

Sharp Rank-Adaptive Bounds for Inverse Self-Commutators

Every dimension, exact sign-cut leakage, and complete extremizers

Lluis Eriksson

30 August 2026

Abstract

Let $F = F^* \in M_d(\mathbb{C})$ be nonzero and traceless, and define the inverse Hermitian-commutator cost

$$\kappa_d(F) = \inf\{\|H\|_{\text{HS}} \|K\|_{\text{HS}} : H = H^*, K = K^*, -i[H, K] = F\}.$$

Writing $\rho = \text{rank } F$ and $\mathcal{P}(F) = \|F\|_1 / 2 = \text{Tr } F_+$, we prove the dimension-free, rank-adaptive sharp bounds

$$\mathcal{P}(F) \leq \kappa_d(F) \leq \frac{\rho}{2} \mathcal{P}(F).$$

The lower equality holds exactly when the nonzero spectrum of F is centrally symmetric. The upper equality holds exactly when, up to positive scaling and sign, the nonzero spectrum is

$$(\rho - 1, -1, \dots, -1).$$

Thus the worst normalized inverse cost at fixed rank is precisely $\rho/2$, including for targets embedded with arbitrary zero padding; both extrema are attained.

The lower endpoint has an exact geometric refinement. If $Q = \mathbf{1}_{(0, \infty)}(F)$ is the positive spectral projection, then

$$\kappa_d(F) - \mathcal{P}(F) = \frac{1}{2} \min_{CC^* - C^*C = 2F} (\|(1 - Q)C\|_{\text{HS}}^2 + \|CQ\|_{\text{HS}}^2).$$

Hence the entire trace-norm tax is the least wrong-way Hilbert–Schmidt leakage across the sign cut of the prescribed target. The upper bound follows from an averaged family of finite weighted shifts. Sharpness at the upper endpoint follows from only $\rho - 1$ explicit Horn–Littlewood–Richardson inequalities, valid in every ambient dimension. The subset-to-partition convention and the resulting Littlewood–Richardson coefficients are written out explicitly, so the certificate is independent of a fixed-dimensional Horn enumeration. A dependency-free exact-arithmetic replay accompanies the manuscript. No closed formula is claimed for a general interior spectrum.

1 The inverse problem and the result

Every trace-zero complex matrix is a commutator, and every finite-dimensional traceless Hermitian matrix is a self-commutator [1, 2, 4]. Existence alone leaves a metric inverse question: for a prescribed target, what is the least Hilbert–Schmidt resource of Hermitian factors that generate it? We use the unnormalized norm

$$\|X\|_{\text{HS}}^2 = \text{Tr}(X^*X)$$

and study

$$\kappa_d(F) := \inf\{\|H\|_{\text{HS}} \|K\|_{\text{HS}} : H = H^*, K = K^*, -i[H, K] = F\}. \quad (1)$$

The target is always assumed Hermitian and traceless. Set

$$\mathcal{P}(F) := \text{Tr } F_+ = \frac{1}{2} \|F\|_1, \quad \rho(F) := \text{rank } F. \quad (2)$$

For a nonzero traceless target, $\rho(F) \geq 2$.

The nearby norm literature mainly treats the forward problem, bounding a commutator in terms of its factors [8, 11, 12, 13, 18], or treats mixed operator–Hilbert–Schmidt costs for an arbitrary trace-zero matrix [14, 17]. Other self-commutator work concerns approximation or comparison with the generating operator rather than the prescribed-target minimum studied here [9, 10, 19]. Weiss studied the inverse HS–HS problem with *arbitrary*, generally non-Hermitian, factors [15]. This is a genuinely different optimization: for

$$F_4 = \text{diag}(-1, 1/3, 1/3, 1/3),$$

Weiss’s unrestricted minimum product is $4/3$ (equivalently, after balancing, each factor has norm $\sqrt{4/3}$), whereas Theorem 1.1 gives the Hermitian-factor cost $\kappa_4(F_4) = 2$. Thus neither result subsumes the other by a change of notation.

