hexagon:2611.01066v1

EUCLIDEAN TENSOR PRICES AND SHARP MULTIQUADRATIC COMPOSITION BOUNDS

SUVRIT SRA*

Abstract. Conditional on two recent results, the Euclidean price theorem announced by OpenAI on 24 November 2026 and Khoruzhii's integral crossing lemma, we determine the exponent of the sum-of-squares composition problem for every fixed number of factors. Over an algebraically closed field of characteristic different from two, a product of k sums of squares, the i-th in $\lceil n^{\beta_i} \rceil$ variables, is a sum of $n^{\|\beta\|_2+o(1)}$ squares of multilinear forms, and of no fewer; for k blocks of equal size the exponent is $\sqrt{k}$. The same exponent holds for the partial-transpose rank of the identity matrix over every field, and for the border versions of both quantities. The lower bound comes from the price theorem and holds at finite sizes with explicit constants; the upper bound composes crossing decompositions recursively. Combined with an argument of Liu, the price theorem also shows that noncommutative circuits for the permanent and the determinant of order N have size at least $2^{N/2-O(\sqrt{N \log N})}$ over every field, even in the border model. We have not verified the proof of the price theorem, on which every lower bound here depends.

1. Introduction

For block sizes $\boldsymbol{n} = (n_1, \ldots, n_k)$ with $n_i \geq 1$, and independent blocks of variables $x^{(1)}, \ldots, x^{(k)}$ of these sizes, consider the product of k sums of squares

$$Q_{\boldsymbol{n}}(x) = \prod_{i=1}^{k} \left(\sum_{j=1}^{n_i} (x_j^{(i)})^2 \right).$$

The composition problem asks for $\mathrm{SOS}_{\mathbb{F}}(\boldsymbol{n})$, the least number of squares that sum to $Q_{\boldsymbol{n}}$ when each is the square of a form over the field $\mathbb{F}$ that is linear in every block. For $k = 2$ the problem is Hurwitz's, and Hrubeš et al. (2010) recount its history. They also showed why its asymptotics matter in complexity theory: a bound $\mathrm{SOS}_{\mathbb{F}}(n, n) \geq \Omega(n^{1+\varepsilon})$ over an algebraically closed field would force noncommutative circuits of size $2^{\Omega(N)}$ for the permanent of order N (Hrubeš et al., 2010, Theorem 1.4). For general k we use the partial-transpose framework of Rossman and Zhu (2026).

Upper bounds for $\mathrm{SOS}_{\mathbb{C}}(n, n)$ have improved recently. A subquadratic bound is due to Hrubeš (2026), and Khoruzhii (2026, Corollary 3) lowered it to $n^{\sqrt{2}+o(1)}$, building on the crossing construction in Anthropic's algorithm for 3-SAT (Claude, 2026a). On 24 November 2026, OpenAI announced a theorem that prices tensors by Euclidean geometry (OpenAI, 2026a, Theorem 1.1). When the indices of a tensor are placed at points of a Euclidean space, a wire of dimension q between two points at distance ℓ has price $\ell \log q$; on each finite set of points, the price is a multiple of the logarithm of a point of the asymptotic spectrum of tensors. On the same day, Liu (2026, Section 2) showed that $\mathrm{SOS}_{\mathbb{C}}(n, n) = n^{\sqrt{2}+o(1)}$ if no planar crossing admits a strict shortening, a hypothesis on tensor networks in the unit disk that was isolated by OpenAI (2026c).

Date: November 25, 2026.

* TU Munich, Dept. of Mathematics, School of CIT; Garching, Germany. Email: s.sra@tum.de.

We derive the composition exponents from the price theorem directly, for every fixed number k of factors, for blocks of unequal sizes, and with explicit bounds at finite sizes.[1] Liu's derivation treats two factors of equal size over $\mathbb{C}$ and passes through the no-shortening hypothesis, whereas ours applies the price theorem to a placement of the indices that is weighted by the block sizes. Two external statements enter, both recalled in Section 2: (EP), the price theorem, and (CR), the integral crossing lemma of Khoruzhii (2026, Lemma 1). A statement below depends on one of them exactly when it begins by assuming it.

To state the result, write $\mathrm{BF}_{\mathbb{F}}(\boldsymbol{n})$ for the least number of products of two forms, each linear in every block, that sum to $Q_{\boldsymbol{n}}$; the two factors of a product need not be equal. Both $\mathrm{SOS}_{\mathbb{F}}$ and $\mathrm{BF}_{\mathbb{F}}$ refer to identities of formal polynomials, not to equality on Boolean inputs. Write $P_{\mathbb{F}}(\boldsymbol{n})$ for the partial-transpose rank (PT rank) of the identity matrix on $\bigotimes_i \mathbb{F}^{n_i}$, the least number of matrices that sum to the identity when each has rank one after transposing some of the k tensor factors (Definition 3.1). Underlines denote the border versions of these quantities over the algebraic closure $\overline{\mathbb{F}}$. All logarithms are to base two.

Theorem 1.1 (Sharp fixed-arity exponent)**.** *Assume (EP) and (CR). Fix $k \geq 1$ and $\beta_1, \dots, \beta_k > 0$; let $n_i = \lceil n^{\beta_i} \rceil$ as $n \to \infty$. Then over every field $\mathbb{F}$,*

$$\underline{P}_{\mathbb{F}}(\boldsymbol{n}),\ P_{\mathbb{F}}(\boldsymbol{n}) = n^{\|\boldsymbol{\beta}\|_2 + o(1)}.$$

If $\operatorname{char}\mathbb{F} \neq 2$, the same holds for $\mathrm{BF}_{\mathbb{F}}$ and its border version. If moreover -1 is a sum of finitely many squares in $\mathbb{F}$, the same holds for $\mathrm{SOS}_{\mathbb{F}}$ and its border version. The border SoS conclusion holds over the algebraic closure in every characteristic different from two, even if the original field is formally real. Over all fields of characteristic different from two, the stated exponent is a lower bound for exact SoS.

The equality in Theorem 1.1 is one of exponents, that is, the logarithm of each rank is $(\|\boldsymbol{\beta}\|_2 + o(1)) \log n$. In particular, the exponent over $\mathbb{C}$ is $\sqrt{1+\beta^2}$ for two blocks of sizes n and $\lceil n^{\beta} \rceil$, and it is $\sqrt{k}$ for k blocks of size n, which for $k = 2$ is the classical composition problem.

The lower bound compares lengths. Place the two indices of the i-th identity wire at the points $\pm e_i$ of $\mathbb{R}^k$, so that the identity has price $2\sum_i \log n_i$. A matrix of rank r across a cut factors through two auxiliary points, one for each side of the cut, joined by a wire of dimension r. Moving these two points a distance t from the origin toward the two sides of the cut lengthens the central wire to $2t$ and shortens the wires to the $2k$ indices. To first order in t, the saving is at most $2t\,(\sum_i (\log n_i)^2)^{1/2}$, attained in the direction $\sum_i \pm(\log n_i) e_i$, which is the source of the Euclidean norm in Theorem 1.1. Under (EP), a sum of J such networks has price at most that of its most expensive summand plus $W \log J$, where W is the mean width of the set of all points used (Lemma 2.1). Since the identity has a fixed price, the ranks r must be large, and optimizing t gives Theorem 3.2.

