

Polynomial vertex reduction and optimal stability at balanced inertia (4,4)

Lluis Eriksson

Independent researcher. Research manuscript, 5 September 2026.

Abstract

The optimal forward stability constant at a balanced inertia ceiling is an optimization over two ordered spectral simplexes. We reduce that optimization, for every multiplicity N and fixed ambient dimension d , to a rational list of $O(N^5)$ spectral pairs, independent of d . The reduction uses an explicit description of the vertices and edges of the ordered simplex cut by its uniform coordinates. Each remaining value is an exact Horn linear program; the polynomial bound counts programs, not their size or total running time. For inertia (4,4), the reduction has 89 ordered vertices. Rational common-spectrum witnesses satisfy all 8,752 Horn inequalities at each vertex and prove the optimal bound $\frac{4}{7}D^* \leq \delta$. A two-row Horn argument proves the matching boundary cost in every ambient dimension, and an explicit perturbation establishes sharpness for every prescribed number of zeros. We also classify all equality pairs on the compact spectral closure: the two one-spike, opposite-flat pairs and two additional boundary pairs. In particular the forward bound is strict on every exact inertia-(4,4) target. Together with the previously sharp reverse bound, this gives $\frac{4}{7}D^* \leq \delta \leq \frac{3}{2}D^*$. The general optimal constants for $N \geq 5$ are not evaluated here. The classical Horn theorem is imported; all finite certificates for the new inertia-(4,4) result are supplied.

1. Definitions, antecedents, and statements

For a nonzero traceless Hermitian d by d matrix F , put

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

Assume F has exactly N positive and N negative eigenvalues, with $d - 2N$ further zeros. Its decreasing positive and negative-magnitude lists, divided by P , are a, b in

$$S_N = \{ a_1 \geq \dots \geq a_N \geq 0 : \sum_i a_i = 1 \}.$$

The strict inertia stratum has $a_N, b_N > 0$; its compact closure allows zeros. Set

$$e = (1, 0, \dots, 0), \quad u = (1/N, \dots, 1/N), \quad H_N = (N + 1)/2,$$

$$E(a) = \|a - e\|_1 = 2(1 - a_1), \quad U(a) = \|a - u\|_1,$$

$$D^*(a, b) = \min \{ E(a) + U(b), U(a) + E(b) \}, \quad \delta_d(a, b) = H_N - \kappa_d(F)/P.$$

The constant H_N remains fixed on the closure, even when the actual inertia of a boundary target drops. The preceding consolidated paper [1] proves the sharp inertia ceiling, convexity and continuity on the ordered chamber, the two optimal constants for unequal multiplicities, and the sharp reverse bound

$$\delta_d \leq \frac{N-1}{2} D^*.$$

Its balanced forward constants are $1/2$ for $N = 2$ and $17/36$ for $N = 3$. It leaves $N \geq 4$ open. The present work adds a general finite reduction and closes $N = 4$; it does not reclaim those antecedents.

Theorem A (polynomial vertex reduction). For $N \geq 2$ and $d \geq 2N$, the largest coefficient $\gamma_{N,d}$ satisfying $\gamma_{N,d} D^* \leq \delta_d$ on the exact $(N, N, d - 2N)$ stratum is a positive rational number. There is an explicit rational set W_N , independent of d , such that

$$\gamma_{N,d} = \min_{(a,b) \in W_N, D^*(a,b) > 0} \frac{H_N - \kappa_d(a, 0, -b)}{D^*(a,b)}. \quad (1)$$

Here κ_d uses the decreasing target $(a_1, \dots, a_N, 0, \dots, 0, -b_N, \dots, -b_1)$. With

$$V_N = 1 + \frac{N(N-1)}{2}, \quad L_N = \frac{N(N-1)(2N-1)}{6},$$

the set can be constructed using at most

$$|W_N| \leq V_N^2 + 2V_N L_N = O(N^5) \quad (2)$$

candidate pairs. Evaluating (1) therefore requires at most this many rational Horn linear programs. This bound is on the number of programs; the number of Horn inequalities and the cost of solving a program depend strongly on d .

Theorem B (sharp stability at balanced inertia (4,4)). For every fixed $d \geq 8$ and every F of exact inertia $(4, 4, d - 8)$,

$$\frac{4}{7} D^* \leq \delta_d \leq \frac{3}{2} D^*. \quad (3)$$

Both coefficients are optimal for every fixed d . The right coefficient is the result of [1]; the new coefficient is $4/7$.