Fan and Fong’s compact self-commutator theorem uses arrangements with nonnegative partial sums to establish existence [4]. Beltiřă, Patnaik, and Weiss survey that setting, record the finite weighted-shift construction, and discuss the same one-spike energy comparison [16]. Those antecedents supply constructions and historical context, but not the minimum in (1), its rank-adaptive upper constant, or either equality classification proved here.

Exact formulas in dimensions three, four, and five can be obtained by solving specialized Horn linear programs [20, 21, 22]. Their chamber geometry grows with dimension. The present result takes a different direction: it determines the two extremal faces and the exact worst-target constant in every dimension without enumerating the Horn cone.

Theorem 1.1 (Rank-adaptive trace tax and complete extremizers). *Let $0 \neq F = F^* \in M_d(\mathbb{C})$ satisfy $\text{Tr } F = 0$, and put $\rho = \text{rank } F$. Then*

$$\mathcal{P}(F) \leq \kappa_d(F) \leq \frac{\rho}{2} \mathcal{P}(F). \quad (3)$$

Moreover:

- (i) $\kappa_d(F) = \mathcal{P}(F)$ if and only if, including multiplicities, the nonzero spectrum is

$$(a_1, \dots, a_m, -a_m, \dots, -a_1), \quad a_1 \geq \dots \geq a_m > 0. \quad (4)$$

In particular, lower equality forces $\rho = 2m$.

- (ii) $\kappa_d(F) = \rho \mathcal{P}(F)/2$ if and only if, up to multiplication by a positive scalar and replacement of F by $-F$, the nonzero spectrum is

$$(\rho - 1, -1, \dots, -1). \quad (5)$$

Zero eigenvalues may occur with arbitrary multiplicity.

Two immediate consequences make the sharpness precise.

Corollary 1.2 (Worst target at fixed rank and dimension). *For $2 \leq r \leq d$,*

$$\sup_{\substack{F=F^* \neq 0, \text{Tr } F=0 \\ \text{rank } F=r}} \frac{\kappa_d(F)}{\mathcal{P}(F)} = \frac{r}{2}.$$

Consequently the unrestricted supremum in $M_d(\mathbb{C})$ is $d/2$.

Remark 1.3 (No dimension salami). Theorem 1.1 is not an extrapolation of a five-level max formula. It is an arbitrary-dimensional extremal theorem, is adaptive to the support rank rather than the ambient dimension, and uses no enumeration of fixed- d Horn facets. The five-level result supplies one specialization of (3), while retaining separate information about the full interior chamber function.

2 One matrix and the exact sign-cut leakage identity

The first step replaces two Hermitian factors by one unrestricted matrix.

Lemma 2.1 (Balanced self-commutator reduction). *For every traceless Hermitian F ,*

$$\kappa_d(F) = \frac{1}{2} \min\{\|C\|_{\text{HS}}^2 : CC^* - C^*C = 2F\}. \quad (6)$$

The minimum is attained.

Proof. Given feasible Hermitian H, K , rescale them as $H \mapsto tH$ and $K \mapsto t^{-1}K$. Their commutator and norm product are unchanged, and t may be chosen so that the two Hilbert–Schmidt norms agree. For $C = H + iK$, direct calculation gives

$$CC^* - C^*C = -2i[H, K] = 2F, \quad \|C\|_{\text{HS}}^2 = \|H\|_{\text{HS}}^2 + \|K\|_{\text{HS}}^2 = 2\|H\|_{\text{HS}}\|K\|_{\text{HS}}.$$

This proves that the right side of (6) is at most $\kappa_d(F)$.

Conversely, write a feasible C as $C = X + iY$ with $X = X^*$ and $Y = Y^*$. Then $-i[X, Y] = F$ and

$$\|X\|_{\text{HS}}\|Y\|_{\text{HS}} \leq \frac{1}{2}(\|X\|_{\text{HS}}^2 + \|Y\|_{\text{HS}}^2) = \frac{1}{2}\|C\|_{\text{HS}}^2.$$

Thus the reverse inequality holds. Feasibility follows, for example, from the weighted-shift construction in Section 4. The constraint is closed, and every norm sublevel set is compact in finite dimension, so the minimum is attained. $\square$

The cost above the trace-norm floor has an exact interpretation. It is useful both for the equality proof and for applications in which the positive spectral subspace defines a desired output sector.