For the upper bound, (CR) gives the exponent $\sqrt{a^2+b^2}$ for two factors of sizes n^a and n^b (Lemma 4.1). Splitting the k factors into two groups whose PT ranks are n^a and n^b exhibits the identity as a sum of boundedly many linear images of a two-factor identity of these sizes (Lemma 4.2), and Euclidean norms compose under this recursion. A tensor product of two-factor decompositions, by comparison, gives only the exponent $\sqrt{2}\lfloor k/2 \rfloor + (k \bmod 2)$ for k blocks of size n.

The lower bound is not only asymptotic. Theorem 3.2 holds for all block sizes and for border rank, and in characteristic different from two it gives $\log \mathrm{SOS}_{\mathbb{F}}(n,m) \geq \sqrt{(\log n)^2 + (\log m)^2} - \sqrt{2\sqrt{2}\log(nm)}$ for all n and m (Corollary 3.3). The same comparison of prices applies to noncommutative circuits. Liu (2026, Theorem 1.1) proved that the row-ordered permanent and determinant of

[1] The composition and partial-transpose bounds of this paper were derived independently of Liu (2026), whose preprint is dated 24 November 2026; the overlap in the case of two factors of equal size was found during preparation. We make no claim of priority. Section 6, by contrast, uses Liu's argument.

order N have border circuit size s with $\log(s+1) \geq N/2 - O(\sqrt{N \log(N+1)})$, provided that no planar crossing admits a strict shortening. We show that (EP) implies this hypothesis (Proposition 6.1), and Theorem 6.2 records the resulting bound over every field.

Three limits of these results should be stated at the outset. Every lower bound in the paper rests on (EP), and the upper bounds for $k \geq 2$ rest on (CR); we take both as stated in the cited sources and have verified neither proof. The circuit bounds concern noncommutative computation. On commuting inputs the row-ordered determinant is the ordinary determinant, so Theorem 6.2 gives no lower bound for commutative circuits or for Boolean circuits that compute zero-one permanents, no proof of $\mathsf{P} \neq \mathsf{NP}$, and no independent proof of $\mathsf{VP} \neq \mathsf{VNP}$. Finally, over formally real fields we prove the lower bound for $\mathrm{SOS}_{\mathbb{F}}$ but no matching upper bound.

The formal statements are the following.

(1) The lower bound for PT rank at finite sizes, valid for Zariski closures (Theorem 3.2), and its consequence for squares and products of multilinear forms (Corollary 3.3).

(2) The matching upper bound in every fixed arity over every field (Lemmas 4.1 and 4.2), which completes Theorem 1.1.

(3) The exponents for Hurwitz's problem over $\mathbb{C}$ and their form as a rank minimization over an affine space of matrices (Corollary 5.1 and Proposition 5.2).

(4) The derivation of Liu's no-shortening hypothesis from (EP), and the resulting bounds for noncommutative circuits (Proposition 6.1 and Theorem 6.2).

2. The two external inputs and the price of a sum

Every lower bound below compares the price of a target tensor with the prices of networks that sum to it. This section states the two external theorems, (EP) for the lower bounds and (CR) for the upper bounds, and proves the comparison in the form we use. We have not verified the proof of (EP), the proof of (CR), or their formalizations. The companion paper (OpenAI, 2026b) of OpenAI (2026a) is not used.

(EP): Euclidean prices. We take OpenAI (2026a, Theorem 1.1) as stated. Let S be a finite set of distinct points of a Euclidean space, of positive diameter; its points are called *sites*. Over any field there is a spectral point F_S on tensors whose indices are placed at sites of S, and such a tensor T has the scaled logarithmic price $\mathsf{H}(T) = W(S) \log F_S(T)$, where

$$W(S) = \sqrt{2\pi}\, \mathbb{E} \max_{x \in S} \langle g, x \rangle$$

for a standard Gaussian vector g in the Euclidean span of S. The mean width W is normalized so that a segment has mean width equal to its length. We use the following properties, and no formula for the price of a general tensor beyond them. The spectral point F_S is multiplicative under tensor products, additive on direct sums of tensors, normalized on the unit, and monotone under linear maps applied at the sites. A *wire* of dimension q from a site x to a site y is the identity tensor I_q with one index at each of the two sites; its price is $\|x - y\|_2 \log q$, and since F_S is multiplicative, the prices of independent wires add. The price is compatible with enlarging the set of sites and with grouping indices placed at the same site. In particular we do not assume that a geometric mean of flattening ranks is itself a spectral point.

(CR): Integral crossing. The upper bounds rest on one construction for two factors. We take the following statement from the version of 4 November 2026 of Khoruzhii (2026, Lemma 1 and Section 4). Let $Q(T)$ be the number of sets of distinct half-integer energy modes of total energy at most T. For integers $H \geq 1$ and $E \geq 0$ with $Q(H) > 2Hu$, the tensor $\delta_{ii'}\delta_{jj'}$, with $i, i' \in [u]$ and $j, j' \in [Q(E)]$, has a decomposition with integer coefficients into at most $4HQ(H+E)$ terms, each of rank one across one of the two transversal cuts. Moreover $\log Q(T) = \kappa\sqrt{T} + o(\sqrt{T})$ with a positive constant κ. Because the coefficients are integers, the decomposition specializes to every field. No matrix multiplication exponent is used in the rank argument below.

The price of a sum. Under (EP), a sum of J networks costs at most as much as its most expensive summand plus $W(S)\log J$. To state this bound, call a tensor product of wires along the edges e of a graph on S, with dimensions q_e, a *graph tensor*, and call the image of a tensor under linear maps applied at the sites a *restriction* of it. The introduction of OpenAI (2026a) notes the bound for sums. We include its proof together with a version for Zariski closures, which is what gives the border versions of all later bounds. The closure argument is the one sketched after OpenAI (2026c, Definition 2.1) and used by Liu (2026, Section 2).

Lemma 2.1 (Network price bound)**.** *Assume (EP). Let T be a sum of restrictions of graph tensors $G_1,\dots,G_J$ on the same finite set of sites S (allow unused sites). With $\ell_j=\sum_{e\in G_j}\mathrm{length}(e)\log q_e$,*

$$2^{\mathsf{H}(T)/W(S)}\le\sum_{j=1}^{J}2^{\ell_j/W(S)},\qquad \mathsf{H}(T)\le\max_j\ell_j+W(S)\log J.$$

The same bounds apply to the Zariski closure of all such sums over an algebraically closed field, with the graphs, their placement on S, and the dimensions q_e fixed.

