Equicorrelation and the permanental cofactor ratio: local rigidity, Walsh families and exact small orders

Scott Schubr
Luminary Oddities

November 2026

Abstract

For a complex correlation matrix B let $C(B)$ be its raw permanental cofactor matrix and $\Lambda(B) = \lambda_{\max}(C(B))/\operatorname{per} B$. The optimal constant $\gamma_n = \sup \Lambda$ over correlation matrices of order n is known to satisfy $(n+1)^2/(4n) < \gamma_n \le (n^2+2n+3)/(4n)$, and at order three every extremizer is a diagonal-unitary conjugate of a positive equicorrelation matrix. We give strong evidence that equicorrelation is the extremal structure in every order. First, we optimize exactly over the real equicorrelation family; the optimum, described by a conditioned Poisson law, exceeds $(n+1)^2/(4n)$ at every order (for example 1.6452 instead of 1.5625 at $n = 4$) and coincides with the exact value of γ_3. Second, the optimal equicorrelation matrix is a strict local maximizer of Λ on the full complex correlation space modulo diagonal-unitary conjugation, for all large orders, with Hessian kernel exactly the phase-orbit tangent and four transverse curvatures asymptotic to $-n/8$, $-n/16$, $-1/8$, $-1/8$. Third, among all one-factor correlation matrices $\operatorname{diag}(1-|z_i|^2)+zz^*$, with arbitrary loadings and possibly singular residuals, the optimal equicorrelation matrix is the unique global maximizer up to diagonal-unitary conjugation for all large orders. Fourth, for all unit-diagonal matrices whose off-diagonal part is a bounded Gram correction of fixed rank r, the maximum of Λ is $n/4 + 1/2 + 1/(4n) + O(n^{-3})$ and exceeds the exact equicorrelation optimum by at most $O(n^{-3})$; in particular the coefficient of $1/n$ is $1/4$ throughout this class, and every near-maximizing sequence converges, after diagonal gauges, to equicorrelation. The same coefficient $1/4$, with excess $O(n^{-3/2})$ over the equicorrelation optimum, holds with no rank restriction for every correlation matrix with the spectral floor $A \succeq (1-B^2/n)I$. Fifth, for correlation matrices built from a fixed finite set of Walsh factors with bounded strengths, the scalar equicorrelation maximizer and its character gauges are the only global maximizers for all large sizes, so the coefficient of $1/n$ in this class is $1/4$; at strength one a coupled quartic law describes the transition to growing cofactors, and nearly equal supercritical strengths produce a torus phase distribution with an explicit Bessel ratio. Finally, unrestricted numerical optimization over complex correlation matrices of orders three, four and five attains exactly the optimal equicorrelation value, and the order-six value $2.1044 > 2$ is an explicit equicorrelation witness against the raw cofactor bound with constant two.

Overview

Sections 1–5 prove the exact equicorrelation optimum (Proposition 2) and local rigidity (Theorem 1). Section 6 compares the optimum with the exact order-three constant and with unrestricted numerical optimization at small orders. Section 7 proves global maximality among one-factor correlation matrices (Theorem 3), and Sections 11–15 determine the asymptotics for bounded fixed-rank corrections (Theorem 4), and Sections 16–19 remove the rank restriction (Theorem 6). Sections 20–24 treat Walsh families (Theorems 7–10).

1 Introduction and main result

For a matrix $B = (B_{ij})$ of order n, write $\operatorname{per} B = \sum_\sigma \prod_i B_{i,\sigma(i)}$ and $C(B)_{ij} = \operatorname{per} B(i,j)$, where $B(i,j)$ deletes row i and column j. A correlation matrix is Hermitian positive semidefinite with diagonal entries one. On this space $C(B)$ is Hermitian and $\operatorname{per} B \geq 1$ by Lieb's permanent inequality [4]. Define

$$\Lambda(B) = \frac{\lambda_{\max}(C(B))}{\operatorname{per} B}.$$

The raw cofactor convention matters for phase conjugation and differentiation. It also differs from the entry-weighted cofactor matrix that arises in a conjecture of Bapat and Sunder [6].

The all-order bounds in [7] put the supremum of Λ between $(n+1)^2/(4n)$ and $(n^2+2n+3)/(4n)$. The lower construction is a positive equicorrelation matrix with off-diagonal entry $2/(n+1)$. Its exact scalar quotient and the exact order-three classification in that work motivate the study of perturbations in every complex correlation direction. Lavi [3, Section 4] bounds the same normalized raw cofactor spectrum at order four. Marcus and Merris [5] prove $C(B)^{\mathsf{T}} \preceq n(\operatorname{per} B)B^{-1}$ for positive definite B. The scalar permanent derivative, expressed through its cofactors, also enters the fixed-spectrum extremal analysis of Grone, Johnson, Sá and Wolkowicz [2]. Bhatia, Grover and Jain [1, Section 4] give higher derivatives of the scalar permanent. Here the simple-eigenvalue derivative and the combinatorics of two directed edges give the second variation of the normalized raw spectrum.

Let I be the identity, $J = \mathbf{1}\mathbf{1}^*$, and $B(b) = (1-b)I + bJ$ for $-1/(n-1) \leq b \leq 1$. Choose b_n to maximize $\Lambda(B(b))$ on this interval, and put $B_n = B(b_n)$. The space of affine perturbations consists of Hermitian matrices X with diagonal zero. We use the real Frobenius inner product and write $\|X\|_F^2 = \sum_{i,j} |X_{ij}|^2$.

Theorem 1. *For all sufficiently large n, b_n is interior and positive. The matrix B_n is a strict local maximizer of Λ modulo diagonal-unitary conjugation. The Hessian kernel is precisely the space of imaginary gradients $X_{ij} = i(\theta_i - \theta_j)$, where θ_i are real. On its Frobenius orthogonal complement the Hessian has the following four orthogonal sectors. On each sector its quadratic form equals the indicated coefficient times $\|X\|_F^2$.*

constant real off-diagonals	$-n/8 + o(n)$
real stars $X_{ij} = (v_i + v_j)/2$ $(i \neq j)$, $\sum_i v_i = 0$	$-n/16 + o(n)$
real symmetric, zero row sums	$-1/8 + o(1)$
purely imaginary Hermitian, zero row sums	$-1/8 + o(1)$.

Proposition 2. *There is an absolute $c > 0$ such that every maximizing b_n satisfies*

$$t_n = \frac{1-b_n}{b_n} = \frac{n-1}{2} + O(ne^{-cn/2}), \qquad \Lambda(B_n) = \frac{(n+1)^2}{4n} + O(ne^{-cn})$$

for all sufficiently large n.

The proposition optimizes the previously known scalar formula over the full real family, and the theorem treats perturbations in the full complex correlation space near the optimizing orbit. The exact optimum of the scalar formula (Section 2) improves the known lower bound at every finite order; Table 1 lists representative values.

For $4 \leq n \leq 8$ we also evaluated the second difference of Λ at B_n along 60 random complex directions per order, after removing the phase-orbit component. Every value was strictly negative (the largest was -0.27, at $n = 8$), in agreement with Theorem 1 already at these small orders.

n	t_n	$b_n = 1/(1+t_n)$	$\Lambda(B_n)$	$(n+1)^2/(4n)$
3	1.79549	0.35772	1.42285	1.33333
4	2.19358	0.31313	1.64523	1.56250
5	2.60559	0.27735	1.87277	1.80000
6	3.02927	0.24818	2.10437	2.04167
8	3.90485	0.20388	2.57665	2.53125
12	5.73673	0.14844	3.54389	3.52083
20	9.57853	0.09453	5.51829	5.51250
50	24.50076	0.03921	13.00503	13.00500

Table 1: The optimal real equicorrelation value $\Lambda(B_n)$, computed from the exact formula (1) at 40 digits, compared with the earlier lower bound. The values agree with direct Ryser evaluation of the cofactor spectrum for $n \le 8$.

2 Optimizing the real equicorrelation family

Write $a = 1 - b$, and for $b > 0$ set $t = a/b$. Expansion in the diagonal part gives $\operatorname{per}(aI + bJ) = n! b^n \sum_{j=0}^{n} t^j/j!$. The raw cofactor matrix has positive off-diagonal entries, so its largest eigenvalue is its common row sum. This row sum is $n^{-1}\partial_b \operatorname{per}(aI + bJ)$ with a fixed, since the derivative adds every raw cofactor. Accordingly, for N_t Poisson with mean t, the quotient is

$$\Lambda(B(b)) = (1+t)\left(1 - \frac{t}{n} + \frac{t}{n}\frac{\mathbb{P}(N_t = n)}{\mathbb{P}(N_t \le n)}\right) = \frac{1+t}{n}\mathbb{E}[n - N_t \mid N_t \le n]. \tag{1}$$

This is the scalar formula underlying the lower construction of [7].

For $0 \le t \le 3n/4$, Chernoff's inequality gives $\mathbb{P}(N_t \ge n) \le e^{-cn}$ with $c = \log(4/3) - 1/4 > 0$. Thus the Poisson correction in (1) is $O(ne^{-cn})$ uniformly on this range. The quadratic part is maximized at $t_* = (n-1)/2$, with maximum $(n+1)^2/(4n)$ and exact defect $(t - t_*)^2/n$.

For $3n/4 \le t \le n$, the quadratic is at most $3n/16 + O(1)$, and $\mathbb{P}(N_t = n) \le \mathbb{P}(N_n = n) = O(n^{-1/2})$. Also $\mathbb{P}(N_t \le n) \ge \mathbb{P}(N_n \le n) \ge c_0 > 0$: Stirling's formula gives an atom of order $n^{-1/2}$ at n for N_n, and the $\lfloor\sqrt{n}/2\rfloor$ atoms immediately below it are each comparable to it by the exact ratios of Poisson masses. The correction is consequently $O(\sqrt{n})$. For $n \le t \le 3n/2$ the conditional law $N_t \mid N_t \le n$ is stochastically increasing by monotone likelihood ratio. At $t = n$ its expected deficiency is $O(\sqrt{n})$ by Cauchy–Schwarz and the same conditioning bound. For $t \ge 3n/2$ successive deficiency masses have ratio at most $q = n/t$. If $p_j = \mathbb{P}(n - N_t = j \mid N_t \le n)$, then $\sum_j p_{j+k} \le q^k$, so the conditional mean deficiency is at most $n/(t-n)$. The quotient on this last range is $O(1)$. These bounds place the positive-branch maximum inside $0 \le t < 3n/4$ for large n. Comparison to t_* there proves the two exponential estimates.

For the negative branch, set $b = -r$ and $a = 1 + r$ with $0 \le r \le 1/(n-1)$. The identity $\operatorname{per}(aI + bJ) = \int_0^\infty (a + bu)^n e^{-u}\,du$ follows by expanding the polynomial and hence is valid for negative b as well. The cofactor diagonal d and row sum w have the integral expressions

$$d = \int_0^\infty (a - ru)^{n-1} e^{-u}\,du, \qquad w = \int_0^\infty u(a - ru)^{n-1} e^{-u}\,du.$$

The other eigenvalue is $(nd - w)/(n-1)$. On $0 \le u \le a/r$, one uses $a^{n-1} \le e$. Translating the two tails for $r > 0$ gives absolute bounds $e^{-a/r} r^{n-1}(n-1)!$ and that quantity times $a/r + n$, respectively. Since $r^{n-1}(n-1)! \le 1$ and $a/r \ge n$, both eigenvalues are uniformly bounded. The case $r = 0$ is the identity, and $\operatorname{per} B(b) \ge 1$ on this branch. Thus its quotient is $O(1)$, proving Proposition 2.

3 Spectral variation and directed edges

For the remaining proof write $B = aI + bJ$, $P = \operatorname{per} B$, $Q = C(B)/P$, and $e = n^{-1/2}\mathbf{1}$. We evaluate at $b = b_n, a = 1 - b_n$ after taking partial derivatives at fixed a where specified. The top eigenvalue of Q is simple and its eigenvector is e. Denote the gap to the other, common eigenvalue by $\Delta > 0$. Simplicity gives smoothness near B_n. Permutation invariance and invariance under complex conjugation imply that the first variation in a zero-diagonal Hermitian direction is a scalar multiple of the sum of the real off-diagonal entries. Since b_n is an interior branch maximizer, this scalar is zero.

For a Hermitian zero-diagonal X with total sum zero, write $P'' = D_X^2 \operatorname{per}(B)$. The first variation $D_X \operatorname{per}(B)$ is zero by the same invariance of the permanent. The identity $e^* C(B) e = n^{-1}\partial_b P$ holds also along $B + sX$ if a, X are held fixed. The second variation of the simple eigenvalue is therefore

$$D_X^2 \Lambda = \frac{1}{n}\partial_b (P''/P) + \frac{2}{\Delta}\|(I - ee^*) D_X Q\, e\|^2. \tag{2}$$

Write $r_i = \sum_j X_{ij}$ for its row sums. Equivariance makes $D_X Q e$ zero when $r = 0$. Indeed, for a real symmetric input the coefficient of each unordered off-diagonal edge at a given output coordinate depends only on incidence with that coordinate. When the total sum is zero, the image is a scalar multiple of r. For a purely imaginary Hermitian input, exchanging the ends of an edge disjoint from the output coordinate fixes the coordinate and negates the edge. Such edges contribute zero. Every incident edge contributes a constant times the entry oriented from the output coordinate. These statements also identify the mixing term in each pure sector of total sum zero as a scalar multiple of $\|r\|^2$.

Expand the permanent by the diagonal part:

$$\operatorname{per}(aI + bJ + sX) = \sum_{S \subseteq [n]} a^{n-|S|} \operatorname{per}((bJ + sX)[S]).$$

A second derivative chooses two ordered directed X-edges with distinct rows and distinct columns. When their union contains $m = 2, 3, 4$ vertices, the coefficient of that ordered product in P''/P is

$$D_m = \frac{\mathbb{E}[(K-2)_{m-2}\, \mathbf{1}_{\{K \geq 2\}}]}{b^2 (n)_m}, \qquad K = n - (N_t \mid N_t \leq n). \tag{3}$$

Here $(x)_j$ is a falling factorial. A principal set of size k in the expansion has aggregate weight proportional to $\binom{n}{k} a^{n-k} k! b^k$, so its size has law K. For $k \geq m$ the chance it contains the designated vertices is $(k)_m/(n)_m$, and the two specified edges leave $(k-2)!$ permutations instead of $k!$. The coefficient also contains b^{-2}. Events with $2 \leq k < m$ give zero through the falling factorial in (3).

Write $F_2 = \|X\|_F^2$ and $s_0 = \sum_i r_i$. The sums of ordered edge products of each type are

$$\begin{aligned} &\Sigma_2 = F_2, && \Sigma_3 = 2(\|r\|^2 - F_2), \\ &\Sigma_4 = s_0^2 - 2\operatorname{Re}\sum_i r_i^2 + \sum_{i,j} X_{ij}^2 - \Sigma_2 - \Sigma_3. \end{aligned} \tag{4}$$

The two-vertex pairs are precisely the reciprocal edges. The three-vertex pairs have one common endpoint, outgoing for one edge and incoming for the other. The oriented sum for each order is $\|r\|^2$ before removing reciprocal edges. For the total of all admissible pairs, subtract common-source and common-target pairs from s_0^2, then add identical edges once. This yields (4), and $P''/P = \sum_{m=2}^{4} D_m \Sigma_m$. At zero row sums the relative coefficients Σ_m/F_2 are $(1, -2, 2)$ for real X and $(1, -2, 0)$ for imaginary X.

At $t = t_n$, the conditioning in (3) and the truncation at $K \geq 2$ alter the polynomial moments by exponentially small errors times powers of n. This follows from Poisson exponential tail estimates for each fixed polynomial degree. The conclusion holds after one t-derivative by the identity $\partial_t \mathbb{E}H(N_t) = \mathbb{E}[H(N_t + 1) - H(N_t)]$, including for the truncated polynomials and the conditioning probability. The three unconditioned polynomial numerators are

$$1, \qquad n - t - 2, \qquad (n - t - 2)(n - t - 3) + t.$$

Since $\partial_b t = -t/b$ at fixed a, differentiation gives

$$\frac{\partial_b D_2}{n} \longrightarrow -\frac{1}{4}, \qquad \frac{\partial_b D_3}{n} \longrightarrow -\frac{1}{16}, \qquad \frac{\partial_b D_4}{n} \longrightarrow 0.$$

For D_4, the leading terms from differentiating the numerator and b^{-2} cancel, using $t_n \sim n/2$, $b_n \sim 2/n$. When $r = 0$, the mixing term in (2) vanishes. The two triples in (4) then give the two $-1/8 + o(1)$ coefficients.

4 Real star directions

Let v be a real vector of sum zero. Set $w_i = b_n + \eta v_i$ and

$$B(\eta) = \operatorname{diag}(1 - w_i) + xx^*, \qquad x_i = \sqrt{w_i}.$$

For η near zero this is a correlation matrix. Its tangent is the real star $X_{ij} = (v_i + v_j)/2$ for $i \neq j$, with $\|X\|_F^2 = (n - 2)\|v\|^2/2$. Stationarity at B_n means the second derivative along this family equals the affine Hessian on X.

For $T \geq 0$ put $f_i = 1 + (T - 1)w_i$, and use the probability density proportional to $e^{-T} \prod_i f_i$. The normalized raw cofactor matrix is the expectation of

$$R(T) = \operatorname{diag}(h(w_i)) + Tgg^*, \qquad h(w) = \frac{1 - w}{[1 + (T - 1)w]^2}, \qquad g_i = g(w_i) = \frac{\sqrt{w_i}}{1 + (T - 1)w_i}. \quad (5)$$

To see this, expand the permanent along the diagonal part and represent the factorials of the rank-one principal submatrices by gamma integrals. A diagonal cofactor omits one f_i. An off-diagonal cofactor has an unmatched row and column, producing the factor $\sqrt{w_i w_j}T$ and omitting f_i, f_j. The diagonal of (5) equals $1/f_i$, as required.

At $\eta = 0$ write $b = b_n$, $a = 1 - b$, and $f = a + bT$. Primes on h, g denote w-derivatives evaluated at b. The logarithm of the unnormalized density has first derivative zero pointwise and second derivative $-\|v\|^2(T - 1)^2/f^2$. The row sum of $R(T)$ at this point is $\lambda_T = h + nTg^2$, and $\Delta = n\mathbb{E}(Tg^2)$. Differentiating (5) shows that the first derivative of $R(T)e$ is $(h' + nTgg')v/\sqrt{n}$. The projection of the second derivative on e gives $\|v\|^2(h''/n + 2Tgg'')$. Differentiating the expectation and applying the simple-eigenvalue formula thus give

$$\frac{D_X^2 \Lambda}{\|v\|^2} = \mathbb{E}[h''/n + 2Tgg''] - \operatorname{Cov}\big(\lambda_T, (T - 1)^2/f^2\big) + \frac{2}{n\Delta}\big(\mathbb{E}[h' + nTgg']\big)^2. \quad (6)$$

The random variable $T + t_n$ has the $\Gamma(n+1, 1)$ law conditioned to be at least t_n, since $f = b(T + t_n)$. Proposition 2 and gamma moments imply $T/n \to 1/2$ in probability and $\mathbb{E}|T - t_n|^k = O_k(n^{k/2})$ for $k = 2, 4$. Also

$$\frac{1}{4a} - Tg^2 = \frac{(a - bT)^2}{4a(a + bT)^2}.$$

This identity and the fourth moment give $\operatorname{Var}\lambda_T = O(1)$, while $\Delta \asymp n$. The function $(T-1)^2/f^2$ has derivative $2(T-1)/f^3$, uniformly $O(n)$, giving variance $O(n^3)$ using the second moment about t_n. The covariance in (6) is thus $O(n^{3/2})$. Direct differentiation gives $h''/n = O(n)$ and $h' = O(n)$ uniformly. Moreover

$$g' = \frac{1-(T-1)b}{2\sqrt{b}f^2}.$$

The numerator has second moment $O(n^{-1})$, and $T/f^3 = O(n)$, so $\mathbb{E}|Tgg'| = O(\sqrt{n})$. The last term in (6) is $O(n)$.

