matching trace, Horn's theorem says that Hermitian matrices with spectra α, β can have sum spectrum γ precisely when

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

for all triples $(I, J, K) \in T_r^d$, $1 \le r < d$, with positive Littlewood–Richardson coefficient.

Theorem 2.2 (Inverse-commutator Horn program)**.** *Let $\lambda_1 \ge \cdots \ge \lambda_d$ be the eigenvalues of a traceless Hermitian F. Put*

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

Then

$$\boxed{\kappa_d(F) = \frac{1}{2} \min_{p_1 \ge \cdots \ge p_{d-1} \ge 0} \left\{ \sum_{j=1}^{d-1} p_j : (\alpha(p), \beta(p), \gamma) \text{ obey } (3) \right\}.} \tag{4}$$

Proof. For feasible C, the positive matrices CC^* and C^*C are isospectral and have difference $2F$. If their least common eigenvalue p_d is positive, subtracting $p_d\mathbf{1}$ from both preserves the difference and lowers the trace; hence $p_d = 0$ at a minimum. Applying Horn's theorem to $CC^* + (-C^*C) = 2F$ gives (4).

Conversely, Horn feasibility supplies positive isospectral matrices $P = U \operatorname{diag}(p) U^*$ and $Q = V \operatorname{diag}(p) V^*$ with $P - Q = 2F$. Then $C = U \operatorname{diag}(\sqrt{p}) V^*$ has $CC^* = P$, $C^*C = Q$, and $\|C\|_{\mathrm{HS}}^2 = \sum p_j$. Lemma 2.1 finishes the proof. $\square$

3 The twelve-facet formula

Let $F \in M_5(\mathbb{C})$ be traceless Hermitian. Write

$$g_i = \lambda_i - \lambda_{i+1} \ge 0, \qquad g = (g_1, g_2, g_3, g_4)^T, \tag{5}$$

and let $\rho(q_1, q_2, q_3, q_4) = (q_4, q_3, q_2, q_1)$. Define the six rows

$$\begin{aligned} q_1 &= (2, 0, -1, -1), & q_2 &= \tfrac{1}{5}(6, 7, 3, -1), & q_3 &= \tfrac{1}{5}(2, 9, 1, -2), \\ q_4 &= \tfrac{1}{5}(8, 6, -1, -3), & q_5 &= \tfrac{1}{5}(4, 8, 2, 1), & q_6 &= \tfrac{1}{5}(4, 3, 7, 1), \end{aligned} \tag{6}$$

and $\mathcal{Q} = \{q_1, \ldots, q_6, \rho q_1, \ldots, \rho q_6\}$.

Theorem 3.1 (Exact five-level inverse cost)**.** *For every traceless Hermitian $F \in M_5(\mathbb{C})$,*

$$\boxed{\kappa_5(F) = \max_{q \in \mathcal{Q}} q \cdot g.} \tag{7}$$

All twelve forms are exposed: each is uniquely maximal on a nonempty open cone. The minimum is attained on all chambers, walls, and degeneracies.

Table 2 translates the six representative rows into trace-zero spectral forms. The reflected forms are evaluated on $F^\sharp = -JFJ$, where J is the antidiagonal permutation matrix; its ordered gap vector is ρg.

Table 2: Six representative facets; the other six are their reflections.

r	$q_r \cdot g$	an interior gap witness
1	$2\lambda_1 - 2\lambda_2 - \lambda_3 + \lambda_5$	$(7, 1, 1, 1)$
2	$\lambda_1 - \lambda_3 - \lambda_4$	$(2, 1, 1, 1)$
3	$\lambda_2 - 2\lambda_3 - \lambda_4$	$(1, 7, 1, 1)$
4	$2\lambda_1 - \lambda_3 + \lambda_5$	$(4, 1, 1, 1)$
5	$\lambda_1 + \lambda_2 - \lambda_3$	$(1, 2, 1, 1)$
6	$\lambda_1 + \lambda_3 - \lambda_4$	$(2, 1, 2, 1)$

