UPPER AND LOWER BOUNDS FOR THE UNIVERSAL ASYMPTOTIC RANK EXPONENT

JOSH ALMAN, MATEUSZ MICHAŁEK, AND VIRGINIA VASSILEVSKA WILLIAMS

Abstract. Asymptotic rank is the exponential growth rate of ranks of tensor powers. Over a field k, let α_k be the least exponent a such that every tensor in $(k^n)^{\otimes 3}$ has asymptotic rank at most n^a, for every n. We prove

$$\frac{2}{\sqrt{3}} \leq \alpha_k \leq \frac{10161}{8503} < 1.195$$

over every field. The lower bound is an exact calculation on an equilateral triangle. For the upper bound, we minimize a logarithmic tensor invariant over radial placements, then let two radii readjust as the third changes. A resulting endpoint inequality is compared with Euclidean lower bounds that retain the joint inner-product data. An exact finite verification completes the proof.

hexagon:2611.01211v1

1. Introduction

Fix a field k. The *rank* $\mathrm{R}_k(T)$ of a three-tensor $T \in U \otimes V \otimes W$, with finite-dimensional factors, is the least number of decomposable tensors $u \otimes v \otimes w$ whose sum is T; put $\mathrm{R}_k(0) = 0$. The product $S \boxtimes T$ groups corresponding factors of the ordinary tensor product. Submultiplicativity of rank gives the limit

$$\widetilde{\mathrm{R}}_k(T) = \lim_{m\to\infty} \mathrm{R}_k(T^{\boxtimes m})^{1/m} = \inf_{m\geq 1} \mathrm{R}_k(T^{\boxtimes m})^{1/m}, \tag{1.1}$$

called the *asymptotic rank*. It measures the exponential growth of the ranks of tensor powers.

The *universal exponent* is

$$\alpha_k = \sup_{n\geq 2} \sup_{0\neq T\in(k^n)^{\otimes 3}} \log_n \widetilde{\mathrm{R}}_k(T). \tag{1.2}$$

Thus α_k is the least a such that $\widetilde{\mathrm{R}}_k(T) \leq n^a$ in every $n \times n \times n$ tensor space, or *cubic format*. Unless otherwise specified, log denotes $\log_2$.

Theorem 1.1. *For every field k,*

$$\frac{2}{\sqrt{3}} \leq \alpha_k \leq \beta_* := \frac{10161}{8503}. \tag{1.3}$$

More explicitly:

1. *For every integer $n \geq 1$ and every $T \in (k^n)^{\otimes 3}$, one has $\widetilde{\mathrm{R}}_k(T) \leq n^{\beta_*}$.*
2. *For every $\varepsilon > 0$ and every sufficiently large integer n, there exists $T \in (k^n)^{\otimes 3}$ such that $\widetilde{\mathrm{R}}_k(T) > n^{2/\sqrt{3}-\varepsilon}$.*

The endpoints are $1.154700538\ldots$ and $\beta_* \approx 1.1949900035$. Equivalently, the upper bound gives, for each fixed $T \in (k^n)^{\otimes 3}$ and every $\varepsilon > 0$,

$$\mathrm{R}_k(T^{\boxtimes m}) \leq n^{(\beta_*+\varepsilon)m}$$

for all sufficiently large m.

The elementary range is $1 \leq \alpha_k \leq 2$. Expansion in bases of two factors gives rank at most n^2. If $e_1, \ldots, e_n$ is the standard basis of k^n, the diagonal $\Delta_3(n) = \sum_{j=1}^n e_j \otimes e_j \otimes e_j$ has asymptotic rank n: grouping tensor indices as matrix rows and columns, a *flattening*, gives matrix rank n^m on its mth power.

Date: 26 November 2026.

The four main ideas. *1. An actual minimizing extension.* Strassen's theorem expresses asymptotic rank as the maximum of *spectral points*: tensor invariants ϕ that are multiplicative on grouped products and monotone under linear maps on the factors. Their full definition is in Section 2. The identity tensor $E_{ab}(q) = \sum_{j=1}^{q} e_j \otimes e_j$ between positions a, b, with one-dimensional factors elsewhere, defines a distance $d(a,b) = \log_2 \phi(E_{ab}(2))$, after identifying positions at distance zero. An exact extension adds an auxiliary position p while preserving all old tensor values. Let a_1, a_2, a_3 be the original positions. In the product of the identities $E_{pa_i}(n)$, apply at p the linear functional whose basis values are the coefficients of T. This gives

$$\phi(T) \leq n^{d(p,a_1)+d(p,a_2)+d(p,a_3)}.$$

The exact-extension theorem supplies a position attaining the smallest possible exponent.

2. A radial minimum and an optimized neighboring profile. Near that minimizing position, the three original factors determine three rays. We place a binary diagonal—four indices sharing a label with two possible values—at their origin and at one point on each ray. Minimize its logarithmic value with the three radial lengths summing to one. Polynomial multiplication proves differentiability at this minimum and bounds the variation of its derivatives. Allowing two radii to readjust as the third changes gives a two-variable minimum, the *neighboring profile*. A reciprocal convexity inequality supplies an explicit upper bound for its value at a chosen endpoint.

3. Compatible Euclidean comparisons. The geometric input identifies Euclidean diagonal values with *mean width*: the average length of projections, normalized to equal length on a segment. Our lower bounds keep all pairwise inner products, the *Gram matrix*, together. A concavity argument then reduces an angular minimum to four vertices of one polygon. Estimates referring to a common comparison position are imposed together at that position. The origin remains in every configuration; it can increase mean width beyond half the perimeter of the other three points.

4. An exact finite cover. The remaining real parameters are covered by *boxes*, products of closed intervals. Each box is excluded by an inequality valid throughout it or subdivided further. Integer arithmetic verifies all numerical enclosures and the complete cover [25]. The final step uses the proved bound $\alpha_k \leq 239/200$ to bound the radial invariant from below; Section 11 records this dependency and the earlier certificates it requires.

Inputs and related work. The extended asymptotic rank conjecture predicted $\alpha_{\mathbb{C}} = 1$ [5, Problem 15.5]; see also [14]. Our lower bound contradicts it, using OpenAI's Euclidean spectral values [12] and an equilateral triangle.

For infinite k, Kaski and Michałek [10, Theorem 6] constructed an explicit sequence of zero-one tensors whose normalized asymptotic-rank exponents tend to α_k. Thus the worst exponent over all tensors is captured by one explicit family.

For comparison, write $e_{ij} = e_i \otimes e_j$. The matrix multiplication tensor is

$$\mathrm{MM}_q = \sum_{i,j,\ell=1}^{q} e_{ij} \otimes e_{j\ell} \otimes e_{\ell i},$$

and its exponent ω_k is the infimum of w for which $\mathrm{R}_k(\mathrm{MM}_q) = O(q^w)$. Strassen's general bound is $\alpha_k \leq 2\omega_k/3$; see [10, Lemma 25]. Alman and Li [1, Theorem 5] improved the corresponding bound $n^{2\omega_k/3}$ for every fixed $n \geq 3$. Their saving in the exponent tends to zero as n grows; they asked for an improvement uniform in n [1, Problem 41].

Independently, Liu [11, Theorem 1.1] proved, over every field,

$$\alpha_k \leq \frac{\omega_k}{\sqrt{3}} \leq \frac{3\sqrt{3}}{4}.$$

The second inequality uses $\omega_k \leq 9/4$ [3, Theorem 1.1]. We also use the quadratic constraint on spectral points in [3, Proposition 4.14].

Anthropic's midpoint inequality controls distances to an intermediate factor position [8, Theorem C.1]. OpenAI's exact extensions preserve all old tensor values, and its Euclidean theorem identifies diagonal values with mean width [12]. Section 3 states these inputs and explains their meaning. Coefficient tensors of polynomial multiplication provide radial regularity and curvature [17, 18]. The upper proof works first over an algebraically closed field; invariance under algebraic extension then gives every field.

A short history of the swarm bounds. The swarm recorded an exponent below 1.38585 on 22 November 2026, $\omega_k/\sqrt{3}$ on 24 November, and 1.1969 on 25 November [15, 21]. On 26 November, the completed bounds $2391/2000 = 1.1955$ and $239/200 = 1.195$ preceded the present $\frac{10161}{8503}$ [23, 24, 25]. Section 11 lists the bounds needed as inputs.

AI assistance. The ideas and this manuscript were developed with the assistance of Astra.

2. Spectral points and identity tensors

A tensor on a finite label set J is an element of $\bigotimes_{i\in J} V_i$. A factor $V_i = k$ is a *scalar factor*.

Definition 2.1. A tensor T is a *restriction* of S, written $T \le S$, if linear maps on the individual factors take S to T. The *direct sum* $S \oplus T$ places the arrays in separate diagonal blocks, in $V_i \oplus W_i$ at every position i. The product $S \boxtimes T$ groups corresponding factors.

Local changes of basis and *padding* by zero coordinates preserve rank and asymptotic rank: both give restrictions in both directions, using inclusion and coordinate projection for padding.

Definition 2.2. A *spectral point* on J, where $|J| \ge 2$, is a function ϕ from tensors on J to $\mathbb{R}_{\ge 0}$ satisfying

$$\phi(0) = 0, \qquad \phi(1) = 1,$$
$$\phi(S \oplus T) = \phi(S) + \phi(T), \qquad \phi(S \boxtimes T) = \phi(S)\phi(T),$$

and $\phi(T) \le \phi(S)$ whenever $T \le S$. Here 1 is any nonzero decomposable tensor, up to mutual restriction.

Matrix rank after flattening across a fixed partition of J into two nonempty parts is an example. For three positions, Alman, Li and Pratt [2, Theorem 1.1] constructed explicit spectral points beyond these flattening ranks over every field. The full diagonal

$$\Delta_J(q) = \sum_{a=1}^{q} \bigotimes_{i\in J} e_a$$

is the direct sum of q copies of 1, so $\phi(\Delta_J(q)) = q$. For $T \ne 0$, the restrictions $1 \le T \le \Delta_J(\mathrm{R}_k(T))$, applied also to powers, give

$$1 \le \phi(T) \le \widetilde{\mathrm{R}}_k(T) \le \mathrm{R}_k(T).$$

Strassen proved the converse after maximizing over spectral points.

Theorem 2.3 (Strassen's spectral characterization [13])**.** *For every tensor on a finite label set with at least two elements,*

$$\widetilde{\mathrm{R}}_k(T) = \max_{\phi} \phi(T),$$

where ϕ ranges over the spectral points on that set.

Thus an upper bound must hold for every spectral point; a lower bound needs only one.

Identity tensors define distances. For nonempty $S \subseteq J$, let $\Delta_S(q)$ share one q-valued diagonal index among S, with scalar factors elsewhere. For distinct positions a, b, write

$$E_{ab}(q) = \Delta_{\{a,b\}}(q) = \sum_{j=1}^{q} e_j \otimes e_j.$$

Contraction with $E_{ab}(q)$ transfers a q-dimensional factor between a and b without changing its coefficients. Put

$$H(T) = \log_2 \phi(T), \qquad \delta(S) = H(\Delta_S(2)), \qquad d(a,b) = H(E_{ab}(2)),$$

for $T \neq 0$, and set $d(a,a) = 0$. Multiplicativity makes H additive on grouped products. A full binary diagonal has value one; a singleton diagonal is a nonzero vector and has value zero.

Lemma 2.4. *The function d is a pseudometric: it is symmetric, nonnegative, zero on the diagonal, and satisfies the triangle inequality. For every $q \geq 1$ and nonempty S,*

$$H(\Delta_S(q)) = \delta(S) \log_2 q. \tag{2.1}$$

Writing $\operatorname{diam}(S) = \max_{a,b \in S} d(a,b)$ *for its diameter, for any tree whose vertices lie in J and include S,*

$$\operatorname{diam}(S) \leq \delta(S) \leq \sum_{ab \text{ an edge of the tree}} d(a,b). \tag{2.2}$$

Proof. Contraction at the common position gives $E_{ac}(2) \leq E_{ab}(2) \boxtimes E_{bc}(2)$, proving the triangle inequality; symmetry and nonnegativity are immediate. Tensoring diagonals multiplies dimensions, while projection reduces them. Comparing powers of q with consecutive powers of two and passing to the limit proves (2.1).

A diagonal restricts to an identity tensor between any two of its positions, giving the lower tree bound. For the upper bound, put a binary identity on each edge. At every vertex, kill unequal incident labels and identify equal ones. Retain this label at vertices of S and contract it elsewhere. The result is $\Delta_S(2)$, so monotonicity and multiplicativity give the bound. □

Forgetting a diagonal factor shows that δ is monotone in its set of positions. In particular, $0 \leq d(a,b) \leq 1$ for a spectral point on three labels.

Placements and normalization. A *placement* $x = (x_1, \ldots, x_r)$ assigns factor i of $T \in \bigotimes_{i=1}^{r} V_i$ to position x_i. Factors at the same position are grouped by tensor product; unoccupied positions carry k. Its *support* is $\{x_1, \ldots, x_r\}$. Write $H(T; x)$ for the logarithmic value at this placement. Moving each factor, of dimension $q_i = \dim V_i$, by an identity tensor in either direction gives the *transport estimate*

$$|H(T;x) - H(T;y)| \leq \sum_i d(x_i, y_i) \log_2 q_i. \tag{2.3}$$

Thus positions at distance zero can be identified without changing values. A nonzero tensor placed at one position is a vector and has value zero.

For example, flattening across $a \mid \{b, c\}$ gives $d(a,b) = d(a,c) = 1$ and $d(b,c) = 0$: the identity between b and c lies on one side of the flattening and has matrix rank one.

For $S \subseteq J$, write $\oplus_S$ for direct sum on S and $\iota_* T$ for insertion of scalar factors outside S. The tensors $\iota_*(U \oplus_S V)$ and $\iota_* U \oplus_J \iota_* V$ differ: the latter also has a two-dimensional direct-sum factor at each position outside S. Consequently restriction to fewer positions requires normalization.

Lemma 2.5 (Normalization on a subset of positions)**.** *Let ϕ be a spectral point on J and $\varnothing \neq S \subseteq J$. Put*

$$G(T) = \phi(\iota_* T), \qquad e = \log_2 G(\Delta_S(2)) = \delta(S).$$

If $e > 0$, then $\phi_S(T) := G(T)^{1/e}$ is a spectral point on S, and

$$H(\iota_* T) = e \log_2 \phi_S(T) \qquad (T \neq 0). \tag{2.4}$$

If $e = 0$, then $G(T) = 1$ and $H(\iota_* T) = 0$ for every nonzero tensor on S.

For instance, $G(1 \oplus_S 1) = 2^e$, whereas $G(1) + G(1) = 2$. The power $1/e$ restores additivity, as follows.

Proof. Scalar insertion preserves restrictions and products, so G is monotone and multiplicative, with $G(\Delta_S(q)) = q^e$. For $e > 0$ and $U, V \neq 0$, put

$$A = G(U)^{1/e}, \qquad B = G(V)^{1/e}, \qquad C = G(U \oplus_S V).$$

Group the expansion of $(U \oplus_S V)^{\boxtimes m}$ according to the number j of copies of U. Up to local coordinates, the jth group is

$$\Delta_S\left(\binom{m}{j}\right) \boxtimes U^{\boxtimes j} \boxtimes V^{\boxtimes(m-j)},$$

with zeroth powers interpreted as 1. Its value is $\binom{m}{j}^e A^{je} B^{(m-j)e}$. Each group restricts from the full tensor, which in turn restricts from the direct sum of these groups formed on J. Additivity of ϕ on J gives, with $M_m = \max_{0 \le j \le m} \binom{m}{j} A^j B^{m-j}$,

$$M_m^e \le C^m \le (m+1) M_m^e.$$

The binomial theorem gives $(A+B)^m/(m+1) \le M_m \le (A+B)^m$. Taking mth roots and letting $m \to \infty$ yields $C = (A+B)^e$, proving additivity. Zero summands cause no difficulty; the other spectral axioms follow from those of G.

If $e = 0$, the restrictions $1 \le T \le \Delta_S(\mathrm{R}_k(T))$ give $1 \le G(T) \le \mathrm{R}_k(T)^0 = 1$ for $T \neq 0$. □

Definition 2.6. A *spectral space* over k consists of a metric space (X, d) with logarithmic values $H(T; x)$ at finite placements. Grouping factors at one position and inserting scalar factors do not change values, and $H(E_{ab}(2)) = d(a, b)$. For every finite $S \subseteq X$ of positive diameter, the number $\delta(S) = H(\Delta_S(2); S)$ is positive and there is a spectral point ϕ_S such that

$$H(T; x) = \delta(S) \log_2 \phi_S(T) \qquad (T \neq 0)$$

whenever x is supported in S, with T viewed as a tensor on S by grouping factors and inserting scalar factors. Nonzero tensors at a single point have value zero. Set $H(0; x) = -\infty$; all finite logarithmic estimates concern nonzero tensors.

A spectral point gives such a space after identifying positions at distance zero. The tree bounds (2.2) hold in every spectral space. On finite sets $S \subseteq J$ of positive diameter, compatibility reads

$$\delta(J) \log_2 \phi_J(\iota_* T) = \delta(S) \log_2 \phi_S(T) \qquad (T \neq 0). \tag{2.5}$$

The scale factors give a common logarithmic value even though the normalized spectral points may differ. Multiplying d and H by the same positive number rescales δ and leaves every ϕ_S unchanged.

Passing to an algebraic closure. The analytic upper proof uses algebraically closed fields. The following descent argument gives the result over every field.

Lemma 2.7 (Invariance under algebraic extension)**.** *If E/k is an algebraic field extension and T is a three-tensor defined over k, then $\widetilde{\mathrm{R}}_E(T) = \widetilde{\mathrm{R}}_k(T)$.*

Proof. A rank decomposition over k remains one over E, giving one inequality. Conversely, fix $m \ge 1$ and an r-term decomposition of $T^{\boxtimes m}$ over E. Its coefficients lie in a finite extension K/k, of degree d. Choose a k-linear map $\lambda : K \to k$ with $\lambda(1) = 1$ and a basis $b_1, \dots, b_d$ of K. The form $\lambda(xyz)$ is a sum of d^2 products of linear forms, since

$$\lambda(xyz) = \sum_{i,j=1}^{d} x_i y_j \lambda(b_i b_j z), \qquad x = \sum_i x_i b_i, \quad y = \sum_j y_j b_j.$$

Raise the decomposition to its sth power and apply λ to its coefficients. The left side is unchanged; each of the r^s decomposable K-tensors becomes a sum of at most d^2 decomposable k-tensors. Hence

$$\mathrm{R}_k(T^{\boxtimes ms}) \le d^2 r^s.$$

Letting $s \to \infty$ with m, K fixed gives $\widetilde{\mathrm{R}}_k(T) \le r^{1/m}$. Minimize over decompositions and let $m \to \infty$. □

3. Spectral extensions and their geometry

An *isometric embedding* preserves every pairwise distance. An *exact extension* of a spectral space adds positions while preserving $H(T; x)$ for every tensor at every old placement.

Lemma 3.1 (Restriction through an auxiliary factor)**.** *Let $x_1, \dots, x_r, p$ be positions in a spectral space. For every nonzero $T \in \bigotimes_{i=1}^r k^{q_i}$,*

$$H(T; x_1, \dots, x_r) \le \sum_i d(p, x_i) \log_2 q_i. \tag{3.1}$$

The position p may coincide with some of the x_i.