The remaining contribution is determined by

$$gg'' = -\frac{1}{4bf^2} - \frac{T-1}{f^3} + \frac{2b(T-1)^2}{f^4}.$$

Set $q = (T-1)b$; then $q \to 1$ in probability. At $q = 1$ the right side equals $-1/(16b)$. The rational expression for $2Tgg''/n^2$ is uniformly bounded for $q \geq -b$, so $T/n \to 1/2$ and $nb \to 2$ give convergence in mean to $-1/32$. Equation (6) is therefore $-n^2/32 + o(n^2)$. Dividing by $(n-2)/2$ in the norm identity yields the $-n/16 + o(n)$ coefficient.

5 The constant direction and the phase orbit

On a fixed relative neighborhood of $t_n \sim n/2$, (1) is $(1+t)(1-t/n)$ plus an exponentially small correction times a polynomial in n, also after up to two t-derivatives. In fact the atom $p_n(t) = \mathbb{P}(N_t = n)$ satisfies $p_n' = (n/t - 1)p_n$, and the conditioning probability has derivative $-p_n(t)$. The scalar derivative at b_n is zero, and $dt/db = -b^{-2}$ along the correlation branch. Therefore its second derivative in b is $-2/(nb_n^4) + o(n^3)$. The constant off-diagonal direction with entries one has $\|X\|_F^2 = n(n-1)$, giving $-n/8 + o(n)$.

For $U = \operatorname{diag}(u_i)$ diagonal unitary, the retained row and column factors give

$$C(UBU^*)_{ij} = \overline{u}_i u_j C(B)_{ij}.$$

The permanent is unchanged and $C(UBU^*) = U^* C(B) U$. The orbit at B_n has tangent the imaginary gradients. Differentiating the invariance at the stationary point shows that the Hessian vanishes when either input is tangent to this orbit. The imaginary gradients are Frobenius orthogonal to the purely imaginary Hermitian matrices with zero row sums. The constant real direction has zero mixed terms with its orthogonal complement by permutation averaging. Complex conjugation separates the real and imaginary sectors. In the total-sum-zero real sector, (2) and (4) express the Hessian as a linear combination of Frobenius and row-sum squares. These forms have zero mixed terms between real stars and real zero-row-sum directions. These subspaces exhaust the tangent space, giving the claimed Hessian and its kernel.

An explicit local slice to the phase orbit is obtained by making the nonzero first-row off-diagonal entries real and positive by diagonal-unitary conjugation. Its tangent intersects the imaginary gradients only at zero. Restricting the Hessian to the slice therefore gives a negative definite form. The matrix B_n is positive definite, so the slice lies locally in the interior of the correlation domain. Taylor's theorem on the slice proves strict local maximality modulo the phase orbit and completes the proof of Theorem 1.

6 Exact and numerical maximality at small orders

The order-three constant is known exactly: by [7, order-three theorem], every maximizer of Λ among 3×3 complex correlation matrices is a diagonal-unitary conjugate of a positive equicorrelation matrix, so γ_3 is the maximum of formula (1) at $n = 3$,

$$\gamma_3 = \Lambda(B_3) = 1.4228496379\ldots, \qquad t_3 = 1.7954900\ldots, \quad b_3 = 0.3577190\ldots,$$

in agreement with Table 1. We also maximized Λ numerically over all complex correlation matrices $B = D^{-1/2}XX^*D^{-1/2}$ of orders four and five, using random complex X, Ryser permanents for every cofactor, and Nelder–Mead ascent from 15 (order four) and 10 (order five) independent starting points. The largest values found were

$$1.645226619865\ (n = 4), \qquad 1.872765356\ (n = 5),$$

which agree with the optimal equicorrelation values 1.645226619865 and 1.872766516 of Table 1 to the accuracy of the ascent. No start produced a larger value. At order six the optimal equicorrelation matrix has $\Lambda(B_6) = 2.10437\ldots > 2$; it is therefore an explicit equicorrelation witness that the raw cofactor bound $C(B) \preceq 2(\operatorname{per} B)I$ fails at order six, a failure established in [7] by a different exact obstruction.

These computations, the order-three theorem, the local rigidity of Theorem 1, the global maximization among one-factor matrices in Theorem 3 and the exact Walsh-family maximization of Theorem 8 all identify the optimal equicorrelation matrix as the extremizer of the cofactor ratio.

7 One-factor correlation matrices

Let $C(A)_{ij}$ be the permanent of the submatrix obtained by deleting row i and column j. We study $\Lambda(A) = \lambda_{\max}(C(A))/\operatorname{per} A$. This is the raw cofactor ratio. The adjoint in [5] uses the transpose of this cofactor convention. Marcus and Merris give an inverse-matrix upper bound for the cofactor adjoint. Lavi [3, Section 4] bounds the raw cofactor ratio in small orders. Schubr [7] obtains all-order bounds whose first discrepancy is the coefficient of $1/n$, with a scalar equicorrelation matrix giving the lower bound. The unrestricted extremal coefficient remains open.

Feinsilver and McSorley [14, Theorem 5.3 and Lemma 5.4] give exponential moment formulas for equicorrelation subpermanents. The scalar evaluation here is consistent with those formulas. For a diagonal plus a positive rank-one factor, the positive residual case, after removing loading phases, has an inverse with nonpositive off-diagonal entries by the rank-one inverse formula. Sra's inequality [15] thus applies to its principal permanents. The present objective also uses all nonprincipal cofactors and their joint top eigenvalue. Our main step is a global concentration argument for this eigenvalue over arbitrary independent loadings. The proof includes the scalar calculation and uses no separate local Hessian theorem.

Theorem 3. *Let $A = \operatorname{diag}(1 - |z_i|^2) + zz^*$, where $z \in \mathbb{C}^n$ and $0 \le |z_i| \le 1$. Write $\Lambda(A) = \lambda_{\max}(C(A))/\operatorname{per} A$ for the raw permanental cofactor ratio. For all sufficiently large n, this class has a unique maximizing matrix up to diagonal-unitary conjugation. It is the positive equicorrelation matrix of scalar strength $b_n = 2/(n+1) + O(e^{-cn})$. The maximum is $(n+1)^2/(4n) + O(ne^{-cn})$ for an absolute $c > 0$. The class includes every diagonal-plus-one-positive-factor correlation matrix, without a small-loading assumption on any entry.*

8 Gaussian formula and the global new coordinate

Gauge the loadings to be nonnegative, put $r_i = |z_i|^2$, $a_i = 1 - r_i$, $d_i(T) = a_i + r_i T$, and $x_i(T) = r_i T/d_i(T)$ for $T > 0$. Expectations are with respect to the density proportional to $e^{-T}\prod_i d_i(T)$. Let G_i, Z_0 be independent standard circular complex Gaussians. The variables $\sqrt{a_i}G_i + \sqrt{r_i}Z_0$ have covariance A. By the complex Gaussian moment identity [8], the expectation of the product of squared moduli equals the permanent. For an off-diagonal cofactor the two omitted variables leave one conditional mean and its conjugate, while each remaining index contributes d_i. For a diagonal cofactor the factor d_i is simply omitted. Dividing by the scalar permanent and conditioning on $T = |Z_0|^2$ gives

$$M = C(A)/\operatorname{per} A = \mathbb{E}[\operatorname{diag}(a_i/d_i^2) + vv^{\mathsf{T}}], \qquad v_i = \sqrt{r_i T}/d_i.$$

All identities are for $T > 0$. The formulas below prove integrability even at singular residuals. They follow from $1/d_i = 1 - x_i + x_i/T$:

$$v_i^2 = x_i(1 - x_i) + x_i^2/T, \qquad a_i/d_i^2 = (1 - x_i)^2 + x_i(1 - x_i)/T.$$

Introduce an integer K by expanding $\prod_i d_i(T) = \sum_{k=0}^n c_k T^k$ with $c_k \geq 0$. In the joint law $K = k$ has probability proportional to $c_k k!$, and $T \mid K = k$ is Gamma of shape $k + 1$, rate one. Given T, the generating function of K is $\prod_i[(1 - x_i) + x_i u]$, so it is a sum of conditionally independent Bernoulli variables. In particular

$$\mathbb{E}\sum_i x_i/T = \mathbb{E}(K/T) = \Pr(K > 0) \leq 1.$$

Therefore the expected diagonal operator norm is at most 2, the total rank-component trace is at most $n/4 + 1$, and

$$\mathbb{E}\sum_i (x_i - 1/2)^2 = O(1) \quad \text{if} \quad \lambda_{\max}(M) \geq n/4 + 1/2 - o(1).$$

For the last implication use $\operatorname{tr}\mathbb{E}[vv^{\mathsf{T}}] = n/4 - \mathbb{E}\sum(x_i - 1/2)^2 + \mathbb{E}\sum x_i^2/T$. Every sequence of maximizers meets the hypothesis by the scalar test in the last section.

By conditional Bernoulli variance and Cauchy–Schwarz, $\mathbb{E}(K - n/2)^2 \leq n/4 + n\mathbb{E}\sum_i (x_i - 1/2)^2 = O(n)$. The Gamma variance and mean also give $\mathbb{E}(T - n/2)^2 = O(n)$. Write $T_0 = n/2$ and use the bijective coordinates $p_i = x_i(T_0) \in [0, 1]$, even at $r_i = 0, 1$. In detail

$$r_i = \frac{p_i}{T_0(1 - p_i) + p_i}, \qquad x(p_i, t) = \frac{p_i t}{1 + p_i(t - 1)}, \qquad t = T/T_0.$$

Choose one T with $|T - T_0| \leq C\sqrt{n}$ and $\sum_i (x_i(T) - 1/2)^2 \leq C$ by the two expectation bounds and a union bound. The derivative $\partial_t x = p(1-p)/[1 + p(t-1)]^2$ is uniformly bounded on $1/2 \leq t \leq 3/2$ for every $p \in [0, 1]$. Thus $\|p - \frac{1}{2}\mathbf{1}\|_2 = O(1)$ on any near-maximizing sequence, uniformly including its outlier coordinates.

9 The two-term law with strong loading outliers

The next calculation applies to every sequence $\|p - \frac{1}{2}\mathbf{1}\|_2 = O(1)$, including sequences with some p_i equal to 0 or 1. Put $b = n^{-1}\sum p_i$, $\delta_i = p_i - b$, $u = \sqrt{n}(b - 1/2)$. Then $u, \|\delta\|_2 = O(1)$, and we claim

$$\lambda_{\max}(M) = \frac{n}{4} + \frac{1}{2} - 4u^2 - \|\delta\|_2^2 + o(1). \tag{7}$$

The error is uniform for each fixed bound on $\|p - \frac{1}{2}\mathbf{1}\|_2$. This follows from the estimates below, or by applying sequential convergence to every sequence within the bound.

The tilt density, after canceling $\prod_i d_i(T_0)$, uses $\prod_i[1 + p_i(t-1)]$. Its ratio to the scalar unnormalized density at b is e^ℓ with

$$\ell = \sum_i \log(1 + p_i(t-1)) - n\log(1 + b(t-1)) \le 0.$$

On $1/2 \le t \le 3/2$, Taylor cancellation gives $|\ell| \le C(t-1)^2\|\delta\|_2^2$. For the scalar density, $T + s_b$ is Gamma of shape $n+1$, conditioned to exceed $s_b = T_0(1-b)/b$. Here $s_b = n/2 + O(\sqrt{n})$, so truncation is exponentially small. The centered variable $Y = (T - (n - s_b))/\sqrt{n}$ converges with all fixed moments to $N(0,1)$, and its moments are uniformly bounded. This holds first for Gamma by the independent-exponential representation and exponential moments, then for the heterogeneous tilt because $e^\ell \le 1$ globally, tends to one uniformly for bounded Y, and its normalization tends to one. All fixed polynomial $\tau = T/n$ moments outside $[1/4, 3/4]$ are exponentially small by the Gamma Chernoff bound.

We also need the contribution of $1/T$ near zero. Write $F(T) = \prod_i[1 + p_i(t-1)]$, $Z = \int e^{-T}F(T)\,dT$. Since $F'(T) = F(T)\sum_i x_i/T$, integration by parts on $[0, n/4]$ gives

$$\mathbb{E}\left[\mathbf{1}_{T<n/4}\sum_i x_i/T\right] \le e^{-n/4}F(n/4)/Z + \Pr(T < n/4).$$

The Lebesgue density at $n/4$ is exponentially small by comparison with the scalar shifted Gamma density, using $F \le [1 + b(t-1)]^n$, the bounded normalization ratio, and Stirling's formula there. For $T > 3n/4$, $\sum_i x_i/T \le 4/3$. Thus the expected squared norm of v and the expected diagonal operator on the complement of $[n/4, 3n/4]$ are exponentially small.

On the central interval let $D(p,t) = (x(p,t) - 1/2)^2$. At $t = 1$, $D(p,1) = (p - 1/2)^2$ exactly. Taylor cancellation in strength, using bounded mixed derivatives on $[0,1] \times [1/2, 3/2]$, shows

$$\sum_i D(p_i, t) = nD(b,t) + \|\delta\|_2^2 + O(|t-1|\|\delta\|_2^2).$$

The expected error tends to zero. Also

$$\sqrt{n}(x(b,t) - 1/2) = 2u + Y/2 + o_{L^2}(1).$$

Indeed $n - s_b = n/2 + 2u\sqrt{n} + O(1)$; expand x at $(1/2, 1)$, where its derivatives are 1 and $1/4$, and use the uniformly bounded fourth moments of Y and bounded u. Thus $\mathbb{E}\sum_i D(p_i,t) = 4u^2 + 1/4 + \|\delta\|_2^2 + o(1)$. Meanwhile $\mathbb{E}\sum_i x_i^2/T \to 1/2$: on the central interval $\sum_i x_i^2/n \to 1/4$ in mean by the just-bounded square loss and Cauchy–Schwarz, while $T/n \to 1/2$ in mean and $n/T \le 4$; the tails were handled above.

For the rank and diagonal spectral correction, use the exact expression

$$v_i(T) = \frac{\sqrt{p_i(1 - p_i + p_i/T_0)}\sqrt{t}}{1 + p_i(t-1)}.$$

Choose $e = v(T_0)/\|v(T_0)\|$, $P = I - ee^\mathsf{T}$ and $c(t) = \sqrt{t}/[1 + b(t-1)]$. The difference $v_i(T) - c(t)v_i(T_0)$ equals

$$v_i(T_0)\sqrt{t}\left(\frac{1}{1 + p_i(t-1)} - \frac{1}{1 + b(t-1)}\right).$$

On the central interval this gives $\|Pv\| \le C|t-1|\|\delta\|_2$ without differentiating the square root at the endpoints. Hence $\mathbb{E}\|Pv\|^2 = O(n^{-1})$. The function $v_i(T_0)$ is uniformly Lipschitz on a fixed neighborhood of $1/2$; outside that neighborhood its difference from the scalar value is bounded by a constant times $|p_i - b|$. Thus $\|e - \mathbf{1}/\sqrt{n}\| = O(n^{-1/2})$. The expected diagonal operator is bounded. Its quadratic form on e tends to $1/4$ since replacing e by $\mathbf{1}/\sqrt{n}$ loses $o(1)$, the average of $(1-x_i)^2$ tends to $1/4$, and the remaining average is bounded by $1/n$. The lower block PMP is $O(1)$ in norm. The off-diagonal block PMe has bounded norm by Cauchy–Schwarz from $\mathbb{E}\|Pv\|^2 = O(1/n)$ and $\mathbb{E}\|v\|^2 = O(n)$. The exact scalar Schur complement then shows $\lambda_{\max}(M) = e^{\mathsf{T}}Me + O(1/n) = \mathbb{E}\|v\|^2 + 1/4 + o(1)$. This gives (7).

In particular maximizing sequences satisfy $\|p - \frac{1}{2}\mathbf{1}\|_2 \to 0$. The conclusion rules out even a single loading bounded away from $1/2$ in the p coordinates.

10 Uniform local differentiation and exactness

We first prove the dimension-free estimate $\|D_p^3\Lambda\| \le C$ on a fixed small Euclidean ball around $p = \frac{1}{2}\mathbf{1}$. The auxiliary variables $\beta_i = nr_i = 2p_i/(1-p_i+2p_i/n)$ have bounded Euclidean differentials through order three, since the map is entrywise and a product of unit Euclidean directions has norm at most one. It suffices to work on a small Euclidean ball in β about $\beta_*\mathbf{1}$, where $\beta_* = 2/(1+2/n)$, and transfer the derivative bounds back to p. The local estimates use $b = n^{-1}\sum_i \beta_i$ and $\delta_i = \beta_i - b$ on the ball $\|\beta - \beta_*\mathbf{1}\|_2 \le \varepsilon$, where $\varepsilon > 0$ is fixed and small. With $\tau = T/n$ the diagonal residuals are $a_i = 1 - \beta_i/n$ and $d_i = 1 + \beta_i(\tau - 1/n)$. Write

$$D_n(\beta_i, \tau) = \frac{(\beta_i(\tau + 1/n) - 1)^2}{4a_i d_i^2}.$$

The pointwise trace identity $\|v\|^2 = \sum_i (4a_i)^{-1} - \sum_i D_n(\beta_i, \tau)$ follows by subtraction. The scalar β-tilt translates T by $n/b - 1 = n/2 + O(\sqrt{n})$. The likelihood ratio of the heterogeneous law against this scalar reference is controlled as in the preceding section by Taylor expansion of $\sum_i \log d_i$ around the scalar mean.

For the tilt, put $X = \sqrt{n}(\tau - 1/2)$. The shifted-Gamma comparison from the preceding section, in the equivalent strength variables, applies uniformly since $b = 2 + O(n^{-1/2})$ and $\|\delta\|_2 \le \varepsilon$. To be explicit about the uniform denominator, on $\tau \in [1/4, 3/4]$ the scalar law has mass at least $1/2$ for large n and the likelihood ratio there is bounded below. Thus all fixed moments of X are uniformly bounded, and all polynomial moments of τ outside that interval are exponentially small. Before differentiating the normalized tilted measure in the strengths, subtract its log-density value at $\tau_0 = 1/2$; this does not change the normalized law. The first log-derivative in any unit direction h is

$$\sum_i h_i \left(\frac{\tau - 1/n}{1 + \beta_i(\tau - 1/n)} - \frac{\tau_0 - 1/n}{1 + \beta_i(\tau_0 - 1/n)} \right).$$

Its bound is $C|X|$ because the τ derivative of each fraction is $d_i^{-2} \le 4$ and $\|h\|_1 \le \sqrt{n}$. For higher mixed derivatives through order three, $\sum_i |h_i^{(1)} \cdots h_i^{(k)}| \le 1$ for $k \ge 2$ for unit directions. The derivative formulas for $\log(1 + \beta_i(\tau - 1/n))$ and $d_i \ge 1/2$ give uniform polynomial bounds in $|X|$. Differentiating the quotient of unnormalized integrals gives a finite sum of expectations of observable derivatives times log derivatives and of products of such expectations. It therefore preserves uniform moment bounds, even when integrands are matrix-valued, by taking operator norms and using Hölder. Pointwise bounds on the central interval with exponentially small tails yield the same n-dependent estimates after differentiation.