For clarity, the complete twelve-term spectral version of (7) is

$$\kappa_5(F) = \max\{2\lambda_1 - 2\lambda_2 - \lambda_3 + \lambda_5, -\lambda_1 + \lambda_3 + 2\lambda_4 - 2\lambda_5,$$

$$\begin{aligned}&\lambda_1-\lambda_3-\lambda_4, \lambda_2+\lambda_3-\lambda_5,\\&\lambda_2-2\lambda_3-\lambda_4, \lambda_2+2\lambda_3-\lambda_4,\\&2\lambda_1-\lambda_3+\lambda_5, -\lambda_1+\lambda_3-2\lambda_5,\\&\lambda_1+\lambda_2-\lambda_3, \lambda_3-\lambda_4-\lambda_5,\\&\lambda_1+\lambda_3-\lambda_4, \lambda_2-\lambda_3-\lambda_5\}.\end{aligned} \tag{8}$$

4 Exact lower and upper certificates

Put $\mu = 2\lambda$, $p = (p_1, p_2, p_3, p_4, 0)$, $\alpha = p$, $\beta = (0, -p_4, -p_3, -p_2, -p_1)$, and $s = p_1 + p_2 + p_3 + p_4$. For a Horn triple define its nonnegative residual

$$R_{IJK} = \sum_{i\in I} \alpha_i + \sum_{j\in J} \beta_j - \sum_{k\in K} \mu_k. \tag{9}$$

4.1 Six dual identities and reflection

Table 3 gives sparse exact identities. We abbreviate subsets by digit strings; for example,

$$123; 145; 145 \quad \text{means} \quad (I, J, K) = (\{1, 2, 3\}, \{1, 4, 5\}, \{1, 4, 5\}).$$

In every row the sum of the displayed residuals and the displayed multiple of p_4 equals

$$s - 2q_r \cdot g. \tag{10}$$

It is therefore nonnegative. Reflection gives the other six lower bounds.

Table 3: Sparse Horn–LR dual certificates for the six representative facets.

r	Horn residuals in (10)	extra term
1	$R_{1;1;1} + R_{12;15;15} + R_{123;145;145} + R_{1234;1345;1345}$	0
2	$R_{1;1;1} + R_{123;125;125}$	$2p_4$
3	$R_{2;1;2} + R_{124;124;125} + R_{1234;1245;1245}$	p_4
4	$R_{1;1;1} + R_{12;15;15} + R_{1234;1245;1245}$	p_4
5	$R_{12;12;12} + R_{1234;1245;1245}$	$2p_4$
6	$R_{13;12;13} + R_{1234;1235;1235}$	$2p_4$

Every displayed triple belongs to the exact LR list. Thus Theorem 2.2 gives $\kappa_5 \geq \max_{q\in\mathcal{Q}} q \cdot g$ without numerical optimization.

4.2 Twelve attaining spectra

For the six representative facets define $P_r = M_r/5$, where

$$M_1 = \begin{pmatrix} 8 & 6 & 4 & 2\\ 6 & 2 & -2 & -6\\ 4 & -2 & -8 & -4\\ 2 & -6 & -4 & -2 \end{pmatrix}, \qquad M_2 = \begin{pmatrix} 8 & 6 & 4 & 2\\ 2 & 4 & 6 & -2\\ 2 & 4 & -4 & -2\\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$M_3 = \begin{pmatrix} 4 & 8 & 2 & -4\\ -2 & 6 & 4 & 2\\ 2 & 4 & -4 & -2\\ 0 & 0 & 0 & 0 \end{pmatrix}, \qquad M_4 = \begin{pmatrix} 8 & 6 & 4 & 2\\ 6 & 2 & -2 & -6\\ 2 & 4 & -4 & -2\\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$M_5 = \begin{pmatrix} 8 & 6 & 4 & 2\\ -2 & 6 & 4 & 2\\ 2 & 4 & -4 & -2\\ 0 & 0 & 0 & 0 \end{pmatrix}, \qquad M_6 = \begin{pmatrix} 8 & 6 & 4 & 2\\ 2 & 4 & 6 & -2\\ -2 & -4 & 4 & 2\\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{11}$$

