Moment certificates for homogeneous stable count laws

Scott Schubr
Luminary Oddities

Abstract

Let P be a probability law on nonnegative integer count vectors of total size n, with a homogeneous real-stable probability generating polynomial and positive mean vector t. We compare P with $M = \mathrm{Mult}(n, t/n)$ using the weighted covariance deficit

$$W = \frac{1}{2}\left(m - 1 - \sum_{i=1}^{m} \frac{\mathrm{Var}(A_i)}{t_i}\right).$$

We prove the sharp real elementary-symmetric-polynomial inequality $|e_k(z)|^2 \le \binom{n}{k}$ whenever $\sum_i z_i = 0$ and $\sum_i z_i^2 = n$, simultaneously in every degree and coefficient order. For a homogeneous real-stable count law this gives the covariance certificate

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

Both universal scalar and linear-TV coefficients are optimal. The argument uses classical same-degree root extrema, the centered coefficient and orthogonal tensor framework of Han and Niles-Weed, a unimodal contour estimate, exact reciprocal coefficient recurrences, and complete rational polynomial certificates. A finite boundary classification and two infinite-degree strips cover all real coefficients. We identify the equivalent binomial-orthonormal Krawtchouk bound at an integer mean and retain a Gaussian mutual-TV limit along powers of each fixed stable law. Standalone standard-library code verifies every finite certificate. The count-law results allow arbitrary category count and positive retained means, nonfactorizable stable polynomials and singular covariance deficit.

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. Let

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

be homogeneous and real stable, with integer $n \ge 1$. 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

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

We will prove $B \succeq 0$. The identity $\sum_i t_i = n$ gives the formula for W in the abstract.

Multinomial approximation for dependent random variables is classical; Loh's Stein-method treatment [7] is a central predecessor. Ehm's binomial approximation [3] already treats the bivariate stable subclass, whose generating polynomials factor into linear forms. Roos [9] 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 [5]. 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 [10]; later centered-vector and matrix inequalities appear in [6]. 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).

The contribution here is the complete real scalar coefficient-one certificate, its finite chi-square covariance bound for the full stable class, the optimal TV coefficient two, and a fixed-law Gaussian mutual-TV limit along convolution powers. Banach's theorem, the predecessor centered ESP method, hyperbolic-slice extremality, and the orthogonal-expansion architecture are credited inputs rather than new general mechanisms. Fixed degree and category count are separate parameters throughout.

The all-real two-moment estimate of Derksen, Ivanov, Lee and Viola [14, 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. [15, 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.

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

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

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

Theorem 2. *For every integer $n \ge 1$ and every normalized homogeneous real-stable count law with nonnegative coefficients,*

$$\chi^2(P\|M) \le \sum_{k=2}^{n}(2W)^k, \qquad \mathrm{TV}(P,M) \le \min\{1, 2W\}. \tag{3}$$

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 $n \geq 2$.

Lemma 3. *There are homogeneous polynomials R_k of degree k, $2 \leq k \leq n$, invariant under $x \mapsto x + a\mathbf{1}$, such that*

$$p(x) = L(x)^n + \sum_{k=2}^{n} R_k(x) L(x)^{n-k}. \tag{4}$$

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 $\lambda_1, \ldots, \lambda_n$, 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 (4).

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 $-x^\top Bx$. On the other hand the root identity gives

$$R_2(x) = -\tfrac{1}{2} \sum_j \lambda_j^2, \qquad \sum_j \lambda_j^2 = 2x^\top Bx. \tag{5}$$

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 (5) 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 $R_k(x) = T_k(x, \ldots, x)$. Our ordered-index convention means that a monomial coefficient at α is $(k!/\alpha!)T_{k,\alpha}$.