For the block estimates choose $u_0 = v(\tau_0)$, $e = u_0/\|u_0\|$ and $P = I - ee^{\mathsf{T}}$ as before, and write $w(\tau) = \sqrt{\beta_*\tau}/[1 + \beta_*(\tau - 1/n)]$ for the scalar component at the base strength. The residual $W = v(\tau) - [w(\tau)/w(\tau_0)]u_0$ vanishes at τ_0, and at $\beta = \beta_*\mathbf{1}$. Each component depends only on its own strength. On the central interval, coordinate differentiation and the mean value theorem in τ give $\|D^j W\| \le C|\tau - \tau_0|$ through order three, interpreted in mixed unit directions for $j > 0$; for $j = 0$ use the bound with additional factor $\|\beta - \beta_*\mathbf{1}\|_2$. The norm of an entrywise product of unit Euclidean directions is at most one. The normalization of u_0, whose norm is comparable to $\sqrt{n}$, gives bounded derivatives of e and P to order three. Thus $v_\perp = Pv = PW$ and its derivatives are $O(|X|/\sqrt{n})$, and $e^{\mathsf{T}}v$ and its derivatives are $O(\sqrt{n})$ on the central interval. On its complement strength derivatives in this ball grow at most polynomially in n, τ, since β_i is bounded below and $d_i \ge 1/2$. The estimate for the tilt therefore applies on the full interval.

Denote $D = \operatorname{diag}(a_i/d_i^2)$, $A_0 = e^{\mathsf{T}}Me$, $q = PMe$ and $R = PMP$. The preceding estimates show $\mathbb{E}_\beta\|v_\perp\|^2 = O(1/n)$ through three derivatives and $q, R = O(1)$ through three derivatives. For q use $(e^{\mathsf{T}}v)v_\perp$ and for R use $v_\perp v_\perp^{\mathsf{T}}$, along with their bounded diagonal terms. For A_0 use

$$A_0 = \sum_i (4a_i)^{-1} - \mathbb{E}_\beta \sum_i D_n(\beta_i, \tau) - \mathbb{E}_\beta\|v_\perp\|^2 + \mathbb{E}_\beta e^{\mathsf{T}}De.$$

The first term minus $n/4$ has bounded derivatives to order three, using $\|h\|_1 \le \sqrt{n}$ at order one and the product bound at higher orders. In the loss, $k_i = \beta_i(\tau + 1/n) - 1$ has ℓ^2 norm $O(1 + |X|)$ from the local ball, and its first directional derivative has norm at most $|\tau + 1/n|$. Its higher derivatives vanish. Write the loss as $k^{\mathsf{T}} \operatorname{diag}[(4a_i d_i^2)^{-1}]k$. The central diagonal derivatives are bounded in operator norm. This shows $A_0 - n/4 = O(1)$ through order three.

The top eigenvalue is uniformly simple since $A_0 \sim n/4$ and the lower block R is bounded. It satisfies the exact scalar equation

$$\lambda = A_0 + q^{\mathsf{T}}(\lambda I - R)^{-1}q.$$

Here R can be treated on the full space because $Re = 0$, $q \perp e$, and $\lambda > 0$. The correction function and its first three partial derivatives in β at fixed $\lambda \sim n/4$ are $O(1/n)$, and each additional λ-derivative gains one factor $O(1/n)$. This follows by the resolvent product rule using the bounds on q, R. The implicit derivative coefficient at $d\lambda$ is $1 + q^{\mathsf{T}}(\lambda I - R)^{-2}q \ge 1$. Differentiating successively therefore bounds all three total derivatives uniformly.

We next evaluate the scalar maximum. All $r_i = r \in (0, 1)$ give $s = (1 - r)/r$ and $U = T + s$. By symmetry and positivity the all-ones vector is the top eigenvector. Its Rayleigh quotient is

$$\Lambda = (s + 1)\big[n\mathbb{E}(1/U) - (n - 1)s\mathbb{E}(1/U^2)\big].$$

Here U has the $\Gamma(n + 1, 1)$ law conditioned on $U > s$. For s/n in a fixed neighborhood of $1/2$, replacing the two moments by their unconditioned values $1/n$ and $1/[n(n-1)]$ gives $\Lambda = (s+1)(1 - s/n) + O(ne^{-cn})$. For the omitted inverse moments, their integrals factor as the corresponding lower-shape Gamma probabilities times the exact inverse moments, so each relative error is exponentially small. The quadratic is maximized at $s = (n - 1)/2$ with value $(n + 1)^2/(4n)$. In particular this scalar test supplies $n/4 + 1/2 + o(1)$ for localization.

Let p_n be any maximizing scalar coordinate in $[0, 1]$. It is near $1/2$ by (7) and the prior localization, and therefore interior. The quadratic estimate shows $\Lambda(p_n\mathbf{1}) = (n + 1)^2/(4n) + O(ne^{-cn})$ and $s(p_n) = (n - 1)/2 + O(ne^{-cn/2})$, which gives the asserted off-diagonal strength after decreasing c. For $F_n(p) = \Lambda(A(p))$, permutation invariance and scalar stationarity give $DF_n(p_n\mathbf{1}) = 0$.

Let h be a unit Euclidean direction, $m = n^{-1/2}\sum_i h_i$. For a fixed small $t > 0$, apply (7) at $p_n\mathbf{1}$ and $p_n\mathbf{1} \pm th$. The symmetric second difference divided by t^2 is $-8m^2 - 2(1-m^2) + o(1)$ uniformly in h as $n \to \infty$; the mean offset cancels. The uniform C^3 bound gives an additional error at most Ct when replacing that difference by the Hessian at the center. First fix t small and then take n large. This proves $D^2F_n \preceq -I$ at the center. The third derivative bound makes F_n strictly concave on a fixed smaller ball around the center. Global localization puts every full-class maximizing sequence in this ball, so the stationary scalar profile is the unique maximizer of the full cube in sufficiently large order. Gauge restoration accounts for precisely the stated matrix equalities.

11 Bounded fixed-rank corrections: setup and main result

For a correlation matrix A, let $C(A)_{ij}$ be the permanent with row i and column j deleted. Write $\Lambda(A) = \lambda_{\max}(C(A))/\operatorname{per} A$. Bounds of Marcus and Merris [5] and Lavi [3, Section 4] concern this raw cofactor observable. The all-order bounds in [7] agree up to and including the constant term and leave the $1/n$ coefficient open. A local Hessian calculation Theorem 1 gives strict local maximality at the scalar optimum modulo diagonal-unitary conjugation. The cap-free one-factor result Theorem 3 gives exact eventual maximization for independently variable single loadings. The fixed balanced Walsh result Theorem 8 treats certain additional factors with symmetry.

Here the latent dimension r is fixed. The full matrix has diagonal one and off-diagonal entries $u_i^* u_j/n$, where $u_i \in \mathbb{C}^r$ and $\|u_i\| \le B$. Equivalently one adds a positive diagonal to a rank-at-most-r Gram matrix in all sufficiently large orders. The restriction is on the correction as well as every row at its natural scale. The cofactor spectrum depends on the whole correlation matrix.

Low-rank permanent methods and Gaussian representations include Barvinok [16, 17] and Fassino, Pistone, and Rogantin [8]. McCullagh [10] studies determinant approximations for scalar permanents. Our comparison concerns the top eigenvalue of the entire normalized cofactor matrix. The proof bounds losses due to variation of both the amplitudes and phases of the conditional cofactor vector.

Theorem 4. *Fix $r \ge 1$ and $B > \sqrt{2}$. Let $\gamma_{n,r,B}$ be the maximum of Λ in the class just defined, and let $A_{*,n}$ denote the maximizing positive equicorrelation matrix. Then*

$$0 \le \gamma_{n,r,B} - \Lambda(A_{*,n}) = O_{r,B}(n^{-3}), \qquad \gamma_{n,r,B} = \frac{n}{4} + \frac{1}{2} + \frac{1}{4n} + O_{r,B}(n^{-3}).$$

Any sequence with $\Lambda \ge n/4 + 1/2 - o(1)$ admits diagonal row gauges and a latent unit direction h such that, for $s = \|n^{-1}\sum_i u_i\|$, $\ell_i = u_i^* h$ and $w_i = P_{h^\perp} u_i$,

$$\sum_i w_i = 0, \qquad \rho_n = \frac{1}{n}\sum_i \|w_i\|^2 \longrightarrow 0, \qquad \sqrt{n}(s^2 - 2) \longrightarrow 0, \qquad \sum_i |\ell_i - s|^2 \longrightarrow 0.$$

In particular the gauged matrices converge in Frobenius norm to equicorrelation with off-diagonal $2/n$. For any near-maximizing sequence with $\Lambda > \Lambda(A_{*,n})$ the transverse energy ρ_n in this frame is $O_{r,B}(n^{-1})$.

12 Localizing a near-maximizing profile

Use $a_i = 1 - \|u_i\|^2/n$, $m_i = u_i^* y$, $d_i = a_i + |m_i|^2$, $v_i = m_i/d_i$. The latent Gaussian divided by $\sqrt{n}$ has tilted density proportional to $e^{nf(y)}$ with $f(y) = -\|y\|^2 + n^{-1}\sum_i \log d_i$. By the complex Gaussian moment identity [8] and conditioning on y, the normalized cofactor matrix, up to transpose, is

$M = \mathbb{E}[D + vv^*]$, $D_{ii} = a_i/d_i^2$. Choose a unit top eigenvector e, gauge each row to make its coordinate of e nonnegative, and put $P = I - ee^*$. The exact identity

$$\Lambda = \sum_i (4a_i)^{-1} + \mathbb{E}e^*De - \mathbb{E}\|Z\|^2 - \mathbb{E}\|Pv\|^2, \qquad Z_i = \frac{|m_i|^2 - a_i}{2\sqrt{a_i}d_i} \tag{8}$$

gives bounded expected squared losses, since $a_i \geq 1/2$, the sum is $n/4 + O(1)$ and $D \preceq 2I$. Also $|v_i| \leq 1/(2\sqrt{a_i})$, and $(1/(2\sqrt{a_i}) - |v_i|)^2 \leq Z_i^2$. Thus $\mathbb{E}\||v| - \mathbf{1}/2\|^2 = O(1)$.

Let $\alpha = e^*v$. Pointwise $\||v| - |\alpha|e\| \leq \|Pv\|$. Choose one realization with both losses bounded by a constant, which implies $\||\alpha|e - \mathbf{1}/2\| = O(1)$. Taking norms and subtracting gives $\|e - \mathbf{1}/\sqrt{n}\| = O(n^{-1/2})$. Set $\phi = \arg\alpha$, choosing it arbitrarily at zero. The same triangle inequalities in expectation now give $\mathbb{E}\|v - e^{i\phi}\mathbf{1}/2\|^2 = O(1)$.

The tilted law has exponentially small fixed polynomial tails outside a sufficiently large fixed ball. Its original Gaussian normalization is at least $\prod_i a_i \geq e^{-2B^2}$, and the numerator is bounded by a fixed-dimensional Gaussian factor times $(1 + B^2\|y\|^2)^n$. Taking the ball large enough gives any needed exponential tail. On the ball, $m_i - 2v_i = m_i(d_i - 2)/d_i$ and $d_i - 2 = (|m_i|^2 - a_i) + 2(a_i - 1)$. Thus $\sum_i |m_i - 2v_i|^2 \leq C\sum_i Z_i^2 + O(1/n)$ on the ball. The tail contribution is exponential. It follows that

$$\mathbb{E}\|m - e^{i\phi}\mathbf{1}\|^2 = O(1). \tag{9}$$

The score gives more information than this projected alignment:

$$\partial_{\bar{y}} f = -y + \frac{1}{n}\sum_i u_i v_i, \qquad \mathbb{E}\|\partial_{\bar{y}} f\|^2 = -\frac{1}{4n}\mathbb{E}\,\Delta f = O(n^{-1}).$$

The score identity follows by integrating the real gradient by parts and using $\|\nabla f\|^2 = 4\|\partial_{\bar{y}} f\|^2$. The same score identity underlies [18, Theorem 1]. There is no boundary term by the Gaussian tail. The Laplacian equals $-4r + (4/n)\sum_i \|u_i\|^2 a_i/d_i^2$, so the expected squared score is at most r/n. With $\bar{u} = n^{-1}\sum_i u_i$, the spectral alignment of v and Cauchy–Schwarz therefore yield

$$\mathbb{E}\|y - e^{i\phi}\bar{u}/2\|^2 = O(n^{-1}). \tag{10}$$

The operator with rows u_i^* has norm at most $B\sqrt{n}$. From (9) and (10), $\sum_i |u_i^*\bar{u}/2 - 1|^2 = O(1)$. Averaging gives $s^2 = 2 + O(n^{-1/2})$ and, with $h = \bar{u}/s$, $\|\ell - s\mathbf{1}\|_2 = O(1)$. The definition of h gives $\sum_i w_i = 0$. These conclusions were obtained without assuming an initial transverse covariance gap.

12.1 A nondegenerate limit from tightness

Take the common-phase section $y = e^{i\theta}(\sqrt{x}h + \eta)$, $\theta = \arg h^*y$, $\eta \in h^\perp$. The measure is a constant times $dx\,d\eta$ after removing θ, and the conditional means in the section are $m_i = \ell_i\sqrt{x} + w_i^*\eta$. Put $\beta = s^2$ and $x_0 = 1 - 1/\beta$. The identity $\beta/4 - x_0 = (\beta - 2)^2/(4\beta) = O(1/n)$ and (10) imply tightness of $(\zeta, \xi) = (\sqrt{n}(x - x_0), \sqrt{n}\eta)$. Indeed $\sqrt{n}\,\big||h^*y| - s/2\big|$ and $\sqrt{n}\|\eta\|$ have bounded second moments, and $|h^*y| + s/2$ is tight; the product formula gives the first tightness assertion. Identify the varying $h^\perp$ with a fixed complex space of dimension $r - 1$ by a unitary frame.

On every compact set of scaled coordinates, the log density after subtracting its value at $(x_0, 0)$ is

$$-\frac{1}{2}\zeta^2 + \ell_n^{\mathrm{lin}} \cdot \xi - \xi^{\mathsf{T}} H_n \xi + o(1), \qquad H_n = I - Q_n/2, \quad \xi^{\mathsf{T}} Q_n \xi = \frac{1}{n}\sum_i (\operatorname{Im} w_i^*\xi)^2. \tag{11}$$

Here ℓ_n^{lin} is a bounded real vector. For the Taylor verification, $\sum(\ell_i - s) = 0$, $\sum|\ell_i - s|^2 = O(1)$, and $a_i = 1 + O(1/n)$ yield $\partial_x f(x_0, 0) = O(1/n)$, by subtracting the scalar derivative and canceling its

linear term. The η gradient is $O(n^{-1/2})$ by subtracting the constant scalar coefficient of w_i, using $\sum w_i = 0$ and $\sum |\ell_i - s| = O(\sqrt{n})$. The Hessian at the center differs by $o(1)$ from its value with all $\ell_i = s$, $a_i = 1$: coordinatewise derivatives are uniformly smooth on the bounded parameters, and their average perturbation is $O(n^{-1/2})$. In that scalar array the x Hessian is -1, mixed Hessian vanishes by centering, and transverse Hessian tends to $-2I + Q_n$ up to $o(1)$ from $\beta \to 2$. This uses $\log(1 + |1+z|^2) = \log 2 + \operatorname{Re} z + (\operatorname{Im} z)^2/2 + O(|z|^3)$. The third derivatives in a fixed neighborhood are uniformly bounded, which proves (11).

Take a convergent subsequence of the bounded matrices Q_n and vectors ℓ_n^{lin}. Let N_n be the total normalization after subtracting the local exponent and rescaling volume. From (11), a fixed box yields a positive lower bound on N_n. By tightness choose a large fixed box with probability at least $1/2$; its unnormalized mass is bounded by (11), giving an upper bound on N_n. Any subsequential limit of N_n is finite and positive. Local convergence and tightness then give the limiting probability density proportional to the limiting exponential in (11). Its integrability forces $H = I - Q/2$ positive definite, even though there was no gap assumption. Thus ζ is centered Gaussian of variance one and independent of ξ, whose mean possibly shifts but whose covariance is $H^{-1}/2$.

12.2 Transverse variance and collapse

Let R_n be the corresponding real-form matrix for $\operatorname{Re} w_i^* \xi$. Take further subsequences $R_n \to R$; $\operatorname{tr} R_n = \operatorname{tr} Q_n = n^{-1} \sum \|w_i\|^2$. On any fixed compact scaled set,

$$Z = Z^0 + \frac{1}{2\sqrt{n}} \Big(\zeta \mathbf{1} + (\operatorname{Re} w_i^* \xi)_i \Big) + o_{\ell^2}(1), \qquad Pv = Pv^0 + \frac{i}{2\sqrt{n}} (\operatorname{Im} w_i^* \xi)_i + o_{\ell^2}(1). \tag{12}$$

Superscript zero means $(x, \eta) = (x_0, 0)$. In the second identity use $\sum w_i = 0$, $\|e - \mathbf{1}/\sqrt{n}\| = O(n^{-1/2})$, and the fact that the real differential of $m/(1+|m|^2)$ at $m = 1$ vanishes on the real axis and equals $i \operatorname{Im}/2$ on the imaginary axis. In the first identity the differential of Z at $m = 1, a = 1$ is $\operatorname{Re}/2$, and $\partial_x m = s/(2\sqrt{x_0}) \to 1$. To justify replacing the derivative arrays by their common scalar-center values, their pointwise errors are bounded by $C(|\ell_i - s| + |\beta - 2| + 1/n)$; the squared sum is $O(1)$ and each derivative is multiplied by $1/\sqrt{n}$. Uniform Taylor remainders also have Euclidean norm $O(n^{-1/2})$ on each fixed scaled box. The offsets Z^0, Pv^0 are bounded in Euclidean norm by $\|\ell - s\mathbf{1}\|_2 = O(1)$ and the uniform comparison of e.

The two nonnegative quadratic expressions in (12) have bounded coefficients as polynomials in the fixed-dimensional scaled variables. On a further subsequence these coefficients converge. Their Gaussian expectations are at least the variances of their linear parts, since the squared norm of the deterministic offset plus mean is nonnegative. Equivalently this follows for each n under the fixed limiting Gaussian and then by coefficient convergence. Local density convergence on increasing boxes gives, by nonnegativity and monotone convergence on the limiting polynomials,

$$\liminf \mathbb{E}\|Z\|^2 \geq \frac{1}{4} + \frac{1}{8} \operatorname{tr}(RH^{-1}), \qquad \liminf \mathbb{E}\|Pv\|^2 \geq \frac{1}{8} \operatorname{tr}(QH^{-1}).$$

No uniform integrability of the full scaled law beyond tightness is needed for these lower bounds.