Proof. Write $T = \sum t_{a_1 \dots a_r} e_{a_1} \otimes \cdots \otimes e_{a_r}$, and put $I = \{i : x_i = p\}$ and $J = \{1, \dots, r\} \setminus I$. At position p in $\boxtimes_{j \in J} E_{p x_j}(q_j)$, apply

$$\bigotimes_{j \in J} k^{q_j} \longrightarrow \bigotimes_{i \in I} k^{q_i}, \qquad \bigotimes_{j \in J} e_{a_j} \longmapsto \sum_{(a_i)_{i \in I}} t_{a_1 \dots a_r} \bigotimes_{i \in I} e_{a_i}.$$

With identity maps elsewhere, this gives T, grouping factors at equal positions. Empty tensor products mean k: when I is empty, the displayed map is the coefficient functional of T; when J is empty, it sends 1 to T. Monotonicity and additivity of H on products prove the inequality, since $d(p, x_i) = 0$ for $i \in I$. □

For fixed positions $x_1, \dots, x_r$ and nonnegative integer weights $w = (w_1, \dots, w_r)$, define the *format exponent*

$$g_H(w) = \sup_{\substack{m \ge 1,\ m \in \mathbb{Z} \\ 0 \ne T \in \bigotimes_i k^{2^{m w_i}}}} \frac{H(T; x_1, \dots, x_r)}{m}. \tag{3.2}$$

Weight zero gives a scalar factor. Write g_ϕ when $H = \log_2 \phi$. The auxiliary restriction gives $g_H(w) \le \sum_i w_i d(p, x_i)$ in every exact extension; the following theorem attains equality.

Theorem 3.2 (Exact extension [12, Theorem 3.1 and its extension to arbitrary spectral spaces])**.** *For a spectral space X, positions $x_1, \dots, x_r \in X$, and nonnegative integer weights w_i, there is an exact extension with a position p such that*

$$g_H(w) = \sum_i w_i d(p, x_i). \tag{3.3}$$

For a finite spectral point ϕ, the extension can be a spectral point ϕ' satisfying $\phi'(\iota_* T) = \phi(T)$ for every old tensor T.

The auxiliary minimum is attained even if the tensor supremum is not. The result includes infinite spaces, notably the Euclidean plane: the cited proof extends larger finite sets and passes to a simultaneous limit preserving all old values. A positive weight bounds the auxiliary position's distance by (3.3); if all weights vanish, both sides are zero.

Midpoints and convexity. A *geodesic segment* $[x,y]_t$, $0 \le t \le 1$, satisfies $d([x,y]_s,[x,y]_t) = |s-t|d(x,y)$. A metric space is *complete* if Cauchy sequences converge. A set is *convex* if it contains each segment between its points; the segments in the spaces below are unique. The closed convex hull is the smallest closed convex set containing a given set.

Theorem 3.3 (Geodesic extension and curvature [12, Corollary 3.4 and Theorem 4.3])**.** *Every spectral space has an exact extension to a complete geodesic spectral space X whose distances satisfy*

$$d(z,[x,y]_t)^2 \le (1-t)d(z,x)^2 + td(z,y)^2 - t(1-t)d(x,y)^2. \tag{3.4}$$

For each fixed tensor, logarithmic values are jointly convex:

$$H(T;[x_1,y_1]_t,\dots,[x_r,y_r]_t) \le (1-t)H(T;x_1,\dots,x_r) + tH(T;y_1,\dots,y_r). \tag{3.5}$$

A complete geodesic space satisfying (3.4) is called *CAT(0)*, a metric form of nonpositive curvature. In Euclidean space this inequality is an identity. In general it bounds distances to intermediate points by the Euclidean comparison. At $t = 1/2$ it is the squared midpoint inequality proved in the spectral setting by Claude (Anthropic) [8, Theorem C.1]. For example, if three pairwise distances equal one, a side's midpoint is at distance at most $\sqrt{3}/2$ from the third point. The second inequality says that moving all tensor factors together along their segments gives at most the weighted average of the two endpoint values.

Lemma 3.4 (Projection onto a closed convex set)**.** *In a complete CAT(0) space X, every point p has a unique nearest point q in each nonempty closed convex subset Y. For every $v \in Y$,*

$$d(p,v)^2 \ge d(p,q)^2 + d(q,v)^2. \tag{3.6}$$

Thus replacing p by q does not increase distances to points of Y.

Proof. Put $a = \inf_{y\in Y} d(p,y)$ and choose $y_n \in Y$ with $d(p,y_n) \to a$. The midpoint inequality gives

$$d(y_n,y_m)^2 \le 2d(p,y_n)^2 + 2d(p,y_m)^2 - 4a^2 \longrightarrow 0.$$

Hence y_n converges to a nearest point q; the same inequality gives uniqueness. For $v \in Y$, minimality and (3.4) give

$$d(p,q)^2 \le d(p,[q,v]_t)^2 \le (1-t)d(p,q)^2 + td(p,v)^2 - t(1-t)d(q,v)^2.$$

Rearrange, divide by $t > 0$, and let $t \downarrow 0$. □

Lemma 3.5 (The format exponent in a complete space)**.** *For positions $x_1,\dots,x_r$ in a complete CAT(0) spectral space X and nonnegative integer weights w_i,*

$$g_H(w) = \min_{y\in X} \sum_i w_i d(y,x_i).$$

Proof. The auxiliary restriction bounds $g_H(w)$ by every displayed sum. Theorem 3.2 extends the whole space by a point attaining equality. Complete this extension using Theorem 3.3. The old copy of X is closed by completeness and convex by uniqueness of geodesics: its old geodesics remain geodesics in the extension. Project the new point to X by Lemma 3.4. No relevant distance increases, so the projected point still attains equality. □

Euclidean mean width. Write $\langle x,y\rangle = x\cdot y$ for the Euclidean inner product and $|x| = \|x\| = \sqrt{\langle x,x\rangle}$ for length. For finite nonempty $A \subset \mathbb{R}^d$, $d \ge 1$, and a unit vector u, the *width in direction u* is its projection length

$$\max_{a\in A}\langle u,a\rangle - \min_{a\in A}\langle u,a\rangle.$$

The *normalized mean width* $W(A)$ is its average over directions, divided by the corresponding average for a unit segment. Segments thus have their ordinary lengths, independently of the ambient dimension.

Let g be a standard Gaussian vector, with independent real normal coordinates of mean zero and variance one, and let $\mathbb{E}$ denote expectation. Equivalently,

$$W(A) = \sqrt{2\pi}\,\mathbb{E}\max_{a\in A}\langle g, a\rangle. \tag{3.7}$$

The Gaussian direction is uniform and independent of length; symmetry turns the expected maximum into half the expected width, and $\mathbb{E}|g_1| = \sqrt{2/\pi}$ gives the factor. Mean width is translation invariant and depends only on the convex hull, so we also write $W(\operatorname{conv} A)$. For a planar polygon it is half the perimeter [12, Section 5], counting a segment's perimeter as twice its length; a singleton has width zero.

Theorem 3.6 (Euclidean logarithmic spectral values [12, Theorems 1.1 and 5.3])**.** *Over every field, Euclidean space of every finite dimension admits compatible logarithmic spectral values with its usual distances and*

$$\delta(S) = W(S)$$

for each finite nonempty S. Thus on a set of positive diameter, $H(T; x) = W(S)\log_2 \phi_S(T)$ for a spectral point ϕ_S. Moreover, whenever a whole Euclidean space embeds isometrically in a spectral space, its inherited diagonal values equal $W(S)$.

The existence assertion supplies the lower bound's spectral point; the last assertion evaluates the Euclidean comparisons used for the upper bound.

4. The lower bound from an equilateral triangle

Theorem 4.1. *For every field k, one has $\alpha_k \geq 2/\sqrt{3}$. More precisely, for every $\varepsilon > 0$, there are an integer $m \geq 1$ and a tensor $T \in (k^{2^m})^{\otimes 3}$ with*

$$\widetilde{\mathrm{R}}_k(T) > (2^m)^{2/\sqrt{3}-\varepsilon}.$$

Proof. Equip the Euclidean plane with the logarithmic values from Theorem 3.6, and take an equilateral triangle A, B, C of side one. The function

$$y \longmapsto |y - A| + |y - B| + |y - C|$$

is convex and unchanged by the triangle's three rotations. Averaging the rotated images of y gives the centroid, so it minimizes this sum. Each centroid-to-vertex distance is $1/\sqrt{3}$, giving minimum $\sqrt{3}$.

The triangle's mean width is half its perimeter, $3/2$. Its normalized spectral point therefore satisfies

$$H(T; A, B, C) = \frac{3}{2}\log_2 \phi(T).$$

Lemma 3.5, with weights $(1, 1, 1)$, gives

$$\sup_{\substack{m\geq 1\\ 0\neq T\in(k^{2^m})^{\otimes 3}}} \frac{\log_2 \phi(T)}{m} = \frac{\sqrt{3}}{3/2} = \frac{2}{\sqrt{3}}. \tag{4.1}$$

Since $\phi(T) \leq \widetilde{\mathrm{R}}_k(T)$, finite tensors approaching this supremum prove the theorem. □

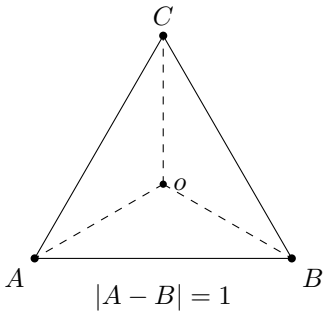

FIGURE 1. The unit equilateral triangle and its centroid.

Large formats and generic tensors.

Corollary 4.2. *Fix k and $\varepsilon > 0$. For every sufficiently large n,*

$$\{T \in (k^n)^{\otimes 3} : \widetilde{\mathrm{R}}_k(T) > n^{2/\sqrt{3}-\varepsilon}\}$$

is nonempty and Zariski open. It is dense when k is infinite; over $\mathbb{R}$ or $\mathbb{C}$, its complement has Lebesgue measure zero.

Proof. Choose a tensor T of side length $a \geq 2$ with $\log_a \widetilde{\mathrm{R}}_k(T) > 2/\sqrt{3} - \varepsilon/2$. For larger n, pad $T^{\boxtimes m}$ to side length n, where $m = \lfloor \log_a n \rfloor$. Powers multiply asymptotic rank and padding preserves it, so its exponent is

$$\frac{m \log \widetilde{\mathrm{R}}_k(T)}{\log n} \longrightarrow \log_a \widetilde{\mathrm{R}}_k(T) > 2/\sqrt{3} - \varepsilon/2.$$

This proves nonemptiness for all sufficiently large n. Asymptotic rank sublevel sets in a fixed format are Zariski closed [7, Theorem 2]. The displayed complement is therefore open; density and the measure assertion follow from the corresponding facts for a proper algebraic subset of affine space. □

5. A minimizing extension and its tangent cone

Fix a spectral point ϕ on the original positions a_1, a_2, a_3. Extend it to a complete CAT(0) spectral space. By Lemma 3.5, there is a point o attaining

$$L := g_\phi(1,1,1) = \min_x \sum_{i=1}^{3} d(x, a_i). \tag{5.1}$$

Such a minimizer is a *Fermat point*. Projection lets us choose o in the closed convex hull of the three positions. The auxiliary restriction gives $\phi(T) \leq n^L$ for every $T \in (k^n)^{\otimes 3}$, so we must prove $L \leq \beta_*$.

Lemma 5.1 (Adjoining a position in a convex hull)**.** *For a finite nonempty set A in a geodesic spectral space with jointly convex logarithmic values, and o in its closed convex hull, $\delta(A \cup \{o\}) = \delta(A)$. In particular, the diagonal on $\{o, a_1, a_2, a_3\}$ has logarithmic value one.*

Proof. Keep the old diagonal factors fixed and move the extra factor. Its value is convex and continuous by (3.5) and (2.3). At an old position, a repeated diagonal label can be inserted or removed by a local linear map, so the value is $\delta(A)$. It is therefore at most $\delta(A)$ on the closed convex hull. Forgetting the extra factor proves the reverse inequality. □

Angles and the tangent cone. For $x, y \neq o$, the Euclidean triangle with side lengths $d(o,x), d(o,y), d(x,y)$ defines the *comparison angle*

$$\cos \overline{\angle}_o(x,y) = \frac{d(o,x)^2 + d(o,y)^2 - d(x,y)^2}{2d(o,x)d(o,y)}.$$

As x, y approach o along fixed geodesics, these angles converge to their *angle at o*. Identify directions of angle zero and complete the angular metric space. The result is the *space of directions*, or *link*, $\Sigma = \Sigma_o X$, with metric d_Σ taking values in $[0, \pi]$.

The *tangent cone* $C = T_o X$ consists of an apex 0 and pairs (r, e), written re, with $r > 0$ and $e \in \Sigma$. Its metric is

$$d_C(re, sf)^2 = r^2 + s^2 - 2rs \cos d_\Sigma(e, f). \tag{5.2}$$

Write $|v| = d_C(0, v)$ for radial length. Scalar multiplication means radial dilation; addition is used only in identified Euclidean subspaces.

For complete CAT(0) spaces, Σ is CAT(1) and C is CAT(0) [4, II.3.18–19]. CAT(1) compares triangles of perimeter less than 2π with triangles on the unit sphere. In particular, directions at distance less than π have a unique shortest arc; at distance π an arc need not exist. The tangent-cone construction also applies to CAT(1) spaces [9, Definition 2.60 and Theorem 2.61].

Limits of logarithmic spectral values. We use limits both when magnifying near o and when moving a basepoint far along a ray. Each limit retains distances and all logarithmic tensor values simultaneously. Choose a generalized limit on bounded real sequences that agrees with ordinary limits and preserves finite sums, products and inequalities. A *nonprincipal ultrafilter* provides such a choice and ignores changes to finitely many terms.

For spaces with chosen basepoints, consider sequences at bounded distance from those basepoints. Take generalized limits of their pairwise distances and identify sequences at limiting distance zero. This is the *pointed ultralimit*. Transport bounds fixed-tensor values and makes their limits independent of the representatives. On each finite support of positive limiting diameter, the tree bounds keep the diagonal normalization positive and bounded, so the spectral axioms pass to the limit.

Lemma 5.2 (Logarithmic values on the tangent cone)**.** *The cone $C = T_o X$ carries compatible logarithmic spectral values H_o. For each nonzero tensor T, they are jointly convex along geodesics and homogeneous under dilation:*

$$H_o(T; cv_1, \dots, cv_r) = cH_o(T; v_1, \dots, v_r) \qquad (c \geq 0). \tag{5.3}$$

If v_i is represented by a geodesic γ_i starting at o with speed $|v_i|$, then

$$H_o(T; v_1, \dots, v_r) = \lim_{t \downarrow 0} \frac{H(T; \gamma_1(t), \dots, \gamma_r(t))}{t}. \tag{5.4}$$

The tensor coefficients are fixed throughout this limit. The proof is in Appendix A.1.

The diagonal bound and normalization. Put $\rho_i = d(o, a_i)$. When $\rho_i > 0$, let e_i be the direction towards a_i and put $\widehat{v}_i = \rho_i e_i \in C$; otherwise put $\widehat{v}_i = 0$. Thus $L = \rho_1 + \rho_2 + \rho_3$. Write δ_C for the binary-diagonal values of H_o. The diagonal on (o, a_1, a_2, a_3) has value one by Lemma 5.1. Contracting its positions towards o and using (3.5) and (5.4) gives

$$\delta_C(\{0, \widehat{v}_1, \widehat{v}_2, \widehat{v}_3\}) \leq 1. \tag{5.5}$$

Lemma 5.3 (Angles at a minimizing position)**.** *If o is an original position, then $L \leq 2/\sqrt{3}$. Otherwise all ρ_i are positive and*

$$\sum_{i=1}^{3} \cos d_\Sigma(b, e_i) \leq 0 \qquad (b \in \Sigma). \tag{5.6}$$

Every pair angle $\theta_{ij} = d_\Sigma(e_i, e_j)$ then belongs to $[2\pi/3, \pi]$. Moreover, $d_C(\widehat{v}_i, \widehat{v}_j) \leq d_X(a_i, a_j)$.

Proof. When all $\rho_i > 0$, moving o a distance t in a unit direction b changes its distance to $a_i \neq o$ by $-t \cos d_\Sigma(b, e_i) + o(t)$ [4, II.3.6]. Minimality gives (5.6) for geodesic directions and hence, by continuity, for every direction. If a pair angle were less than $2\pi/3$, its arc midpoint would contribute more than one from those two cosines, contradicting the inequality since the third is at least -1.

If $o = a_1$ and $\rho_2, \rho_3 > 0$, the first distance instead has derivative 1. The same bisector argument gives $d_\Sigma(e_2, e_3) \geq 2\pi/3$, and hence

$$d_C(\widehat{v}_2, \widehat{v}_3) \geq \sqrt{\rho_2^2 + \rho_3^2 + \rho_2\rho_3} \geq \frac{\sqrt{3}}{2}(\rho_2 + \rho_3).$$

Convexity of distance under contraction gives $d_C(\widehat{v}_i, \widehat{v}_j) \leq d_X(a_i, a_j) \leq 1$ for every pair, proving the asserted bound on L. If another radius vanishes, $L \leq 1$. □

The vertex case is below the target. Otherwise write $H = H_o$ and $\delta = \delta_C$, and normalize by

$$r_i = \frac{\rho_i}{L}, \qquad v_i = \frac{\widehat{v}_i}{L} = r_i e_i, \qquad \eta = \frac{1}{L}.$$

Homogeneity gives

$$r_1 + r_2 + r_3 = 1, \qquad \delta(\{0, v_1, v_2, v_3\}) \leq \eta. \tag{5.7}$$

The angles and (5.6) are unchanged. Denote this initial radial vector by $r^0 = (r_1, r_2, r_3)$. We next vary the lengths while keeping the directions fixed.

6. Radial minimization and derivative bounds

Keep the three directions fixed and vary their radii within the same spectral cone. At a minimum of the normalized diagonal value, the three radial derivatives will agree. Polynomial multiplication supplies the regularity and derivative bounds needed to prove this.

Take k algebraically closed; Section 2 supplies the field reduction. Radii and derivatives are real, while tensor coefficients and restrictions are over k.

One diagonal function on the three rays. Fix the complete homogeneous spectral cone from Section 5, its apex 0, and the three unit directions e_1, e_2, e_3. For $r = (r_1, r_2, r_3) \in [0, \infty)^3$, define

$$F(r) = \delta(\{0, r_1 e_1, r_2 e_2, r_3 e_3\}).$$

This is the binary diagonal's logarithmic value at the four positions; repeating a position does not change it.

The function F is continuous, convex, homogeneous of degree one, coordinatewise nondecreasing, and separately 1-Lipschitz. Convexity and homogeneity follow from their placement versions. A shorter point on a ray lies between its longer point and 0; convex-hull insertion proves monotonicity. Transport of the one moving binary factor proves the Lipschitz bound.