Proof. ***Sums.*** A sum of tensors is a restriction of their direct sum, obtained by adding the coordinate blocks at each site. Since F_S is monotone under restriction and additive on direct sums, $F_S(T)\le\sum_j F_S(G_j)$. Multiplicativity and the price of a wire give $F_S(G_j)=2^{\ell_j/W(S)}$, which proves both inequalities for sums.

Closures. The closure claim holds for every spectral point F. Let $\mathcal{X}$ be the set of all sums of restrictions of $G_1,\dots,G_J$, and let $B=\sum_j F(G_j)$, so that $F(X)\le B$ for every $X\in\mathcal{X}$. Suppose that T lies in the Zariski closure of $\mathcal{X}$ in a space of tensors of finite dimension D. For every positive integer h, the power $T^{\otimes h}$ belongs to the linear span of the powers $X^{\otimes h}$ with $X\in\mathcal{X}$, because a linear equation that vanishes on all these powers is a polynomial equation in the coordinates of X, and so it also vanishes at T. After the canonical reordering of tensor factors, this span has dimension at most $\binom{h+D-1}{D-1}$, since the entries of an h-th power are monomials of degree h in D coordinates, grouped by their exponent vectors. The count is valid in every characteristic.

We can therefore express $T^{\otimes h}$ as a linear combination of at most $\binom{h+D-1}{D-1}$ powers $X^{\otimes h}$, absorbing the scalar coefficients into the map at one site. Subadditivity under sums and multiplicativity give

$$F(T)^h\le\binom{h+D-1}{D-1}B^h.$$

Taking h-th roots and letting $h\to\infty$ with D fixed proves $F(T)\le B$, which we apply to $F=F_S$. The argument concerns closures only; it does not assert that spectral points are continuous. □

3. Weighted lower bounds for partial-transpose rank

This section proves the lower bound of Theorem 1.1 in a form valid for all block sizes. All cuts receive networks on one set of sites, the two hubs of each cut are moved a distance t toward the sides of that cut, and Lemma 2.1 compares the prices. We first fix the tensor and its decompositions.

Definition 3.1. The tensor $I_{\boldsymbol{n}}$ is the product of k identity wires I_{n_i}, one on each pair of index spaces V_i^+,V_i^- of dimension n_i. A *transversal cut* contains one member of each pair on each side; the two orientations of a cut are identified. Write $P_{\mathbb{F}}(\boldsymbol{n})$ for the least number of tensors, each of rank one across a transversal cut, whose sum is $I_{\boldsymbol{n}}$; different terms may choose different cuts. This number is the partial-transpose rank of the identity matrix with these index blocks, in the sense of Rossman and Zhu (2026). A *grouped decomposition* is a representation $I_{\boldsymbol{n}}=\sum_\epsilon T_\epsilon$ with one summand for each transversal cut ϵ, where T_ϵ has flattening rank r_ϵ across ϵ. By rank factorization of each T_ϵ, minimizing $\sum_\epsilon r_\epsilon$ is equivalent to minimizing the number of individual terms. The border version $\underline{P}_{\mathbb{F}}(\boldsymbol{n})$ takes the Zariski closure, over $\overline{\mathbb{F}}$, of the set of tensors that have a grouped decomposition of bounded total rank.

The bounds are stated through the logarithms a_i of the block sizes, their sum A and Euclidean norm L, the number p of transversal cuts, and a mean width W_k:

$$a_i = \log n_i, \qquad A = \sum_i a_i, \qquad L = \Big(\sum_i a_i^2\Big)^{1/2}, \qquad p = 2^{k-1}, \qquad W_k = \sqrt{2\pi}\,\mathbb{E}\max_{1\le i\le k}|g_i|,$$

where g_i are independent standard normal variables. The number W_k is the mean width of the set of sites $\{\pm e_i : 1 \le i \le k\}$ on which the proof places the indices. In particular $W_1 = 2$ and $W_2 = 2\sqrt{2}$; for W_2, polar coordinates give $\mathbb{E}\|g\|_2 = \sqrt{\pi/2}$, and the angular mean of $\max(|\cos\theta|, |\sin\theta|)$ is $2\sqrt{2}/\pi$. In general $W_k \le \sqrt{2\pi}\sqrt{2\ln(2k)}$. Indeed, bounding the exponential of the maximum by $\sum_i (e^{tg_i} + e^{-tg_i})$ and using Jensen's inequality and the Gaussian moment generating function gives $\mathbb{E}\max|g_i| \le t^{-1}\ln(2k) + t/2$, which we optimize in t.

The main estimate bounds the total rank R of a grouped decomposition, which is what PT rank measures, and also its largest rank $r_{\max}$, which is what a sum of squares controls (Corollary 3.3). Its free parameter t is the distance of the hubs from the origin.

Theorem 3.2 (Finite weighted bound)**.** *Assume (EP) and $A > 0$. For a grouped decomposition of $I_{\boldsymbol{n}}$, let $r_{\max} = \max_\epsilon r_\epsilon$ and $R = \sum_\epsilon r_\epsilon$. For every $0 < t \le L/A$,*

$$\log r_{\max} \ge L - \frac{At}{2} - \frac{W_k \log p}{2t}, \tag{3.1}$$

$$\log R \ge L - \frac{At}{2} - \frac{W_k \log p}{2t} + \log p. \tag{3.2}$$

The same estimates hold when $I_{\boldsymbol{n}}$ lies in the Zariski closure of the set of tensors that have a grouped decomposition with ranks at most $(r_\epsilon)_\epsilon$. In particular, for all $\boldsymbol{n}$,

$$\log \underline{P}_{\mathbb{F}}(\boldsymbol{n}) \ge \max\left\{\max_i a_i,\ L - \sqrt{AW_k(k-1)}\right\}. \tag{3.3}$$

When $L^2 \ge AW_k(k-1) > 0$, the second term in the maximum for $\log \underline{P}$ can be increased by $k-1$.

Proof. ***The networks.*** To apply Lemma 2.1, we realize all summands T_ϵ by networks on one set of sites. Place the two indices of the i-th pair at the sites $+e_i$ and $-e_i$, which we call *terminals*. Each wire I_{n_i} then has length two, so the target $I_{\boldsymbol{n}}$ has price $2A$. A cut is a class $\epsilon \in \{\pm 1\}^k/\{\pm 1\}$, and the terminals on one of its sides are the $\epsilon_i e_i$. For each cut add two sites, the *hubs* $\pm t v_\epsilon$, where

$$v_\epsilon = L^{-1}\sum_i \epsilon_i a_i e_i, \qquad \|v_\epsilon\|_2 = 1.$$

Both hubs lie in the convex hull of the terminals, because $\|tv_\epsilon\|_1 = tA/L \le 1$. The union of all sites, over all cuts, therefore has mean width exactly W_k, and the networks of all cuts are defined on this one finite set of sites.