Finally $\mathbb{E} e^* D e \to 1/4$. Replacing the unit vector by $\mathbf{1}/\sqrt{n}$ costs $o(1)$ because $D \preceq 2I$; (9) and the bounded derivative of $a/(a+|m|^2)^2$ in m then give the limiting average. Also

$$\sum_i (4a_i)^{-1} = n/4 + 1/2 + \operatorname{tr} Q_n/4 + o(1).$$

Substitution in (8) gives

$$\limsup(\Lambda - n/4 - 1/2) \leq -\frac{1}{8} \operatorname{tr}\big[(Q+R)(H^{-1} - I)\big] \leq -\frac{1}{16} \operatorname{tr}(Q^2).$$

The assumed near-optimality forces $Q = 0$, hence $\operatorname{tr} Q_n \to 0$ along the full sequence since all convergent subsequences obey this conclusion. The boundedness of the longitudinal error gives $\|\ell\ell^*/n - s^2\mathbf{1}\mathbf{1}^*/n\|_{\mathrm{F}} = O(n^{-1/2})$, allowing conjugation of ℓ if needed by Gram convention. The transverse Gram contribution is positive and its Frobenius norm is at most its trace $\operatorname{tr} Q_n$. Subtracting the diagonal does not increase Frobenius norm. This proves the transverse and Frobenius assertions of Theorem 4.

12.3 Longitudinal convergence

The transverse variance inequality gives $Q_n \to 0$ under $\Lambda \geq n/4 + 1/2 - o(1)$. The transverse linear coefficient in the scaled density expansion then tends to zero: after subtracting the scalar derivative coefficient, its norm is bounded by

$$Cn^{-1/2} \sum_i \|w_i\|(|\ell_i - s| + n^{-1}) = o(1).$$

This follows from Cauchy–Schwarz, $\|\ell - s\mathbf{1}\|_2 = O(1)$, and $n^{-1}\sum \|w_i\|^2 = o(1)$. Keep the squared offsets in the two variance bounds on any convergent subsequence of their polynomial coefficients. The limiting Gaussian is now centered, $Q = R = 0$, and the longitudinal variance already consumes the full $1/4$ loss budget. Hence $\|Z^0\|_2 + \|Pv^0\|_2 \to 0$. Precisely, their liminf sum of squares plus $1/4$ is a lower bound on the sum of the expected losses by the same restriction to increasing boxes used in Section 12.2, while the exact identity and near-optimality bound that total above by $1/4 + o(1)$.

The center moduli satisfy $\||v^0| - \mathbf{1}/2\|_2 = o(1)$ from the defect bound, and $Pv^0 = o_{\ell^2}(1)$. With $e \geq 0$ this implies $\|v^0 - e^{i\phi_n}\mathbf{1}/2\|_2 = o(1)$ for $\phi_n = \arg(e^* v^0)$, by taking coordinate moduli as in the preceding localization. The identity for $m - 2v$ at the bounded center gives $\|m^0 - e^{i\phi_n}\mathbf{1}\|_2 = o(1)$. Its mean $s\sqrt{x_0} = \sqrt{s^2 - 1}$ is real positive, so by averaging and taking moduli, $\sqrt{n}|s^2 - 2| \to 0$ and $\|\ell - s\mathbf{1}\|_2 \to 0$. In particular, the longitudinal vector is $o(1)$ in Euclidean norm from the scalar optimizing longitudinal vector $c_n\mathbf{1}$, where $c_n = \sqrt{nb_n} = \sqrt{2} + O(1/n)$.

Gauge the longitudinal loading real positive coordinatewise, which is legitimate for all sufficiently large n since each complex coordinate differs from s by $o(1)$. Write the resulting loadings as the real longitudinal vector $c_n\mathbf{1} + t_n h$, where $t_n = o(1)$, h is real and $\|h\|_2 = 1$, and row transverse matrix $\sqrt{n\rho_n}\,V$, where $\rho_n = n^{-1}\sum_i \|w_i\|^2 = o(1)$ and $\|V\|_{\mathrm{F}} = 1$ when $\rho_n > 0$. Use conjugate rows if needed to match the Gram convention. Before the gauge change, $V^*\mathbf{1} = 0$. Afterward $\|V^*\mathbf{1}\| \leq \|(\operatorname{diag}(e^{i\theta_i}) - I)\mathbf{1}\| = o(1)$ because the original longitudinal phases have Euclidean norm $o(1)$ and $\|V\| \leq 1$. If $\rho_n = 0$, use the one-factor maximum Theorem 3, or the scalar/local comparison below with $\alpha = 0$.

13 Uniform differentiation in the covariance chart

Fix a real unit h and a complex n-by-k matrix V with $k \leq r - 1$, $\|V\|_{\mathrm{F}} = 1$, and $\|V^*\mathbf{1}\| \leq 1$. For t near zero and $\alpha \geq 0$ near zero let $A(t, \alpha)$ have diagonal one and off-diagonal entries from

$$n^{-1}\ell\ell^* + \alpha VV^*, \qquad \ell = c_n\mathbf{1} + th.$$

We prove $\Lambda(A(t,\alpha)) - n/4 = O(1)$ with all mixed t, α derivatives through total order three uniformly bounded on a fixed small half-rectangle. The covariance parameter α is nonnegative. We first allow both signs of its square-root amplitude.

Use the amplitude $s = \sqrt{\alpha}$ at first and allow either sign. Conditional on the longitudinal latent radial variable $\tau = T/n$, T initially exponential of mean one, and an independent transverse

standard complex Gaussian $\xi \in \mathbb{C}^k$, set

$$z = V\xi, \quad L_i = \|V_{i,:}\|^2, \quad m_i = \ell_i\sqrt{\tau} + sz_i, \quad a_i = 1 - \ell_i^2/n - s^2 L_i, \quad d_i = a_i + |m_i|^2.$$

The removed longitudinal phase may be absorbed into the independent circular ξ. For small $|t|, |s|$ and large n, $a_i \geq 1/2$ since $L_i \leq 1$, $\sum_i L_i = 1$. The tilt is by $\prod_i d_i$. For scalar mean strength $b = \|\ell\|^2/n = 2 + O(n^{-1/2})$, let $d_b = 1 + b(\tau - 1/n)$. Jensen gives

$$\sum_i \log d_i - n \log d_b \leq C\big(|s|\sqrt{\tau}\|\xi\| + s^2\|\xi\|^2\big). \tag{13}$$

Indeed the linear term in $\sum_i d_i - nd_b$ is $2s\sqrt{\tau}\ \mathrm{Re}\langle \ell, z\rangle$, bounded by $C|s|\sqrt{\tau}\|\xi\|$ since $\|V^*\mathbf{1}\| \leq 1$ and $\|V^*h\| \leq 1$. The quadratic term is $s^2(\|z\|^2 - 1)$, and $d_b \geq 1/2$. On $\tau \in [1/4, 3/4]$, $\|\xi\| \leq 1$, the same log ratio is bounded below: Taylor the logarithm at d_b using $d_i, d_b \geq 1/2$, the bounded linear sum, and $\|(d_i - d_b)_i\|_2 = O(1)$. This last bound follows from $\|\ell - c_n\mathbf{1}\|_2 \leq |t|$, $\|z\| \leq \|\xi\|$ and $\|L\|_2 \leq 1$. Thus the tilted normalization compared with the scalar reference is bounded above and below. For the upper bound, absorb the first term in (13) into $\|\xi\|^2/4 + Cs^2\tau$, choose $|s|$ small so $Cs^2\|\xi\|^2 \leq \|\xi\|^2/4$, and use scalar shifted-Gamma domination for τ.

In particular the tilted law has uniform moments of every fixed degree in $X = \sqrt{n}(\tau - 1/2)$, $\sqrt{\tau}$, and $\|\xi\|$. The scalar law is the translate by $n/b - 1$ of Gamma$(n + 1, 1)$ conditioned to be positive after translation. Under the same upper bound an additional fixed factor $e^{Cs^2\tau}$ changes its Gamma rate by $O(n^{-1})$, which preserves these moments and bounded normalization. The lower bound uses the positive scalar mass of $[1/4, 3/4]$ and the fixed-dimensional Gaussian ball. These facts apply uniformly in h, V.

Subtract the value at $(\tau, \xi) = (1/2, 0)$ from the log tilt before differentiating its normalized expectation. The first log derivatives in t, s are bounded by a polynomial in the variables with uniform moments. To see this, compare m_i/d_i to m_0/d_0, $m_0 = c_n\sqrt{\tau}$, $d_0 = 1 - c_n^2/n + |m_0|^2$. The comparison has bounded Euclidean norm up to a polynomial in $\sqrt{\tau}, \|\xi\|$: the differences in m_i, a_i have such norms and the rational map $m/(a + |m|^2)$ is uniformly Lipschitz for $a \geq 1/2$. Thus the error in the first derivative sum involving $\partial_t m_i = h_i\sqrt{\tau}$ or $\partial_s m_i = z_i$ is bounded by Cauchy–Schwarz. The a_i derivative arrays have bounded ℓ^1 norms. The scalar component of the s derivative is bounded since $|\sum z_i| \leq \|\xi\|$. The scalar component of the t derivative is $2c_n\tau \sum h_i/d_0$. Subtracting its value at $\tau = 1/2$ bounds it by $C|X|$: the scalar fraction has bounded derivative in τ and $|\sum h_i| \leq \sqrt{n}$.

Higher mixed log derivatives, through at least order six, have the same polynomial bounds. The first derivatives of d_i form arrays of bounded Euclidean norm up to such polynomials. The second derivatives form arrays of bounded ℓ^1 norm: the m_i derivatives multiply two Euclidean-bounded arrays and the second a_i derivatives use h_i^2/n or L_i. All higher derivatives of d_i vanish. Each chain-rule term for $\sum_i \log d_i$ with order at least two has either two first-derivative factors, summable by Cauchy–Schwarz, or a second-derivative factor, summable in ℓ^1. Remaining factors and d_i^{-1} have polynomial sup bounds. The values subtracted at $(1/2, 0)$ have the same bounds. Differentiation of normalized expectations to order six now preserves any uniform polynomial observable bound, with its derivatives, by the quotient rule and uniform moments. Differentiation under the integral is justified on a slightly larger fixed parameter rectangle by the same scalar/Gaussian exponential domination for each fixed n.

Let $e = \mathbf{1}/\sqrt{n}$, $P = I - ee^*$ in this local chart. For $v_i = m_i/d_i$, the Euclidean norm of Pv and all parameter derivatives through order six are bounded by polynomials in $|X|, \sqrt{\tau}, \|\xi\|$. At order zero compare v_i to m_0/d_0 and apply P. At higher orders each coordinate derivative contains at least one factor from a derivative of m_i or a_i, an array of bounded Euclidean norm, and other factors have

polynomial sup bounds. The same reasoning applies to $Z_i = (|m_i|^2 - a_i)/(2\sqrt{a_i}d_i)$: at order zero the common scalar baseline has Euclidean norm at most $C(1+|X|)$ since its numerator is affine in τ with value $O(1/n)$ at $1/2$, and its remainder has bounded Euclidean norm. The diagonal matrix D and its derivatives have polynomial operator bounds. Finally the scalar $\sum(4a_i)^{-1} - n/4$ and its derivatives are uniformly bounded by bounded ℓ^1 norms of the zero-order deficits $1 - a_i$ and the parameter derivative arrays.

It follows that for $M = \mathbb{E}[D + vv^*]$, $A_0 = e^*Me - n/4$ and $R = PMP$ are bounded with derivatives through order six. For A_0 use the exact trace-defect formula and subtract $\mathbb{E}\|Pv\|^2$. The off-diagonal block $q = PMe$ and its derivatives are $O(\sqrt{n})$, since e^*v and derivatives are bounded by $\sqrt{n}$ times a polynomial and Pv is bounded as just shown. The simple top eigenvalue $\lambda \sim n/4$ satisfies

$$\lambda = n/4 + A_0 + q^*(\lambda I - R)^{-1}q.$$

The correction and its parameter derivatives at fixed $\lambda \sim n/4$ are $O(1)$, and each extra λ-derivative gains $O(1/n)$. Differentiate the equation successively, whose implicit coefficient is $1 + q^*(\lambda I - R)^{-2}q \geq 1$, to obtain uniform derivatives through order six in t, s. The eigenvalue is even in s by symmetry of ξ. For an even function f, the derivative in $\alpha = s^2$ is $\frac{1}{2}\int_0^1 f''(u\sqrt{\alpha})\,du$, extending continuously to zero. Iterating this identity with bounded derivatives and powers of u proves mixed derivative bounds in t, α whenever the weighted order $i + 2j \leq 6$ for $\partial_t^i\partial_\alpha^j$. This gives the required uniform ordinary C^3 estimates on the half-rectangle.

14 Comparison with the affine Hessian

At $(0,0)$, the full matrix gradient vanishes by the scalar stationarity and permutation and conjugation symmetry. Thus the chart Hessian is the affine Hessian evaluated on the tangent

$$tX + \alpha Y, \qquad X_{ij} = c_n(h_i + h_j)/n, \qquad Y = VV^* - \operatorname{diag}(L_i).$$

The affine Hessian Theorem 1 has the following spectral sectors. The constant real mode and real stars have negative coefficients of order n; real and imaginary zero-row-sum sectors have negative coefficients bounded away from zero; the imaginary gradients are the gauge kernel. The tangent X is entirely in the constant plus real-star spaces, with $n\|X\|_{\mathrm{F}}^2$ bounded above and below away from zero uniformly in unit h. For Y, $\|Y\|_{\mathrm{F}} \leq 1$ and $\|Y\mathbf{1}\| \leq \|V^*\mathbf{1}\| + \|L\|_2 \leq 2$. The projections of Y onto the constant and real-star spaces and the imaginary gradient space have Frobenius norm $O(n^{-1/2})$. Indeed the projection onto constant and real-star spaces is determined by the real row sums, with row-sum to Frobenius singular values $\sqrt{n-1}$ for constants and $\sqrt{(n-2)/2}$ for stars, obtained directly from $X_{ij} = a_i + a_j$. The imaginary gradient singular value is $\sqrt{n/2}$ from $X_{ij} = i(a_i - a_j)$, $\sum a_i = 0$. Thus the weighted Hessian norm of Y is bounded.

If $\max_i L_i \leq 1/[16(r-1)]$ for $r > 1$, then $\|VV^*\|_{\mathrm{F}} \geq 1/\sqrt{r-1}$ by rank and trace, while $\|\operatorname{diag} L\|_{\mathrm{F}} \leq \sqrt{\max_i L_i}$. The projection of Y onto the two zero-row-sum sectors therefore has Frobenius norm bounded below by a positive constant depending on r. The negative Hessian on $tX + \alpha Y$ then bounds $t^2 + \alpha^2$ below up to a constant. To see this, the normal part bounds α^2, and in the weighted real-star plus constant part the norm of tX is bounded below by $c|t|$ while the norm of αY is bounded above by $C|\alpha|$. A triangle inequality controls t^2 by the weighted combined norm plus α^2. The uniform C^3 remainder therefore gives a strictly smaller objective for any sufficiently small nonzero chart displacement under this leverage condition.

For an arbitrary V, the longitudinal Hessian entry remains strictly negative. The uniform derivative bounds give the upper Taylor estimate $-ct^2 + C|t|\alpha + C\alpha^2 + C(|t| + \alpha)^3$ for the excess. Thus if $|t| \geq K\alpha$ for a sufficiently large fixed K and the neighborhood is sufficiently small, the

excess is nonpositive. If $|t| < K\alpha$, use the full negative semidefiniteness of the Hessian at the origin, including the gauge kernel. The quadratic term is nonpositive, while the cubic remainder is $O(\alpha^3)$. Together these give

$$\Lambda(A(t,\alpha)) \le \Lambda(A_{*,n}) + C\alpha^3. \tag{14}$$

If the original row cap is B and $\alpha = \rho_n \ge 16(r-1)B^2/n$, then $L_i = \|w_i\|^2/(n\rho_n) \le 1/[16(r-1)]$, so the strict comparison applies. Otherwise (14) gives excess $O_{r,B}(n^{-3})$.

A maximizing profile exists by compactness and includes the scalar test $A_{*,n}$ for sufficiently large n. Every such sequence is near-optimal for the Gaussian-score localization, so $(t,\alpha) \to (0,0)$ and $\|V^*\mathbf{1}\| \le 1$ eventually. Apply either case above uniformly; a contradiction sequence handles any purported failure for infinitely many n. The scalar expansion now gives the claimed sharp $1/n$ coefficient and the $O(n^{-3})$ error. Thus any improvement over the scalar maximum lies in the small-energy, relatively high-leverage regime in Theorem 4.

15 A quantitative centered transverse model

Theorem 5. *Fix $r \ge 2$, $B < \infty$, and $\epsilon > 0$. For every n, take $w_i \in \mathbb{C}^{r-1}$ with $\sum_i w_i = 0$, $\max_i \|w_i\| \le B$, and $\Sigma_n = n^{-1}\sum_i w_i w_i^* \preceq (2-\epsilon)I$. Put $\beta = 2 + u_n/\sqrt{n}$ with bounded real u_n. For sufficiently large n, let A be the correlation matrix with Gaussian factors $(\sqrt{\beta}, w_i)$ multiplied by $n^{-1/2}$ and diagonal residual $a_i = 1 - (\beta + \|w_i\|^2)/n$. Define real symmetric Q_n, R_n on the realification of $\mathbb{C}^{r-1}$ by*

$$\eta^{\mathsf{T}} Q_n \eta = \frac{1}{n}\sum_i (\operatorname{Im} w_i^*\eta)^2, \qquad \eta^{\mathsf{T}} R_n \eta = \frac{1}{n}\sum_i (\operatorname{Re} w_i^*\eta)^2.$$

Let $H_n = I - Q_n/2$, which is uniformly positive. Then

$$\Lambda(A) = \frac{n}{4} + \frac{1}{2} - \frac{u_n^2}{16} - \frac{1}{8}\operatorname{tr}\left[(Q_n + R_n)(H_n^{-1} - I)\right] + o(1)$$

uniformly for the given caps and bounded u_n. The transverse loss is at least $\operatorname{tr}(Q_n^2)/16$, so an order-one empirical transverse covariance forces a strict order-one deficit. If one secondary real Walsh mode has strength $\kappa < 2$, the transverse loss reduces to $\kappa^2/[8(2-\kappa)]$.