Lemma 4. *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*

$$\frac{P}{M} - 1 = \sum_{k=2}^{n} U_k, \qquad \chi^2(P\|M) = \sum_{k=2}^{n} \mathbb{E}_M U_k^2, \tag{6}$$

and

$$\mathbb{E}_M U_k^2 = \binom{n}{k}^{-1} \sum_{i_1, \ldots, i_k} \frac{T_k(e_{i_1}, \ldots, e_{i_k})^2}{q_{i_1} \cdots q_{i_k}}. \tag{7}$$

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 (7), while (4) gives (6). All operations are finite sums. □

3 Hilbert tensor transfer and a universal certificate

We use Banach's classical real Hilbert-space theorem [1, 2]: the norm of a symmetric multilinear form on arbitrary unit-vector tuples equals its norm on diagonal tuples. It is also the input in [5, Lemma 4.2].

Lemma 5. *If $|R_k(x)| \le c(x^\top B x)^{k/2}$ for all real x, then*

$$\mathbb{E}_M U_k^2 \le \frac{c^2 (nW)^k}{\binom{n}{k}}. \tag{8}$$

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 $x^\top B x$. 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 (7) 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 [5, 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 [4, Corollary 5.3], which credits Gilardoni's Proposition 2.1. We give the short proof needed here.

Lemma 6. *For probability measures $P \ll M$, write $v = \chi^2(P\|M)$ and $\delta = \mathrm{TV}(P, M)$. If $\delta > 1/2$, then $v \geq \delta/(1-\delta)$.*

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

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

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

We first establish the contour and finite-coefficient tools. The central and boundary estimates then prove the scalar theorem in every order, after which the orthogonal transfer and density positivity give the count-law bounds.

4 Coefficient bounds and finite-degree certificates

The coefficient extremum reduction used next is classical: see Riener and Schabert [8, Theorem 2.8 and Corollary 2.10], as well as the real-vector reduction in [5]. We include the needed special proof.

Lemma 7. *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. □

4.1 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 [5, Appendix B.1]; the change is to use this estimate instead of their endpoint/interior split and to retain the normalization factors.

Lemma 8. *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 9. *For every centered real vector z with $\sum_i z_i^2 = n$,*

$$|e_k(z)|^2 \leq C_0 \binom{n}{k} \quad (0 \leq k \leq n), \qquad C_0 = \frac{\pi e^{1/4}}{\sqrt{2}} < \frac{20}{7}. \tag{9}$$

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 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 8,

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

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 (9). 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 (5) and applying Lemma 5 now gives the all-order estimate

$$\mathbb{E}_M U_k^2 \le C_0 (2W)^k \le \frac{20}{7} (2W)^k. \tag{10}$$

4.2 Small-degree scalar certificates

Lemma 10. *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. 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. □

5 Exceptional profiles and fixed low orders

We normalize centered real root vectors by $\sum_i z_i^2 = n$. If

$$\frac{|e_k(z)|^2}{\binom{n}{k}} \le b_k, \tag{11}$$

then scaling (5) and applying Lemma 5 gives $\mathbb{E}_M U_k^2 \le 2^k b_k W^k$. All root extrema below use Lemma 7 one coefficient at a time; their maximizing profiles need not coincide.

5.1 Discrete exceptional profiles

Lemma 11. *If a centered, variance-normalized two-valued vector has either root value with multiplicity one or two, then* (11) *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

$$\frac{e_k^2}{\binom{n}{k}} = \binom{n}{k}\left(\frac{2}{n-2}\right)^k F(n,k)^2, \qquad F(n,k) = \frac{(k-1)((k-4)n+4)}{4(n-1)}. \tag{12}$$

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 (12) reduces the assertion to

$$16k!(n-1)^2(n-2)^k - 2^k(n)_k(k-1)^2((k-4)n+4)^2 \ge 0. \tag{13}$$

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 (13) by integer linear-factor multiplication and checks every coefficient, so this proves the inequality for all real $x \geq 0$, not just a finite list of n.