The supported diagonal satisfies *intersecting submodularity*: for finite sets of positions A, B with $A \cap B \neq \varnothing$,

$$\delta(A) + \delta(B) \geq \delta(A \cup B) + \delta(A \cap B).$$

This is an algebraic consequence of multiplication of several polynomials, with all input factors kept separate [18, notes C989–C990]. Applied to radial placements, it gives

$$F(r) + F(s) \geq F(r \vee s) + F(r \wedge s). \tag{6.1}$$

Here $\vee, \wedge$ mean coordinatewise maximum and minimum. Indeed, freely insert the smaller radial points into both finite sets. They share 0, their union has value $F(r \vee s)$, and their intersection contains the set for $F(r \wedge s)$. We write $F_i = \partial F/\partial r_i$ and $F_{ij} = \partial^2 F/(\partial r_i \partial r_j)$ when these derivatives exist; the matrix $\nabla^2 F = (F_{ij})$ is the Hessian. In particular, wherever second derivatives exist,

$$F_{ij} \leq 0 \qquad (i \neq j).$$

The apex remains a Fermat point when the three positive radii vary: the directional derivatives of the distance sum depend only on the directions, so remain nonnegative; convexity proves minimality. Its cost is $\sum_i r_i$. Projection shows that the apex belongs to the closed convex hull of the tips; free insertion makes $F(r)$ their diagonal value as well.

Lemma 3.5 consequently identifies the format exponent of the normalized spectral point at these tips as

$$\frac{r_1 + r_2 + r_3}{F(r)}. \tag{6.2}$$

Thus any already proved universal bound $\alpha_k \le \alpha_0$ gives the pointwise estimate

$$F(r) \ge \frac{r_1 + r_2 + r_3}{\alpha_0}. \tag{6.3}$$

The tensor input: a profile of polynomial multiplication. Fix $r_i > 0$, and write $D = F(r)$. Place three polynomial inputs at $r_i e_i$ and their output at 0. For integers $L_i \ge 1$, the coefficient tensor of multiplication is

$$M(L_1, L_2, L_3) = \sum_{0 \le d_i < L_i} e_{d_1} \otimes e_{d_2} \otimes e_{d_3} \otimes e_{d_1+d_2+d_3}.$$

In the standard coefficient bases, this tensor multiplies three polynomials of degrees below L_1, L_2, L_3, retaining every output coefficient. Its three input factors remain separate.

In the formulas about this tensor, use natural logarithms: put $\mathcal{H} = (\ln 2)H$. The distance normalization remains fixed, so a q-dimensional identity between two positions has value $d(x, y) \ln q$, and a common q-label at this placement has value $D \ln q$. The continuum theorem for polynomial multiplication [17, subsection "The continuum profile"] gives a function $P = P_r$ on $x_i > 0$, defined by

$$\mathcal{H}(M(L_1, L_2, L_3)) = D \ln N + P(x) + o(1) \quad \text{as } N \to \infty, \quad L_i/N \longrightarrow x_i > 0. \tag{6.4}$$

The widths x_i describe polynomial degrees; the radii r_i describe factor positions. They are independent variables.

For positive coefficients f_i, define

$$Q(f) = D + \sup_{x_i > 0} \big(P(x) - f \cdot x\big). \tag{6.5}$$

Here $f \cdot x = \sum_i f_i x_i$ penalizes large widths; Q records the best remaining value and is finite and convex.

Theorem 6.1 (Polynomial multiplication input [17, 18])**.** *At every positive radial placement, the profile P and energy Q defined above have the following properties.*

1. *P is finite and concave on the positive orthant, and $P(cx) = P(x) + D \ln c$. With the other widths fixed, $P(x)$ tends to infinity as any one width tends to infinity.*
2. *With only x_i varying, the function $\exp(P(x)/r_i)$ is concave.*
3. *Writing $\mathbf{1} = (1, 1, 1)$, the same placement has*

$$Q(f) - Q(\mathbf{1}) = \int_0^\infty \left[F\left(\frac{r_1}{f_1 + t}, \frac{r_2}{f_2 + t}, \frac{r_3}{f_3 + t} \right) - \frac{D}{1 + t} \right] dt. \tag{6.6}$$

No regularity of F is presumed. Synchronized coefficient blocks give concavity of $\exp(P/D)$, hence of P; comparison of successive degrees in one input gives the stronger scale r_i. The corresponding large-input comparison gives positive growth. The energy identity is the homogeneous radial specialization in [18, note C1004] of the continuum energy construction in [17, subsection "From degrees to widths"], with our logarithm convention. Its integrand is $O((1 + t)^{-2})$ by the radial Lipschitz bound, so the integral converges.

Regularity from a supporting linear function. A vector p is a *subgradient* of the convex function F at r if

$$F(z) \ge F(r) + p \cdot (z - r) \qquad (z_i > 0).$$

It gives a supporting affine function from below. For a concave function, the reverse inequality defines a *supergradient*. Write $\partial F(r)$ for the set of subgradients at r. The radial properties above imply $0 \le p_i \le 1$ for every subgradient, and homogeneity gives $p \cdot r = F(r)$.

For a real symmetric matrix B, we write $B \succeq 0$ when $x^{\mathsf{T}} B x \ge 0$ for every real vector x; such a matrix is *positive semidefinite*. It is *positive definite* when the inequality is strict for $x \ne 0$. The notation $B \succeq C$ means $B - C \succeq 0$, and $\operatorname{diag}(b_i)$ denotes the diagonal matrix with entries b_i.

Lemma 6.2 (Radial regularity)**.** *Suppose $r_i > 0$ and F has a subgradient p at r with*

$$p_i > 0, \qquad p_1 + p_2 + p_3 > 2.$$

Then F is differentiable at r. It is $C^{1,1}$ on a neighborhood of r: its gradient exists there and is locally Lipschitz. At almost every point of this neighborhood, put

$$p_i = F_i, \qquad g_i = r_i p_i, \qquad D_r = \operatorname{diag}(r_1, r_2, r_3).$$

Then the positive definite matrix

$$A = \operatorname{diag}(g_1, g_2, g_3) + \tfrac{1}{2} D_r (\nabla^2 F) D_r \tag{6.7}$$

satisfies

$$(A^{-1})_{ii} \geq \frac{1}{r_i} \qquad (i = 1, 2, 3). \tag{6.8}$$

The proof combines the three coordinate concavity bounds into a strict quadratic estimate, then uses convex duality and the energy identity [19]; see Appendix A.2.

Choosing the radii.

Proposition 6.3 (A stationary minimizing configuration)**.** *A counterexample to a proposed universal exponent $\beta > 2/\sqrt{3}$ gives three Fermat rays and positive radii r_i, with $\sum_i r_i = 1$, for which*

$$h := F(r) = \min_{\substack{s_i \geq 0 \\ \sum_i s_i = 1}} F(s) < \frac{1}{\beta}, \qquad F_1(r) = F_2(r) = F_3(r) = h.$$

Proof. The initial tangent configuration has $F(r^0) \leq \eta < 1/\beta$ by (5.7). Minimize F on the closed radius simplex. On a face, the surviving tips have mutual angle at least 120 degrees. Restricting their diagonal to the identity between them gives

$$F(0, b, c) \geq \sqrt{b^2 + bc + c^2} \geq \frac{\sqrt{3}}{2}(b + c).$$

At a simplex vertex the value is one. Since $1/\beta < \sqrt{3}/2$, the minimum occurs at positive radii.

Homogeneity and the simplex minimum give $F(s) \geq h \sum_i s_i$ for every positive s, with equality at the chosen r. Thus $h\mathbf{1}$ is a subgradient. The previously proved bound $\alpha_k \leq 3\sqrt{3}/4 < 3/2$ [11], applied through (6.3), gives $h > 2/3$. Lemma 6.2 proves differentiability, and its unique subgradient is $h\mathbf{1}$. □

We sort the radii and their angular labels only after minimizing. Equal sorted radii are therefore interior cases of this argument.

An inequality for the weighted derivatives. At a differentiable positive placement write

$$p_i = F_i, \qquad g_i = r_i p_i, \qquad u_i = 1 - p_i.$$

The quantities g_i are the *weighted derivatives*; Euler's identity is $\sum_i g_i = F(r)$. The nonnegative u_i measure the difference between the derivative and its upper bound one.

Lemma 6.4. *In an open region with $p_i > 0$ and $\sum_i p_i > 2$, F is locally $C^{1,1}$. At almost every point put*

$$\lambda_{ij} = -\tfrac{1}{2} r_i r_j F_{ij} \geq 0 \quad (i \neq j).$$

For distinct i, j, k,

$$r_i u_i \geq \frac{\lambda_{ij} g_j}{\lambda_{ij} + g_j} + \frac{\lambda_{ik} g_k}{\lambda_{ik} + g_k}. \tag{6.9}$$

If $g_j, g_k \geq B > r_i u_i$, then

$$0 \leq F_{ii} \leq \frac{2B u_i}{r_i (B - r_i u_i)}. \tag{6.10}$$

The corresponding integrated inequalities hold on every coordinate segment with an open neighborhood in this region and with the indicated strict denominator conditions.

Proof. Homogeneity gives $(\nabla^2 F)r = 0$. Thus, writing $\mathbf{e}_i$ for the standard coordinate vectors,

$$\tfrac{1}{2} D_r (\nabla^2 F) D_r = \sum_{i<j} \lambda_{ij} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^{\mathsf{T}}.$$

Delete the edge term λ_{jk} from this matrix. This decreases the positive definite matrix A, so increases its inverse in positive-semidefinite order. Taking the Schur complement of the j, k entries gives

$$\frac{1}{(A^{-1})_{ii}} \geq g_i + \frac{\lambda_{ij} g_j}{\lambda_{ij} + g_j} + \frac{\lambda_{ik} g_k}{\lambda_{ik} + g_k}.$$

By (6.8), the left side is at most r_i, proving (6.9).

Put $X = \lambda_{ij} + \lambda_{ik} = r_i^2 F_{ii}/2$. Monotonicity in g_j, g_k and direct fraction addition give

$$r_i u_i \geq \frac{B\lambda_{ij}}{B + \lambda_{ij}} + \frac{B\lambda_{ik}}{B + \lambda_{ik}} \geq \frac{BX}{B + X}.$$

Solving this inequality proves (6.10). Integrate first on almost every parallel coordinate line, then pass to any prescribed compact segment by continuity of the gradient and the strict denominator conditions. □

A convenient sufficient condition for this regular region is

$$F(r) > \max\{r_1 + r_2, r_1 + r_3, r_2 + r_3\}.$$

Each omitted-input diagonal is bounded by its two remaining radii, so every subgradient has positive coordinates. Also,

$$(\min_i r_i) \sum_i (1 - p_i) \leq \sum_i r_i (1 - p_i) = \sum_i r_i - F(r) < \min_i r_i.$$

Hence $\sum_i p_i > 2$. The endpoint tests verify these strict inequalities on every segment using derivative bounds.

Controlling the integral along a shortening. Hold radii r_j, r_k fixed and let the selected radius be t. Write

$$u(t) = 1 - F_i(t, r_j, r_k), \qquad g_j(t) = r_j F_j(t, r_j, r_k), \qquad g_k(t) = r_k F_k(t, r_j, r_k).$$

Assume the whole segment has an open neighborhood in the regular region above. Submodularity says that the neighboring weighted derivatives increase when t decreases.

Lemma 6.5 (A monotone product)**.** *On every such segment with positive neighboring weighted derivatives, the function*

$$t \longmapsto \frac{t^2 u(t)}{g_j(t) g_k(t)} \tag{6.11}$$

is nondecreasing.

Proof. At a second-order point put $x = \lambda_{ij}$, $y = \lambda_{ik}$, $d = tu(t)$. Equation (6.9) gives

$$d \geq \frac{x g_j}{x + g_j} + \frac{y g_k}{y + g_k} \geq x + y - d(x/g_j + y/g_k).$$

For the second inequality, each term in the first sum is at most d, so $x/(x + g_j) \leq d/g_j$, and similarly for y. Since $g_j' = -2x/t$, $g_k' = -2y/t$, and $u' = -2(x + y)/t^2$, the derivative of (6.11) is

$$\frac{2}{g_j g_k} \big[d - x - y + d(x/g_j + y/g_k) \big] \geq 0.$$

No division by u was used. Integration on nearby lines and continuity give the assertion on the prescribed segment. □

Start at $t = a > 0$, first supposing $q = u(a) > 0$. Write the two initial neighboring weighted derivatives as b and $\rho_{\mathrm{act}}^2 b$, where $b > 0$ and $\rho_{\mathrm{act}} \geq 1$. For a shortening to $t < a$, put $J(t) = \int_t^a u(v)\,dv$. Integration of $F_i = 1 - u$ gives

$$F(a, r_j, r_k) = F(t, r_j, r_k) + (a - t) - J(t). \tag{6.12}$$

A geometric lower bound at the shortened endpoint and an upper bound for J therefore bound the initial value from below. The following estimate retains both neighboring derivatives [21].

Lemma 6.6 (The two initial weighted derivatives)**.** *Let $B \leq b$ be positive and choose $1 \leq \rho \leq \rho_{\mathrm{act}}$. Put $\mu = aq/B$, and suppose $0 < \mu < 1$. Define*

$$z_1 = \frac{\rho\mu}{\rho - 1 + \mu}, \qquad m_1 = \frac{\mu}{z_1} = \frac{\rho - 1 + \mu}{\rho}.$$

For an endpoint $z = t/a$, the integral is bounded as follows:

(1) *If $z_1 \leq z \leq 1$, then*

$$J(az) \leq aq\,\frac{(1-\mu)(1-z)}{z-\mu}. \tag{6.13}$$

(2) *If $z < z_1$ and the equation*

$$zve^{-v} = \mu e^{-m_1} \tag{6.14}$$

has a solution $m_1 \leq v < 1$, then

$$J(az) \leq B\left\{\rho^2\left[\frac{zv(2-v)}{\mu} - 1\right] - 1 + \mu\right\}. \tag{6.15}$$

A known upper bound for q may replace its actual value, provided the displayed branch conditions are checked. If $q = 0$, then $J = 0$ on every inward segment satisfying the hypotheses.

Proof. Let Z be the sum of the two neighboring derivative increments. For fixed Z, their product is largest by increasing the smaller one until they agree, then increasing both equally. The ratio to the initial product is at most

$$\Pi_\rho(w) = \begin{cases} 1 + w, & 0 \leq w \leq \rho^2 - 1, \\ (1 + \rho^2 + w)^2/(4\rho^2), & w \geq \rho^2 - 1, \end{cases} \qquad w = Z/b.$$

This increases in w and decreases in ρ, so the lower bounds B, ρ preserve the upper estimate. Put

$$R(z) = \sqrt{u(az)/q}, \qquad I(z) = J(az)/(aq) = \int_z^1 R(v)^2\,dv.$$

Euler gives $Z = aq[zR(z)^2 - 1 + I(z)]$. Lemma 6.5 therefore yields

$$z^2 R(z)^2 \leq \Pi_\rho(\mu[zR(z)^2 - 1 + I(z)]). \tag{6.16}$$

The first branch of its equality solution is

$$R_+(z) = \frac{1-\mu}{z-\mu}, \qquad I_+(z) = \frac{(1-\mu)(1-z)}{z-\mu}.$$

At $z = z_1$ it reaches $zR_+ = \rho$. Thereafter set $R_+ = \rho v/\mu$, where (6.14) gives $v' = -v/[z(1-v)]$. The solutions agree at z_1; integrating $I_+' = -R_+^2$ gives the two stated bounds.

Here is why the equality solution is an upper bound. For $zR \geq 1$, invert Π_ρ in (6.16) and move the term $\mu z R^2$ to the left. The left side becomes

$$\begin{cases} z(z-\mu)R^2 - 1, & zR \leq \rho, \\ 2\rho zR - \mu z R^2 - 1 - \rho^2, & zR \geq \rho. \end{cases}$$

These pieces and their derivatives agree at the join and increase strictly with R for $z > \mu$, $R < \rho/\mu$. Write $x_+ = \max\{x, 0\}$ for the positive part. On compact parts of this branch, comparison with equality bounds $(R - R_+)_+$ by a constant times its integral to one. Gronwall's inequality makes it zero. Continuity prevents crossing $R = \rho/\mu$ while the certified root has

$v < 1$. When $zR < 1$, the bound is automatic because $zR_+ \geq 1$. This covers the joining point as well.

The profile R_+ is nondecreasing in μ: differentiate its rational first branch, and on the second use

$$\partial_\mu \ln R_+ = \frac{v - \mu/\rho}{\mu(1-v)} \geq 0.$$

Continuity preserves this direction across the join. Since $J \leq aq \int_z^1 R_+^2$, increasing the initial deficit is safe when the branch conditions hold. If $q = 0$, the nonnegative product (6.11) is at most its initial value zero on the inward segment, so $u = J = 0$. □

At the stationary placement of Proposition 6.3,

$$q = 1 - h, \qquad g_j = r_j h, \qquad g_k = r_k h.$$

A prior lower bound for h bounds q above and the two weighted derivatives below; their ratio is the ratio of the radii. Table 1 records these prior bounds.

Updating the permissible height. Suppose $h_- \leq h \leq h_+$ and a shortening by $T > 0$ has endpoint lower bound E. Under the trial $q = 1 - h \leq q_0$, Lemma 6.6 supplies $\overline{J}(q_0)$ when its regularity and branch conditions hold. Equation (6.12) gives

$$h \geq E + T - \overline{J}(q_0).$$

If this exceeds h_+, then $q > q_0$ and $h < 1 - q_0$. Replace h_+ by $\min\{h_+, 1 - q_0\}$ in subsequent tests. The common initial derivative h makes the update valid for every radius; each update is a finite implication.

7. Euclidean lower bounds for the diagonal

We now obtain explicit lower bounds for the same diagonal function F. Put $v_i = r_i e_i$. The rays stay fixed; unless a normalization is stated, the radii may be any positive numbers.

The five parameters and a quadratic inequality. Relabel a minimizing radial vector, and the directions with it, so that $r_1 \leq r_2 \leq r_3$. The third radius is $r_3 = 1 - r_1 - r_2$. For each pair define its angular weight

$$w_{ij} = -\cos\theta_{ij} - \frac{1}{2} \in [0, 1/2]. \tag{7.1}$$

The three distances between the tips are

$$d_{ij}(r) = \sqrt{r_i^2 + r_j^2 + (1 + 2w_{ij}) r_i r_j}. \tag{7.2}$$

Thus the normalized parameters lie in the closed region

$$\mathcal{D} = \left\{ (r_1, r_2, w_{12}, w_{13}, w_{23}) : \begin{array}{l} 0 \leq r_1 \leq r_2 \leq 1 - r_1 - r_2, \\ 0 \leq w_{12}, w_{13}, w_{23} \leq 1/2 \end{array} \right\}. \tag{7.3}$$

In particular, $r_1 \leq 1/3$ and $r_2 \leq 1/2$. We allow boundary points and parameters not realized by a spectral space: excluding this larger domain suffices. A *box* is a product of closed intervals.

Theorem 7.1 (Amanov–van den Berg–Zuiddam [3, Proposition 4.14])**.** *Let ψ be a spectral point on three positions, and put $b_{ij} = \log_2 \psi(E_{ij}(2))$. Then, for all $a_i \geq 0$,*

$$4 \sum_{i<j} b_{ij} a_i a_j \leq (a_1 + a_2 + a_3)^2.$$

Over $\mathbb{C}$, Uznański [26, Proposition 5.2, equation (22)] proved the equivalent angular condition $\sum_{i<j} \arccos \sqrt{b_{ij}} \geq \pi/2$. We use the all-field quadratic form above.

Normalize the logarithmic values on our three tips by their diagonal value $F(r)$. Their spectral point has $b_{ij} = d_{ij}(r)/F(r)$. Consequently, for every $a_i \geq 0$ with $\sum_i a_i = 1$,

$$4 \sum_{i<j} a_i a_j d_{ij}(r) \leq F(r). \tag{7.4}$$

For $a_i = 1/3$ this gives $\sum_{i<j} d_{ij}(r) \le 9F(r)/4$. If all three pair angles are π, then the distance sum is $2\sum_i r_i$, so $F(r) \ge 8\sum_i r_i/9$. This case already exceeds the trial height in the upper-bound proof, so the remaining comparisons may choose a pair with angle less than π.