15.1 Laplace law with the common phase removed

Write the latent vector divided by $\sqrt{n}$ as $y = (y_0, \eta')$, with complex Gaussian density proportional to $\exp(-n\|y\|^2)$. Remove the common phase by taking $y_0 = \sqrt{x}\, e^{i\theta}$, $\eta' = e^{i\theta}\eta$, $x \ge 0$. The measure after integration in θ is constant times $dx\, d\eta$. Each conditional mean modulus comes from $m_i = \sqrt{\beta x} + w_i^*\eta$. Put $d_i = a_i + |m_i|^2$. The tilted density is proportional to $\exp(nf_n)$ with

$$f_n(x,\eta) = -x - \|\eta\|^2 + \frac{1}{n}\sum_i \log d_i.$$

Center at $x_0 = 1 - 1/\beta$ and $\eta = 0$. For global domination set $\bar{a} = n^{-1}\sum a_i$, $q = x + \eta^*\Sigma_n\eta/\beta$, and $\phi_n(q) = -q + \log(\bar{a} + \beta q)$. Jensen and centering give

$$f_n(x,\eta) \le \phi_n(q) - (1 - \|\Sigma_n\|/\beta)\|\eta\|^2.$$

The scalar maximum is at $q_n = 1 - \bar{a}/\beta = x_0 + O(n^{-1})$. There, the Jensen gap at the point $(q_n, 0)$ is $O(n^{-2})$ because $a_i - \bar{a} = O(n^{-1})$ and the linear terms cancel. The change of variables

$x \mapsto q$ has Jacobian one and extends to $q \geq 0$ for bounding. On a fixed interval about q_n, $\phi_n(q) \leq \phi_n(q_n) - c(q - q_n)^2$, with exponential gaps/tails on its complement and a linear bound at large q. The transverse bound has a fixed quadratic gap. On a fixed box in $\zeta = \sqrt{n}(x - x_0)$, $\xi = \sqrt{n}\eta$ the exact exponent stays within $O(1)$ of $n\phi_n(q_n)$ by the Taylor calculation below, giving the required lower normalizer of order $n^{-r+1/2}\exp(n\phi_n(q_n))$. The change from q to x costs $O(\|\eta\|^2)$, so these bounds give bounded moments of every fixed order in (ζ, ξ), uniform bounds on the scaled densities on compact sets, and exponentially small tails outside sufficiently small fixed neighborhoods of $(x_0, 0)$ after any polynomial weight in $x, \|\eta\|$. For the convergence of moments it suffices to apply the uniform higher moment bounds after truncating at a fixed scaled radius.

In scaled variables the exponent relative to $f_n(x_0, 0)$ is

$$-\frac{\zeta^2}{2} - \xi^{\mathsf{T}} H_n \xi + o(1)$$

uniformly on compact scaled sets. To check, at the center the leading common conditional mean is 1 and leading $d_i = 2$. For a perturbation $m_i = 1 + \delta m_i$, the first and second terms of $\log(1 + |m_i|^2) - \log 2$ are $\operatorname{Re}\delta m_i + \frac{1}{2}(\operatorname{Im}\delta m_i)^2$. The radial Taylor expansion of $-x + \log(1 + \beta x)$ at x_0 has second derivative -1 exactly. The cross derivative with η averages to zero by $\sum w_i = 0$. At bounded u_n, replacing the common center $\sqrt{\beta - 1}$ by 1 in the transverse Hessian costs $O(n^{-1/2})$; the possible transverse linear term from a_i variations is $O(n^{-1})$ before scaling because the order-one term cancels. These estimates and the bounded third derivatives on a fixed neighborhood prove the expansion. Thus all fixed polynomial moments are given up to $o(1)$ by independent $\zeta \sim N(0, 1)$ and real $\xi \sim N(0, \frac{1}{2}H_n^{-1})$, uniformly. This is also obtained by subsequences $Q_n \to Q$, which exist by fixed dimension and uniform bounds.

15.2 Amplitude and phase losses

The conditional cofactor identity gives $M = \mathbb{E}[D + vv^*]$, with $D_{ii} = a_i/d_i^2$ and $v_i = m_i/d_i$ up to conjugating the full matrix. The common phase has dropped out. Let $e = \mathbf{1}/\sqrt{n}$, $P = I - ee^*$, and let $L_n\xi$ be the vector with entries $\operatorname{Im} w_i^*\xi$. Expand on the central neighborhood:

$$\sqrt{n}(|m_i|^2 - a_i) = u_n + \beta\zeta + 2\operatorname{Re} w_i^*\xi + O\big(n^{-1/2}(1 + |\zeta| + \|\xi\|)^2\big).$$

The same bound with a constant independent of i, n follows from $\max\|w_i\| \leq B$ and the positivity of x_0. The exact trace identity

$$\|v\|^2 = \sum_i \frac{1}{4a_i} - \sum_i \frac{(|m_i|^2 - a_i)^2}{4a_i d_i^2}$$

therefore gives

$$\mathbb{E}\|v\|^2 = \frac{n}{4} + \frac{\beta + \operatorname{tr} Q_n}{4} - \frac{u_n^2 + 4}{16} - \frac{1}{8}\operatorname{tr}(R_n H_n^{-1}) + o(1).$$

In this computation $\sum_i \operatorname{Re} w_i^*\xi = 0$, $\operatorname{tr} Q_n = \operatorname{tr} R_n = \operatorname{tr}\Sigma_n$, and the radial Gaussian has mean zero and variance one. All denominator replacements are controlled in mean by bounded scaled moments, and the omitted tails are exponential since $a_i \geq 1/2$ eventually and the trace defect is pointwise bounded per coordinate.

At the common center the map $m \mapsto m/(1 + |m|^2)$ has differential $z \mapsto i \operatorname{Im} z/2$ at $m = 1$. Centering therefore gives

$$Pv = \frac{i}{2\sqrt{n}} L_n\xi + o_{L^2}(1), \qquad e^*v/\sqrt{n} = \tfrac{1}{2} + o_{L^2}(1).$$

For the first formula one can subtract the constant scalar v evaluated at the perturbed x with $\eta = 0, a_i = 1$, then use uniform Taylor expansion; the residual $a_i - 1$, second-order terms, and the derivative errors from the moving radial center have total Euclidean norm $O(n^{-1/2}(1+|\zeta|+\|\xi\|)^2)$. Consequently

$$\mathbb{E}\|Pv\|^2 = \tfrac{1}{8}\operatorname{tr}(Q_n H_n^{-1}) + o(1).$$

The diagonal Rayleigh term tends to $1/4$. The lower block PMP has $O(1)$ norm. To see that the upper eigenvalue equals the e-Rayleigh quotient up to $o(1)$, the off-diagonal block must be $o(\sqrt{n})$, not only $O(\sqrt{n})$. Its division by $\sqrt{n}$ and the last two expansions reduce it to $iL_n\mathbb{E}\xi/(4\sqrt{n}) + o(1) = o(1)$, because $\|L_n\| \le C\sqrt{n}$ and the limiting transverse Gaussian is centered. The Schur complement for the $n/4 + O(1)$ Rayleigh quotient now gives $o(1)$. Subtract the transverse rank loss from the trace, add the diagonal term, and use $2\operatorname{tr} Q_n = \operatorname{tr}(Q_n + R_n)$. This proves the formula.

Finally $H_n^{-1} - I \succeq Q_n/2$, and $\operatorname{tr}[(Q_n+R_n)Q_n] \ge \operatorname{tr}(Q_n^2)$ by positivity. For a real one-dimensional secondary direction, Q_n and R_n have disjoint ranges and the single eigenvalue of Q_n is the empirical secondary strength, giving the claimed Walsh reduction.

16 Bounded Gram corrections of any rank

For a correlation matrix A, let $C(A)_{ij}$ be the permanent obtained by deleting row i and column j. We study $\Lambda(A) = \lambda_{\max}(C(A))/\operatorname{per} A$. Marcus and Merris [5] and Lavi [3, Section 4] give earlier spectral inequalities for this raw cofactor observable. Schubr [7] bounds its unrestricted maximum up to constant order, and Theorem 1 gives the full affine Hessian at the scalar equicorrelation optimum. Exact one-factor maximization Theorem 3 includes singular diagonal residuals. The bounded fixed-rank result Theorem 4 gives an $O(n^{-3})$ upper bound on possible excess, using a local chart normalized by the transverse trace.

Here the entries of A off diagonal are $u_i^* u_j/n$ with $\|u_i\| \le B$. The diagonal entries of A are one. There is no bound on the rank of the Gram correction. Equivalently, $A = I + G - \operatorname{diag} G$, with $G \succeq 0$ and $G_{ii} \le B^2/n$. For fixed B this is also exactly the spectral condition $A \succeq (1 - B^2/n)I$. One direction follows from $G \succeq 0$ and its diagonal bound. For the reverse, take $G = A - (1 - B^2/n)I$. The sequence of sharp upper bounds in this paper concerns this bounded-diagonal Gram model.

Theorem 6. *Fix $B \ge \sqrt{2}$. Let $\gamma_{n,B}$ be the maximum of Λ over the bounded Gram correction class, and $A_{*,n}$ the positive equicorrelation optimizer with off-diagonal entry b_n. Then*

$$0 \le \gamma_{n,B} - \Lambda(A_{*,n}) = O_B(n^{-3/2}), \qquad \gamma_{n,B} = \frac{n}{4} + \frac{1}{2} + \frac{1}{4n} + O_B(n^{-3/2}).$$

Any sequence with $\Lambda \ge n/4 + 1/2 - o(1)$ admits row gauges and a unit latent direction h. With $s = \|n^{-1}\sum_i u_i\|$, $\ell_i = u_i^* h$, and $w_i = P_{h^\perp} u_i$, in these gauges

$$\sqrt{n}(s^2 - 2) \to 0, \qquad \|\ell - s\mathbf{1}\|_2 \to 0, \qquad \|(w_i^* w_j/n)_{ij}\|_{\mathrm{F}} \to 0.$$

The gauged matrices converge in Frobenius norm to equicorrelation with off-diagonal entry $2/n$. Any possible excess over the exact scalar optimum in a near-maximizing sequence occurs at transverse Frobenius norm $O_B(n^{-1/2})$ in this decomposition.

The score argument uses Gaussian integration by parts [18] with a trace-bounded weight. The quadratic comparison controls the permanent-tilted law itself by a Gaussian with an adjusted covariance. It derives the gap for the comparison quadratic form from near-optimality. General Gaussian total variation estimates such as [20] compare Gaussian laws given their covariances. The

local regularity uses Gaussian rotation [19] and the heat differentiation for a Gaussian expectation [21]. Its trace cancellation is specific to the conditional cofactor vector and removes a term of size up to $\sqrt{n}$ in the local chart.

17 Rank-independent localization

17.1 Weighted score estimates

Restrict the latent space to the span of the loadings so its dimension is at most n. Use a latent Gaussian y with covariance I/n before tilting. Write $m_i = u_i^* y$, $d_i = a_i + |m_i|^2$, $a_i = 1 - \|u_i\|^2/n$, $v_i = m_i/d_i$, and $Z_i = (|m_i|^2 - a_i)/(2\sqrt{a_i} d_i)$. The Gaussian moment identity [8], followed by conditioning on y, gives the normalized cofactor matrix up to transpose as $M = \mathbb{E}[D + vv^*]$ with $D_{ii} = a_i/d_i^2$, under the tilt by $\prod_i d_i$. Let U be the matrix of rows u_i^*, choose a unit top eigenvector e of M, and write $P = I - ee^*$. Make $e \geq 0$ by row gauges. The exact deficit identity

$$\Lambda = \sum_i (4a_i)^{-1} + \mathbb{E} e^* D e - \mathbb{E}\|Z\|^2 - \mathbb{E}\|Pv\|^2$$

gives both expected losses $O(1)$. The pointwise identity $Z_i^2 = (4a_i)^{-1} - |v_i|^2$ yields $\mathbb{E}\||v| - \mathbf{1}/2\|^2 = O(1)$, since $a_i = 1 + O(1/n)$ and the squared gap to the maximal modulus is bounded by Z_i^2. Let $\xi_0 = e^* v$. Pointwise $\|\,|v| - |\xi_0| e\| \leq \|Pv\|$. Choose a realization where both losses are bounded. It gives $\|\,|\xi_0| e - \mathbf{1}/2\| = O(1)$, and $\|e - \mathbf{1}/\sqrt{n}\| = O(n^{-1/2})$ follows by taking norms. The same inequalities in expectation give $\mathbb{E}\|v - e^{i\phi}\mathbf{1}/2\|^2 = O(1)$ with $\phi = \arg e^* v$. At a zero argument choose any phase.

For $\|y\|$ large, the Gaussian tail estimate stays uniform because its original Gamma shape is at most n. The tilted normalizer relative to the Gaussian is at least $\prod_i a_i \geq e^{-2B^2}$. The tilt is at most $(1 + B^2\|y\|^2)^n$. With R large, on $\|y\| > R$ this is at most $e^{n\|y\|^2/4}$. Absorbing polynomial factors with a further $e^{n\|y\|^2/16}$, the Gaussian Chernoff bound gives exponential decay since $\mathbb{E}_0 e^{n\|y\|^2/2} \leq 2^n$. Increasing R lets the residual e^{-cnR^2} beat 2^n. On the ball, use $m_i - 2v_i = m_i(d_i - 2)/d_i$ and the Z loss to obtain

$$\mathbb{E}\|m - e^{i\phi}\mathbf{1}\|^2 = O(1). \tag{15}$$

The realified potential $f = -\|y\|^2 + n^{-1}\sum \log d_i$ has uniformly bounded real Hessian in operator norm, independently of rank: each real Hessian of a log term is dominated in absolute quadratic form by $C|u_i^* h'|^2$ for a real direction h', using $a_i \geq 1/2$. For any deterministic complex positive semidefinite K, realify it and integrate by parts to get the weighted score identity (cf. [18, Theorem 1])

$$\mathbb{E}\langle \partial_{\bar{y}} f, K \partial_{\bar{y}} f\rangle = -\frac{1}{4n}\mathbb{E}\,\mathrm{tr}(\mathcal{R}(K) D_{\mathbb{R}}^2 f) \leq \frac{C}{n}\,\mathrm{tr}\,K.$$

The Gaussian tails justify the boundary. Here $\partial_{\bar{y}} f = -y + U^* v/n$. Choose $K = U^* U/n$, whose trace is at most B^2, and combine with the spectral alignment of v and $\|U\| \leq B\sqrt{n}$. This gives

$$\mathbb{E}\|m - e^{i\phi} U\bar{u}/2\|^2 = O(1), \qquad \bar{u} = n^{-1}\sum_i u_i.$$

For the replacement of v use the operator norm $\|UU^*/n\| \leq B^2$. With (15) this proves $\|U\bar{u}/2 - \mathbf{1}\| = O(1)$, hence $s^2 = 2 + O(n^{-1/2})$, $h = \bar{u}/s$, $\|\ell - s\mathbf{1}\| = O(1)$, and $\sum w_i = 0$.

Choose $K = hh^*$ for the radial direction. It gives $\mathbb{E}|h^* y - e^{i\phi} s/2|^2 = O(1/n)$: the replacement of v costs at most $\|Uh\|/n = O(n^{-1/2})$ times its $O(1)$ Euclidean error. Take the phase section

$y = e^{i\theta}(\sqrt{x}\,h + \eta)$, $\theta = \arg h^*y$, with $x = |h^*y|^2$. Write $z = (w_i^*\eta)_i$, $\beta = s^2$, $x_0 = 1 - 1/\beta$, and $\zeta = \sqrt{n}(x - x_0)$. The radial estimate and $\beta/4 - x_0 = O(1/n)$ give tightness of ζ. From (15) and the radial estimate, $\mathbb{E}\|z\|^2 = O(1)$: the section vector $e^{-i\theta}m$ is within bounded mean-square distance of $e^{i(\phi-\theta)}\mathbf{1}$; $\sqrt{n}|e^{i(\phi-\theta)} - 1|$ is bounded in mean square because the complex number h^*y is close to $e^{i\phi}s/2$, with s bounded below; and $\sqrt{x}\,\ell - \mathbf{1}$ is bounded in mean square by the radial estimate and the deterministic ℓ bound. For the phase bound use $|e^{i\theta} - e^{i\phi}| \le 4|h^*y - e^{i\phi}s/2|/s$ even if $h^*y = 0$. We now have joint tightness in $|\zeta|, \|z\|_2$, without tightness in the full latent vector at scale $\sqrt{n}$.

17.2 Quadratic comparison of the projected field

Under the untilted Gaussian in the phase section, the radial variable and z are independent, and z is circular complex Gaussian with covariance $\Omega = (w_i^* w_j/n)_{ij}$. It satisfies $\Omega_{ii} \le B^2/n$, $\operatorname{tr}\Omega \le B^2$. Write z also as a real vector when needed, with real covariance $\Gamma = \mathcal{R}(\Omega)/2$. Take the reference law ν_n that tilts only the radial variable by the homogeneous scalar denominator $d_b(x)^n$, $d_b = 1 - \beta/n + \beta x$. Then ζ converges to $N(0,1)$ with uniform moments via the shifted Gamma formula and z is unchanged. The full law μ_n has likelihood proportional to

$$L_n(x,z) = \frac{\prod_i (a_i + |\ell_i\sqrt{x} + z_i|^2)}{d_b(x)^n}.$$

On the domain $E_K = \{|\zeta|, \|z\|_2 \le K\}$, for every fixed K, $\log L_n$ is bounded in absolute value by a constant depending on K. Indeed if $b_i = a_i + |\ell_i\sqrt{x} + z_i|^2$, its difference array from d_b has bounded Euclidean norm there. Its sum is bounded since $\sum(\ell_i - s) = 0$, $\sum z_i = 0$, $\|\ell - s\mathbf{1}\|_2 = O(1)$, and the diagonal residual difference has bounded sum. Taylor expansion of the logarithms at d_b with $b_i, d_b \ge 1/2$ proves the assertion. Both measures have positive probability on some fixed E_K, by their respective tightness bounds. Consequently the likelihood normalizer $\int L_n\, d\nu_n$ is bounded above and below by positive constants. The same holds for $L_n(x_0, 0)$.

The reference Gaussian satisfies $\Pr(\max_i |z_i| > n^{-1/4}) \to 0$ by its diagonal bound and a union bound. The same holds under μ_n: first restrict to a high-probability E_K, on which the normalized likelihood is bounded, and then let $K \to \infty$. On $E_K \cap \{\|z\|_\infty \le n^{-1/4}\}$ there is the uniform expansion

$$\log[L_n(x,z)/L_n(x_0,0)] = \tfrac{1}{2}\|\operatorname{Im} z\|_2^2 + c_n^{\mathsf{T}} z_{\mathbb{R}} + o(1), \qquad \|c_n\|_2 = O(1). \tag{16}$$

For the radial part use the log ratio at $z = 0$, whose derivative near x_0 is $O(1)$: cancel terms linear in $\ell_i - s$ and bound the diagonal residual discrepancy in ℓ^1. Over $|x - x_0| \le K/\sqrt{n}$ this costs $o(1)$. For the linear transverse term, subtract the scalar coefficient of z_i, whose sum vanishes exactly. Its residual coefficient vector at x_0 has bounded Euclidean norm by the ℓ discrepancy and the $O(n^{-1})$ diagonal residual deviations. Changing x costs $o(1)$ in Euclidean norm after this same scalar subtraction: derivatives of the residual coefficient are bounded by $C(|\ell_i - s| + 1/n)$. The leading quadratic term at $m = 1, a = 1$ is $(\operatorname{Im} z_i)^2/2$. Errors from diagonal and loading differences are at most $C\sum_i(|\ell_i - s| + n^{-1/2})|z_i|^2 = o(1)$ on this set by Cauchy–Schwarz and its sup bound. Third-order Taylor errors are bounded by $C\|z\|_\infty\|z\|_2^2 = o(1)$.