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

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

Indeed $t\log(t/(t-2))$ decreases for $t > 2$, and at twelve $6^{12} < 9 \cdot 5^{12}$. Thus the ratio in (12) 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. □

5.2 Orders zero through four

Lemma 12. *For every centered real vector with squared norm n, $|e_k|^2 \leq \binom{n}{k}$ for $0 \leq k \leq \min(4, n)$.*

Proof. Orders zero and one are immediate. At order two $e_2 = -n/2$, so its squared ratio is $n/[2(n-1)] \leq 1$ for $n \geq 2$. At order three, Lemma 7 gives a binary extremizer with $c = (s-r)/\sqrt{rs}$ and $e_3 = nc/3$. Since $rs \geq n-1$, $c^2 \leq (n-2)^2/(n-1)$, 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 $z_i^2 \leq n-1$, whence $n \leq p_4 \leq n(n-1)$ and $|e_4| \leq n(n-2)/8$. Its squared ratio is at most

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

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

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

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

Consequently, for every law under consideration,

$$\mathbb{E}_M U_k^2 \leq 2^k W^k, \qquad 7 \leq k \leq \min(12, n). \tag{14}$$

Proof. By Lemma 7, it suffices to treat two-valued profiles. Reflect and permute so their multiplicities satisfy $s \geq r$, $n = r + s$. 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

$$D_k(r, s) = (n)_k (rs)^k - k! E_k(r, s)^2 \geq 0. \tag{15}$$

For each $k = 5, \ldots, 12$, 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 (15) 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 (15) holds for the whole real quadrant $u, x \geq 0$, covering every integer $s \geq r \geq \lceil k/2 \rceil$, 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 $0 \leq u \leq 3k$, $0 \leq x \leq 2k$, 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 11. 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

$$k! r^{2r} (n)_r^2 (n-r)^k - r^k (n)_k N_{r,k}(n)^2 \geq 0. \tag{16}$$

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 5 finish the proof. □

5.4 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 14. *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 $0 < A \le 3/10$, then $I \le \Phi(A)$, where $\Phi(a) = \arccos(1-a)/\pi$, and

$$I^2 \le \frac{2A}{\pi^2}(1 + A/10)^2. \tag{17}$$

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 $I \le \int_0^1 \Phi(\ell(h))\,dh$. Although Φ is concave only on $[0,1]$ and convex on $[1,2]$, its tangent T_A at every $0 < A \le 3/10$ 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 $(\arccos(7/10) + \sqrt{17/3})/\pi > 1$: $\cos(77/100) \ge 1 - (77/100)^2/2 > 7/10$, $\sqrt{17/3} > 119/50$, 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 (17). □

Lemma 15. *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 $\gamma \geq 8$, 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 16. *Every centered real vector z with $\sum_i z_i^2 = n$ satisfies*

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

More precisely, for binary profiles with $r, s \geq 11$ and $v_0 = k(n-k)/n \geq 40$, 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 9. 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

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

where $E_a = (1/r + 1/s)/6 \leq 1/33$ and $E_b = 1/(12v_0) \leq 1/480$. Their sum is at most $57/1760 < 1/30$. The estimate $e^x \leq (1-x)^{-1}$ gives $e^{E_a} \leq 33/32$ 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 14 applies. Multiplying its bound with the two displayed estimates yields

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

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 $\min(k, n-k) \geq 80$. Then $v_0 \geq \min(k, n-k)/2 \geq 40$ and $k \geq 80$. Apply the same-n extremum reduction of Lemma 7, reflecting its binary profile so $r \leq s$. If $r \leq 10$, $k \geq 80 \geq 8r$ and Lemma 15 applies. If $r \geq 11$, 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. □

6 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.