Mean width on a Euclidean halfspace. The comparisons below may occupy only a halfspace. The Euclidean value theorem still gives the lower bound we need.

Lemma 7.2. *Suppose a Euclidean halfspace $E_+ \subseteq E = \mathbb{R}^d$ embeds isometrically in a geodesic spectral space with jointly convex logarithmic values. For every finite nonempty $A \subseteq E_+$,*

$$\delta(A) \ge W(A).$$

Proof. Let n be the inward unit normal. For any fixed nonzero tensor T and placement $x = (x_1, \dots, x_r)$ in E, all positions $x_i + tn$ lie in E_+ for sufficiently large t. Their logarithmic value is a convex function of t. If the index dimensions are q_i, transport to the first position bounds it by

$$0 \le H(T; x + tn) \le \sum_{i=2}^{r} \|x_i - x_1\| \log_2 q_i.$$

A convex function bounded above on a half-line is nonincreasing. It therefore has a limit $H_\infty(T; x)$.

These limits give a spectral space on E. Indeed, on a finite support S of positive diameter the diagonal normalizations lie between $\operatorname{diam} S$ and a fixed spanning-tree length. The normalized spectral axioms pass to the limit, and using one translation on each finite union of supports proves compatibility. Theorem 3.6 now gives $\delta_\infty(A) = W(A)$. For $A \subseteq E_+$ the translated logarithmic value is nonincreasing from its original value, proving the inequality. □

Comparison by moving far along a ray. Moving far along one ray separates the radial and transverse directions and decreases every logarithmic tensor value.

Fix a unit direction $b \in \Sigma$. The product $\mathbb{R} \times T_b\Sigma$ is equipped with the metric

$$d\big((a,u),(a',u')\big)^2 = (a-a')^2 + d_{T_b\Sigma}(u,u')^2.$$

For $v = re \in C$, put $\vartheta = d_\Sigma(b,e)$. If $0 < \vartheta < \pi$, let u be the unit tangent direction at b along the shortest arc to e, and define

$$Q_b(v) = (r\cos\vartheta,\ r\sin\vartheta\, u). \tag{7.5}$$

For $\vartheta = 0$ or π, the second coordinate is zero; also put $Q_b(0) = 0$.

Lemma 7.3 (Comparison along a direction)**.** *The whole product $\mathbb{R} \times T_b\Sigma$ admits compatible, jointly convex logarithmic spectral values H_b for which*

$$|Q_b(v)| = |v|, \qquad d(Q_b(v), Q_b(w)) \le d_C(v,w),$$

and, for every nonzero tensor T and every finite placement,

$$H_b(T; Q_b(v_1), \dots, Q_b(v_r)) \le H(T; v_1, \dots, v_r). \tag{7.6}$$

Proof. For $R > 0$, define $Q_R(v) = 2[Rb, v]_{1/2}$, where 2 means dilation about the cone origin. Joint convexity, homogeneity, and the zero value when all factors are placed at Rb give

$$H(T; Q_R(v_1), \dots, Q_R(v_r)) \le H(T; v_1, \dots, v_r). \tag{7.7}$$

Convexity of distance gives $d_C(Q_R(v), Q_R(w)) \le d_C(v,w)$.

Take the pointed limit as $R \to \infty$, with basepoint Rb. To realize the whole product in this limit, suppose $u \in T_b\Sigma$ is represented by an arc γ_u starting at b with speed $|u|$, represent (a,u) by

$$z_R(a,u) = (R+a)\gamma_u(1/R).$$

For sufficiently large R this is defined, and the cone distance formula gives

$$\begin{aligned} d_C(z_R(a,u), z_R(a',u'))^2 &= (a-a')^2 + 2(R+a)(R+a')\big(1 - \cos d_\Sigma(\gamma_u(1/R), \gamma_{u'}(1/R))\big) \\ &\longrightarrow (a-a')^2 + d_{T_b\Sigma}(u,u')^2. \end{aligned} \tag{7.8}$$

Completion embeds the whole product. Its geodesics are the unique ambient geodesics, so its image is convex and the inherited values are jointly convex.

For $0 < \vartheta < \pi$, the cone over the arc from b to e is a Euclidean sector. In that sector, $Q_R(re)$ is the vector sum of Rb and re. Relative to Rb, its limiting coordinates are precisely (7.5). At angles zero and π, the two rays form a ray or a line, and the same formula follows directly. Thus the limit preserves each distance from the origin and does not increase other distances. Passing to the same limit in (7.7) proves (7.6) for every fixed tensor. □

In particular, for $w_i = Q_b(v_i)$,

$$d(w_i, w_j) \le d_C(v_i, v_j), \qquad \delta(\{0, w_1, w_2, w_3\}) \le F(r). \tag{7.9}$$

The Euclidean model along a link arc. Choose a pair i, j with $\theta = \theta_{ij} < \pi$, and let k be the remaining index. Let $\gamma : [0, \theta] \to \Sigma$ be the unit-speed arc from e_i to e_j. Put

$$c(t) = d_\Sigma(\gamma(t), e_k), \qquad f(t) = \cos c(t).$$

The triangle inequality makes c 1-Lipschitz, hence differentiable at almost every interior parameter; call these parameters *regular*. The angles from e_k to both arc endpoints are at least $2\pi/3$, so $c(t) \ge 2\pi/3 - \min\{t, \theta - t\} \ge \pi/6 > 0$.

Lemma 7.4 (A Euclidean model at a regular parameter)**.** *At each regular $t \in (0, \theta)$ there is a Euclidean model*

$$\begin{aligned} w_i &= r_i(\cos t, -\sin t, 0), \\ w_j &= r_j(\cos(\theta - t), \sin(\theta - t), 0), \\ w_k &= r_k(f(t), f'(t), \sqrt{1 - f(t)^2 - f'(t)^2}). \end{aligned} \tag{7.10}$$

The square root is real. These points satisfy

$$\begin{aligned} W(\{0, w_i, w_j, w_k\}) &\le F(r), \\ |w_i - w_j| &= d_C(v_i, v_j), \\ |w_\ell - w_k| &\le d_C(v_\ell, v_k) \qquad (\ell = i, j). \end{aligned} \tag{7.11}$$

Furthermore,

$$f(t) \le -\cos t - \cos(\theta - t). \tag{7.12}$$

Proof. Write $b = \gamma(t)$. The arc gives two opposite unit directions $+$ and $-$ in $T_b\Sigma$. If $c(t) < \pi$, let u point towards e_k. First variation, applied in the CAT(1) link, gives

$$\cos \angle(+, u) = -c'(t), \qquad \cos \angle(-, u) = c'(t).$$

Their two angles sum to π. If both are positive, the shortest arcs from $+$ to u and from u to $-$ in the space of directions of Σ concatenate to an interval of length π: its length equals the distance of its endpoints. The cone on that interval is a Euclidean half-plane in $T_b\Sigma$. If one angle is zero, all three directions lie on its boundary line. Consequently the relevant part of $\mathbb{R} \times T_b\Sigma$ is a Euclidean halfspace, or a Euclidean plane in the degenerate case.

Choose its transverse coordinates so that $+$ is the positive second coordinate and the third coordinate of u is nonnegative. Formula (7.5) gives the displayed coordinates, since $\sin c(t) \cos \angle(+, u) = -\sin c(t) c'(t) = f'(t)$. The remaining component has the stated squared length. If $c(t) = \pi$, differentiability gives $c'(t) = 0$ and $f'(t) = 0$; the third image is purely radial, and the radial coordinate together with the arc tangent line already gives a Euclidean plane. Thus the same formulas cover this antipodal case.

Apply Lemma 7.2, or Theorem 3.6 in the planar case, to (7.9). This proves the mean-width and distance bounds. The base distance is preserved by (5.2), because the two displayed base vectors make angle θ. Finally, (7.12) is (5.6) at b. □

At a nonregular interior parameter t_0, take regular $t_\nu \to t_0$ and a subsequence on which the bounded derivatives converge. The model and its inequalities pass to the limit by continuity. This includes the midpoint. No base arc is assumed when $\theta_{ij} = \pi$; every arc comparison requires $\theta_{ij} < \pi$.

Why contracting distances decreases mean width. Contracting a model preserves a valid lower bound by the following lemma.

Lemma 7.5 (Sudakov–Fernique comparison [6, Theorem 1.1])**.** *Let $a_1, \dots, a_m$ and $b_1, \dots, b_m$ be Euclidean configurations, possibly in different dimensions. If $|a_i - a_j| \le |b_i - b_j|$ for every i, j, then $W(\{a_i\}) \le W(\{b_i\})$.*

Proof. Take independent standard Gaussian vectors g, g' in the two ambient spaces. The variables $X_i = \langle g, a_i \rangle$ and $Y_i = \langle g', b_i \rangle$ satisfy $\mathbb{E}(X_i - X_j)^2 \le \mathbb{E}(Y_i - Y_j)^2$. For $a > 0$, the smooth maximum $f_a(z) = a^{-1} \ln \sum_i \exp(a z_i)$ lies between $\max_i z_i$ and $\max_i z_i + (\ln m)/a$, where ln is the natural logarithm. For a smooth function F, Gaussian integration by parts reads $\mathbb{E}[g_i F(g)] = \mathbb{E}[\partial_i F(g)]$. Applying it to $Z_t = \sqrt{1-t}X + \sqrt{t}Y$ gives

$$\frac{d}{dt}\mathbb{E} f_a(Z_t) = \frac{a}{4} \sum_{i,j} \mathbb{E}(p_i p_j)(\mathbb{E}(Y_i - Y_j)^2 - \mathbb{E}(X_i - X_j)^2) \ge 0,$$

where $p_i = e^{a(Z_t)_i} / \sum_j e^{a(Z_t)_j}$. Integrate from zero to one and let $a \to \infty$. Equation (3.7) gives the assertion. □

The contribution of the retained origin. The auxiliary position remains the origin of every comparison. Retaining it can improve the half-perimeter bound for the three tips.

Lemma 7.6 (An additional mean-width bound)**.** *Let $K = \operatorname{conv}\{x_1, x_2, x_3\} \subset \mathbb{R}^3$ be a nondegenerate triangle of perimeter P. Let d be the distance from the origin to its affine plane, and suppose $0 \le l \le d$ and $P \le P_+$. Then*

$$W(\{0, x_1, x_2, x_3\}) \ge \frac{P}{2} + \frac{5l^3}{3(5(P_+/(2\pi))^2 + 3l^2)} \ge \frac{P}{2} + \frac{l^3}{3((P_+/(2\pi))^2 + l^2)}. \tag{7.13}$$

When $l = 0$, the second term is interpreted as zero. For a degenerate triangle, the lower bound $W(\{0, x_1, x_2, x_3\}) \ge P/2$ is used.

Proof. Assume $l > 0$; the zero and degenerate cases follow from monotonicity of mean width. Let c be the perpendicular projection of the origin to the plane of K. For a unit vector u parallel to that plane, put $\rho(u) = \max_i \langle x_i - c, u \rangle$. This is the support function of the triangle relative to c. Its circle average is $v = P/(2\pi)$. The function ρ may be negative when c lies outside the triangle; the calculation below allows this.

Integrating the added support in the perpendicular direction in the spherical formula for mean width gives

$$W(\{0, x_1, x_2, x_3\}) - \frac{P}{2} = \mathbb{E} J_d(\rho), \qquad J_d(s) = d - s \operatorname{atan2}(d, s). \tag{7.14}$$

Here the expectation is the uniform circle average, and $\operatorname{atan2}(d, s) \in (0, \pi)$ has sine $d/\sqrt{d^2 + s^2}$ and cosine $s/\sqrt{d^2 + s^2}$. Measure the polar angle θ from the perpendicular direction pointing from c towards the origin. At a fixed planar direction the added support is the positive part of $d \cos\theta - t \sin\theta$; its integral against $\sin\theta \, d\theta$ is $J_d(t)/2$, and the normalization of mean width gives the displayed identity.

For $d > 0$,

$$\frac{\partial J_d(s)}{\partial d} = \frac{d^2}{d^2 + s^2} \ge 0, \qquad J_d''(s) = \frac{2d^3}{(d^2 + s^2)^2} > 0. \tag{7.15}$$

We may therefore replace d by l and apply convexity to the circle average. Since $v > 0$, this gives

$$\mathbb{E} J_d(\rho) \ge J_l(v) = \int_0^l \frac{t^2}{v^2 + t^2} \, dt \ge \frac{5l^3}{3(5v^2 + 3l^2)}.$$

For the last step, apply convexity of $x \mapsto 1/x$ to the probability density $3t^2/l^3$ on $[0, l]$. The mean of t^2 under this density is $3l^2/5$. Now use $v \le P_+/(2\pi)$; the second inequality in the statement is immediate. □

8. Allowing the other two radii to vary

We strengthen the endpoint comparison by allowing the other two radii to redistribute while keeping their sum fixed [25]. The resulting upper bound applies to a minimizing split. To contradict it, Section 10 will exclude every split, including those with a zero radius.

Throughout this section, F is the same diagonal function on the three Fermat rays. Write

$$k_{ij} = -\cos\theta_{ij}, \qquad \tfrac{1}{2} \le k_{ij} < 1.$$

All three pairs must therefore be nonantipodal; other bounds handle boxes reaching an antipodal pair.

The exact value on a two-ray face. The face values and their affine lower bounds will supply regularity at positive splits and control the zero-radius cases directly.

Lemma 8.1 (Face values and an affine lower bound)**.** *For $t, z \ge 0$, after relabeling the surviving pair as $1, 2$, the value with the third radius zero is*

$$F(t, z, 0) = \frac{t + z + \sqrt{t^2 + z^2 + 2k_{12}tz}}{2}. \tag{8.1}$$

For distinct indices i, j, k and every nonnegative radius vector r,

$$F(r) \ge F(r \text{ with } r_i = 0) + (1 - k_{jk})r_i. \tag{8.2}$$

Proof. For any three positions a, b, c in a spectral space, one has

$$2\delta(\{a, b, c\}) \le d(a, b) + d(b, c) + d(c, a). \tag{8.3}$$

To prove this, place a dimension-n identity on each edge, with indices i, j, k. The progression-free construction cited in [12, paragraph following Proposition 5.2] gives $B \subseteq \{1, \dots, \lfloor n/3 \rfloor\}$ of size $n^{1-o(1)}$ with no nonconstant three-term progression. Project the three local coordinate pairs onto $j - i \in B$, $k - j \in B$, and $k - i \in 2B := \{2b : b \in B\}$. The consecutive differences must agree. Each surviving triple is determined by any local pair, so these projections produce a diagonal of size at least $\lfloor n/3 \rfloor |B| = n^{2-o(1)}$: take $i \le \lfloor n/3 \rfloor$ and the common difference in B. Taking logarithmic values and dividing by $\log_2 n$ proves (8.3) over every coefficient field.

Apply Lemma 7.4 with the nonantipodal pair $1, 2$ as base and let the remaining radius tend to zero. Its mean width is half the perimeter of the apex-tip triangle, giving the reverse inequality. The tip distance is $\sqrt{t^2 + z^2 + 2k_{12}tz}$, so (8.1) follows, including zero radii by continuity. For an antipodal surviving pair the apex lies on the tip segment, giving the same formula directly by free insertion.

For (8.2), use the midpoint comparison on the opposite pair j, k. Put $q^2 = (1 - k_{jk})/2$ and $s^2 = 1 - q^2$. The two base directions are $u_j = (q, s, 0)$ and $u_k = (q, -s, 0)$; the third unit direction v has first coordinate $f \le -2q$. Discarding the apex from the Euclidean comparison gives half the perimeter of the triangle with tips $r_i v, r_j u_j, r_k u_k$. The inequalities

$$|r_j u_j - r_i v| \ge r_j - r_i \langle u_j, v \rangle, \qquad |r_k u_k - r_i v| \ge r_k - r_i \langle u_k, v \rangle$$

show that this half-perimeter is at least

$$\frac{r_j + r_k + |r_j u_j - r_k u_k|}{2} - q f r_i \ge F(r \text{ with } r_i = 0) + (1 - k_{jk})r_i.$$

This proves the affine bound. Its conclusion for an antipodal opposite pair is simply coordinatewise monotonicity. □

Minimizing over the split. Relabel the chosen moving ray as the first ray and define

$$g(t, z) = \min_{\substack{v, w \ge 0 \\ v + w = z}} F(t, v, w), \qquad t, z > 0. \tag{8.4}$$

The minimum exists by compactness. Interpolating attaining triples proves convexity; scaling them proves homogeneity. Monotonicity and separate 1-Lipschitz continuity also pass from F to g.

Lemma 8.2 (Regularity of the minimized diagonal function)**.** *Fix $z > 0$ and a positive interval of values of t. Suppose that, at every point of the interval,*

$$t \geq z/2 \qquad \text{or} \qquad g(t,z) > z. \tag{8.5}$$

Then $g(\,\cdot\,, z)$ is $C^{1,1}$ locally on the interval and

$$g_{tt}(t,z) \leq \frac{2(1-g_t(t,z))}{t} + \frac{2(1-g_t(t,z))^2}{g(t,z)-t} \tag{8.6}$$

almost everywhere. The conclusion includes points at which a minimizing split has a zero radius or is not unique. Moreover,

$$g(t,z) - t > z/2. \tag{8.7}$$

Proof. We first prove the strict gap without a regularity assumption. At any split, relabel the neighbors so that $v \geq w$, and put $k = k_{12} < 1$ for the pair formed by the moving ray and the neighbor of radius v. The face value and affine bound give

$$F(t,v,w) \geq \frac{t + v + \sqrt{t^2 + v^2 + 2ktv}}{2} + (1-k)w.$$

Since the square root is at least $t + kv$,

$$F(t,v,w) - t - z/2 \geq \frac{k(v-w) + (1-k)w}{2} > 0.$$

The last inequality also holds when $w = 0$, since $v = z > 0$. Taking the minimum over the compact split interval proves (8.7).

At a positive attaining split, every subgradient (p, λ) of g pulls back under $(s, v', w') \mapsto (s, v' + w')$ to a subgradient of F. Hence

$$(p, \lambda, \lambda) \in \partial F(t,v,w), \qquad 0 \leq p, \lambda \leq 1, \qquad g = tp + z\lambda.$$

Comparison with the affine face bounds gives $p_i \geq 1 - k_{jk} > 0$ for every component. With $d = 1 - p \in [0,1]$, Euler's identity becomes

$$z(p + 2\lambda - 2) = 2(g - t) - z + (2t - z)d. \tag{8.8}$$

The right side is positive when $t \geq z/2$, and is at least $2(g - z) > 0$ under the other guard. Lemma 6.2 makes F locally $C^{1,1}$ there. Its unique gradient then makes the subgradient of g unique as well.

At a boundary split, g has a smooth upper contact with (8.1), again forcing a unique subgradient. Thus g is differentiable throughout the guarded region, which is locally open: at $t = z/2$ the strict gap implies $g > z$. Convexity makes its gradient continuous.