Let P_I be orthogonal projection to the imaginary coordinates of the realification, and $J_n = \Gamma^{1/2} P_I \Gamma^{1/2}$. In the reference representation $z_{\mathbb{R}} = \Gamma^{1/2} g$ with real standard Gaussian g, the comparison tilt in (16) is $\exp(\frac{1}{2} g^{\mathsf{T}} J_n g + b_n^{\mathsf{T}} g)$ with $b_n = \Gamma^{1/2} c_n$, uniformly bounded in norm. We claim a strict uniform gap $\|J_n\| \le 1 - \delta$ holds eventually along a near-maximizing sequence. Otherwise take a subsequence with leading eigenvalue $\lambda_n \ge 1 - o(1)$, and split $g = qv + g'$ into the unit leading mode v and its independent orthogonal Gaussian. Restrict $|q| \le L$, $|\zeta| \le K_0$, $\|\Gamma^{1/2} g'\| \le K_0$ and $\|\Gamma^{1/2} g'\|_\infty \le n^{-1/4}/2$ in the complex-coordinate norm. Fixed large K_0 gives uniform positive

Gaussian mass for the latter two restrictions by the trace and diagonal bounds. The event in g' is even, so integrating its exponential quadratic plus linear tilt costs at least its probability by the even average and positivity of the quadratic. The radial event has positive mass. For fixed L the full event is within the domain of (16) with K depending on L, since each coordinate of $\Gamma^{1/2}v$ is $O(n^{-1/2})$ and its Euclidean norm is bounded. The q integral is bounded below by a fixed multiple of

$$\int_{-L}^{L} \exp\bigl[-(1-\lambda_n)q^2/2\bigr]\,dq,$$

using even averaging in q as well. As $n \to \infty$ this has liminf at least $2L$. The normalized likelihood integral was bounded above, so letting $L \to \infty$ gives a contradiction.

Let $\nu_n^{(2)}$ be the reference law with the normalized quadratic tilt from (16). It keeps the radial law and makes $z_{\mathbb{R}}$ a Gaussian with covariance

$$\Gamma_2 = \Gamma^{1/2}(I - J_n)^{-1}\Gamma^{1/2}$$

and mean $\mu_z = \Gamma^{1/2}(I - J_n)^{-1}b_n$. Its normalization is uniformly bounded above and below: use the determinant formula, the gap, $\operatorname{tr} J_n \le B^2$, and $\|b_n\| \le C$. It has uniform moments of $\|z\|$ by its trace and mean bounds. Its coordinate variances are $O(n^{-1})$ by $\Gamma_2 \preceq \delta^{-1}\Gamma$, and its coordinate means are $O(n^{-1/2})$ by the diagonal bound of Γ. Thus for large K the high-probability sets of (16) have high probability under both μ_n and $\nu_n^{(2)}$. There the unnormalized likelihood ratios agree up to $1 + o(1)$ after using the factor $L_n(x_0, 0)$. This proves total variation convergence of the two laws: first bound the normalization ratio by the ratios of the restricted masses, then take $n \to \infty$ at fixed K followed by $K \to \infty$.

17.3 The cofactor deficit and Frobenius convergence

On the sets of (16), uniform Taylor expansion in the section gives

$$Z = Z^0 + \frac{1}{2}(\zeta\,\mathbf{1}/\sqrt{n} + \operatorname{Re} z) + o_{\ell^2}(1), \qquad Pv = Pv^0 + \frac{i}{2}\operatorname{Im} z + o_{\ell^2}(1).$$

Indeed the second residuals in each coordinate are bounded by $C(|\zeta|/\sqrt{n} + |z_i|)^2$, which tend to zero in Euclidean norm on those sets. The first differential coefficients may be replaced by those at $m = 1, a = 1$: the radial errors are bounded in Euclidean norm before multiplying $\zeta/\sqrt{n}$, and the transverse coefficient errors from $|\ell_i - s|$ cost at most $C\|\ell - s\mathbf{1}\|\|z\|_\infty$. Also $\sum z_i = 0$, so replacing the projection onto $e^\perp$ by the projection onto the constant vector's complement costs $o(1)$. Both offsets Z^0, Pv^0 are bounded in Euclidean norm.

The two nonnegative squared losses under μ_n, restricted to the high-probability sets, can now be compared with their polynomial approximations under $\nu_n^{(2)}$ by total variation. On each fixed E_K the polynomials are bounded. Outside the sets their comparison-law expectations are uniformly small as $K \to \infty$ by uniform higher moments, while the true losses are nonnegative. Therefore their sum is bounded below by

$$\tfrac{1}{4} + \tfrac{1}{4}\operatorname{tr}\Gamma_2 + \|Z^0 + \tfrac{1}{2}\operatorname{Re}\mu_z\|^2 + \|Pv^0 + \tfrac{i}{2}\operatorname{Im}\mu_z\|^2 - o(1). \tag{17}$$

The negligible radial mean is absorbed into $o(1)$, and its variance tends to one. The radial and transverse comparison variables are independent. The diagonal Rayleigh contribution tends to $1/4$ by (15) and the uniform gauge comparison of e. The trace term in the deficit identity is $n/4 + 1/2 + \operatorname{tr}\Omega/4 + o(1)$, as seen by its first-order expansion in the diagonal residuals. Note $\operatorname{tr}\Gamma = \operatorname{tr}\Omega$.

The excess of comparison covariance is nonnegative and controls the Frobenius norm:

$$\operatorname{tr}(\Gamma_2 - \Gamma) \geq \operatorname{tr}(\Gamma J_n) = \operatorname{tr}(P_I \Gamma^2) = \tfrac{1}{4}\|\Omega\|_{\mathrm{F}}^2.$$

The first inequality uses $(I - J_n)^{-1} - I \succeq J_n$; the last equality is the circular realification identity $\Gamma = \mathcal{R}(\Omega)/2$, whose imaginary diagonal block in the square has half the full realified trace. Thus the near-maximal loss budget forces $\|\Omega\|_{\mathrm{F}} \to 0$ and both shifted offsets in (17) tend to zero. Moreover $\|\mu_z\| \leq \delta^{-1}\|\Gamma\|\|c_n\| = o(1)$, so the offsets themselves tend to zero. Taking moduli and then aligning v^0 to a common phase via Pv^0 gives m^0 converging in Euclidean norm to a unimodular constant vector. In detail, $\||v^0| - \mathbf{1}/2\| \to 0$ by the defect identity, and $\|Pv^0\| \to 0$ then gives $\|v^0 - e^{i\phi_n}\mathbf{1}/2\| \to 0$ by the same modulus comparison used above. The bounded center satisfies $\|m^0 - 2v^0\| \to 0$ by the identity preceding (15). Its mean is $s\sqrt{x_0} = \sqrt{s^2 - 1}$, real positive. Thus $\sqrt{n}(s^2 - 2) \to 0$ and $\|\ell - s\mathbf{1}\| \to 0$. The longitudinal Gram, the vanishing transverse Frobenius term, and diagonal subtraction give the convergence to the equicorrelation orbit. The transverse trace may persist: for $0 < \tau_* < B^2 - 2$, take $\Omega = \tau_*(I - \mathbf{1}\mathbf{1}^*/n)/n$ and a common longitudinal loading of squared size $nb_n + \tau_*/n$, where b_n is the scalar optimum off diagonal. The resulting off-diagonals are exactly b_n, the row cap holds for large n, and $\operatorname{tr}\Omega \to \tau_*$ while $\|\Omega\|_{\mathrm{F}} \to 0$.

18 The Frobenius covariance chart

By Section 17, near-maximizers admit gauged longitudinal loadings $\ell = c_n\mathbf{1} + th$, where ℓ is real positive, $c_n = \sqrt{nb_n} = \sqrt{2} + O(1/n)$ is at the scalar optimum, h is real unit, and $t = o(1)$. The transverse Gram perturbation is αW, $\alpha = o(1)$ its Frobenius norm, $W = VV^* \succeq 0$, $\|W\|_{\mathrm{F}} = 1$. When $\alpha = 0$, use $W = 0$ in the chart. The strictly negative longitudinal Hessian below gives the scalar comparison. The original transverse directions are centered before the longitudinal gauge. After making ℓ_i positive, the gauge phases are $o(1)$ in Euclidean norm by the strong longitudinal convergence, so $\|V^*\mathbf{1}\| = o(1)$ from $\|V\| \leq 1$.

We work uniformly for all real unit h, all such W with $\|V^*\mathbf{1}\| \leq 1$, small real t, s, and $\alpha = s^2$. Use scalar radial variable $\tau = T/n$ with T originally standard exponential, and independent circular complex Gaussian field z with covariance W. Put

$$L_i = W_{ii}, \quad m_i = (c_n + th_i)\sqrt{\tau} + sz_i, \quad a_i = 1 - (c_n + th_i)^2/n - s^2 L_i, \quad d_i = a_i + |m_i|^2.$$

Then $a_i \geq 1/2$, $\sum_i L_i^2 \leq 1$, $\sum_i L_i \leq \sqrt{n}$. The normalized cofactor matrix is $M = \mathbb{E}[D + vv^*]$ under the tilt by $\prod_i d_i$, with $D_{ii} = a_i/d_i^2$, $v_i = m_i/d_i$. We prove all derivatives through order six in (t, s) of $\Lambda - n/4$ are uniformly bounded. Evenness in s then gives ordinary C^3 uniformity in (t, α) for $\alpha \geq 0$, by the even differentiation formula proved below.

18.1 Gaussian moment estimates

We use Gaussian rotation and its iteration (cf. [19]). Let $z_{\mathbb{R}} = \Gamma^{1/2}g$, g real standard Gaussian, $\Gamma = \mathcal{R}(W)/2$. For a sum $F = \sum_i f_i(z_i)$ with real or complex local functions, the Gaussian rotation estimate is

$$\|F - \mathbb{E}F\|_p \leq C\sqrt{p}\,\big\|\|\nabla_g F\|\big\|_p \quad (p \geq 2).$$

Indeed use an independent copy g', Jensen, and the path $g_\theta = g\cos\theta + g'\sin\theta$, $0 \leq \theta \leq \pi/2$; its derivative is independent standard Gaussian for each θ. Minkowski and the conditional Gaussian moment give the estimate. The same proof for a Euclidean-vector-valued function bounds the right side by $C\sqrt{p}$ times the L^p Frobenius norm of its Jacobian, since a matrix applied to a standard

Gaussian has L^p Euclidean norm at most $C\sqrt{p}$ times its Frobenius norm. Applying this to $\nabla_g F$ gives

$$\|F - \mathbb{E}F\|_p \le C\sqrt{p}\,\|\mathbb{E}\nabla_g F\| + Cp\,\big\|\|D_g^2 F\|_{\mathrm{F}}\big\|_p. \tag{18}$$

Apply (18) to the additive sums of $\log d_i$ and v_i, and their t, s derivatives through order six. Here are uniform bounds on both terms, for fixed p, polynomial in $1 + \sqrt{\tau}$. The Hessian is $\Gamma^{1/2}\,\mathrm{diag}_{\mathrm{block}}(D_i')\Gamma^{1/2}$, where each 2-by-2 block D_i' is the real-coordinate local Hessian. Each block has a uniform L^p bound of the required polynomial size because $d_i \ge 1/2$, $L_i \le 1$, $|h_i| \le 1$, and the local expressions and their finitely many derivatives are rational with polynomial numerators. For real components the squared Frobenius norm of the Hessian is at most

$$\sum_{ij} \|D_i'\|\,\|D_j'\|\,\|\Gamma_{ij}\|_{\mathrm{F}}^2,$$

by the block trace formula. Apply Minkowski and Hölder using $\sum_{ij}\|\Gamma_{ij}\|_{\mathrm{F}}^2 = \|\Gamma\|_{\mathrm{F}}^2 \le 1$. This avoids a maximum over coordinate magnitudes.

For the mean gradient write it as $\Gamma^{1/2}$ times the array of the expected local gradients in z_i. Compare that array to the same deterministic gradient at the scalar baseline $h_i = 0, L_i = 0, z_i = 0$. The local means differ by at most $C_p(1 + \sqrt{\tau})^k(|h_i| + L_i)$. For the deterministic parameters use smoothness and $\ell_i - c_n = th_i$; for the Gaussian evaluation Taylor expand the local gradient at $z_i = 0$, use its mean-zero linear term and the expected quadratic remainder $O(L_i)$, with polynomially bounded higher factors. This remains valid for the parameter derivatives: a positive t-order contains factors h_i, and s-derivatives use the same local Gaussian variable and L_i through $a_i = 1 - \ell_i^2/n - s^2 L_i$. The difference array has bounded Euclidean norm by $\|h\|_2 = 1, \|L\|_2 \le 1$. Multiplying a constant two-dimensional array by $\Gamma^{1/2}$ is bounded because $\|V^*\mathbf{1}\| \le 1$. This proves the asserted polynomial moment bounds. All sufficiently many local derivatives used in this comparison exist uniformly since $a_i \ge 1/2$.

For the tilt itself we also need exponential, not only fixed polynomial, moments. For $F = \sum \log d_i$, its Hessian in g has operator norm at most Cs^2 uniformly, because the second derivative of $\log(a + |m|^2)$ in m is uniformly bounded. The variance estimate for its gradient from (18)'s vector predecessor implies $(\mathbb{E}\|\nabla_g F\|^2)^{1/2} \le C(1 + \sqrt{\tau})$. In this bound the expected scalar local gradient itself is uniformly bounded by Cs; its deviation from the constant for h_i, z_i, L_i can be bounded in Euclidean norm by $C(1 + \sqrt{\tau})$ using the Lipschitz first derivative and Gaussian second-order Taylor with uniformly bounded derivatives. The Gaussian rotation estimate for the scalar gradient norm, which is Cs^2-Lipschitz, gives

$$\|F - \mathbb{E}F\|_p \le C(1 + \sqrt{\tau})\sqrt{p} + Cs^2 p, \qquad p \ge 2. \tag{19}$$

For a prescribed fixed $q \ge 1$, this implies $\mathbb{E}_g e^{q(F - \mathbb{E}_g F)} \le C_q e^{C_q \tau}$ on a small fixed amplitude interval. To see this, optimizing Markov over integers $p \ge 2$ gives a tail bounded by a constant times the sum of a Gaussian tail $\exp[-cx^2/(1 + \tau)]$ and an exponential tail $\exp[-cx/s^2]$, allowing a larger constant for small x. Integrate e^{qx} against those tails and choose $s^2 < c/(2q)$. This proof of (19) is independent of rank.

18.2 Tilt normalization and radial anchoring

Use a fixed scalar reference at c_n, with $a_* = 1 - c_n^2/n$, $m_* = c_n\sqrt{\tau}$, $d_* = a_* + |m_*|^2$; denote by ν the reference radial tilted law times the original Gaussian field. Let $F_{\mathrm{rel}} = \sum \log d_i - n \log d_*$. Anchor it at its Gaussian expectation at $\tau_0 = 1/2$. This preserves the normalized tilt. For $\tau \in [1/4, 3/4]$, the difference of its Gaussian mean from the anchor and its parameter derivatives through six are

$O(1+|X|)$, $X=\sqrt{n}(\tau-1/2)$. For other τ they are bounded by a polynomial in $1+|X|, 1+\sqrt{\tau}$. Here is a direct verification of the mean estimates. For a smooth local observable $f(a,m)$, define

$$\mathcal{T}_q f(a,m)=\mathbb{E}_{G\sim CN(0,1)} f(a-q, m+\sqrt{q}\,G), \qquad \mathcal{L}=-\partial_a+\tfrac{1}{4}\Delta_m.$$

For $q=s^2L_i$ and $a=1-\ell_i^2/n$, the expectation of the local observable is $\mathcal{T}_q f(a,\ell_i\sqrt{\tau})$. On the chart $a-q\geq 1/2$, $\partial_q^j\mathcal{T}_q f=\mathcal{T}_q\mathcal{L}^j f$ by Gaussian heat differentiation, including at zero (cf. [21]). For $f=\log(a+|m|^2)$, the difference from $q=0$ and all its needed amplitude derivatives are bounded by CL_i times a polynomial in $1+\sqrt{\tau}$; it follows from $q=s^2L_i$, the fundamental theorem and the local rational derivatives. The same bound applies after a bounded number of longitudinal derivatives and one radial derivative on the central interval. The baseline longitudinal scalar difference from c_n costs $C|h_i|$ times such a polynomial, and this statement also covers its positive-order derivatives. Thus the total raw mean differences and their derivatives are bounded by $C\sqrt{n}$ on the central interval with radial derivative, and by $C\sqrt{n}$ times a polynomial globally. This proves the anchored estimates by the mean value theorem centrally, and by $|X|\geq\sqrt{n}/4$ on the complement.

For the undifferentiated mean, a sharper global bound $C\sqrt{n}(1+\sqrt{\tau})$ suffices for exponential integrability outside the central interval. The log function has uniformly bounded first derivatives in m,a and $\mathcal{L}\log(a+|m|^2)$ is uniformly bounded when $a\geq 1/2$. This proves the sharper bound by the same rowwise comparison. The centered conditional exponential moment from (19), the central bound $C|X|$, and this global bound show the anchored likelihood has bounded $L^q(\nu)$ for a prescribed $q>1$, on a fixed small amplitude interval. Indeed, for the scalar τ law $T+(n/c_n^2-1)$ has the Gamma$(n+1,1)$ law conditioned to exceed n/c_n^2-1. It has bounded exponential moments in any fixed multiple of $|X|$ for sufficiently large n; outside the central interval a factor $\exp[C_q\sqrt{n}(1+\sqrt{\tau})+C_q\tau]$ is absorbed by its exponential tails. The tilted normalization is bounded below by Jensen on $|X|\leq 1$, using the Gaussian-mean anchored bound there. The normalized log likelihood derivatives have bounded moments after using (18) and the anchored mean estimates, with polynomial radial factors. Hence the quotient differentiation rule preserves uniform bounds on expectation derivatives of observables with such L^p estimates. Differentiation under the integral for each n follows directly by differentiating the finite polynomial likelihood and the rational observables with $d_i\geq 1/2$ on a slightly larger parameter rectangle.

18.3 Trace compensation in the cofactor Rayleigh quotient

Put $v_*=m_*/d_*$ and $S=\sum_i(v_i-v_*)$. Its centered conditional Gaussian fluctuations, through six parameter derivatives, have already been bounded by (18). It also has rank-independent conditionally expected size and derivatives polynomial in $1+|X|, 1+\sqrt{\tau}$. For the mean estimate put $a_{0,i}=1-\ell_i^2/n$. The scalar baseline difference $\sum_i[v(a_{0,i},\ell_i\sqrt{\tau})-v_*]$ has uniform bound of this size: Taylor in ℓ_i about c_n, using the first derivative

$$\frac{d}{d\ell}\frac{\ell\sqrt{\tau}}{1+\ell^2(\tau-1/n)}=\frac{\sqrt{\tau}\,[1-\ell^2(\tau-1/n)]}{[1+\ell^2(\tau-1/n)]^2}.$$

At c_n this is $O((1+\sqrt{\tau})(|\tau-1/2|+1/n))$. The difference at order one and the base linear term sum with $\|h\|_1\leq\sqrt{n}$; second and higher longitudinal terms sum with $\sum_i h_i^2=1$.

