

Ten Advances in Mathematics and Theoretical Computer Science

OpenAI

Abstract

We present a collection of results obtained by an internal OpenAI model, spanning mathematics and theoretical computer science:

1. **High-dimensional sphere packing.** The asymptotic strength of the Cohn–Elkies linear program is determined exactly. This gives an improved general packing bound in high dimensions and settles the corresponding Fourier sign-uncertainty problem asymptotically.
2. **Binary and spherical codes.** Classical upper bounds for fixed-distance binary and spherical codes are improved by exponential factors for all parameters. The spherical construction also recovers the sphere-packing exponent of Chapter 1.
3. **Non-sofic groups.** An explicit non-sofic group is constructed, resolving the question of whether every countable group admits finite permutation approximations. The argument uses property- (T) expanders and the binary Leavitt algebra.
4. **Connes’s rigidity conjecture.** Infinitely many pairwise nonisomorphic property- (T) groups are constructed with the same group von Neumann algebra, disproving Connes’s conjecture and answering a related finite-to-one question of Popa.
5. **Arithmetic circuit complexity.** For the permanent, division-free circuits require $\Omega(n^2 \log \log n)$ gates, while formulas require $\Omega(n^4 / \log n)$ leaves.
6. **Quantum parallel repetition.** Exponential parallel repetition is proved for every finite two-player entangled game, extending the classical repetition principle beyond previously treated special classes of quantum games.
7. **Closest vector problem.** A direct reduction from 3SAT gives $n^{1/400}$ -factor hardness for the Euclidean closest vector problem, with related consequences for binary decoding and other lattice norms.
8. **Ehrhart’s volume conjecture.** The sharp bound $(n + 1)^n / n!$ is proved in every dimension for convex bodies whose barycenter is their only interior lattice point.
9. **Multicolor Ramsey numbers.** A superexponential lower bound proves $R_k(3) = k^{\Theta(k)}$.
10. **Compactness and degeneracy.** Separate bipartite graph constructions disprove two conjectures in extremal graph theory: the compactness conjecture of Erdős and Simonovits and a degeneracy conjecture of Erdős.

CHAPTER 1

Exponential Growth Rate of the Cohn–Elkies Sphere Packing Linear Program

ABSTRACT. We determine the exact exponential decay rate of the Cohn–Elkies sphere-packing linear program. If Δ_d denotes the maximal sphere-packing density in $\mathbb{R}^d$ and LP_d denotes the optimal density bound furnished by this program, then

$$\limsup_{d \rightarrow \infty} \Delta_d^{1/d} \leq \lim_{d \rightarrow \infty} \text{LP}_d^{1/d} = \sqrt{\frac{e}{2\pi}}.$$

The positive- and negative-eigenvalue Fourier sign-uncertainty radii are both $(1/\pi + o(1))\sqrt{d}$.

CONTENTS

1. Introduction
 2. Fourier-analytic preliminaries
 3. The universal Cohn–Elkies lower bound
 4. The admissible primal upper bound
- Appendix A. Comparison of the sign-uncertainty constants
- References

1. INTRODUCTION

Sphere packing asks how densely congruent balls can fill Euclidean space. For centers separated by at least 1, let Δ_d be the supremum of the upper densities of radius-1/2 balls in $\mathbb{R}^d$. The Fourier-analytic linear-programming method of Gorbachev and Cohn–Elkies [Gor00, Coh02, CE03] bounds this density using an auxiliary function and its Fourier transform. Viazovska used this framework to prove that the E_8 lattice gives the optimal packing in dimension eight [Via17]; Cohn, Kumar, Miller, Radchenko, and Viazovska subsequently established the optimality of the Leech lattice in dimension twenty-four [CKMRV17]. Related Fourier interpolation formulas reveal further structure behind these exceptional configurations [RV19, CKMRV22].

By contrast, the behavior of the linear program in high dimensions remained poorly understood. Cohn and Zhao proved that it is always at least as strong as the Kabatianskii–Levenshtein spherical-code bound [CZ14], but whether it improves the classical high-dimensional sphere-packing exponent remained open. Motivated by the modular bootstrap [HMR19], Afkhami-Jeddi, Cohn, Hartman, de Laat, and Tajdini conjectured its high-dimensional rate [AJCHLT20, § 3].

We use the Fourier convention and unit-ball volume

$$\widehat{f}(\xi) = \int_{\mathbb{R}^d} f(x) e^{-2\pi i x \cdot \xi} dx, \quad v_d = \frac{\pi^{d/2}}{\Gamma(d/2 + 1)}. \quad (1)$$

A ball of radius 1/2 has volume $v_d/2^d$. For the real Schwartz space $\mathcal{S}(\mathbb{R}^d; \mathbb{R})$, define

$$\mathcal{A}_d = \left\{ f \in \mathcal{S}(\mathbb{R}^d; \mathbb{R}) : \widehat{f}(0) > 0, \quad \widehat{f} \geq 0 \text{ on } \mathbb{R}^d, \quad f \leq 0 \text{ on } \{|x| \geq 1\} \right\}, \quad (2)$$

and the associated linear-programming bound by

$$\text{LP}_d = \frac{v_d}{2^d} \inf_{f \in \mathcal{A}_d} \frac{f(0)}{\widehat{f}(0)}. \quad (3)$$

Fourier inversion gives $f(0) > 0$. The Gorbachev–Cohn–Elkies bound [Gor00, CE03] assigns every $f \in \mathcal{A}_d$ the density upper bound $v_d f(0)/(2^d \widehat{f}(0))$, so (3) yields

$$\Delta_d \leq \text{LP}_d. \quad (4)$$

Our main theorem proves the exponential rate conjectured in [AJCHLT20, Conjs. 3.1–3.2].

Theorem 1.1. *As $d \rightarrow \infty$, $\text{LP}_d^{1/d} \rightarrow \sqrt{e/(2\pi)}$.*

By (4), the theorem gives $\Delta_d \leq \text{LP}_d = 2^{-(\alpha_* + o(1))d}$, where $\alpha_* = \frac{1}{2} \log_2(2\pi/e) = 0.6044\dots$. This is the first improvement since 1978 to the general sphere-packing exponent. The classical Kabatianskii–Levenshtein exponent was 0.59905576\dots [KL78, Lev79]; subsequent spherical-code refinements had improved only lower-order factors [CZ14, SZ24, Z24]. The matching lower bound shows that no Cohn–Elkies auxiliary function can improve this exponent.

Our companion paper in Chapter 2 recovers the upper-bound direction of Theorem 1.1 as a small-distance limit of spherical-code bounds.

The packing sign conditions are related to the Bourgain–Clozel–Kahane uncertainty principle for eventually nonnegative Fourier eigenfunctions [BCK10, GOSS17, GOSR21]. Cohn and Gonçalves [CG19] introduced its anti-self-Fourier counterpart and connected it with sphere packing. If $g \in L^1(\mathbb{R}^d; \mathbb{R})$ satisfies $\widehat{g} = \varsigma g$ for $\varsigma \in \{-1, +1\}$, then $\widehat{g} \in L^1$ and Fourier inversion supplies a continuous representative of g . Using this representative for all pointwise values and sign conditions, define

$$r(g) = \inf\{R \geq 0 : g(x) \geq 0 \text{ for } |x| \geq R\}, \quad (5)$$

$$A_\varsigma(d) = \inf\{r(g) : 0 \neq g \in L^1(\mathbb{R}^d; \mathbb{R}), \widehat{g} = \varsigma g, g(0) = 0\}. \quad (6)$$

Here $r(g) = \infty$ when no such radius exists. The signs $+1$ and -1 give the original and complementary uncertainty problems, respectively.

The negative-eigenvalue problem also arises in the spinless modular bootstrap, where modular S -antisymmetry produces anti-self-Fourier test functions whose eventual sign controls the

spectral gap [HMR19, AJCHLT20]. General lower bounds for sign uncertainty and limitations of Gaussian–polynomial test functions of sublinear degree appear in [Edw25] and [CDG24, Thm. 1.2], respectively. The next theorem proves the asymptotic equality conjectured in [CG19, Conj. 1.5] and [AJCHLT20, (3.5)], with the common value predicted in [AJCHLT20, (3.3)]. Although these asymptotics coincide, Appendix A observes that $A_+(d) < A_-(d)$ for all d .

Theorem 1.2. *The Fourier-eigenfunction sign-uncertainty constants satisfy*

$$\lim_{d \rightarrow \infty} \frac{A_+(d)}{\sqrt{d}} = \lim_{d \rightarrow \infty} \frac{A_-(d)}{\sqrt{d}} = \frac{1}{\pi}.$$

The proof reduces both problems to the last sign changes of Fourier eigenfunctions. In § 3, Proposition 3.1 shows that a radial Schwartz function g satisfying $\widehat{g} = \pm g$ and $g(0) = 0$ has exponentially little L^1 mass inside $B(0, c\sqrt{d})$ whenever $c < 1/\pi$. For an admissible packing function F , let $a = (\widehat{F}(0)/F(0))^{1/d}$ and $h(x) = F(ax)$. Then $h(0) = \widehat{h}(0)$, while the nonzero function $g = \widehat{h} - h$ is anti-self-Fourier, vanishes at the origin, and is nonnegative outside $B(0, 1/a)$. Since $\int g = 0$, its negative part has mass $\|g\|_1/2$, all of which lies in this ball. The mass estimate therefore forces $1/a \geq (1/\pi - o(1))\sqrt{d}$. In § 4, Theorem 4.1 modifies the Mellin transform of a Gaussian to construct a Fourier pair and a self-Fourier function whose required exterior sign conditions begin at $(1/\pi + o(1))\sqrt{d}$. Stirling’s formula converts the matching radius bounds into the packing exponent $\sqrt{e/(2\pi)}$.

2. FOURIER-ANALYTIC PRELIMINARIES

We reduce the packing and sign-uncertainty problems to radial functions and establish the Mellin–Fourier identities used in both bounds. We repeatedly use the gamma recurrence and reflection formulas [DLMF, § 5.5], together with their consequences for real $b \neq 0$:

$$\begin{aligned} \Gamma(z+1) &= z\Gamma(z), & \Gamma(z)\Gamma(1-z) &= \frac{\pi}{\sin(\pi z)}, \\ |\Gamma(ib)|^2 &= \frac{\pi}{b \sinh(\pi b)}, & |\Gamma(1/2 + ib)|^2 &= \frac{\pi}{\cosh(\pi b)}. \end{aligned} \tag{7}$$

2.1. Radial reduction. Throughout, Fourier transforms use the convention (1), and we denote a radial function and its one-variable profile by the same symbol. With normalized Haar measure on $O(d)$, define the rotational average

$$\mathcal{R}f(x) = \int_{O(d)} f(Ux) dU.$$

Write $\mathcal{S}_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$ for the real Schwartz functions depending only on $|x|$, and $L^1_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$ for the radial real integrable functions. Set $\mathcal{A}_d^{\text{rad}} = \mathcal{A}_d \cap \mathcal{S}_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$. For every continuous integrable f , rotational invariance gives

$$\widehat{\mathcal{R}f} = \mathcal{R}\widehat{f}, \quad (\mathcal{R}f)(0) = f(0), \quad \widehat{\mathcal{R}f}(0) = \widehat{f}(0).$$

Consequently, rotational averaging preserves every sign condition in (2), so the infimum in (3) is unchanged when $\mathcal{A}_d$ is replaced by $\mathcal{A}_d^{\text{rad}}$ [CE03, CG19]. If $\widehat{g} = \varsigma g$, then

$$\widehat{\mathcal{R}g} = \varsigma \mathcal{R}g, \quad r(\mathcal{R}g) \leq r(g).$$

Moreover, $\mathcal{R}g \neq 0$ whenever $g \neq 0$ is eventually nonnegative. Indeed, if $\mathcal{R}g = 0$, the nonnegative values of g outside a ball have zero rotational average, so g vanishes outside that ball. Its Fourier transform is then entire, and $\widehat{g} = \varsigma g$ makes this entire function vanish on the same exterior region, forcing $g = 0$.