The summand T_ϵ is a restriction of a graph tensor with two kinds of wires (Figure 1): a *spoke* of dimension n_i from each terminal to the hub on its side of the cut, and a central wire of dimension r_ϵ between the two hubs. The maps at the two hubs are the two factors of a rank factorization of T_ϵ, with outputs of dimension one, since T_ϵ has no index at a hub. Zero summands may be omitted. The central wire has length $2t$, and the weighted length of the spokes is independent of the cut:

$$\begin{aligned} B(t) &= 2\sum_i a_i\sqrt{1 - 2ta_i/L + t^2} \\ &\le 2A - 2tL + At^2. \end{aligned} \tag{3.4}$$

The inequality uses $\sqrt{1+z} \le 1 + z/2$ for $z \ge -1$.

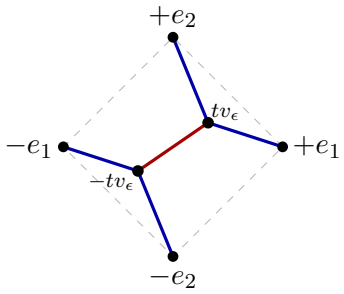

FIGURE 1. The network of one transversal cut for $k = 2$: four spokes (blue) from the terminals to the hubs $\pm tv_\epsilon$, and the central wire (red). The hubs lie in the convex hull of the terminals (dashed), so they do not change the mean width.

The comparison. Lemma 2.1, applied to the grouped decomposition, gives

$$2^{(2A-B(t))/W_k} \le \sum_\epsilon r_\epsilon^{2t/W_k}. \tag{3.5}$$

By the closure part of the lemma, the inequality (3.5) also holds in the border case, for fixed bounds on the ranks r_ϵ. Bounding the sum by $p\, r_{\max}^{2t/W_k}$ and using the spoke bound (3.4) gives (3.1). For the total rank, the exponent $\theta = 2t/W_k$ lies in $(0, 1]$, because $W_k \ge 2$ (the hull contains a segment of length two) and $t \le L/A \le 1$. Concavity gives $\sum_\epsilon r_\epsilon^\theta \le p^{1-\theta} R^\theta$, zero ranks included, hence (3.2).

The choice of t**.** Let $k \ge 2$, so that $W_k \log p = W_k(k-1) > 0$. If $L^2 \ge AW_k(k-1)$, the choice $t = \sqrt{W_k(k-1)/A}$ is allowed, and the two penalty terms in (3.1) and (3.2) sum to $\sqrt{AW_k(k-1)}$; the term $\log p = k-1$ in (3.2) is the increase stated at the end of the theorem. If $L^2 < AW_k(k-1)$, the term $L - \sqrt{AW_k(k-1)}$ is negative and the bound is trivial. For $k = 1$ the rank is n_1, including in the border case, and if $A = 0$, all dimensions are one and the rank is one.

The bound $\max_i a_i$**.** The first term of (3.3) does not use prices. Fixing both indices of every pair except the i-th to their first basis coordinate maps each transversal term to a matrix of rank at most one and maps the target to I_{n_i}. Hence $R \ge n_i$ by subadditivity of matrix rank, and the same holds for border rank because matrix minors give closed conditions. The two bounds together prove (3.3). In the border model with bounded total rank there are finitely many allocations of integer ranks $(r_\epsilon)_\epsilon$, so the closure of their finite union is the union of their closures, and the proof applies to the relevant allocation. □

Polarization turns a sum of squares into a grouped decomposition whose ranks are bounded cut by cut, so Theorem 3.2 bounds the number of squares and of products.

Corollary 3.3 (Multilinear composition lower bound)**.** *Assume (EP) and* $\operatorname{char} \mathbb{F} \ne 2$*. Then*

$$\log \underline{\mathrm{SOS}}_{\mathbb{F}}(\boldsymbol{n}) \ge \max\{\max_i a_i, L - \sqrt{AW_k(k-1)}\},$$
$$\log \underline{\mathrm{BF}}_{\mathbb{F}}(\boldsymbol{n}) \ge L - \sqrt{AW_k(k-1)} - 1.$$

In particular, for two factors,

$$\log \mathrm{SOS}_{\mathbb{F}}(n, m) \ge \max\left\{\log n, \log m, \sqrt{(\log n)^2 + (\log m)^2} - \sqrt{2\sqrt{2}\log(nm)}\right\}. \tag{3.6}$$

The bound also holds for border SoS over the algebraic closure.

Proof. ***From squares to cuts.*** The link between squares and transversal cuts is polarization, which is independent of the price theorem. Separately polarize the quadratic dependence in each

block: replace $x^{(i)}$ by $u^{(i)} + v^{(i)}$ and subtract the two single-block evaluations. The target becomes $2^k \prod_i \langle u^{(i)}, v^{(i)} \rangle$, with bilinear, not conjugate-linear, pairings. A multilinear square $f(x)^2$ contributes

$$\sum_{\epsilon \in \{\pm 1\}^k} f(z_\epsilon) f(z_{-\epsilon}),$$

where z_ϵ selects $u^{(i)}$ or $v^{(i)}$ in block i according to the sign ϵ_i. The two orientations of a cut give identical products. After division by 2^k, a sum of s squares therefore supplies a grouped decomposition of $I_{\boldsymbol{n}}$ with $r_\epsilon \le s$ for every cut. In a product $f(x)g(x)$ the two orientations need not give the same product, so each cut receives two terms, and b products supply a grouped decomposition with $r_\epsilon \le 2b$. Nonzero scalar factors from polarization can be absorbed in the coefficient tables. Conversely, setting $v = u$ in a decomposition of $I_{\boldsymbol{n}}$ into R terms of rank one across transversal cuts gives a sum of at most R products of multilinear forms. Thus

$$\mathrm{BF}_{\mathbb{F}} \le P_{\mathbb{F}} \le 2^k \, \mathrm{BF}_{\mathbb{F}}, \qquad P_{\mathbb{F}} \le 2^{k-1} \, \mathrm{SOS}_{\mathbb{F}} \,. \tag{3.7}$$

Polarization and the substitution $v = u$ are linear maps of coefficient spaces, so they map closures into closures, and the bounds $r_\epsilon \le s$ and $r_\epsilon \le 2b$ also hold for border identities.