Theorem 2.2 (Exact sign-cut leakage formula). *Let $Q = \mathbf{1}_{(0,\infty)}(F)$ be the positive spectral projection of a traceless Hermitian F . Then*

$$\kappa_d(F) - \mathcal{P}(F) = \frac{1}{2} \min_{CC^* - C^*C = 2F} \left(\|(\mathbf{1} - Q)C\|_{\text{HS}}^2 + \|CQ\|_{\text{HS}}^2 \right). \quad (7)$$

The minimizers in (7) are exactly the cost minimizers in (6).

Proof. Fix a feasible C and put $R = CC^*$ and $S = C^*C$. Since $R - S = 2F$ and Q is the positive spectral projection of F ,

$$2\mathcal{P}(F) = \text{Tr}(2F)_+ = \text{Tr } Q(R - S).$$

Also $\text{Tr } R = \text{Tr } S = \|C\|_{\text{HS}}^2$. Therefore

$$\begin{aligned} \frac{1}{2}\|C\|_{\text{HS}}^2 - \mathcal{P}(F) &= \frac{1}{2}(\text{Tr } R - \text{Tr } Q(R - S)) \\ &= \frac{1}{2}(\text{Tr}((\mathbf{1} - Q)R) + \text{Tr}(QS)) \\ &= \frac{1}{2}(\|(\mathbf{1} - Q)C\|_{\text{HS}}^2 + \|CQ\|_{\text{HS}}^2). \end{aligned} \quad (8)$$

This identity holds pointwise on the feasible set. Minimizing and using Lemma 2.1 proves the result. $\square$

Corollary 2.3 (Quantitative near-pairing). *If $\kappa_d(F) \leq (1 + \varepsilon)\mathcal{P}(F)$, there is an optimal factor C with*

$$\|(\mathbf{1} - Q)C\|_{\text{HS}}^2 + \|CQ\|_{\text{HS}}^2 \leq 2\varepsilon\mathcal{P}(F).$$

Thus a small relative trace tax quantitatively forces an almost one-way factor across the sign cut.

3 The lower endpoint: central spectral symmetry

The lower bound and its equality classification now follow without Horn theory.

Proposition 3.1 (Trace floor and equality). *For every nonzero traceless Hermitian F , $\kappa_d(F) \geq \mathcal{P}(F)$. Equality holds exactly for the spectra in (4).*

Proof. The lower bound is immediate from the nonnegative right side of (7). Suppose equality holds and choose an optimal C . The two nonnegative terms in (7) vanish, so

$$(\mathbf{1} - Q)C = 0, \quad CQ = 0.$$

Thus $R = CC^*$ is supported in Q and $S = C^*C$ is supported in $\mathbf{1} - Q$. On $\ker F$, however, $R - S = 2F = 0$ while $R = 0$, so $S = 0$ there as well. Hence R and S are supported on the positive and negative spectral subspaces respectively, and in particular $RS = 0$. They have the same nonzero eigenvalues, namely the squared singular values of C . Since $2F = R - S$, the positive and negative nonzero eigenvalues of F occur in equal pairs. This gives (4).

Conversely, suppose F has an orthonormal eigenbasis $e_1, \dots, e_m, f_1, \dots, f_m$ on its support, with

$$Fe_j = a_j e_j, \quad Ff_j = -a_j f_j.$$

Define

$$C = \sum_{j=1}^m \sqrt{2a_j} |e_j\rangle\langle f_j|. \quad (9)$$

Then $(CC^* - C^*C)/2 = F$ and

$$\frac{1}{2} \|C\|_{\text{HS}}^2 = \sum_{j=1}^m a_j = \mathcal{P}(F).$$

Lemma 2.1 proves equality. □

Remark 3.2. The equality face is spectral, not entrywise: F need not be presented in paired block form. Unitary covariance transports (9) to any matrix with the same spectrum.