We also need approximation by Schwartz functions preserving the Fourier eigenvalue and vanishing at the origin. Suppose that $g \in L^1_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$ satisfies $\widehat{g} = \varsigma g$ and $g(0) = 0$, where $\varsigma \in \{-1, +1\}$. For $n \geq 1$, set

$$\kappa_n(x) = n^d e^{-\pi n^2 |x|^2}, \quad \eta_n(x) = e^{-\pi |x|^2/n^2}, \quad q_n = \eta_n(g * \kappa_n).$$

Gaussian mollification gives $q_n \in \mathcal{S}_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$, $q_n \rightarrow g$, and $\widehat{q}_n = \kappa_n * (\eta_n \widehat{g}) \rightarrow \widehat{g}$ in L^1 . Moreover, $q_n(0) \rightarrow g(0) = 0$ and $\widehat{q}_n(0) = \int q_n \rightarrow \widehat{g}(0) = 0$. Define

$$p_n = \frac{q_n + \varsigma \widehat{q}_n}{2}, \quad \psi_+(x) = e^{-\pi|x|^2}, \quad \psi_-(x) = \left(|x|^2 - \frac{d}{4\pi}\right) e^{-\pi|x|^2}.$$

Since $\widehat{p}_n = \varsigma p_n$, $\widehat{\psi}_\varsigma = \varsigma \psi_\varsigma$, and $\psi_\varsigma(0) \neq 0$, the correction

$$g_n = p_n - \frac{p_n(0)}{\psi_\varsigma(0)} \psi_\varsigma$$

satisfies $g_n \in \mathcal{S}_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$, $\widehat{g}_n = \varsigma g_n$, $g_n(0) = 0$, and $g_n \rightarrow g$ in L^1 .

2.2. The radial Mellin transform. For a radial Schwartz function g , write $r = |x|$ and $\rho = |\xi|$ for the spatial and frequency radii, respectively, and set

$$\lambda = \frac{d}{2}, \quad S_d = \frac{2\pi^{d/2}}{\Gamma(d/2)}.$$

Here S_d is the area of the unit sphere, and polar integration gives

$$\int_{\mathbb{R}^d} g(x) dx = S_d \int_0^\infty g(r) r^{d-1} dr.$$

For $\rho > 0$, the Fourier transform has the Hankel representation

$$\widehat{g}(\rho) = 2\pi \rho^{1-d/2} \int_0^\infty g(r) J_{d/2-1}(2\pi r \rho) r^{d/2} dr.$$

Because its Bessel kernel depends on $r\rho$, the radial Fourier transform becomes particularly simple after a Mellin transform: it reflects the Mellin variable and multiplies by an explicit gamma factor. For $\text{Re } z > 0$, $t \in \mathbb{R}$, and $r > 0$, the Mellin transform, its restriction to the critical line, and Mellin inversion [PK01] are

$$\begin{aligned} M_g(z) &= \int_0^\infty g(r) r^{z-1} dr, & X_g(t) &= M_g(\lambda - it), \\ g(r) &= \frac{r^{-\lambda}}{2\pi} \int_{\mathbb{R}} X_g(t) r^{it} dt. \end{aligned} \tag{8}$$

In particular, the radial integration normalization is

$$\widehat{g}(0) = \int_{\mathbb{R}^d} g(x) dx = S_d M_g(d).$$

In logarithmic radius $v = \log r$, set $\Phi_g(v) = e^{\lambda v} g(e^v)$. Smoothness at $r = 0$ gives exponential decay of Φ_g and all its derivatives as $v \rightarrow -\infty$, while the Schwartz decay of g gives rapid decay as $v \rightarrow +\infty$. Thus $\Phi_g \in \mathcal{S}(\mathbb{R}; \mathbb{R})$, and ordinary one-dimensional Fourier inversion gives

$$X_g(t) = \int_{\mathbb{R}} \Phi_g(v) e^{-itv} dv, \quad \Phi_g(v) = \frac{1}{2\pi} \int_{\mathbb{R}} X_g(t) e^{itv} dt.$$

The standard Mellin transform of the Hankel kernel gives, for $0 < \text{Re } z < d$, the Mellin–Hankel functional equation [SW71, Ch. IV], [DLMF, (10.22.43)]:

$$M_{\widehat{g}}(z) = \pi^{\lambda-z} \frac{\Gamma(z/2)}{\Gamma((d-z)/2)} M_g(d-z). \tag{9}$$

If $g(0) = \widehat{g}(0) = 0$, smooth radiality gives $g(r), \widehat{g}(r) = O(r^2)$ at the origin. Both Mellin transforms therefore extend holomorphically to $\text{Re } z > -2$, and (9) continues to $\text{Re } z = 0$:

$$M_g(d) = S_d^{-1} \widehat{g}(0) = 0, \quad M_g(d-z) = -z M'_g(d) + O(z^2), \quad \Gamma(z/2) = \frac{2}{z} + O(1),$$

so the apparent gamma pole at $z = 0$ cancels. The line $\text{Re } z = d/2$ is fixed by the reflection $z \mapsto d - z$. On this line the Fourier transform acts, for $t \in \mathbb{R}$, by

$$X_{\widehat{g}}(t) = m_\lambda(t) X_g(-t), \quad m_\lambda(t) = \pi^{it} \frac{\Gamma((\lambda - it)/2)}{\Gamma((\lambda + it)/2)}. \tag{10}$$

For $t \in \mathbb{R}$, its multiplier satisfies

$$m_\lambda(-t) = \overline{m_\lambda(t)} = m_\lambda(t)^{-1}, \quad |m_\lambda(t)| = 1.$$

3. THE UNIVERSAL COHN–ELKIES LOWER BOUND

By the radial reduction in § 2.1, assume that the admissible packing function F is radial. Let $a = (\widehat{F}(0)/F(0))^{1/d}$ and define $h(x) = F(ax)$. Then $h(0) = \widehat{h}(0)$, and $g = \widehat{h} - h$ is anti-self-Fourier, vanishes at the origin, and is nonnegative for $|x| \geq 1/a$; see (29)–(30). The proof of Theorem 3.8 verifies that $g \neq 0$. Since $\int g = 0$, its negative part has mass $\|g\|_1/2$, all of which lies in $B(0, 1/a)$.

Proposition 3.1 shows that every radial Schwartz Fourier eigenfunction vanishing at the origin, of either eigenvalue, has exponentially little L^1 mass in $B(0, c\sqrt{d})$ when $c < 1/\pi$. Thus the negative half-mass of g cannot fit inside this ball, forcing $1/a \geq (1/\pi - o(1))\sqrt{d}$.

The obstruction comes from the Mellin–Fourier identity (9). Lemma 3.2 bounds a normalized Mellin transform Z on the strip $|\operatorname{Im} t| < d/2$: total L^1 mass controls the upper boundary, while the Fourier functional equation controls the lower boundary. Writing $\lambda = d/2$, Poisson interpolation gives

$$\log |Z(s + i\sigma\lambda)| \leq H_\sigma(s) \leq \lambda M_\sigma(\log(2\pi c^2) + J_\sigma) + O_\sigma(\log \lambda), \quad M_\sigma = \frac{1 - \sigma}{2},$$

by Lemma 3.3. The sharp constant enters through Lemma 3.4: $J_\sigma \rightarrow \log(\pi/2)$ as $\sigma \uparrow 1$, so the parenthesized rate tends to $\log(\pi^2 c^2)$. It is negative exactly when $c < 1/\pi$. An all-frequency bound and shifted Mellin inversion, in Lemmas 3.5 and 3.6, turn this negativity into the interior-mass estimate; Stirling’s formula then yields the packing lower bound.

3.1. The Mellin-strip obstruction.

Proposition 3.1. *For every $0 < c < 1/\pi$, there exist $C_c, \gamma_c > 0$ and $d_0(c) \in \mathbb{N}$ such that, for every $d \geq d_0(c)$, every $\varsigma \in \{-1, +1\}$, and every nonzero $g \in \mathcal{S}_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$ satisfying $\widehat{g} = \varsigma g$ and $g(0) = 0$, one has*

$$\int_{|x| < c\sqrt{d}} |g(x)| dx \leq C_c e^{-\gamma_c d} \|g\|_1. \quad (11)$$

Fix $0 < c < 1/\pi$, $d \in \mathbb{N}$, $\lambda = d/2$, $R = c\sqrt{d}$, and a nonzero $g \in \mathcal{S}_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$ with $\widehat{g} = \varsigma g$, $\varsigma \in \{-1, +1\}$, and $g(0) = 0$. Let $S_d = 2\pi^\lambda/\Gamma(\lambda)$ be the area of the unit sphere. In the logarithmic coordinate $r = Re^v$, normalize the Mellin inversion formula (8) by setting

$$\varphi(v) = \frac{S_d}{\|g\|_1} (Re^v)^d g(Re^v), \quad Z(t) = \frac{S_d}{\|g\|_1} R^{\lambda+it} X_g(t), \quad (12)$$

$$\|\varphi\|_1 = 1, \quad \int_{\mathbb{R}} \varphi = 0, \quad \int_{-\infty}^0 |\varphi(v)| dv = \frac{1}{\|g\|_1} \int_{|x| < R} |g(x)| dx. \quad (13)$$

The second line follows from polar integration and $\widehat{g}(0) = \varsigma g(0) = 0$.

