Covariance-Spectrum Certificates for Homogeneous Stable Count Laws

Scott Schubr

Luminary Oddities, Illinois, USA
E-mail address: scottschubr@luminaryoddities.com
URL: https://orcid.org/0009-0009-7357-4439

Abstract. We compare a homogeneous real-stable count law with the multinomial distribution having the same mean. The normalized covariance deficit determines the quadratic orthogonal component of the likelihood ratio and controls its higher-order remainder through explicit spectral bounds. These estimates give a uniform quadratic approximation in chi-square distance as the trace deficit vanishes, and convergence for diffuse covariance-deficit spectra with bounded trace, allowing the degree and number of categories to vary. They also yield a sharp covariance-matching exponent, while nonfactorizable stable cubics show why a stronger Frobenius-power estimate for categorical products cannot extend to the full stable class. An optimal all-degree centered-real elementary-symmetric inequality, with equality classification, supplies the coefficient input and gives an optimal linear total-variation certificate. Separately, a sharp centered-complex inequality in six coordinates, including its complete equality orbit, yields the optimal uniform factor $3/4$ in the complete-homogeneous bound for the squared L^2 norm of the fifth likelihood component of degree-six stable laws. Ternary determinantal representations transfer further complete-homogeneous coefficient bounds, and a fixed-law convolution theorem identifies the limiting Gaussian total variation. The arguments adapt classical tensor, polynomial and orthogonal-expansion methods; standalone exact-arithmetic code reconstructs the finite certificates used in the proofs.

1. Setting and relation to prior work

A nonzero polynomial is *real stable* if it has real coefficients and no zero when every variable has positive imaginary part. Delete categories of mean zero, whose counts vanish almost surely, and let m be the number of retained categories. Let

$$p(x) = \sum_{|a|=n} P(a)x^a, \qquad P(a) \geq 0, \qquad p(\mathbf{1}) = 1$$

be homogeneous and real stable, with integer $n \geq 1$. Here $|a| = \sum_i a_i$ and A has law P. Put $t_i = \mathbb{E}A_i > 0$, $q = t/n$, $L(x) = q \cdot x$, and $M = \mathrm{Mult}(n, q)$. Total variation is $\mathrm{TV}(P, M) = \frac{1}{2}\sum_a |P(a) - M(a)|$. Define

$$B = \frac{1}{2}\left(\mathrm{diag}(t) - \frac{tt^\top}{n} - \mathrm{Cov}(A)\right), \qquad W = \mathrm{Tr}(\mathrm{diag}(t)^{-1}B). \tag{1.1}$$

We will prove $B \succeq 0$. The identity $\sum_i t_i = n$ gives $W = \frac{1}{2}(m - 1 - \sum_i \mathrm{Var}(A_i)/t_i)$.

2020 *Mathematics Subject Classification.* 60E15, 60F05.

Key words and phrases. real-stable count laws, multinomial approximation, covariance deficit, elementary symmetric inequalities.

The author received no external research funding. The author owns Luminary Oddities, which operates the Luminary Research Archive hosting the preprint, and declares no additional relevant competing interests.

The covariance-spectrum theorem. The density components U_k are the explicit degree-k orthogonal multinomial components constructed in Lemma 2.2. In particular their sum is $P/M - 1$; all norms below refer to the mean-matched multinomial law. Put $T = \mathrm{diag}(t)$ and

$$D = 2T^{-1/2}BT^{-1/2}, \quad x = \mathrm{Tr}\, D = 2W, \quad s_2 = \mathrm{Tr}(D^2), \quad a_n = \frac{n}{2(n-1)} \quad (n \geq 2).$$

The matrix D annihilates $\sqrt{t}$ and satisfies $0 \preceq D \preceq I - \sqrt{t}\sqrt{t}^{\top}/n$, by covariance positivity. For $r \geq 0$ define the nonnegative finite spectral polynomials

$$g_r(D) = [z^r] \det(I - 2zD)^{-1/2}, \qquad b_r(D) = \min\{x^r, g_r(D)\}.$$

These coefficients are defined for every PSD D, regardless of convergence at $z = 1$.

Theorem 1.1. *For $n \geq 2$, the exact quadratic density in Lemma 2.2 satisfies*

$$\|U_2\|_2^2 = a_n s_2. \tag{1.2}$$

For $3 \leq k \leq n$,

$$\|U_k\|_2^2 \leq \min\{s_2 x^{k-2},\ s_2 g_{k-2}(D),\ g_k(D)\}. \tag{1.3}$$

Consequently, with $J_n(D) = \sum_{r=1}^{n-2} b_r(D)$,

$$\chi^2(P\|M) = a_n s_2 + \|P/M - 1 - U_2\|_2^2, \tag{1.4}$$

$$0 \leq \|P/M - 1 - U_2\|_2^2 \leq s_2 J_n(D), \tag{1.5}$$

$$a_n s_2 \leq \chi^2(P\|M) \leq s_2(a_n + J_n(D)). \tag{1.6}$$

All norms are in $L^2(M)$ and all sums are finite; no smallness assumption on x is needed. If $x = 0$, all terms vanish by rigidity. Empty sums apply at $n = 2$.

Relation to earlier work. Multinomial approximation for dependent random variables is classical; Loh's Stein-method treatment (1992) is a central predecessor. Ehm's binomial approximation (1991) already treats the bivariate stable subclass, whose generating polynomials factor into linear forms. Roos (2010) develops approximation of convolution products, including independent categorical sums.

Loh's Theorems 5 and 6 treat, respectively, independent categorical summands and dependent multivariate hypergeometric sampling. The present result supplies a uniform covariance-deficit certificate.

The closest methodological predecessor is Han and Niles-Weed (2026a). Their analysis of permutation mixtures combines centered elementary symmetric polynomials, Banach transfer and orthogonal chi-square expansions. Their Theorem 4.3(2), equation (27), supplies the real coefficient estimate refined below. Related Banach and coefficient methods occur in Roos (2014); later centered-vector and matrix inequalities appear in Han and Niles-Weed (2026b). Our argument adapts these tools to the centered hyperbolic expansion of an arbitrary homogeneous stable count law. For categorical observations, a single permutation mixture of independent coordinates has the same count generating polynomial as the underlying product of categorical linear forms. A general exchangeable mixture representation instead yields a mixture of independent reference laws; it does not directly give our one mean-matched multinomial or the deficit in (1.1).

The main refinement is the covariance-spectrum bound and uniform quadratic likelihood normal form in Theorem 1.1, together with the sharp matching exponent and the full-stable cubic obstruction in Theorem 6.1. The complete real scalar coefficient-one certificate, its equality classification, optimal TV coefficient two and fixed-law Gaussian mutual-TV limit supply the surrounding certification theory. Banach's theorem, the predecessor centered ESP method, hyperbolic-slice extremality, and the orthogonal-expansion architecture are classical inputs used throughout. Fixed degree and category count are separate parameters throughout. The sharp complex-six theorem

also determines its complete equality orbit, linking the finite-dimensional coefficient certificate to an explicit extremal geometry. Its rigidity argument is integrated with the same weighted-median and real-rooted proof below.

Spectral and Frobenius improvements are already present in Han–Niles-Weed (2026a, Section 5) and (2026b, Corollary 1.3) for permanent/permutation-product models. Roos (2002) already studies signed multivariate Krawtchouk corrections for independent categorical sums, and multinomial orthogonal projections have classical reproducing-kernel descriptions Diaconis and Griffiths (2019). Our bounds cover arbitrary nonfactorizable homogeneous stable count laws, with explicit spectral-square remainder and a sharp obstruction to extending a product-class Frobenius power. The Gaussian tensor-norm and hyperbolicity tools in that obstruction, and quantitative AM–GM at the scalar endpoint, are also classical and explicitly attributed below.

The all-real two-moment estimate of Derksen, Ivanov, Lee and Viola (2024, Lemma 36) is another important predecessor. On the centered squared-norm-n slice it gives $|e_k| \le \binom{n}{k}((k-1)/(n-1))^{k/2}$. Its squared ratio to $\binom{n}{k}$ is $32/25$ at $n = 6, k = 3$, so that displayed estimate alone does not give the value-one bound proved below. Their real-rooted truncation argument and its earlier ancestry are established methods. For binomial-orthonormal biased Krawtchouk polynomials, Hsieh et al. (2026, Lemma 3.12) give an $O(k^{1/4})e^{y^2/4}$ bound under fixed-bias and spatial-window hypotheses. Our statement concerns the exact finite bound one at every admissible integer mean and every order. The ESP–Krawtchouk correspondence, root reduction, recurrence and Hermite picture are classical inputs. The fixed complex case $(n, k) = (6, 5)$ is proved separately below.

Theorem 1.2. *Let $n \ge 1$ and let $z \in \mathbb{R}^n$ satisfy $\sum_i z_i = 0$. For $1 \le k \le n$,*

$$|e_k(z)|^2 \le \binom{n}{k}\left(\frac{\sum_i z_i^2}{n}\right)^k. \tag{1.7}$$

Also $e_0 = 1$. The universal squared coefficient one is optimal.

Theorem 1.3. *For every integer $n \ge 1$ and every normalized homogeneous real-stable count law with nonnegative coefficients,*

$$\chi^2(P\|M) \le \sum_{k=2}^{n}(2W)^k, \qquad \mathrm{TV}(P, M) \le \min\{1, 2W\}. \tag{1.8}$$

Mean-zero categories are deleted before defining W using the retained positive means. The conclusion allows arbitrary category count, repeated roots, zero coefficients and singular covariance deficit. The coefficient two is optimal uniformly over the degree.

The proof is computer-assisted: the exact polynomial identities, infinite-domain reductions and exhaustive finite domains are specified below, with standalone code reconstructing their complete certificates. The scalar theorem is proved by the complete domain partition below, rather than assumed from finite samples.

2. Centered expansion and the exact density identity

Degree one is exactly multinomial. In the lemmas below we therefore assume $n \ge 2$.

Lemma 2.1. *There are homogeneous polynomials R_k of degree k, $2 \le k \le n$, invariant under $x \mapsto x + a\mathbf{1}$, such that*

$$p(x) = L(x)^n + \sum_{k=2}^{n} R_k(x)L(x)^{n-k}. \tag{2.1}$$

Moreover $B \succeq 0$, $B\mathbf{1} = 0$, and $R_2(x) = -x^\top Bx$.

Proof: For real x, set $y = x - L(x)\mathbf{1}$. The polynomial $p(z\mathbf{1} + y)$ is monic of degree n. Every root is real: an upper-half-plane root contradicts stability, and a lower-half-plane root has an upper conjugate. Its coefficient of z^{n-1} is $t \cdot y = 0$. Write its roots as $\lambda_1, \dots, \lambda_n$, so $\sum_j \lambda_j = 0$, and let $R_k(x)$ be its coefficient of z^{n-k}. This defines homogeneous polynomials, invariant under addition of $a\mathbf{1}$ because y is unchanged. Substitution $z = L(x)$ gives (2.1).

Write $H = \nabla^2 p(\mathbf{1})$. The coefficient of z^{n-2} is $\frac{1}{2} y^\top H y$. Euler's identity gives $H\mathbf{1} = (n-1)t$ and $\mathbf{1}^\top H \mathbf{1} = n(n-1)$, so

$$R_2(x) = \tfrac{1}{2} x^\top \left(H - \frac{n-1}{n} t t^\top \right) x.$$

Since $H = \mathrm{Cov}(A) + tt^\top - \mathrm{diag}(t)$, this is $-x^\top B x$. On the other hand the root identity gives

$$R_2(x) = -\tfrac{1}{2} \sum_j \lambda_j^2, \qquad \sum_j \lambda_j^2 = 2x^\top B x. \tag{2.2}$$

Thus $B \succeq 0$; the displayed formula or translation invariance gives $B\mathbf{1} = 0$. □

If $W = 0$, positive t and $B \succeq 0$ imply $B = 0$. Equation (2.2) then makes every centered root zero for every real x, so all R_k vanish and $P = M$. This proves rigidity without dividing by W. The cases $n = 1$ or $m = 1$ are also exact multinomial laws. Zero-mean categories can be deleted because their nonnegative counts vanish almost surely.

For a homogeneous polynomial R_k, let T_k denote its symmetric k-linear form, with $R_k(x) = T_k(x, \dots, x)$. Our ordered-index convention means that a monomial coefficient at α is $(k!/\alpha!)T_{k,\alpha}$.

Lemma 2.2. *Let ν_k be the signed coefficient measure of $R_k L^{n-k}$. There exist mutually orthogonal mean-zero densities $U_k = d\nu_k/dM$ satisfying*

$$\frac{P}{M} - 1 = \sum_{k=2}^{n} U_k, \qquad \chi^2(P\|M) = \sum_{k=2}^{n} \mathbb{E}_M U_k^2, \tag{2.3}$$

and

$$\mathbb{E}_M U_k^2 = \binom{n}{k}^{-1} \sum_{i_1, \dots, i_k} \frac{T_k(e_{i_1}, \dots, e_{i_k})^2}{q_{i_1} \cdots q_{i_k}}. \tag{2.4}$$

Proof: Translation invariance of R_k gives $T_k(\mathbf{1}, v_2, \dots, v_k) = 0$: first differentiate the diagonal polynomial and then polarize its remaining arguments. Let $X_1, \dots, X_n$ be independent categorical variables of law q and set

$$H_k(i_1, \dots, i_k) = \frac{T_k(e_{i_1}, \dots, e_{i_k})}{q_{i_1} \cdots q_{i_k}}, \qquad U_k = \binom{n}{k}^{-1} \sum_{|S|=k} H_k(X_S).$$

Symmetry makes U_k a function of the count vector. Summing any one argument of H_k against q gives zero by the contraction identity. Hence the kernel is completely conditionally degenerate. For each fixed subset S,

$$\mathbb{E}[H_k(X_S) x^{\mathrm{counts}(X)}] = R_k(x) L(x)^{n-k}.$$

Coefficient comparison after symmetrization proves $\nu_k(a) = M(a) U_k(a)$. All $M(a)$ on $|a| = n$ are positive. In a product of two kernels associated with different subsets, an index in the symmetric difference occurs in only one kernel; integrate that independent categorical observation to obtain zero. This applies both within an order and between distinct orders. Thus only identical subsets contribute to the square norm, giving (2.4), while (2.1) gives (2.3). All operations are finite sums. □

3. Hilbert tensor transfer and a universal certificate

We use Banach's classical real Hilbert-space theorem Banach (1938); Carando and Rodríguez (2018): the norm of a symmetric multilinear form on arbitrary unit-vector tuples equals its norm on diagonal tuples. It is also the input in Han and Niles-Weed (2026a, Lemma 4.2).

Lemma 3.1. *If* $|R_k(x)| \le c(x^\top Bx)^{k/2}$ *for all real* x, *then*

$$\mathbb{E}_M U_k^2 \le \frac{c^2(nW)^k}{\binom{n}{k}}. \tag{3.1}$$

Proof: For $v \in \ker B$, the polynomial $R_k(x+sv)$ is bounded independently of real s, hence is constant. Differentiation and polarization give $T_k(v, v_2, \dots, v_k) = 0$. Thus T_k descends to the real Hilbert quotient by $\ker B$, with squared norm $x^\top Bx$. Banach's theorem there implies

$$|T_k(v_1, \dots, v_k)| \le c \prod_j \sqrt{v_j^\top B v_j}.$$

Apply this to coordinate vectors in (2.4) and sum the products to obtain

$$\mathbb{E}_M U_k^2 \le \binom{n}{k}^{-1} c^2 \left(\sum_i \frac{B_{ii}}{q_i} \right)^k = \binom{n}{k}^{-1} c^2 (nW)^k.$$

No inverse of B is used. □

Han and Niles-Weed (2026a, Theorem 4.3(2)) prove the uniform real centered coefficient bound $|e_k(z)|^2 \le 10\binom{n}{k}$ when $\sum_j z_j^2 = n$. Below we refine their binary contour and Beta-integral argument to a constant below three. The classical extremum reduction needed for that refinement is stated first; the bound with constant ten serves as a comparison, rather than as a tail input in our final theorem.

The next elementary bound retains positivity of the actual density. It is the upper-distance branch of the classical sharp relation between chi-square divergence and total variation; see Guntuboyina, Saha and Schiebinger (2013, Corollary 5.3), which credits Gilardoni's Proposition 2.1. We give the short proof needed here.

Lemma 3.2. *For probability measures* $P \ll M$, *write* $v = \chi^2(P\|M)$ *and* $\delta = \mathrm{TV}(P, M)$. *If* $\delta > 1/2$, *then* $v \ge \delta/(1-\delta)$.

Proof: Let $Y = dP/dM - 1$ and $a = M(Y < 0)$. Then $\mathbb{E}Y_+ = \mathbb{E}Y_- = \delta$ and $Y_- \le 1$, so $a \ge \delta$. Cauchy–Schwarz on the two parts of the partition gives

$$v \ge \frac{\delta^2}{a} + \frac{\delta^2}{1-a} = \frac{\delta^2}{a(1-a)}.$$

For $a \ge \delta > 1/2$, $a(1-a) \le \delta(1-\delta)$. The degenerate partition cases force $\delta = 0$ and do not occur here. □

The probability consequences are grouped together below; the detailed real coefficient proof and its finite-domain certificates follow afterward. The finer coefficient estimates retain their separate degree and rank hypotheses.

4. Proof of the covariance-spectrum bounds

The derivation of Theorem 1.1 uses the sharp real coefficient bound in Theorem 1.2. That bound is proved independently in Section 7, without using the count-law conclusions of this section. This order keeps the likelihood estimates and their probability consequences together.

The gain over trace-only control retains s_2/x^2, the reciprocal effective rank when $x > 0$. Spectral permanent bounds already occur in Han–Niles-Weed (2026a, Section 5); their later matrix theorem and Corollary 1.3 also retain Frobenius information Han and Niles-Weed (2026b). The result here concerns the full homogeneous stable class and its explicit covariance-controlled remainder.

Lemma 4.1. *If S is a symmetric real k-linear form of multilinear norm at most one and $\Lambda \succeq 0$ has eigenvalues λ_i, then, for $k \geq 2$,*

$$\sum_{i_1,\dots,i_k} \lambda_{i_1} \cdots \lambda_{i_k} S(e_{i_1}, \dots, e_{i_k})^2 \leq \operatorname{Tr}(\Lambda^2)(\operatorname{Tr}\Lambda)^{k-2}.$$

At $k = 2$ the right side means $\operatorname{Tr}(\Lambda^2)$.

Proof: Fix the last $k-2$ unit basis arguments. The symmetric matrix $C_{ij} = S(e_i, e_j, \dots)$ is an operator contraction. Thus

$$\sum_{i,j} \lambda_i \lambda_j C_{ij}^2 \leq \tfrac{1}{2} \sum_{i,j} (\lambda_i^2 + \lambda_j^2) C_{ij}^2 = \sum_i \lambda_i^2 \sum_j C_{ij}^2 \leq \sum_i \lambda_i^2.$$

Symmetry justifies the equality and the row norms are at most one. Sum the remaining nonnegative weights. This is an elementary weighted Hilbert–Schmidt contraction argument. □

Proof of Theorem 1.1: The scalar theorem gives $|R_k(v)|^2 \leq \binom{n}{k}(2/n)^k (v^\top B v)^k$. Use the quotient and Banach argument of Lemma 3.1. The weighted coordinate family has covariance

$$\Lambda = B^{1/2} \operatorname{diag}(q)^{-1} B^{1/2}.$$

Its nonzero spectrum equals that of $(n/2)D$ by the $XX^\top, X^\top X$ identity. Lemma 4.1 and the exact tensor identity cancel the factors $\binom{n}{k}$ and $(2/n)^k (n/2)^k$, giving $\|U_k\|_2^2 \leq s_2 x^{k-2}$. For $k = 2$, use $T_2 = -B$ directly:

$$\|U_2\|_2^2 = \frac{n^2}{\binom{n}{2}} \operatorname{Tr}\left((T^{-1/2} B T^{-1/2})^2\right) = a_n s_2.$$

No inverse of B is used.