For the curvature bound, we must also handle optimizer sets of three-dimensional measure zero. A *smooth lower test* is a twice continuously differentiable function that agrees at the point of contact and lies below the function in a neighborhood. Lemma B.1 transfers the inverse-matrix inequality to such tests at every positive attaining split. If φ touches $g(\,\cdot\,, z)$ from below at t, homogeneity gives the lower test

$$\psi(s, v', w') = \frac{v' + w'}{z}\, \varphi\!\left(\frac{zs}{v' + w'}\right).$$

Its weighted Hessian at (t, v, w), meaning $\frac{1}{2} D_{(t,v,w)} (\nabla^2 \psi) D_{(t,v,w)}$, equals $\gamma a a^{\mathsf{T}}$, where

$$\gamma = \tfrac{1}{2} t^2 \varphi''(t), \qquad a = (1, -v/z, -w/z).$$

The test matrix is $\operatorname{diag}(tp, v\lambda, w\lambda) + \gamma a a^{\mathsf{T}}$. For $\gamma > 0$ it is positive definite, and the rank-one inverse formula, using $(v/z)^2/(v\lambda) + (w/z)^2/(w\lambda) = 1/(z\lambda)$, gives

$$tp + \frac{\gamma z \lambda}{\gamma + z\lambda} \leq t. \tag{8.9}$$

Writing $u = 1 - p$ and $K = g - t$, one has $z\lambda = K + tu$. Equation (8.9) becomes $\gamma K \leq tu(K + tu)$, which is (8.6) with φ'' in place of g_{tt}. For $\gamma \leq 0$ this upper bound is automatic. At a boundary split, differentiating the exact face function gives the same bound. Appendix B proves this

calculation and the final step: $g(\cdot, z)$ is *semiconcave*, meaning that locally $g(t,z) - Ct^2$ is concave for some constant C. Together with convexity, this gives $C^{1,1}$ regularity and (8.6) almost everywhere. □

A rational upper bound at any guarded endpoint. The following calculus lemma also applies to the fixed-neighbor comparison in Section 9.

Lemma 8.3 (Reciprocal gap)**.** *Let f be $C^{1,1}$ on a positive interval, with $0 \le f' \le 1$. For a fixed constant c, put $K(t) = f(t) - t - c > 0$ and $u(t) = 1 - f'(t)$. If*

$$f''(t) \le \frac{2u(t)}{t} + \frac{2u(t)^2}{K(t)} \quad \text{almost everywhere,}$$

then $t^2u(t)/K(t)^2$ is nondecreasing and $y \mapsto 1/K(1/y)$ is convex. In particular, for any a, t in the interval,

$$\frac{1}{K(t)} \ge \frac{1}{K(a)} + \frac{a^2u(a)}{K(a)^2}\left(\frac{1}{a} - \frac{1}{t}\right). \tag{8.10}$$

Proof. Since $K' = -u$ and $u' = -f''$, differentiation gives $(t^2u/K^2)' \ge 0$ almost everywhere. These functions are locally absolutely continuous, so the monotonicity holds everywhere. The derivative of $1/K(1/y)$ at $y = 1/t$ is $-t^2u(t)/K(t)^2$, which is nondecreasing in y. Its tangent inequality is (8.10). □

Let r be a positive minimum of F on the simplex $r_1 + r_2 + r_3 = 1$, and write $h = F(r)$. By Proposition 6.3, $F_i(r) = h$. For the selected radius r_1, rescale the sum of the other two radii to one, and put

$$x = \frac{r_1}{1 - r_1}, \qquad g(x,1) = h(1+x), \qquad g_t(x,1) = h. \tag{8.11}$$

Indeed global minimality and homogeneity give $F(s,v,w) \ge h(s+v+w)$ for every nonnegative triple. The rescaled original split attains equality at $(x,1)$. Every subgradient of g there pulls back to a subgradient of F at this positive triple. Its gradient is (h,h,h) by Proposition 6.3 and homogeneity. Thus the subgradient of g is uniquely (h,h), proving the displayed derivative without any further regularity guard.

Proposition 8.4 (Profiled endpoint bound)**.** *Suppose an already proved uniform bound supplies $F(r) \ge h_0 \sum_i r_i$, with $0 < h_0 \le h \le h_* < 1$. Fix $t > 0$ and assume*

$$\min(x,t) \ge \tfrac{1}{2} \quad \text{or} \quad h_0(1 + \min(x,t)) > 1. \tag{8.12}$$

If $h - (1-h)x^2/t > 0$, then

$$g(t,1) \le \mathcal{U}(x,t,h) := t + \frac{[h - (1-h)x]^2}{h - (1-h)x^2/t}. \tag{8.13}$$

The endpoint t may lie on either side of x.

Proof. The automatic guard licenses Lemma 8.2 throughout the interval from x to t. Apply Lemma 8.3 to $f(s) = g(s,1)$ with $c = 0$. Here $K(x) = h - (1-h)x =: B > 0$ and $u(x) = 1 - h$, so (8.10) gives

$$\frac{1}{K(t)} \ge \frac{1}{B} - \frac{x^2(1-h)}{B^2}\left(\frac{1}{t} - \frac{1}{x}\right) = \frac{h - (1-h)x^2/t}{B^2}.$$

Inverting its positive right side proves (8.13). □

Corollary 8.5 (Four-corner upper bound)**.** *Suppose $x \in [x_-, x_+]$ and $h \in [h_0, h_*] \subset (0,1)$. For a fixed $t > 0$, assume*

$$h_0 - (1-h_0)x_+ > 0, \qquad h_0 - (1-h_0)x_+^2/t > 0,$$

and assume (8.12) with x_- in place of x. Then

$$g(t,1) \le \max_{\substack{\xi \in \{x_-, x_+\} \\ \eta \in \{h_0, h_*\}}} \mathcal{U}(\xi, t, \eta). \tag{8.14}$$

Every expression on the right is rational when its arguments are rational.

Proof. The two positivity conditions hold throughout the parameter rectangle because their left sides increase with h and decrease with x. For fixed x, the fraction in (8.13) is a square of an affine function of h divided by a positive affine function of h. Such a function is convex, so its maximum on an interval is at an endpoint. For fixed h, direct differentiation gives

$$\frac{\partial \mathcal{U}}{\partial x} = \frac{2h(1-h)[h-(1-h)x](x/t-1)}{[h-(1-h)x^2/t]^2}.$$

Thus $\mathcal{U}$ decreases up to $x = t$ and increases after it; its maximum on the x interval is also at an endpoint. These two observations prove (8.14). □

The endpoint bound concerns $g(t, 1)$, not the original split $F(t, v, w)$. It asserts that at least one split has a small value. This distinction is why the subsequent lower bounds must exclude the entire closed split interval.

9. Outward comparisons with consistent angular parameters

Let r be the positive simplex minimum, with $F(r) = F_i(r) = h$. We give a fixed-neighbor counterpart to the profiled comparison, then construct Euclidean lower bounds keeping the angular parameters together [21]. Throughout, $h \le h_* < 1$ is the candidate inequality to exclude.

An outward endpoint inequality.

Lemma 9.1. *Relabel the radii as (a, b, c), and suppose*

$$a + b + c = 1, \qquad F(a, b, c) = h \le h_* < 1, \qquad F_1(a, b, c) = F_2(a, b, c) = F_3(a, b, c) = h.$$

Let $h_0 \le h$ be a known lower bound. Fix $T \ge 0$, and write $R = a + T$ and $m = \max\{b, c\}$. Suppose that $W_{\mathrm{end}} \le F(R, b, c)$ and that

$$W_{\mathrm{end}} > R + m, \qquad \max\{W_{\mathrm{end}} - T, h_0\} > b + c. \tag{9.1}$$

Set

$$C = W_{\mathrm{end}} - R - h_* m > 0, \qquad \Delta = h_* + T - W_{\mathrm{end}}.$$

Then $\Delta \ge 0$ and

$$\Delta(C + \Delta) \ge C(1 - h_*)\frac{aT}{a + T}. \tag{9.2}$$

Without the conditions in (9.1), the inequality $\Delta < 0$ already excludes $h \le h_*$.

Proof. The Lipschitz bound gives $W_{\mathrm{end}} \le h + T \le h_* + T$, so $\Delta \ge 0$ without any regularity premise. Under the guards, put $f(t) = F(t, b, c)$ for $a \le t \le R$. Monotonicity and the radial Lipschitz bound give

$$f(t) \ge W_{\mathrm{end}} - (R - t) > t + m, \qquad f(t) \ge \max\{W_{\mathrm{end}} - T, h_0\} > b + c.$$

Thus F exceeds every two-radius sum throughout the segment, and Lemma 6.2 applies there.

Submodularity bounds the two fixed-coordinate derivatives by their initial value $h \le h_*$. Put

$$u(t) = 1 - f'(t), \qquad K(t) = f(t) - t - h_* m.$$

Then $K(t) \ge K(R) \ge C > 0$, and Euler's identity gives

$$g_j(t) \ge K(t) + tu(t),$$

for either fixed coordinate. For example, $bF_2 = f - tf' - cF_3 \ge f - t + tu - h_* c \ge K + tu$. Lemma 6.4, with $B = K + tu$, supplies the curvature hypothesis of Lemma 8.3. Apply its tangent inequality with the fixed offset $h_* m$ and $u(a) = 1 - h \ge 1 - h_*$:

$$\frac{1}{K(R)} - \frac{1}{K(a)} \ge \frac{(1 - h_*)aT}{RK(a)^2}.$$

For $J = K(a) - K(R) = h + T - f(R)$, one has $0 \le J \le \Delta$, and the last display is equivalent to

$$J + \frac{J^2}{K(R)} \ge (1 - h_*)\frac{aT}{R}.$$

Now use $K(R) \ge C$ and $J \le \Delta$ to obtain (9.2). □

The endpoint has total radius $1 + T$; F keeps its original scale.

Choosing one position on a base arc. Choose an ordered pair of directions e_j, e_k with angle $\theta = \theta_{jk} < \pi$, and let i be the remaining index. Put

$$q = \cos(\theta/2) > 0, \qquad s = \sin(\theta/2) > 0, \qquad w = w_{jk} = \frac{1}{2} - 2q^2. \tag{9.3}$$

The base directions have coordinates $(q, -s, 0)$ and $(q, s, 0)$. Choose a number $\kappa > 0$ and define

$$X_\kappa = \frac{s}{q}\frac{\kappa - 1}{\kappa + 1}, \qquad t = \frac{\theta}{2} - \arctan X_\kappa. \tag{9.4}$$

Since $|X_\kappa| < \tan(\theta/2)$, this selects an interior point of the base arc.

At this point, rotate the model of Lemma 7.4 so that its base has the fixed coordinates above. Write its third unit vector as (ξ, Y, Z), where $Z \ge 0$. The stationarity inequality in that lemma becomes

$$\xi \le -2q + X_\kappa Y. \tag{9.5}$$

Indeed, if $f(t)$ is the first coordinate before rotation and $d = t - \theta/2$, then $\xi + \tan(d)Y = f(t)/\cos d \le -2q$. At a nonregular parameter use the limiting model from that lemma. Fix its unit vectors before varying the radii: the comparison then bounds the same $F(\rho)$ at every positive ρ.

There are two alternatives for the sign of Y. If $Y \ge 0$, set $a = j$, $b = k$ and $\lambda = \kappa$. If $Y < 0$, reflect the second coordinate and set $a = k$, $b = j$ and $\lambda = 1/\kappa$. In either case write $y = |Y|$; the base endpoint indexed by a now has second coordinate $-s$. The identity $X_{1/\kappa} = -X_\kappa$ shows that these alternatives describe the same chosen position on the arc.

Put $v = w_{ia}$. The fact that the comparison does not increase the distance between directions i, a gives

$$q\xi - sy \ge -\frac{1}{2} - v.$$

Together with (9.5), in the chosen orientation, this implies

$$0 \le y \le \frac{w + v}{s - qX_\lambda} = \frac{(\lambda + 1)(w + v)}{2s}.$$

We can therefore write

$$y = \frac{(\lambda + 1)z(w + v)}{2s}, \qquad 0 \le z \le 1. \tag{9.6}$$

If $w + v = 0$, then $y = 0$, and we take $z = 0$. Define also

$$f = -2q + \frac{(\lambda - 1)z(w + v)}{2q}. \tag{9.7}$$

The actual third direction has $\xi \le f$.

Lemma 9.2 (The three comparison coefficients)**.** *For the sign and parameter z just defined, set*

$$c_{ab} = \frac{1}{2} + w, \qquad c_{ia} = \frac{1}{2} - w + z(w+v), \qquad c_{ib} = \frac{1}{2} - w - \lambda z(w+v). \tag{9.8}$$

At every positive radius vector ρ, the squared side lengths of the actual Euclidean comparison are at least

$$\rho_\ell^2 + \rho_m^2 + 2c_{\ell m}\rho_\ell\rho_m \qquad (\ell < m). \tag{9.9}$$

If $f \le 0$, there is also a Euclidean comparison with third unit vector

$$(f, y, \sqrt{1 - f^2 - y^2}), \tag{9.10}$$

whose squared side lengths equal (9.9). Its mean width, including the origin, is at most $F(\rho)$.

Proof. The negative dot products of the third direction with the two base directions are at least $-qf + sy$ and $-qf - sy$. Substitution of (9.6) and (9.7) gives the last two coefficients in (9.8). The first is the negative dot product of the base directions. The cosine law proves the side-length assertion.

If $\xi \le f \le 0$, then $f^2 + y^2 \le \xi^2 + y^2 \le 1$, so (9.10) is real. Replacing ξ by f decreases the two moving side lengths and preserves the distances to the origin. Gaussian comparison, Lemma 7.5, proves the mean-width bound. The side-length assertion itself does not require $f \le 0$. □

We call (9.10) the *saturated model*: its first coordinate equals the upper bound f for the original coordinate $\xi \le f$.

Keeping the angular coefficients together. We will enclose w, v in closed intervals $[w_-, w_+], [v_-, v_+] \subseteq [0, 1/2]$, with $w_+ < 1/2$, and enclose z in a closed subinterval $[z_-, z_+]$ of $[0, 1]$. For fixed radii define

$$P_\lambda(\rho; w, v, z) = \frac{1}{2} \sum_{\ell<m} \sqrt{\rho_\ell^2 + \rho_m^2 + 2c_{\ell m}(w, v, z)\rho_\ell\rho_m}. \tag{9.11}$$

At the parameters of the actual comparison, whenever the radicals are real, Lemma 9.2 and the half-perimeter formula for a triangle give $F(\rho) \ge P_\lambda(\rho; w, v, z)$.

Lemma 9.3 (Consistent angular corners)**.** *Suppose the three coefficients in (9.8) are at least -1 at each of the eight triples*

$$(w, v, z) \in \{w_-, w_+\} \times \{v_-, v_+\} \times \{z_-, z_+\}.$$

Then, for every fixed positive ρ,

$$P_\lambda(\rho; w, v, z) \ge \min_{\substack{w' \in \{w_-, w_+\},\ v' \in \{v_-, v_+\} \\ z' \in \{z_-, z_+\}}} P_\lambda(\rho; w', v', z'). \tag{9.12}$$

Proof. For fixed z, each radicand is affine in the pair (w, v). For fixed w, v, it is affine in z. The corner assumptions therefore ensure nonnegative radicands throughout the box. Concavity of the square root shows first that the minimum over (w, v) is bounded below by the minimum at the four corners of that rectangle, and then that the minimum on each z interval is bounded below by its two endpoint values. This proves (9.12). □

Both occurrences of w and of v use the same values in each corner. This retains the relation between the two moving sides. All three coefficients are at most one on the enclosing box. The enclosing rectangle need not consist entirely of realizable angles. If a corner coefficient is below -1, we use the lower bound zero for the entire half-perimeter on that interval. This convention is valid without applying the corner lemma.

Lower bounds over a radius interval. The radii at an outward endpoint vary in a set of the form

$$\mathcal{R}_S = \left\{\rho : \rho_j^- \le \rho_j \le \rho_j^+ \ (j=1,2,3),\ \rho_1+\rho_2+\rho_3 = S\right\}, \qquad S>0. \tag{9.13}$$

The superscripts $-$ and $+$ denote lower and upper bounds. At an outward endpoint, $S = 1+T$. For each of the eight angular corners we lower-bound (9.11) on this same set.

Here is the supporting linear bound used for that purpose. The angular coefficients in the preceding valid corners are also at most one, by (9.8). For $-1 \le c \le 1$, put

$$N_c(x,y) = \sqrt{x^2+y^2+2cxy}.$$

It is a norm, or a seminorm when $c = \pm 1$, because its defining quadratic form is positive semidefinite. At a reference point (x_0, y_0) with $N_c(x_0, y_0) > 0$, convexity and homogeneity give

$$N_c(x,y) \ge \frac{(x_0+cy_0)x + (y_0+cx_0)y}{N_c(x_0,y_0)}. \tag{9.14}$$

If the denominator is zero, the zero linear function is used. Adding the three supporting functions and dividing by two gives a linear lower bound for the half-perimeter. To minimize it on $\mathcal{R}_S$, start with every radius at its lower bound and distribute the remaining total in increasing order of the three linear coefficients, stopping at each upper bound. An exchange of mass from a larger coefficient to a smaller one can only decrease the value, which proves this rule.

An independent lower bound comes from the three radicals directly. For example, replace x^2, y^2 by their lower interval endpoints; in $2cxy$, use the lower product when $c \ge 0$ and the upper product when $c < 0$. Replace a negative resulting radicand by zero. Take the larger of this bound and the supporting linear bound at each angular corner, and then the smallest value over the eight corners. Denote the result by P_*. The preceding lemma gives $F(\rho) \ge P_*$ uniformly on the chosen angular and radius intervals.

Adding the contribution of the origin. When $f \le 0$ throughout the angular interval, we can use the real model (9.10). Let K be the triangle formed by its three scaled tips. Its perimeter satisfies $P(K) \le 2S$, since each side is at most the sum of its endpoint radii. Suppose d_* is a uniform lower bound for the distance from the origin to the affine plane of K. Lemma 7.6 then supplies

$$G_* = \frac{d_*^3}{3((S/\pi)^2 + d_*^2)}, \qquad F(\rho) \ge P_* + G_*. \tag{9.15}$$

The gain is computed on the full angular interval and is added after the minimum over corners. If the sign condition or a positive plane-distance bound is unavailable, take $G_* = 0$.

The general plane-distance estimates are proved in Appendix C. In the present notation, let d_{ab} be the base length and put

$$b_0 = \frac{2\rho_a\rho_b qs}{d_{ab}}, \qquad \nu = \frac{(s(\rho_a+\rho_b), q(\rho_a-\rho_b))}{d_{ab}}.$$

Here b_0 is the distance to the base line and ν the unit normal pointing from the origin towards it. Given $|f| \le f_+$ and $y \le y_+$, put $Z_- = \sqrt{\max\{0, 1-f_+^2-y_+^2\}}$ and choose $J_+ \ge f_+ + y_+|\nu_2|$. Equation (C.14) gives the lower bound

$$\frac{b_0\rho_i Z_-}{\sqrt{(b_0+\rho_i J_+)^2 + (\rho_i Z_-)^2}}. \tag{9.16}$$

Use lower bounds for b_0, ρ_i, Z_- and an upper bound for J_+. The appendix also gives $3\sqrt{3}D_-/(2S^2)$ when D_- bounds $2\rho_a\rho_b\rho_i qsZ$ from below. Each estimate defaults to zero if a required positive input is unavailable; take their maximum. No degenerate area is divided out.

