

Optimal stability and all equality cases at balanced inertia (5,5)

Lluis Eriksson

Independent researcher. ARR v1 - ARR-2026-6ZJY1SSSJA98WBWX. 6 September 2026.

Abstract

We determine the optimal forward stability coefficient for the inverse Hilbert-Schmidt self-commutator cost at balanced inertia (5,5). The coefficient is 113/152, for every prescribed number of ambient zero eigenvalues. The previously established polynomial vertex reduction gives 267 ordered spectral pairs. Rational primal and dual certificates compute all their costs in dimension ten; the forward estimate uses only the primal witnesses and therefore extends to every larger dimension. An elementary extension of the preceding size-two Horn obstruction supplies the matching boundary cost. We also classify the entire equality locus on the compact spectral closure: it consists of exactly four ordered pairs. A rational midpoint witness excludes the two candidate contact segments. Thus the forward inequality is strict at every target of exact inertia (5,5). The finite verifier reconstructs 191,353 Horn triples and checks 51,282,604 integer inequalities. Classical Horn sufficiency and the previously proved sharp reverse estimate are attributed explicitly. The constants at balanced multiplicity at least six remain open.

1. The cost and the new result

For a nonzero traceless Hermitian matrix F of dimension d define

$$\kappa_d(F) = \frac{1}{2} \min \{ \|C\|_{HS}^2 : CC^* - C^*C = 2F \}, \quad P = \operatorname{tr} F_+.$$

Write its normalized positive and negative-magnitude spectra as decreasing lists a, b in the ordered simplex

$$S_5 = \{ a_1 \geq \dots \geq a_5 \geq 0 : \sum_i a_i = 1 \}.$$

The target spectrum is P times the concatenation of a , $d-10$ zeros, and the reversed negative list. The exact inertia stratum requires $a_5, b_5 > 0$; its compact closure permits zeros. Throughout this closure the reference ceiling is held fixed at 3. Put

$$e = (1, 0, 0, 0, 0), \quad u = (1/5, 1/5, 1/5, 1/5, 1/5),$$

$$E(a) = 2(1 - a_1), \quad U(a) = \sum_{i=1}^5 |a_i - 1/5|,$$

$$D_*(a, b) = \min \{ E(a) + U(b), U(a) + E(b) \}, \quad \delta_d(a, b) = 3 - \kappa_d(F)/P. \quad (1)$$

Theorem 1. For every fixed $d \geq 10$ and every exact inertia-(5,5,d-10) target,

$$\frac{113}{152}D^* \leq \delta_d \leq 2D^*. \quad (2)$$

Both constants are optimal for every fixed d. The right inequality and its sharpness are the prior reverse stability theorem [1]. The new assertion is the optimal left constant.

Theorem 2. On the compact domain $S_5 \times S_5$, equality in the left inequality of (2) holds precisely at

$$(e, u), \quad (u, e), \quad (A, B), \quad (B, A), \quad (3)$$

where

$$A = (13/20, 7/80, 7/80, 7/80, 7/80), \quad B = (1/2, 1/2, 0, 0, 0). \quad (4)$$

Consequently the left inequality is strict at every target of exact inertia (5,5). The equality classification and sharpness hold in every ambient dimension d at least ten, without assuming general invariance of the exact cost under zero padding.

The predecessor [2] proves the finite reduction for arbitrary balanced multiplicity and evaluates multiplicity four as 4/7. The earlier consolidated article [1] gives the inertia ceiling, the sharp reverse coefficient, and balanced constants 1/2 and 17/36 at multiplicities two and three. We use those results with their original normalization. This article evaluates the previously open next multiplicity; it does not introduce the spectral-sum machinery or the polynomial reduction.

2. Spectral optimization and complete finite coverage

Put $R = CC^*/2$ and $S = C^*C/2$. They have a common nonnegative spectrum s and satisfy $R-S=F$. Conversely, if $R = U\text{diag}(s)U^*$ and $S = V\text{diag}(s)V^*$, then $C = \sqrt{2}U\text{diag}(\sqrt{s})V^*$ realizes their difference. Thus the cost is the minimum of $\sum_i s_i$ under the common-spectrum condition.