4 A rank-adaptive weighted-shift upper bound

Let the positive eigenvalues of F be $a_1, \dots, a_m > 0$, and write the negative eigenvalues as $-b_1, \dots, -b_n < 0$. Tracelessness gives

$$\sum_{j=1}^m a_j = \sum_{k=1}^n b_k = \mathcal{P}(F) =: P, \quad m + n = \rho. \quad (10)$$

Zero eigenvalues play no role in the construction and will be placed last.

Lemma 4.1 (Excursion shift). *Let $\mu_1, \dots, \mu_d$ be an ordering of the eigenvalues of F whose partial sums*

$$s_k = \sum_{j=1}^k \mu_j$$

are nonnegative. In the corresponding eigenbasis, define

$$Ce_{k+1} = \sqrt{2s_k} e_k \quad (1 \leq k < d), \quad Ce_1 = 0. \quad (11)$$

Then

$$\frac{1}{2}(CC^* - C^*C) = \text{diag}(\mu_1, \dots, \mu_d), \quad \frac{1}{2} \|C\|_{\text{HS}}^2 = \sum_{k=1}^{d-1} s_k.$$

Proof. The diagonal entries of $(CC^* - C^*C)/2$ are

$$s_1, s_2 - s_1, \dots, s_{d-1} - s_{d-2}, -s_{d-1},$$

which are $\mu_1, \dots, \mu_d$ because the total sum is zero. The norm identity follows by summing the squared shift weights. $\square$

Put all positive eigenvalues first, all negative eigenvalues next, and all zeros last. Every partial sum is nonnegative, independently of the orders chosen inside the positive and negative blocks. Let $A(\pi, \sigma)$ be the cost in Lemma 4.1 for permutations $\pi \in S_m$ and $\sigma \in S_n$.

Lemma 4.2 (Exact permutation average). *The block-order shift costs have average*

$$\frac{1}{m!n!} \sum_{\pi \in S_m} \sum_{\sigma \in S_n} A(\pi, \sigma) = \frac{m+n}{2} P = \frac{\rho}{2} P. \quad (12)$$

Proof. The zero tail has zero partial sums, so compute with the first $\rho = m+n$ positions. An entry at position j has coefficient $\rho - j$ in $\sum_{k=1}^{\rho-1} s_k$. Averaging over the positive block, each a_i has mean coefficient

$$\rho - \frac{m+1}{2}.$$

Each negative magnitude b_j has mean coefficient

$$\frac{n-1}{2}.$$

Using (10), the averaged cost is

$$P \left(\rho - \frac{m+1}{2} - \frac{n-1}{2} \right) = \frac{\rho}{2} P.$$

$\square$

At least one block order has cost no larger than the average. Lemmas 2.1, 4.1, and 4.2 prove

$$\kappa_d(F) \leq \frac{\rho}{2} \mathcal{P}(F). \quad (13)$$

This argument is rank-adaptive because zero eigenvalues contribute neither mass nor steps to the excursion.

5 Necessity at the upper endpoint

The averaging proof carries enough rigidity to classify every possible upper extremizer.

Proposition 5.1 (Flat sign blocks and a single spike). *If $0 \neq F$ and $\kappa_d(F) = \rho \mathcal{P}(F)/2$, then one of F_+ and F_- has rank one, and the nonzero eigenvalues on the other sign are all equal. Equivalently, the nonzero spectrum has the form (5), up to scale and sign.*

Proof. Every block-order shift is feasible, hence

$$\kappa_d(F) \leq A(\pi, \sigma)$$

for all π, σ . If $\kappa_d(F) = \rho P/2$, equality with the average (12) forces every $A(\pi, \sigma)$ to equal $\rho P/2$.