Lemma 3.2. *For every $-1 < \sigma < 1$, the function Z is bounded and holomorphic on a neighborhood of the strip $|\operatorname{Im} t| \leq \lambda$. Its boundary values satisfy $|Z(y + i\lambda)| \leq 1$ and $\log |Z(y - i\lambda)| \leq h_\lambda(y)$ for $y \neq 0$, where the lower-boundary majorant is*

$$h_\lambda(y) = \lambda \log(\pi R^2) + \log |\Gamma(-iy/2)| - \log |\Gamma(\lambda + iy/2)|. \quad (14)$$

Writing $\theta = \pi(1 + \sigma)/2$, define

$$P_\sigma(T) = \frac{\sin \theta}{4(\cosh(\pi T/2) - \cos \theta)}, \quad M_\sigma = \int_{\mathbb{R}} P_\sigma = \frac{1 - \sigma}{2}. \quad (15)$$

Then

$$|Z(s + i\sigma\lambda)| \leq \exp \left(\int_{\mathbb{R}} P_\sigma(T) h_\lambda(s - \lambda T) dT \right) \quad (s \in \mathbb{R}).$$

Proof. Since $g(0) = \hat{g}(0) = 0$, the holomorphic continuation and pole cancellation following (9) show that Z is holomorphic whenever $\text{Im } t > -\lambda - 2$. In particular, it is holomorphic on a neighborhood of the closed strip. Its upper-boundary values follow from (12), whereas the Mellin–Fourier functional equation (9) gives the lower-boundary values:

$$Z(y + i\lambda) = \int_{\mathbb{R}} \varphi(v) e^{-iyv} dv, \quad |Z(y + i\lambda)| \leq 1, \quad (16)$$

$$Z(y - i\lambda) = \varsigma(\pi R^2)^{\lambda+iy} \frac{\Gamma(-iy/2)}{\Gamma(\lambda + iy/2)} Z(-y + i\lambda). \quad (17)$$

Taking absolute values and using (16) gives $\log |Z(y - i\lambda)| \leq h_\lambda(y)$, independently of ς . At $y = 0$, the zero $Z(i\lambda) = 0$ from (13) cancels the gamma pole in (17). Hence $Z(y - i\lambda)$ remains bounded there, although

$$h_\lambda(y) = -\log |y| + O_\lambda(1) \quad (y \rightarrow 0).$$

For $t_0 = s + i\sigma\lambda$, the conformal map $t \mapsto \exp(\pi(t + i\lambda)/(2\lambda))$ and the upper-half-plane Poisson formula [Ahl79] give the lower-edge harmonic measure $\lambda^{-1} P_\sigma((s - y)/\lambda) dy$, with P_σ and M_σ as in (15). The complementary upper-edge measure is $(1 + \sigma)/2$, and its boundary bound (16) contributes nothing.

The logarithmic singularity of h_λ requires a bounded truncation before applying the Poisson principle. The actual lower-boundary values satisfy

$$\sup_{y \in \mathbb{R}} |Z(y - i\lambda)| \leq \frac{S_d R^d}{\|g\|_1} \int_0^\infty \frac{|g(r)|}{r} dr < \infty.$$

Indeed, $g(r) = O(r^2)$ at zero and is Schwartz at infinity. More generally, for $-\lambda \leq \eta \leq \lambda$, (12) gives

$$|Z(s + i\eta)| \leq \frac{S_d R^{\lambda-\eta}}{\|g\|_1} \int_0^\infty |g(r)| r^{\lambda+\eta-1} dr.$$

Splitting at $r = 1$ bounds this expression uniformly in s and η , so Z is bounded on the closed strip. Choose $D > \max\{0, \sup_y \log |Z(y - i\lambda)|\}$. The strip Poisson principle [Ahl79], applied to $\log |Z|$ with lower-boundary majorant $\min\{h_\lambda, D\}$ and upper-boundary majorant 0, gives

$$\log |Z(t_0)| \leq \int_{\mathbb{R}} \frac{1}{\lambda} P_\sigma \left(\frac{s - y}{\lambda} \right) \min\{h_\lambda(y), D\} dy.$$

The kernel in (15) decays exponentially, the singularity of h_λ is locally integrable, and Stirling's formula gives $h_\lambda(y) = -\lambda \log |y| + O_\lambda(1)$ at infinity [DLMF, § 5.11]. Dominated convergence therefore permits $D \rightarrow \infty$, and the substitution $T = (s - y)/\lambda$ yields

$$|Z(s + i\sigma\lambda)| \leq e^{H_\sigma(s)}, \quad H_\sigma(s) = \int_{\mathbb{R}} P_\sigma(T) h_\lambda(s - \lambda T) dT. \quad \square \quad (18)$$

Lemma 3.3. *For every $-1 < \sigma < 1$, there exists $C_\sigma > 0$, independent of d , c , and g , such that*

$$\int_{\mathbb{R}} P_\sigma(T) \left| h_\lambda(\lambda T) - \lambda \left(\log(2\pi c^2) - \int_0^1 \log \sqrt{x^2 + T^2/4} dx \right) \right| dT \leq C_\sigma \log(2 + \lambda). \quad (19)$$

Define

$$J_\sigma = -\frac{1}{M_\sigma} \int_{\mathbb{R}} P_\sigma(T) \int_0^1 \log \sqrt{x^2 + T^2/4} dx dT.$$

Then, for every $s \in \mathbb{R}$,

$$H_\sigma(s) \leq H_\sigma(0) = \lambda M_\sigma (\log(2\pi c^2) + J_\sigma) + O_\sigma(\log(2 + \lambda)). \quad (20)$$

Proof. To estimate the lower-boundary majorant (14) with $R = c\sqrt{2\lambda}$, put

$$f_T(x) = \log \sqrt{x^2 + T^2/4}, \quad b = \frac{\lambda T}{2}.$$

If $d = 2n$, so that $\lambda = n$, the gamma recurrence in (7) and $|\Gamma(-ib)| = |\Gamma(ib)|$ give, for $T \neq 0$,

$$h_n(nT) = n \log(2\pi c^2) - \sum_{k=0}^{n-1} f_T(k/n).$$

Monotonicity of f_T bounds the Riemann-sum error by $f_T(1) - f_T(0)$; thus

$$0 \leq h_n(nT) - n \left(\log(2\pi c^2) - \int_0^1 f_T(x) dx \right) \leq \frac{1}{2} \log \left(1 + \frac{4}{T^2} \right).$$

If $d = 2n + 1$, so that $\lambda = n + \frac{1}{2}$, the gamma identities (7) instead give

$$h_\lambda(\lambda T) = \lambda \log(2\pi c^2) - \sum_{k=0}^{n-1} f_T \left(\frac{k + 1/2}{\lambda} \right) + \frac{1}{2} \log \lambda + \frac{1}{2} \log \frac{\coth(\pi|b|)}{|b|}.$$

The midpoint Riemann-sum error on $[0, n/\lambda]$ is at most $f_T(n/\lambda) - f_T(0)$; the remaining interval $[n/\lambda, 1]$ has length $1/(2\lambda)$. Together, these contributions are bounded by

$$C \left(1 + \log(2 + |T|) + \log(2 + |T|^{-1}) \right).$$

Adding $C \log(2 + \lambda)$ also bounds the gamma endpoint correction $\frac{1}{2} \log \lambda + \frac{1}{2} \log(\coth(\pi|b|)/|b|)$. Since (15) gives $P_\sigma(T) \ll_\sigma e^{-\pi|T|/2}$, and $\log(2 + |T|^{-1})$ is locally integrable, integrating the even- and odd-dimensional bounds proves (19).

It remains to identify the maximum of the Poisson extension H_σ . Both h_λ and P_σ are even, and the digamma series [DLMF, § 5.7] gives, for $y > 0$,

$$h'_\lambda(y) = \frac{1}{2} (\operatorname{Im} \psi(\lambda + iy/2) - \operatorname{Im} \psi(iy/2)) < 0, \quad \operatorname{Im} \psi(a + ib) = \sum_{k=0}^{\infty} \frac{b}{(k+a)^2 + b^2}.$$

Here $\psi = \Gamma'/\Gamma$. The explicit kernel (15) is likewise decreasing on $(0, \infty)$. For $N > 0$, the function $q_N = (h_\lambda + N)_+$ is nonnegative, integrable, even, and decreasing there. The convolution of two such functions is largest at zero. Subtracting NM_σ and letting $N \rightarrow \infty$, using the exponential decay of P_σ , therefore gives $H_\sigma(s) \leq H_\sigma(0)$. Evaluating at zero with (19) proves (20). $\square$

Lemma 3.4. *For J_σ defined in Lemma 3.3,*

$$\lim_{\sigma \uparrow 1} J_\sigma = \log \frac{\pi}{2}.$$

Consequently, for every $0 < c < 1/\pi$, there exists $\sigma = \sigma(c) \in (-1, 1)$ such that $\log(2\pi c^2) + J_\sigma < 0$.

Proof. With $T = 2u$, divide the lower-edge harmonic measure $P_\sigma(T) dT$ from (15) by its total mass M_σ . The corresponding probability density in u is

$$p_\sigma(u) = \frac{2P_\sigma(2u)}{M_\sigma} = \frac{\sin \theta}{(1 - \sigma)(\cosh(\pi u) - \cos \theta)}.$$

As $\sigma \uparrow 1$, $\sin \theta/(1 - \sigma) \rightarrow \pi/2$. The densities p_σ are uniformly bounded by $Ce^{-\pi|u|}$ and therefore converge in $L^1(\mathbb{R})$ to

$$p(u) = \frac{\pi}{4} \operatorname{sech}^2 \left(\frac{\pi u}{2} \right).$$

The characteristic function of this density is

$$\int_{\mathbb{R}} p(u) e^{itu} du = \frac{t}{\sinh t}, \quad t \in \mathbb{R}, \quad (21)$$

interpreted as 1 at $t = 0$. For $t > 0$, a contour shift by $2i$ across the double pole at i gives $(1 - e^{-2t}) \int_{\mathbb{R}} p(u) e^{itu} du = 2te^{-t}$; evenness handles $t < 0$. Define $I(x) = \int_{\mathbb{R}} p(u) \log \sqrt{x^2 + u^2} du$. For $x > 0$, the Laplace representation of $x/(x^2 + u^2)$, (21), and the trigamma integral [DLMF, § 5.9] yield

$$I'(x) = \int_0^\infty e^{-xt} \frac{t}{\sinh t} dt = \frac{1}{2} \psi' \left(\frac{x+1}{2} \right).$$

Since both $I(x)$ and $\psi((x+1)/2) + \log 2$ equal $\log x + o(1)$ at infinity [DLMF, § 5.11], their integration constants agree. Local integrability at $u = 0$ extends the identity to $x = 0$:

$$\int_{\mathbb{R}} p(u) \log \sqrt{x^2 + u^2} du = \psi\left(\frac{x+1}{2}\right) + \log 2 \quad (x \geq 0). \quad (22)$$

The uniform exponential bound on p_σ and local integrability of $\log |u|$ justify dominated convergence in J_σ . Since

$$\int_0^1 \psi\left(\frac{x+1}{2}\right) dx = 2 \log \frac{\Gamma(1)}{\Gamma(1/2)} = -\log \pi,$$

(22) gives the exact threshold

$$\lim_{\sigma \uparrow 1} J_\sigma = \log \frac{\pi}{2}, \quad \log(2\pi c^2) + \log \frac{\pi}{2} = \log(\pi^2 c^2).$$

If $c < 1/\pi$, choose $\sigma = \sigma(c) < 1$ sufficiently close to 1 that

$$\delta_c = -(\log(2\pi c^2) + J_{\sigma(c)}) > 0. \quad \square \quad (23)$$

Fix $\sigma = \sigma(c)$ and $\delta_c > 0$ as in (23). We first bound the Mellin transform Z on the horizontal line $\text{Im } t = \sigma\lambda$, and then use that bound to control the mass of g inside $B(0, c\sqrt{d})$.

Lemma 3.5. *There exist $\gamma_c, C_c, B_c > 0$, depending only on c , such that, for every sufficiently large d ,*

$$H_\sigma(s) \leq -\gamma_c \lambda \quad (s \in \mathbb{R}), \quad \int_{\mathbb{R}} |Z(s + i\sigma\lambda)| ds \leq C_c \lambda e^{-\gamma_c \lambda}.$$

Moreover, after increasing B_c if necessary,

$$H_\sigma(\lambda S) \leq -\frac{M_\sigma \lambda}{2} \log \frac{|S|}{C_c} \quad (|S| \geq B_c).$$

Proof. The maximum estimate for H_σ in (20), together with $\log(2\pi c^2) + J_\sigma = -\delta_c$, gives

$$H_\sigma(s) \leq -\lambda M_\sigma \delta_c + O_\sigma(\log \lambda) \leq -\gamma_c \lambda \quad (s \in \mathbb{R})$$

for all sufficiently large d , with $\gamma_c > 0$ independent of s, g , and the Fourier eigenvalue ς .