For clarity we give the Gaussian steps, with all normalizations. The whitened polynomial $F_k(y) = R_k(\operatorname{diag}(q)^{-1/2} y)$ has ordered symmetric tensor A_k and

$$|F_k(y)|^2 \leq \binom{n}{k} (y^\top D y)^k, \qquad \|U_k\|_2^2 = \|A_k\|_{\mathrm{F}}^2 / \binom{n}{k}.$$

For a standard real Gaussian G, the degree-k Hermite component of $F_k(G)$ has squared norm $k! \|A_k\|_{\mathrm{F}}^2$. Indeed the coefficient of the product probabilists' Hermite H_α is $(k!/\alpha!) A_{k,\alpha}$, and $\mathbb{E} H_\alpha^2 = \alpha!$. Lower-degree chaos squares are nonnegative. Hence

$$\|U_k\|_2^2 \leq \mathbb{E}(G^\top D G)^k / k! = g_k(D).$$

The last identity follows by integrating independent Gaussian squares after diagonalization, initially near $z = 0$, and comparing finite coefficients.

For the stronger hybrid bound, kernel descent expresses $A_k = \sqrt{\binom{n}{k}}\, S \circ (D^{1/2})^{\otimes k}$ with S symmetric of multilinear norm at most one on $\operatorname{ran} D$. Work in an eigenbasis and put $r = k - 2$ and

$$P_{ij}(G) = S(e_i, e_j, D^{1/2} G, \dots, D^{1/2} G).$$

Its highest Hermite chaos has squared norm $r!$ times the remaining r-index weighted tensor square. Multiply by $\lambda_i \lambda_j$, sum, and discard the nonnegative lower-chaos norms. Pointwise Lemma 4.1's matrix argument bounds the inner bilinear sum by $s_2 (G^\top D G)^r$. Thus

$$r! \|U_k\|_2^2 \leq s_2 \mathbb{E}(G^\top D G)^r, \qquad \|U_k\|_2^2 \leq s_2 g_r(D).$$

For $r = 0$ this is just the matrix argument, while the exact quadratic identity is sharper. Orthogonality proves the remaining assertions.

□

Corollary 4.2. *If $\rho = \|D\|_{\mathrm{op}} < 1/2$, then*

$$\chi^2(P\|M) \le s_2\{a_n + \det(I-2D)^{-1/2} - 1\} \le s_2\{a_n + \exp(x + s_2/(1-2\rho)) - 1\}.$$

For any sequence in the full stable class with bounded x and $s_2 \to 0$, one has $\chi^2(P\|M) \to 0$ and $\mathrm{TV}(P, M) \to 0$, even when degree, category count and positive means vary.

Proof: The convergent generating series sums the hybrid coefficients. Also

$$\log\det(I-2D)^{-1/2} = x + \sum_{j\ge 2} \frac{2^{j-1}}{j}\,\mathrm{Tr}(D^j) \le x + \frac{s_2}{1-2\rho},$$

since $\mathrm{Tr}(D^j) \le \rho^{j-2} s_2$ and $2^{j-1}/j \le 2^{j-2}$. Finally $\rho \le \sqrt{s_2}$, so the required spectral-radius condition holds eventually. □

For the product/permutation-mixture subclass, convergence from small Frobenius norm already follows from Han and Niles-Weed (2026b, Corollary 1.3). Our assertion extends this convergence criterion to the full homogeneous stable class. The determinant above is not the Gaussian-likelihood expression $\det(I-D^2)^{-1/2} - 1$: for $p = (XY + XZ + YZ)/3$ the latter is $1/3$, whereas the actual chi-square is $1/2$ and D has spectrum $(1/2, 1/2, 0)$.

4.1. *Quadratic approximation and covariance matching.* Put $H_n(x) = \sum_{r=1}^{n-2} x^r$. Since $J_n(D) \le H_n(x)$,

$$0 \le \frac{\chi^2(P\|M)}{a_n s_2} - 1 \le \frac{H_n(x)}{a_n} \le \frac{2x}{1-x} \quad (s_2 > 0,\ 0 \le x < 1).$$

This is a small-deficit likelihood normal form, uniform in degree, categories and positive means. The exact quadratic is

$$U_2(a) = -\frac{n}{n-1}\left[(a-t)^\top T^{-1} B T^{-1}(a-t) - \sum_i \frac{B_{ii} a_i}{t_i^2}\right]. \tag{4.1}$$

To verify it, the ordered distinct observation pairs in the degenerate kernel have multiplicities $a_i a_j - \delta_{ij} a_i$. Their contraction is divided by $n(n-1)$; substitute $q = t/n$ and use $B\mathbf{1} = 0$ to center a. The signed measure $M(1 + U_2)$ need not be nonnegative. Signed multivariate Krawtchouk approximations for independent categorical sums are classical Roos (2002), as are multinomial orthogonal projections Diaconis and Griffiths (2019); the new estimate controls this correction's full-stable remainder.

With $\delta_2 = \frac{1}{2}\mathbb{E}_M|U_2|$, reverse triangle and Cauchy give

$$|\,\mathrm{TV}(P, M) - \delta_2| \le \tfrac{1}{2}\sqrt{s_2 J_n(D)}, \qquad \mathrm{TV}(P, M) \le \tfrac{1}{2}\sqrt{s_2(a_n + J_n(D))}.$$

If P, Q have the same n, t and covariance, they share U_2, D and M, whence

$$\mathrm{TV}(P, Q) \le \min\{1, \sqrt{s_2 J_n(D)}\} \le \frac{x^{3/2}}{\sqrt{1-x}} \quad (0 \le x < 1). \tag{4.2}$$

The first bound follows by applying Cauchy to each higher-chaos remainder; the last uses $J_n \le x/(1-x)$ and $s_2 \le x^2$.

The exponent $3/2$ in this uniform matching conclusion is sharp. For $0 < \epsilon < 1/4$, take three independent Bernoulli probabilities $(1/2 + \epsilon, 1/2 + \epsilon, 1/2 - 2\epsilon)$, and let P_- reverse the sign of ϵ in P_+. With $L = (X + Y)/2$,

$$p_\pm = L^3 - 3\epsilon^2 (X-Y)^2 L \mp 2\epsilon^3 (X-Y)^3.$$

They have $t = (3/2, 3/2)$, $B = 3\epsilon^2 \left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$, $x = 8\epsilon^2$, and $s_2 = x^2$. Direct use of the exact tensor norm gives

$$\|U_2\|_2^2 = 48\epsilon^4, \quad \|U_3\|_2^2 = 256\epsilon^6 = x^3/2, \quad \frac{\chi^2(P_+\|M)}{a_3 s_2} - 1 = \tfrac{2}{3}x.$$

The absolute coefficient sum of $p_+ - p_- = -4\epsilon^3(X-Y)^3$ is $32\epsilon^3$, so $\mathrm{TV}(P_+, P_-) = 16\epsilon^3 = x^{3/2}/\sqrt{2}$. Thus neither the relative $O(x)$ remainder nor the matching $O(x^{3/2})$ exponent improves uniformly.

4.2. *Rare cells and a nonproduct illustration.* Small deficit does not uniformly imply the local Taylor ratios $\mathrm{KL}/\chi^2 \to 1/2$ and $H^2/\chi^2 \to 1/4$, where $H^2 = \sum(\sqrt{P} - \sqrt{M})^2$. For $n = 2$ and $0 < q < 1/2$, independent Bernoulli probabilities $0, 2q$ give

$$P = (1-2q, 2q, 0), \quad M = ((1-q)^2, 2q(1-q), q^2), \quad x = q/(1-q), \quad \chi^2 = x^2.$$

Here

$$\mathrm{KL}(P\|M) = (1-2q)\log\frac{1-2q}{(1-q)^2} + 2q\log\frac{1}{1-q} = q^2 + O(q^3),$$

and $H^2 = q^2 + O(q^3)$ as well. The latter follows by expanding the affinity $(1-q)\sqrt{1-2q} + 2q\sqrt{1-q}$; the absent count-two mass provides its leading deficit. Both ratios therefore tend to one. By contrast, probabilities $1/2 \pm \epsilon$ have a fixed positive reference law and yield the ordinary ratios $1/2, 1/4$ by Taylor expansion. This is a barrier to a uniform local expansion, not to classical divergence inequalities.

For a diffuse nonproduct example, let $p = e_d(X_1, \dots, X_m)/\binom{m}{d}$, $2 \le d < m$. This classical polynomial is real stable, with $t_i = d/m$ and D eigenvalue $(d-1)/(m-1)$ repeated $m-1$ times. Thus $x = d-1$ and $s_2 = (d-1)^2/(m-1)$. For fixed $d \ge 3$, $x \ge 2$ yet Corollary 4.2 applies as $m \to \infty$. Directly, with $c = (m)_d/m^d$, the multinomial conditional on no collisions is P, giving the classical birthday identities

$$\mathrm{TV}(P, M) = 1 - c, \qquad \chi^2(P\|M) = c^{-1} - 1 \sim d(d-1)/(2m) \sim a_d s_2.$$

No categorical-product representation exists: multiaffinity forces distinct linear factors to have disjoint variable supports, which cannot generate every d-subset when $m > d$. These classical birthday formulas illustrate the full stable domain.

4.3. *The complete-homogeneous target in low orders.* Let $h_k(D) = [z^k]\det(I - zD)^{-1}$. The following partial conclusion makes its present domain precise.

Proposition 4.3. *The target holds at $k = 2, 3$ for every admissible degree, and at $k = 4$ for every $n \ge 5$. At $k = 3$, or $k = 4, n \ge 5$, it is strict when $D \ne 0$. Rank-one D is covered at every order. The top-order theorem below supplies the remaining quartic degree $n = 4$.*

Proof: The exact quadratic energy is at most $s_2 \le h_2 = (x^2 + s_2)/2$. Newton gives $h_3 - xs_2 = e_3(\lambda(D)) \ge 0$, so Theorem 1.1 already proves the non-strict cubic assertion. For strictness, the centered scalar cubic maximum from the binary calculation in Lemma 7.6 is $\beta_{n,3} = 2n(n-2)/[3(n-1)^2] < 2/3$. A squared scalar constant β multiplies both spectral-square and hybrid bounds by β, since the normalized form's norm is at most $\sqrt{\beta}$. Thus $\|U_3\|_2^2 \le \beta_{n,3} s_2 x < h_3$ for nonzero D.

For every nonzero PSD spectrum,

$$h_4(D) > \tfrac{3}{5} s_2 g_2(D), \qquad g_2 = x^2/2 + s_2. \tag{4.3}$$

To prove it normalize $x = 1$, and put $p = \mathrm{Tr}\, D^2$, $q = \mathrm{Tr}\, D^3$, $v = \mathrm{Tr}\, D^4$. Here $0 < p \le 1$, $q \ge p^2$ and $v \ge p^3$ by Cauchy/Jensen with weights given by the eigenvalues. Newton gives $24h_4 = 1 + 6p + 3p^2 + 8q + 6v$. Therefore five times the lower bound for $24[h_4 - (3/5)p(1/2 + p)]$ is

$$5 - 6p - 17p^2 + 30p^3 = \tfrac{3}{2} - \tfrac{1}{2}w + (28 + 30w)w^2 \ge 13(w - 1/52)^2 + \tfrac{3}{2} - \tfrac{1}{208} > 0, \quad w = p - \tfrac{1}{2} \in [-1/2, 1/2].$$

Homogeneity restores (4.3).

For $n \geq 6$ the centered scalar quartic bound of Lemma 7.6 is at most $3/5$: its comparison reduces to $9n^2 - 66n + 72 \geq 0$, zero at six and increasing thereafter. At $n = 5$ the two reflected multiplicities $r = 1, 2$ give squared scalar ratios $45/256$ and $5/9$, respectively, both below $3/5$; these follow directly from the raw generating products. Scaling the hybrid bound by these scalar constants and applying (4.3) proves the quartic assertion. If rank $D = 1$, then $h_k(D) = x^k$ and the trace theorem suffices at every order. □

4.4. *Top-order complete-homogeneous bound by ternary restriction.*

Theorem 4.4. *For every degree $n \geq 2$ and every admissible covariance deficit, including singular D,*

$$\|U_n\|_2^2 \leq h_n(D).$$

Consequently the complete-homogeneous target holds at $k = 4$ for every admissible degree. More strongly, for every $w \in \mathbb{C}^m$,

$$|R_n(w)|^2 \leq \left(\frac{2}{n} w^* B w\right)^n. \tag{4.4}$$

Proof: We use the classical Helton–Vinnikov theorem: every degree-n ternary hyperbolic form has a definite real symmetric degree-sized linear determinantal representation. See Hanselka (2015, Theorem 1), including repeated-root and nonsmooth forms. This deep ternary theorem is an input; no global determinantal representation of p is assumed.

Fix $w = a + ib$. Put $a_0 = a - L(a)\mathbf{1}$, $b_0 = b - L(b)\mathbf{1}$ and $F(t, s, u) = p(t\mathbf{1} + sa_0 + ub_0)$. Real stability and conjugation imply hyperbolicity in $(1, 0, 0)$; its leading t coefficient is one even if these directions are dependent. The centered expansion is

$$F(t, s, u) = t^n + \sum_{j=2}^{n} R_j(sa_0 + ub_0) t^{n-j}.$$

The classical theorem and congruence normalization give $F = \det(tI_n + sA + uC)$ with real symmetric A, C. Indeed the definite coefficient at $(1, 0, 0)$ may be chosen positive, and its determinant is $F(1, 0, 0) = 1$, so the congruence introduces no scalar factor. Vanishing of the first centered coefficient gives $\operatorname{Tr} A = \operatorname{Tr} C = 0$. Comparison of the second coefficient yields

$$-\tfrac{1}{2} \operatorname{Tr}(sA + uC)^2 = -(sa_0 + ub_0)^\top B (sa_0 + ub_0).$$

In particular $\operatorname{Tr} A^2 = 2a^\top Ba$ and $\operatorname{Tr} C^2 = 2b^\top Bb$, using $B\mathbf{1} = 0$. Translation invariance and the determinant identity at $(t, s, u) = (0, 1, i)$ imply $R_n(w) = \det(A + iC)$. For its singular values σ_j, AM–GM gives

$$|R_n(w)|^2 = \prod_{j=1}^{n} \sigma_j^2 \leq \left(\frac{\sum_j \sigma_j^2}{n}\right)^n, \qquad \sum_j \sigma_j^2 = \operatorname{Tr}[(A - iC)(A + iC)] = 2w^* B w.$$

The commutator has zero trace. This proves (4.4), including zero right side, without inverting B.

Put $Q = \operatorname{diag}(q)$. Let Z be standard circular complex Gaussian with independent coordinates of density $\pi^{-1} e^{-|z|^2}$. For the ordered symmetric tensor A_n of $R_n(Q^{-1/2} z)$, complex monomial orthogonality gives

$$\mathbb{E}|R_n(Q^{-1/2} Z)|^2 = n! \|A_n\|_F^2, \qquad \|U_n\|_2^2 = \|A_n\|_F^2.$$

The second identity uses the exact density denominator $\binom{n}{n} = 1$. Since $D = (2/n) Q^{-1/2} B Q^{-1/2}$, integration of (4.4) yields

$$\|U_n\|_2^2 \leq \frac{\mathbb{E}(Z^* D Z)^n}{n!} = h_n(D).$$

For the last equality, unitary diagonalization gives independent unit-mean exponential squared moduli, with generating function $\prod_i (1 - z\lambda_i)^{-1}$. This is a finite coefficient identity at arbitrary

PSD D, not a Gaussian approximation or a spectral-radius restriction. The quartic conclusion now follows from Proposition 4.3 and the case $n = 4$ just proved. At $n = 1$ the centered density is zero. □

For clarity, applying the same argument to the normalized degree-k derivative $(k!/n!)D_{\mathbf{1}}^{n-k}p$ only gives

$$\|U_k\|_2^2 \le \binom{n}{k}\left(\frac{k-1}{n-1}\right)^k h_k(D).$$

Its top coefficient is $R_k/\binom{n}{k}$ and its normalized deficit is $(k-1)D/(n-1)$. The displayed factor can exceed one. The degree-six case is closed below. The real-scalar and complex-vector branches use their respective coefficient hypotheses.

Theorem 4.5 (All-order comparison from the known complex factor)**.** *For every $2 \le k \le n$,*

$$\|U_k\|_2^2 \le \min\{24, \beta_{n,k}^{\mathbb{C}}\}h_k(D), \qquad \beta_{n,k}^{\mathbb{C}} = \binom{n}{k}\left(\frac{k-1}{n-1}\right)^k. \tag{4.5}$$

The factor 24 is the known complex-vector constant of Han–Niles-Weed. All coefficients are finite at arbitrary PSD D. In particular, if $\rho(D) < 1$,

$$\begin{aligned}\chi^2(P\|M) &\le 24\{\det(I-D)^{-1} - 1 - x\},\\ \chi^2(P\|M) &\le a_n s_2 + 24\{\det(I-D)^{-1} - 1 - x - h_2(D)\}.\end{aligned} \tag{4.6}$$

Proof: For each complex direction w, the ternary pencil in the preceding proof gives $M_w = A + iC$ with $\operatorname{Tr} M_w = 0$ and $\|M_w\|_F^2 = 2w^*Bw$. Its characteristic coefficients give $R_k(w) = e_k(\lambda(M_w))$, even if the matrix is not diagonalizable. Schur triangularization gives $\sum_j |\lambda_j(M_w)|^2 \le \|M_w\|_F^2$. Apply the scaled complex centered-vector theorem of Han–Niles-Weed (2026b, Theorem 1.1):

$$|R_k(w)|^2 \le 24\binom{n}{k}\left(\frac{2}{n}w^*Bw\right)^k.$$

If all eigenvalues vanish the assertion is immediate. For standard circular complex Gaussian Z, substitute $w = Q^{-1/2}Z$. Complex monomial orthogonality and the exact density normalization give

$$k!\|A_k\|_F^2 \le 24\binom{n}{k}\mathbb{E}(Z^*DZ)^k = 24\binom{n}{k}k!h_k(D), \qquad \|U_k\|_2^2 = \frac{\|A_k\|_F^2}{\binom{n}{k}}.$$

Cancellation proves the factor-24 branch, and the derivative argument above proves the other branch. Thus finite orthogonality gives $\chi^2(P\|M) \le 24\sum_{k=2}^n h_k(D)$. If $\rho(D) < 1$, extending the nonnegative generating series to all orders proves the first determinant bound. Keeping the exact quadratic energy $a_n s_2$ and extending only orders $k \ge 3$ proves the second. The determinant has power minus one, distinct from the real-Gaussian envelope's power minus one half. □

One may combine all these branches order by order:

$$\chi^2(P\|M) \le a_n s_2 + \sum_{k=3}^n \min\{s_2 x^{k-2},\ s_2 g_{k-2}(D),\ g_k(D),\ 24h_k(D),\ \beta_{n,k}^{\mathbb{C}}h_k(D)\}.$$

The representation is still chosen separately for each ternary restriction.

4.5. *A sharper full-spectrum bound at order five.* Order five admits a stronger uniform comparison in general degree. The proof uses the same real scalar-to-hybrid transfer, with an explicit fifth-order scalar constant, and elementary symmetric-polynomial inequalities.