2. The spectral program and the continuity used below

Put $R = CC^*/2$ and $S = C^*C/2$. They are positive semidefinite with the same spectrum S , satisfy $R - S = F$, and $\text{tr}(R) = \|C\|_{HS}^2/2$. Conversely, two positive semidefinite matrices with the same spectrum admit decompositions $R = U \text{diag}(s) U^*$, $S = V \text{diag}(s) V^*$. Taking $C = \sqrt{2} U \text{diag}(\sqrt{s}) V^*$ realizes their difference. Thus the optimization is exactly over common positive spectra.

For decreasing $\lambda(F)$, Horn sufficiency [2,3] gives the linear constraints

$$s_1 \geq \dots \geq s_d \geq 0, \quad \sum_{i \in I} s_i - \sum_{j \in J} s_{d+1-j} \geq \sum_{k \in K} \lambda_k \quad (4)$$

for every recursive Horn triple (I, J, K) . The objective is $\sum_i s_i$. Subtracting s_d from every entry preserves (4) and decreases the objective unless $s_d=0$. We may impose $s_d=0$. Rational λ gives a rational optimal value and rational primal and dual certificates. Feasibility also follows from the standard partial-sum weighted-shift construction; existence of an optimum follows by restricting to a bounded objective sublevel.

The epigraph in (λ, cost) is a projection of a finite polyhedron, so it is polyhedral. The value is finite on the ordered trace-zero chamber and is the maximum of finitely many affine linear forms there. In particular it is convex and continuous on that closed chamber. Convexity here does not assert convexity on the vector space of all Hermitian matrices. Appending zeros to a realizing factor gives an upper bound in larger dimensions; it does not assert general zero-padding invariance of the exact optimum.

These facts were established in the earlier inverse-commutator papers and in [1]. They are recalled to make the domain and boundary arguments explicit.

3. The entire simplex geometry in arbitrary multiplicity

For $r = 1, \dots, N-1$, define

$$Q_r = \{a \in S_N : a_r \geq 1/N \geq a_{r+1}\}.$$

These cells cover S_N . On Q_r ,

$$U(a) = 2 \left(\sum_{i=1}^r a_i - r/N \right), \quad (5)$$

so both U and $q = E - U$ are affine.

Lemma 1 (vertices and edges). In addition to u , the vertices of Q_r are

$$v_{kl} = (A_{kl} \mathbf{1}_k, N^{-1} \mathbf{1}_{l-k}, \mathbf{0}_{N-l}), \quad A_{kl} = \frac{N-l+k}{Nk}, \quad 1 \leq k \leq r \leq l < N. \quad (6)$$

Here the bold vectors have the indicated lengths and consist entirely of ones or zeros.

Every vertex v_{kl} is joined to u by an edge. Two distinct vertices v_{kl} and v_{hj} in the same cell are joined by an edge precisely when $k = h$ or $l = j$. Consequently, the union of the cell vertices has V_N elements and the union of their edges has L_N elements as in Theorem A.

Proof. Write $z_i = a_i - 1/N$. In Q_r the first r entries are nonnegative, the remaining entries are nonpositive, and z is decreasing with sum zero. Parameterize the nonnegative part by its decreasing-step coefficients $\alpha_k \geq 0$ ($k = 1, \dots, r$), so $z_i = \sum_{k \geq i} \alpha_k$. Parameterize the negative part by coefficients $\beta_l \geq 0$ ($l = r, \dots, N-1$), where $-z_i = \sum_{l < i} \beta_l$. The two remaining conditions are

$$\sum_{k=1}^r k\alpha_k = \sum_{l=r}^{N-1} (N-l)\beta_l, \quad \sum_{l=r}^{N-1} \beta_l \leq 1/N. \quad (7)$$

The apex $\alpha = \beta = 0$ is u . A nonzero vertex must saturate the last inequality: otherwise scaling all coefficients both up and down gives a nontrivial segment. At this base, if α and β have respectively p and q positive coefficients, the local face has dimension $p + q - 2$; the two equalities are independent. Thus a vertex has $p = q = 1$, giving $\beta_l = 1/N$ and $\alpha_k = (N-l)/(Nk)$, exactly (6). The same face count proves the stated base-edge rule. The rays from the apex to base vertices are all remaining edges. This also proves the cell is combinatorially a pyramid over the product of two simplexes.

Across the cells, all pairs $1 \leq k \leq l < N$ occur. There are $N(N-1)/2$, plus u . The edges from u contribute the same number. Pairs with fixed k contribute $\binom{N}{3}$ and pairs with fixed l contribute another $\binom{N}{3}$; no pair is counted in both. This yields L_N . QED.