For a reflected facet use $P_r J_4$, where $J_4 g = \rho g$. Direct summation gives

$$\mathbf{1}^T P_r g = 2q_r \cdot g. \tag{12}$$

Let

$$\mathcal{C}_r = \{g \in \mathbb{R}^4_+ : q_r \cdot g \geq q \cdot g \text{ for every } q \in \mathcal{Q}\}. \tag{13}$$

The key finite statement is

$$g \in \mathcal{C}_r \quad \Longleftrightarrow \quad p = P_r g \text{ obeys } p_1 \geq p_2 \geq p_3 \geq p_4 \geq 0 \text{ and all 142 Horn inequalities.} \tag{14}$$

This is an exact polyhedral identity. Each side reduces to the same four-row simplicial cone $H_r g \geq 0$. The six representative matrices H_r are

$$H_1 = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & -3 & -2 & -1 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \qquad H_2 = \begin{pmatrix} 2 & -1 & 1 & -2 \\ 1 & 2 & -2 & -1 \\ -2 & 1 & 4 & 2 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$
$$H_3 = \begin{pmatrix} -2 & 1 & -1 & -3 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \qquad H_4 = \begin{pmatrix} 2 & -1 & -4 & -2 \\ -1 & 3 & 2 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$
$$H_5 = \begin{pmatrix} 3 & 1 & -1 & -3 \\ 2 & -1 & 1 & 3 \\ -2 & 1 & -1 & 2 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \qquad H_6 = \begin{pmatrix} 3 & 1 & -1 & -3 \\ -1 & -2 & 2 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \tag{15}$$

The reflected cones use reversed columns.

For transparency about the computer-assisted step, the ancillary verifier performs three exact checks. First it recursively enumerates all 142 LR triples. Second, it reduces every one of the 142 Horn residuals after $p = P_r g$ to a nonnegative rational combination of the four rows of H_r. Third, it reduces every dominance inequality in (13) to the same cone and proves the reverse containments. All identities are checked over $\mathbb{Q}$; the canonical Horn tuple list has SHA-256

`75e617faca4d8007c5c96a82e2efaac9dedcc2d8c87d3212dd5085e88d791a34.`

An independent expanded verifier checks 5048 nonzero rational Farkas multipliers directly.

Every $g \geq 0$ has a maximizing row in $\mathcal{Q}$. Equation (14) then makes $p = P_r g$ Horn feasible, and (12) makes its cost equal to the dual bound. Horn existence reconstructs C exactly. This proves Theorem 3.1.

Remark 4.1 (Attainment versus an entrywise formula). The theorem gives an exact singular spectrum and an exact reconstruction procedure through the Horn spectral-sum problem. It does not claim one universal entry-by-entry radical expression for the optimal 5×5 matrix C. Attainment uses the classical Horn existence theorem, not a numerical eigensolver.

5 Optimal rank on every spectral stratum

Proposition 5.1 (Exact optimal-rank law)**.** *For every nonzero traceless Hermitian* $F \in M_5(\mathbb{C})$*,*

$$\boxed{\min\{\operatorname{rank} C : CC^* - C^*C = 2F,\ \|C\|^2_{\mathrm{HS}} = 2\kappa_5(F)\} = \max\{n_+(F), n_-(F)\}.} \tag{16}$$

Proof. If $\operatorname{rank} C = r$, both CC^* and C^*C have rank r. Their difference has at most r strictly positive and at most r strictly negative eigenvalues. Thus every feasible factor obeys $r \geq \max\{n_+(F), n_-(F)\}$.