The bounds. Use the bound (3.1) with these bounds on $r_{\max}$ and with the choice of t made in the proof of Theorem 3.2. The lower bound $\mathrm{SOS} \ge n_i$ follows by setting every other block to a standard basis vector: the symmetric matrix of a sum of s squares of linear forms has rank at most s, while the matrix of the target is the identity, and the argument survives closure. Substituting $W_2 = 2\sqrt{2}$ gives (3.6). □

Two features of the proofs enter the constants in (3.6): the mean width $W_2 = 2\sqrt{2}$ of the diamond, which is smaller than the bound π valid for arbitrary sites in the unit disk, and the fact that s squares give at most s terms for every cut, so that the bound (3.1) applies to s itself. For given sizes one can further optimize the spoke cost $B(t)$ in (3.5) instead of its upper bound (3.4), and no asymptotic limit in t is needed. Pairs with $a_i = 0$ have dimension one and may be omitted, and for one factor the rank lower bound is exact. The restriction on the characteristic in Corollary 3.3 is necessary. In characteristic two polarization is not invertible, and $Q_{\boldsymbol{n}} = (\prod_i \sum_j x_j^{(i)})^2$ is one square, so the lower bound for PT rank over every field is not a lower bound for SoS over every field.

4. Matching upper bounds in arbitrary fixed arity

Two steps give the upper bound of Theorem 1.1. The first is (CR), which settles two factors of sizes n^a and n^b. The second is an elementary composition lemma, valid over every field, which reduces k factors to two, because a grouped decomposition of each of two groups of factors writes $I_{\boldsymbol{n}}$ as a sum of linear images of two-factor identities.

Lemma 4.1 (Rectangular crossing exponent)**.** *Assume (CR). For fixed $a, b > 0$, over every field,*

$$P_{\mathbb{F}}(\lceil n^a \rceil, \lceil n^b \rceil) \le n^{\sqrt{a^2+b^2}+o(1)}.$$

Proof. Fix a slack $\eta > 0$ and choose integer cutoffs $H \sim ((a+\eta)\log n/\kappa)^2$ and $E \sim ((b+\eta)\log n/\kappa)^2$. Then $Q(H) = n^{a+\eta+o(1)} > 2H\lceil n^a \rceil$ and $Q(E) \ge \lceil n^b \rceil$ for large n, so (CR) applies with $u = \lceil n^a \rceil$; restrict the second index space of its decomposition to the required size. Its total rank is at most $4HQ(H+E) = n^{\sqrt{(a+\eta)^2+(b+\eta)^2}+o(1)}$, and the coefficients specialize from $\mathbb{Z}$ to $\mathbb{F}$. Letting $\eta \downarrow 0$ bounds the limit superior of the exponent, and monotonicity under restricting pairs handles rounding. □

Together with Theorem 3.2, Lemma 4.1 also determines the PT-rank exponent of two factors in characteristic two, where polarization is not available. To pass from two factors to k, we combine decompositions of two groups of factors.

Lemma 4.2 (Composition of grouped PT decompositions)**.** *Split the set of paired factors into two nonempty disjoint groups A_0, B_0 of sizes k_A, k_B. With $P_A = P_{\mathbb{F}}((n_i)_{i\in A_0})$ and $P_B = P_{\mathbb{F}}((n_i)_{i\in B_0})$,*

$$P_{\mathbb{F}}(\boldsymbol{n}) \le 2^{k_A+k_B-2}\, P_{\mathbb{F}}(P_A, P_B). \tag{4.1}$$

Proof. Take PT decompositions of the identities of the two groups with P_A and P_B terms, and group each by its cuts. For one cut of A_0, the corresponding summand is a tensor $\sum_{\alpha=1}^{r} U_\alpha \otimes V_\alpha$ of rank r, where U_α occupies one member of every pair in A_0 and V_α occupies the complementary members. For a cut of B_0 write the summand as $\sum_{\beta=1}^{s} X_\beta \otimes Y_\beta$. The tensor product of these two summands is a linear image of the two-factor identity $I_r \otimes I_s$, whose four index spaces we map to the four groups of indices occupied by $U_\alpha, X_\beta, V_\alpha, Y_\beta$, respectively, by the maps that take basis vectors to these tensors. A term of rank one across the first transversal cut of $I_r \otimes I_s$ maps to a term of rank one across the union of the two selected cuts; across the second cut it maps to one across the cut of A_0 united with the complement of the cut of B_0. In both cases the resulting cut is transversal on all pairs. The four maps need not factor further inside their groups of indices, which is harmless because each group lies entirely on one side of the resulting cut.

Thus the product of the two summands has a PT decomposition of length at most $P_{\mathbb{F}}(r, s) \le P_{\mathbb{F}}(P_A, P_B)$, by monotonicity under index restriction. Summing over the at most $2^{k_A-1}2^{k_B-1}$ pairs of cuts gives (4.1). The proof uses neither division nor signs, and is valid over every field. □

Proof of Theorem 1.1. Lower bound. Here $a_i = \beta_i \log n + o(1)$. With k fixed, the penalty in Theorem 3.2 is $O_{k,\beta}(\sqrt{\log n}) = o(\log n)$; hence both exact and border PT rank have logarithmic exponent at least $\|\boldsymbol{\beta}\|_2$. The same reasoning applies to Corollary 3.3.

Upper bound for PT rank. Induct on k using Lemma 4.2. At $k = 1$ the rank is exactly n_1. For larger k, split the factors into two nonempty groups A_0 and B_0. By induction the PT ranks of the two groups are at most $n^{a+o(1)}$ and $n^{b+o(1)}$, where $a = \|\beta_{A_0}\|_2$ and $b = \|\beta_{B_0}\|_2$. For any fixed $\eta > 0$, enlarge these ranks to $\lceil n^{a+\eta} \rceil$ and $\lceil n^{b+\eta} \rceil$, and use Lemma 4.1 and monotonicity. The number of pairs of cuts is constant in n. Letting $\eta \downarrow 0$ yields the exponent $\sqrt{a^2+b^2} = \|\beta\|_2$ as a limit superior. Exact upper bounds also bound the border quantities.

Products and squares. By the comparison (3.7), exact product rank has the same upper exponent. If $-1 = \sum_{j=1}^{h} c_j^2$ in the field, then in characteristic different from two,

$$UV = \left(\frac{U+V}{2}\right)^2 + \sum_{j=1}^{h} \left(c_j \frac{U-V}{2}\right)^2.$$

The identity replaces each product by at most $1 + h$ squares of multilinear forms, a factor that is constant for a fixed field. Such a representation of -1 exists in algebraically closed fields of characteristic different from two, in $\mathbb{Q}(\mathrm{i})$, and in all finite fields of odd characteristic (use the fixed prime characteristic to express -1 as a sum of ones). The border upper bound over $\overline{\mathbb{F}}$ follows there. For formally real fields the lower bound holds by the first step, but we assert no matching exact SoS upper bound. □

5. Consequences for composition and matrix completion

Theorem 1.1 directly gives the exponents for Hurwitz's problem over $\mathbb{C}$, and the same exponents for a rank minimization over an affine space of matrices. We also record what Theorem 3.2 gives when the number of factors grows.