Lemma 2 (the common finite set). Let V be the union of vertices and L the union of edges in Lemma 1. Start W_N with V times V . For every v in V and every edge $[x, y]$ in L such that $q(v)$ lies strictly between $q(x)$ and $q(y)$, add (v, z) and (z, v) , where

$$z = (1-t)x + ty, \quad t = \frac{q(v) - q(x)}{q(y) - q(x)}. \quad (8)$$

The resulting set is precisely the union of vertices of all Q_r times Q_s cut by $q(a) \leq q(b)$ or $q(a) \geq q(b)$. In particular it obeys (2).

Proof. The two distance orientations differ by $q(a) - q(b)$. A hyperplane cut of a polytope adds only intersections with its edges; a vertex of the cut lying in the interior of a face of dimension at least two would remain on a segment. Every edge of a product fixes a vertex in one factor and follows an edge in the other. Formula (8) lists exactly these transverse intersections. An edge lying in the cutting hyperplane introduces no new vertex. Duplicate pairs are removed. Rationality follows from (6), (5), and (8). QED.

4. Proof of the general optimal-constant reduction

On each of the cut product cells, D^* is affine and K_d is convex. Let γ be the minimum on the right side of (1). The only points with $D^* = 0$ are (e, u) and (u, e) ; their costs equal H_N by the one-spike formula in [1]. Therefore

$$K_d(v) + \gamma D^*(v) \leq H_N$$

holds at every vertex of every cut cell. Convexity bounds the same expression throughout each cell, and their union is the whole compact domain. This proves the proposed coefficient everywhere.

Conversely, choose a minimizing vertex (a,b) with $D^* > 0$. The lists

$$a_\varepsilon = (1 - \varepsilon)a + \varepsilon u, \quad b_\varepsilon = (1 - \varepsilon)b + \varepsilon u$$

belong to the exact stratum for $0 < \varepsilon < 1$ and converge to the vertex. Continuity from Section 2 shows their deficit/distance ratios converge to the vertex ratio. No larger coefficient can hold on that stratum. The positivity of $\mathcal{V}$ also follows from the positive balanced stability bound in [1]. Every vertex and every LP value is rational, proving Theorem A.

This is a finite evaluation procedure for each specified N,d . It does not supply a simple closed formula for all N , a polynomial-time algorithm in d , or constancy with respect to added zeros. Theorem B establishes the latter property for $N = 4$ by a separate lower bound.

5. The rational certificate for $N = 4$

Here the seven one-factor vertices are

$$e_1 = (1, 0, 0, 0), \quad e_2 = (1/2, 1/2, 0, 0), \quad e_3 = (1/3, 1/3, 1/3, 0), \quad u = (1/4, 1/4, 1/4, 1/4),$$

$$v_{12} = (3/4, 1/4, 0, 0), \quad v_{13} = (1/2, 1/4, 1/4, 0), \quad v_{23} = (3/8, 3/8, 1/4, 0).$$

The three cells have respectively (4,6), (5,8), and (4,6) vertices and edges. Their union has 14 edges. Lemma 2 gives exactly 89 ordered pairs, or 48 pairs up to exchanging a and b . The file `witnesses_N4.json` supplies one rational common spectrum for each representative. The verifier reconstructs all cell vertices by exact active-constraint systems and all edges by their face ranks; it does not trust the list of proposed pairs.

For each of the 89 ordered pairs, it checks mass, ordering, nonnegativity, all inequalities (4) in dimension eight, and

$$\sum_i s_i + \frac{4}{7}D^*(a, b) \leq \frac{5}{2}. \tag{9}$$

The recursive Horn generator finds

$$36, 462, 2120, 3516, 2120, 462, 36$$

triples at subset sizes one through seven: 8,752 in total. Thus 778,928 rational Horn inequalities are checked. Each row is scaled to integers by the exact least common multiple of its denominators; no tolerance, floating-point rank, or numerical optimizer enters the verification. The proposed spectra originated in exploratory LPs, whose success is not used as evidence without these checks.

Horn sufficiency supplies $\kappa_8 \leq \sum_i s_i$ at every vertex. Appending zero blocks supplies that upper bound for every $d \geq 8$. The convex-cell argument used in Section 4 now proves $\delta_d \geq \frac{4}{7}D^*$ throughout the compact domain, with $H = 5/2$ fixed on its boundary.