The attaining spectra in (11) reach this bound. The two reflected outer chambers associated with q_1 have

$$p = (2\lambda_1, -2(\lambda_2 + \lambda_3 + \lambda_4), -2(\lambda_2 + \lambda_3), -2\lambda_2, 0) \tag{17}$$

or its reflection. If $\lambda_2 < 0$, all four entries are positive and the inertia is $(1,4)$. If $\lambda_2 = 0$ but $\lambda_3 < 0$, exactly three remain and $n_- = 3$; if also $\lambda_3 = 0$ but $\lambda_4 < 0$, exactly two remain and $n_- = 2$. Finally $\lambda_2 = \lambda_3 = \lambda_4 = 0$ gives the rank-one ray $(x,0,0,0,-x)$. Reflection covers the opposite signs.

On the ten inner chambers, $p_4 = 0$ and $p_3 = 2|\lambda_3|$. If $\lambda_3 \neq 0$, the rank is three; feasibility and ordering then force inertia $(2,3)$ or $(3,2)$. If $\lambda_3 = 0$, the rank is at most two. It equals two whenever either middle side of the spectrum is nonzero, by the inertia lower bound already proved; if both sides vanish, the target is the rank-one ray just treated. Hence the displayed attaining spectrum has rank $\max(n_+, n_-)$ on every stratum. Horn reconstruction preserves this rank. □

Thus five levels exhibit a new optimal resource: rank four is forced on the two outer interiors. It is not an artifact of a particular constructor; inertia already gives the matching lower bound.

6 The sharp trace-norm tax

For nonzero traceless Hermitian F, let

$$P(F) = \frac{\|F\|_1}{2} = \sum_{\lambda_j > 0} \lambda_j. \tag{18}$$

Theorem 6.1 (Tax interval and equality loci)**.** *For every nonzero traceless Hermitian $F \in M_5(\mathbb{C})$,*

$$\boxed{1 \le \frac{\kappa_5(F)}{P(F)} \le \frac{5}{2}.} \tag{19}$$

The lower equality holds exactly when

$$\operatorname{spec} F = (x, y, 0, -y, -x), \qquad x \ge y \ge 0. \tag{20}$$

The upper equality holds exactly on the rays generated by

$$(4,-1,-1,-1,-1), \qquad (1,1,1,1,-4). \tag{21}$$

Proof. Let E be the positive spectral projection of F and let $CC^* - C^*C = 2F$. Then

$$2P(F) = \operatorname{Tr}\big(E(CC^* - C^*C)\big) \le \operatorname{Tr}(C^*EC) \le \|C\|_{\mathrm{HS}}^2.$$

Minimizing and using Lemma 2.1 gives $P(F) \le \kappa_5(F)$.

For the upper bound, split the ordered chamber into the four possible inertia cones

$$g_i \ge 0, \qquad \lambda_r \ge 0 \ge \lambda_{r+1}, \qquad 1 \le r \le 4.$$

On each cone $P = \lambda_1 + \cdots + \lambda_r$. For every one of the twelve rows q, exact elimination gives

$$\frac{5}{2}P - q \cdot g = \sum_{i=1}^{4} a_i g_i + b\lambda_r - c\lambda_{r+1}, \qquad a_i, b, c \ge 0,$$

with rational coefficients. The 48 identities are stored and rechecked by the certificate. Theorem 3.1 proves the upper bound.

Intersecting the exact equality cones for the lower bound leaves precisely the cone generated by gaps $(1,0,0,1)$ and $(0,1,1,0)$, equivalently $g_1 = g_4$ and $g_2 = g_3$; trace zero gives (20). For the upper bound, the only surviving rays are $g = (1,0,0,0)$ and $g = (0,0,0,1)$, which give (21). These ray computations use exact rational row reduction, not tolerance tests. □

The spike $F = \operatorname{diag}(4,-1,-1,-1,-1)$ has $P(F) = 4$ but $\kappa_5(F) = 10$. It therefore disproves both the qutrit adjacent-gap rule and a naive continuation of the four-level tax constant.