6.1 Bottom orders thirteen through forty

Lemma 17. *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, 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

$$F_k(r, s)^2 \le \binom{n}{k}(rs)^k. \tag{18}$$

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

$$H_k(u, x) = k!(n)_k(rs)^k - E_k^2 \tag{19}$$

and checks every monomial coefficient is nonnegative. Hence $H_k \ge 0$ throughout the quadrant, which is (18) 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

$$F_k(r, s) = \sum_a \binom{r}{a}\binom{s}{k-a} s^a (-r)^{k-a}. \tag{20}$$

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

Multiplicities one and two are covered by Lemma 11. For each remaining $3 \le r < b_k$, define

$$N_{r,k}(n) = \sum_{a=0}^{r} (-1)^a \binom{r}{a} r^{r-a} (k)_a n^a (n-a)_{r-a}. \tag{21}$$

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 (18) is equivalent, after clearing positive denominators, to

$$G_{r,k}(n) = k!r^{2r}(n)_r^2(n-r)^k - r^k(n)_k N_{r,k}(n)^2 \geq 0. \tag{22}$$

For $k = 13$ the four rows $r = 3, 4, 5, 6$ use $n = 13 + y$. For $14 \leq k \leq 40$ every row $3 \leq r < b_k$ uses $n = \max(2k, 2r) + y$. All $y \geq 0$ 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 (20) at degree-plus-one distinct nonnegative integer nodes in each row. The cases $r \geq b_k$, $3 \leq r < b_k$ and $r = 1, 2$ exhaust all reflected profiles, proving the lemma. □

6.2 Bottom orders forty-one through seventy-nine

Lemma 18. *For $41 \leq k \leq 79$, $n \geq 2k$, and every centered real vector of squared norm n, $|e_k|^2 \leq \binom{n}{k}$.*

Proof. Use the same binary reduction. Multiplicities one and two are already covered. If $3 \leq r \leq 10$ and $k \geq 8r$, apply Lemma 15. The remaining pairs are exactly

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

There are 115 such pairs. For each pair, construct (21) and (22) 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 $y \geq 0$.

For $r \geq 11$ the joint assertion of Theorem 16 applies as soon as

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

The finite complement is exactly $2k \leq n < L_k$, $11 \leq r \leq \lfloor n/2 \rfloor$. 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

$$F_0 = 1, \quad F_1 = 0, \quad (h+1)F_{h+1} = -dhF_h - rs(n - h + 1)F_{h-1}. \tag{23}$$

The program checks (18) 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 19. *For every centered real vector of squared norm n and every $0 \leq k \leq \lfloor n/2 \rfloor$, $|e_k|^2 \leq \binom{n}{k}$.*

Proof. Lemmas 12 and 13 cover $k \leq 12$. Lemmas 17 and 18 cover the remaining $13 \leq k \leq 79$ in their exact bottom-half domains. If $k \geq 80$ and $n \geq 2k$, then $n - k \geq k \geq 80$, so the full-vector assertion of Theorem 16 applies. This partition covers every bottom order at the same n. □

7 Top boundary bridges

7.1 Exact reciprocal normalization

For a reflected binary profile put $d = s - r \geq 0$. 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

$$\begin{aligned} |e_{n-j}|^2 &= \left(\frac{r}{s}\right)^d \frac{T_j^2}{(rs)^j}, \\ |e_{n-j}|^2 \leq \binom{n}{j} \quad &\Longleftrightarrow \quad T_j^2 r^d \leq \binom{n}{j} (rs)^j s^d. \end{aligned} \tag{24}$$

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

$$T_0 = 1, \quad T_1 = -nd, \quad (h+1)T_{h+1} = -d(n-h)T_h - rs(n-h+1)T_{h-1}. \tag{25}$$

These identities apply only after reduction to nonzero binary roots. The original vector may have zeros; Lemma 7 preserves its original n.

Lemma 20. *For $41 \leq j \leq 79$, $n \geq 2j$, and every centered real vector of squared norm n, $|e_{n-j}|^2 \leq \binom{n}{j}$.*

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 $n \geq L_j$ and $r \geq 11$, then $j(n-j)/n \geq 40$ and the joint central theorem applies. If $r \leq 10$, then $k \geq 80 \geq 8r$ and the sparse lemma applies.