Suppose two positive eigenvalues a and a' occupy block positions $u < v$. Interchanging only these entries changes the weighted sum of partial sums by

$$(v-u)(a-a').$$

Since all block orders have the same cost, $a = a'$. Thus all positive eigenvalues are equal. The same transposition argument shows that all negative magnitudes are equal. The spectrum is therefore

$$\left(\underbrace{P/m, \dots, P/m}_m, \underbrace{-P/n, \dots, -P/n}_n \right), \quad (14)$$

apart from zeros.

Assume now that $m, n \geq 2$. Start with all m positive entries followed by all n negative entries. Swap the last positive entry and the first negative entry. Just after the swapped negative entry, the partial sum is

$$P \left(1 - \frac{1}{m} - \frac{1}{n} \right) \geq 0,$$

and after the following positive entry the path rejoins the original partial-sum path. Thus the swapped order remains feasible. Its cost is smaller than the block-order cost by

$$P \left(\frac{1}{m} + \frac{1}{n} \right) > 0,$$

contradicting upper equality. Hence $\min\{m, n\} = 1$, which gives (5). $\square$

The only remaining issue is sufficiency: a one-spike target might conceivably admit a coherent non-shift factor below the displayed weighted-shift cost. A short universal Horn certificate rules this out.

6 A universal Horn certificate for the one-spike family

We use only a special elementary face of the classical Horn–Littlewood–Richardson theorem [3, 5, 6, 7]. If Hermitian matrices with decreasing spectra α, β have sum with decreasing spectrum γ , then every Horn triple (I, J, K) of equal-size subsets satisfies

$$\sum_{k \in K} \gamma_k \leq \sum_{i \in I} \alpha_i + \sum_{j \in J} \beta_j. \quad (15)$$

For completeness, fix the convention that an increasing ℓ -subset

$$A = \{a_1 < \dots < a_\ell\} \subseteq \{1, \dots, d\} \quad \text{corresponds to} \quad \pi(A) = (a_\ell - \ell, a_{\ell-1} - (\ell - 1), \dots, a_1 - 1).$$

Thus the initial subset $I = \{1, \dots, \ell\}$ corresponds to the zero partition. Consequently

$$I_\ell = \{1, \dots, \ell\}, \quad J_\ell = K_\ell = \{1, d - \ell + 2, \dots, d\} \quad (16)$$

is a Horn triple: $J_\ell = K_\ell$ literally, and the unit property of the zero partition gives

$$c_{\pi(I_\ell), \pi(J_\ell)}^{\pi(K_\ell)} = c_{0, \pi(J_\ell)}^{\pi(J_\ell)} = 1.$$

This verifies the convention-sensitive part of the certificate directly for every d and ℓ used below.

Lemma 6.1 (Removing a common singular floor). *At a minimizer in (6), the least squared singular value of C is zero.*

Proof. Put $S = C^*C$ and use the polar decomposition $C = US^{1/2}$. If the least eigenvalue of S were $t > 0$, then U would be unitary and

$$C' = U(S - t\mathbf{1})^{1/2}$$

would satisfy

$$C'C'^* = CC^* - t\mathbf{1}, \quad C'^*C' = C^*C - t\mathbf{1}.$$

The self-commutator would be unchanged while $\|C'\|_{\text{HS}}^2 = \|C\|_{\text{HS}}^2 - dt$, contradicting minimality. $\square$

Proposition 6.2 (Sharp one-spike lower certificate). *Let the nonzero spectrum of F be*

$$\left(P, -\frac{P}{\rho-1}, \dots, -\frac{P}{\rho-1}\right), \quad P > 0,$$

with any number of zeros between the positive and negative levels. Then

$$\kappa_d(F) = \frac{\rho}{2}P.$$

Proof. Choose an optimal C , and let

$$p_1 \geq \dots \geq p_d = 0$$

be the decreasing squared singular values, using Lemma 6.1. The spectra of CC^* and $-C^*C$ are

$$\alpha = (p_1, \dots, p_{d-1}, 0), \quad \beta = (0, -p_{d-1}, \dots, -p_1).$$

Apply (15) to the triple (16) for each Horn subset size

$$1 \leq \ell \leq \rho - 1.$$

The right side telescopes exactly:

$$\sum_{i \in I_\ell} \alpha_i + \sum_{j \in J_\ell} \beta_j = (p_1 + \dots + p_\ell) - (p_1 + \dots + p_{\ell-1}) = p_\ell.$$