6. A dimension-independent obstruction and sharpness

Lemma 3. If F has at least four positive eigenvalues $a_1 \geq \dots$ and at least two negative magnitudes $b_1 \geq b_2 \geq \dots$, where these lists denote the actual eigenvalues of F before division by P , then

$$\kappa_d(F) \geq 3b_2 - a_2 + a_4. \quad (10)$$

Equivalently, for the normalized lists used elsewhere in the paper, $\kappa_d(F)/P \geq 3b_2 - a_2 + a_4$.

Proof. For any realizing common spectrum s , apply (4) to $S - R = -F$ and the size-two triple

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

It is admissible for every $d \geq 4$. The index-sum equality is $d + 4 = d + 1 + 3$, and its three size-one recursive conditions reduce to $2 \leq 3$, $d \leq d$, and $4 \leq d$. The resulting inequality is

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

Also $F \leq R$ and $-F \leq S$ imply $s_2 \geq b_2$ and $s_4 \geq a_4$ by the min-max principle. Summing gives

$$\sum_i s_i \geq s_1 + s_2 + s_3 + s_4 \geq 3b_2 - a_2 + a_4 + s_d \geq 3b_2 - a_2 + a_4.$$

This argument neither restricts the number of added zeros nor assumes $s_d=0$. QED.

Use the boundary pair

$$a = (5/8, 1/8, 1/8, 1/8), \quad b = (1/2, 1/2, 0, 0). \quad (11)$$

Lemma 3 gives $\kappa_d \geq 3/2$. Its rational upper witness is

$$s = (3/4, 1/2, 1/8, 1/8, 0, 0, 0, 0),$$

which has trace $3/2$ and appears in the verified certificate. Therefore $\kappa_d = 3/2$ for every $d \geq 8$. Both orientations of D^* equal $7/4$, so the ratio is exactly $4/7$.

For a sequence inside the exact $(4,4)$ stratum, retain a and put

$$b_\eta = (1/2 - \eta, 1/2 - \eta, \eta, \eta), \quad 0 < \eta < 1/8.$$

The two distances are $7/4 - 4\eta$ and $7/4 + 2\eta$. Lemma 3 and the bound already proved yield the explicit sandwich

$$\frac{3}{2} - 3\eta \leq \kappa_d(a, -b_\eta) \leq \frac{3}{2} + \frac{16}{7}\eta. \quad (12)$$

Hence δ_d/D^* tends to $4/7$ for every fixed $d \geq 8$. The left coefficient in (3) is optimal. The right coefficient $3/2$ and its sharp family are Theorem 3 and Section 5.4 of [1]. This completes Theorem B.

7. All equality pairs and strictness on the exact stratum

Write $A = (5/8, 1/8, 1/8, 1/8)$ and $B = (1/2, 1/2, 0, 0)$.

Theorem C (complete contact set). For every fixed $d \geq 8$, equality in the forward bound on $S_4 \times S_4$ holds precisely at

$$\delta_d(a, b) = \frac{4}{7}D^*(a, b) \iff (a, b) \in \{(e, u), (u, e), (A, B), (B, A)\}. \quad (13)$$

Consequently the forward inequality is strict at every target of exact inertia $(4, 4, d - 8)$. The best coefficient is an infimum approached at the boundary, as in (12).

Proof. Define the nonnegative gap

$$g_d(a, b) = \frac{5}{2} - \kappa_d(a, 0, -b) - \frac{4}{7}D^*(a, b).$$

It is concave on every cut product cell of Lemma 2. For dimension eight the additional rational witnesses in `optimal_constant_N4_d8.json` give a strictly positive lower bound for g_8 at every one of the 89 vertices except the four displayed pairs. Their zero gap follows from the one-spike formula and Section 6. This assertion uses only feasible upper witnesses and the two exact boundary costs; a solver's declaration of optimality is unnecessary. The supplied contact verifier checks these slacks exactly.

If a point in a cell has zero gap, express it as a convex combination of that cell's vertices. Concavity and nonnegativity force every vertex with positive weight to have zero gap. It remains to inspect the convex hulls of the contact vertices in each cell. Here e and A belong only to Q_1 , B only to Q_2 , and u belongs to all three one-factor cells. Moreover $q(e) = -3/2$, $q(u) = 3/2$, and $q(A) = q(B) = 0$. The only cells containing two contact vertices therefore give the segment from (e, u) to (A, B) and its factor exchange. All other nonempty contact hulls are singletons. The exact cell-membership check is included in the verifier.