Classical Horn sufficiency [3,4] makes that condition a finite system. For decreasing target spectrum lambda it is

$$s_1 \geq \dots \geq s_d \geq 0, \quad \sum_{i \in I} s_i - \sum_{j \in J} s_{d+1-j} \geq \sum_{k \in K} \lambda_k \quad (5)$$

for every recursive Horn triple (I,J,K). Subtracting s_d from every entry preserves the inequalities and decreases the objective, so an optimum has $s_d = 0$. These facts also give convexity and continuity of the value on the closed ordered trace-zero chamber: its finite dual representation is a maximum of finitely many linear forms. This convexity is used only on that chamber. Changing the order of arbitrary Hermitian matrices is outside the argument.

For completeness we recall how the finite set in [2] covers the whole domain. Define

$$Q_r = \{a \in S_5 : a_r \geq 1/5 \geq a_{r+1}\}, \quad 1 \leq r \leq 4.$$

On Q_r , U is the affine function $2(\sum_{i=1}^r a_i - r/5)$. Besides u , the cell vertices are

$$v_{kl} = (A_{kl}\mathbf{1}_k, (1/5)\mathbf{1}_{l-k}, \mathbf{0}_{5-l}), \quad A_{kl} = \frac{5-l+k}{5k},$$

with $1 \leq k \leq r \leq l < 5$. The bold vectors consist of the indicated numbers of ones or zeros. Every such vertex is joined to u ; two base vertices share an edge precisely when their k indices or their l indices agree.

Here is a direct explanation of this description. Subtract $1/5$ from the entries and write the nonnegative first r entries as decreasing steps with coefficients $\alpha_k \geq 0$. Write the negatives of the remaining entries as increasing steps with coefficients $\beta_l \geq 0$. The constraints become

$$\sum_{k=1}^r k\alpha_k = \sum_{l=r}^4 (5-l)\beta_l, \quad \sum_{l=r}^4 \beta_l \leq 1/5.$$

The zero point is the apex. Any other vertex must saturate the last inequality, since otherwise radial scaling yields a segment. At that base, p positive alpha coefficients and q positive beta coefficients give face dimension $p+q-2$. Vertices therefore have $p+q=1$; edges have $p+q=3$. This proves the formulas and the edge rule.

The union of these cells has 11 vertices and 30 edges. Set $q=E-U$. Start with all ordered pairs from the full eleven-vertex set $\{u\} \cup \{v_{kl} : 1 \leq k \leq l < 5\}$, including pairs that contain the apex u . For a vertex v and an edge $[x,y]$, whenever $q(v)$ lies strictly between $q(x)$ and $q(y)$, add the two ordered pairs $(v,z),(z,v)$, where

$$z = (1-t)x + ty, \quad t = \frac{q(v) - q(x)}{q(y) - q(x)}. \quad (6)$$

Remove duplicates. This gives the set W of 267 ordered pairs. It is exactly the union of vertices of all product cells cut by $q(a) \leq q(b)$ or the opposite inequality: cutting a polytope adds vertices only on its edges, and an edge of a product fixes a vertex in one factor. Edges lying in the cutting hyperplane introduce no new vertices. On each resulting cell D^* is affine. Therefore an upper bound for $\kappa_d + (113/152)D^*$ at all these vertices proves that upper bound on every point by convexity.

The verifier reconstructs the vertices by solving active constraint systems over the rationals and the edges by exact face ranks. It compares the full resulting set against the supplied witnesses. Thus a successful run does not merely verify a selected sample of spectral pairs.

3. The exact certificates in dimension ten

The supplied data contain one primal common spectrum and one dual certificate for each of the 139 pairs up to swapping a,b . Sign symmetry of the cost permits this reduction; the verifier still checks primal feasibility at all 267 ordered pairs explicitly.