To control all Mellin frequencies, apply the gamma identities (7) to (14). In both parities,

$$h_\lambda(\lambda U) \leq \lambda \log \frac{4\pi c^2}{|U|} + E_\lambda(U) \quad (U \neq 0),$$

where

$$E_\lambda(U) = \begin{cases} 0, & \lambda \in \mathbb{N}, \\ \frac{1}{2} \log \coth\left(\frac{\pi\lambda|U|}{2}\right), & \lambda \in \mathbb{N} + \frac{1}{2}. \end{cases}$$

Each gamma-recurrence factor has modulus at least $|\lambda U|/2$; in odd dimension, the remaining ratio at $b = \lambda U/2$ contributes $E_\lambda(U)$. Expanding $\log \coth x$ in its convergent odd-exponential series gives, in odd dimension,

$$\int_{\mathbb{R}} E_\lambda(U) dU = \frac{2}{\pi\lambda} \int_0^\infty \log \coth x dx = \frac{\pi}{4\lambda}.$$

Thus in both parities the P_σ -convolution of E_λ is at most $\pi\|P_\sigma\|_\infty/(4\lambda)$. Consequently, (18) gives

$$H_\sigma(\lambda S) \leq \lambda M_\sigma \log(4\pi c^2) - \lambda \int_{\mathbb{R}} P_\sigma(T) \log |S - T| dT + O_\sigma(\lambda^{-1}).$$

Split the logarithmic integral at $|T| = |S|/2$. The exponential decay of (15) gives mass $M_\sigma + O_\sigma(e^{-\pi|S|/4})$ on $|T| \leq |S|/2$; the possible negative contribution from $|S - T| < 1$ is $O_\sigma(e^{-\pi|S|/2})$, by local integrability of $\log |S - T|$. Hence, for some $B_c, C'_c > 0$,

$$\int_{\mathbb{R}} P_\sigma(T) \log |S - T| dT \geq \frac{M_\sigma}{2} \log |S| - C'_c \quad (|S| \geq B_c).$$

After increasing C'_c , the estimates for H_σ on the whole line and at large frequencies become

$$H_\sigma(s) \leq -\gamma_c \lambda \quad (s \in \mathbb{R}), \quad (24)$$

$$H_\sigma(\lambda S) \leq -\frac{M_\sigma \lambda}{2} \log \frac{|S|}{C'_c} \quad (|S| \geq B_c). \quad (25)$$

Choose $B > \max\{B_c, C'_c\}$ and $q = M_\sigma \lambda / 2 > 1$. The two bounds in (24)–(25) yield

$$\begin{aligned} \int_{|s| \leq B\lambda} e^{H_\sigma(s)} ds &\leq 2B\lambda e^{-\gamma_c \lambda}, \\ \int_{|s| > B\lambda} e^{H_\sigma(s)} ds &\leq \frac{2\lambda C'_c}{q-1} \left(\frac{B}{C'_c}\right)^{1-q}. \end{aligned}$$

The second bound decays at rate $(M_\sigma/2) \log(B/C'_c) > 0$. Decreasing γ_c if necessary, and applying (18), we obtain

$$\int_{\mathbb{R}} |Z(s + i\sigma\lambda)| ds \leq C_c \lambda e^{-\gamma_c \lambda}. \quad \square \quad (26)$$

The global L^1 estimate (26) for Z on $\text{Im } t = \sigma\lambda$ now controls $\varphi(v)$ for $v < 0$, corresponding by (13) to $|x| < c\sqrt{d}$.

Lemma 3.6. *For every $0 < c < 1/\pi$, there exist $C_c, \gamma_c > 0$ and $d_0(c) \in \mathbb{N}$, independent of g and ς , such that*

$$\int_{-\infty}^0 |\varphi(v)| dv \leq C_c e^{-\gamma_c d} \quad (d \geq d_0(c)).$$

Proof. Let $G(v) = e^{(\sigma-1)\lambda v} \varphi(v)$. The normalized profile φ from (12) satisfies

$$\int_{\mathbb{R}} |G(v)| dv = \frac{S_d R^{(1-\sigma)\lambda}}{\|g\|_1} \int_0^\infty |g(r)| r^{(1+\sigma)\lambda-1} dr < \infty.$$

Thus $G \in L^1(\mathbb{R})$, and (12) identifies its angular-frequency Fourier transform $\int_{\mathbb{R}} G(v) e^{-isv} dv$ with $Z(s + i\sigma\lambda)$. The integrability of this transform follows from (26), so Fourier inversion gives

$$Z(s + i\sigma\lambda) = \int_{\mathbb{R}} e^{(\sigma-1)\lambda v} \varphi(v) e^{-isv} dv, \quad \varphi(v) = \frac{e^{(1-\sigma)\lambda v}}{2\pi} \int_{\mathbb{R}} Z(s + i\sigma\lambda) e^{isv} ds.$$

Taking absolute values and integrating over $v < 0$ contributes $\int_{-\infty}^0 e^{(1-\sigma)\lambda v} dv = ((1-\sigma)\lambda)^{-1}$. Combining this with (26) and $\lambda = d/2$, and decreasing γ_c if necessary, gives

$$\int_{-\infty}^0 |\varphi(v)| dv \leq \frac{1}{2\pi(1-\sigma)\lambda} \int_{\mathbb{R}} |Z(s + i\sigma\lambda)| ds \leq C_c e^{-\gamma_c d}. \quad \square \quad (27)$$

Proof of Proposition 3.1. By (13), the left side of (27) is precisely $\|g\|_1^{-1} \int_{|x| < c\sqrt{d}} |g(x)| dx$. Thus (27) proves (11), uniformly in g and its Fourier eigenvalue ς . $\square$

3.2. The packing lower bound.

Proposition 3.7. *For every $0 < c < 1/\pi$, there exists $d_0(c) \in \mathbb{N}$ such that, for every $d \geq d_0(c)$ and $\varsigma \in \{-1, +1\}$, no nonzero $g \in L^1(\mathbb{R}^d; \mathbb{R})$ satisfies $\hat{g} = \varsigma g$, $g(0) = 0$, and $g(x) \geq 0$ for $|x| \geq c\sqrt{d}$. Here g denotes its continuous Fourier-inversion representative.*

Proof. First, suppose that g is a radial Schwartz eigenfunction. Since $\int g = \hat{g}(0) = \varsigma g(0) = 0$, its negative part $g_- = \max\{-g, 0\}$ has integral $\|g\|_1/2$. The assumption $g(x) \geq 0$ for $|x| \geq c\sqrt{d}$ forces g_- to vanish outside $B(0, c\sqrt{d})$, contradicting (11) once $C_c e^{-\gamma_c d} < 1/2$.

For general $g \in L^1$, the radial reduction in § 2.1 shows that $h = \mathcal{R}g \neq 0$ and has the same Fourier eigenvalue, origin value, and exterior sign. The approximation constructed there gives radial Schwartz eigenfunctions $h_n \rightarrow h$ in L^1 with $\hat{h}_n = \varsigma h_n$ and $h_n(0) = 0$. Writing $(h_n)_- = \max\{-h_n, 0\}$ and $R = c\sqrt{d}$, the exterior nonnegativity of h gives

$$\int_{\mathbb{R}^d} (h_n)_- \leq \int_{|x| < R} |h_n(x)| dx + \int_{|x| \geq R} |h_n(x) - h(x)| dx.$$

Applying (11) to h_n therefore implies

$$\frac{1}{2}\|h_n\|_1 = \int_{\mathbb{R}^d} (h_n)_- \leq C_c e^{-\gamma_c d} \|h_n\|_1 + \|h_n - h\|_1.$$

Letting $n \rightarrow \infty$ gives $\|h\|_1/2 \leq C_c e^{-\gamma_c d} \|h\|_1$, which contradicts $h \neq 0$ for all sufficiently large d . $\square$

Theorem 3.8. *There is a sequence $\epsilon_d \rightarrow 0$ such that, for every sufficiently large d and every $F \in \mathcal{A}_d$,*

$$\frac{F(0)}{\widehat{F}(0)} \geq \frac{2^d}{v_d} \left(\sqrt{\frac{e}{2\pi}} - \epsilon_d \right)^d. \quad (28)$$

Proof. By radial reduction, replace F with its rotational average, which preserves $F(0)$, $\widehat{F}(0)$, and admissibility. Define $a = (\widehat{F}(0)/F(0))^{1/d} > 0$, $h(x) = F(ax)$, and $g = \widehat{h} - h$. The Fourier convention (1) and the admissibility conditions (2) give

$$\widehat{h}(\xi) = a^{-d} \widehat{F}(\xi/a), \quad h(0) = \widehat{h}(0) = F(0), \quad (29)$$

$$\widehat{g} = -g, \quad g(0) = 0, \quad g(x) = a^{-d} \widehat{F}(x/a) - F(ax) \geq 0 \quad (|x| \geq 1/a) \quad (30)$$

Moreover, $g \neq 0$: otherwise $h = \widehat{h} \geq 0$, while $h(x) = F(ax) \leq 0$ for $|x| \geq 1/a$. Thus h would be a nonzero compactly supported self-Fourier function, contradicting Fourier analyticity. Hence Proposition 3.7 gives $1/a > c\sqrt{d}$, uniformly in F , for each fixed $c < 1/\pi$ and all sufficiently large d . Consequently,

$$\liminf_{d \rightarrow \infty} \frac{1}{\sqrt{d}} \inf_{F \in \mathcal{A}_d} \left(\frac{F(0)}{\widehat{F}(0)} \right)^{1/d} \geq c \quad (0 < c < 1/\pi).$$

Taking the supremum over $c < 1/\pi$ yields

$$\inf_{F \in \mathcal{A}_d} \left(\frac{F(0)}{\widehat{F}(0)} \right)^{1/d} \geq \left(\frac{1}{\pi} - o(1) \right) \sqrt{d}. \quad (31)$$

Combining (31) with Stirling's formula [DLMF, § 5.11], $v_d^{1/d} = (1 + o(1))\sqrt{2\pi e/d}$ gives (28), for a sequence $\epsilon_d \rightarrow 0$ independent of F . $\square$

4. THE ADMISSIBLE PRIMAL UPPER BOUND

Section 3 established

$$\min \left\{ \inf_{F \in \mathcal{A}_d} \left(\frac{F(0)}{\widehat{F}(0)} \right)^{1/d}, A_-(d), A_+(d) \right\} \geq \left(\frac{1}{\pi} - o(1) \right) \sqrt{d}.$$

We explain how the upper bounds in the introduction reduce to constructing functions whose sign changes occur at the same radius. Suppose that real radial Schwartz functions f_-, f_+ satisfy

$$\widehat{f_-} = f_+ > 0, \quad f_-(0) = f_+(0), \quad f_-(x) < 0 \quad (|x| \geq R).$$

For $F(x) = f_-(Rx)$, Fourier scaling gives

$$\widehat{F}(\xi) = R^{-d} f_+(\xi/R) > 0, \quad F(x) < 0 \quad (|x| \geq 1), \quad \frac{F(0)}{\widehat{F}(0)} = R^d.$$

Consequently $F \in \mathcal{A}_d$ and $\text{LP}_d \leq v_d(R/2)^d$. The difference $g_- = f_+ - f_-$ is anti-self-Fourier, vanishes at the origin, and is positive for $|x| \geq R$, so it also proves $A_-(d) \leq R$. A self-Fourier function f_0 satisfying $f_0(0) = 0$ and $f_0(x) > 0$ for $|x| \geq R$ similarly proves $A_+(d) \leq R$. Thus all the upper bounds reduce to producing one Fourier pair, one self-Fourier function, and $R = (1/\pi + o(1))\sqrt{d}$.