Lemma 4.6. *On the centered real squared-norm-n slice, let $\beta_{n,5} = \max |e_5|^2/\binom{n}{5}$. For every $n \geq 6$,*

$$\beta_{n,5} \leq \tfrac{3}{4}.$$

The uniform constant is sharp, attained at $n = 6$ and binary multiplicity $r = 2$. For $n \geq 8$ one also has

$$\beta_{n,5} \leq B_n := \frac{4n(5n-12)^3}{1215(n-1)(n-2)(n-3)(n-4)}. \tag{4.7}$$

Proof: By the classical two-root reduction in Lemma 7.1, take multiplicities $r, s = n - r$. Put $c = (s-r)/\sqrt{rs}$ and $u = c^2 \in [0, U]$, where $U = (n-2)^2/(n-1)$. Newton identities give

$$e_5 = \frac{nc(6c^2 + 12 - 5n)}{30}, \qquad \beta = \frac{n^2 u(6u + 12 - 5n)^2}{900\binom{n}{5}}.$$

For $A = 5n - 12$, the derivative of $u(6u - A)^2$ is $(6u - A)(18u - A)$. Its maximum on $[0, U]$ is therefore bounded by the larger of its value at U and its local-maximum value at $A/18$, even when the latter lies outside the interval. Their normalized contributions are

$$E_n = \frac{2n(n-2)(n-3)(n-4)}{15(n-1)^4} < \tfrac{2}{15}, \qquad B_n.$$

For $n \geq 8$, $5n - 12 \geq (7/2)n$ and the denominator product in (4.7) is less than n^4, so $B_n > 343/2430 > 2/15$. This proves (4.7). Moreover $B_n < 3/4$ for $n \geq 8$: after clearing positive denominators, the difference is

$$\begin{aligned} &3645(n-1)(n-2)(n-3)(n-4) - 16n(5n-12)^3 \\ &\quad = 251944 + 468998y + 195495y^2 + 30590y^3 + 1645y^4 > 0, \quad y = n - 8 \geq 0. \end{aligned}$$

At $n = 6$, reflected multiplicities $r = 1, 2, 3$ give $96/3125, 3/4, 0$;
at $n = 7$ they give $7/162, 6776/9375, 7/64$. These direct rational evaluations finish the proof. □

Lemma 4.7 (Sharp fifth spectral comparison)**.** *For every PSD D,*

$$h_5(D) \geq \tfrac{3}{8} s_2 g_3(D). \tag{4.8}$$

For nonzero D, equality holds exactly when its positive spectrum consists of two equal eigenvalues. Thus $3/8$ is the best universal constant.

Proof: The zero case is immediate. Normalize $x = 1$ and write $p = p_2, q = p_3, r = p_4, s = p_5$, where $p_j = \mathrm{Tr}(D^j)$. Newton identities give

$$g_3 = \frac{1 + 6p + 8q}{6}, \qquad 120h_5 = 1 + 10p + 15p^2 + 20q + 20pq + 30r + 24s.$$

Hence $E = 120[h_5 - (3/8)pg_3]$ equals

$$E = 1 + \tfrac{5}{2}p - 30p^2 + (20 - 40p)q + 30r + 24s. \tag{4.9}$$

If $p \leq 1/2$, Cauchy/Jensen with the eigenvalues as probability weights yields $q \geq p^2, r \geq p^3, s \geq p^4$, and the q coefficient is nonnegative. Thus

$$E \geq (1 - 2p)(1 + \tfrac{9}{2}p - p^2 - 12p^3) \geq 0.$$

The second factor is concave on $[0, 1/2]$ and has endpoint values $1, 3/2$, so this is strict if $p < 1/2$.

For $p \geq 1/2$, let e_j be the elementary symmetric functions of the normalized eigenvalues. Another Newton rearrangement of (4.9) gives

$$E = \tfrac{15}{2}(2p - 1)(4 - 3p) + (240 - 60p)e_3 - 240e_4 + 120e_5.$$

Since $e_3 = xe_3 \geq 4e_4$ and $p \leq 1$,

$$E \geq \tfrac{15}{2}(2p-1)(4-3p) + (180-60p)e_3 + 120e_5 \geq 0.$$

Equality forces $p = 1/2$ and $e_3 = 0$, hence exactly two equal positive eigenvalues. Conversely that spectrum gives equality. For eigenvalues (a, a), $h_5 = 6a^5$, $s_2 = 2a^2$ and $g_3 = 8a^3$, proving sharpness. Homogeneity restores arbitrary x. □

Theorem 4.8 (Fifth-order full-spectrum refinement)**.** *For every admissible stable count law of degree $n \geq 6$,*

$$\|U_5\|_2^2 \leq \min\{2, \beta_{n,5}^{\mathbb{C}}\} h_5(D). \tag{4.10}$$

For $n \geq 8$ the sharper coefficient is

$$\|U_5\|_2^2 \leq C_n h_5(D), \qquad C_n = \min\{2, \beta_{n,5}^{\mathbb{C}}, (8/3)B_n\}. \tag{4.11}$$

In particular, $\|U_5\|_2^2 \leq (10/9)h_5(D)$ for every $n \geq 256$ and every spectrum. These bounds include singular D and impose no spectral-radius restriction. The third branch tends to $800/729 > 1$; no unrestricted factor-one conclusion follows. At $n = 5$ the top-order factor one is stronger.

Proof: A scalar squared constant β multiplies the hybrid bound by β, as in Proposition 4.3. Lemmas 4.6 and 4.7 give respectively the factors 2 and $(8/3)B_n$. The derivative branch is already established in (4.5). At $n = 6$ it is $6144/3125 < 2$, so it is retained.

To obtain the stated high-degree corollary, $B_n \leq 5/12$ is equivalent to

$$P(n) = 6075(n-1)(n-2)(n-3)(n-4) - 48n(5n-12)^3 \geq 0.$$

Its exact shift at $n = 256 + y$ is

$$P(256+y) = 34765845384 + 1638253434y + 16121745y^2 + 59250y^3 + 75y^4.$$

Every coefficient is positive, proving the assertion on the entire real ray $n \geq 256$ and hence every integer degree there. The limit of $(8/3)B_n$ is $(8/3)(500/1215) = 800/729$. The threshold is a sufficient one. The separate scalar and spectral inputs are sharp. □

Corollary 4.9 (Two fifth-order factor-one regions)**.** *Let $n \geq 6$ and $D \neq 0$. The factor-one bound holds if either*

$$\frac{s_2}{x^2} \leq \frac{1}{11} \quad \text{or} \quad \frac{\lambda_{\max}(D)}{x} \geq \frac{7}{8}.$$

The corresponding gaps $h_5 - \|U_5\|_2^2$ are at least $17x^5/53240$ and $x^5/4096$, respectively. For general degrees $n \geq 7$, an intermediate spectral region remains uncovered.

Proof: Normalize $x = 1$. For $p \leq 1/11$, use $\|U_5\|_2^2 \leq (3/4)pg_3$ and

$$\begin{aligned} 120[h_5 - (3/4)pg_3] &= 1 - 5p - 75p^2 + (20-100p)q + 30r + 24s \\ &\geq 1 - 5p - 55p^2 - 70p^3 \geq 51/1331. \end{aligned}$$

Here the q coefficient is positive, $q \geq p^2, r \geq p^3$, and the last polynomial decreases for $p \geq 0$. For $\alpha = \lambda_{\max}(D) \geq 7/8$, the spectral-square branch and the monomials of h_5 give

$$\|U_5\|_2^2 \leq \tfrac{3}{4}p \leq \tfrac{3}{4}[\alpha^2 + (1-\alpha)^2], \qquad h_5 \geq \alpha^5 + \alpha^4(1-\alpha) = \alpha^4.$$

Their difference is $f(\alpha) = \alpha^4 - (3/2)\alpha^2 + (3/2)\alpha - 3/4$, with $f(7/8) = 1/4096$ and $f'(\alpha) > 0$ on $[7/8, 1]$. Rescale by x^5. □

In any orderwise chi-square bound, the $k = 5$ summand may therefore be replaced by these stronger branches. The ingredients are the credited real root reduction and tensor transfer, together with the explicit fifth-degree calculation and sharp spectral comparison.

5. Sharp degree-six coefficients and stable-law transfer

Theorem 5.1. *For every $z \in \mathbb{C}^6$ with $\sum_i z_i = 0$,*

$$|e_5(z)|^2 \le \frac{\left(\sum_i |z_i|^2\right)^5}{1728}. \tag{5.1}$$

The constant is optimal. Equality holds precisely when $z = 0$ or when z is a permutation of $c(2,2,-1,-1,-1,-1)$ for some $c \in \mathbb{C} \setminus \{0\}$. Consequently every degree-six homogeneous stable count law, in every covariance-deficit rank, satisfies the sharp bound

$$\|U_5\|_2^2 \le \tfrac{3}{4} h_5(D). \tag{5.2}$$

In particular, together with the lower-order and top-order results, every order in degree six satisfies the complete-homogeneous factor-one target.

Proof: The zero vector is immediate. Otherwise homogeneity gives $\sum |z_i|^2 = 1$. Two zero coordinates give $e_5 = 0$; exactly one gives $|e_5|^2 \le 5^{-5} < 1/1728$ by AM–GM. Otherwise write $z_i = r_i u_i$, where $r_i > 0$, $|u_i| = 1$ and $\sum r_i u_i = 0$. With $P_6 = \prod r_i$, for every real c,

$$|e_5(z)| = P_6 \left| \sum_i \frac{\overline{u_i}}{r_i} \right| \le P_6 \sum_i |r_i^{-1} - c r_i|.$$

The right side is minimized at a weighted median $c = t^{-2}$ of the numbers r_i^{-2} with weights r_i, where $t = r_j$ is one radius. Choose $\sigma_i = 1$ for $r_i < t$ and $\sigma_i = -1$ for $r_i > t$. Signs at ties are free. If there are m tied radii, including j, the median condition bounds the signed strict-group sum J by $|J| \le mt$. The other $m-1$ tie signs give offsets $-(m-1)t, -(m-3)t, \dots, (m-1)t$; intervals of radius t around these offsets cover $[-mt, mt]$. Hence the signs can be chosen so that, for the five real numbers $x_i = \sigma_i r_i$ $(i \ne j)$,

$$L = \sum x_i, \qquad |L| \le t, \qquad t^2 + \sum x_i^2 = 1. \tag{5.3}$$

The minimized majorant is exactly

$$P_6 \sum_i |r_i^{-1} - r_i/t^2| = |t e_4(x) - (L/t) e_5(x)|\,. \tag{5.4}$$

Indeed $e_4(x) = e_5(x) \sum 1/x_i$, and every tied sign contributes zero. We may discard all remaining sign and ordering restrictions on x.

Fix $L, R = \sum x_i^2$ and t. The maximum of the absolute affine coefficient functional in (5.4) on the compact real-rooted slice with fixed $e_1 = L$ and $e_2 = (L^2 - R)/2$ occurs with at most two distinct roots. This is the classical hyperbolic two-slice principle, directly covered by Riener–Schabert (2023, Theorem 2.8 and Corollary 2.10). For completeness, maximize either sign and select a maximizing quintic f with the largest number d of distinct roots. If $3 \le d < 5$, write $f = hg$, with g containing each distinct root once. The perturbations $h(g + \epsilon) = f + \epsilon h$ preserve real-rootedness for both sufficiently small signs and preserve the first two coefficients because $\deg h \le 2$. A nonzero change of the functional improves one sign. A zero change preserves the maximum but splits a repeated root, increasing d. Both are contradictions. If $d = 5$, perturb by ϵX; both small signs preserve real-rootedness, while the nonzero coefficient t of e_4 improves one sign. Thus $d \le 2$. Since the selection maximized the distinct-root count over all maximizers, this excludes every maximizing vector with three or more distinct values, including a flat maximizing stratum.

Label the two values a, b, with minority multiplicity $r \in \{1, 2\}$ and majority multiplicity $s = 5 - r$. The one-value case is included by $a = b$. If $b = 0$, at least three coordinates vanish and the objective is zero. Otherwise reversing every sign allows $b > 0$. Put

$$v = a/b, \quad \tau = t/b, \quad L = rv + s,\ R = rv^2 + s,\ P = v^r,\ A = rv^{r-1} + sv^r.$$

(Here L, R are the rescaled quantities.) The normalized objective squared is

$$\frac{(\tau A - LP/\tau)^2}{(\tau^2 + R)^5}, \qquad \tau > 0, \quad \tau^2 \geq L^2.$$

It suffices to prove, for every real v,

$$f_r(v, X) = X(X + R)^5 - 1728(AX - LP)^2 \geq 0 \quad (X \geq L^2). \tag{5.5}$$

Write $X = L^2 + y$, $V = L^2 + R > 0$, and $f_r = \sum_{j=0}^{6} c_j y^j$. Then

$$\begin{aligned}
c_0 &= L^2 J, & J &= V^5 - 1728(LA - P)^2, \\
c_1 &= V^4(6V - 5R) - 3456AL(LA - P), & c_2 &= 5V^3(3V - 2R) - 1728A^2, \\
c_3 &= 10V^2(2V - R), & c_4 &= 5V(3V - R), \\
c_5 &= 6V - R, & c_6 &= 1.
\end{aligned}$$

In particular $c_3, c_4, c_5, c_6 > 0$. For $r = 1$, putting $w = (v + 2)^2$ gives

$$J = 32w(w^4 + 30w^3 + 360w^2 + 1296w + 11664) \geq 0.$$

For $r = 2$,

$$J = 7776(v + 2)^2 Q(v + 1), \qquad Q(w) = 1 - 2w + 8w^2 - 14w^3 + 22w^4 - 14w^5 + 8w^6 - 2w^7 + w^8.$$

For $w < 0$, $Q(w) > 0$ term by term, and $Q(0) = 1$. For $w > 0$, $Q(w)/w^4 = z^4 - 2z^3 + 4z^2 - 8z + 8$, where $z = w + w^{-1} \geq 2$; substitution $z = 2 + h$ gives $8 + 16h + 16h^2 + 6h^3 + h^4 > 0$. Thus $c_0 \geq 0$ in both cases.

The remaining five univariate positivity certificates are given explicitly below. For $r = 1$ they prove $c_1, c_2 > 0$ on $\mathbb{R}$; hence every coefficient of f_1 is nonnegative. For $r = 2$ they prove $c_1, c_3, \Delta = 4c_1c_3 - c_2^2 > 0$ on $\mathbb{R}$. Therefore

$$f_2 = c_0 + y(c_1 + c_2 y + c_3 y^2) + c_4 y^4 + c_5 y^5 + y^6 \geq 0 \quad (y \geq 0).$$

This proves (5.5), then the radial bound and (5.1). Sharpness follows from $z = (2, 2, -1, -1, -1, -1)$: $\sum |z_i|^2 = 12$ and $e_5(z) = -12$.

We now determine equality in the complex inequality. The zero-coordinate cases above are strict, so all radii are positive. Equality in the radial majorant forces the selected real five-vector to maximize the appropriate signed coefficient functional on its own fixed-L, R slice: an improving point on that slice would preserve t, the normalization and $|L| \leq t$, contradicting the sharp radial bound. The all-maximizer argument therefore applies to this actual vector. Use its binary parameters above; $b = 0$ would make the objective zero, so $b > 0$ after simultaneous sign reversal. For $r = 1$, the strict positivity of c_1 gives $f_1(v, L^2 + y) > 0$ whenever $y > 0$. For $r = 2$, the strictly positive quadratic $c_1 + c_2 y + c_3 y^2$ gives the same conclusion. Equality thus requires $y = 0$. Since $X = t^2/b^2 = L^2 > 0$, the identity $c_0 = L^2 J = 0$ forces $J = 0$. The strict factors displayed above leave exactly $v = -2$. For $r = 1$ the five signed radii are $(-2b, b, b, b, b)$ and $t = 2b$; for $r = 2$ they are $(-2b, -2b, b, b, b)$ and $t = b$. In either case the six original radii are two copies of $2b$ and four copies of b.

Equality must also hold in the triangle inequality in the exact identity

$$|e_5(z)| = P_6 \left| \sum_i \overline{u_i}(r_i^{-1} - r_i/t^2) \right|.$$

If $t = 2b$, the two large-radius coefficients vanish and the other four equal $3/(4b) > 0$. Thus the four small phases agree, say with a unit λ. Centering gives $u_{L,1} + u_{L,2} = -2\lambda$, forcing both large phases to equal $-\lambda$. If $t = b$, the small-radius coefficients vanish and the two large coefficients equal $-3/(2b)$. They force the large phases to agree; centering makes the sum of the four unit small phases have modulus four in the opposite direction, so those phases also agree. This includes

all median ties: the selected endpoint radius is either b or $2b$, and signs at the tied radii contributed zero to the majorant. Hence equality gives precisely the stated complex orbit. The sharpness calculation above, complex homogeneity and permutation invariance prove the converse.

For the stable-law assertion, the existing ternary restriction supplies, separately for every complex direction w, real symmetric matrices A, C such that $R_5(w) = e_5(A + iC)$, $\mathrm{Tr}(A + iC) = 0$ and $\|A + iC\|_F^2 = 2w^*Bw$. The six eigenvalues are centered and their squared-modulus sum is at most this Frobenius square by Schur triangularization. Consequently

$$|R_5(w)|^2 \le (2w^*Bw)^5/1728.$$

For standard circular Gaussian Z, substitute $w = Q^{-1/2}Z$ and use $D = \frac{1}{3}Q^{-1/2}BQ^{-1/2}$. Gaussian coefficient orthogonality and the exact density normalization give

$$\|U_5\|_2^2 \le \frac{6^5\, 5!\, h_5(D)}{1728\, 5!\binom{6}{5}} = \tfrac{3}{4}h_5(D).$$

Matrices may depend on w; no global determinantal representation in arbitrarily many variables is used. Singular deficits cause no division. □

Exact certificates on the full real line. The following table specifies complete Bernstein partitions for the five polynomials used above. The variables in the partition column are denoted w and belong to $[0, 1]$.

r	polynomial	degree	partition endpoints
1	c_1	10	$0, \frac{1}{2}, \frac{5}{8}, \frac{11}{16}, \frac{3}{4}, 1$
1	c_2	8	$0, \frac{1}{2}, \frac{3}{4}, 1$
2	c_1	10	$0, \frac{1}{2}, \frac{5}{8}, \frac{11}{16}, \frac{3}{4}, 1$
2	c_3	6	$0, \frac{1}{2}, \frac{3}{4}, 1$
2	$4c_1c_3 - c_2^2$	16	$0, \frac{1}{2}, \frac{5}{8}, \frac{3}{4}, 1$

Here is the exact verification method, implemented in the accompanying standard-library program `verify_complex_six.py`. Reconstruct each $p(v) = \sum_{i=0}^{d} a_i v^i$ from the displayed formulas. All a_i are nonnegative and $a_0 > 0$, proving positivity for $v \ge 0$. For $v < 0$ set $v = -w/(1 - w)$; the Bernstein controls of $(1 - w)^d p(-w/(1 - w))$ are

$$\beta_i = \frac{(-1)^i a_i}{\binom{d}{i}}, \qquad 0 \le i \le d.$$

Apply exact rational midpoint de Casteljau subdivision, retaining both children, to the listed intervals. Every entry of every resulting control array is strictly positive. These are twenty full arrays, and their intervals cover $[0, 1]$. Since the Bernstein basis is nonnegative and sums to one, this certifies positivity for every negative real v, without a cutoff or sampling argument. The program also verifies the two exact factorizations of J. Thus the computational part consists only of finite integer and rational identities and strict sign checks.