Corollary 5.1 (Sharp Hurwitz exponents)**.** *Under (EP) and (CR), over $\mathbb{C}$,*

$$\mathrm{SOS}_{\mathbb{C}}(n, \lceil n^{\beta} \rceil) = n^{\sqrt{1+\beta^2}+o(1)}, \qquad \mathrm{SOS}_{\mathbb{C}}(n, \ldots, n) = n^{\sqrt{k}+o(1)} \quad (k \text{ fixed}).$$

These exponents also hold for the respective border ranks. In particular, the square two-factor exponent is $\sqrt{2}$.

Proof. Apply Theorem 1.1. □

The exponent $\sqrt{2}$ for two factors of equal size is the one that Liu (2026, Section 2) obtains when no planar crossing over $\mathbb{C}$ admits a strict shortening; Proposition 6.1 derives that hypothesis from (EP), whereas the proof above uses (EP) directly.

The numbers $\mathrm{SOS}_{\mathbb{F}}$ and $\mathrm{BF}_{\mathbb{F}}$ are also minimal ranks in an affine space of matrices, so the exponents above answer a matrix completion problem.

Proposition 5.2 (Gram completion)**.** *Over an algebraically closed field of characteristic different from two,* $\mathrm{SOS}_{\mathbb{F}}(\boldsymbol{n})$ *equals the smallest matrix rank of a symmetric matrix* M *on the product index space whose average over all partial transposes is the identity. The analogous minimum over arbitrary, not necessarily symmetric, matrices equals* $\mathrm{BF}_{\mathbb{F}}(\boldsymbol{n})$*.*

Proof. Let $z = \bigotimes_i x^{(i)}$. The form $z^{\mathsf{T}}Mz$ is unchanged by partial transposition in any pair, and averaging M over all 2^k partial transposes produces a tensor symmetric separately in each pair. The map from this symmetric subspace to multiquadratic polynomials is injective: each coefficient is a tensor entry multiplied by 2^j, where j counts the pairs with two distinct indices, and all these factors are invertible. Hence $z^{\mathsf{T}}Mz = Q_{\boldsymbol{n}}$ holds exactly when the averaged matrix is $I_{\boldsymbol{n}}$. A rank-r matrix factors into r outer products, yielding r bilinear products in z, and every such product identity yields that matrix. A symmetric rank-r matrix over the stated field has a factorization $M = C^{\mathsf{T}}C$ with r rows: diagonalize by congruence using a nonzero quadratic value as a pivot and take square roots of its nonzero diagonal entries. Conversely a sum of r squares gives such a factorization. The two descriptions prove both equalities. □

Under (EP) and (CR), Proposition 5.2 and Theorem 1.1 show that these affine spaces of matrices have minimum rank $n^{\|\beta\|_2+o(1)}$ over algebraically closed fields of characteristic different from two. The minimum is over matrices with unrestricted coefficients and cancellations, so it says nothing about entrywise-positive factorizations or about PSD rank. Over the reals, a PSD Gram matrix gives a decomposition into real squares, so the lower bound applies to it, but the complex upper bound does not give a real PSD upper bound. The border version rules out approximating the polynomial $Q_{\boldsymbol{n}}$ by too few squares; it does not concern approximating its values on a finite grid.

The formula lower bounds of Rossman and Zhu (2026) need the number of factors to grow, so we record what happens when k grows with n. For equal dimensions the bound (3.3) reads

$$\log \underline{P} \geq \sqrt{k}\log n - \sqrt{k \log n \, W_k(k-1)}.$$

Under (EP), the right-hand side is $(1-o(1))\sqrt{k}\log n$ when $k\sqrt{\ln(2k)} = o(\log n)$. We assert no matching exponent for variable arity. An upper bound that follows from (CR) alone is

$$P_{\mathbb{F}}(n,\dots,n) \leq P_{\mathbb{F}}(n,n)^{\lfloor k/2 \rfloor} n^{k \bmod 2}.$$

It follows by taking products of PT decompositions on disjoint pairs of factors and collecting the resulting transversal cuts, and for growing k its exponent is at most $(\sqrt{2}+o(1))\lfloor k/2 \rfloor + 1$. Consequently the identity matrices cannot meet the hypothesis of PT rank $n^{k-o(\log k)}$ in their criterion (Rossman and Zhu, 2026, Theorem 1.5 and Corollary 1.6), even in characteristic two, and Theorem 1.1 should not be read as a new separation for commutative formulas or circuits.

6. Noncommutative circuits

Liu (2026) works under the hypothesis that no planar crossing admits a strict shortening in the sense of OpenAI (2026c, Definition 2.1), proves an interval decomposition of noncommutative circuits under that hypothesis, and uses an encoding by two increasing runs to transfer a lower bound for a repeated-word polynomial to the row-ordered permanent and determinant. This section shows that (EP) supplies the hypothesis, and then gives the resulting bound with the mean width

W_2 and the closure argument of Lemma 2.1. The mechanism is Liu's; the reductions from repeated words and from sums of squares go back to Hrubeš et al. (2010).

The first step is to derive the hypothesis. In the next statement, the *crossing with two r-bit wires* is the tensor product of two wires of dimension 2^r, one from e_1 to $-e_1$ and one from e_2 to $-e_2$.

Proposition 6.1 (No strict planar shortening)**.** *Assume (EP) over the field. Place the four terminals at $e_1, e_2, -e_1, -e_2$. The crossing with two r-bit wires has price $4r$. It cannot be a sum of restrictions of J networks inside the Euclidean unit disk, each of wire cost at most w, if $w + 4\log J < 4r$. The same holds for the closure of such sums with the networks fixed.*

Proof. The mean width of any finite set of sites in the unit disk is at most $\sqrt{2\pi}\,\mathbb{E}\|g\|_2 = \pi < 4$ for a planar standard Gaussian vector g, whose mean radius in polar coordinates is $\sqrt{\pi/2}$. Including all auxiliary sites simultaneously, which the compatibility in (EP) permits, Lemma 2.1 gives $4r \le w + W(S)\log J \le w + \pi \log J$. The proof for closures is the same. Since $\pi < 4$, the conclusion is stronger than the absence of a strict shortening assumed by Liu (2026, Section 2). □

Liu's conditional bound therefore holds under (EP). We state it as a theorem and write out the proof, to make the parameters and the circuit model explicit.

Theorem 6.2 (Exact and border noncommutative lower bounds)**.** *Assume (EP) over every field. For the row-ordered permanent and row-ordered (Cayley) determinant of order N, with unrestricted binary addition, subtraction, ordered multiplication and arbitrary scalar constants,*

$$\log(1 + \underline{C}_{\mathrm{nc}}(\mathrm{Perm}_N)), \quad \log(1 + \underline{C}_{\mathrm{nc}}(\mathrm{Det}_N)) \;\ge\; \frac{N}{2} - O(\sqrt{N\log(N+1)}).$$

In particular, their exact and border circuit sizes are $2^{\Theta(N)}$. The border model takes the closure over $\overline{\mathbb{F}}$ of degree-N output components of size-bounded circuits, with arbitrary intermediate degrees.