The number of recursive Horn triples at sizes one through nine is

$$55, 1287, 12140, 46208, 71973, 46208, 12140, 1287, 55. \quad (7)$$

Their sum is 191,353. For each proposed common spectrum the verifier checks all inequalities (5), ordering, normalization, nonnegativity, and

$$\sum_i s_i + \frac{113}{152} D^*(a, b) \leq 3. \quad (8)$$

The 525 nonzero dual terms are also checked: each is an admissible Horn or ordering inequality with nonnegative rational weight, their total coefficient in each nonnegative optimization variable is at most one, and their right-side total equals the primal objective exactly. This computes each vertex cost, although feasibility and (8) alone suffice for the global estimate.

The optimized left sides contain only 2,188 distinct coefficient vectors after fixing $s_{10} = 0$. The exploratory solver merges identical left sides by keeping their strongest right side. This is an exact reduction of duplicate inequalities. The verification stage reconstructs the unmerged system and uses no optimization routine. Integer arrays accelerate arithmetic; explicit magnitude checks rule out integer overflow. Dual objectives use rational arithmetic.

Of the 267 vertices, exactly the four pairs in (3) have zero slack in (8). The other 263 have positive slack, at least $7/80$. The sharp pair (A,B) has the witness

$$s = (59/80, 1/2, 7/40, 7/80, 7/80, 0, 0, 0, 0, 0), \quad (9)$$

whose trace is $127/80$. Horn sufficiency and the convex-cell argument prove the forward inequality for dimension ten. Append zero blocks to a realizing factor to get the same upper bound in every larger dimension. This proves validity of the forward bound in Theorem 1, including the entire compact closure.

4. A dimension-independent obstruction

Lemma 3. Suppose F has positive eigenvalues $a_1 \geq \dots \geq a_m > 0$ with $m \geq 4$ and at least two negative magnitudes $b_1 \geq b_2 > 0$. Here a, b denote the actual eigenvalues, without dividing by P . Then

$$\kappa_d(F) \geq 3b_2 - a_2 + \sum_{i=4}^m a_i. \quad (10)$$

Proof. Apply the common-spectrum Horn inequality to $-F$ with

$$I = (1, 3), \quad J = (1, d-1), \quad K = (2, d-1).$$

This size-two triple is admissible for every $d \geq 4$: the index sum equality holds, and its three recursive size-one tests reduce to $2 \leq 3$, $d \leq d$, and $4 \leq d$. The inequality is

$$s_1 + s_3 - s_2 - s_d \geq b_2 - a_2.$$

The matrix inequalities $F \leq R$ and $-F \leq S$ give $s_2 \geq b_2$ and $s_i \geq a_i$ by the min-max principle. Hence

$$\sum_i s_i \geq s_1 + s_2 + s_3 + \sum_{i=4}^m s_i \geq 3b_2 - a_2 + \sum_{i=4}^m a_i + s_d.$$

Discard $s_d \geq 0$ and minimize. QED.

The size-two mechanism is the earlier obstruction [2, Lemma 3]; the present extension retains every positive tail term rather than only the fourth. At (A,B), (10) gives 127/80. Combined with (9) and zero padding, it proves the exact cost in every $d \geq 10$. Directly $D^*(A, B) = 19/10$, and therefore the ratio of deficit to distance is 113/152.

To put sharpness inside exact inertia, leave A fixed and set

$$B_\eta = (1/2 - 3\eta/2, 1/2 - 3\eta/2, \eta, \eta, \eta), \quad 0 < \eta < 1/5.$$

This is a decreasing strictly positive list of mass one. Exact computation gives

$$D^*(A, B_\eta) = 19/10 - 6\eta,$$

$$127/80 - 9\eta/2 \leq \kappa_d(A, 0, -B_\eta) \leq 127/80 + (339/76)\eta. \quad (11)$$

The lower inequality is (10), and the upper inequality is the already proved forward bound. Both converge to 127/80, so the ratios tend to 113/152. This proves sharpness for each prescribed number $d-10$ of zeros, without interchanging ambient dimension and the limit.