Theorem 4.1. *There is $\epsilon_0 > 0$ such that, for every fixed $0 < \epsilon < \epsilon_0$ and every sufficiently large dimension d , there exist real radial Schwartz functions f_-, f_+, f_0 and a radius $R_{\epsilon,d} > 0$ such that $f_+ > 0$ everywhere, $f_- < 0 < f_0$ whenever $|x| \geq R_{\epsilon,d}$, and*

$$\hat{f}_- = f_+, \quad \hat{f}_0 = f_0, \quad f_-(0) = f_+(0) > 0, \quad f_0(0) = 0.$$

Moreover,

$$\lim_{\epsilon \downarrow 0} \lim_{d \rightarrow \infty} \frac{R_{\epsilon,d}}{\sqrt{d}} = \frac{1}{\pi}.$$

4.1. Outline of the construction. We need a Fourier pair $f_-, f_+ = \hat{f}_-$ and a self-Fourier function f_0 , with $f_+ > 0$ everywhere and $f_- < 0 < f_0$ outside a common radius. In Mellin coordinates, (10) reduces these Fourier symmetries to reflection of the frequency. The Gaussian

$$g_G(r) = 2\pi^{\lambda/2} e^{-\pi r^2}, \quad E_\lambda^G(t) = \pi^{it/2} \Gamma\left(\frac{\lambda - it}{2}\right), \quad \lambda = \frac{d}{2},$$

already obeys $m_\lambda(t) E_\lambda^G(-t) = E_\lambda^G(t)$. Multiplication by an even factor therefore preserves Fourier symmetry. We choose $E_\lambda(t) = E_\lambda^G(t) e^{\lambda h_\epsilon(t/\lambda)}$, where the even perturbation h_ϵ in (36) is determined by a signed density w .

The variable a in that density parametrizes radial dilations: the Mellin multiplier $\cos(at/\lambda) - 1$ corresponds to

$$g(r) \mapsto \frac{e^a g(re^{a/\lambda}) + e^{-a} g(re^{-a/\lambda})}{2} - g(r).$$

A *shell* is an interval of these dilation parameters supporting one component of w . Our negative component w_s moves the sign radius inward, and a positive component w_B , supported on much larger dilation parameters, restores decay at every nonzero Mellin frequency.

We multiply the common envelope by polynomials P_-, P_+, P_0 . Reflection exchanges P_- and P_+ and fixes P_0 , giving the desired Fourier symmetries. The first gamma pole, at normalized frequency $\zeta = -i$, controls the value at the origin: the polynomial P_0 cancels this pole, while P_- and P_+ retain equal positive residues. On the imaginary axis, the polynomials are chosen so that $P_+(iu) > 0$ for $u > -1$, whereas $P_-(iu) < 0$ and $P_0(iu) > 0$ just above $u = 1$; see (39). Thus the essential contour is $t = \lambda(T + iu)$ with $u \approx 1$.

On this contour, the radius $r = e^{v(u)}$ for which the Mellin integrand is stationary at $T = 0$ satisfies (44). For fixed $u > -1$, the digamma asymptotic gives

$$\frac{e^{v(u)}}{\sqrt{d}} \rightarrow \sqrt{\frac{1+u}{4\pi}} \exp\left(\int_0^\infty w(a) a \sinh(ua) da\right).$$

Consequently, negative w decreases the logarithm of the stationary radius when $u > 0$. However, it also reduces the decay of the Mellin integrand away from $T = 0$. At $u = 1$, after dividing the damping by λ , the limiting gamma contribution has density $e^{-2a}/(2a^2)$, while the perturbation contributes $w(a) \cosh a$. A sufficient pointwise condition for nonnegative total damping is therefore

$$|w(a)| \cosh a \leq \frac{e^{-2a}}{2a^2}.$$

Within this pointwise constraint, the greatest inward displacement is achieved formally by

$$w_*(a) = -\frac{e^{-2a}}{2a^2 \cosh a}, \quad \int_0^\infty w_*(a) a \sinh a da = -\frac{1}{2} \log \frac{\pi}{2}, \quad (32)$$

where the integral evaluation follows from the gamma duplication formula [DLMF, § 5.5]. The Gaussian stationary radius at $u = 1$ is $(2\pi)^{-1/2} \sqrt{d}$, so the displacement in (32) gives

$$\frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \log \frac{\pi}{2}\right) = \frac{1}{\pi}. \quad (33)$$

Although its Mellin perturbation is well defined, the formal density w_* has infinite mass at zero and saturates the gamma damping. Its Mellin factor therefore lacks the strict decay needed for a Schwartz function. Instead, we truncate it to a bounded interval and taper it slightly, obtaining a negative shell w_s that preserves the displacement up to $O(\epsilon)$ while leaving a definite

damping margin. This margin suffices near $u = 1$, but not on contours with arbitrarily large u . A second shell w_B , supported on $[B, B + 1]$, is chosen to have negligible effect on the stationary radius at $u = u_0$ while dominating the negative shell at every frequency when $u \geq U > u_0$. Using an interval rather than a single dilation avoids frequencies at which its damping would vanish.

The remaining proof has two analytic parts. Uniform saddle estimates show that the Mellin integral at $r = e^{v(u)}$ has the sign of $P_j(iu)$, which gives the exterior signs and positivity of f_+ for $r \geq r_*$. A downward contour shift then expresses f_+ on $[0, r_*]$ as a perturbed exponential series, proving positivity at the remaining radii. The negative-shell displacement in (32) finally yields the sharp radius (33).

4.2. The Mellin ansatz. We turn the ideal negative density w_* in (32) into two compactly supported shells. The negative shell w_s retains its logarithmic-radius displacement up to $O(\epsilon)$, while the positive shell w_B is negligible at $u = u_0$ and dominates the damping for $u \geq U$.

Put $\lambda = d/2$, and throughout the construction fix ϵ before taking $d \rightarrow \infty$. Constants in $O_\epsilon(\cdot)$ and $\ll_\epsilon$ may depend on ϵ , but are independent of d and of the saddle parameter whenever a bound is stated uniformly in that parameter; unadorned implicit constants are absolute.

Introduce cutoffs $0 < a_0 < A < B$, a positive-shell amplitude $Q > 0$, and saddle parameters

$$u_0 = 1 + \frac{\epsilon}{4}, \quad U = 1 + \frac{\epsilon}{2}, \quad C_0 = A + a_0^{-1}.$$

As $\epsilon \downarrow 0$, the cutoffs should satisfy

$$a_0 = o(\epsilon), \quad \frac{e^{-2A}}{A} = o(\epsilon), \quad \epsilon A = o(1), \quad A = o(B).$$

The first two conditions make the omitted portions of the ideal saddle displacement negligible, and the third permits an $O(\epsilon(1 + a))$ taper on the negative shell. For the positive shell, of amplitude Q , the required separation is

$$BQe^{(u_0-1)B} = o(1), \quad C_0Q^{-1}e^{-(U-1)(B-A)} = o(1).$$

The bound on $BQe^{(u_0-1)B}$ keeps the positive shell from moving the target saddle. The bound on $C_0Q^{-1}e^{-(U-1)(B-A)}$ makes that shell dominate the negative shell once $u \geq U$. Because $u_0 - 1 < U - 1$, both conditions are compatible: writing $Q = e^{-q_\epsilon B}$, choose $u_0 - 1 < q_\epsilon < U - 1$ with enough separation that ϵB dominates $\log B + \log C_0$.

One convenient realization of all these requirements is

$$\begin{aligned} a_0 &= \epsilon^2, & A &= \log(1/\epsilon), & B &= \epsilon^{-3}, \\ q_\epsilon &= \frac{(u_0 - 1) + (U - 1)}{2}, & Q &= e^{-q_\epsilon B}, \\ b(a) &= 1 - 2\epsilon(1 + a), & \beta &= u_0 - 1. \end{aligned} \tag{34}$$

The exponential slope q_ϵ is the midpoint between the two contour heights $u_0 - 1$ and $U - 1$. The taper leaves a damping margin of order ϵ , and $\beta = u_0 - 1$ places the sign transition at the target saddle. For sufficiently small ϵ , we have $b > 0$ on $[a_0, A]$ and $B > A + 1$. Define

$$\begin{aligned} w_s(a) &= -\frac{b(a)e^{-2a}}{2a^2 \cosh a} \mathbf{1}_{[a_0, A]}(a), \\ w_B(a) &= \frac{Q}{\cosh a} \mathbf{1}_{[B, B+1]}(a), & w &= w_s + w_B. \end{aligned} \tag{35}$$

The signed density w determines the even entire Mellin perturbation

$$h_\epsilon(\zeta) = \int_0^\infty w(a)(\cos(a\zeta) - 1) da. \tag{36}$$

For $\eta > 0$, the positive density describing the unperturbed gamma damping is

$$\mu_{\lambda, \eta}(a) = \frac{e^{-\eta a}}{a(1 - e^{-2a/\lambda})} \quad (a > 0). \tag{37}$$

Lemma 4.2. *There are absolute constants $\epsilon_0, c, C > 0$ such that, for every $0 < \epsilon < \epsilon_0$, the shells in (35) have the following properties. For every $\lambda > 0$, $-1 < u \leq U$, and $a \in [a_0, A]$,*

$$\lambda |w_s(a)| \cosh(ua) \leq (1 - c\epsilon) \mu_{\lambda, 1+u}(a).$$

At the target saddle,

$$\begin{aligned} \int_{a_0}^A w_s(a) a \sinh(u_0 a) da &= -\frac{1}{2} \log \frac{\pi}{2} + O(\epsilon), \\ 0 &\leq \int_B^{B+1} w_B(a) a \sinh(u_0 a) da \leq C e^{-c/\epsilon^2}. \end{aligned}$$

Finally, the target-saddle contribution and the remote-saddle domination ratio satisfy

$$BQ e^{(u_0-1)B} \leq C e^{-c/\epsilon^2}, \quad C_0 Q^{-1} e^{-(U-1)(B-A)} \leq C e^{-c/\epsilon^2}.$$

The implicit constant in $O(\epsilon)$ is absolute.