The set K_ℓ selects the positive spike and the lowest $\ell - 1$ negative eigenvalues of $2F$. Therefore

$$p_\ell \geq 2P \left(1 - \frac{\ell - 1}{\rho - 1}\right) = \frac{2P(\rho - \ell)}{\rho - 1}. \quad (17)$$

Summing (17) gives

$$\|C\|_{\text{HS}}^2 = \sum_{j=1}^{d-1} p_j \geq \sum_{\ell=1}^{\rho-1} p_\ell \geq \frac{2P}{\rho-1} \sum_{q=1}^{\rho-1} q = \rho P.$$

Lemma 2.1 yields $\kappa_d(F) \geq \rho P/2$. The weighted-shift upper bound (13) gives equality. $\square$

Propositions 3.1, 5.1, and 6.2 complete the proof of Theorem 1.1. Corollary 1.2 follows by choosing a zero-padded one-spike spectrum at every rank.

7 What the theorem adds

Table 1 separates the classical ingredients from the new conclusions.

The result also clarifies the relation between low-dimensional exact formulas and arbitrary dimension. A fixed-dimensional formula resolves the complete polyhedral interior of the inverse cost. Theorem 1.1 instead resolves the global trace floor, the worst normalized target, and both equality loci uniformly over all dimensions. Neither statement contains the other.

8 Reproducibility, scope, and limitations

The analytic proof is contained in the manuscript. The accompanying replay script uses only the Python standard library and exact rational arithmetic to check:

1. the weighted-shift commutator identity in the free coefficient module;
2. the exact block-permutation average on rational examples;

Input	Role in this paper
Trace-zero/self-commutator existence	Guarantees the inverse problem is feasible; not claimed as new.
Weighted shifts and positive partial sums	Supply a family of explicit feasible factors; the exact permutation average and rank-adaptive optimization are the new use.
Horn–Littlewood–Richardson inequalities	Classical spectral-sum theorem; only the $\rho - 1$ triples (16) are used to certify the sharp endpoint.
Finite-dimensional trace variationality	Underlies the trace floor; the pointwise identity (8) identifies its complete defect.
New theorem package	Sharp rank constant $\rho/2$; both equality classifications; exact sign-cut leakage formula; zero-padded endpoint certificates at every rank.

Table 1: Novelty boundary and proof architecture.

- the adjacent-swap strict deficit for every two-sided flat multiplicity through a configurable dimension;
- the zero-padded one-spike shift cost and the summed Horn lower certificate at every rank through a configurable dimension.

These computations audit algebra and indexing; they are not substitutes for the arbitrary-dimensional proofs.

The exact release is preserved under ARR record [ARR-2026-1D2QV1RP1292JREW](#). Its standard-library replay is `src/repro/verify_extremal_tax.py`, with SHA-256 `e0570dbd13215e84f0449e200f1f4c6e32a1351cf8e7359da31e51a411a3cd2a`; the release also preserves the exact JSON output and reproduction instructions.

The theorem is finite-dimensional and uses the unnormalized Hilbert–Schmidt norm. It does not address infinite factors, unbounded canonical commutation relations, operator-norm minimization, graph-local or gate-local factors, fixed control algebras, or laboratory time and energy. The Horn theorem is a classical imported result, not formalized in the replay; the manuscript does make the convention-sensitive Littlewood–Richardson coefficient in its special certificate explicit. No entrywise optimal constructor is claimed for a general interior spectrum; the shift supplies the sharp universal upper bound, while the exact interior cost remains a richer Horn optimization.

The exact upper equality classification is qualitative. A natural open problem is to determine a sharp modulus of stability: for fixed rank ρ , quantify how close the normalized nonzero spectrum must be, up to sign, to the one-spike ray when $\kappa_d(F)/\mathcal{P}(F)$ lies within ε of $\rho/2$. No such upper-endpoint stability estimate is claimed here.