Excluding a closed parameter box. We finally state how the comparisons are used uniformly. Let $\mathcal{R}_1$ be a nonempty radius set of the form (9.13), with total one and positive lower radii. Fix an extended coordinate i and $T \geq 0$. Suppose the preceding construction gives a lower bound W_{end} for $F(r_1, \dots, r_i + T, \dots, r_3)$ on the entire parameter box and the chosen sign and z interval. For $0 \leq \gamma \leq 1$, define

$$M_\gamma = \max_{r \in \mathcal{R}_1} (r_i + \gamma \max\{r_j, r_k\}), \qquad N = \max_{r \in \mathcal{R}_1} (r_j + r_k),$$

where j, k are the other two indices. These are maxima of linear functions on the intersection of a box with a simplex, so they can be evaluated by the allocation rule above, in decreasing coefficient order. Let $a_- = r_i^-$ and set

$$C_- = W_{\mathrm{end}} - T - M_{h_*}, \qquad \Delta = h_* + T - W_{\mathrm{end}}.$$

If $\Delta < 0$, the box is excluded by the Lipschitz bound. Otherwise the conditions

$$W_{\mathrm{end}} > T + M_1, \qquad \max\{W_{\mathrm{end}} - T, h_0\} > N, \qquad C_- > 0 \tag{9.17}$$

make Lemma 9.1 applicable at every stationary vector in the box. Its necessary inequality implies

$$\Delta(C_- + \Delta)(a_- + T) \geq C_-(1 - h_*)a_- T. \tag{9.18}$$

Indeed, $C \geq C_-$ and $a/(a+T) \geq a_-/(a_- + T)$ in (9.2). A strict reverse inequality in (9.18) therefore excludes all stationary vectors in the box.

For a fixed base and κ, certify both signs ($\lambda = \kappa, 1/\kappa$) on closed loss intervals covering $[0, 1]$. The moving coordinate and T may vary between intervals, since the same fixed model bounds F at every radius vector. The actual configuration has one of these signs and loss values, so one certified inequality excludes it. This is the complete comparison formalized in Proposition 10.4.

10. Euclidean comparisons using the full Gram matrix

We now bound F from below for every neighboring split allowed by the profiled endpoint estimate. Keeping the three directions together leads to four consistent Gram matrices and four supporting linear forms, as used in the completed certificate [25].

From angular parameters to four matrices. Retain the ordered base a, b, opposite index i, ratio $\lambda > 0$, and parameters w, v, z from Lemma 9.2. Put $\zeta = z(w + v)$. The *Gram matrix* of unit vectors is the matrix of their pairwise inner products. In the order (a, b, i), the saturated comparison has

$$G(w, \zeta) = \begin{pmatrix} 1 & -\frac{1}{2} - w & -\frac{1}{2} + w - \zeta \\ -\frac{1}{2} - w & 1 & -\frac{1}{2} + w + \lambda\zeta \\ -\frac{1}{2} + w - \zeta & -\frac{1}{2} + w + \lambda\zeta & 1 \end{pmatrix}. \tag{10.1}$$

Conversely, every positive semidefinite matrix is a Gram matrix: diagonalize it and take square roots of its eigenvalues.

Suppose the three parameters lie in the closed intervals

$$w \in [w_-, w_+], \qquad v \in [v_-, v_+], \qquad z \in [z_-, z_+],$$

where $0 \leq w_- \leq w_+ < 1/2$, $0 \leq v_- \leq v_+ \leq 1/2$, and $0 \leq z_- \leq z_+ \leq 1$. The possible pairs (w, ζ) lie in the polygon

$$\mathcal{P} = \left\{ (w, \zeta) : \begin{array}{l} w_- \leq w \leq w_+, \\ z_-(w + v_-) \leq \zeta \leq z_+(w + v_+) \end{array} \right\}. \tag{10.2}$$

This polygon is the convex hull of the four points obtained by taking $w = w_-$ or w_+ and then taking either displayed endpoint for ζ. Some of the points may coincide.

Lemma 10.1 (Four consistent Gram matrices)**.** *Assume that the saturation condition $f \le 0$ from Lemma 9.2 holds throughout the parameter intervals, and that the four corner Gram matrices $G^{(1)}, \dots, G^{(4)}$ are positive definite. Then the saturated Euclidean comparison of Lemma 9.2 has Gram matrix*

$$G = \sum_{\nu=1}^{4} \mu_\nu G^{(\nu)}, \qquad \mu_\nu \ge 0, \qquad \sum_{\nu=1}^{4} \mu_\nu = 1,$$

and for every nonnegative radius vector ρ,

$$F(\rho) \ge W(\{0, \rho_a u_a, \rho_b u_b, \rho_i u_i\}), \qquad \langle u_\ell, u_m \rangle = G_{\ell m}.$$

The coefficients μ_ν depend only on the actual angular parameters, not on ρ.

Proof. The matrix (10.1) is affine in (w, ζ). Express the actual point of that polygon as a convex combination of its vertices. Positive definiteness is preserved by this combination. Lemma 9.2 supplies the Euclidean comparison and its mean-width bound, simultaneously for all positive radii. Continuity includes zero radii. □

The saturation check has a simple sufficient form. It is automatic when $\lambda \le 1$. When $\lambda > 1$, it follows from

$$1 - 2w_+ - (\lambda - 1) z_+ (w_+ + v_+) \ge 0.$$

Indeed $2qf = -1 + 2w + (\lambda - 1)\zeta$ in the notation of Section 9. The finite verification checks the matrix conditions as exact inequalities. A failed check leaves this comparison unavailable.

Fixing a Gaussian partition. Write Z for a standard Gaussian vector in $\mathbb{R}^3$. A *Gaussian partition* here means four disjoint measurable sets A_0, A_a, A_b, A_i whose union is $\mathbb{R}^3$ up to a set of Gaussian measure zero. Define their first-moment vectors by

$$m_j = \sqrt{2\pi}\, \mathbb{E}[Z\, \mathbf{1}_{A_j}] \qquad (j = a, b, i),$$

where $\mathbf{1}_A$ is one on A and zero outside it. Put these three vectors into the columns of a matrix M.

For a matrix, tr denotes the sum of its diagonal entries. The square root of a positive semidefinite matrix is its unique positive semidefinite square root, obtained by taking square roots of its eigenvalues.

Proposition 10.2 (Supporting forms from one Gaussian partition)**.** *Under the hypotheses of Lemma 10.1, realize each $G^{(\nu)}$ by unit vectors $u_j^{(\nu)}$. For one fixed Gaussian partition with first moments m_j, put $a_j^{(\nu)} = \langle u_j^{(\nu)}, m_j \rangle$. Then*

$$F(\rho) \ge \min_{1 \le \nu \le 4} \sum_j a_j^{(\nu)} \rho_j \qquad (\rho_j \ge 0). \tag{10.3}$$

Downward replacements of the coefficients preserve this bound.

Proof. Fix $\rho_j \ge 0$. For any unit-vector Gram matrix $G = U^{\mathsf{T}} U$, write u_j for the columns of U and $D_\rho = \operatorname{diag}(\rho_a, \rho_b, \rho_i)$. Assign each Gaussian vector in A_j to its jth linear form in (3.7), and A_0 to zero. The maximum is at least the assigned value. Rotating the partition by an orthogonal matrix O, meaning $O^{\mathsf{T}} O = I$, preserves Gaussian measure; hence

$$W(\{0, \rho_a u_a, \rho_b u_b, \rho_i u_i\}) \ge \sum_j \rho_j \langle u_j, O m_j \rangle.$$

Write $U D_\rho M^{\mathsf{T}} = V \operatorname{diag}(s_j) Q^{\mathsf{T}}$, with V, Q orthogonal and $s_j \ge 0$ the singular values. The maximum over O is $\sum_j s_j$, attained at $O = V Q^{\mathsf{T}}$. Thus

$$W(\{0, \rho_a u_a, \rho_b u_b, \rho_i u_i\}) \ge \Psi_M(\rho, G) := \operatorname{tr} \sqrt{M D_\rho G D_\rho M^{\mathsf{T}}}. \tag{10.4}$$

This lower bound is concave in G. Indeed, for positive semidefinite B,

$$\operatorname{tr} \sqrt{B} = \inf_{Y > 0} \frac{1}{2} \operatorname{tr}(Y + B Y^{-1}),$$

where $Y > 0$ means real symmetric positive definite. The inequality is $\|Y^{1/2} - Y^{-1/2}\sqrt{B}\|_{\mathrm{F}}^2 \geq 0$, with $\|C\|_{\mathrm{F}}^2 = \operatorname{tr}(C^{\mathsf{T}}C)$ the sum of squared entries. Equality holds at $Y = \sqrt{B}$ when $B > 0$ and follows by a limit otherwise. An infimum of affine functions of B is concave.

Apply this to the common combination $G = \sum_\nu \mu_\nu G^{(\nu)}$ from Lemma 10.1. Using the identity rotation at each corner,

$$F(\rho) \geq \Psi_M(\rho, G) \geq \sum_\nu \mu_\nu \Psi_M(\rho, G^{(\nu)}) \geq \sum_\nu \mu_\nu \sum_j a_j^{(\nu)} \rho_j \geq \min_\nu \sum_j a_j^{(\nu)} \rho_j.$$

Nonnegative radii permit downward rounding of every coefficient. □

The same partition is used at all corners. Concavity belongs to its lower bound Ψ_M; no concavity of mean width in the Gram matrix is assumed.

Computing the first moments. Choose a nondegenerate reference tetrahedron with vertices $P_0 = 0, P_a, P_b, P_i$ in $\mathbb{R}^3$. For each j let A_j be the set of Gaussian vectors at which $\langle Z, P_j\rangle$ is the largest of the four linear forms. Ties have Gaussian measure zero. This gives a partition, and differentiating the Gaussian maximum with respect to the three coordinates of P_j gives

$$m_j = \nabla_{P_j} W(\{0, P_a, P_b, P_i\}).$$

The reference tetrahedron may be chosen freely; it is used only to construct a lower bound.

For completeness, write $\ell_{jk} = |P_j - P_k|$. The *exterior dihedral angle* η_{jk} is the angle between the outward normals of the two faces meeting along edge jk. The mean-width and gradient formulas are

$$W(\{P_0, P_a, P_b, P_i\}) = \frac{1}{2\pi} \sum_{j<k} \ell_{jk}\eta_{jk}, \tag{10.5}$$

$$m_j = \frac{1}{2\pi} \sum_{k\neq j} \eta_{jk} \frac{P_j - P_k}{|P_j - P_k|}. \tag{10.6}$$

The first follows by integrating the support function over the unit sphere: along the boundary between the regions exposing P_j and P_k, the derivative jump is ℓ_{jk} and the boundary length is η_{jk}. Differentiation gives the second, because the classical Euclidean Schläfli identity $\sum_{j<k} \ell_{jk}\, d\eta_{jk} = 0$ cancels the derivatives of the angles [16]. In particular,

$$W(\{0, Q_a, Q_b, Q_i\}) \geq \sum_j \langle m_j, Q_j\rangle \tag{10.7}$$

for arbitrary vectors Q_j. This also follows directly from the fixed partition.

Reference vertices with rational coordinates permit directed enclosures of every m_j. Integer square-root tests and rational bounds for the dihedral angles then give the downward coefficient bounds in Proposition 10.2. Approximate calculations may choose a useful reference, but the lower forms used in the certificate are checked by exact inequalities.

Excluding every redistribution at an endpoint. Choose a moving index $\ell \in \{a, b, i\}$ and an endpoint $t > 0$. Let j, k be the other two indices. Normalize their total radius to one and write

$$\rho_\ell = t, \qquad \rho_j = \tau, \qquad \rho_k = 1 - \tau, \qquad 0 \leq \tau \leq 1.$$

Relabel the arguments in the definition of g so that the selected index ℓ is its first argument. Suppose the profiled estimate of Section 8 gives $g(t, 1) \leq C$. By the definition of g, some minimizing split τ therefore has $F(\rho(\tau)) \leq C$. This same split must obey every valid lower bound for F.

Proposition 10.3 (Finite exclusion of endpoint splits)**.** *Suppose the following bounds are valid throughout one signed angular interval:*

1. *finitely many individual linear bounds $F(\rho) \geq L_s(\rho)$;*
2. *finitely many families of four forms satisfying $F(\rho) \geq \min_\nu L_{m,\nu}(\rho)$.*

A minimizing split with $g(t,1) \le C$ must belong to

$$[0,1] \ \cap \ \bigcap_{s} \{\tau : L_s(\rho(\tau)) \le C\} \ \cap \ \bigcap_{m} \ \bigcup_{\nu=1}^{4} \{\tau : L_{m,\nu}(\rho(\tau)) \le C\}. \tag{10.8}$$

If this set is empty, the candidate stationary configuration is excluded.

Proof. Every individual lower bound is at most the actual value. For each four-form family, at least one of its forms must be at most C. These are precisely the necessary conditions in (10.8). A minimizing split exists because F is continuous on the closed interval. □

All inequalities in (10.8) are linear in τ. With rational coefficients their intervals can be computed exactly. The verifier starts with $[0,1]$, intersects with an ordinary interval or a union of four intervals, and may replace a surviving union by its smallest enclosing interval. This replacement only enlarges the possible set. An empty interval is therefore a valid exclusion, including its boundary points. Several reference partitions can be used in succession; each supplies another necessary condition on the same possible split.

One comparison position for both signs. Choosing a base and a ratio in Section 9 selects one position on the base arc. We call this choice a *common return*. It has two sign alternatives, with reciprocal oriented ratios, and in either case a loss parameter $z \in [0,1]$. The position and ratio are fixed before considering these alternatives.

Proposition 10.4 (A complete common-return comparison)**.** *Fix a base and a positive ratio for a box of original radii and angular parameters. For each sign, cover $[0,1]$ by finitely many closed loss intervals. Suppose that on each interval a valid endpoint comparison excludes every candidate in the box. The moving radius and endpoint may vary from one interval to another. Then the whole box is excluded.*

Proof. The actual Euclidean comparison at the chosen position has one of the two signs and one loss value. That value belongs to a closed interval in the corresponding cover. Its endpoint comparison then contradicts the candidate configuration. Changing the moving radius or endpoint does not change the original radial invariant, the stationary minimum, or the chosen position on the base arc. □

A subdivision of the original parameter box may give each child its own complete common-return comparison, since both children are retained. A subdivision of a conditional loss interval must retain its base, ratio, and sign, and must cover both subintervals. These elementary distinctions ensure that each local contradiction applies to the same candidate throughout its proof. They will be used explicitly in the finite verification.

11. The finite verification and the upper bound

We now assemble the proof. Set

$$\beta_* = \frac{10161}{8503}, \qquad h_* = \frac{8503}{10161}, \qquad h_0 = \frac{200}{239}.$$

The new target height is h_*; h_0 will come from a completed earlier bound. By Proposition 6.3, a counterexample would supply a positive radial minimum with

$$\sum_i r_i = 1, \qquad h := F(r) = \min_{\substack{s_i \ge 0 \\ \sum_i s_i = 1}} F(s) < h_*, \qquad F_1(r) = F_2(r) = F_3(r) = h. \tag{11.1}$$

We exclude the larger closed condition $h \le h_*$. Sorting the radii places its parameters in $\mathcal{D}$ from (7.3), which also retains zero-radius and antipodal boundaries.

The lower bound used to start the comparison. By (6.3), the completed bound $\alpha_k \le 239/200$ supplies the pointwise estimate

$$F(s) \ge \frac{200}{239} \sum_i s_i \qquad (s_i \ge 0), \tag{11.2}$$

including zero coordinates by continuity. It therefore applies at every endpoint and every neighboring split. Table 1 records the prerequisites; each row uses only earlier rows, although some tests retain weaker earlier floors.

TABLE 1. The numerical dependency chain. The middle column gives the principal lower bound for $F(s)/\sum_i s_i$ used in that stage.

Proved exponent	Earlier radial lower bound	Source
$753/625 = 1.2048$	direct geometry	[20]
$481/400 = 1.2025$	$625/753$	[20]
$1201/1000 = 1.2010$	$400/481$	[20]
$11999/10000 = 1.1999$	$1000/1201$	[20]
$599/500 = 1.198$	$10000/11999$	[20]
$11969/10000 = 1.1969$	$500/599$	[21]
$2393/2000 = 1.1965$	$10000/11969$	[22]
$2391/2000 = 1.1955$	$2000/2393$	[23]
$239/200 = 1.195$	$2000/2391$	[24]
$10161/8503$	$200/239$	[25]

The first row uses Liu's analytic bound $3\sqrt{3}/4$ for the regularity premise in Proposition 6.3 and direct geometric exclusions. Thus no row assumes its own conclusion.

From analytic comparisons to a finite cover. Table 2 summarizes the logical steps. The additional comparisons used by inherited tests are proved in Appendix C.

TABLE 2. The logical role of the analytic statements.

Required fact	Statement supplying it
A differentiable minimum of the actual radial invariant	Lemma 6.2 and Proposition 6.3
Curvature control when one radius changes	Lemma 6.4
Upper bound after optimizing the other two radii	Proposition 8.4 and Corollary 8.5
Lower bound retaining the joint angular data	Lemma 10.1 and Proposition 10.2
Exclusion of every endpoint split and both signs	Propositions 10.3 and 10.4
Fixed-neighbor, quadratic, midpoint, and retained-origin bounds	Sections 6–9 and Appendix C

A *control* specifies a base, ratio, loss intervals, moving radius, endpoint and reference tetrahedra. It is accepted only after its inequalities and hypotheses hold throughout the assigned box. Arc tests require a nonantipodal base; profiled tests require all three original angles below π, regularity and positive denominators; Gram tests require saturation and positive definiteness. A failed test is inconclusive. Guards proved on a box remain valid on its subdivisions.

Theorem 11.1 (Completed finite exclusion)**.** *Using the pointwise bound* (11.2)*, the finite certificate of* [25] *excludes every positive stationary minimum satisfying the other conditions of* (11.1) *and* $h \le 8503/10161$*. Its original closed partition has* 8 155 *leaves. All these leaves are completely covered by verified exclusions, with no unresolved region.*

Verification of the finite implication. Let $U = 2^{30}$. The initial integer box $[0, \lceil U/3 \rceil] \times [0, U/2]^4$, rescaled by $1/U$, contains $\mathcal{D}$. The checker reconstructs its entire binary partition: 5 950 leaves have ordinary records and 2 205 have signed records. These are record formats, not different tensor classes. Both closed children are retained at each bisection. A leaf at depth d has formal weight 2^{-d}, and all leaf weights sum exactly to one before clipping by the ordered radius simplex. Empty radius slices are discarded.

Each terminal proof excludes its entire box or conditional loss interval. Proposition 10.4 justifies the two subdivision rules: children of an original parameter box may choose different complete comparisons; children of a conditional interval retain its base, ratio and sign. A complete comparison covers both signs and all loss values in $[0, 1]$.

Every numerical enclosure is *directed*: lower bounds round down and upper bounds round up. Integer inequalities certify square roots, and finite series with explicit remainders certify angles and logarithms. Rational interval operations implement Proposition 10.3. Thus an accepted inequality holds on the whole closed assigned region.