Attainment by positive stable count laws. For $0 < \epsilon < 1/4$ consider

$$p_\epsilon(X, Y) = [(\tfrac{1}{2} + 2\epsilon)X + (\tfrac{1}{2} - 2\epsilon)Y]^2[(\tfrac{1}{2} - \epsilon)X + (\tfrac{1}{2} + \epsilon)Y]^4.$$

This is a normalized stable product law with positive coefficients, means $(3, 3)$ and $q = (1/2, 1/2)$. Direct calculation gives

$$B = 6\epsilon^2 \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \quad D = 4\epsilon^2 \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \quad R_5 = -12\epsilon^5(X - Y)^5.$$

The nonzero eigenvalue of D is $8\epsilon^2$, so

$$h_5(D) = 32768\epsilon^{10}, \qquad \|U_5\|_2^2 = \frac{144\epsilon^{10}}{6}(2 + 2)^5 = 24576\epsilon^{10} = \tfrac{3}{4}h_5(D).$$

Thus the constant is attained within the stated count-law class.

Corollary 5.2. *For every degree $n \geq 6$, the fifth-order spectral bound has the additional branch*

$$\|U_5\|_2^2 \leq b_n h_5(D), \qquad b_n = \frac{3125}{8}\frac{\binom{n}{5}}{(n-1)^5}.$$

This branch is sharp at $n = 6$.

Proof: Use the classical complex derivative-root contraction, in the credited derivative argument of Han–Niles-Weed (2026b, Section 1.1). For centered $z \in \mathbb{C}^n$ with $\sum |z_i|^2 = n$, differentiate its monic root polynomial $n-6$ times and normalize to obtain six roots w. Coefficient differentiation and the iterated contraction give

$$e_5(w) = \frac{6}{\binom{n}{5}} e_5(z), \qquad \sum_{j=1}^{6} |w_j|^2 \leq \frac{30}{n-1}.$$

Theorem 5.1 yields

$$|e_5(z)|^2 \leq \left(\frac{\binom{n}{5}}{6}\right)^2 \frac{(30/(n-1))^5}{1728} = b_n \binom{n}{5}.$$

Apply the same pointwise ternary/Schur and Gaussian transfer as above, now with

$$D = (2/n) Q^{-1/2} B Q^{-1/2},$$

to obtain the spectral assertion. □

This improves the previous degree-five derivative branch by the factor 3125/8192. The higher-degree bounds satisfy $b_7 = 21875/20736 > 1$ and $\lim_{n\to\infty} b_n = 625/192 > 1$. The branch can be combined by minimum with the other established bounds.

5.1. *An exact rank-two matrix formulation.* The degree-six, rank-two sharp bound has an exact finite-dimensional matrix formulation. The converse below realizes every admissible matrix pair by an actual positive stable count law, so this reformulation does not enlarge the problem to arbitrary symmetric tensors.

Theorem 5.3 (Exact rank-two matrix reformulation)**.** *The following assertions are equivalent.*

(1) *Every admissible stable count law of degree 6 with* $\operatorname{rank} D = 2$ *satisfies* $\|U_5\|_2^2 \leq (3/4) h_5(D)$.

(2) *For a 6×6 matrix M, write $e_5(M) = [t]\det(tI_6 + M)$, the ordinary fifth elementary function of its eigenvalues. For every pair of real symmetric 6×6 matrices H, K with*

$$\operatorname{Tr} H = \operatorname{Tr} K = \operatorname{Tr}(HK) = 0, \qquad \operatorname{Tr} H^2 = \operatorname{Tr} K^2 = 1,$$

write

$$e_5(zH + wK) = \sum_{j=0}^{5} c_j z^{5-j} w^j.$$

Then, for every $a, b \geq 0$,

$$S_{H,K}(a,b) := \sum_{j=0}^{5} c_j^2 (5-j)!\, j!\, a^{5-j} b^j \;\leq\; \frac{5}{72} h_5(a,b), \qquad h_5(a,b) = \sum_{j=0}^{5} a^{5-j} b^j. \tag{5.6}$$

It is enough to establish the second assertion for $a, b > 0$. Both assertions hold by Theorem 5.1. The equivalence proof also realizes every admissible matrix pair by a positive stable count law.

Proof: First suppose the matrix assertion holds. Each centered coefficient R_j is constant along $\ker B$. To see this directly, fix $v \in \ker B$. The centered roots associated with $x + sv$ have squared norm $2x^\top Bx$, independent of real s, by (2.2). Their elementary coefficients are therefore uniformly bounded in s. The polynomial $R_j(x + sv)$ is consequently constant. This argument only requires a finite root-norm bound, not the desired spectral inequality.

Let $a, b > 0$ be the two nonzero eigenvalues of D, with orthonormal eigenvectors u_1, u_2, and put $Q = \operatorname{diag}(q)$. Since $D\sqrt{q} = 0$, the vectors $Q^{-1/2}u_i$ are centered for L. The kernel descent just proved shows that all whitened R_j depend only on these two coordinates. The full ternary restriction

$$F(t, z, w) = p\big(t\mathbf{1} + Q^{-1/2}(zu_1 + wu_2)\big)$$

is monic of degree six in t, hyperbolic in $(1, 0, 0)$, and has zero t^5 coefficient and quadratic coefficient $-3(az^2 + bw^2)$. The classical ternary representation used in Theorem 4.4 gives

$$F(t, z, w) = \det(tI + zA + wC)$$

with real symmetric matrices satisfying

$$\operatorname{Tr} A = \operatorname{Tr} C = \operatorname{Tr}(AC) = 0, \qquad \operatorname{Tr} A^2 = 6a, \quad \operatorname{Tr} C^2 = 6b.$$

Thus $H = A/\sqrt{6a}$ and $K = C/\sqrt{6b}$ satisfy the matrix hypotheses.

For independent standard circular complex Gaussians Z, W, monomial orthogonality and the exact density identity give

$$\|U_5\|_2^2 = \frac{\mathbb{E}|e_5(\sqrt{6a}ZH + \sqrt{6b}WK)|^2}{5!\binom{6}{5}} = \frac{6^5}{720}S_{H,K}(a, b) = \frac{54}{5}S_{H,K}(a, b).$$

Here other Gaussian coordinates disappear by kernel descent, and $h_5(D) = h_5(a, b)$. Equation (5.6) proves the first assertion.

Conversely, fix any matrices in the second assertion and $a, b > 0$. Set

$$\begin{aligned} H_1 &= \sqrt{a}\, H, \\ H_2 &= -\tfrac{1}{2}\sqrt{a}\, H + \tfrac{1}{2}\sqrt{3b}\, K, \\ H_3 &= -\tfrac{1}{2}\sqrt{a}\, H - \tfrac{1}{2}\sqrt{3b}\, K. \end{aligned}$$

For sufficiently small $\epsilon > 0$, all three matrices $A_i = I/3 + \epsilon H_i$ are positive definite and sum to I. Define

$$p(X_1, X_2, X_3) = \det\left(\sum_{i=1}^{3} X_i A_i\right).$$

This is homogeneous of degree six and normalized at $\mathbf{1}$. Its stability follows because $\sum_i (\Im X_i)A_i$ is positive definite when every $\Im X_i > 0$. Its coefficients are nonnegative by Cauchy–Binet applied to concatenated square-root columns of the A_i. They are in fact strictly positive: write $A_i = \delta I + B_i$ with a common $\delta > 0$ and $B_i \succeq 0$; the same expansion contains the positive contribution $\delta^6(X_1 + X_2 + X_3)^6$.

The means are $t_i = \operatorname{Tr} A_i = 2$, so $q_i = 1/3$. Direct second-coefficient comparison gives

$$B = \frac{\epsilon^2}{2}\operatorname{Gram}(H_1, H_2, H_3), \qquad D = B.$$

In the orthonormal H, K coordinates, the three frame vectors are $(\sqrt{a}, 0)$, $(-\sqrt{a}/2, \sqrt{3b}/2)$ and $(-\sqrt{a}/2, -\sqrt{3b}/2)$. Their frame operator is $\operatorname{diag}(3a/2, 3b/2)$. Hence D has exactly two positive eigenvalues $(3\epsilon^2/4)a, (3\epsilon^2/4)b$.

Whitening multiplies the category coordinates by $\sqrt{3}$. An orthogonal change among independent circular Gaussian coordinates therefore makes the whitened fifth-coefficient matrix equal in law to

$$\epsilon\sqrt{9/2}\,(\sqrt{a}\, ZH + \sqrt{b}\, WK).$$

Consequently

$$\|U_5\|_2^2 = \frac{\epsilon^{10}(9/2)^5}{720} S_{H,K}(a,b), \qquad h_5(D) = \epsilon^{10}(3/4)^5 h_5(a,b).$$

Their ratio is $(54/5)S_{H,K}(a,b)/h_5(a,b)$. The first assertion applied to this genuine rank-two stable law proves (5.6) for positive a, b. Polynomial continuity proves the nonnegative endpoints as well; both sides vanish at $a = b = 0$. □

The realization retains the simultaneous matrix constraints. The coupled weighted inequality (5.6) is therefore a consequence of the full stable-law theorem, with the stronger coefficient $5/72$.

Corollary 5.4 (Two coupled pencil subclasses)**.** *Let H, K satisfy the matrix hypotheses of Theorem 5.3. Suppose either*

(1) $\det(zH + wK)$ *is identically zero as a polynomial in z, w; or*
(2) *there is a common orthogonal decomposition $\mathbb{R}^6 = V \oplus W$, with $\dim V = 4$, $\dim W = 2$, preserved by both matrices, and every restricted block $H|_V, K|_V, H|_W, K|_W$ has trace zero.*

Then, for all $a, b \geq 0$,

$$S_{H,K}(a,b) \leq \frac{120}{3125} h_5(a,b).$$

In particular, a degree-six rank-two stable law whose centered ternary pencil has either property satisfies

$$\|U_5\|_2^2 \leq \frac{1296}{3125} h_5(D) < h_5(D) \qquad (D \neq 0).$$

Neither condition requires H and K to commute. The singular-pencil condition concerns the matrix polynomial, not a reduction of the covariance rank below two. The same count-law bound also holds, without a restriction on covariance rank, for $p = p_4 p_2$ where the normalized homogeneous stable count-law factors have degrees four and two and the same mean proportions q.

Proof: Let Z, W be independent standard circular complex Gaussians and set $M = \sqrt{a}\, ZH + \sqrt{b}\, WK$. The global orthonormality conditions give

$$\|M\|_F^2 = a|Z|^2 + b|W|^2.$$

In the first case, M has a zero eigenvalue. If it has at least two zero eigenvalues counted algebraically, $e_5(M) = 0$. Otherwise $e_5(M)$ is the product of its other five eigenvalues. Schur triangularization and AM–GM therefore give

$$|e_5(M)|^2 \leq \left(\frac{\|M\|_F^2}{5} \right)^5.$$

This argument does not require normality or a common kernel for the pencil.

In the second case write $M = A \oplus B$ with trace-free complex 4×4 and 2×2 blocks. Characteristic-coefficient multiplication gives $e_5(M) = e_3(A) \det B$. For the four eigenvalues of A, their zero sum permits the Hadamard parameterization

$$\lambda(A) = \tfrac{1}{2}(u + v + w,\ u - v - w,\ -u + v - w,\ -u - v + w).$$

It satisfies $e_3(A) = uvw$ and $\sum_j |\lambda_j(A)|^2 = |u|^2 + |v|^2 + |w|^2$. Thus, with $X = \|A\|_F^2$ and $Y = \|B\|_F^2$, Schur triangularization and AM–GM imply

$$|e_3(A)|^2 \leq \frac{X^3}{27}, \qquad |\det B|^2 \leq \frac{Y^2}{4}.$$

Consequently

$$|e_5(M)|^2 \leq \frac{X^3 Y^2}{108} \leq \frac{(X+Y)^5}{3125} = \frac{\|M\|_F^{10}}{3125}.$$

The middle inequality is weighted AM–GM, with the maximum $X : Y = 3 : 2$ for fixed $X + Y$.

In either case complex monomial orthogonality and exponential moments give

$$S_{H,K}(a,b) = \mathbb{E}|e_5(M)|^2 \le \frac{\mathbb{E}(a|Z|^2 + b|W|^2)^5}{3125} = \frac{120}{3125} h_5(a,b).$$

The exact factor $54/5$ in Theorem 5.3 gives the count-law conclusion. Positive $h_5(D)$ makes the final inequality strict whenever $D \neq 0$.

For the last assertion, let B_4, B_2 be the two factors' covariance deficits. Their common $L = q \cdot x$ makes the centered product coefficient exactly $R_5 = R_3^{(4)} R_2^{(2)}$, and covariance additivity gives $B = B_4 + B_2$. On each complex direction w, apply the justified ternary representation separately to each factor. The size-four estimate above and the quadratic top bound give

$$|R_3^{(4)}(w)|^2 \le \frac{8}{27}(w^* B_4 w)^3, \qquad |R_2^{(2)}(w)|^2 \le (w^* B_2 w)^2.$$

With $X = w^* B_4 w$ and $Y = w^* B_2 w$, weighted AM–GM yields

$$|R_5(w)|^2 \le \frac{32}{3125}(w^* B w)^5.$$

For total degree six, $D = \frac{1}{3} Q^{-1/2} B Q^{-1/2}$. Substituting $w = Q^{-1/2} Z$ and using the exact fifth-density identity therefore gives

$$\|U_5\|_2^2 \le \frac{32 \cdot 3^5}{3125 \cdot 6} h_5(D) = \frac{1296}{3125} h_5(D).$$

The separate ternary matrices may depend on w; no global determinantal model in an arbitrary number of variables is used. □

The second condition retains a genuine common block constraint without imposing simultaneous diagonalization. The first includes the case $R_6 \equiv 0$ in the centered expansion. These are sufficient subclasses. The proof uses classical characteristic identities, Schur triangularization and AM–GM. The product extension requires matching mean proportions of its two factors; arbitrary products with different means are not covered.

6. A nonfactorizable cubic obstruction to a Frobenius-power extension

The preceding spectral-square bound cannot generally be replaced by the stronger Frobenius power available for permutation-product models. We give actual stable count laws, rather than only a generic tensor obstruction.

Theorem 6.1. *For every integer $r \ge 5$, there are degree-three real-stable count laws P_r, P_r^0 on $2r$ categories with strictly positive coefficients, identical means and covariance, and nonzero D eigenvalue $2\delta^2/(3r)$ repeated r times, where any fixed $0 < \delta \le 1/4$ is allowed. With $c = 1/4$,*

$$x = 2\delta^2/3, \quad s_2 = 4\delta^4/(9r), \quad \|U_3(P_r)\|_2^2 \ge \frac{c^2\delta^6}{1000r} = \frac{27c^2}{8000} s_2 x.$$

The law P_r^0 has $U_3 = 0$. Consequently no dimension-uniform finite K gives $\|U_3\|_2^2 \le K s_2^{3/2}$ on the full stable class, even at fixed degree three and fixed $x > 0$. Both laws are nonfactorizable into nonnegative categorical linear forms.

Norm-one real symmetric cubic tensors with squared Frobenius norm of order r^2 are established tensor-norm phenomena; see Cao, He, Li and Wang (2023, Theorem 4.6 and Corollary 4.9). We include a self-contained probabilistic existence proof with sufficient constants. This is an analytic existence argument, not a numerical search or an algorithm outputting a tensor in every dimension.

Proof: Let $N = \dim \mathrm{Sym}^3(\mathbb{R}^r) = \binom{r+2}{3}$ and choose an isotropic Gaussian G in that Euclidean space, with the ordered Frobenius metric. Then $\|G\|_{\mathrm{F}}^2$ has mean N, variance $2N$, and

$$\Pr(\|G\|_{\mathrm{F}}^2 < N/2) \le 8/N < 1/4 \quad (r \ge 5).$$

For a fixed unit u, $G(u,u,u)$ is a standard normal because $\|u^{\otimes 3}\|_{\mathrm{F}} = 1$. A sphere $1/6$-net has at most 13^r points, by disjoint radius-$1/12$ ball volume comparison. Banach's theorem identifies the diagonal supremum M with the full trilinear norm; telescoping gives $|G(u,u,u) - G(v,v,v)| \le 3M\|u-v\|$, hence $M \le 2\max_{\mathrm{net}} |G(v,v,v)|$. The Gaussian tail and union bound give

$$\Pr(M > \sqrt{32r}) \le 2\,13^r e^{-4r} < 2e^{-r} < 1/4.$$

Here $\log 13 < 3$ and $e^5 > 8$, both certified by lower Taylor sums. A realization outside both failures has $M > 0$ and, upon normalization $A = G/M$,

$$\sup_{\|u\|=1} |A(u,u,u)| = 1, \qquad \|A\|_{\mathrm{F}}^2 \ge N/(64r) \ge r^2/384 \ge r^2/1000. \tag{6.1}$$

Put $C(u) = A(u,u,u)$ and $F_c(z,u) = z^3 - z\|u\|^2 + cC(u)$, with $c = 1/4$. The depressed-cubic discriminant is

$$4\|u\|^6 - 27c^2C(u)^2 \ge (4 - 27/16)\|u\|^6 \ge 0.$$

Thus F_c is homogeneous and hyperbolic in $e = (1,0)$ with $F_c(e) = 1$, also at $u = 0$; the same holds for $c = 0$. Let $m = 2r$ and index u_i by $+e_1, -e_1, \dots, +e_r, -e_r$. Then $\sum_i u_i = 0$ and $\sum_i u_i u_i^\top = 2I$. Set $v_i = (1, \delta u_i)$. Every real root of $F_c(z, \delta u_i)$ has modulus less than 2δ: for $|z| \ge 2\delta$, $|z^3 - \delta^2 z| \ge 6\delta^3 > c\delta^3|C(u_i)|$. Therefore the characteristic roots of v_i relative to e are positive, being one minus these roots. The classical Gårding cone theorem gives an open convex cone of such directions and hyperbolicity in every direction in it Bauschke et al. (2001, Fact 2.7).

Define

$$p_c(X) = F_c\left(\frac{\sum_i X_i}{m}, \frac{\delta}{m}\sum_i X_i u_i\right) = m^{-3} F_c(\sum_i X_i v_i).$$

It is normalized at $\mathbf{1}$. If all $\Im X_i > 0$, the imaginary vector $b = \sum_i (\Im X_i) v_i$ belongs to that cone. With $a = \sum_i (\Re X_i) v_i$, the polynomial $t \mapsto F_c(a + tb)$ has only real roots and nonzero leading coefficient. Hence $F_c(a + ib) \ne 0$: p_c is real stable. This deduction is the standard hyperbolicity-to-tube argument, not a coordinatewise discriminant shortcut.

Coefficient positivity is separate and explicit. The symmetric polarization at v_i, v_j, v_k equals

$$1 - \frac{\delta^2}{3}(u_i \cdot u_j + u_i \cdot u_k + u_j \cdot u_k) + c\delta^3 A(u_i, u_j, u_k) \ge 1 - \delta^2 - c\delta^3 > 0.$$

Banach bounds the last absolute value by one even with repeated indices. Each count monomial coefficient is a positive multinomial multiplicity times this expression divided by m^3. Thus p_c and p_0 are genuine strictly positive count laws.

Write $L = \sum X_i/m$ and $U = \sum X_i u_i$. Exactly,

$$p_c = L^3 - \frac{\delta^2}{m^2}\|U\|^2 L + \frac{c\delta^3}{m^3} C(U).$$

Since $U(\mathbf{1}) = 0$, the last term has zero first and second derivatives there. Both laws have $t_i = 3/m$ and $B = (\delta^2/m^2)\,\mathrm{Gram}(u_i)$, independently of c; the displayed terms are their centered R_2, R_3. The Gram spectrum is two repeated r times and zero otherwise. Thus $D = (2m/3)B$ has the asserted spectrum.