Proof. Write

$$w_*(a) = -\frac{e^{-2a}}{2a^2 \cosh a}.$$

The negative shell agrees with w_* , up to its taper, on $[a_0, A]$. Its omitted saddle contributions are

$$\int_0^{a_0} |w_*(a)| a \sinh a da = O(a_0), \quad \int_A^\infty |w_*(a)| a \sinh a da = O\left(\frac{e^{-2A}}{A}\right),$$

because $\tanh a \leq \min(a, 1)$. The taper changes the same integral by only

$$\int_{a_0}^A |1 - b(a)| |w_*(a)| a \sinh a da = O\left(\epsilon \int_0^\infty (1+a) \frac{e^{-2a} \tanh a}{a} da\right) = O(\epsilon).$$

Moving from $u = 1$ to $u = u_0$ contributes another $O(\epsilon)$: the mean-value theorem gives

$$\int_{a_0}^A |w_s(a)| a |\sinh(u_0 a) - \sinh a| da \ll \epsilon \int_0^\infty \frac{e^{-2a} \cosh(u_0 a)}{\cosh a} da \ll \epsilon.$$

Since $a_0 = o(\epsilon)$ and $e^{-2A}/A = o(\epsilon)$, the identity $\int_0^\infty w_*(a) a \sinh a da = -\frac{1}{2} \log(\pi/2)$ now gives the asserted negative-shell saddle displacement.

To compare the negative shell with gamma damping, divide its density by $\mu_{\lambda, 1+u}$:

$$\frac{\lambda |w_s(a)| \cosh(ua)}{\mu_{\lambda, 1+u}(a)} = b(a) \Theta_\lambda(a) e^{(u-1)a} \frac{\cosh(ua)}{\cosh a}, \quad \Theta_\lambda(a) = \frac{1 - e^{-2a/\lambda}}{2a/\lambda}.$$

The finite-dimensional correction satisfies $0 < \Theta_\lambda(a) \leq 1$, by $1 - e^{-x} \leq x$. When $-1 < u \leq 1$, both $e^{(u-1)a}$ and $\cosh(ua)/\cosh a$ are at most 1, so the ratio is at most $b(a)$. When $1 \leq u \leq U$, the elementary inequality $\cosh(ua) \leq e^{(u-1)a} \cosh a$ bounds it by $b(a) e^{2(u-1)a} \leq b(a) e^{\epsilon a}$. Since $\epsilon A = o(1)$, uniformly on the negative shell,

$$b(a) e^{\epsilon a} = (1 - 2\epsilon(1+a))(1 + \epsilon a + O(\epsilon^2 a^2)) \leq 1 - c\epsilon.$$

Thus the taper retains a damping margin of order ϵ on every contour $-1 < u \leq U$.

It remains to check that the positive shell has opposite effects at the two saddle locations. At u_0 , $\sinh(u_0 a)/\cosh a \leq e^{(u_0-1)a}$, and hence

$$0 \leq \int_B^{B+1} w_B(a) a \sinh(u_0 a) da \leq (B+1) Q e^{(u_0-1)(B+1)}.$$

The amplitude in (34) has exponential slope strictly between $u_0 - 1$ and $U - 1$. Hence

$$Q e^{(u_0-1)B} = e^{-c\epsilon B}, \quad Q e^{(U-1)B} = e^{c\epsilon B}$$

for an absolute $c > 0$. Since $\epsilon B = \epsilon^{-2}$, while B and C_0 grow only polynomially in $1/\epsilon$, both separation quantities are $O(e^{-c'/\epsilon^2})$ for another absolute $c' > 0$. The same estimate controls the positive-shell saddle displacement. $\square$

The shells now determine the common envelope; it remains to impose the Fourier symmetries and select the signs. The first gamma pole occurs at $t = -i\lambda$, or $\zeta = -i$. Thus P_0 should vanish at $-i$, while $P_{\pm}$ should take the same positive value there. Reflection $\zeta \mapsto -\zeta$ should also exchange P_- with P_+ . These requirements lead to the following envelope and polynomials:

$$\begin{aligned} E_{\lambda}(t) &= \pi^{it/2} \Gamma\left(\frac{\lambda - it}{2}\right) e^{\lambda h_{\epsilon}(t/\lambda)}, \\ P_{\pm}(\zeta) &= 1 + \zeta^2 + \beta \pm i\zeta(1 + \zeta^2), \quad P_0(\zeta) = -(1 + \zeta^2), \\ X_{f_j}(t) &= E_{\lambda}(t) P_j(t/\lambda), \quad f_j(r) = \frac{r^{-\lambda}}{2\pi} \int_{\mathbb{R}} X_{f_j}(t) r^{it} dt \quad (j \in \{-, 0, +\}). \end{aligned} \quad (38)$$

Indeed, $P_-(-\zeta) = P_+(\zeta)$, P_0 is even, and

$$P_{\pm}(-i) = \beta > 0, \quad P_0(-i) = 0.$$

On the imaginary saddle branch,

$$\begin{aligned} P_+(iu) &= \beta + (1 - u)^2(1 + u), \\ P_-(iu) &= \beta + (1 - u)(1 + u)^2, \\ P_0(iu) &= u^2 - 1. \end{aligned} \quad (39)$$

Consequently $P_+(iu) > 0$ for every $u > -1$, whereas $\beta = u_0 - 1$ gives

$$P_-(iu_0) = \frac{\epsilon}{4} \left(1 - \left(2 + \frac{\epsilon}{4}\right)^2\right) < 0, \quad P_0(iu_0) = \left(1 + \frac{\epsilon}{4}\right)^2 - 1 > 0.$$

The saddle calculation below will transfer precisely these polynomial signs to the radial functions.

Lemma 4.3. *For every sufficiently small $\epsilon > 0$ and every integer $d \geq 1$, with $\lambda = d/2$, the inverse Mellin integrals in (38), initially defined for $r > 0$, extend to $f_j \in \mathcal{S}_{\text{rad}}(\mathbb{R}^d; \mathbb{R})$ for $j \in \{-, 0, +\}$. These extensions satisfy $\hat{f}_- = f_+$, $\hat{f}_0 = f_0$, $f_-(0) = f_+(0) > 0$, and $f_0(0) = 0$.*

Proof. Fix d and ϵ . Compact support of w makes h_{ϵ} entire, and on each horizontal line $t = s + i\tau$ it gives

$$|h_{\epsilon}((s + i\tau)/\lambda)| \leq 2 \int_0^{\infty} |w(a)| \cosh(a\tau/\lambda) da.$$

Thus the perturbation is bounded on every fixed horizontal strip. Uniformly for τ in compact pole-free intervals, the standard gamma asymptotics [DLMF, § 5.11] give

$$|X_{f_j}(s + i\tau)| \leq C_{d,\epsilon,\tau} (1 + |s|)^{(\lambda + \tau - 1)/2 + 3} e^{-\pi|s|/4}.$$

In particular, the vertical sides of rectangular contour shifts tend to zero. Shifting the Mellin contour upward arbitrarily far proves rapid decay as $r \rightarrow \infty$, also after differentiation. Shifting downward past the gamma poles

$$t = -i(\lambda + 2n), \quad n = 0, 1, 2, \dots,$$

gives an expansion in even powers r^{2n} , with a remainder of arbitrarily high order. Thus each f_j extends to a smooth radial Schwartz function on $\mathbb{R}^d$. Conjugate symmetry of its real-line Mellin data makes this extension real.

Because h_{ϵ} is even, the envelope obeys $m_{\lambda}(t)E_{\lambda}(-t) = E_{\lambda}(t)$. The polynomial reflection identities and the multiplier (10) therefore give

$$\hat{f}_- = f_+, \quad \hat{f}_0 = f_0. \quad (40)$$

Finally, the value at the origin is determined by the first pole of the downward-shifted contour. More generally, $\text{Res}_{z=-n} \Gamma(z) = (-1)^n/n!$ shows that the pole $t = -i(\lambda + 2n)$ contributes

$$2\pi^{\lambda/2+n} r^{2n} \frac{(-1)^n}{n!} e^{\lambda h_{\epsilon}(-i(1+2n/\lambda))} P_j(-i(1 + 2n/\lambda)). \quad (41)$$

At $n = 0$, evenness gives $h_\epsilon(-i) = h_\epsilon(i)$, and $P_\pm(-i) = \beta$, whereas $P_0(-i) = 0$. Consequently,

$$f_+(0) = f_-(0) = 2\pi^{\lambda/2} e^{\lambda h_\epsilon(i)} \beta > 0, \quad f_0(0) = 0. \quad \square \quad (42)$$

4.3. Saddle geometry. Recall u_0 and U from (34), and set

$$u_* = -1 + \frac{\log \lambda}{4\lambda}. \quad (43)$$

To determine the signs of the Mellin integrals (38), we associate each contour $t = \lambda(T + iu)$ with the radius $r = e^{v(u)}$ at which $E_\lambda(t)r^{it}$ is stationary in T . The decrease $D_u(T)$ of its logarithmic modulus must remain positive away from $T = 0$. We establish this first using the gamma contribution when $u_* \leq u \leq U$, and then using the positive shell w_B when $u \geq U$. The resulting saddle estimates will identify the sign of each integral with the sign of $P_j(iu)$.

Put

$$\eta = 1 + u > 0, \quad m = \frac{\lambda\eta}{2}, \quad \psi = (\log \Gamma)', \quad \psi^{(1)} = \psi'.$$

Take the branch of $\log \Gamma$ on $\operatorname{Re} z > 0$ which is real on the positive axis. On the contour $t = \lambda(T + iu)$, the logarithm of $E_\lambda(t)r^{it}$ is

$$\frac{i\lambda(T + iu)}{2} \log \pi + \log \Gamma\left(m - \frac{i\lambda T}{2}\right) + \lambda h_\epsilon(T + iu) + i\lambda(T + iu) \log r.$$

Its derivative with respect to T at $T = 0$ equals

$$i\lambda \left(\frac{1}{2} \log \pi - \frac{1}{2} \psi(m) - \int_0^\infty w(a) a \sinh(ua) da + \log r \right).$$

Thus $T = 0$ is stationary precisely when $r = e^{v(u)}$, where

$$\begin{aligned} v(u) &= -\frac{1}{2} \log \pi + \frac{1}{2} \psi(m) + \int_0^\infty w(a) a \sinh(ua) da, \\ V(u) = v'(u) &= \frac{\lambda}{4} \psi^{(1)}(m) + \int_0^\infty w(a) a^2 \cosh(ua) da. \end{aligned} \quad (44)$$

Differentiating the saddle log-radius gives $V(u) = v'(u)$. Thus $V(u) > 0$ makes $e^{v(u)}$ increase with u ; the quadratic decrease of the logarithmic modulus at $T = 0$ is $\lambda V(u) T^2/2$.

To separate the gamma and shell contributions to this decrease, recall the positive density $\mu_{\lambda,\eta}$ from (37). The standard log-gamma integral [DLMF, § 5.9], after substituting $a = \lambda s/2$, gives

$$\begin{aligned} G_{\lambda,\eta}(T) &:= \log \Gamma(m - i\lambda T/2) - \log \Gamma(m) + \frac{i\lambda T}{2} \psi(m) \\ &= \int_0^\infty (e^{iaT} - 1 - iaT) \mu_{\lambda,\eta}(a) da, \\ D_\gamma(T) &:= -\operatorname{Re} G_{\lambda,\eta}(T) = \int_0^\infty (1 - \cos(aT)) \mu_{\lambda,\eta}(a) da. \end{aligned} \quad (45)$$

In particular, the gamma function in the Mellin envelope always damps the integrand away from $T = 0$.