7 Balanced triangular synthesis

The matrix optimum has a direct finite-control interpretation. Let h_1, h_2, h_3 be Hermitian five-level kicks balanced modulo the center. Removing scalar parts gives traceless sides $x_1 + x_2 + x_3 = 0$. For one orientation, define

$$H_{\text{geo}} = \frac{i}{2}[x_1, x_2], \qquad S_2 = \sum_{j=1}^{3} \text{dist}_{\text{HS}}(h_j, \mathbb{R}\mathbf{1}) = \sum_{j=1}^{3} \|x_j\|_{\text{HS}} . \tag{22}$$

Corollary 7.1 (Exact five-level triangular perimeter)**.** *Among balanced three-kick loops with $H_{\text{geo}} = F$,*

$$\boxed{\inf S_2^2 = 12\sqrt{3}\,\kappa_5(F).} \tag{23}$$

Proof. The three sides form a triangle in the real Hilbert space of traceless Hermitian matrices. If Δ is its Euclidean area, the commutator constraint implies $\kappa_5(F) \le \Delta$. The sharp triangular isoperimetric inequality gives $S_2^2 \ge 12\sqrt{3}\,\Delta$.

Conversely, for any $-i[A, B] = F$, remove the component of B parallel to A and rescale so that A, B are orthogonal and have equal norm. With $\alpha^2 = 4/\sqrt{3}$, set

$$x_1 = \alpha A, \quad x_2 = \alpha\left(-\frac{A}{2} - \frac{\sqrt{3}B}{2}\right), \quad x_3 = \alpha\left(-\frac{A}{2} + \frac{\sqrt{3}B}{2}\right).$$

This is an equilateral balanced triangle, has $(i/2)[x_1, x_2] = F$, and satisfies $S_2^2 = 12\sqrt{3}\,\|A\|_{\text{HS}}\,\|B\|_{\text{HS}}$. Taking the infimum proves (23). ☐

The constant depends on conventions: S_2 is the perimeter $\sum \|x_j\|$, not the quadratic action $\sum \|x_j\|^2$, and F is H_{geo} rather than twice that coefficient.

8 Reproducibility and falsifiability

Two exact scripts and one independent numerical campaign accompany the manuscript.

1. The compact verifier recursively constructs the 142 Horn triples, checks twelve sparse dual identities, verifies the twelve P_j maps, proves equality of dominance and feasibility cones, and derives the tax equality rays. Its JSON digest is

   `87532b63c195a5b87c73eaab21e79bee100f4f4969ef755aa10ab98e2b0fe6b1.`

2. The expanded verifier independently expresses every specialized Horn and ordering residual as a nonnegative rational combination of chamber generators. It checks 5048 nonzero Farkas multipliers. Its JSON digest is

   `82263d74b96694d59d810b091936d2b4217cad828c3681efea98f22b775a5a7c.`

3. A deterministic campaign samples 4096 ordered spectra and solves the full Horn LP independently. All twelve facets appear. The maximum formula–LP and primal-feasibility errors are both 4.44×10^{-16}; these floating checks are validation, not proof.