The exact union of accepted exclusions covers all original leaves, with zero remaining weight. The domain audit checks the subdivisions, interval unions and their membership in the original partition. This finite closed cover includes boundaries; zero numerical volume alone would not suffice. Hence no candidate in $\mathcal{D}$ survives. □

Conclusion over every field. First let k be algebraically closed. If a spectral point had $g_\phi(1, 1, 1) > \beta_*$, Proposition 6.3 would produce the witness excluded by Theorem 11.1. Hence every spectral point has $g_\phi(1, 1, 1) \le \beta_*$. The auxiliary-factor restriction and Strassen's theorem give

$$\widetilde{\mathrm{R}}_k(T) \le n^{10161/8503} \qquad (T \in (k^n)^{\otimes 3}).$$

For an arbitrary field, pass to its algebraic closure and use Lemma 2.7. The cases $T = 0$ and $n = 1$ follow directly from the rank definition. Together with the lower bound from Section 4, this proves Theorem 1.1.

Corollary 11.2 (Largest flattening rank)**.** *Let T be a nonzero three-tensor over any field, and let d be the largest of its three flattening ranks. Then*

$$\widetilde{\mathrm{R}}_k(T) \le d^{10161/8503}.$$

Proof. The image of the flattening map into each factor is the smallest subspace of that factor needed to express T. Restrict to these three subspaces, and then pad them to dimension d. Inclusions and projections give restrictions in both directions, so asymptotic rank is unchanged. Apply the upper bound in format $d \times d \times d$. □

Reproduction and scope. The certificate is frozen at revision `97958f33` in [25]. Its manifest has 1 520 inputs. A fresh execution on 26 November 2026 rebuilt all five native libraries with fatal undefined-behavior checks and passed all 16 stages. A separate audit also passed: with search disabled, it reconstructed coverage and checked 7 169 saved repair trees and 1 313 static controls.

The companion contains the pinned sources, prerequisite and execution records, and mathematical audits. The two executions share the exact geometric predicates. This is an audited computational proof with the stated mathematical inputs, rather than a proof-assistant formalization.

Appendix A. Limits and radial regularity

We supply the two technical proofs used in the main reduction.

A.1. **Compatibility in the tangent limit.**

Proof of Lemma 5.2. For $T \neq 0$ with index dimensions q_i, joint convexity makes the quotient in (5.4) nondecreasing in $t > 0$. Transport bounds it between zero and $\sum_i |v_i| \log_2 q_i$, so the limit exists. Rescaled geodesic distances tend to the cone distances; transport therefore makes the value independent of the representing geodesics and continuous on the completion. Replacing t by ct proves homogeneity.

The represented vectors embed isometrically in a pointed ultralimit of $(X, t^{-1}d, t^{-1}H, o)$; complete this embedding to obtain C. Its inherited values are precisely (5.4). On every finite support of positive diameter, the diagonal value is bounded between that diameter and a spanning-tree length. Thus the normalization in Definition 2.6 survives. The complete geodesic image of C is convex in the CAT(0) limit space, by uniqueness of geodesics, so the values remain jointly convex. □

A.2. **The radial regularity criterion.**

Proof of Lemma 6.2. Duality and a logarithmic lower bound. Use the profile $P = P_r$ at the fixed placement.

At each interior point P has a supergradient, and all its coordinates are positive since P grows unboundedly in each coordinate. Consequently (6.5) gives

$$P(x) = \inf_{f_i>0} (Q(f) - D + f \cdot x). \tag{A.1}$$

This duality does not presume an optimizer at every cost. Write $F'(r; v)$ for the one-sided directional derivative. Differentiating (6.6) gives

$$Q'(\mathbf{1}; v) = F'(r; -D_r v).$$

Homogeneity and the Lipschitz bound dominate the quotient by a constant times $(1 + t)^{-2}$, whose integral is one. Using $F'(r; v) = \max_{p \in \partial F(r)} p \cdot v$, and similarly for Q, we obtain

$$\partial Q(\mathbf{1}) = -D_r \, \partial F(r). \tag{A.2}$$

Set $g_i = r_i p_i > 0$. Since $\sum_i g_i = D$, duality gives $P(g) = Q(\mathbf{1})$ and makes g a maximizer of $P(x) - \sum_i x_i$. Thus $\mathbf{1}$ is a supergradient at g. Insert $F(z) \geq p \cdot z$ into the energy identity to get

$$Q(f) - Q(\mathbf{1}) \geq -\sum_i g_i \ln f_i.$$

Minimizing in (A.1) yields

$$P(x) \geq P(g) + \sum_i g_i \ln(x_i/g_i). \tag{A.3}$$

A strict quadratic remainder. Define the nonnegative convex remainder

$$R(v) = P(g) + \sum_i v_i - P(g + v), \qquad g + v > 0,$$

and put $v_i = \sqrt{g_i}\, w_i$, $\psi(w) = 2R(v)$. Taylor expansion of (A.3) gives $\psi(w) \leq |w|^2 + O(|w|^3)$. Let $\mathbf{e}_i$ denote the ith standard coordinate vector in the width space $\mathbb{R}^3$. The one-coordinate concavity of $\exp(P/r_i)$, with the supporting slope supplied by $\mathbf{1}$, gives

$$P(g + \sqrt{g_i}\, t\mathbf{e}_i) - P(g) \leq r_i \ln(1 + \sqrt{g_i}\, t/r_i).$$

For both signs of small t, consequently,

$$\psi(t\mathbf{e}_i) \geq p_i t^2 - O(|t|^3).$$

To combine these bounds, choose $0 < \theta < 1$ so that $S = \sum_i [1 - (1 - \theta)p_i] < 1$, and put $s_i = 1 - (1 - \theta)p_i \geq \theta$. Given w, choose i maximizing w_i^2/s_i, set $t = \theta w_i/s_i$, and $w' = (t\mathbf{e}_i - \theta w)/(1 - \theta)$. Convexity at $t\mathbf{e}_i = \theta w + (1 - \theta)w'$ gives

$$\psi(w) \geq \frac{\theta}{1-\theta} \left(\frac{w_i^2}{s_i} - |w|^2 \right) - O(|w|^3) \geq \frac{\theta}{1-\theta} (S^{-1} - 1)|w|^2 - O(|w|^3).$$

Thus $R(v) \geq c|v|^2$ near zero for some $c > 0$.

Another subgradient p' would give positive subgradients $p_\varepsilon = (1-\varepsilon)p + \varepsilon p'$ and distinct dual maximizers $D_r p_\varepsilon \to g$, where R vanishes. This contradicts the strict quadratic bound. The subgradient is unique. Compactness in $[0,1]^3$ and upper semicontinuity keep all nearby subgradients positive with sum greater than two. The same argument gives differentiability there; convexity makes the gradient continuous.

Local Lipschitz continuity of the gradient. The quadratic constants are uniform on a smaller neighborhood. Convexity extends $R(v) \geq c|v|^2$, valid for $|v| \leq \varepsilon$, to $R(v) \geq c\varepsilon|v|$ outside that ball. For small b, duality therefore gives

$$Q(\mathbf{1}+b) - Q(\mathbf{1}) + g \cdot b = \sup_{g+v>0} \{-b \cdot v - R(v)\} \leq \frac{|b|^2}{4c}. \tag{A.4}$$

Write

$$B_F(z) = F(r+z) - F(r) - p \cdot z \geq 0, \qquad z_t = D_r \left[\left(\mathbf{1} + \frac{b}{1+t} \right)^{-1} - \mathbf{1} \right],$$

where inversion is coordinatewise. Equation (6.6) gives exactly

$$Q(\mathbf{1}+b) - Q(\mathbf{1}) + g \cdot b = \int_0^\infty \frac{B_F(z_t)}{1+t}\, dt + \sum_i g_i(b_i - \ln(1+b_i)).$$

The last sum is nonnegative. Set $b_i = -2z_i/r_i$. For $0 \leq t \leq 1$, $z_t = 2z/(1+t) + O(|z|^2)$. The continuous gradient makes the error in the corresponding remainder $o(|z|^2)$, uniformly in this interval. Convexity gives $B_F(\lambda z) \geq \lambda B_F(z)$ for $\lambda \geq 1$, and $\int_0^1 2(1+t)^{-2} dt = 1$. Thus (A.4) bounds $B_F(z)$ by a locally uniform constant times $|z|^2$. For a sufficiently large C, the function $F(r) - C|r|^2$ therefore has an affine upper support at every point of a smaller convex neighborhood, and is concave there. This semiconcavity, together with convexity of F, gives the asserted $C^{1,1}$ regularity.

The inverse-matrix inequality. At almost every point F has a second-order Taylor expansion. Expanding (6.6) at unit costs gives the Hessian $Q''(\mathbf{1}) = A$ in (6.7); the factor $1/2$ comes from $\int_0^\infty (1+t)^{-3} dt = 1/2$. The matrix is positive definite because $g_i > 0$ and $\nabla^2 F$ is positive semidefinite. Dual minimization in (A.1) gives

$$P(g+v) = P(g) + \sum_i v_i - \tfrac{1}{2} v^{\mathsf{T}} A^{-1} v + o(|v|^2).$$

Indeed the positive quadratic form localizes the minimizing cost perturbation to size $O(|v|)$, after which ordinary quadratic minimization applies. Comparison with the preceding one-coordinate logarithmic upper bound proves (6.8). □

Appendix B. Regularity of the optimized profile

We complete Lemma 8.2, using its definition of a smooth lower test.

Lemma B.1 (Transfer to a smooth lower test)**.** *Suppose F is $C^{1,1}$ near a positive triple r and satisfies the inverse inequality (6.8) almost everywhere there. If a smooth function ψ touches F from below at r and*

$$A_\psi = \operatorname{diag}(r_i F_i(r)) + \tfrac{1}{2} D_r \nabla^2 \psi(r) D_r$$

is positive definite, then $(A_\psi^{-1})_{ii} \geq 1/r_i$ for each i.

Proof. Choose a closed ball about r on which $F - \psi \geq 0$. For $\varepsilon > 0$ and sufficiently small vectors v, minimize on this ball the function

$$F(x) - \psi(x) + \varepsilon|x-r|^2 - v \cdot (x-r).$$

Its minima lie in the interior. Let N be the null set where the second-order expansion of F or its inverse inequality is unavailable. The gradient of $F - \psi + \varepsilon|x-r|^2$ is Lipschitz, so its

image of N is a null set. Choose $v_n \to 0$ outside that image, and let x_n be a corresponding minimizer. Then $x_n \notin N$. Moreover $\varepsilon|x_n - r|^2 \le |v_n||x_n - r|$, so $x_n \to r$. At these minima,

$$\nabla^2 F(x_n) \succeq \nabla^2 \psi(x_n) - 2\varepsilon I.$$

Here $B \succeq C$ means that $B - C$ is positive semidefinite. Form the weighted matrices using the same radii x_n and the actual gradient of F at x_n. For sufficiently small ε and sufficiently large n, the matrix formed from the right side is positive definite and is at most the actual matrix $A_F(x_n)$. Inverting positive definite matrices reverses this order. Its inverse diagonal is therefore at least $1/(x_n)_i$. First let $n \to \infty$, then let $\varepsilon \downarrow 0$. Continuity of inversion proves the claim. □

The rank-one calculation in (8.9) now applies even if all attaining splits lie in a set of three-dimensional measure zero.

At a boundary split the upper contact is

$$f_k(t,z) = \frac{t + z + D}{2}, \qquad D = \sqrt{t^2 + z^2 + 2ktz}, \qquad \tfrac{1}{2} \le k < 1.$$

Its derivatives are

$$(f_k)_t = \frac{1 + (t + kz)/D}{2}, \qquad (f_k)_{tt} = \frac{(1 - k^2)z^2}{2D^3}.$$

Since $D \ge t + kz$ and $D \ge t$,

$$\frac{t(f_k)_{tt}}{2(1 - (f_k)_t)} = \frac{t(D + t + kz)}{2D^2} \le 1.$$

A smooth lower test for g is also a lower test for f_k at this contact. Its first derivative equals $(f_k)_t = g_t$ and its second derivative is at most $(f_k)_{tt}$. This proves the lower-test version of (8.6) at a boundary split as well.

On a compact positive interval, $0 \le g_t \le 1$ and $g - t > z/2$ bound the right side of (8.6) by a constant C. Thus every lower test for $g(t,z) - Ct^2/2$ has nonpositive second derivative. This function is concave: otherwise a shallow strictly convex quadratic through two chord endpoints lies above it somewhere inside. Translating that quadratic down to its first lower contact gives an interior contact of positive curvature, a contradiction. Convexity of g and this semiconcavity make g_t locally Lipschitz. At a second-order point, subtracting a small positive quadratic from the Taylor polynomial and then letting its coefficient tend to zero proves (8.6).

Appendix C. Supplementary comparisons

This appendix supplies the short-arc and scalar estimates used alongside the comparisons of Sections 8 and 10. Fix a trial exponent $\beta > 1$ and a positive minimizing radius vector r of sum one with $F(r) \le h_* := 1/\beta$. Write $p_{ij} = d_{ij}(r)$ as in (7.2).

C.1. **Midpoints and short arcs.** Choose an ordered base i, j with $\theta = \theta_{ij} < \pi$, and let k be the remaining index. Put

$$q = \cos(\theta/2) = \sqrt{\frac{1}{4} - \frac{1}{2}w_{ij}}, \qquad s = \sin(\theta/2) = \sqrt{\frac{3}{4} + \frac{1}{2}w_{ij}}, \tag{C.1}$$
$$A_i = w_{ij} + w_{ik}, \qquad A_j = w_{ij} + w_{jk}.$$

Rotate the models of Lemma 7.4 to fix the base:

$$0, \quad r_i(q, -s, 0), \quad r_j(q, s, 0), \quad r_k(\xi, Y, \sqrt{1 - \xi^2 - Y^2}). \tag{C.2}$$

Denote their mean width by $\Psi^{ij}(\xi, Y)$. Every supplied model has a real square root and satisfies

$$\Psi^{ij}(\xi, Y) \le h_*. \tag{C.3}$$

The two moving sides contract precisely when their dot products increase:

$$q\xi - sY \ge -\frac{1}{2} - w_{ik}, \qquad q\xi + sY \ge -\frac{1}{2} - w_{jk}. \tag{C.4}$$

In particular, raising a nonpositive ξ towards zero, with Y fixed and unit norm restored, contracts these sides. Gaussian comparison then decreases the width. We use this replacement below.

Let $\gamma(t)$ be the unit-speed direction-space arc from e_i to e_j, and put $t_0 = \theta/2$ and $f(t) = \cos d_\Sigma(\gamma(t), e_k)$. Stationarity gives

$$f(t) \le -2q\cos(t - t_0). \tag{C.5}$$

The *midpoint slack* is $\sigma = -2q - f(t_0) \ge 0$. At regular parameters, rotation through $\eta = t - t_0$ gives

$$\xi = f(t)\cos\eta - f'(t)\sin\eta, \qquad Y = f(t)\sin\eta + f'(t)\cos\eta. \tag{C.6}$$

A convergent sequence of such models approaching t_0 gives a midpoint model with $\xi = -2q - \sigma$ and width at most h_*. Its dot-product constraints become

$$-A_j + q\sigma \le sY \le A_i - q\sigma. \tag{C.7}$$

Lemma C.1 (An upper bound for the slack)**.** *If $\sigma \ge \tau \ge 0$, there is a real Y with*

$$\frac{-A_j + q\tau}{s} \le Y \le \frac{A_i - q\tau}{s}, \qquad (-2q - \tau)^2 + Y^2 \le 1, \qquad \Psi^{ij}(-2q - \tau, Y) \le h_*. \tag{C.8}$$

Excluding this set proves $\sigma < \tau$; at $\tau = 0$ it excludes the configuration.

Proof. Raise the midpoint coordinate from $-2q - \sigma$ to $-2q - \tau$. This preserves a real unit direction and decreases the width; (C.7) gives the asserted interval. □

Lemma C.2 (Selection on a short arc)**.** *For $0 < \ell < \tan t_0$, some regular parameter with $0 \le \tan(t - t_0) \le \ell$ satisfies*

$$Y \le \frac{\sigma}{\ell}, \qquad \xi \le -2q + \ell(-Y)_+, \qquad (-Y)_+ = \max\{-Y, 0\}. \tag{C.9}$$

Proof. Set $u = \tan(t - t_0)$ and $g(u) = f(t_0 + \arctan u)/\cos(t - t_0)$. Then $g' = Y$ almost everywhere, $\xi = g - uY$, $g \le -2q$, and $g(0) = -2q - \sigma$. Absolute continuity on $[0, \ell]$ forces $Y \le \sigma/\ell$ at some regular point: otherwise integration would give $g(\ell) > -2q$. The identity for ξ proves the second inequality. □

Lemma C.3 (The two signs of the transverse coordinate)**.** *Suppose $\sigma \le \bar\sigma$ and*

$$0 < \ell < \tan t_0, \qquad s - q\ell > 0, \qquad \frac{\ell A_j}{s - q\ell} \le 2q.$$

At least one of the following sets is nonempty:

$$0 \le Y \le \min\left(\frac{A_i}{s}, \frac{\bar\sigma}{\ell}\right), \quad 4q^2 + Y^2 \le 1, \quad \Psi^{ij}(-2q, Y) \le h_*; \tag{C.10}$$

$$-\frac{A_j}{s - q\ell} \le Y \le 0, \quad (-2q - \ell Y)^2 + Y^2 \le 1, \quad \Psi^{ij}(-2q - \ell Y, Y) \le h_*. \tag{C.11}$$

Proof. Use Lemma C.2. If $Y \ge 0$, its inequalities and (C.4) give the first interval; raise ξ to $-2q$. If $Y = -v \le 0$, the same inequalities give $(s - q\ell)v \le A_j$. The hypotheses allow us to raise ξ to $-2q + \ell v \le 0$. Both replacements decrease the width. □

Thus both signs must be excluded for one admissible ℓ. Reversing the base examines the other side of the midpoint.

C.2. **Width and perimeter bounds.** In these tests $\xi = b + cY$ is affine in Y, and $z = \sqrt{1 - (b + cY)^2 - Y^2}$. The dihedral formula (10.5) and the linear supports (10.7) bound the width from below. If the coefficient of $r_k z$ in a chosen support is nonnegative, that support is concave in Y on each valid interval, so the smaller endpoint value is a lower bound. The outer parameters are still bounded uniformly on their box.

A bound that also covers degenerate models is

$$W(0, P_1, P_2, P_3) \ge W(0, Q_1, Q_2, Q_3) - \sum_{a=1}^{3} e_a, \qquad \|P_a - Q_a\| \le e_a. \tag{C.12}$$

Indeed, the second hull lies in the first plus the segments $[0, Q_a - P_a]$. This addition takes all sums of one point from each set. Mean width is monotone under inclusion and additive under set addition, and a segment has width equal to its length.

For perimeter bounds, write $P = p_{ij} + p_{ik} + p_{jk}$. At the actual midpoint the dot-product increases are

$$\ell_i = A_i - q\sigma - sY, \qquad \ell_j = A_j - q\sigma + sY, \qquad \ell_i + \ell_j = B - 2q\sigma \geq 0, \quad B = A_i + A_j.$$

The moving lengths L_{uk} satisfy

$$L_{uk}^2 = p_{uk}^2 - 2r_u r_k \ell_u, \qquad \frac{p_{uk} - L_{uk}}{2} = \frac{r_u r_k \ell_u}{p_{uk} + L_{uk}} \quad (u = i, j).$$

If c_{loss} bounds both coefficients and G is a uniform lower bound for the retained-origin gain, then

$$h_* \geq \frac{P}{2} - c_{\mathrm{loss}}(B - 2q\sigma) + G. \tag{C.13}$$

Only positive denominators are used; the first coordinate gives $L_{uk} \geq q(r_u + 2r_k)$ at an actual or saturated midpoint.