Since $q_i = 1/m$ and $\binom{3}{3} = 1$, the exact ordered cubic density identity and $UU^\top = 2I$ give

$$\|U_3\|_2^2 = \frac{c^2\delta^6}{m^3} 2^3 \|A\|_{\mathrm{F}}^2 = \frac{c^2\delta^6}{r^3}\|A\|_{\mathrm{F}}^2.$$

Equation (6.1) proves the lower bound. Since $s_2^{3/2} = 8\delta^6/(27r^{3/2})$, its ratio is at least $27c^2\sqrt{r}/8000$ and diverges.

Finally a product of three normalized categorical forms with vectors w_a and mean row q has $B = \frac{1}{2}\sum_{a=1}^{3}(w_a - q)(w_a - q)^\top$ and rank at most two, because the deviations sum to zero. Our B has rank $r \geq 5$ for both laws. Any nonnegative linear factors can be normalized, so neither polynomial admits such a factorization. □

The relative quadratic chi-square increment here is at least $3c^2\delta^2/1000$, since $\chi^2(P_r^0\|M) = a_3 s_2$ and the cubic energy adds orthogonally. It need not vanish as $s_2 \to 0$ at fixed trace; absolute chi-square still tends to zero by Corollary 4.2. This does not infer a TV lower bound from L^2, nor contradict the product-subclass results of Han–Niles-Weed. Generic real tensors cannot prove that target: $\Re(X + iY)^4 = X^4 - 6X^2Y^2 + Y^4$ has diagonal norm one and ordered squared tensor norm eight, whereas $h_4(1, 1) = 5$. This last tensor is an abstract symmetric-tensor example.

7. Sharp real coefficients and equality

The complete domain partition is the guide to the following proof: degrees 2–50 use Lemma 7.4; degrees 51–599 use Lemma 7.15. For $n \geq 600$, the bottom half uses Corollary 7.13. In the top half, codimension zero uses AM–GM, codimensions 1–19 use Theorem 7.16, 19–40 use Theorem 7.17, 41–79 use Lemma 7.14, and the remaining central range uses Theorem 7.10. The shared boundary at 19 and the 50/51, 599/600 degree boundaries are included. Equality is handled after the strictness gates.

7.1. *Coefficient bounds and finite-degree certificates.* The coefficient extremum reduction used next is classical: see Riener and Schabert (2023, Theorem 2.8 and Corollary 2.10), as well as the real-vector reduction in Han and Niles-Weed (2026a). We include the needed special proof.

Lemma 7.1. *For fixed $n \geq 3$ and $3 \leq k \leq n$, the maximum of $|e_k(\lambda)|$ on $\sum_i \lambda_i = 0$, $\sum_i \lambda_i^2 = 2$ is attained at a vector with two distinct values.*

Proof: The slice is nonempty and compact. Maximize either sign of e_k, and among maximizers select one with the largest number d of distinct roots of $f(z) = \prod_i (z - \lambda_i)$. If $3 \leq d < n$, write $f = hg$, where g contains each distinct linear factor once and h retains multiplicity minus one. For either sufficiently small sign of ε, $g + \varepsilon$ has d simple real roots near the old roots. Thus $f + \varepsilon h = h(g + \varepsilon)$ is real-rooted and preserves the first two coefficients, since $\deg h = n - d \leq n - 3$. If the target coefficient of h is nonzero, one sign improves the objective. Otherwise the objective is constant. Every repeated old root remains through h and also supplies a new nearby root through $g + \varepsilon$; no new root equals an old one. Taking disjoint neighborhoods of distinct old roots shows the distinct-root count strictly increases, a contradiction. If $d = n$, perturbing by εz^{n-k} preserves the first two coefficients and real-rootedness for both small signs, but changes the objective. Hence $d \leq 2$. The positive square sum rules out $d = 1$. Apply this argument to both signs to obtain the absolute-value assertion. □

7.1.1. *A uniform real coefficient refinement.* The following mass-budget estimate is an elementary rearrangement inequality for the arcsine weight. The contour, binary reduction and enlarged Beta integral that follow are the ingredients of Han and Niles-Weed (2026a, Appendix B.1); the change is to use this estimate instead of their endpoint/interior split and to retain the normalization factors.

Lemma 7.2. *If f is measurable and $0 \leq f \leq 1$ on $[-1, 1]$ and $A = \int_{-1}^{1} f(t)\,dt$, then*

$$\left(\frac{1}{\pi}\int_{-1}^{1} \frac{f(t)}{\sqrt{1 - t^2}}\,dt\right)^2 \leq \frac{A}{2}.$$

Proof: The endpoints $A = 0, 2$ are immediate. Otherwise put $b = 1 - A/2$ and $E = [-1, -b] \cup [b, 1]$. Its length is A. For $w(t) = (1 - t^2)^{-1/2}$ one has $(\mathbf{1}_E - f)(w - w(b)) \geq 0$ on both E and its complement. Since $\int (\mathbf{1}_E - f) = 0$, integration gives $\int fw \leq \int_E w = 2 \arccos b$. For $0 \leq u \leq \pi/2$, the chord bound for sine on $[0, \pi/4]$ gives $\sin(u/2) \geq \sqrt{2}u/\pi$, so $1 - \cos u \geq 4u^2/\pi^2$. Apply this with $u = \arccos(1 - A/2)$ to obtain the assertion. □

Theorem 7.3. *For every centered real vector z with $\sum_i z_i^2 = n$,*

$$|e_k(z)|^2 \leq C_0 \binom{n}{k} \quad (0 \leq k \leq n), \qquad C_0 = \frac{\pi e^{1/4}}{\sqrt{2}} < \frac{20}{7}. \tag{7.1}$$

Proof: The normalized domain is empty at $n = 1$. The orders zero, one and two follow from $e_0 = 1, e_1 = 0, e_2 = -n/2$; order n also follows from AM–GM on $\prod z_i^2$. By Lemma 7.1 and scaling, the other orders reduce to a vector with r entries $\sqrt{s/r}$ and s entries $-\sqrt{r/s}$, where $r + s = n$. For $1 \leq k \leq n - 1$ put

$$\rho = \sqrt{\frac{k}{n-k}}, \qquad v_0 = \frac{k(n-k)}{n}, \qquad H_0 = \frac{(1+\rho^2)^{n/2}}{\rho^k}.$$

Define, for $-1 \leq t \leq 1$,

$$f(t) = \frac{(1 + \rho^2 s/r + 2\rho\sqrt{s/r}\, t)^{r/2} (1 + \rho^2 r/s - 2\rho\sqrt{r/s}\, t)^{s/2}}{(1+\rho^2)^{n/2}}.$$

The affine brackets are nonnegative. Their weighted average with weights $r/n, s/n$ is $1 + \rho^2$, so weighted AM–GM gives $0 \leq f \leq 1$. Cauchy's coefficient formula, followed by the triangle inequality and $t = \cos\theta$, yields

$$|e_k(z)| \leq \frac{H_0}{\pi} \int_{-1}^{1} \frac{f(t)}{\sqrt{1-t^2}}\, dt.$$

Let t_- and t_+ be the zeros of the first and second brackets. They satisfy $t_- \leq -1$ and $t_+ \geq 1$ by AM–GM. The affine substitution $t = t_- + (t_+ - t_-)u$, extending the nonnegative unweighted integrand to this larger interval, gives

$$A = \int_{-1}^{1} f(t)\, dt \leq \frac{1+\rho^2}{2\rho} \frac{n^{n/2+1}}{r^{(r+1)/2} s^{(s+1)/2}} \frac{\Gamma(r/2+1)\Gamma(s/2+1)}{\Gamma(n/2+2)}.$$

Here the Beta integral is $\int_0^1 u^{r/2}(1-u)^{s/2}\, du$. To bound it, use the standard Stirling inequalities, for $x > 0$,

$$\sqrt{2\pi x}(x/e)^x \leq \Gamma(x+1) \leq \sqrt{2\pi x}(x/e)^x e^{1/(12x)}.$$

Apply upper bounds at $r/2, s/2$ and the lower bound at $n/2$, keeping $\Gamma(n/2+2) = (n/2+1)\Gamma(n/2+1)$. Cancellation, and $(1+\rho^2)/\rho = \sqrt{n/v_0}$, give

$$A \leq \frac{\sqrt{\pi}}{\sqrt{v_0}} \frac{n}{n+2} e^{(1/r+1/s)/6}.$$

Independently, factorial Stirling bounds at $k, n-k, n$ give

$$\frac{H_0^2}{\binom{n}{k}} \leq \sqrt{2\pi v_0}\, e^{(1/k+1/(n-k))/12}.$$

Combining both estimates with Lemma 7.2,

$$\frac{|e_k(z)|^2}{\binom{n}{k}} \leq \frac{\pi}{\sqrt{2}} \frac{n}{n+2} \exp\left(\frac{1/r+1/s}{6} + \frac{1/k+1/(n-k)}{12}\right).$$

Each reciprocal sum is at most $n/(n-1)$, since its positive integer arguments sum to n. The exponent is therefore at most $n/[4(n-1)]$. Since

$$\log(1+2/n) \ge \frac{2}{n+2} \ge \frac{1}{4(n-1)} \qquad (n \ge 2),$$

the last bound is at most C_0.

For completeness we certify the rational ceiling in (7.1). The elementary identity

$$\frac{22}{7} - \pi = \int_0^1 \frac{x^4(1-x)^4}{1+x^2}\,dx > 0$$

follows by division: its integrand equals $x^6 - 4x^5 + 5x^4 - 4x^2 + 4 - 4/(1+x^2)$. Also $\sqrt{2} > 140/99$ since $19600 < 19602$. The exponential series has consecutive-term ratio at most $1/20$ after its order-four term, hence

$$e^{1/4} \le \sum_{j=0}^{3} \frac{4^{-j}}{j!} + \frac{4^{-4}}{4!(1-1/20)} < \frac{257}{200}.$$

Consequently

$$C_0 < \frac{22}{7}\frac{257}{200}\frac{99}{140} = \frac{559746}{196000} < \frac{20}{7}.$$

These rational comparisons are also checked by the accompanying code. □

Scaling the roots in (2.2) and applying Lemma 3.1 now gives the all-order estimate

$$\mathbb{E}_M U_k^2 \le C_0 (2W)^k \le \frac{20}{7}(2W)^k. \tag{7.2}$$

7.1.2. *Small-degree scalar certificates.*

Lemma 7.4. *For $2 \le n \le 50$, every centered real vector with squared norm n satisfies $|e_k|^2 \le \binom{n}{k}$ for every $0 \le k \le n$.*

Proof: Orders zero and one are immediate, and $e_2 = -n/2$ gives squared ratio $n/[2(n-1)] \le 1$. For $k \ge 3$, apply Lemma 7.1. The raw binary roots s repeated r times and $-r$ repeated s times, $r+s=n$, have squared norm nrs. Their coefficient is

$$F_{n,r,k} = \sum_a \binom{r}{a}\binom{s}{k-a} s^a (-r)^{k-a}.$$

The standalone program checks the integer inequality

$$F_{n,r,k}^2 \le \binom{n}{k}(rs)^k$$

for every $2 \le n \le 50$, every $1 \le r < n$ and every $2 \le k \le n$. All 40,425 comparisons pass. Each label and total count is checked; only exact integers decide the inequalities. Dividing the roots by $\sqrt{rs}$ and applying the same-n extremum lemma proves the full-real statement. This is a direct scalar certificate, not an inference from a sufficient finite-degree TV criterion. □

7.2. *Exceptional profiles and fixed low orders.* We normalize centered real root vectors by $\sum_i z_i^2 = n$. If

$$\frac{|e_k(z)|^2}{\binom{n}{k}} \le b_k, \tag{7.3}$$

then scaling (2.2) and applying Lemma 3.1 gives $\mathbb{E}_M U_k^2 \le 2^k b_k W^k$. All root extrema below use Lemma 7.1 one coefficient at a time; their maximizing profiles need not coincide.

7.2.1. *Discrete exceptional profiles.*

Lemma 7.5. *If a centered, variance-normalized two-valued vector has either root value with multiplicity one or two, then (7.3) holds with $b_k = 1$ at every order.*

Proof: Negation and permutation allow us to treat the smaller specified block. For multiplicity one the values are $\sqrt{n-1}$ once and $-1/\sqrt{n-1}$ repeated $n-1$ times. Expansion gives

$$e_k = (-1)^k(1-k)\binom{n}{k}(n-1)^{-k/2}, \qquad \frac{e_k^2}{\binom{n}{k}} = \frac{(k-1)^2\binom{n}{k}}{(n-1)^k}.$$

The ratios at $k = 0, 1$ are one and zero. At $k = 2$ the ratio is $n/[2(n-1)] \le 1$. For fixed $k \ge 3$, regard the last expression as a function of real $n \ge k$. Its logarithmic derivative is

$$-\frac{1}{n(n-1)} + \sum_{j=2}^{k-1} \frac{j-1}{(n-j)(n-1)} > 0,$$

because the $j = 2$ term exceeds the negative term. It therefore increases to $(k-1)^2/k! \le 2/3$. This last sequence starts at $2/3$ for $k = 3$ and decreases thereafter, since its next-term ratio is $k^2/[(k+1)(k-1)^2] < 1$.

For multiplicity two, assume $n \ge 4$; the case $n = 3$ is already multiplicity one after reflection. The values are $\sqrt{(n-2)/2}$ twice and $-\sqrt{2/(n-2)}$ repeated $n-2$ times. Direct expansion gives

$$\frac{e_k^2}{\binom{n}{k}} = \binom{n}{k}\left(\frac{2}{n-2}\right)^k F(n,k)^2, \qquad F(n,k) = \frac{(k-1)((k-4)n+4)}{4(n-1)}. \tag{7.4}$$

The orders zero and one are immediate. At order two the general centered identity gives ratio $n/[2(n-1)] \le 1$. For order three, two-valued extrema have $e_3 = nc/3$, with $c = (s-r)/\sqrt{rs}$, and $c^2 \le (n-2)^2/(n-1)$. Hence the ratio is at most $\frac{2}{3}n(n-2)/(n-1)^2 < 1$. For order four, the centered fourth-moment bound proved below gives ratio at most $3n(n-2)/[8(n-1)(n-3)] \le 1$ for $n \ge 4$; the cleared difference $5n^2 - 26n + 24$ is zero at four and increases thereafter.

For $5 \le k \le 11$, clearing positive denominators in (7.4) reduces the assertion to

$$16k!(n-1)^2(n-2)^k - 2^k(n)_k(k-1)^2((k-4)n+4)^2 \ge 0. \tag{7.5}$$

Put $n = k + x$. The companion exact program expands these seven integer polynomials. Every coefficient is strictly positive; their minimum coefficients, in order $k = 5, \dots, 11$, are

$$1408,\ 5120,\ 39168,\ 444416,\ 4986880,\ 55074816,\ 628633600.$$

The program constructs both products in (7.5) by integer linear-factor multiplication and checks every coefficient, so this proves the inequality for all real $x \ge 0$, not just a finite list of n.

For $n \ge k \ge 12$, the factor F is nonnegative and is at most $nk(k-1)/[4(n-1)] \le (3/11)k^2 < k^2/3$. Also $\binom{n}{k} \le n^k/k!$ and

$$\left(\frac{n}{n-2}\right)^k \le \left(\frac{k}{k-2}\right)^k < 9.$$

Indeed $t\log(t/(t-2))$ decreases for $t > 2$, and at twelve $6^{12} < 9 \cdot 5^{12}$. Thus the ratio in (7.4) is at most $2^k k^4/k! < 1$. This holds at twelve because $2^{12}12^4 < 12!$ and decreases thereafter, since its next-term ratio is $2(k+1)^3/k^4 < 1$. This completes the all-degree exceptional-profile proof. □

7.2.2. *Orders zero through four.*

Lemma 7.6. *For every centered real vector with squared norm n, $|e_k|^2 \le \binom{n}{k}$ for $0 \le k \le \min(4, n)$.*

Proof: Orders zero and one are immediate. At order two $e_2 = -n/2$, so its squared ratio is $n/[2(n-1)] \le 1$ for $n \ge 2$. At order three, Lemma 7.1 gives a binary extremizer with $c = (s-r)/\sqrt{rs}$ and $e_3 = nc/3$. Since $rs \ge n-1$, $c^2 \le (n-2)^2/(n-1)$, and the squared ratio is at most $2n(n-2)/[3(n-1)^2] < 1$.

For order four write $p_h = \sum_i z_i^h$. Newton's identity gives $e_4 = (n^2 - 2p_4)/8$. Cauchy on the other $n-1$ coordinates implies $z_i^2 \le n-1$, whence $n \le p_4 \le n(n-1)$ and $|e_4| \le n(n-2)/8$. Its squared ratio is at most

$$\frac{3n(n-2)}{8(n-1)(n-3)} \le 1 \qquad (n \ge 4),$$

since the cleared difference $5n^2 - 26n + 24$ equals $14y + 5y^2$ at $n = 4 + y$. □

Fixed orders through twelve on an infinite multiplicity domain. We next prove a scalar coefficient-one estimate uniformly in both integer root multiplicities. It supplies the scalar estimates needed at orders five through twelve. The finite polynomial computations below, including all exceptional rows, are generated by the accompanying standalone code.

Lemma 7.7. *For every real vector z with $\sum_i z_i = 0$, $\sum_i z_i^2 = n$, and $5 \le k \le \min(12, n)$,*

$$|e_k(z)|^2 \le \binom{n}{k}.$$

Consequently, for every law under consideration,

$$\mathbb{E}_M U_k^2 \le 2^k W^k, \qquad 7 \le k \le \min(12, n). \tag{7.6}$$

Proof: By Lemma 7.1, it suffices to treat two-valued profiles. Reflect and permute so their multiplicities satisfy $s \ge r$, $n = r + s$. Let $E_k(r, s)$ be the elementary coefficient of the integer root vector consisting of r copies of s and s copies of $-r$. Its squared norm is nrs, so the desired inequality is precisely

$$D_k(r, s) = (n)_k (rs)^k - k! E_k(r, s)^2 \ge 0. \tag{7.7}$$

For each $k = 5, \dots, 12$, put $a = \lceil k/2 \rceil$, $r = a + u$, $s = a + u + x$. Newton's identities reconstruct the exact polynomial E_k from

$$p_j = rs^j + s(-r)^j, \qquad E_0 = 1, \qquad E_h = \frac{1}{h} \sum_{j=1}^{h} (-1)^{j-1} E_{h-j} p_j.$$

The program expands (7.7) under this substitution and checks that every coefficient in u, x is nonnegative. The numbers of nonzero coefficients, in increasing order of k, are

$$121,\ 168,\ 225,\ 288,\ 361,\ 440,\ 529,\ 624.$$

Thus (7.7) holds for the whole real quadrant $u, x \ge 0$, covering every integer $s \ge r \ge \lceil k/2 \rceil$, with no upper cutoff on either multiplicity or on n.

The construction also has an independent exact polynomial-identity check. Since $p_1 = 0$, each Newton expansion term in E_k has at most $\lfloor k/2 \rfloor$ power-sum factors. Its total degree is at most $k + \lfloor k/2 \rfloor$ and its x-degree at most k. Therefore D_k has u-degree at most $3k$ and x-degree at most $2k$. On the complete grid $0 \le u \le 3k$, $0 \le x \le 2k$, the program independently compares the polynomial with the direct formula

$$E_k(r, s) = \sum_j \binom{r}{j} \binom{s}{k-j} s^j (-r)^{k-j}.$$

These 4068 exact comparisons determine the eight polynomials by interpolation. Coefficient nonnegativity, rather than those finite evaluations alone, proves the infinite-domain inequality.

It remains to cover $r < \lceil k/2 \rceil$. Multiplicities one and two are already covered by Lemma 7.5. The remaining pairs are exactly

k	7	8	9	10	11	12
r	3	3	3, 4	3, 4	3, 4, 5	3, 4, 5.