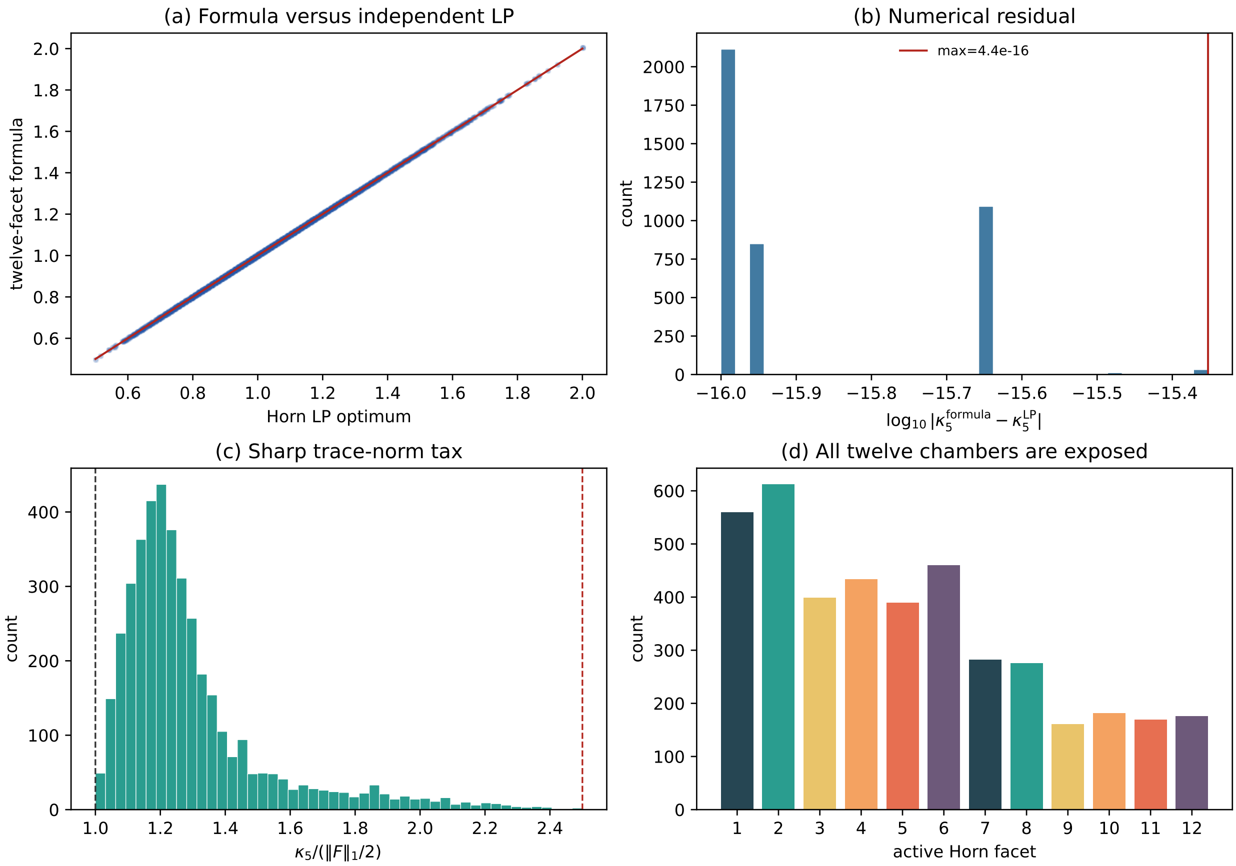

Figure 1: Independent numerical audit. (a) The twelve-facet formula agrees with the full Horn LP. (b) Residuals remain at floating roundoff. (c) Random spectra lie inside the proved tax interval; the two dashed values are attained by exact targeted examples. (d) Every one of the twelve Horn facets is exposed.

The result is readily falsifiable. A proposed counterexample must either violate one of the twelve sparse lower identities or provide a negative residual for its maximizing P_j map. Both checks use rational arithmetic. The only imported non-elementary theorem is the classical equivalence between the Horn–LR inequalities and Hermitian spectral-sum existence. No claim about infinite factors, unrestricted non-Hermitian targets, or an entrywise radical constructor is made.

9 Conclusion

The five-level inverse commutator problem has an exact polyhedral answer. Its 142 Horn–Littlewood–Richardson inequalities collapse, after optimization, to twelve exposed spectral forms. Exact dual identities prove every lower facet; rational singular-spectrum maps and cone identities prove global attainment. The resulting geometry fixes the minimum optimal rank on every stratum and enlarges the sharp trace-norm resource tax from the four-level value 2 to 5/2.

The calculation also shows what genuinely changes with dimension. New Horn combinations are not reducible to adjacent gaps or spectral diameters, and the outer five-level targets force rank-four factors. At the same time, the formula remains auditable: every chamber, multiplier, equality ray, and numerical comparison is supplied in machine-readable form.