Proof. We follow the argument of Liu (2026, Lemmas 2.1 and 3.1) in the case where no strict shortening exists, with Lemma 2.1 in place of the no-shortening hypothesis.

Repeated words. The bound is first proved for a polynomial whose coefficient tensor is a crossing. In two noncommuting letters z_0, z_1, let

$$D_m = \sum_{w \in \{0,1\}^m} z_{w_1} \cdots z_{w_m} z_{w_1} \cdots z_{w_m}, \qquad m = 2r.$$

The polynomial D_m is homogeneous of degree $d = 2m = 4r$. Its coefficient tensor on the d positions is a product of m binary wires, connecting position i to position $i + m$. Assign the successive positions the unit vectors $e_1, e_2, -e_1, -e_2$, each repeated r times, so that the product of wires has cost $2m$. The prefix sums $p_0, \ldots, p_d$ of those vectors trace the boundary of a square of side r.

Interval decomposition. Lemma 2.1 needs the coefficient tensor of a small circuit as a sum of few tensors, each of low rank across a cut that a short network can realize. Following Liu, we claim that any circuit of size s, projected to its degree-d component, yields a sum of at most $J = (d+1)^2$ tensors, one for each interval $I = (i, j]$ of positions whose chord length $\|p_j - p_i\|$ is at least r, and each of rank at most $G = O((d+1)^3(s+1))$ across I and its complement. For the projection, retain the components of degrees $0, \ldots, d$ at each gate; multiplication becomes a convolution, at cost $O((d+1)^2)$ per gate. Copy each homogeneous gate for every possible starting position. At a degree-one input use the corresponding position-tagged letter, so that a word becomes a coefficient monomial with one variable in each position. Multiplication concatenates intervals and addition preserves the interval, and the result is an interval-labeled circuit of size at most G.

The intervals with long chords are found through a property of convex hulls. The convex hull of a consecutive boundary arc $(p_i, \ldots, p_j)$ splits at p_k into the hulls of the two shorter arcs and the triangle with vertices p_i, p_k, p_j; to see the splitting, cut the convex boundary polygon along the two

chords, deleting redundant collinear vertices when endpoints are collinear or repeated. Moreover, any triangle of boundary points that contains the center has an edge of length at least r. Indeed, if all its edges are shorter, no two vertices lie on opposite sides of the square, so the vertices lie on at most two adjacent sides. With the common corner of these sides as origin, the triangle lies in the half-square $x + y \le r$, and its intersection with the line $x + y = r$ can contain the center only if both endpoints $(r, 0)$ and $(0, r)$ of the two sides are vertices. Their distance is $r\sqrt{2}$, a contradiction.

Now replace every interval gate whose chord has length at least r by a formal placeholder of the same interval. Setting all placeholders to zero makes the output zero, because by induction through the circuit every gate that is not replaced and whose interval hull contains the center is zero. At an addition the claim holds by induction for the two inputs. At a product, a replaced input is already zero; if neither chord is long, the triangle misses the center, so the hull of at least one input contains it, and that input is zero by induction. Leaves with short chords cannot contain the center, and the hull of the root does. Every output word after the substitution therefore contains at least one placeholder. Group the words by their first placeholder and then substitute the original polynomials. The factor for the chosen gate is independent of the remaining positions, so every such group contributes a tensor of rank one across that interval and its complement. No placeholder occurs twice within a word, because its interval would overlap itself. Grouping by interval gives rank at most G for each interval. The decomposition applies uniformly to circuits of bounded size, regardless of reuse or cancellations, so the closure of their projected outputs is contained in the closure of the set of sums of restrictions of the fixed networks defined next.

Price estimate. For each interval I use hubs $\pm t(p_j - p_i)/\|p_j - p_i\|$, with $0 < t \le 1/\sqrt{2}$, which lie inside the diamond of terminals. Put one binary spoke from each position to the hub on its side, and a wire of dimension G between the hubs. The direction sums on the two sides are opposite and of magnitude at least r. By the square-root estimate used for (3.4), the spoke price is at most $2m - 2tr + mt^2$. All networks share the mean width $W_2 = 2\sqrt{2}$. Lemma 2.1, including its closure version, therefore gives

$$2m \le 2m - 2tr + mt^2 + 2t \log G + W_2 \log J.$$

Taking $t = \sqrt{W_2 \log J/m}$ for all sufficiently large m yields

$$\log(1 + \underline{C}_{\mathrm{nc}}(D_m)) \ge m/2 - O(\sqrt{m \log(m+1)}),$$

since $\log G \le \log(1 + s) + O(\log(m + 1))$. Finitely many small m are absorbed in the constant.

Two increasing runs. It remains to obtain D_N from the permanent and the determinant of order N, which is Liu's transfer. For $1 \le j < N$, retain from a row-ordered permanent of order N only the permutations whose first j columns increase strictly and whose last $N - j$ columns increase strictly. Such a permutation consists of the increasing list of a subset $S \subset [N]$ of size j followed by the increasing list of its complement. A matrix substitution of polynomial overhead both filters and encodes this condition. Use states $p = 0, \dots, N$ recording the previous column, starting at zero in each run. At row i and column c, an allowed transition requires $c > p$. For the first run its weight is $z_0^{c-p-1} z_1$, and for the second it is $z_1^{c-p-1} z_0$. The next state is c, except at the last row of the run, when the appropriate trailing power (z_0^{N-c} or z_1^{N-c}) is appended and the state is reset to zero. All other transition weights are zero. These rules specify a matrix for each variable, and the $(0, 0)$ entry of the resulting polynomial emits the incidence word of S twice. For the determinant, in the first run multiply the transition weight at row i, column c by $(-1)^{c-i}$. The total sign of the shuffle is exactly $(-1)^{\sum_{i=1}^{j}(c_i - i)}$, since each $c_i - i$ counts the smaller complementary columns after c_i, and the weights cancel it, including in characteristic two. Each accepted subset occurs once.

Overhead and closure. Sum the $N - 1$ matrix evaluations and add the two constant-bit words z_0^{2N} and z_1^{2N}. The resulting polynomial is exactly D_N. Matrix evaluation on matrices of size $N + 1$ increases circuit size by $O((N + 1)^3)$ per original gate, plus a polynomial preparation of the entries;

the $N-1$ evaluations and the homogeneous extraction cost only further polynomial factors. The overhead is therefore polynomial, which is all the logarithmic lower bound needs. For border circuits, first project to the degree-N homogeneous component, then substitute and take the indicated entries, sum, add the two fixed words, and project to degree $2N$. On the finite-dimensional coefficient space of the degree-N component this composite is an affine map, so the image of a closure point belongs to the closure of the image of the family of circuits of bounded size, with the same polynomial overhead. For even N, insert the lower bound on D_N. For odd N, an identity padding projects the preceding even order from Perm_N and Det_N with no superpolynomial cost. The two stated lower bounds follow.