We give the cleared polynomial construction for these twelve cases. Write

$$E_k = (-r)^k \binom{n}{k} \frac{N_{r,k}(n)}{r^r (n)_r},$$

$$N_{r,k}(n) = \sum_{j=0}^{\min(r,k)} (-1)^j \binom{r}{j} n^j r^{r-j} (k)_j (n-j)_{r-j}.$$

To derive this formula, express the generating function as $(1-rT)^n[1+nT/(1-rT)]^r$ and compare coefficients; then use $\binom{n-j}{k-j}/\binom{n}{k} = (k)_j/(n)_j$. Dividing by $\sqrt{rs}$ normalizes the roots to squared norm n. Clearing positive denominators in the scalar inequality gives

$$k! r^{2r} (n)_r^2 (n-r)^k - r^k (n)_k N_{r,k}(n)^2 \geq 0. \tag{7.8}$$

For every listed pair the code substitutes $n = \max(k, 2r) + x$, expands both products, and checks that all coefficients are nonnegative. Since $x \geq 0$ is exactly the admissible range, this covers all remaining degrees. It independently checks the displayed coefficient factorization against the direct binomial formula on a degree-sized grid. The two-root reduction and Lemma 3.1 finish the proof. □

7.2.3. *Scalar constant one in a uniform central region.* The binary contour modulus has more structure than the area bound alone records: its logarithm is concave. Retaining its interval superlevel sets and splitting sparse from dense multiplicities yields an infinite central scalar-one region without extending a finite-order table.

Lemma 7.8. *Let $f : [-1, 1] \to [0, 1]$ be measurable with interval strict superlevel sets up to null sets. Put*

$$A = \int_{-1}^{1} f(t)\, dt, \qquad I = \frac{1}{\pi} \int_{-1}^{1} \frac{f(t)}{\sqrt{1-t^2}}\, dt.$$

If $0 < A \leq 3/10$, then $I \leq \Phi(A)$, where $\Phi(a) = \arccos(1-a)/\pi$, and

$$I^2 \leq \frac{2A}{\pi^2}(1 + A/10)^2. \tag{7.9}$$

The assertions hold trivially when $A = 0$.

Proof: The weight $w(t) = (1-t^2)^{-1/2}$ is convex, with $w''(t) = (1+2t^2)/(1-t^2)^{5/2} > 0$ in the open interval. Its integral over an interval of fixed length a is convex in the interval's left endpoint. Its maximum is thus attained at either endpoint interval; symmetry gives the same value, $\pi\Phi(a)$, at both. Integrable endpoint singularities follow by limits.

Let $\ell(h)$ be the length of $\{f > h\}$. Layer cake gives $A = \int_0^1 \ell(h)\, dh$ and $I \leq \int_0^1 \Phi(\ell(h))\, dh$. Although Φ is concave only on $[0, 1]$ and convex on $[1, 2]$, its tangent T_A at every $0 < A \leq 3/10$ majorizes it throughout $[0, 2]$. Concavity proves this on $[0, 1]$. On $[1, 2]$, tangent minus Φ is concave, so it suffices to check both endpoints. The endpoint one is already covered. Also

$$T_A(2) = \Phi(A) + (2-A)\Phi'(A), \qquad \frac{d}{dA} T_A(2) = (2-A)\Phi''(A) < 0.$$

At $A = 3/10$ its value is $(\arccos(7/10) + \sqrt{17/3})/\pi > 1$: $\cos(77/100) \geq 1 - (77/100)^2/2 > 7/10$, $\sqrt{17/3} > 119/50$, and their lower bounds sum to $63/20 > 22/7 > \pi$. Thus $T_A(2) > \Phi(2)$.

Integrating the global tangent and using the mean of ℓ proves $I \le \Phi(A)$ without assuming global concavity.

For $0 \le v \le 3/20$,

$$(1+3v/5)^2(1-v)-1 = v(1/5-21v/25-9v^2/25) \ge 0,$$

since the decreasing bracket equals $659/10000 > 0$ at $3/20$. Squaring positive factors and integrating gives $\arcsin x \le x + x^3/5$ for $x^2 \le 3/20$. Substituting $\arccos(1-A) = 2\arcsin\sqrt{A/2}$ proves (7.9). □

Lemma 7.9. *For a two-valued normalized centered profile with multiplicities r, s, scalar one holds whenever $k \ge 8\min(r,s)$.*

Proof: Reflect so $r \le s$, and put $a = \sqrt{s/r}$, $b = -\sqrt{r/s}$, $n = r+s$. Expanding the r exceptional factors relative to $(1+bT)^n$ gives

$$e_k = b^k \binom{n}{k} F_r(n,k), \qquad F_r(n,k) = \sum_{j=0}^{r} \binom{r}{j} (-n/r)^j \frac{(k)_j}{(n)_j},$$

where terms with $j > k$ vanish. Because $k \le n$, each falling-factorial ratio is at most $(k/n)^j$, so $|F_r| \le (1+k/r)^r$ and

$$\frac{e_k^2}{\binom{n}{k}} \le \frac{r^k}{k!}\left(\frac{n}{n-r}\right)^k (1+k/r)^{2r}.$$

For $n \ge k \ge 2r$, the middle factor is at most 4^r: it decreases in n, and the resulting function $t\log(t/(t-r))$ decreases for $t > r$, so its value at $t = k$ is at most that at $2r$. With $k! \ge (k/e)^k$, $e < 3$ and $\gamma = k/r \ge 8$, the last bound is at most

$$[4(3/\gamma)^\gamma(1+\gamma)^2]^r.$$

The bracket decreases for $\gamma \ge 8$, since its log derivative is $\log(3/\gamma) - 1 + 2/(1+\gamma) < 0$. At eight it equals $2125764/16777216 < 1$. This proves the assertion in every degree. □

Theorem 7.10. *Every centered real vector z with $\sum_i z_i^2 = n$ satisfies*

$$|e_k(z)|^2 \le \binom{n}{k} \qquad \text{if } \min(k, n-k) \ge 80.$$

More precisely, for binary profiles with $r, s \ge 11$ and $v_0 = k(n-k)/n \ge 40$, the squared normalized ratio is at most $3150873/3171875 < 1$.

Proof: For the joint binary assertion, use the modulus f, area A and factor H_0 from Theorem 7.3. The logarithm of f is a sum of positive multiples of logarithms of positive affine functions on $(-1,1)$, hence is concave. Therefore f has interval superlevel sets, including cases with endpoint zeros.

The exact Beta-area and factorial bounds already proved give

$$A \le \sqrt{\frac{\pi}{v_0}}\frac{n}{n+2} e^{E_a}, \qquad \frac{H_0^2}{\binom{n}{k}} \le \sqrt{2\pi v_0}\, e^{E_b},$$

where $E_a = (1/r + 1/s)/6 \le 1/33$ and $E_b = 1/(12v_0) \le 1/480$. Their sum is at most $57/1760 < 1/30$. The estimate $e^x \le (1-x)^{-1}$ gives $e^{E_a} \le 33/32$ and $e^{E_a+E_b} < 30/29$. Thus

$$A^2 < \frac{22}{7}\left(\frac{33}{32}\right)^2 \frac{1}{40} = \frac{11979}{143360} < \frac{9}{100}.$$

So $A < 3/10$ and Lemma 7.8 applies. Multiplying its bound with the two displayed estimates yields

$$\frac{|e_k(z)|^2}{\binom{n}{k}} \le \frac{2\sqrt{2}}{\pi}\frac{n}{n+2}e^{E_a+E_b}(1+A/10)^2$$
$$< \frac{792}{875}\frac{30}{29}\left(\frac{103}{100}\right)^2 = \frac{3150873}{3171875} < 1.$$

Here $\sqrt{2} < 99/70$ and $\pi > 25/8$; the first 64 terms of the alternating arctangent series give a rational lower bound on π exceeding $25/8$. The accompanying code checks all rational comparisons.

Now let $\min(k, n-k) \ge 80$. Then $v_0 \ge \min(k, n-k)/2 \ge 40$ and $k \ge 80$. Apply the same-n extremum reduction of Lemma 7.1, reflecting its binary profile so $r \le s$. If $r \le 10$, $k \ge 80 \ge 8r$ and Lemma 7.9 applies. If $r \ge 11$, both multiplicities are at least eleven and the joint assertion applies. These cases exhaust all profiles. No reduction of dimension, extrapolation of a finite table or continuous-multiplicity assumption enters the lift. □

7.3. *Complete bottom-half coefficient certificates.* We use $(n)_h = n(n-1)\cdots(n-h+1)$, with $(n)_0 = 1$. Every polynomial certificate in this section has nonnegative coefficients in nonnegative variables, so it proves an infinite-domain inequality. Exact determining grids check the polynomial identities; sample values alone never establish positivity.

7.3.1. *Bottom orders thirteen through forty.*

Lemma 7.11. *For $13 \le k \le 40$, $n \ge 2k$, and every centered real vector of squared norm n,*

$$|e_k|^2 \le \binom{n}{k}.$$

Proof: By Lemma 7.1, use binary multiplicities $r \le s$, $r + s = n$. Let $F_k(r, s)$ be the elementary coefficient of r copies of s and s copies of $-r$. These raw roots have squared norm nrs, so the desired inequality is exactly

$$F_k(r,s)^2 \le \binom{n}{k}(rs)^k. \tag{7.10}$$

The following table specifies a base b_k for each order.

k	13	14	15	16	17	18	19	20	21	22
b_k	7	8	8	9	9	10	10	11	11	12
k	23	24	25	26	27	28	29	30	31	32
b_k	12	14	13	16	15	18	18	21	20	23
k	33	34	35	36	37	38	39	40		
b_k	22	26	25	29	28	32	31	35		

For $r \ge b_k$ put $r = b_k + u$, $s = r + x$ with $u, x \ge 0$. Form the integer power-sum polynomials $p_h = rs^h + s(-r)^h$ and the factorial-scaled Newton polynomials

$$E_0 = 1, \qquad E_h = \sum_{a=1}^{h}(-1)^{a-1}\frac{(h-1)!}{(h-a)!}E_{h-a}p_a.$$

Thus $E_h = h!F_h$. The exact program constructs all twenty-eight polynomials

$$H_k(u,x) = k!(n)_k(rs)^k - E_k^2 \tag{7.11}$$

and checks every monomial coefficient is nonnegative. Hence $H_k \ge 0$ throughout the quadrant, which is (7.10) after division by $(k!)^2$.

Here and below the coefficient construction has an additional exact identity check. Since $p_1 = 0$, a Newton term for F_k contains at most $\lfloor k/2 \rfloor$ power-sum factors, so the total degree is at most

$k + \lfloor k/2 \rfloor$. Consequently H_k has u-degree at most $3k$ and x-degree at most $2k$. On a complete rectangular grid of one more than each degree, the program compares it with

$$F_k(r,s) = \sum_a \binom{r}{a}\binom{s}{k-a} s^a (-r)^{k-a}. \tag{7.12}$$

These comparisons determine the full polynomial; its coefficient signs prove positivity on the infinite quadrant.

Multiplicities one and two are covered by Lemma 7.5. For each remaining $3 \le r < b_k$, define

$$N_{r,k}(n) = \sum_{a=0}^{r} (-1)^a \binom{r}{a} r^{r-a} (k)_a n^a (n-a)_{r-a}. \tag{7.13}$$

The generating identity $(1+sT)^r(1-rT)^s = (1-rT)^n[1+nT/(1-rT)]^r$ gives

$$F_k = (-r)^k \binom{n}{k} \frac{N_{r,k}(n)}{r^r (n)_r}.$$

Thus (7.10) is equivalent, after clearing positive denominators, to

$$G_{r,k}(n) = k! r^{2r} (n)_r^2 (n-r)^k - r^k (n)_k N_{r,k}(n)^2 \ge 0. \tag{7.14}$$

For $k = 13$ the four rows $r = 3,4,5,6$ use $n = 13 + y$. For $14 \le k \le 40$ every row $3 \le r < b_k$ uses $n = \max(2k, 2r) + y$. All $y \ge 0$ cover the required integer domains, and every cleared factor is positive there. The program checks every power coefficient in these 409 ray polynomials is nonnegative. It also checks the coefficient factorization and the gap identity against (7.12) at degree-plus-one distinct nonnegative integer nodes in each row. The cases $r \ge b_k$, $3 \le r < b_k$ and $r = 1, 2$ exhaust all reflected profiles, proving the lemma. □

7.3.2. *Bottom orders forty-one through seventy-nine.*

Lemma 7.12. *For $41 \le k \le 79$, $n \ge 2k$, and every centered real vector of squared norm n, $|e_k|^2 \le \binom{n}{k}$.*

Proof: Use the same binary reduction. Multiplicities one and two are already covered. If $3 \le r \le 10$ and $k \ge 8r$, apply Lemma 7.9. The remaining pairs are exactly

$$41 \le k \le 79, \qquad \lfloor k/8 \rfloor + 1 \le r \le 10.$$

There are 115 such pairs. For each pair, construct (7.13) and (7.14) at $n = 2k + y$. Every power coefficient is nonnegative, as checked in full by the program, with the same complete univariate identity checks. This proves each entire ray $y \ge 0$.

For $r \ge 11$ the joint assertion of Theorem 7.10 applies as soon as

$$\frac{k(n-k)}{n} \ge 40, \quad \text{equivalently} \quad n \ge L_k := \left\lceil \frac{k^2}{k-40} \right\rceil.$$

The finite complement is exactly $2k \le n < L_k$, $11 \le r \le \lfloor n/2 \rfloor$. To compute the integer coefficient, logarithmically differentiate $(1+sT)^r(1-rT)^s$. With $d = s - r$ this gives

$$(1 + dT - rsT^2)F'(T) = -nrsTF(T),$$

and hence

$$F_0 = 1, \quad F_1 = 0, \quad (h+1)F_{h+1} = -dhF_h - rs(n-h+1)F_{h-1}. \tag{7.15}$$

The program checks (7.10) with exact integers at every profile using this recurrence. All 39 labels and 1,231,692 profiles pass. Each per-label count is also checked against $\sum_{n=2k}^{L_k-1}(\lfloor n/2 \rfloor - 10)$. The largest cutoff is $L_{41} = 1681$, so the finite interval ends at 1680, while its next endpoint is analytically covered. There is no extrapolation in n and no coefficient-reversal symmetry. □

Corollary 7.13. *For every centered real vector of squared norm n and every $0 \le k \le \lfloor n/2 \rfloor$, $|e_k|^2 \le \binom{n}{k}$.*

Proof: Lemmas 7.6 and 7.7 cover $k \le 12$. Lemmas 7.11 and 7.12 cover the remaining $13 \le k \le 79$ in their exact bottom-half domains. If $k \ge 80$ and $n \ge 2k$, then $n - k \ge k \ge 80$, so the full-vector assertion of Theorem 7.10 applies. This partition covers every bottom order at the same n. □

7.4. *Top boundary bridges.*

7.4.1. *Exact reciprocal normalization.* For a reflected binary profile put $d = s - r \ge 0$. Its normalized coordinates are $\sqrt{s/r}$ repeated r times and $-\sqrt{r/s}$ repeated s times. Their product has squared modulus $(r/s)^d$ and their reciprocals are $r/\sqrt{rs}$ repeated r times and $-s/\sqrt{rs}$ repeated s times. Define integers T_h by

$$F(T) = (1 + rT)^r (1 - sT)^s = \sum_{h=0}^{n} T_h T^h.$$

Elementary reciprocal duality gives

$$\begin{aligned} |e_{n-j}|^2 &= \left(\frac{r}{s}\right)^d \frac{T_j^2}{(rs)^j}, \\ |e_{n-j}|^2 \le \binom{n}{j} \quad &\Longleftrightarrow \quad T_j^2 r^d \le \binom{n}{j} (rs)^j s^d. \end{aligned} \tag{7.16}$$

In particular $d = 0$ causes no exceptional convention beyond $r^0 = s^0 = 1$. The differential identity

$$(1 - dT - rsT^2)F'(T) = -n(d + rsT)F(T)$$

yields the exact reverse recurrence

$$T_0 = 1, \quad T_1 = -nd, \quad (h+1)T_{h+1} = -d(n-h)T_h - rs(n-h+1)T_{h-1}. \tag{7.17}$$

These identities apply only after reduction to nonzero binary roots. The original vector may have zeros; Lemma 7.1 preserves its original n.

Lemma 7.14. *For $41 \le j \le 79$, $n \ge 2j$, and every centered real vector of squared norm n, $|e_{n-j}|^2 \le \binom{n}{j}$.*

Proof: Set $k = n - j$ and

$$L_j = \max\left\{ \left\lceil \frac{j^2}{j - 40} \right\rceil, j + 80 \right\} = \left\lceil \frac{j^2}{j - 40} \right\rceil.$$

The equality holds because $j^2/(j-40) - (j+80) = 40(80-j)/(j-40) > 0$. If $n \ge L_j$ and $r \ge 11$, then $j(n-j)/n \ge 40$ and the joint central theorem applies. If $r \le 10$, then $k \ge 80 \ge 8r$ and the sparse lemma applies.

For $2j \le n < L_j$ every $r \le \lfloor (n-j)/8 \rfloor$ is also sparse. The program checks the exact integer inequality in (7.16) for every remaining

$$\lfloor (n-j)/8 \rfloor + 1 \le r \le \lfloor n/2 \rfloor.$$

All 39 labels and 1,006,861 profiles pass. Exact per-label counts are compared with $\sum_{n=2j}^{L_j - 1} (\lfloor n/2 \rfloor - \lfloor (n-j)/8 \rfloor)$. Again $L_{41} = 1681$ is the maximum cutoff and its endpoint is analytically covered. The coefficient order $n - j \ge j \ge 41$ permits the same-n extremum lift, completing the full-real assertion. □

Lemma 7.15. *For $51 \le n \le 599$, every centered real vector of squared norm n satisfies $|e_k|^2 \le \binom{n}{k}$ at every order $0 \le k \le n$.*

Proof: For $k \le n/2$ use Corollary 7.13. Otherwise put $j = n - k < n/2$. At $j \ge 80$ both $k, j \ge 80$, so Theorem 7.10 applies; at $41 \le j \le 79$ use Lemma 7.14. The only remaining labels are $0 \le j \le 40$ with $n \ge 2j$.

The program checks (7.16) for every one of these forty-one labels and every reflected multiplicity, in exactly the domain

$$0 \le j \le 40, \qquad \max(51, 2j) \le n \le 599, \qquad 1 \le r \le \lfloor n/2 \rfloor.$$

There is no pruning. The reverse recurrence shares its coefficient computation across all j at each (n, r); the zero-order coefficient is $T_0 = 1$. All 3,645,330 scalar comparisons pass across 89,075 shared profiles. For each j the exact count is checked against $\sum_{n=\max(51,2j)}^{599} \lfloor n/2 \rfloor$. The coefficient order $k = n - j \ge \lceil n/2 \rceil \ge 26$ is in Lemma 7.1's range. Exact arithmetic and that lemma give the stated full-vector bounds. □

7.5. *An exact reciprocal strip in every degree at least six hundred.*

Theorem 7.16. *For $n \ge 600$, $1 \le j \le 19$, and every centered real vector of squared norm n,*

$$|e_{n-j}|^2 \le \binom{n}{j}.$$

Proof: Use Lemma 7.1 at order $k = n - j \ge 581$ and the binary notation of (7.16). Put

$$d = s - r, \quad u = r/s, \quad \eta = 1/n, \quad y = d^2/n, \quad a = 1 - y\eta = 4rs/n^2.$$

Define $A_h = h!2^h T_h / n^{3h/2}$. The reverse recurrence gives exactly

$$\begin{aligned} A_0 &= 1, \qquad A_1 = -2\sqrt{y}, \\ A_{h+1} &= -2\sqrt{y}(1 - h\eta)A_h - h(1 - y\eta)(1 - (h-1)\eta)A_{h-1}. \end{aligned} \tag{7.18}$$

There is no asymptotic error term. With

$$D_j(y, \eta) = (1 - y\eta)^j \prod_{i=0}^{j-1} (1 - i\eta),$$

the squared ratio in (7.16) is

$$\frac{|e_{n-j}|^2}{\binom{n}{j}} = \frac{u^d A_j^2}{j! D_j(y, \eta)}. \tag{7.19}$$

All denominator factors are positive. For $v = d/n \in [0, 1)$, $\log((1-v)/(1+v)) \le -2v$, so $u^d \le e^{-2y}$; at $d = 0$ equality holds with $u^0 = 1$. It therefore suffices to bound A_j^2 by $j! D_j e^{2y}$.