For $2j \leq n < L_j$ every $r \leq \lfloor (n-j)/8 \rfloor$ is also sparse. The program checks the exact integer inequality in (24) for every remaining

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

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 $n - j \geq j \geq 41$ permits the same-n extremum lift, completing the full-real assertion. □

Lemma 21. *For $51 \leq n \leq 599$, every centered real vector of squared norm n satisfies $|e_k|^2 \leq \binom{n}{k}$ at every order $0 \leq k \leq n$.*

Proof. For $k \le n/2$ use Corollary 19. Otherwise put $j = n - k < n/2$. At $j \ge 80$ both $k, j \ge 80$, so Theorem 16 applies; at $41 \le j \le 79$ use Lemma 20. The only remaining labels are $0 \le j \le 40$ with $n \ge 2j$.

The program checks (24) 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's range. Exact arithmetic and that lemma give the stated full-vector bounds. □

8 An exact reciprocal strip in every degree at least six hundred

Theorem 22. *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 at order $k = n - j \ge 581$ and the binary notation of (24). 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

$$\begin{aligned} A_0 &= 1, \qquad A_1 = -2\sqrt{y}, \\ A_{h+1} &= -2\sqrt{y}(1 - h\eta)A_h - h(1 - y\eta)(1 - (h-1)\eta)A_{h-1}. \end{aligned} \tag{26}$$

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 (24) is

$$\frac{|e_{n-j}|^2}{\binom{n}{j}} = \frac{u^d A_j^2}{j! D_j(y, \eta)}. \tag{27}$$

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 $j! D_j e^{2y}$.

8.1 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

$$B_{h+1} = -2y^{h \bmod 2}(1 - h\eta)B_h - h(1 - y\eta)(1 - (h-1)\eta)B_{h-1}. \tag{28}$$

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

$$P_j = N_j! \left[j! D_j(y, \eta) S_{N_j}(2y) - y^{j \bmod 2} B_j(y, \eta)^2 \right] \tag{29}$$

has nonnegative complete tensor-product Bernstein arrays on each of

$$\left[0,\frac{25}{4}\right],\quad \left[\frac{25}{4},\frac{25}{2}\right],\quad \left[\frac{25}{2},25\right],\quad [25,50],\qquad \eta\in[0,1/600]. \tag{30}$$

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

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_yl_η, is

$$\sum_{b\le i,\ c\le h} c_{bc}\binom{i}{b}\frac{l_y}{\binom{M}{b}}\binom{h}{c}\frac{l_\eta}{\binom{Z}{c}}. \tag{31}$$

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 (30). 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.

8.2 An analytic tail with translated polynomial certificates

If $r\le k/8$, Lemma 15 already proves the result at the original order $k=n-j$. Otherwise

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

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

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

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

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

In (26) every coefficient factor apart from $2\sqrt{y}$ and h lies in $[0,1]$ on the actual domain. Induction yields $|A_j|\le C_j(2\sqrt{y})$. Thus (27) is at most

$$\frac{10}{7}\left(\frac{5}{2}\right)^j e^{-2y}\frac{C_j(2\sqrt{y})^2}{j!}. \tag{32}$$

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 $1 \le j \le 19$. 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 $[t^b]P(50+t) = \sum_{h\ge b}[y^h]P(y)\binom{h}{b}50^{h-b}$. All 19 complete translated arrays pass. Hence $Q_j(y) \ge 0$ for every $y \ge 50$, and $e^{2y} \ge S_N(2y)$ makes (32) 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 completes the full-real conclusion without reciprocal operations on original zero coordinates. The proof controls finite n exactly. □

9 A global affine contour budget for top codimensions nineteen through forty

Theorem 23. *For $n \ge 600$, $19 \le j \le 40$, and every centered real vector of squared norm n, $|e_{n-j}|^2 \le \binom{n}{j}$.*