Alternatively, on $\ell_i, \ell_j \geq 0$, $\ell_i + \ell_j \leq L$, the sum

$$\sqrt{p_{ik}^2 - 2r_i r_k \ell_i} + \sqrt{p_{jk}^2 - 2r_j r_k \ell_j}$$

is concave and decreasing in each loss. If both full-loss radicands are nonnegative, its minimum occurs at $(L, 0)$ or $(0, L)$. Such auxiliary endpoints need not be realized: the gain added to their perimeter bound must hold throughout the actual model family. The same endpoint rule applies when upper dot products are affine in Y and the radicands remain nonnegative. An endpoint upper dot product below -1 may be raised to -1, which only lowers the resulting bound.

For an upper dot product $\alpha \geq -1$, put $\zeta = |r_u - r_v|/(r_u + r_v)$. The side ratio satisfies

$$\left(\frac{L_{uv}}{r_u + r_v}\right)^2 = \frac{2 - 2\alpha + (2 + 2\alpha)\zeta^2}{4}.$$

Equivalently $\alpha = -1/2 - a$, where $a = w_{uv} - \ell$ is the effective angular weight. The ratio decreases with α and increases with ζ^2. After checking nonnegative radicands and positive denominators, these monotonicities give uniform side-ratio bounds. The resulting linear form in the radii is minimized by allocating the available mass to its smallest coefficients, subject to the radius intervals.

At a saturated midpoint, write $w = w_{ij}$. On the half $0 \leq Y \leq (w + w_{ik})/s$, both moving weights start at $-w$; at the far endpoint ik has its original weight and jk has weight $-(2w + w_{ik})$. The negative half exchanges i, j. On either half,

$$z^2 \geq \left[2w - \frac{(w + w_f)^2}{3/4 + w/2}\right]_+, \qquad [a]_+ = \max\{a, 0\},$$

where w_f is the original weight of the unchanged endpoint edge. The bracket increases with w and decreases with w_f on $[0, 1/2]^2$. Each half uses its own perimeter and gain bounds, and both must be excluded.

A slack refinement assumes $\sigma \geq \tau$ and uses $L = B - 2q\tau$ as a total-loss upper bound, with outward rounding. Negative L is impossible. Otherwise the endpoint test applies, using the gain justified by the previous upper bound on σ. A contradiction proves $\sigma < \tau$.

C.3. **Plane-distance bounds.** For the vertices in (C.2), let K be their nondegenerate triangle, P_K its perimeter, and d its plane distance from the origin. Write $\widehat{p}_{ij} = \|x_i - x_j\|$. The base-line distance and its unit normal towards the line are

$$b = \frac{2r_i r_j q s}{\widehat{p}_{ij}}, \qquad \nu = \frac{(s(r_i + r_j), q(r_i - r_j))}{\widehat{p}_{ij}}.$$

For $|\xi| \le f_+$, $|Y| \le y_+$, put $J = f_+ + y_+|\nu_y|$. The determinant of the tips is $\widehat{p}_{ij} b r_k z$, while their double area is $\widehat{p}_{ij}\sqrt{(b - r_k(\xi\nu_x + Y\nu_y))^2 + (r_k z)^2}$. Taking their quotient gives

$$d \ge D(b, r_k, z, J) := \frac{b r_k z}{\sqrt{(b + r_k J)^2 + (r_k z)^2}}. \tag{C.14}$$

This increases with b, r_k, z and decreases with J, since

$$D^{-2} = b^{-2} + \frac{(r_k^{-1} + J/b)^2}{z^2}.$$

A vanishing lower input gives the bound zero.

The same calculation applies without the normalization $\sum r_i = 1$. If $P_K \le 2S$ and $\mathcal{D} = |\det(x_i, x_j, x_k)| = 2 r_i r_j r_k q s z$, then

$$2\,\mathrm{area}(K) \le \frac{P_K^2}{6\sqrt{3}} \le \frac{2S^2}{3\sqrt{3}}, \qquad d \ge \frac{3\sqrt{3}\,\mathcal{D}}{2S^2}.$$

The area inequality follows from Heron's formula: with side lengths a, b, c and $t = P_K/2$, the three factors $t - a, t - b, t - c$ sum to t, so their product is at most $(t/3)^3$. One may take $S = \sum r_i$ at any endpoint, or $S = h_*$ for the normalized models satisfying (C.3).

At an actual midpoint, (C.7) and $\sigma \le \bar{\sigma}$ imply

$$z^2 \ge \left[1 - (2q + \bar{\sigma})^2 - \left(\frac{\max\{A_i, A_j\}}{s}\right)^2\right]_+.$$

For a saturated midpoint or signed-window replacement with $|\xi| \le 2q$, use instead $z^2 \ge [2w_{ij} - y_+^2]_+$.

A sharper area estimate retains one short side. If A_1, A_2, C are the squared sides, $\overline{S} \ge A_1 + A_2$, $\overline{C} \ge C$, and $C_* = \min(\overline{C}, \overline{S})$, Heron's identity gives

$$16\,\mathrm{area}(K)^2 = 2C(A_1 + A_2) - (A_1 - A_2)^2 - C^2 \le 2C_*\overline{S} - C_*^2.$$

Here C_* maximizes $2x\overline{S} - x^2$ on $[0, \overline{C}]$. A positive determinant lower bound D_- therefore gives

$$d \ge \frac{2D_-}{\sqrt{2C_*\overline{S} - C_*^2}}. \tag{C.15}$$

If no positive determinant bound is available, use zero.

For the saturated model of Lemma 9.2, put $\zeta = z(w + v)$ and denote its third unit vector by (f, y, Z). The formulas in (9.6) and (9.7) yield

$$Z^2 = 2w + 2(\lambda - 1)\zeta - \zeta^2\left(\frac{(\lambda - 1)^2}{4q^2} + \frac{(\lambda + 1)^2}{4s^2}\right).$$

This identity and $Z^2 = 1 - f^2 - y^2$ give two lower bounds; their maximum can be used in $\mathcal{D} = 2 r_a r_b r_i q s Z$.

On a redistribution rectangle the radii are $(a + T, (1 - a)\rho, (1 - a)(1 - \rho))$, with fixed T, $0 \le a, \rho \le 1$, and $a + T \ge 0$ throughout. Their product is minimized at a corner: $\rho(1 - \rho)$ is concave, and $(a + T)(1 - a)^2$ has derivative $(1 - a)(1 - 3a - 2T)$ and no interior minimum. After checking $|c_{\ell m}| \le 1$ at every angular corner, each squared side $r_\ell^2 + r_m^2 + 2c_{\ell m} r_\ell r_m$ is a nonnegative quadratic form, separately convex in a, ρ and affine separately in w, v, z. Its upper bound, and that of $A_1 + A_2$, therefore occur at consistent radius and angular corners. For the sum $A_1 + A_2$, both sides are evaluated at one corner before taking the maximum. These observations justify the determinant and area inputs to (C.15).

C.4. **Variance and rational bounds for the retained-origin gain.** Let x_1, x_2, x_3 form a nondegenerate comparison triangle, and write $W = W(\{0, x_1, x_2, x_3\})$, so that $F(r) \geq W$. In this subsection P denotes its *half-perimeter*, $d > 0$ its plane distance, and $R_0 \geq \max_i \|x_i\|$. Let h be its planar support function relative to the projected origin, $\mathbb{E}$ the circle average, and $\mu = \mathbb{E}h = P/\pi$. Equation (7.14) gives $W = P + \mathbb{E}J_d(h)$. We refine this lower bound using the variance of h and a rational estimate of J_d, as required by the certificate [25].

Lemma C.4 (Triangle support variance)**.** *For any nondegenerate planar triangle and any planar origin,* $\mathbb{E}(h-\mu)^2 \geq \mu^2/32$.

Proof. The directions supported at vertex i form an interval of length $2u_i$, where $0 < u_i < \pi/2$ and $\sum_i u_i = \pi$. Centering the angular coordinate on that interval gives $h = A_i \cos\theta + B_i \sin\theta$. Put $c(u) = u + \sin u \cos u$ and $Q(u) = 4\sin^2 u/c(u)$. Integration and Cauchy–Schwarz give, for either sign of A_i,

$$2\pi\mu = \sum_i 2A_i \sin u_i, \qquad \int_0^{2\pi} h^2 \geq \sum_i A_i^2 c(u_i) \geq \frac{(2\pi\mu)^2}{\sum_i Q(u_i)}.$$

The function Q is concave: the derivative

$$(2u - Q(u))' = 2\left(\frac{u - \sin u\cos u}{u + \sin u \cos u}\right)^2$$

is nondecreasing because $\sin(2u)/(2u)$ decreases on $[0, \pi/2]$. Thus $\sum_i Q(u_i) \leq 3Q(\pi/3) = 9/(\pi/3 + \sqrt{3}/4)$, whence

$$\mathbb{E}(h-\mu)^2 \geq \left(\frac{2\pi^2}{27} + \frac{\pi\sqrt{3}}{18} - 1\right)\mu^2 > \frac{\mu^2}{32}.$$

The last inequality follows already from $\pi > 157/50$ and $\sqrt{3} > 433/250$. □

Lemma C.5 (A rational retained-origin bound)**.** *For* $0 < p \leq P$, $0 < \delta \leq d$, *set*

$$A = (50p/157)^2, \quad B = \delta^2, \quad D_0 = R_0^2 - \frac{2}{3}\left(\frac{4}{27} - (50/157)^2\right)p^2.$$

Then $D_0 > 0$ *and*

$$W \geq p + \mathcal{B}(p,\delta) + \frac{\delta^3(7p/22)^2}{32D_0^2}, \tag{C.16}$$

$$\mathcal{B}(p,\delta) = \frac{\delta^3(5005A^2 + 5082AB + 969B^2)}{35(429A^3 + 693A^2B + 315AB^2 + 35B^3)}. \tag{C.17}$$

Replacing D_0^2 *by* R_0^4 *gives a weaker valid correction. When* $p = 0$ *or* $\delta = 0$, *both corrections may be set to zero; for degenerate triangles use the half-perimeter alone.*

Proof. Write $\eta = h - \mu$ and $V = \mathbb{E}\eta^2$. Taylor's formula and (7.15) give

$$\mathbb{E}J_d(h) - J_d(\mu) = 2d^3\mathbb{E}\left[\eta^2 \int_0^1 \frac{1-t}{((\mu + t\eta)^2 + d^2)^2}\,dt\right].$$

The measure $2(1-t)\eta^2\,dt\,d\theta/(2\pi V)$ has mass one. The projected vertices y_i have norm at most $\rho = \sqrt{R_0^2 - d^2}$. For their side lengths a, b, c,

$$4P^2 \leq 3(a^2+b^2+c^2) = 9\sum_i \|y_i\|^2 - 3\left\|\sum_i y_i\right\|^2 \leq 27\rho^2.$$

Consequently, with $\kappa = 4/27 - 1/\pi^2 > 0$, $\mu^2 + d^2 \leq R_0^2 - \kappa P^2$ and $h^2 + d^2 \leq R_0^2$. Convexity of the square bounds the Taylor denominator before squaring by $R_0^2 - (1-t)\kappa P^2$. The above measure gives $1-t$ mean $2/3$; Jensen's inequality for $s \mapsto s^{-2}$ yields

$$\mathbb{E}J_d(h) - J_d(\mu) \geq \frac{d^3 V}{[R_0^2 - \frac{2}{3}\kappa P^2]^2}.$$

The bracket lies between $R_0^2/3$ and D_0, so in particular $D_0 > 0$. The variance lemma and $P \geq p$, $d \geq \delta$, $157/50 < \pi < 22/7$ now give the last term of (C.16).

For the other term, $P + J_d(P/\pi)$ increases with P, d: (7.15) gives monotonicity in d, and, for $s > 0$,

$$J_d'(s) = -\arctan(d/s) + \frac{ds}{d^2+s^2} \geq -\frac{\pi}{2}.$$

Its P-derivative is therefore at least $1/2$. This permits simultaneous replacement of P, d by their lower bounds p, δ. We now bound $J_\delta(p/\pi)$. The positive-support integral of Lemma 7.6 gives

$$J_\delta(p/\pi) = \delta^3 \int_0^1 \frac{t^2}{(p/\pi)^2 + \delta^2 t^2}\, dt \geq \delta^3 \int_0^1 \frac{t^2}{A + Bt^2}\, dt.$$

The identity

$$\frac{1}{A+Bt^2} - 2q(t) + (A+Bt^2)q(t)^2 = \frac{[1-(A+Bt^2)q(t)]^2}{A+Bt^2} \geq 0$$

gives a rational lower bound by choosing $q(t) = c_0 + c_1 t^2 + c_2 t^4$. With

$$v_i = \frac{1}{2i+3}, \qquad K_{ij} = \frac{A}{2i+2j+3} + \frac{B}{2i+2j+5} \quad (0 \leq i,j \leq 2),$$

integration against t^2 gives $2v^{\mathsf{T}}c - c^{\mathsf{T}}Kc$. The matrix K is positive definite, since its quadratic form is $\int_0^1 t^2(A+Bt^2)q(t)^2\, dt$. Maximization gives $v^{\mathsf{T}}K^{-1}v$; solving this three-dimensional system gives exactly $\mathcal{B}(p,\delta)/\delta^3$. □

The isolated gain $\mathcal{B}(p,\delta)$ decreases with p and increases with δ. Indeed, writing $K = AK_0 + \delta^2 K_1$, both K_0, K_1 are positive definite by the same integral argument. Increasing A decreases K^{-1} in the positive-semidefinite order, whereas

$$\delta^3 K^{-1} = \delta\,[K_1 + (A/\delta^2)K_0]^{-1}$$

increases with δ. Thus if $P \leq p^+$ and $P - Q(r) \geq c$, then

$$W - Q(r) \geq c + \mathcal{B}(p^+,\delta) + \frac{\delta^3(7p/22)^2}{32D_0^2}.$$

The isolated gain uses the *upper* half-perimeter bound p^+; the variance correction uses the lower bound p. One available upper bound is $P \leq \sum_i \|x_i\|$.

References

[1] J. Alman and B. Li, *Asymptotic rank speedup theorems, revisited*, CCC 2026, LIPIcs **383**, art. 36, 36:1–36:42, Theorem 5 and Problem 41.

[2] J. Alman, B. Li and K. Pratt, *The edge of the asymptotic spectrum of tensors*, arXiv:2604.01386v2 (2026), Theorem 1.1 and Corollary 8.7.

[3] A. Amanov, M. van den Berg and J. Zuiddam, *Multiplying matrices via polynomials: exposition and barriers*, arXiv:2610.30054v1 (2026), especially Proposition 4.14.

[4] M. R. Bridson and A. Haefliger, *Metric Spaces of Non-Positive Curvature*, Grundlehren der mathematischen Wissenschaften **319**, Springer, 1999.

[5] P. Bürgisser, M. Clausen and M. A. Shokrollahi, *Algebraic Complexity Theory*, Grundlehren der mathematischen Wissenschaften **315**, Springer, 1997.

[6] S. Chatterjee, *An error bound in the Sudakov–Fernique inequality*, arXiv:math/0510424v1 (2005), Theorem 1.1 and Section 2.

[7] M. Christandl, K. Hoeberechts, H. Nieuwboer, P. Vrana and J. Zuiddam, *Asymptotic tensor rank is characterized by polynomials*, Proceedings of STOC 2025, Theorem 2. Author manuscript.

[8] Claude (Anthropic), *VP ≠ VNP over every field*, hexagon:2611.00954v1, 23 November 2026, especially Theorem C.1. Pinned edition, repository revision `11b6c7e3`.

[9] V. Kapovitch and C. Ketterer, *CD meets CAT*, arXiv:1712.02839v3, 20 January 2018 edition, Definition 2.60 and Theorem 2.61. Author manuscript.

[10] P. Kaski and M. Michałek, *A universal sequence of tensors for the asymptotic rank conjecture*, ITCS 2025, LIPIcs **325**, art. 64, 64:1–64:24.

[11] Y. Liu, *A uniform exponent bound for asymptotic tensor rank*, hexagon:2611.00972v1, 24 November 2026, Theorem 1.1.

[12] OpenAI, *Euclidean geometry in the asymptotic spectrum of tensors*, 24 November 2026 edition. Pinned preprint, OpenAI Math repository revision `92b516ef`; Theorems 1.1, 3.1, 4.3, 5.3, Corollary 3.4 and Section 4.4.
[13] V. Strassen, *The asymptotic spectrum of tensors*, J. reine angew. Math. **384** (1988), 102–152.
[14] V. Strassen, *Algebra and complexity*, First European Congress of Mathematics, vol. II, Progr. Math. **120**, Birkhäuser, 1994, 429–446.
[15] Asymptotic Rank Research Swarm, *Dated progress records*, 22–25 November 2026. 22 November report; PR 1131, PR 1145, PR 1149, PR 1150, PR 1152. Dates refer to recorded results, not a priority claim against independent work.
[16] Asymptotic Rank Research Swarm, Groups B and C, *Tangent, recentering, retained-apex, and affine-width proofs*, 24–25 November 2026. Source collection.
[17] OpenAI, *The matrix multiplication exponent is nine fourths*, 24 November 2026 edition. Pinned preprint, revision `3344da59`; the continuum profile and radial energy construction in `build/contact.tex`.
[18] Asymptotic Rank Research Swarm, Group B, *Radial diagonal values, multiplier profiles and curvature*, 24–25 November 2026. Pinned source collection, revision `d78a03a8`, notes C989–C990, C1004–C1006, C1008, C1011–C1012 and C1020. Exact files and source identities are listed in the companion's primary-source manifest.
[19] Asymptotic Rank Research Swarm, *Radial-simplex minimization and regularity without an essentiality premise*, 25 November 2026. Pinned proof, revision `fd52d355`; with its independent mathematical review.
[20] Asymptotic Rank Research Swarm, *Completed radial certificates through the exponent* 599/500, 25 November 2026. Pinned source archive, including the preceding completed bounds 1.2048, 1.2025, 1.2010 and 1.1999, and their full finite-verification records.
[21] Asymptotic Rank Research Swarm, *Completed upper* 1.1969, 25 November 2026. Completion and reproducibility record, revision `fd52d355`. The correlated angular comparison and rational parameter selection are credited there to the preceding 1.1970 construction. The companion to this article includes the pinned sources, mathematical reviews, and a fresh complete execution at 11969/10000.
[22] Asymptotic Rank Research Swarm, Group B, *Universal asymptotic rank exponent at most* 2393/2000, 25 November 2026. Pinned certificate and proof, revision `54fd0c28`.
[23] Asymptotic Rank Research Swarm, Group B, *Universal exponent at* 1.1955 *and the open* 1.18 *frontier*, 26 November 2026, PR 1218. Pinned completion record and source archive, revision `17b44508`. Only the completed unrestricted 2391/2000 certificate is used here.
[24] Asymptotic Rank Research Swarm, Group B, *Universal asymptotic-rank exponent at most* 239/200, 26 November 2026, PR 1220. Pinned complete certificate, revision `79b0c465`; complete replay transcript, revision `488013c7`.
[25] Asymptotic Rank Research Swarm, Group B, *Completed universal exponent* 10161/8503, 26 November 2026, PR 1245. Pinned proof, source distribution and completion records, revision `97958f33`. The complete execution and the fixed-record audit were repeated for the present manuscript.
[26] P. Uznański, *Rectangular matrix multiplication from polynomial products*, arXiv:2610.10494v2 (2026), Proposition 5.2 and equation (22).