7.5.1. *The compact finite-n polynomial certificate.* Parity gives $A_h = y^{(h \bmod 2)/2} B_h(y, \eta)$, where $B_0 = 1, B_1 = -2$ and

$$B_{h+1} = -2y^{h \bmod 2}(1 - h\eta)B_h - h(1 - y\eta)(1 - (h-1)\eta)B_{h-1}. \tag{7.20}$$

Let $N_j = 4j + 40$ and $S_N(2y) = \sum_{h=0}^{N} (2y)^h / h!$. The integer polynomial

$$P_j = N_j! \left[j! D_j(y, \eta) S_{N_j}(2y) - y^{j \bmod 2} B_j(y, \eta)^2 \right] \tag{7.21}$$

has nonnegative complete tensor-product Bernstein arrays on each of

$$\left[0, \frac{25}{4}\right], \quad \left[\frac{25}{4}, \frac{25}{2}\right], \quad \left[\frac{25}{2}, 25\right], \quad [25, 50], \qquad \eta \in [0, 1/600]. \tag{7.22}$$

All 19 polynomials and 76 rectangles pass. The degrees are at most $N_j + j$ and $2j - 1$ in y, η, respectively, with maxima 135 and 37. The program reconstructs every basis entry and accepts a rectangle only if all its entries are nonnegative. Throughout this enlarged continuous rectangle

$1 - y\eta \geq 11/12$, so every denominator factor remains positive. As $e^{2y} \geq S_N(2y)$, the certificate proves the desired ratio bound for every $0 \leq y \leq 50$.

Here is the exact basis conversion and coverage argument. If M, Z are the two degrees, substitute $y = 50t, \eta = w/600$ and multiply P_j by 600^Z to get integer power coefficients c_{bc}. Set $l_y = \operatorname{lcm}_b \binom{M}{b}$ and $l_\eta = \operatorname{lcm}_c \binom{Z}{c}$. The tensor-product Bernstein coefficient at (i, h), multiplied by the positive integer $l_y l_\eta$, is

$$\sum_{b \leq i,\ c \leq h} c_{bc} \binom{i}{b} \frac{l_y}{\binom{M}{b}} \binom{h}{c} \frac{l_\eta}{\binom{Z}{c}}. \tag{7.23}$$

This follows by applying $t^b = \sum_{i=b}^{M} \binom{i}{b} \binom{M}{b}^{-1} \binom{M}{i} t^i (1-t)^{M-i}$ in both variables. All basis functions are nonnegative. Exact midpoint de Casteljau subdivision uses sums and a common positive denominator; common positive gcd factors can be removed. Both children are retained, and an already-positive interval can be refined further to obtain the common partition (7.22). The code explicitly checks those four intervals, the full η interval and total area $1/12$. Thus the finite computation covers the full rectangle, not a set of sample points. Optional output identifies each reconstructed full array by its stream hash; its sign proof is the checked entries themselves, not the hash.

7.5.2. *An analytic tail with translated polynomial certificates.* If $r \leq k/8$, Lemma 7.9 already proves the result at the original order $k = n - j$. Otherwise

$$\frac{d}{n} < \frac{3 + j/n}{4} \leq \frac{1819}{2400} < \frac{19}{25}.$$

Consequently $a > 264/625 > 2/5$, and the elementary product inequality gives

$$\prod_{i=0}^{j-1} (1 - i/n) \geq 1 - \frac{j(j-1)}{2n} \geq \frac{143}{200} > \frac{7}{10}.$$

Let $C_0(X) = 1, C_1(X) = X$ and $C_{h+1}(X) = XC_h(X) + hC_{h-1}(X)$. These are the absolute-coefficient Hermite polynomials, explicitly

$$C_j(X) = \sum_{l=0}^{\lfloor j/2 \rfloor} \frac{j! X^{j-2l}}{2^l l! (j-2l)!}.$$

In (7.18) every coefficient factor apart from $2\sqrt{y}$ and h lies in $[0, 1]$ on the actual domain. Induction yields $|A_j| \leq C_j(2\sqrt{y})$. Thus (7.19) is at most

$$\frac{10}{7} \left(\frac{5}{2}\right)^j e^{-2y} \frac{C_j(2\sqrt{y})^2}{j!}. \tag{7.24}$$

Parity makes $C_j(2\sqrt{y})^2$ an integer polynomial in y. Define

$$Q_j(y) = S_{N_j}(2y) - \frac{10}{7} \left(\frac{5}{2}\right)^j \frac{C_j(2\sqrt{y})^2}{j!}.$$

After multiplication by $7\, 2^j j! N_j!$, the program checks that every power coefficient of $Q_j(50 + t)$ is nonnegative, for every $1 \leq j \leq 19$. More explicitly the unshifted coefficient of y^h is

$$7\, 2^j j! 2^h \frac{N_j!}{h!} - 10\, 5^j N_j!\, [y^h] C_j(2\sqrt{y})^2,$$

with the second term zero above degree j. Translation uses $[t^b]P(50 + t) = \sum_{h \geq b} [y^h]P(y) \binom{h}{b} 50^{h-b}$. All 19 complete translated arrays pass. Hence $Q_j(y) \geq 0$ for every $y \geq 50$, and $e^{2y} \geq S_N(2y)$ makes (7.24) at most one on that entire half-line.

The sparse cases and the compact and tail regions exhaust all binary profiles, including balance at $y = 0$ and the shared boundary $y = 50$. Lemma 7.1 completes the full-real conclusion without reciprocal operations on original zero coordinates. The proof controls finite n exactly. □

7.6. *A global affine contour budget for top codimensions nineteen through forty.*

Theorem 7.17. *For $n \geq 600$, $19 \leq j \leq 40$, and every centered real vector of squared norm n, $|e_{n-j}|^2 \leq \binom{n}{j}$.*

Proof: Use the binary Cauchy modulus f, entropy factor H_0 and area A in the proof of Theorem 7.3; put $k = n - j$, $v = k(n-k)/n$. The interval rearrangement and layer-cake argument in Lemma 7.8 gives, without an area restriction,

$$I := \frac{1}{\pi}\int_{-1}^{1} \frac{f(t)}{\sqrt{1-t^2}}\,dt \leq \int_0^1 \Phi(\ell(h))\,dh, \qquad \Phi(a) = \frac{\arccos(1-a)}{\pi}, \quad \int_0^1 \ell(h)\,dh = A.$$

The tangent at $a_0 = 3/10$ is $T(a) = \alpha + \beta a$, where

$$\beta = \frac{1}{\pi\sqrt{51/100}}, \qquad \alpha = \frac{\arccos(7/10)}{\pi} - \frac{3}{10}\beta.$$

It is a global majorant of Φ on $[0,2]$. Indeed Φ is concave on $[0,1]$, while $T - \Phi$ is concave on $[1,2]$, whose minimum is at an endpoint. Endpoint one is covered by the tangent inequality, and at two

$$\pi T(2) = \arccos(7/10) + \sqrt{17/3} > \frac{77}{100} + \frac{119}{50} = \frac{63}{20} > \frac{22}{7} > \pi.$$

The first bound follows from $\cos(77/100) \geq 1 - (77/100)^2/2 > 7/10$ and $(119/50)^2 < 17/3$. Therefore

$$I \leq \alpha + \beta A \qquad (0 \leq A \leq 2). \tag{7.25}$$

Both coefficients are positive. The following conservative rational upper bounds are certified exactly:

$$\alpha < a := \frac{11947}{100000}, \quad \beta < b_0 := \frac{44573}{100000}, \quad \sqrt{\pi} < p_0 := \frac{177246}{100000}, \quad \sqrt{2\pi} < h_0 := \frac{250663}{100000}.$$

For completeness, Machin's identity

$$\pi = 16\arctan(1/5) - 4\arctan(1/239)$$

follows from $\tan(4\arctan(1/5)) = 120/119$ and the tangent subtraction formula on the positive principal branch. Alternating partial sums with eight/nine terms at $1/5$ and three/four terms at $1/239$ give

$$\frac{314159}{100000} < \pi < \frac{3141593}{1000000}.$$

Also $71414/100000 < \sqrt{51/100} < 71415/100000$, and the alternating degree-eight Taylor upper bound for $\cos(7954/10000)$ is below $7/10$, so $\arccos(7/10) < 7954/10000$. These bounds give the upper β by lower bounding both denominator factors, and the upper α by upper bounding its positive numerator and lower bounding π. The numerator is positive, for instance from $\arccos(.7) > .77$ and $.3/\sqrt{.51} < .3/.7$. Squaring checks the two radical upper bounds. The program verifies all these rational comparisons; no decimal transcendental approximation decides acceptance.

If $r \leq k/8$, use Lemma 7.9. Otherwise

$$r > k/8 \geq (n-40)/8 \geq 7n/60, \qquad rs > 371n^2/3600.$$

The Beta and Stirling bounds proved earlier are

$$A \leq \sqrt{\pi/v}\,\frac{n}{n+2}e^{E_a}, \qquad \frac{H_0^2}{\binom{n}{k}} \leq \sqrt{2\pi v}\,e^{E_b}, \qquad E_a = \frac{n}{6rs}, \quad E_b = \frac{1}{12v}.$$

Here $E_a < 1/371$, while $11039/600 \le v \le 40$ gives $E_b \le 50/11039$. The lower endpoint follows because $j(1-j/600)$ increases on $19 \le j \le 40$. Since $e^t \le (1-t)^{-1}$ for $0 \le t < 1$, $e^{E_a} \le 371/370$ and $e^{E_b} \le 11039/10989$.

Drop the favorable $n/(n+2)$ and put

$$b = b_0 p_0 \frac{371}{370}, \qquad C = h_0 \frac{11039}{10989}.$$

Equations (7.25) and the coefficient contour bound yield

$$\frac{|e_k|^2}{\binom{n}{k}} \le C\left(a^2 t + 2ab + \frac{b^2}{t}\right), \qquad t = \sqrt{v}.$$

The right side is convex in $t > 0$, as its second derivative is $2Cb^2/t^3 > 0$. At the two endpoints use

$$\frac{42893}{10000} < \sqrt{11039/600} < \frac{42894}{10000}, \qquad \frac{63245}{10000} < \sqrt{40} < \frac{63246}{10000}.$$

The positive linear term uses the upper enclosure and the reciprocal term the lower. Both resulting rational upper bounds are strictly below one:

$$C\left(a^2 \frac{42894}{10000} + 2ab + b^2 \frac{10000}{42893}\right) < 1,$$
$$C\left(a^2 \frac{63246}{10000} + 2ab + b^2 \frac{10000}{63245}\right) < 1.$$

These are exact fraction comparisons (approximately .99918 and .95378). Convexity proves the entire nonsparse binary domain; the sparse lemma and same-n extrema finish the theorem. □

7.7. *All scalar orders and the covariance certificate.*

Proof of Theorem 1.2: First normalize a positive squared norm to $\sum z_i^2 = n$; the centered normalized slice is empty at $n = 1$. For $2 \le n \le 50$, use Lemma 7.4. For $51 \le n \le 599$, use Lemma 7.15. For $n \ge 600$, every $k \le \lfloor n/2 \rfloor$ follows from Corollary 7.13. In the top half put $j = n - k < n/2$. At $j = 0$, product AM–GM gives $|e_n|^2 \le 1$. At $1 \le j \le 19$, use Theorem 7.16; at $19 \le j \le 40$, use Theorem 7.17. The overlap is harmless. At $41 \le j \le 79$, Lemma 7.14 applies since $n \ge 2j$. At $j \ge 80$, both $j, k \ge 80$ and Theorem 7.10 applies. This exhausts every coefficient and degree at the same n, including the $n = 50/51$ and $599/600$ endpoints and equality $k = n/2$.

Homogeneity of e_k now gives (1.7) for any positive squared norm. At zero norm every coordinate vanishes, so all $e_k = 0$ for $k \ge 1$, while $e_0 = 1$ separately; no 0^0 convention is needed. For even n, the vector with $n/2$ copies of one and $n/2$ copies of minus one has squared norm n and $|e_n|^2 = 1 = \binom{n}{n}$. Thus the universal scalar coefficient one is optimal. Original zeros, repeated coordinates and all balanced or unbalanced profiles are included by the same-n reductions. □

7.8. *Equality and endpoint stability of the scalar bound.*

Theorem 7.18. *On the normalized real slice $\sum_i z_i = 0$, $\sum_i z_i^2 = n$ with $n \ge 2$, equality in $|e_k(z)|^2 \le \binom{n}{k}$ holds exactly at $k = 0$, or at $k = n$ with even n and balanced entries $+1, -1$ up to permutation. Every interior order $1 \le k < n$ is strict.*

Proof: For each fixed $n, k \ge 3$, Lemma 7.1 supplies some binary maximizer at that SAME n. It therefore suffices to prove strictness for every binary interior profile; it is not necessary that every maximizer itself be binary. Orders one and two are immediate from $e_1 = 0$ and the ratio $n/[2(n-1)] < 1$ for $n > 2$. At $k = n$, product AM–GM on z_i^2 gives equality only when all $|z_i| = 1$; centering forces even n and balanced signs.

Here are complete strictness gates on the proof's infinite domains. Multiplicity one has ratio at most $2/3$ for $k \ge 3$. Multiplicity two has strict cubic and quartic bounds except the balanced

$n = k = 4$ endpoint; its seven cleared polynomials have strictly positive coefficients, and its factorial tail is strict. The general cubic bound is strict, and the quartic bound is strict for $n > 4$.

For each fixed order $5 \le k \le 12$, the accepted quadrant gap arrays have nonnegative coefficients, with strictly positive coefficients at both u and x. At odd k the constant is positive. At even k the only zero on the nonnegative quadrant is therefore $u = x = 0$, namely $n = k$, $r = s = k/2$. Every smaller-multiplicity exception ray has positive constant. All 28 bottom13–40 quadrant gaps, all 409 corresponding rays and all 115 bottom41–79 rays have positive constant and nonnegative coefficients. Their cleared denominators are positive. Thus these entire quadrants and rays are strict. The central and sparse bounds have ratios strictly below one, as do both global-tangent endpoint guards.

For the reciprocal strip, if $d = s - r > 0$, then $y > 0$ and $e^{2y} > S_N(2y)$ strictly. The compact and translated-tail certificates, with their positive denominator, therefore give strict ratio below one even if a polynomial gap vanishes. At $d = 0$, the normalized generating polynomial is $(1 - T^2)^{n/2}$. Odd coefficients vanish. For $k = 2l$ interior, Vandermonde gives

$$\binom{n}{2l} = \sum_a \binom{n/2}{a}\binom{n/2}{2l-a} > \binom{n/2}{l}^2,$$

because besides $a = l$, the term $a = l - 1$ is positive for $1 \le l \le n/2 - 1$.

The remaining domains are finite and explicit. The companion verifier checks all 40,425 small scalar labels, 1,231,692 bottom bridge profiles, 1,006,861 top bridge profiles and 3,645,330 finite endpoint labels for STRICT inequality, permitting equality only at original order $k = n$ with $r = s$. It also checks all 572 coefficient strictness/zero-set gates just described on the reconstructed complete arrays. These are the full finite complements of the analytic partition, not a cutoff extrapolation. The partition in the proof of Theorem 1.2 now proves binary strictness at every interior order; Lemma 7.1 completes the full-real assertion. □

The classification completes the strictness analysis of Theorem 1.2. Its elementary endpoint consequences are recorded for clarity. Put $\varrho = |e_n(z)|^2 > 0$. The classical logarithmic AM–GM inequality gives

$$\sum_i (|z_i| - 1)^2 \le -\log \varrho. \tag{7.26}$$

Indeed $\log u \le 2(\sqrt{u} - 1)$ and $\sum z_i^2 = n$ imply $\log \varrho \le 2\sum_i (|z_i| - 1) = -\sum_i (|z_i| - 1)^2$. The square-root variance refinement of AM–GM in Aldaz (2009, Theorem 0.1) also gives the stronger intermediate formula

$$\varrho^{1/n} \le \left(1 - \frac{1}{2n}\sum_i (|z_i| - 1)^2\right)^2.$$

If $\varrho \ge 1 - \epsilon > 0$, let $\sigma_i = \operatorname{sign}(z_i)$. Centering and Cauchy yield

$$|\sum_i \sigma_i| \le \sqrt{n \sum_i (|z_i| - 1)^2}.$$

If $-\log(1 - \epsilon) < 1/n$, the integer sign imbalance is less than one and hence zero. Then n is even, σ is an exact balanced equality vector, and $\|z - \sigma\|_2^2 \le -\log(1 - \epsilon)$. This is endpoint stability only.

For odd $n \ge 3$, classical two-moment product optimization gives the exact endpoint maximum

$$\max |e_n(z)|^2 = \frac{n - 1}{n + 1}. \tag{7.27}$$

For completeness, a maximizer has nonzero product since a binary competitor does. Lagrange multipliers for $\sum \log z_i^2$ are valid: the constraint gradients are independent. The equation $2/z_i = \alpha + 2\beta z_i$, summed after multiplying by z_i, gives $\beta = 1$. All coordinates therefore take at most two values, necessarily of opposite sign. Their reflected multiplicities $r \le s$ have squared product

$(r/s)^{s-r}$. Writing $d = s-r$, its logarithm $d\log((n-d)/(n+d))$ strictly decreases for $d > 0$. Odd n makes the smallest positive d equal one, proving (7.27), with exactly the nearest-balanced two-value profiles up to sign and permutation. This is the displayed elementary endpoint optimization.

8. Completion of the count-law certificate

Proof of Theorem 1.3: For $n \geq 2$, scaling Theorem 1.2 with (2.2) gives

$$|R_k(y)|^2 \leq \binom{n}{k}(2/n)^k (y^\top B y)^k.$$

The quotient tensor transfer therefore yields $\mathbb{E}_M U_k^2 \leq (2W)^k$. Orthogonality proves (1.8)'s chi-square bound for every $W \geq 0$; the sum is finite, so no geometric convergence hypothesis is required. At $2W = 1$ it is $n-1$, and for $2W < 1$ it is at most $(2W)^2/(1-2W)$.

For the TV conclusion let $x = 2W \in (0,1)$ and $\delta = \mathrm{TV}(P, M)$. Then $v := \chi^2(P\|M) \leq \sum_{k=2}^n x^k \leq x^2/(1-x)$. If $x \leq 1/2$, Cauchy gives $\delta \leq \sqrt{v}/2 \leq x$. If $x \geq 1/2$ and $\delta \leq 1/2$ the result is immediate; otherwise $\delta > x$ contradicts Lemma 3.2, since $v \geq \delta/(1-\delta) > x/(1-x) \geq v$. At $W = 0$ use rigidity and at $W \geq 1/2$ use $\mathrm{TV} \leq 1$. Degree one is exactly multinomial, giving an empty chi-square sum; one retained category is deterministic, and mean-zero categories vanish almost surely and can be deleted.