For the covariance difference, use $\mathcal{T}_q v=v+q\mathcal{L}v+O(q^2)$ and the exact cancellation

$$\mathcal{L}v(a,m)=\frac{m(|m|^2-a)}{(a+|m|^2)^3}.$$

Indeed $\frac{1}{4}\Delta_m(m/d) = -2am/d^3$ whereas $-\partial_a(m/d) = m/d^2$. The coefficient vanishes at $a = m = 1$. Thus summing $q\mathcal{L}v$ costs at most a polynomial times

$$\sum_i L_i \left(|h_i| + |\tau - 1/2| + 1/n\right) \le 1 + |X| + n^{-1/2}.$$

Positive longitudinal derivatives of this term use $\sum_i L_i|h_i| \le 1$, and amplitude derivatives retain the compensated coefficient. The $O(q^2)$ remainder and its mixed t, s derivatives sum with $\sum_i L_i^2 \le 1$. More explicitly, write it as $q^2 \int_0^1 (1-\vartheta)\mathcal{T}_{\vartheta q}\mathcal{L}^2 v \, d\vartheta$. Differentiating $q = s^2 L_i$ and this smooth expression at fixed t, s leaves at least L_i^2 in every term, with bounded heat derivatives. This proves all the required estimates for S.

Let $e = \mathbf{1}/\sqrt{n}$, $P = I - ee^*$. For any deterministic unit b the scalar b^*Pv and its parameter derivatives through six have bounded conditional L^p norms polynomial in $1 + \sqrt{\tau}$, uniformly in b. Indeed, b^*P is contracted and annihilates the constant, so subtract v_*. Taylor at $z_i = 0$ with a_i fixed. The array of deterministic baseline differences has bounded Euclidean norm by $\|h\|, \|L\|$. The terms linear in $z_i, \overline{z}_i$ are tested against bounded coefficient Euclidean vectors, so their variance is bounded since $\|W\| \le 1$. The quadratic remainders tested against b have bounded L^p norm by $\sum_i |b_i'|L_i \le 1$ and the uniform local Gaussian moments. The same reasoning applies to all parameter derivatives: take the z Taylor expansion of the local differentiated function, whose coefficients have uniform polynomial bounds and whose diagonal deterministic differences contain either the removed scalar baseline or a factor h_i or L_i.

The exact Rayleigh expression is

$$e^*Me = \mathbb{E}\left[n|v_*|^2 + 2\,\mathrm{Re}(\overline{v}_* S) + |S|^2/n + n^{-1}\sum_i D_{ii}\right].$$

The baseline $n|v_*|^2 - n/4$ is $O(1 + X^2)$ from its exact scalar defect, and the terms involving S are bounded with all expectation derivatives through six. For D, arbitrary unit-vector tests give uniform L^p bounds of operator bilinear forms by the local rational estimates and Cauchy–Schwarz for the deterministic test vectors. For $R = PMP$, use the uniform estimates for scalar tests of Pv twice and Hölder. For $q = PMe$, use a unit test against Pv and $e^*v = \sqrt{n}v_* + S/\sqrt{n}$. Consequently R is $O(1)$ and q is $O(\sqrt{n})$, through six parameter derivatives in operator or Euclidean norm respectively.

The Schur complement expresses the simple top eigenvalue as $e^*Me + q^*(\lambda I - R)^{-1}q$. The correction $q^*(\lambda I - R)^{-1}q$ and its fixed-λ parameter derivatives are bounded, and each λ-derivative gives a further $O(n^{-1})$. Successive implicit derivatives are bounded through order six since the coefficient of the highest derivative of λ is $1 + q^*(\lambda I - R)^{-2}q \ge 1$. The eigenvalue is even in s since the centered field is symmetric. If f is smooth and even, $\partial_\alpha f(\sqrt{\alpha}) = \frac{1}{2}\int_0^1 f''(a\sqrt{\alpha})\,da$. Iterating this formula adds bounded powers of a, so six mixed amplitude derivatives yield three mixed covariance derivatives including at zero. In particular the required C^3 covariance estimates hold with W normalized in Frobenius norm, independently of rank.

19 Affine Hessian comparison

The tangents at the scalar optimum are the real constant/star matrix $X_{ij} = c_n(h_i + h_j)/n$ and $Y = W - \mathrm{diag}(L_i)$. The affine Hessian Theorem 1 is negative semidefinite with imaginary gauge kernel. Its constant and star sectors have negative coefficients of order n, and the other transverse sectors have negative coefficients of order one. Since $\|Y\mathbf{1}\| \le \|VV^*\mathbf{1}\| + \|L\| \le 2$, the projections

on the constant, real-star, and imaginary-gradient spaces are $O(n^{-1/2})$ in Frobenius norm. Indeed these projections are determined by the real and imaginary row sums. The row-sum map multiplies Frobenius norm by $\sqrt{n-1}$ on constant real matrices, $\sqrt{(n-2)/2}$ on real stars, and $\sqrt{n/2}$ on imaginary gradients, as is seen by summing the explicit matrices in these sectors. Therefore the norm of Y with the constant and real-star squared weights multiplied by n, and the ordinary weights on the zero-row-sum sectors, stays bounded. Also X lies in the constant and real-star spaces, with $n\|X\|_{\mathrm{F}}^2$ bounded above and below by positive constants. If $\|\operatorname{diag} W\|_{\mathrm{F}} \le 1/4$, then the normal part of Y is bounded below, since $\|W\|_{\mathrm{F}} = 1$. The negative of the pulled-back Hessian controls $t^2 + \alpha^2$. First its normal term controls α^2, since X has zero normal component. The weighted constant and star terms then control t^2 by the triangle inequality and the bound on the weighted norm of Y. The uniform C^3 estimate makes every sufficiently small nonzero such displacement lose. Without this bound use the strong longitudinal curvature to exclude $|t| \ge C\alpha$ locally and negative semidefiniteness to give possible gain $O(\alpha^3)$ on the remaining $|t| < C\alpha$, by the following Taylor estimate. The full first derivative vanishes by permutation and conjugation symmetry and scalar stationarity. If $|t|$ is larger than a sufficiently large constant times α, the negative longitudinal quadratic dominates the cross term, the α^2 term, and the cubic remainder. Otherwise the entire quadratic is nonpositive and the remainder is $O(\alpha^3)$. By the original row cap, the diagonal of the unnormalized transverse covariance has Frobenius norm at most $B^2/\sqrt{n}$. Therefore a possible gain can only occur for $\alpha \le 4B^2/\sqrt{n}$, and is at most $O_B(n^{-3/2})$.

The feasible class is compact by using latent dimension n. The exact scalar optimum is feasible for sufficiently large n, also at $B = \sqrt{2}$, since $nb_n = 2 - 2/(n+1) + O(ne^{-cn}) < 2$ by Proposition 2. Thus its maximizing sequences satisfy the rank-independent localization and enter this chart, with its uniform neighborhood and derivative constants. The scalar expansion Proposition 2 completes the proof.

The argument also interpolates the upper error as a function of rank. If the transverse correction has rank at most $k \ge 1$, then $\operatorname{tr} W \le \sqrt{k}$ and

$$\|\operatorname{diag} W\|_{\mathrm{F}}^2 \le (\max_i L_i) \operatorname{tr} W \le B^2\sqrt{k}/(n\alpha).$$

Hence possible excess requires $\alpha \le 16B^2\sqrt{k}/n$ as well as $\alpha \le 4B^2/\sqrt{n}$. Its upper bound is $O_B(\min\{\sqrt{k}/n, n^{-1/2}\}^3)$ uniformly in k. The fixed-rank $O(n^{-3})$ bound follows, while Theorem 4 also gives an explicit constant-order formula for centered fixed-gap arrays.

20 Walsh families: the model and the cofactor identity

For a correlation matrix B define $C(B)_{ij} = \operatorname{per} B(i,j)$, with row i and column j deleted, and put $\Lambda(B) = \lambda_{\max}(C(B))/\operatorname{per} B$. This is the raw, unweighted cofactor convention of [7] and Theorem 1. The unrestricted bounds in [7] agree at orders n and one and leave a gap in the coefficient of $1/n$. The local result of Theorem 1 controls all directions near the scalar optimizer. Here we obtain exact maximization in a fixed finite-factor class by treating both its near-optimal limits and its local maximization in strength coordinates.

The raw cofactor bounds of Marcus and Merris [5] and Lavi [3] provide precedents for the spectral observable. Complex Gaussian permanent moments are classical [8], and Gaussian optical photon counts also involve positive semidefinite permanents [9]. McCullagh [10, Section 3] analyzes scalar permanent approximations for matrices with two blocks and shrinking between-block entries. We use the Gaussian law tilted by the permanent to represent the normalized raw cofactor matrix. Its phase covariance gives the limiting cofactor spectrum when the two blocks or more blocks approach decoupling.

Let $G = (\mathbb{Z}/2\mathbb{Z})^k$ be fixed, of size m, and assume n is a multiple of m. Index the rows by equal fibers over $x \in G$. Each character χ gives the sign vector s_χ on these rows and $s_0 = \mathbf{1}$ for the trivial character. We study

$$B_n = a_n I + \frac{1}{n}\sum_\chi \beta_{\chi,n} s_\chi s_\chi^T, \qquad a_n = 1 - \frac{1}{n}\sum_\chi \beta_{\chi,n}, \qquad \beta_{\chi,n} \geq 0.$$

The strengths are uniformly bounded, so $a_n > 0$ for all sufficiently large n. Character gauges are conjugations by diagonal matrices with a character's signs on the fibers. They permute the strength labels and preserve Λ.

Take independent standard circular complex Gaussians G_χ with $\mathbb{E}|G_\chi|^2 = 1$. Set

$$z_x = n^{-1/2}\sum_\chi \sqrt{\beta_{\chi,n}}\chi(x)G_\chi, \qquad d_x = a_n + |z_x|^2.$$

The row variables $X_i = \sqrt{a_n}\epsilon_i + z_{x_i}$, with independent standard circular ϵ_i, have covariance B_n. Wick's formula and conditional integration give $\operatorname{per} B_n = \mathbb{E}\prod_x d_x^{n/m}$. Write $\mathbb{E}_{\mathrm{t}}$ for expectation after tilting the latent Gaussian law by this product. For the raw cofactor the unpaired variables at $i \neq j$ are $\overline{X}_i X_j$ and its diagonal expectation contains $1/d_{x_i}$. It follows that

$$\frac{C(B_n)}{\operatorname{per} B_n} = \mathbb{E}_{\mathrm{t}}\big[\operatorname{diag}(a_n/d_{x_i}^2) + vv^*\big], \qquad v_i = \overline{z}_{x_i}/d_{x_i}. \tag{20}$$

Indeed the rank-one term supplies $\bar{z}_{x_i}z_{x_j}/(d_{x_i}d_{x_j})$ off the diagonal, and the diagonal completes it to $1/d_{x_i}$.

Group translations and within-fiber permutations commute with the cofactor. The fiber-constant space is diagonalized by the characters. Its orthogonal complement sees a diagonal operator of norm $O(1)$ by (20). Write $\langle f\rangle_G = m^{-1}\sum_x f_x$ for a fiber average. The uniform-mode eigenvalue has the exact defect identity

$$\mathbb{E}_{\mathrm{t}}\left\langle a_n/d_x^2\right\rangle_G + \frac{n}{4a_n} - n\mathbb{E}_{\mathrm{t}}\left\langle \frac{(|z_x|^2 - a_n)^2}{4a_n d_x^2}\right\rangle_G - n\mathbb{E}_{\mathrm{t}}\left\langle |v_x - \langle v\rangle_G|^2\right\rangle_G. \tag{21}$$

This follows by subtracting the fiber variance from $\left\langle |v|^2\right\rangle_G$ and using

$$\frac{u}{(a+u)^2} = \frac{1}{4a} - \frac{(u-a)^2}{4a(a+u)^2}.$$

21 A dominant character near strength two

Theorem 7. *Suppose $\beta_{0,n} = 2 + u/\sqrt{n} + o(n^{-1/2})$ and $\beta_{\chi,n} \to \kappa_\chi < 2$ for all $\chi \neq 0$. Then*

$$\Lambda(B_n) = \frac{n}{4} + \frac{1}{2} - \frac{u^2}{16} - \sum_{\chi\neq 0}\frac{\kappa_\chi^2}{8(2-\kappa_\chi)} + o(1).$$

Proof. Rotate by the phase of G_0, let $T = |G_0|^2$, and use the independent rotated secondary Gaussians $H_\chi = X_\chi + iY_\chi$. Write $\beta = \beta_{0,n}, \kappa_{\chi,n} = \beta_{\chi,n}$ and

$$g_0 = \sqrt{\beta T/n}, \quad h_x = \sum_{\chi\neq 0}\sqrt{\kappa_{\chi,n}}\chi(x)H_\chi, \quad z_x = g_0 + n^{-1/2}h_x, \qquad X = (T - n(1-1/\beta))/\sqrt{n}.$$

The common rotation leaves all observables of (20) unchanged. The tilted variables X, X_χ, Y_χ converge with polynomial moments to independent centered Gaussians of variances $1, 1/2, 1/(2-\kappa_\chi)$, respectively. For domination, AM–GM and orthogonality bound the tilted product by $(a_n + [\beta T + \sum \kappa_{\chi,n}|H_\chi|^2]/n)^n$. With $\widetilde{T} = T + \sum(\kappa_{\chi,n}/\beta)|H_\chi|^2$, the Gaussian densities leave a radial factor $e^{-\widetilde{T}}(a_n + \beta\widetilde{T}/n)^n$ and secondary factors $e^{-(1-\kappa_{\chi,n}/\beta)|H_\chi|^2}$. These rates have a uniform positive lower bound. The radial factor is a translated Gamma density of shape $n+1$, with lower cutoff at na_n/β before translation and with center differing by $O(1)$ from $n(1-1/\beta)$ in $\widetilde{T}$. Its rescaled centered moments are bounded by the Gamma moment generating function. The shift from $\widetilde{T}$ to T is controlled by the secondary Gaussian factors. This envelope is an upper bound in integral of the same order as the true normalizer, since the local expansion below gives a comparable lower mass on a fixed compact fluctuation box.

With $q = a_n + \beta T/n = \beta + \beta X/\sqrt{n} + O(n^{-1})$, the radial log density has limit $-X^2/2$, up to its constant. In the fiber log expansion use $d_x = q + 2g_0 \operatorname{Re} h_x/\sqrt{n} + |h_x|^2/n$. Orthogonality cancels the linear term in $\operatorname{Re} h_x$. The two order-one terms give

$$\tfrac{1}{2}\sum_\chi \kappa_\chi (X_\chi^2 + Y_\chi^2) - \tfrac{1}{2}\sum_\chi \kappa_\chi X_\chi^2.$$

With the original Gaussian costs this is precisely the asserted covariance. The expansions are uniform on compact fluctuation boxes, and domination as above gives tightness and uniform integrability of polynomial moments.

On the subspace orthogonal to the uniform vector, the rank expectation in (20) has trace $n\mathbb{E}_{\mathrm{t}} \left\langle |v_x - \langle v\rangle_G|^2 \right\rangle_G = O(1)$. This follows from the uniformly Lipschitz map $z \mapsto \bar{z}/(a_n + |z|^2)$ and the bounded moments of h_x. The uniform-mode eigenvalue is asymptotic to $n/4$ and is therefore the top eigenvalue for large n.

In (21) the average diagonal tends to $1/4$ and $n/(4a_n) = n/4 + (2 + \sum \kappa_\chi)/4 + o(1)$. The first deficit tends to $(u^2 + 4 + 2\sum \kappa_\chi)/16$, since $\sqrt{n}(|z_x|^2 - a_n)$ tends to $u + 2X + 2\operatorname{Re} h_x$ in moments and $d_x \to 2$. For the second deficit, the rescaled centered v_x tends to $-i \operatorname{Im} h_x/2$. This is the derivative at $z = 1, a = 1$, which vanishes in the real direction. Orthogonality and the limiting imaginary variances give the second deficit $\sum \kappa_\chi/[4(2-\kappa_\chi)]$. The positive denominators and moment bounds justify taking expectations in these expressions. Substituting the two deficits gives the formula. □

22 Global rigidity and exact maximization

Theorem 8. *Fix the Walsh group and any cap $B > 2$. For all sufficiently large n divisible by m, the maximum of $\Lambda(B_n)$ over $0 \le \beta_{\chi,n} \le B$ equals the scalar equicorrelation maximum. Every maximizer, after a character gauge, has only the trivial strength nonzero and is the scalar maximizer. In particular, the maximum is $(n+1)^2/(4n) + O(ne^{-cn})$ for some $c > 0$.*

22.1 Localization of two-term near-maximizers

Suppose first that the strengths are uniformly bounded and $\Lambda(B_n) \ge n/4 + 1/2 - o(1)$. The top eigenvector must lie in the fiber-constant space, since its complement has bounded norm. Gauge a maximizing character to the trivial one. Then (21) equals $\Lambda(B_n)$. Its two nonnegative deficits are bounded, since $a_n = 1 - O(n^{-1})$ and the diagonal term is bounded. Hence in tilted probability, for each of the fixed number of fibers,

$$|z_x|^2 = 1 + O_{\mathbb{P}}(n^{-1/2}), \qquad z_x - \langle z\rangle_G = O_{\mathbb{P}}(n^{-1/2}). \tag{22}$$

Indeed the first defect controls $(|z_x|^2 - a_n)/d_x$ at this scale, which also controls $|z_x|^2 - a_n$ by solving for it. The second defect controls $v_x - \langle v \rangle_G$; then use $z_x = d_x \bar{v}_x$ and $d_x = 2 + O_{\mathbb{P}}(n^{-1/2})$.

In rotated latent coordinates as in Theorem 7, (22) implies $\beta T/n = 1 + O_{\mathbb{P}}(n^{-1/2})$ and $\sqrt{\kappa_{\chi,n}} H_\chi = O_{\mathbb{P}}(1)$. The scaled latent vector $(G_\chi/\sqrt{n})$ is tight for bounded strengths. With $R = n^{-1} \sum |G_\chi|^2$, the Gaussian radial cost and AM–GM give the exponential tail bound with exponent $n[-R + \log(1 + BR)]$, up to a fixed-degree radial volume factor. The permanent normalizer is bounded below by $a_n^n \geq c > 0$. In particular β cannot tend to zero along near-maximizers.

Pass to a subsequence on which the strengths converge, with trivial limit $\beta_* > 0$. The scaled latent vector concentrates on the maximizers of

$$-\sum_\chi |w_\chi|^2 + \frac{1}{m} \sum_x \log \left(1 + \left| \sum_\chi \sqrt{\beta_{\chi,*}} \chi(x) w_\chi \right|^2 \right).$$

This follows by uniform convergence of the exponents on compacts, the displayed tail bound, and integration of a small ball near a maximum for a lower bound. A limiting latent vector must give all z_x equal with modulus one by (22). Its positive-strength nontrivial coordinates vanish by character orthogonality, and the zero-strength coordinates vanish at a maximum since they only enter the negative quadratic term. At the trivial-axis maximizer $|w_0|^2 = 1/\beta_*$, so the radial derivative is $-1 + \beta_*/2$. Thus $\beta_* = 2$. No secondary limit can exceed two, since moving to that pure character at the same radius would then increase the exponent.