5. Excluding every further equality point

Let $g_d = 3 - \kappa_d - (113/152)D^*$ on normalized spectra. On each cut cell it is concave and nonnegative. If it vanishes at a point expressed as a convex combination of the cell vertices, every vertex with positive coefficient must itself have zero slack. Thus equality can occur only in the convex hull of the zero-slack vertices belonging to that cell.

The exact cell-incidence check shows that each cell contains at most two of the four vertices (3). The only candidate nontrivial contact sets are the segments joining (e,u) to (A,B), and the sign-swapped segment. It remains to exclude their interiors.

At the midpoint of the first segment the spectra are

$$a = (33/40, 7/160, 7/160, 7/160, 7/160),$$

$$b = (7/20, 7/20, 1/10, 1/10, 1/10).$$

A common-spectrum witness, checked against every dimension-ten Horn inequality, is

$$s = (33/40, 19/40, 1/8, 9/80, 1/10, 7/160, 7/160, 0, 0, 0). \quad (12)$$

It has trace 69/40, while $D^* = 19/20$, giving

$$g_{10} \geq 3 - 69/40 - (113/152)(19/20) = 91/160 > 0. \quad (13)$$

The same lower bound holds for g_d after padding. Concavity on the segment then gives, at its affine parameter $0 < t < 1$,

$$g_d(t) \geq (91/80) \min \{t, 1 - t\} > 0.$$

Sign symmetry covers the other segment. Finally (e,u) and (u,e) have cost 3 by the one-spike formula [1], and Section 4 proves equality at (A,B),(B,A). This proves Theorem 2.

6. Reproducibility and limits

The package supplies the rational vertex certificates, the additional midpoint witness, the exact cell geometry, and the recursive Horn generator. Run `python verify_n5.py` with NumPy installed. No optimizer is used in that check. It verifies 268 times 191,353, or 51,282,604, integer Horn inequalities, including the midpoint, plus the dual and geometric conditions. To regenerate exploratory proposals, install SciPy and run `python horn_fast.py --dimension 10`, `python propose_n5.py`, and `python propose_midpoint.py` in a disposable copy.

The separate LR-tableau verifier can reconstruct every Horn triple without consulting its cache: run `python review/independent_review.py --fresh-lr`. Without this option it reuses a cache only when its generator hash matches, and otherwise regenerates it. Its report records whether the triples were reconstructed or read from cache. Run either mode in a disposable copy because it writes a report and may refresh the cache.

Finite checks complement the analytic proofs above; they do not prove Horn sufficiency or substitute for those proofs. A separate internal review is distributed with the package and records its precise scope. It is not external refereeing. The manuscript and certificate hashes identify the reviewed version.

The optimal constants for balanced multiplicity at least six and the general interior cost remain unresolved here. The generic count of polynomially many Horn programs comes from [2]; neither it nor the present faster proposal implementation proves polynomial total running time. Our bounded literature comparison distinguishes general commutator norm estimates from the fixed-spectrum self-commutator optimization, but does not establish worldwide bibliographic priority.

References

- [1] L. Eriksson, *Sharp inertia ceilings and optimal stability for inverse self-commutators*, consolidated version, ARR-2026-24M24KDPZK8HDBQ9 v1. [Record and sources](#).
- [2] L. Eriksson, *Polynomial vertex reduction and optimal stability at balanced inertia (4,4)*, ARR-2026-54Q3HMFJ0Z8CZB4T v1, 5 September 2026. [Record and certificates](#).
- [3] W. Fulton, *Eigenvalues, invariant factors, highest weights, and Schubert calculus*, Bulletin of the American Mathematical Society 37 (2000), 209-249. In particular Section 1 and Theorem 1. [Primary text](#).
- [4] A. Knutson and T. Tao, *The honeycomb model of $GL(n)$ tensor products I: proof of the saturation conjecture*, Journal of the American Mathematical Society 12 (1999), 1055-1090. [Primary text](#).
- [5] O. Angel and G. Schechtman, *The Hilbert Schmidt version of the commutator theorem for zero trace matrices*, 2015. This treats a general commutator norm problem with a different objective. [Primary text](#).