Literature-search boundary. The searches conducted for this project found self-commutator existence and compactness criteria, forward Frobenius inequalities, mixed operator–HS commutator bounds, and spectral-sum theory, but no prescribed-target inverse HS–HS theorem with the rank-adaptive constant and both equality classifications. This supports the stated novelty boundary but is not a claim of exhaustive bibliographic priority.

AI assistance. The author directed the research program, selected the theorem, reviewed the mathematical claims, and accepts responsibility for the manuscript. OpenAI Codex assisted with repository review, candidate comparison, algebraic exploration, literature triage, exact-arithmetic replay, typesetting, and drafting. AI assistance does not replace the explicit proofs or the author’s responsibility for correctness.

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](https://doi.org/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](https://doi.org/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](https://doi.org/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](https://doi.org/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](https://doi.org/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](https://doi.org/10.1090/S0273-0979-00-00865-X).
- [8] C.-K. Fong, “Norm estimates related to self-commutators,” *Linear Algebra Appl.* **74** (1986), 151–156. [doi:10.1016/0024-3795\(86\)90118-7](https://doi.org/10.1016/0024-3795(86)90118-7).
- [9] P. Maher, “Self-commutator approximants,” *Proc. Amer. Math. Soc.* **134** (2006), 157–165. [doi:10.1090/S0002-9939-05-07871-8](https://doi.org/10.1090/S0002-9939-05-07871-8).
- [10] N. Filonov and Y. Safarov, “On the relation between an operator and its self-commutator,” *J. Funct. Anal.* **260** (2011), 2902–2932. [doi:10.1016/j.jfa.2011.02.011](https://doi.org/10.1016/j.jfa.2011.02.011).
- [11] 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](https://doi.org/10.1016/j.laa.2005.02.012).
- [12] 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](https://doi.org/10.1016/j.laa.2008.05.020).
- [13] 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](https://doi.org/10.1016/j.laa.2009.08.008).
- [14] 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](https://doi.org/10.1073/pnas.1202411109).
- [15] G. Weiss, “Commutators of Hilbert–Schmidt operators II,” *Integral Equations Operator Theory* **3** (1980), 574–600. [doi:10.1007/BF01702316](https://doi.org/10.1007/BF01702316).
- [16] D. Belitiță, S. Patnaik, and G. Weiss, “ $B(H)$ -commutators: a historical survey II and recent advances on commutators of compact operators,” in *The Varied Landscape of Operator Theory*, Theta Series in Advanced Mathematics, vol. 17, Theta, Bucharest, 2014, pp. 57–75. [arXiv:1303.4844](https://arxiv.org/abs/1303.4844).
- [17] 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](https://doi.org/10.1112/blms/bdv045).
- [18] 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](https://doi.org/10.1080/03081087.2016.1273875).
- [19] T. Zhang, “On a conjecture of λ -Aluthge transforms and Hilbert–Schmidt self-commutators,” [arXiv:2603.04655](https://arxiv.org/abs/2603.04655) (2026).
- [20] L. Eriksson, “Sharp costs and exact semigroups from forgotten quantum order: qutrit extremality, tetrahedral depolarization, and echo-critical clock laws,” ARR research paper [ARR-2026-1PT4297HNX9T4RBD](https://arxiv.org/abs/2026.1PT4297HNX9T4RBD) (2026).
- [21] L. Eriksson, “The exact four-level inverse commutator cost: Horn–Littlewood–Richardson facets, rank transitions, and sharp loop synthesis,” ARR research paper [ARR-2026-3M1EEG1T689ADSMW](https://arxiv.org/abs/2026.3M1EEG1T689ADSMW) (2026); [ai.viXra:2608.0031](https://arxiv.org/abs/2608.0031).
- [22] L. Eriksson, “The exact five-level inverse commutator cost,” ARR research paper [ARR-2026-37B8R0QTA894GTFF](https://arxiv.org/abs/2026.37B8R0QTA894GTFF) (2026); [ai.viXra:2608.0049](https://arxiv.org/abs/2608.0049).