The secondary H_χ are tight when their strengths have positive limits, by (22). To see tightness when $\kappa_{\chi,n} \to 0$, condition on T and the other rotated variables. With arbitrarily high probability the trivial latent value $g_0 = \sqrt{\beta T/n}$ is in a fixed compact neighborhood of one, and all other $\sqrt{\kappa_{j,n}} H_j$ are bounded by a fixed M. Denote the field with χ omitted by $z_x^0 = g_0 + O(M/\sqrt{n})$, and write $d_x^0 = a_n + |z_x^0|^2$. The change in the log tilt is bounded above using $\log d_x - \log d_x^0 \leq (d_x - d_x^0)/d_x^0$. Its sum with multiplicities is at most $C_M \sqrt{\kappa_{\chi,n}} |H_\chi| + C \kappa_{\chi,n} |H_\chi|^2$, because the linear term in the constant value of $\bar{z}_x^0/d_x^0$ cancels by the zero character sum. On $|H_\chi| \leq 1$, Taylor expansion and the same cancellation give an absolute bound C_M. Thus the conditional normalizer has a positive lower bound after dividing by the value at $H_\chi = 0$, and the original Gaussian yields a common conditional Gaussian tail. This proves full secondary tightness.

Set $X_n = (T - n/\beta)/\sqrt{n}$, tight by (22). On any compact of its coordinates and the H_χ, the log of the unnormalized density divided by its central value has expansion

$$L_n X_n - \tfrac{1}{2} X_n^2 - \sum_\chi X_\chi^2 - \sum_\chi (1 - \kappa_\chi/2) Y_\chi^2 + o(1), \qquad L_n = \sqrt{n} \left(\frac{\beta}{a_n + 1} - 1 \right),$$

where κ_χ denotes the limiting secondary strengths on the subsequence. This is the same character cancellation as in Theorem 7, with radial center at n/β. If L_n were unbounded, moving any compact box a fixed distance in X_n in its sign direction would make the density on the shifted box arbitrarily larger than on the original, along a subsequence. This would force the original compact mass to vanish, contradicting tightness. Hence $\sqrt{n}(\beta - 2)$ is bounded. Pass further so that L_n converges. If some $\kappa_\chi = 2$, the local limiting unnormalized density is constant in Y_χ, so its full integral is infinite. Integrating over arbitrarily large compact boxes then makes the full normalizer divided by the central density diverge, forcing every fixed compact to have vanishing probability. This again contradicts tightness. Thus all $\kappa_\chi < 2$.

Theorem 7 now applies at every subsequential limit u of $\sqrt{n}(\beta - 2)$. Its nonnegative defect $u^2/16 + \sum_\chi \kappa_\chi^2/[8(2 - \kappa_\chi)]$ must be zero at a near-maximizer. We conclude $\sqrt{n}(\beta - 2) \to 0$ and $\max_{\chi \neq 0} \kappa_{\chi,n} \to 0$ after the chosen gauges.

22.2 A uniform strength neighborhood

Let $\beta_n^* = nb_n$, with b_n the exact scalar maximizer from Theorem 1, so $\beta_n^* = 2+O(n^{-1})$. Parametrize the neighborhood by $\beta = \beta_n^* + u/\sqrt{n}$ and $\kappa_{\chi,n} = \psi_\chi^2$, with u, ψ in fixed sufficiently small symmetric intervals. Let $F_n(u,\psi)$ be (21) minus $n/4$. On this rectangle the uniform mode is the top eigenvalue for large n, uniformly, by the moment and complement bounds of Theorem 7. We prove

$$F_n \;\longrightarrow\; \tfrac{1}{2} - \frac{u^2}{16} - \sum_\chi \frac{\psi_\chi^4}{8(2-\psi_\chi^2)} \quad \text{in } C^4(u,\psi).$$

Take the fixed radial center $T_0 = n(1-1/\beta_n^*)$, and $X = (T-T_0)/\sqrt{n}$. Put $q = a_n + \beta T/n$, $q_0 = a_n + \beta T_0/n$, and write $d_x = q + r_x$, where

$$r_x = 2g_0 \operatorname{Re} h_x/\sqrt{n} + |h_x|^2/n, \qquad h_x = \sum_\chi \psi_\chi \chi(x) H_\chi.$$

Divide the unnormalized density in (X,H) by its value at the center. Its log tilt terms decompose as

$$n\log(q/q_0) + n\langle r_x\rangle_G/q + n\langle \log(q+r_x) - \log q - r_x/q\rangle_G.$$

The first term has polynomially bounded derivatives through order four, uniformly on the rectangle. In fact $q/q_0 = 1 + (\beta/q_0)X/\sqrt{n}$ is bounded below and each positive-order parameter derivative of β/q_0 is $O(n^{-1/2})$. The second term uses $\langle r_x\rangle_G = n^{-1}\sum \psi_\chi^2 |H_\chi|^2$. In the third term use the integral Taylor remainder with numerator r_x^2 and denominators $(q+tr_x)^2$, $0 \le t \le 1$, all bounded below by a_n^2. The factors $\sqrt{n}\, r_x$ and their parameter derivatives are polynomially bounded in (X,H), since $T/n \le C(1+|X|)$ and derivatives of $g_0 = \sqrt{\beta T/n}$ differentiate only $\sqrt{\beta}$. This proves polynomial bounds for all differentiated density factors after taking out the density.

The diagonal contribution and the first defect in (21) have the same bounds after differentiation, since $d_x \ge 1/2$ and $\sqrt{n}(|z_x|^2 - a_n)$ is polynomially bounded with its derivatives. For the last defect, subtract the common value $v(g_0)$ before centering over the fibers. The expression $\sqrt{n}[v(z_x) - v(g_0)]$ is an integral of the real differential of v applied to h_x along the line from g_0 to z_x. The rational map $v(z) = \bar{z}/(a_n + |z|^2)$ and its fixed-order derivatives have polynomial bounds for $a_n \ge 1/2$, which also bounds all differentiated versions of this defect.

The preceding Gamma envelope is uniform on the rectangle. Its integral relative to the density value at the center is bounded in $dX\, dH$ by Stirling's formula, since the translated Gamma center is within $O(\sqrt{n})$ of T_0 and the T-to-X Jacobian is $\sqrt{n}$. Its bounded moments of all orders give uniform integrability of the polynomial derivative bounds. The true relative integral has a uniform positive lower bound on a compact fluctuation box by the local expansion. Differentiation of the normalized integrals is therefore justified, and convergence of the differentiated expansions on compact fluctuation boxes gives uniform convergence through order four. In these fixed-center coordinates, the radial limiting log tilt is $uX/4 - X^2/2$. After recentering X by $u/4$, the dominant-character calculation gives exactly the displayed limit for F_n.

The functions are separately even in every ψ_χ, by changing the sign of the corresponding H_χ. Thus C^4 convergence gives C^2 convergence on the closed strength rectangle $\kappa_\chi = \psi_\chi^2 \ge 0$. For a one-variable even function f, the first derivative of $f(\sqrt{\kappa})$ is $\frac{1}{2}\int_0^1 f''(t\sqrt{\kappa})dt$, and the second is $\frac{1}{4}\int_0^1\int_0^1 t^2 f^{(4)}(st\sqrt{\kappa})ds\, dt$, both continuous at zero. Apply the same formula in each variable for mixed derivatives. The limiting Hessian in the strength coordinates is uniformly negative definite on a small fixed rectangle, so the exact objective is strictly concave there for large n.

Its gradient at $u = 0, \kappa = 0$ is zero exactly. The scalar coordinate is stationary by the interior scalar optimum. At the same equicorrelation point, simplicity and permutation invariance make the

derivative on real symmetric zero-diagonal perturbations proportional to the sum of off-diagonal entries. Since the scalar perturbation has nonzero such sum and derivative zero, the derivative in every secondary direction is zero. Thus the unique strength maximizer in this neighborhood is the scalar point.

The capped family is compact and contains the scalar point for large n. Every maximizing sequence is two-term near-optimal; the localization above places it, after a gauge, inside the strength rectangle for sufficiently large n. The strict concavity then proves Theorem 8, including equality. The scalar value and its error follow from Theorem 1.

23 The critical quartic regime at strength one

The scale in this section concerns growth of the ratio, before the optimization near strength two. The scalar radial transition is related to incomplete Gamma asymptotics [11]. Quartic critical laws for vector spin systems [12] are another precedent for the limiting form of the exponent. Here the law controls the largest eigenvalue of a permanental cofactor.

Theorem 9. *If all $\beta_{\chi,n} \to b_\chi < 1$, then*

$$\Lambda(B_n) \longrightarrow \max_{\chi} (1 - b_\chi)^{-1}.$$

Suppose a nonempty character set J has $\beta_{\chi,n} = 1 + u_\chi/\sqrt{n} + o(n^{-1/2})$, with all other limits $b_\chi < 1$. Let $Y = (Y_\chi)_{\chi \in J}$ be the complex vector of density proportional to

$$\exp\left[\sum_{\chi \in J} u_\chi |Y_\chi|^2 - \frac{1}{2m} \sum_{x \in G} \left| \sum_{\chi \in J} \chi(x) Y_\chi \right|^4 \right].$$

Then

$$n^{-1/2} \Lambda(B_n) \longrightarrow \lambda_{\max}(\mathbb{E}[YY^*]) = \max_{\chi \in J} \mathbb{E}|Y_\chi|^2.$$

For one critical character of detuning u, the value is $u + \varphi(u)/\Phi(u)$, with φ, Φ the standard real normal density and distribution function.

Proof. The exact fiber identity

$$\frac{n}{m} \sum_x (d_x - 1) = \sum_\chi \beta_{\chi,n} (|G_\chi|^2 - 1)$$

follows from orthogonality. In the subcritical regime, the log tilt converges to $\sum b_\chi (|G_\chi|^2 - 1)$ at fixed G_χ. The bound $\log d_x \le d_x - 1$ gives a common Gaussian majorant with all moments, since the strengths are uniformly below one. The latent limits then have independent circular variances $(1 - b_\chi)^{-1}$. In (20) the diagonal tends to the identity. In the orthonormal fiber basis the rank expectation has matrix $(n/m)\mathbb{E}_{\mathrm{t}}[vv^*]$ using the m fiber values. Here $v_x = \bar{z}_x + o(n^{-1/2})$ in tilted L^2, so this rank contribution has limiting eigenvalues $b_\chi/(1 - b_\chi)$ by character orthogonality. The subcritical formula follows.

For critical characters set $G_\chi = n^{1/4} Y_\chi$ and leave the other Gaussians unscaled. At fixed rescaled variables, with $\psi_x = \sum_{\chi \in J} \chi(x) Y_\chi$, one has $n^{1/4} z_x \to \psi_x$. The exact linear sum cancels the order-$\sqrt{n}$ Gaussian weight of the critical variables and gives the detuning terms. The second log term tends to the quartic term in the statement. The unscaled modes have independent limiting Gaussian weights $e^{-(1-b_\chi)|G_\chi|^2}$.

For domination put $T = n^{-1/2}\sum_J \beta_{\chi,n}|Y_\chi|^2 + n^{-1}\sum_{\chi\notin J}\beta_{\chi,n}|G_\chi|^2$. AM–GM bounds the tilt by $(1+T)^n$. Use $\log(1+T) \le T - c\min(T^2, T)$ for $T \ge 0$. After the original Gaussian costs, the joint exponent has the upper bound

$$C\|Y\|^2 - c'\sum_{\chi\notin J}|G_\chi|^2 - c''\min(\|Y\|^4, \sqrt{n}\,\|Y\|^2).$$

This gives a common integrable Gaussian upper envelope for all sufficiently large n, also for any fixed polynomial moment. The local exponents converge uniformly on compact sets and the limit has positive normalizing mass there. The quartic term is coercive, since $m^{-1}\sum_x |\psi_x|^4 \ge \|Y\|^4$ by Jensen and orthogonality. Thus this is the joint limiting law.

After division by $\sqrt{n}$ the bounded diagonal term of (20) vanishes. Since $n^{1/4}v_x \to \bar\psi_x$ in tilted L^2, the normalized rank part on the fiber basis tends to $m^{-1}\mathbb{E}[\bar\psi\,\bar\psi^*]$. Its nonzero eigenvalues are those of the conjugate of $\mathbb{E}[YY^*]$ by the orthonormal Walsh columns. The covariance is diagonal in the character basis, by invariance under group translation of the fiber index in the quartic sum. For one critical character its squared modulus has density proportional to $e^{ut-t^2/2}$ on $t \ge 0$, giving the stated truncated normal mean. □

24 Degenerate supercritical strengths

Theorem 10. *Suppose all m strengths satisfy $\beta_{\chi,n} = \gamma + \tau_\chi/n + o(n^{-1})$ with $\gamma > 1$. Put $W_{x,\chi} = \chi(x)/\sqrt{m}$ and $Q = W\,\mathrm{diag}(\tau_\chi)W^*$. Let $\Xi = (e^{i\theta_x})_{x\in G}$ have density proportional to*

$$\exp\left[\frac{\gamma-1}{m\gamma^2}\Xi^* Q\Xi\right]$$

with respect to uniform product angles on the torus. Then

$$\frac{\Lambda(B_n)}{n} \longrightarrow \frac{\gamma-1}{m\gamma^2}\lambda_{\max}(\mathbb{E}[\Xi\Xi^*]).$$

If $m > 1$ and all detunings are finite, this value is strictly below $(\gamma-1)/\gamma^2$. At equal detunings it is $(\gamma-1)/(m\gamma^2)$. If $m = 2$, it is

$$\frac{\gamma-1}{2\gamma^2}\left(1 + \left|\frac{I_1(\mu)}{I_0(\mu)}\right|\right), \qquad \mu = \frac{\gamma-1}{2\gamma^2}(\tau_0 - \tau_1),$$

where I_0, I_1 are modified Bessel functions.

Proof. In the latent fiber coordinates $z = (z_x)$ the covariance is $(m/n)W\,\mathrm{diag}(\beta_{\chi,n})W^*$. Thus the tilted density is proportional to

$$\exp[-(n/m)z^*W\,\mathrm{diag}(\beta_{\chi,n}^{-1})W^*z]\prod_x (a_n + |z_x|^2)^{n/m}.$$

At reference strength γ, the squared radii have independent translated Gamma densities with rate $n/(m\gamma)$, restricted to positive $t_x = |z_x|^2$. Their centers tend to $\gamma - 1 > 0$. The angles are uniform and independent. Expanding the inverse strengths multiplies this density by $\exp[z^*Qz/(m\gamma^2) + o(1)\|z\|^2]$. The radial reference variables have uniformly bounded exponential moments at each fixed order, from the Gamma formula and the separated lower cutoff. This controls the tilt, and

radial concentration gives the displayed limiting phase density, with convergence of the moments used in the cofactor identity. Phase models of this form are studied in [13].

On the fiber basis the rank term in (20) is $(n/m)\mathbb{E}_{\mathrm{t}}[vv^*]$. Its n^{-1} limit has the spectrum asserted because v_x tends to $\sqrt{\gamma-1}\,\overline{\Xi_x}/\gamma$ in tilted expectations. The diagonal part is bounded. The phase covariance has trace m and is positive definite under the everywhere positive torus density. Thus for $m>1$ its top eigenvalue is strictly below m. Equal detunings give a scalar Q and independent uniform angles. For two fibers the difference angle has density proportional to $e^{\mu\cos\theta}$, giving the stated I_1/I_0 Fourier moment. □

Theorems 8–10 concern different strength regimes in the same fixed Walsh family. In particular Theorem 8 determines the third term only for a fixed bounded-rank correction with balanced Walsh fibers. The unrestricted cofactor extremum remains open.

Acknowledgement

OpenAI Codex and Anthropic Claude assisted with proof development, independent checking, numerical verification and manuscript preparation. The author is responsible for the work.

References

[1] R. Bhatia, P. Grover, and T. Jain, *Derivatives of tensor powers and their norms*, Electron. J. Linear Algebra **26** (2013), 604–619. doi:10.13001/1081-3810.1675.

[2] R. Grone, C. R. Johnson, E. M. Sá, and H. Wolkowicz, *A note on maximizing the permanent of a positive definite Hermitian matrix, given the eigenvalues*, Linear Multilinear Algebra **19** (1986), 389–393. Author copy.

[3] Y. Lavi, *A conjecture on the permanent of positive semidefinite matrices*, 2026. Preprint, version 1.

[4] E. H. Lieb, *Proofs of some conjectures on permanents*, J. Math. Mech. **16** (1966), 127–134.

[5] M. Marcus and R. Merris, *A relation between the permanental and determinantal adjoints*, J. Austral. Math. Soc. **15** (1973), 270–271. doi:10.1017/S1446788700013173.

[6] L. Pioge, K. K. Pietrasz, B. Seron, L. Novo, and N. J. Cerf, *A logical implication between two conjectures on matrix permanents*, Linear Algebra Appl. **725** (2025), 309–318. arXiv:2508.00111.

[7] S. Schubr, *Asymptotically optimal bounds for permanental cofactor matrices*, Luminary Research Archive, ed. 4, 2026. Archive edition 4.

[8] C. Fassino, G. Pistone, and M. P. Rogantin, *Computing the moments of the complex Gaussian: Full and sparse covariance matrix*, Mathematics **7** (2019), 263. doi:10.3390/math7030263.

[9] S. Rahimi-Keshari, A. P. Lund, and T. C. Ralph, *What can quantum optics say about computational complexity theory?*, Phys. Rev. Lett. **114** (2015), 060501. arXiv:1408.3712.

[10] P. McCullagh, *An asymptotic approximation for the permanent of a doubly stochastic matrix*, Author manuscript.

[11] N. M. Temme, *Uniform asymptotic expansions of the incomplete Gamma functions and the incomplete Beta function*, Math. Comp. **29** (1975), 1109–1114. CWI archive.

[12] K. Kirkpatrick and T. Nawaz, *Asymptotics of mean-field $O(N)$ models*, arXiv:1602.03257.

[13] N. Klein, J. Orellana, S. Brincat, E. K. Miller, and R. E. Kass, *Torus graphs for multivariate phase coupling analysis*, arXiv:1910.11044.

[14] P. Feinsilver and J. McSorley. Zeons, permanents, the Johnson scheme, and generalized derangements. *International Journal of Combinatorics*, 2011, Article 539030. doi:10.1155/2011/539030.

[15] S. Sra. A Koteljanskii inequality for permanents. arXiv:2609.39979, 2026.

[16] A. Barvinok. Two algorithmic results for the traveling salesman problem. *Mathematics of Operations Research*, 21(1):65–84, 1996. doi:10.1287/moor.21.1.65.

[17] A. Barvinok. A remark on approximating permanents of positive definite matrices. *Linear Algebra and its Applications*, 608:399–406, 2021. arXiv:2005.06344.

[18] A. Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6:695–709, 2005. Journal article.

[19] R. Adamczak and P. Wolff. Concentration inequalities for non-Lipschitz functions with bounded derivatives of higher order. arXiv:1304.1826.

[20] L. Devroye, A. Mehrabian, and T. Reddad. The total variation distance between high-dimensional Gaussians with the same mean. arXiv:1810.08693.

[21] F. Voigtlaender. A general version of Price's theorem. *Journal of Theoretical Probability*, 34:1474–1485, 2021. doi:10.1007/s10959-020-01017-w.