Normalize $E_\lambda(\lambda(T + iu))r^{i\lambda T}$ by $E_\lambda(i\lambda u)$ and set $r = e^{v(u)}$. The linear contributions cancel by the saddle equation, and (36) gives

$$\begin{aligned} \mathcal{L}_u(T) &:= \log \frac{E_\lambda(\lambda(T + iu))}{E_\lambda(i\lambda u)} + i\lambda T v(u) \\ &= G_{\lambda,\eta}(T) + \lambda \int_0^\infty w(a) \cosh(ua) (\cos(aT) - 1) da \\ &\quad + i\lambda \int_0^\infty w(a) \sinh(ua) (aT - \sin(aT)) da, \end{aligned} \quad (46)$$

$$\begin{aligned} D_u(T) &:= -\operatorname{Re} \mathcal{L}_u(T) \\ &= D_\gamma(T) + \lambda \int_0^\infty w(a) \cosh(ua) (1 - \cos(aT)) da. \end{aligned} \quad (47)$$

Equation (47) isolates the main difficulty: w_s reduces $D_u(T)$, while w_B increases it. We must prove $D_u(T) > 0$ for every $T \neq 0$ and $V(u) > 0$ for every $u \geq u_*$. The gamma density will dominate w_s on $[u_*, U]$, whereas w_B will dominate w_s on $[U, \infty)$.

The quadratic and cubic sizes of the phase are measured by

$$\begin{aligned} V_\gamma &= \frac{1}{\lambda} \int_0^\infty a^2 \mu_{\lambda, \eta}(a) da = \frac{\lambda}{4} \psi^{(1)}(m), \\ M_3 &= \frac{1}{\lambda} \int_0^\infty a^3 \mu_{\lambda, \eta}(a) da + \int_0^\infty (|w_s(a)| + w_B(a)) a^3 \cosh(ua) da. \end{aligned} \quad (48)$$

Here V_γ is the gamma contribution to $V(u)$, and M_3 bounds the cubic error after the signed shell contributions have been replaced by their absolute values.

For later comparisons, write

$$\begin{aligned} D_s(T) &= \lambda \int_{a_0}^A |w_s(a)| \cosh(ua) (1 - \cos(aT)) da, \\ D_B(T) &= \lambda \int_B^{B+1} w_B(a) \cosh(ua) (1 - \cos(aT)) da, \\ V_s &= \int_{a_0}^A |w_s(a)| a^2 \cosh(ua) da, \\ V_B &= \int_B^{B+1} w_B(a) a^2 \cosh(ua) da. \end{aligned} \quad (49)$$

In particular, $D_u = D_\gamma - D_s + D_B$ and $V(u) = V_\gamma - V_s + V_B$.

Lemma 4.4. *There are absolute constants $c, C > 0$ such that, for every $\lambda > 0$, $u > -1$ satisfying $\lambda(1+u) \geq 1$, and $T \in \mathbb{R}$, the phase and quantities defined in (44)–(48) satisfy*

$$\left| \mathcal{L}_u(T) + \frac{\lambda V(u)}{2} T^2 \right| \leq C \lambda M_3 |T|^3. \quad (50)$$

Moreover, writing $\eta = 1 + u$,

$$\frac{1}{2\eta} \leq V_\gamma \leq \frac{C}{\eta}, \quad (51)$$

$$\frac{1}{\lambda} \int_0^\infty a^3 \mu_{\lambda, \eta}(a) da \leq \frac{C}{\eta^2}, \quad (52)$$

$$D_\gamma(T) \geq c\lambda \min\left(\frac{T^2}{\eta}, |T|\right) \quad (T \in \mathbb{R}). \quad (53)$$

Proof. The globally valid Taylor estimates

$$e^{ix} - 1 - ix = -\frac{x^2}{2} + O(|x|^3), \quad x - \sin x = O(|x|^3)$$

applied to (45) and (46), and using $|\sinh(ua)| \leq \cosh(ua)$, give (50).

Since $m = \lambda\eta/2 \geq 1/2$, the variance and third-moment bounds are the standard uniform trigamma and polygamma estimates [DLMF, §§ 5.9, 5.11, 5.15]; indeed,

$$V_\gamma = \frac{\lambda}{4} \psi^{(1)}(m), \quad \frac{1}{\lambda} \int_0^\infty a^3 \mu_{\lambda, \eta}(a) da = -\frac{\lambda^2}{8} \psi^{(2)}(m).$$

To bound $D_\gamma(T)$ at every frequency, use $1 - e^{-x} \leq x$ in (37):

$$\mu_{\lambda, \eta}(a) \geq \frac{\lambda e^{-\eta a}}{2a^2}.$$

For $T \neq 0$, take $L = \min(\eta^{-1}, |T|^{-1})$. On $0 < a < L$, both $e^{-\eta a}$ and $(1 - \cos(aT))/(a^2 T^2)$ are bounded below by absolute positive constants. Therefore $D_\gamma(T) \geq c\lambda T^2 L$, which is (53); the case $T = 0$ is immediate. $\square$

The lower endpoint (43) ensures that the gamma shape parameter grows throughout the saddle analysis:

$$m \geq \frac{\lambda(1+u_*)}{2} = \frac{1}{8} \log \lambda.$$

Lemma 4.4 controls the gamma contribution and cubic remainder. We next show that, for $u_* \leq u \leq U$, the negative shell removes at most a $(1-c\epsilon)$ -fraction of the gamma damping, while w_B contributes nonnegative damping.

Lemma 4.5. *There is $\epsilon_0 > 0$ and an absolute $c > 0$ such that, for every $0 < \epsilon < \epsilon_0$, there are constants $C_\epsilon, \lambda_\epsilon > 0$ with the following property. For every $\lambda \geq \lambda_\epsilon$, every $u_* \leq u \leq U$, and $\eta = 1 + u$,*

$$\lambda |w_s(a)| \cosh(ua) \leq (1 - c\epsilon) \mu_{\lambda, \eta}(a) \quad (a \in [a_0, A]), \quad (54)$$

$$D_u(T) \geq c\epsilon D_\gamma(T) \quad (T \in \mathbb{R}), \quad (55)$$

$$\frac{c\epsilon}{\eta} \leq V(u) \leq \frac{C_\epsilon}{\eta}, \quad (56)$$

$$M_3 \leq \frac{C_\epsilon}{\eta^2}. \quad (57)$$

Moreover, $\lambda\eta \geq (\log \lambda)/4$.

Proof. Lemma 4.2 gives (54). Integrating that inequality against $1 - \cos(aT) \geq 0$ shows that the negative shell removes at most a $(1 - c\epsilon)$ -fraction of the gamma damping. The contribution of w_B is nonnegative. Therefore

$$D_u(T) \geq D_\gamma(T) - (1 - c\epsilon) \int_{a_0}^A (1 - \cos(aT)) \mu_{\lambda, \eta}(a) da \geq c\epsilon D_\gamma(T),$$

proving (55).

For the curvature, integrating (54) against a^2/λ and using (49) gives $V_s \leq (1 - c\epsilon)V_\gamma$, and hence

$$V(u) = V_\gamma - V_s + V_B \geq c\epsilon V_\gamma + V_B \geq \frac{c\epsilon}{\eta}.$$

For $u_* \leq u \leq U$, the positive-shell variance satisfies

$$V_B \leq (B+1)^2 Q e^{(U-1)(B+1)} = O_\epsilon(1).$$

Moreover, $\lambda\eta \geq (\log \lambda)/4$, so (51) gives

$$V(u) \leq V_\gamma + V_B = O_\epsilon(\eta^{-1}),$$

because $\eta \leq 2 + \epsilon/2$. This proves (56).

Similarly, the negative-shell third moment is bounded by the gamma third moment:

$$\int_{a_0}^A |w_s(a)| a^3 \cosh(ua) da \leq \frac{1}{\lambda} \int_0^\infty a^3 \mu_{\lambda, \eta}(a) da.$$

The positive-shell third moment is $O_\epsilon(1)$. Consequently, (52), $\lambda\eta \geq (\log \lambda)/4$, and $\eta \leq 2 + \epsilon/2$ give (57). $\square$

We have proved positive damping and curvature for $u_* \leq u \leq U$. When $u > U$, the factor $\cosh(ua)$ in (47) can make w_s overwhelm the gamma contribution. Put $\delta = u - 1$. For $T \neq 0$, the ratios $D_s(T)/(\lambda \min(T^2, 1))$ and $D_B(T)/(\lambda \min(T^2, 1))$ have respective sizes at most $C_0 e^{\delta A}$ and at least $Q e^{\delta B}$. The separation in Lemma 4.2 therefore makes w_B dominate w_s throughout $u \geq U$.

The interval support $[B, B+1]$ of w_B prevents frequency resonances: a shell concentrated at one a would have $D_B(T) = 0$ whenever $aT \in 2\pi\mathbb{Z}$, whereas

$$\int_B^{B+1} (1 - \cos(aT)) da = 1 - \text{sinc}(T/2) \cos((B + \frac{1}{2})T). \quad (58)$$

Here $\text{sinc}(x) = \sin(x)/x$, with $\text{sinc}(0) = 1$. Since $1 - |\text{sinc}(T/2)| \gg \min(T^2, 1)$, (58) is positive for every $T \neq 0$, uniformly at both small and large frequencies.

Define the separation error

$$\rho_\epsilon = (A + a_0^{-1})Q^{-1} \exp\left(-\frac{\epsilon}{2}(B - A)\right).$$

Lemma 4.2 gives $\rho_\epsilon = O(e^{-c/\epsilon^2}) = o(1)$ as $\epsilon \downarrow 0$. The next lemma bounds $D_s(T)/D_B(T)$ by $O(\rho_\epsilon)$, uniformly for $u \geq U$ and $T \neq 0$; it also gives the curvature and third-moment bounds required for the saddle approximation.

Lemma 4.6. *There are absolute constants $c, C > 0$ and $\epsilon_0 > 0$ such that, for every $0 < \epsilon < \epsilon_0$, there are constants $\lambda_\epsilon, C_\epsilon, c_\epsilon > 0$ with the following property. For every $\lambda \geq \lambda_\epsilon$, every $u \geq U$, and every $T \in \mathbb{R}$, writing $\delta = u - 1$, one has*

$$D_s(T) \leq C\lambda C_0 e^{\delta A} \min(T^2, 1), \quad (59)$$

$$D_B(T) \geq c\lambda Q e^{\delta B} \min(T^2, 1). \quad (60)$$

Consequently,

$$D_s(T) \leq C\rho_\epsilon D_B(T), \quad (61)$$

$$D_u(T) \geq D_\gamma(T) + cD_B(T), \quad (62)$$

$$cV_B \leq V(u) \leq CV_B. \quad (63)$$

The shell variance and third moments obey

$$cB^2 Q e^{\delta B} \leq V_B \leq (B + 1)^2 Q e^{\delta(B+1)}, \quad (64)$$

$$\begin{aligned} \int_{a_0}^A |w_s(a)| a^3 \cosh(ua) da &\leq CA\rho_\epsilon V_B, \\ \int_B^{B+1} w_B(a) a^3 \cosh(ua) da &\leq (B + 1)V_B. \end{aligned} \quad (65)$$

In particular,

$$M_3 \leq C_\epsilon V(u), \quad V(u) \geq c_\epsilon > 0 \quad (u \geq U). \quad (66)$$

Proof. For $u = 1 + \delta$, the explicit shell densities satisfy

$$|w_s(a)| \cosh(ua) \ll \frac{e^{\delta a}}{a^2}, \quad w_B(a) \cosh(ua) \asymp Q e^{\delta a}.$$