Proof. Use the binary Cauchy modulus f, entropy factor H_0 and area A in the proof of Theorem 9; put $k = n - j$, $v = k(n-k)/n$. The interval rearrangement and layer-cake argument in Lemma 14 gives, without an area restriction,

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

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 $\cos(77/100) \ge 1 - (77/100)^2/2 > 7/10$ and $(119/50)^2 < 17/3$. Therefore

$$I \le \alpha + \beta A \qquad (0 \le A \le 2). \tag{33}$$

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 $r \le k/8$, use Lemma 15. Otherwise

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

The Beta and Stirling bounds proved earlier are

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

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 (33) 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. □

10 All scalar orders and the covariance certificate

Proof of Theorem 1. 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 10. For $51 \le n \le 599$, use Lemma 21. For $n \ge 600$, every $k \le \lfloor n/2 \rfloor$ follows from Corollary 19. 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 22; at $19 \le j \le 40$, use Theorem 23. The overlap is harmless. At $41 \le j \le 79$, Lemma 20 applies since $n \ge 2j$. At $j \ge 80$, both $j, k \ge 80$ and Theorem 16 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 (2) 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. □

Proof of Theorem 2. For $n \ge 2$, scaling Theorem 1 with (5) gives

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

The quotient tensor transfer therefore yields $\mathbb{E}_M U_k^2 \le (2W)^k$. Orthogonality proves (3)'s chi-square bound for every $W \ge 0$; 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 $v := \chi^2(P\|M) \le \sum_{k=2}^n x^k \le x^2/(1 - x)$. If $x \le 1/2$, Cauchy gives $\delta \le \sqrt{v}/2 \le x$. If $x \ge 1/2$ and $\delta \le 1/2$ the result is immediate; otherwise $\delta > x$ contradicts Lemma 6, since $v \ge \delta/(1 - \delta) > x/(1 - x) \ge v$. At $W = 0$ use rigidity and at $W \ge 1/2$ use $\mathrm{TV} \le 1$. 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 [9, 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. □

10.1 The equivalent own-mean Krawtchouk bound

In the conventional normalization [16, 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 is equivalent, through the same-n two-value reduction and elementary low orders, to

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

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 $|\widehat{K}_k(x)| \le 1/\sqrt{w_p(x)}$, which does not imply the own-mean bound one. The generating function and norm are classical inputs, not new constructions.

10.2 Reproducibility

Run `verify_scalar_one.py` with its shipped helpers `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 contour, real-root, tensor and power-limit arguments are analytic manuscript proofs rather than numerical limit checks.

11 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 24. *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 $M_N = M^{*N} = \mathrm{Mult}(Nn, t/n)$,*

$$\lim_{N\to\infty} \mathrm{TV}(P_N, M_N) = \mathrm{TV}(\mathcal{N}_H(0, C), \mathcal{N}_H(0, S)). \tag{34}$$

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 $N \ge N_0(P)$. At $W = 0$ all powers are exactly multinomial; at $W \ge 1/2$ the inequality holds at every power.

The support input is classical: Brändén's jump-system theorem [11, Theorem 3.2] extends earlier homogeneous support results of Choe, Oxley, Sokal and Wagner [12]. The finite-support Fourier local limit used below is also classical; see the background and lattice-period discussion in Randles [13]. We give the needed argument, including the l1 passage and singular case, to expose the exact hypotheses.

Proof of Theorem 24. **11.1 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 $m \ge 2$, 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 $\mu = (t_1, \dots, t_d)$.

The matrix S' is positive definite and

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

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

$$\operatorname{Tr}((S')^{-1}B') = W. \tag{35}$$

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.