For optimality, take even n and $p(x,y) = x^{n/2}y^{n/2}$, the classical balanced deterministic/binomial example Roos (2010, Remark 2.2(a)). It has $W = 1/2$ and $\mathrm{TV}(P, \mathrm{Mult}(n, 1/2)) = 1 - 2^{-n}\binom{n}{n/2} \to 1$. Thus no universal linear coefficient below two is possible. □

8.1. *The equivalent own-mean Krawtchouk bound.* In the conventional normalization National Institute of Standards and Technology (n.d., equation 18.23.3 and Table 18.19.1),

$$(1 - [(1-p)/p]t)^x (1+t)^{n-x} = \sum_{k=0}^n \binom{n}{k} K_k(x; p, n) t^k.$$

The squared norm under the binomial probability weight $w_p(x) = \binom{n}{x}p^x(1-p)^{n-x}$ is $[(1-p)/p]^k/\binom{n}{k}$. Thus the binomial-orthonormal polynomial is

$$\widehat{K}_k(x; p, n) = \sqrt{\binom{n}{k}}\,[p/(1-p)]^{k/2} K_k(x; p, n).$$

Put $s = np$, $r = n - s$ with positive integer r, s, and substitute $t = sT$ in the generating function. The raw coefficient of $(1+sT)^r(1-rT)^s$ is $s^k\binom{n}{k}K_k(s; s/n, n)$. After division by $(rs)^{k/2}\sqrt{\binom{n}{k}}$, this yields

$$\frac{e_k}{\sqrt{\binom{n}{k}}} = \widehat{K}_k(np; p, n).$$

Consequently Theorem 1.2 is equivalent, through the same-n two-value reduction and elementary low orders, to

$$|\widehat{K}_k(np; p, n)| \leq 1 \quad (0 < p < 1,\ np \in \mathbb{Z},\ 0 \leq k \leq n).$$

The integer-mean condition is substantive. This normalization is distinct from the counting-measure orthonormal function $\sqrt{w_p(x)}\widehat{K}_k(x)$. Entrywise orthogonal-matrix bounds for that function only give $|\widehat{K}_k(x)| \leq 1/\sqrt{w_p(x)}$, which does not imply the own-mean bound one. The generating function and norm are classical inputs.

8.2. *Reproducibility.* Run `verify_spectral.py` with its shipped helpers `verify_scalar_one.py`, `verify_sharp_two.py` and `verify_refined.py`; the option `-output`, followed by a filename, saves exact certificate data. The program checks the 40,425 small scalar comparisons, all multiplicity-two and fixed-order certificate arrays, the complete bottom and top boundary domains, all 3,645,330 finite endpoint comparisons through 599, the 19 reciprocal-strip polynomials with 76 full compact basis reconstructions and 19 translated tails, and every global-tangent rational guard. Every accepted compact basis entry is regenerated and sign-checked. Saved array hashes identify those computations. Infinite polynomial quadrants/rays and finite exhaustive domains have distinct, explicit coverage arguments. All acceptance gates use exact integers/rationals and remain active under Python optimization. The entry point also verifies all 572 equality strictness gates, the complete finite strictness domains and exact small-law spectral/example identities.

9. Gaussian mutual-TV limits along powers of a fixed law

We now fix one law P throughout, including its dimension, support and probabilities. The following result concerns mutual TV of two discrete convolution laws.

Theorem 9.1. *Let $C = \mathrm{Cov}(P)$, $S = \mathrm{diag}(t) - tt^\top/n$, and $H = \{x : \mathbf{1}^\top x = 0\}$. For $P_N = P^{*N}$ and $M_N = M^{*N} = \mathrm{Mult}(Nn, t/n)$,*

$$\lim_{N\to\infty} \mathrm{TV}(P_N, M_N) = \mathrm{TV}(\mathcal{N}_H(0, C), \mathcal{N}_H(0, S)). \tag{9.1}$$

Degenerate C is allowed. If $0 < W < 1/2$, the right side is strictly below $2W$. Consequently every fixed P with $0 < W < 1/2$ has a threshold $N_0(P)$ such that $\mathrm{TV}(P_N, M_N) < 2W$ for every integer $N \geq N_0(P)$. At $W = 0$ all powers are exactly multinomial; at $W \geq 1/2$ the inequality holds at every power.

The support input is classical: Brändén's jump-system theorem (2007, Theorem 3.2) extends earlier homogeneous support results of Choe, Oxley, Sokal and Wagner (2004). The finite-support Fourier local limit used below is also classical; see the background and lattice-period discussion in Randles (2023). We give the needed argument, including the l1 passage and singular case, to expose the exact hypotheses.

Proof of Theorem 9.1: **Covariance coordinates and support lattice.** For p^N, the mean, covariance, S and B are multiplied by N, so $W_N = W$. Products preserve stability, homogeneity and normalization. If $m = 1$ all laws are identical point masses. Suppose $m \geq 2$, put $d = m - 1$, and delete the final coordinate. The omitted count is determined by the total, so this is a bijection on each support and preserves TV. Denote the principal covariance blocks by C', S', B', and the common projected block mean by $\mu = (t_1, \ldots, t_d)$.

The matrix S' is positive definite and

$$(S')^{-1} = \mathrm{diag}(1/t_1, \ldots, 1/t_d) + \frac{1}{t_m}\mathbf{1}\mathbf{1}^\top.$$

Since $B\mathbf{1} = 0$, $\sum_{i,j\leq d} B_{ij} = B_{mm}$, hence

$$\mathrm{Tr}((S')^{-1}B') = W. \tag{9.2}$$

When $0 < W < 1/2$, the matrix $D = 2(S')^{-1/2}B'(S')^{-1/2}$ is positive semidefinite with trace $2W < 1$. Thus every eigenvalue is below one and $C' = (S')^{1/2}(I - D)(S')^{1/2}$ is positive definite. In particular C is then nondegenerate on H. More generally, C is positive definite on H exactly when C' is positive definite, since centered count vectors are obtained from their first d coordinates by the linear isomorphism $x' \mapsto (x', -\mathbf{1}^\top x')$ onto H.

Assume first that C is positive definite on H. The support $\mathcal{A}$ of p obeys unit exchange: if $a, b \in \mathcal{A}$ and $a_i > b_i$, some j with $a_j < b_j$ has $a - e_i + e_j \in \mathcal{A}$. Indeed the jump-system two-step axiom

applies to the step $-e_i$ toward b; its first step leaves the fixed-total support, and homogeneity forces the second step to be $+e_j$ toward a deficient coordinate. Repeated exchange decreases the distance to b by two and connects every pair of support points.

Make a graph on categories by joining i, j whenever two support points differ by $e_i - e_j$. If it were disconnected, each graph-component total would be constant throughout the support: every exchange path stays within its components. Subtracting a multiple of $\mathbf{1}$ from a proper component indicator would give a nonzero vector in H with zero variance, contradicting positive definiteness of C there. The graph is therefore connected. Its edge differences generate the full root lattice $\{z \in \mathbb{Z}^m : \mathbf{1}^\top z = 0\}$, and all support differences already lie in this lattice. Consequently the projected support-difference group is exactly $\mathbb{Z}^d$.

For a characteristic function $\phi(\theta)$, modulus one occurs precisely when all support phases agree. The full difference group therefore implies $|\phi(\theta)| < 1$ away from the origin of the Fourier torus. The projected multinomial block law has the same aperiodicity: its support contains zero and each unit vector, because all categorical probabilities are positive. Thus both block laws have positive covariance, full difference group and the same mean μ. Positive covariance alone would not exclude lattice periods for an arbitrary discrete law; the stable-support argument is essential here.

Uniform local limit and its l1 consequence. Let X be a fixed finite-support $\mathbb{Z}^d$ law with mean μ, positive definite covariance Σ and full support-difference group. Write $f_N(k) = \Pr(X_1 + \cdots + X_N = k)$ and g_Σ for the centered Gaussian density. Its centered characteristic function satisfies

$$\psi(\theta) = e^{-i\mu\cdot\theta}\phi(\theta) = 1 - \tfrac{1}{2}\theta^\top\Sigma\theta + O(\|\theta\|^3).$$

Positive covariance gives $|\psi(\theta)| \le e^{-c\|\theta\|^2}$ near zero for some fixed $c > 0$. On the compact complement of that neighborhood in the torus, aperiodicity gives $|\phi| \le r < 1$.

Fourier inversion and $\theta = u/\sqrt{N}$ show

$$\sup_{k\in\mathbb{Z}^d}\left|N^{d/2}f_N(k) - g_\Sigma\left(\frac{k-N\mu}{\sqrt{N}}\right)\right| \longrightarrow 0. \tag{9.3}$$

Here is the uniform error argument. On bounded u-sets, $\psi(u/\sqrt{N})^N$ converges uniformly to $e^{-u^\top\Sigma u/2}$. On the expanding small-θ region it is dominated by $e^{-c\|u\|^2}$, so dominated convergence controls the integral error there. The complementary scaled Fourier integral is bounded by a constant times $N^{d/2}r^N$, and the Gaussian integral outside the expanding region also tends to zero. The phase $e^{-i(k-N\mu)\cdot u/\sqrt{N}}$ has modulus one, making the error estimate uniform in k. This proves (9.3) with the usual Gaussian Fourier normalization.

It follows that

$$\sum_{k\in\mathbb{Z}^d}\left|f_N(k) - N^{-d/2}g_\Sigma\left(\frac{k-N\mu}{\sqrt{N}}\right)\right| \longrightarrow 0. \tag{9.4}$$

For fixed R, the box $\|k-N\mu\|_\infty \le R\sqrt{N}$ has $O_R(N^{d/2})$ sites, so (9.3) sums to $o_R(1)$ there. Outside, Chebyshev bounds the probability mass by $\operatorname{Tr}\Sigma/R^2$. Shifted-grid Gaussian sums converge to the corresponding integrals: on compact sets this follows from uniform continuity, and Gaussian decay controls the tails uniformly in the grid shift. Thus their tail also vanishes after first taking $N \to \infty$ and then $R \to \infty$. This proves (9.4).

Apply this to the projected P and M blocks. Their two l1 errors imply

$$\mathrm{TV}(P_N, M_N) - \frac{1}{2}N^{-d/2}\sum_{k\in\mathbb{Z}^d}\left|g_{C'}\left(\frac{k-N\mu}{\sqrt{N}}\right) - g_{S'}\left(\frac{k-N\mu}{\sqrt{N}}\right)\right| \longrightarrow 0.$$

Another shifted-grid Riemann sum gives the Gaussian TV on the right side of (9.1). The coordinate isomorphism onto H preserves TV, proving the claimed coordinate-independent limit when C is nonsingular there.

Singular covariances and the strict Gaussian gap. If C is singular on H, choose nonzero $v \in H$ with $v^\top C v = 0$. Then $v \cdot (A - t) = 0$ almost surely. Thus P_N is supported on $E_N = \{a : v \cdot (a - Nt) = 0\}$. Under M_N, the scalar projection $v \cdot (A - Nt)/\sqrt{N}$ converges by the ordinary scalar CLT to a nondegenerate normal, since $v^\top S v > 0$. For every $\eta > 0$, Portmanteau on the closed interval $[-\eta, \eta]$ gives

$$\limsup_N M_N(E_N) \le \Pr(|\mathcal{N}(0, v^\top S v)| \le \eta).$$

Letting $\eta \downarrow 0$ shows $M_N(E_N) \to 0$. Hence $\mathrm{TV}(P_N, M_N) \ge 1 - M_N(E_N) \to 1$. The degenerate Gaussian with covariance C is supported on the corresponding hyperplane, to which the Gaussian with covariance S assigns zero mass; their TV is one. This proves (9.1) also in the singular case.

Finally suppose $0 < W < 1/2$, so the projected covariances are positive definite. By (9.2), $D \succeq 0$ and $\operatorname{Tr} D = 2W < 1$. For the two centered Gaussian densities,

$$\frac{g_{S'}(x)}{g_{C'}(x)} = \sqrt{\det(I - D)}\, e^{x^\top((C')^{-1} - (S')^{-1})x/2} \ge \sqrt{\det(I - D)}.$$

The quadratic exponent is nonnegative by covariance ordering. Integrating the minimum of the densities bounds their TV by $1 - \sqrt{\det(I - D)}$. If λ_j are the eigenvalues of D, then

$$\det(I - D) = \prod_j (1 - \lambda_j) \ge 1 - \sum_j \lambda_j = 1 - 2W.$$

This elementary product inequality follows by induction. Since $\sqrt{a} > a$ for $0 < a < 1$, the Gaussian TV is strictly below $2W$. Convergence in (9.1) leaves a fixed positive gap, yielding $N_0(P)$ for all subsequent integers. Rigidity handles $W = 0$, while $W \ge 1/2$ is trivial. This completes the proof. □

The coefficient two in this eventual statement is optimal uniformly over fixed base laws. For $p(x, y) = xy$, one has $W = 1/2$ and

$$\mathrm{TV}(P_N, M_N) = 1 - \binom{2N}{N}/4^N \longrightarrow 1.$$

Theorem 1.3 already bounds every power, because $W_N = W$. The limit theorem additionally identifies the limiting value and its strict Gaussian gap, with its threshold attached to the fixed base law. Total variation under powers need not be monotone: $(x/4 + 3y/4)(3x/4 + y/4)$ has TV $1/8$ from its degree-two reference, while its square has TV $11/128$; xy has TV $1/2$, while its square has TV $5/8$. In each example W is unchanged under powers. These exact examples prevent pulling the limiting bound back by an assumed monotonicity.

10. Consequences and interpretation

The covariance deficit controls the likelihood ratio through its exact quadratic component and a spectral remainder. The sharp real coefficient inequality supplies the uniform certificate, while the complex-six theorem gives a separate sharp fifth-component refinement. Their proofs retain the complete analytic reductions, finite-domain coverage and equality arguments. The full-stable cubic obstruction and rare-cell examples delimit stronger conclusions that cannot be inferred from the corresponding product-class or fixed-reference estimates. The fixed-law convolution limit has its own asymptotic quantifiers.

The certificates take a real-stable generating polynomial and its mean vector and covariance as inputs. These hypotheses retain the full count-law structure behind the moment comparison. Singular covariance deficit, repeated roots and zero coefficients are included. The positivity lemma is applied only to the complete density P/M, since individual centered-expansion terms are signed.

Reproducibility and assistance

The five supplied standard-library Python files reconstruct the finite polynomial identities, exhaustive scalar domains and twenty rational Bernstein intervals used in the proof. The weighted-median reduction, real-root perturbation, equality analysis, stable-law transfer and Gaussian limit arguments are given analytically in the manuscript. Reader commands and the complete member inventory are in the source package's README.

OpenAI ChatGPT and Codex assisted mathematical exploration, literature comparison, drafting, exact-code development and proof checks. The author assumes responsibility for the manuscript and its mathematical claims.

References

Aldaz, J. M. Self-improvement of the inequality between arithmetic and geometric means. *Journal of Mathematical Inequalities*, **3** (2), 213–216 (2009). arXiv:0807.1788.

Banach, S. Über homogene Polynome in L^2. *Studia Mathematica*, **7** (1), 36–44 (1938). doi: 10.4064/sm-7-1-36-44.

Bauschke, H. H., Güler, O., Lewis, A. S., and Sendov, H. S. Hyperbolic Polynomials and Convex Analysis. *Canadian Journal of Mathematics*, **53** (3), 470–488 (2001). URL https://people.orie.cornell.edu/aslewis/publications/01-hyperbolic.pdf.

Brändén, P. Polynomials with the half-plane property and matroid theory. *Advances in Mathematics*, **216**, 302–320 (2007). arXiv:math/0605678; version 4.

Cao, S., He, S., Li, Z., and Wang, Z. Extreme Ratio Between Spectral and Frobenius Norms of Nonnegative Tensors. *SIAM Journal on Matrix Analysis and Applications*, **44** (2), 919–944 (2023). doi:10.1137/22M1502951. URL https://pure.port.ac.uk/ws/portalfiles/portal/79417900/norm_extreme_final.pdf. Theorem 4.6 and Corollary 4.9.

Carando, D. and Rodríguez, J. T. Symmetric multilinear forms on Hilbert spaces: where do they attain their norm? (2018). doi:10.48550/arXiv.1810.09373. arXiv:1810.09373.

Choe, Y.-B., Oxley, J. G., Sokal, A. D., and Wagner, D. G. Homogeneous multivariate polynomials with the half-plane property. *Advances in Applied Mathematics*, **32**, 88–187 (2004). doi:10.1016/S0196-8858(03)00078-2. arXiv:math/0202034.

Derksen, H., Ivanov, P., Lee, C. H., and Viola, E. Pseudorandomness, Symmetry, Smoothing: I. In *39th Computational Complexity Conference (CCC 2024)*, volume 300 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pp. 18:1–18:27. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik (2024). doi:10.4230/LIPIcs.CCC.2024.18.

Diaconis, P. and Griffiths, R. C. Reproducing kernel orthogonal polynomials on the multinomial distribution. *Journal of Approximation Theory*, **242**, 1–30 (2019). doi:10.1016/j.jat.2019.01.007. arXiv:1802.07122.

Ehm, W. Binomial approximation to the Poisson binomial distribution. *Statistics and Probability Letters*, **11** (1), 7–16 (1991). doi:10.1016/0167-7152(91)90170-V.

Guntuboyina, A., Saha, S., and Schiebinger, G. Sharp Inequalities for f-divergences (2013). doi: 10.48550/arXiv.1302.0336. arXiv:1302.0336.

Han, Y. and Niles-Weed, J. Approximate independence of permutation mixtures. *The Annals of Statistics*, **54** (3), 1594–1621 (2026a). doi:10.1214/26-AOS2619. arXiv:2408.09341; The coefficient result used here is Theorem 4.3(2) in version 3.

Han, Y. and Niles-Weed, J. Elementary Symmetric Polynomial Inequalities for Centered Vectors and Matrices (2026b). doi:10.48550/arXiv.2607.23836. arXiv:2607.23836.

Hanselka, C. Definite Determinantal Representations of Ternary Hyperbolic Forms (2015). URL https://arxiv.org/abs/1411.1661. Version 2; Theorem 1 (Helton–Vinnikov); arXiv:1411.1661.

Hsieh, J.-T., Kane, D. M., Kothari, P. K., Li, J., Mohanty, S., and Tiegel, S. Rigorous Implications of the Low-Degree Heuristic (2026). doi:10.48550/arXiv.2601.05850. arXiv:2601.05850.

Loh, W.-L. Stein's Method and Multinomial Approximation. *The Annals of Applied Probability*, **2** (3), 536–554 (1992). doi:10.1214/aoap/1177005648.

National Institute of Standards and Technology. Digital Library of Mathematical Functions (n.d.). URL https://dlmf.nist.gov/. Sections 18.19 and 18.23; Table 18.19.1 and equation 18.23.3.

Randles, E. Local limit theorems for complex functions on $\mathbb{Z}^d$. *Journal of Mathematical Analysis and Applications*, **519**, 126832 (2023). doi:10.1016/j.jmaa.2022.126832. arXiv:2201.01319; version 2.

Riener, C. and Schabert, R. Linear slices of hyperbolic polynomials and positivity of symmetric polynomial functions (2023). doi:10.48550/arXiv.2203.08727. arXiv:2203.08727; version 2.

Roos, B. Multinomial and Krawtchouk Approximations to the Generalized Multinomial Distribution. *Theory of Probability and its Applications*, **46** (1), 103–117 (2002). doi:10.1137/S0040585X97978750. Russian original: *Teor. Veroyatnost. i Primenen.* **46**(1) (2001), 117–133.

Roos, B. Closeness of convolutions of probability measures. *Bernoulli*, **16** (1), 23–50 (2010). doi:10.3150/08-BEJ171. arXiv:0811.0622.

Roos, B. On Bobkov's approximate de Finetti representation via approximation of permanents of complex rectangular matrices (2014). doi:10.48550/arXiv.1203.2871. arXiv:1203.2871; version 2.