Upper bound. The claim $2^{\Theta(N)}$ also needs circuits of this size. The exact row-ordered permanent is computed by the subset recursion $f(S)=\sum_{j\in S} f(S\setminus\{j\})x_{|S|,j}$, starting at $f(\varnothing)=1$. For the determinant include the sign $(-1)^{|\{h\in S:h>j\}|}$ at each step, which counts the inversions with the last column. Both circuits have $O(N2^N)$ gates, which gives the matching exponential upper bound, also for border size. □

Sparse symbolic support. A simpler reduction gives exponential lower bounds even for a fixed symbolic support with two entries in each row. On order $2m$, in rows $i, m+i$ and columns $i, m+i$, use the private 2×2 block $\left(\begin{smallmatrix} z_0 & z_1 \\ z_1 & z_0 \end{smallmatrix}\right)$, with all other entries zero. In the row order all first choices are read before all second choices, and selecting one of the two matchings in each block produces exactly D_m. For the determinant replace the lower-left z_1 by $-z_1$ to cancel the swap sign. Hence, under (EP), the corresponding symbolic permanent and determinant with two free variables per row have border log-size at least $m/2-O(\sqrt{m\log(m+1)})$. One may add three rows and columns in block upper-triangular form, fill the degree deficit of each old row using one of the new columns, and put nine independent variables in the new 3×3 block. Setting that block to the identity and specializing the filler entries to zero recovers the original support. The construction yields a fixed support with exactly three symbolic entries per row and an exponential formal noncommutative lower bound. It is not a lower bound for the numerical evaluation of zero-one permanents.

Other established reductions. Corollary 3.3 can also be fed into the reductions of Hrubeš et al. (2010). In characteristic different from two, it yields $\mathrm{BF}(n,n)=n^{\sqrt{2}+o(1)}$, with (CR) for the upper bound. The lower bound alone supplies the premise of their Theorems 1.6 and 1.7 for the constant-degree noncommutative identity polynomial $\sum_{i,j} x_iy_jx_iy_j$: its circuit size is at least $n^{\sqrt{2}-o(1)}$, and the row-ordered permanent has exponential exact circuit complexity. Both conclusions apply their reductions and are not new reductions. The proof of Theorem 6.2 additionally handles characteristic two and border size without using polarization. Liu's manuscript also treats partially commutative models; we assert no new lower bound for those.

7. Discussion

Under (EP) and (CR), the exponent of multiquadratic composition in fixed arity is the Euclidean norm of the logarithmic block sizes. For PT rank the exponent holds over every field, and for products and squares of multilinear forms it holds over the fields named in Theorem 1.1. What makes the lower bound usable is an inequality for all transversal cuts at once, Theorem 3.2, and the noncommutative bounds of Section 6 show how much the price theorem gives once it is accepted.

Three points are left open. Over formally real fields the upper bound for $\mathrm{SOS}_{\mathbb{F}}$ is missing, because the passage from products to squares uses a representation of -1 as a sum of squares. For k growing with n, the lower bound (3.3) and the product upper bound of Section 5 have different exponents. And the finite bounds are not asserted to be optimal for small sizes, where the term $\max_i a_i$ of (3.3) covers small and degenerate inputs; the classical composition identities for small n and m do not contradict them.

We use only the properties of Euclidean prices listed in Section 2. The existence of these prices does not classify the prices that are not Euclidean, and nothing here uses or decides the contracting hypothesis of Claude (2026b, Definition 2.1 and Theorem 2.2) over $\mathbb{Q}$.

Computer checks. As a guard against errors of indexing and sign, the elementary deductions were also checked by computer in finite cases: 180 polarization identities, exactly, by cross differences and by sums over orientations; 80 composition identities with mixed cuts, exactly, by tensor contraction and by an independent entrywise summation; 18,616 boundary triangles that contain the center, by tests on oriented areas and squared distances; 494 words and signs of the encoding by two increasing runs, exactly, together with 21 independent comparisons of permutation counts; and 480 weighted spoke costs in high-precision arithmetic from coordinate distances. All checks passed. They test indexing, signs, geometry and arithmetic; they are not evidence for (EP), and they do not replace the proofs.

AI USAGE

This work is entirely AI-driven. The results, proofs, and computer checks were produced by OpenAI's GPT 6.2, through a long series of prompts and back-and-forth with the author. Anthropic's Claude Fable 5.2 then edited the paper and audited its citations and attributions.

ACKNOWLEDGEMENTS

The author (SS) thanks the Alexander von Humboldt (AvH) Foundation for their generous support via a Humboldt Professorship in AI.

REFERENCES

Claude. A $2^{o(n)}$ time algorithm for 3-SAT. Anthropic, Hexagon 2610.00577v1, October 26, 2026a. URL https://hexagonmath.org/2610.00577v1.

Claude. $\mathsf{VP} \neq \mathsf{VNP}$ over every field. Anthropic, Hexagon 2611.00954v1, November 23, 2026b. URL https://hexagonmath.org/2611.00954v1.

Pavel Hrubeš. A subquadratic upper bound on Hurwitz's problem and related noncommutative polynomials. *SIAM J. Comput.*, 55(2):451–467, 2026. doi: 10.1137/24M1666884.

Pavel Hrubeš, Avi Wigderson, and Amir Yehudayoff. Non-commutative circuits and the sum-of-squares problem. Technical Report TR10-021, Electronic Colloquium on Computational Complexity, 2010. URL https://eccc.weizmann.ac.il/report/2010/021/. Conference version in STOC 2010, pp. 667–676.

Kirill Khoruzhii. Crossing decomposition for subquadratic noncommutative circuits. Lemmata, version of November 4, 2026. URL https://lemmata.iol.berlin/crossing-decomposition.

Yilin Liu. Stretched-exponential circuit lower bounds for the noncommutative permanent and determinant. Hexagon 2611.00994v1, November 24, 2026. URL https://hexagonmath.org/2611.00994v1.

OpenAI. Euclidean geometry in the asymptotic spectrum of tensors. OpenAI Math Release preprint, November 24, 2026a. URL https://mathcommunity.github.io/oai-math-revised/papers/Euclidean-geometry-in-the-asymptotic-spectrum-of-tensors-November-24-2026/.

OpenAI. Circuit lower bounds for Hamiltonian cycle and the permanent. OpenAI Math Release preprint, November 24, 2026b. URL https://github.com/openai/math/blob/main/preprints/Circuit-lower-bounds-for-Hamiltonian-cycle-and-the-permanent-November-24-2026/paper.pdf.

OpenAI. VP and VNP are different. OpenAI Math Release preprint, November 23, 2026c.

Benjamin Rossman and Davidson Zhu. Multi-quadratic sum-of-squares lower bounds imply $\mathsf{VNC}^1 \neq \mathsf{VNP}$. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, volume 362 of *LIPIcs*, 2026. doi: 10.4230/LIPIcs.ITCS.2026.113. Article 113.