At the midpoint of the first segment the lists and a feasible common spectrum are

$$\begin{aligned} a &= (13/16, 1/16, 1/16, 1/16), & b &= (3/8, 3/8, 1/8, 1/8), \\ s &= (13/16, 7/16, 1/8, 1/8, 1/16, 0, 0, 0). \end{aligned} \quad (14)$$

All 8,752 Horn inequalities hold for (14). Since its trace is $25/16$ and $D^* = 7/8$, we obtain $g_8 \geq 7/16$ there. Concavity along the segment, parameterized by $0 \leq t \leq 1$, now yields

$$g_8(t) \geq \frac{7}{8} \min \{t, 1 - t\} > 0 \quad (0 < t < 1). \quad (15)$$

Thus neither segment contains any further contact point. Exchanging the factors leaves the cost and distance unchanged. This proves (13) in dimension eight. In larger dimensions zero padding of a realizing factor gives $\kappa_d \leq \kappa_8$, hence $g_d \geq g_8$. The four contact costs remain exact in all dimensions by the same one-spike formula and Lemma 3. Therefore their set is unchanged for every fixed $d \geq 8$. QED.

8. Reproduction, scope, and next question

The accompanying replay directory is self-contained. Run `python verify_balanced_four.py` for the original 89-vertex bound, `python independent_lr_geometry_review.py` for the separate Littlewood--Richardson and geometry reconstruction, and `python algorithm_certificate_replay.py` for the 48 rational primal--dual optimal-value certificates. The latter supports the tabulated computation but is not needed to prove the upper bound. Run `python verify_contacts.py` for Theorem C: it checks 787,680 exact Horn inequalities at the 89 vertices and the additional midpoint, the vertex slacks, all cut-cell contact sets, and (14). These four replays use only the Python standard library. The separate `python certified_constant.py` proposes solutions using NumPy and SciPy and accepts them only after exact rational primal--dual checks; `requirements.txt` fixes the preparation environment. The general geometry is in `geometry.py`. The scientific theorems also rely on the written convexity, edge-cut, concavity and ambient-dimension arguments; the replays are not formal proof-assistant verification.

The actual ordered-pair counts for $N = 2, \dots, 8$ are 4,22,89,267,628,1354,2611. These are exact geometry checks, not evaluated optimal constants in the unevaluated dimensions. In particular the $N \geq 5$ constants remain open as explicit values in this work. There is no assertion about a general interior matrix-cost formula or about every optimizer's rank.

The polynomial vertex reduction and the coefficient $4/7$ improve the inspected earlier work. The nearest external tools are the classical Horn description, polyhedral projection, and elementary convex maximization on cells. Searches by self-commutator stability, inverse Hilbert-Schmidt cost, balanced inertia, and equivalent Horn formulations did not identify an external source giving these specific statements. The search was bounded; absence of a match is not a proof of bibliographic priority.

The next decisive question is whether the extremizing pairs selected by (1) admit a uniform description for $N \geq 5$, and whether their costs can be certified by bounded-size Horn families independent of the ambient dimension.

References

- [1] Lluís Eriksson. *Sharp inertia ceilings and optimal stability for inverse self-commutators*. ARR-2026-24M24KDPZK8HDBQ9, v1, 5 September 2026. Theorems 1-4 and Sections 2, 5.4, 5.5. <https://arr-research.github.io/papers/ARR-2026-24M24KDPZK8HDBQ9/> . Canonical PDF SHA-256: 15b54435c01a4972af77db0277c7f99fc21afb43be3bbc56b1884b76bb53bf30.
- [2] William Fulton. *Eigenvalues, invariant factors, highest weights, and Schubert calculus*. Bulletin of the American Mathematical Society 37 (2000), 209-249. Theorem 1 and equations (8)-(10). <https://arxiv.org/abs/math/9908012> . Supplies the recursive Horn indexing and spectral-sum sufficiency.
- [3] Allen Knutson and Terence 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. <https://arxiv.org/abs/math/9807160> . Classical sufficiency machinery, not a new result of this manuscript.
- [4] Omer Angel and Gideon Schechtman. *The Hilbert Schmidt version of the commutator theorem for zero trace matrices*. <https://arxiv.org/abs/1503.07980> . This nearby problem uses general commutator factors and a mixed operator/Hilbert-Schmidt estimate; it does not identify the self-commutator objective here.

Lluís Eriksson directed the research. OpenAI Codex assisted with proof development, exact computation, source comparison, internal review and manuscript preparation. Internal AI checks are not independent human refereeing or formal verification.