11.2 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

$$\sup_{k\in\mathbb{Z}^d}\left|N^{d/2}f_N(k) - g_\Sigma\left(\frac{k - N\mu}{\sqrt{N}}\right)\right| \longrightarrow 0. \tag{36}$$

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 (36) with the usual Gaussian Fourier normalization.

It follows that

$$\sum_{k\in\mathbb{Z}^d} \left| f_N(k) - N^{-d/2} g_\Sigma\left(\frac{k-N\mu}{\sqrt{N}}\right)\right| \longrightarrow 0. \tag{37}$$

For fixed R, the box $\|k-N\mu\|_\infty \le R\sqrt{N}$ has $O_R(N^{d/2})$ sites, so (36) 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 (37).

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 (34). The coordinate isomorphism onto H preserves TV, proving the claimed coordinate-independent limit when C is nonsingular there.

11.3 Singular covariances and the strict Gaussian gap

If C is singular on H, choose nonzero $v\in H$ with $v^\top Cv = 0$. 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 $v^\top Sv > 0$. For every $\eta>0$, Portmanteau on the closed interval $[-\eta,\eta]$ gives

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

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 (34) also in the singular case.

Finally suppose $0<W<1/2$, so the projected covariances are positive definite. By (35), $D\succeq 0$ 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 (34) 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 2 already bounds every power, because $W_N = W$. The limit theorem additionally identifies the limiting value and its strict Gaussian gap; it provides no uniform threshold or rate. 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.

12 Scope and remaining questions

The optimal real scalar coefficient is one, and the optimal universal linear TV coefficient is two. The proof establishes every real scalar order by an exhaustive finite/infinite domain partition rather than extrapolating a table or assuming a Hermite limit proves finite degrees. The fixed-law Gaussian limit is an additional asymptotic description and is not used to deduce the first-power theorem.

The mean vector and covariance are treated as known inputs. We prove no sample-complexity result for estimating W, practical testing guarantee, sampler construction or computational speedup. Stability is a substantive hypothesis; matching a few moments without it does not certify the full law. 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.

Computational and AI assistance

The computer-assisted proof is reproducible from the three supplied Python files using only the standard library. The main verifier reconstructs all exact arrays and exhaustive finite domains described above; no private module, external numerical package or random sampling is required. The support-exchange, Fourier local-limit, singular-covariance and Gaussian arguments are analytic proofs. Generative AI systems assisted with mathematical exploration, literature discovery, proof drafting, exact-code development and independent adversarial proof checks. These are internal AI-assisted checks. The author retains responsibility for the manuscript. Classical inputs and the close Han–Niles-Weed methodology are explicitly attributed; the proofs and complete computational domains are provided for independent assessment.

References

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

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

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

[4] A. Guntuboyina, S. Saha and G. Schiebinger, Sharp inequalities for f-divergences (2013). arXiv:1302.0336.

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

[6] Y. Han and J. Niles-Weed, Elementary symmetric polynomial inequalities for centered vectors and matrices (2026). arXiv:2607.23836.

[7] W.-L. Loh, Stein's method and multinomial approximation, *The Annals of Applied Probability* **2**(3) (1992), 536–554. doi:10.1214/aoap/1177005648.

[8] C. Riener and R. Schabert, Linear slices of hyperbolic polynomials and positivity of symmetric polynomial functions (2023). arXiv:2203.08727, version 2.

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

[10] B. Roos, On Bobkov's approximate de Finetti representation via approximation of permanents of complex rectangular matrices (2014). arXiv:1203.2871, version 2.

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

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

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

[14] H. Derksen, P. Ivanov, C. H. Lee and E. Viola, Pseudorandomness, Symmetry, Smoothing: I, in *39th Computational Complexity Conference*, LIPIcs **300** (2024), 18:1–18:27. doi:10.4230/LIPIcs.CCC.2024.18.

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

[16] National Institute of Standards and Technology, *Digital Library of Mathematical Functions*, §18.19 and §18.23, Table 18.19.1 and equation 18.23.3.