AI-assistance disclosure. AI tools assisted with symbolic exploration, exact-certificate generation, literature triage, numerical validation, and manuscript drafting. The author selected the claims and is responsible for the mathematical content. All theorem-critical finite identities are independently replayable in the accompanying artifacts.

References

[1] A. A. Albert and B. Muckenhoupt, "On matrices of trace zero," *Michigan Math. J.* **4** (1957), 1–3. doi:10.1307/mmj/1028990168.

[2] R. C. Thompson, "On matrix commutators," *J. Washington Acad. Sci.* **48** (1958), 306–307.

[3] A. Horn, "Eigenvalues of sums of Hermitian matrices," *Pacific J. Math.* **12** (1962), 225–241. doi:10.2140/pjm.1962.12.225.

[4] P. Fan and C.-K. Fong, "Which operators are the self-commutators of compact operators?," *Proc. Amer. Math. Soc.* **80** (1980), 58–60. doi:10.1090/S0002-9939-1980-0574508-X.

[5] A. A. Klyachko, "Stable bundles, representation theory and Hermitian operators," *Selecta Math. (N.S.)* **4** (1998), 419–445. doi:10.1007/s000290050037.

[6] A. Knutson and T. Tao, "The honeycomb model of $GL_n(\mathbb{C})$ tensor products I: proof of the saturation conjecture," *J. Amer. Math. Soc.* **12** (1999), 1055–1090. doi:10.1090/S0894-0347-99-00299-4.

[7] W. Fulton, "Eigenvalues, invariant factors, highest weights, and Schubert calculus," *Bull. Amer. Math. Soc.* **37** (2000), 209–249. doi:10.1090/S0273-0979-00-00865-X.

[8] A. Knutson, T. Tao, and C. Woodward, "The honeycomb model of $GL_n(\mathbb{C})$ tensor products II: puzzles determine facets of the Littlewood–Richardson cone," *J. Amer. Math. Soc.* **17** (2004), 19–48. doi:10.1090/S0894-0347-03-00441-7.

[9] A. Böttcher and D. Wenzel, "How big can the commutator of two matrices be and how big is it typically?," *Linear Algebra Appl.* **403** (2005), 216–228. doi:10.1016/j.laa.2005.02.012.

[10] G. Nagy, "Abelian self-commutators in finite factors," *Operators and Matrices* **1** (2007), 257–269. doi:10.7153/oam-01-17.

[11] A. Böttcher and D. Wenzel, "The Frobenius norm and the commutator," *Linear Algebra Appl.* **429** (2008), 1864–1885. doi:10.1016/j.laa.2008.05.020.

[12] C.-M. Cheng, S.-W. Vong, and D. Wenzel, "Commutators with maximal Frobenius norm," *Linear Algebra Appl.* **432** (2010), 292–306. doi:10.1016/j.laa.2009.08.008.

[13] W. B. Johnson, N. Ozawa, and G. Schechtman, "A quantitative version of the commutator theorem for zero trace matrices," *Proc. Natl. Acad. Sci. USA* **110** (2013), 19251–19255. doi:10.1073/pnas.1202411109.

[14] O. Angel and G. Schechtman, "The Hilbert–Schmidt version of the commutator theorem for zero trace matrices," *Bull. London Math. Soc.* **47** (2015), 715–719. doi:10.1112/blms/bdv045.

[15] M. Gil', "A sharp bound for the Frobenius norm of self-commutators of matrices," *Linear Multilinear Algebra* **65** (2017), 2333–2339. doi:10.1080/03081087.2016.1273875.

[16] L. Eriksson, "The exact four-level inverse commutator cost: Horn–Littlewood–Richardson facets, rank transitions, and sharp loop synthesis," submitted manuscript (2026).

[17] L. Eriksson, "Sharp costs and exact semigroups from forgotten quantum order: qutrit extremality, tetrahedral depolarization, and echo–critical clock laws," submitted manuscript (2026).