If $|T| \leq 1$, the bound $1 - \cos(aT) = O(a^2 T^2)$ gives

$$D_s(T) \ll \lambda T^2 \int_{a_0}^A e^{\delta a} da \ll \lambda A e^{\delta A} T^2.$$

If $|T| \geq 1$, use $1 - \cos(aT) = O(1)$:

$$D_s(T) \ll \lambda e^{\delta A} \int_{a_0}^A a^{-2} da \ll \lambda a_0^{-1} e^{\delta A}.$$

These two estimates give (59). On the positive shell, (58) yields

$$D_B(T) \gg \lambda Q e^{\delta B} \int_B^{B+1} (1 - \cos(aT)) da \gg \lambda Q e^{\delta B} \min(T^2, 1).$$

For $T \neq 0$, division gives (61), since

$$\frac{D_s(T)}{D_B(T)} \ll C_0 Q^{-1} e^{-\delta(B-A)} \leq \rho_\epsilon = o(1).$$

At $T = 0$, both damping terms vanish. Comparing their quadratic coefficients in (49) gives $V_s = O(\rho_\epsilon V_B)$. Choose ϵ_0 so that the implicit constant times ρ_ϵ is less than $1/2$. Then $D_u = D_\gamma - D_s + D_B$ proves (62), and $V(u) = V_\gamma - V_s + V_B$ gives

$$V_\gamma + cV_B \leq V(u) \leq V_\gamma + V_B.$$

Integrating $w_B(a) \cosh(ua) \asymp Q e^{\delta a}$ against a^2 gives (64). Since $\delta \geq \epsilon/2$,

$$V_B \gg B^2 Q e^{\epsilon B/2} = B^2 e^{c\epsilon B} \gg 1, \quad V_\gamma \ll \eta^{-1} \ll 1.$$

Thus $V_\gamma = O(V_B)$, so the preceding comparison of $V(u)$ with $V_\gamma + V_B$ gives (63). Since $a \leq A$ on the negative shell,

$$\int_{a_0}^A |w_s(a)| a^3 \cosh(ua) da \leq AV_s \ll A\rho_\epsilon V_B,$$

while $a \leq B + 1$ bounds the positive-shell third moment by $(B + 1)V_B$. This proves (65).

Since $\eta \geq 2 + \epsilon/2$, (51) and (52) bound the gamma third moment by $O(V_\gamma)$. The shell third moments are $O_\epsilon(V_B)$, while $V_\gamma = O(V_B)$ and (63) gives $V_B \asymp V(u)$. Hence $M_3 = O_\epsilon(V(u))$. Finally, (64) and (63) give $V(u) \gg V_B \gg B^2 Q e^{\epsilon B/2} =: c_\epsilon > 0$, proving (66). $\square$

Set $T_0 = (2(B + 1))^{-1}$.

Lemma 4.6 controls the total damping for $u \geq U$. To bound the tails of the saddle integral, we sharpen that control on three frequency ranges: $|T| \leq T_0$, where $D_u(T)$ is quadratic; $T_0 \leq |T| \leq \eta$, where w_B supplies a uniform positive floor; and $|T| \geq \eta$, where the gamma contribution also grows linearly.

Lemma 4.7. *There is $\epsilon_0 > 0$ and an absolute $c > 0$ such that, for every $0 < \epsilon < \epsilon_0$, there are constants $c_\epsilon, \lambda_\epsilon > 0$ with the following property. For every $\lambda \geq \lambda_\epsilon$ and $u \geq U$, set $\eta = 1 + u$ and $\delta = u - 1$. Then*

$$D_u(T) \geq c_\epsilon \lambda V(u) T^2 \quad (|T| \leq T_0), \quad (67)$$

$$D_u(T) \geq c_\epsilon \lambda Q e^{\delta B} \quad (T_0 \leq |T| \leq \eta), \quad (68)$$

$$D_u(T) \geq c \lambda |T| + c_\epsilon \lambda Q e^{\delta B} \quad (|T| \geq \eta). \quad (69)$$

Proof. First suppose that $|T| \leq T_0$. Since $|aT| \leq 1/2$ on $[B, B + 1]$, $1 - \cos(aT) \gg a^2 T^2$, and therefore

$$D_B(T) \gg \lambda V_B T^2.$$

Since $\eta \geq 2 + \epsilon/2$ and $|T| \leq \eta$, (53) and (51) also give

$$D_\gamma(T) \gg \frac{\lambda T^2}{\eta} \gg \lambda V_\gamma T^2.$$

Combining these contributions using (62) and (63) proves (67).

Next, if $T_0 \leq |T| \leq \eta$, then $\min(T^2, 1) \geq T_0^2$. Thus (62) and (60) give

$$D_u(T) \gg \lambda T_0^2 Q e^{\delta B} \gg_\epsilon \lambda Q e^{\delta B}.$$

This proves (68).

Finally, if $|T| \geq \eta$, then $\min(T^2, 1) = 1$ and (53) gives $D_\gamma(T) \gg \lambda |T|$. Adding the positive shell through (62) and (60) yields

$$D_u(T) \gg \lambda |T| + \lambda Q e^{\delta B},$$

which proves (69). $\square$

4.4. Global saddle asymptotics. The damping bounds now determine the exterior signs of f_+, f_-, f_0 . On each contour, the centered phase is quadratic near $T = 0$, the factor $P_j(T + iu)$ is asymptotic to $P_j(iu)$, and the remaining contour is negligible. We establish these claims separately for the gamma-controlled range $u_* \leq u \leq U$ and the positive-shell-controlled range $u \geq U$.

Lemma 4.8. *Fix $0 < \epsilon < \epsilon_0$, let $\lambda = d/2$, and recall u_0, U from (34) and u_* from (43). For $u > -1$ and $P \in \{P_+, P_-, P_0\}$, put*

$$I_{\lambda, P}(u) = \int_{\mathbb{R}} e^{\mathcal{L}_u(T)} P(T + iu) dT.$$

As $d \rightarrow \infty$,

$$I_{\lambda, P}(u) = P(iu) \sqrt{\frac{2\pi}{\lambda V(u)}} (1 + o_\epsilon(1)), \quad (70)$$

uniformly for $u \geq u_*$ when $P = P_+$, and uniformly for $u \geq u_0$ when $P = P_-$ or $P = P_0$. More precisely,

$$\sup_{u \geq u_*} \left| \frac{\sqrt{\lambda V(u)} I_{\lambda, P_+}(u)}{\sqrt{2\pi} P_+(iu)} - 1 \right| \rightarrow 0,$$

$$\max_{j \in \{-, 0\}} \sup_{u \geq u_0} \left| \frac{\sqrt{\lambda V(u)} I_{\lambda, P_j}(u)}{\sqrt{2\pi} P_j(iu)} - 1 \right| \rightarrow 0.$$

Proof. The poles of the integrand in (38) are $t = -i(\lambda + 2n)$, $n \geq 0$. Hence Lemma 4.3 allows the contour to be shifted to $t = \lambda(T + iu)$ whenever $u > -1$. At $r = e^{v(u)}$, the shifted Mellin inversion formula is

$$f_j(e^{v(u)}) = \frac{\lambda E_\lambda(i\lambda u)}{2\pi} e^{-(1+u)\lambda v(u)} I_{\lambda, P_j}(u). \quad (71)$$

The prefactor of $I_{\lambda, P_j}(u)$ is positive. By (39), $P_+(iu) > 0$ for $u > -1$, while $P_-(iu) < 0$ and $P_0(iu) > 0$ for $u \geq u_0$. Their fixed degrees and uniform lower bounds on those ranges imply

$$\frac{|P(T + iu)|}{|P(iu)|} \ll_\epsilon 1 + |T|^3, \quad \frac{P(T + iu)}{P(iu)} = 1 + O_\epsilon(|T| + |T|^3). \quad (72)$$

To compare $I_{\lambda, P}(u)$ with its Gaussian approximation, choose $K \rightarrow \infty$, and set

$$T_* = \frac{K}{\sqrt{\lambda V(u)}}.$$

By (50), the phase on $|T| \leq T_*$ is

$$\mathcal{L}_u(T) = -\frac{\lambda V(u)}{2} T^2 + O(\lambda M_3 |T|^3).$$

The approximation is uniform if its central interval shrinks and its cubic error tends to zero:

$$T_* = o_\epsilon(1), \quad \frac{K^3 M_3}{\sqrt{\lambda} V(u)^{3/2}} = o_\epsilon(1).$$

Under these conditions, (72) and the substitution $x = \sqrt{\lambda V(u)} T$ give

$$\int_{|T| \leq T_*} e^{\mathcal{L}_u(T)} P(T + iu) dT = P(iu) \sqrt{\frac{2\pi}{\lambda V(u)}} (1 + o_\epsilon(1)).$$

It remains to show that the integral over $|T| > T_*$ is $o_\epsilon(|P(iu)|/\sqrt{\lambda V(u)})$. We verify the central approximation and this tail bound separately on $[u_*, U]$ and $[U, \infty)$.

First suppose $u_* \leq u \leq U$, and put $\eta = 1 + u$, $L = \lambda\eta$. Equations (56) and (57) give

$$L \geq \frac{1}{4} \log \lambda, \quad V(u) \asymp_\epsilon \eta^{-1}, \quad M_3 \ll_\epsilon \eta^{-2}.$$

Choosing $K = L^{1/12}$, so that $K \rightarrow \infty$ and $K^3 = o(\sqrt{L})$, gives

$$\frac{T_*}{\eta} \ll_\epsilon L^{-5/12}, \quad \frac{K^3 M_3}{\sqrt{\lambda} V(u)^{3/2}} \ll_\epsilon L^{-1/4}.$$

The damping bounds (55) and (53) are quadratic for $|T| \leq \eta$ and linear for $|T| \geq \eta$. Consequently,

$$\sup_{|T| \leq T_*} \left| \mathcal{L}_u(T) + \frac{\lambda V(u)}{2} T^2 \right| \ll_\epsilon L^{-1/4},$$

$$\sqrt{\lambda V(u)} \int_{T_* \leq |T| \leq \eta} (1 + |T|^3) e^{-D_u(T)} dT \ll_\epsilon e^{-c_\epsilon K^2},$$

$$\sqrt{\lambda V(u)} \int_{|T| \geq \eta} (1 + |T|^3) e^{-D_u(T)} dT \ll_\epsilon e^{-c_\epsilon L}.$$

Both tail estimates tend to zero uniformly because $L \geq (\log \lambda)/4$, proving (70) for $u_* \leq u \leq U$.

Now suppose $u \geq U$, and write $\delta = u - 1$. By (63) and (66), the remote shell controls both the curvature and the third moment:

$$V(u) \asymp_\epsilon V_B \gg_\epsilon 1, \quad M_3 \ll_\epsilon V(u).$$

Take $K = \lambda^{1/12}$. Then

$$T_* \ll_\epsilon \lambda^{-5/12}, \quad \frac{K^3 M_3}{\sqrt{\lambda} V(u)^{3/2}} \ll_\epsilon \lambda^{-1/4}.$$

In particular, $T_* < T_0$ for all sufficiently large λ . The quadratic bound (67) on $|T| \leq T_0$ yields

$$\begin{aligned} \sup_{|T| \leq T_*} \left| \mathcal{L}_u(T) + \frac{\lambda V(u)}{2} T^2 \right| &\ll_\epsilon \lambda^{-1/4}, \\ \sqrt{\lambda V(u)} \int_{T_* \leq |T| \leq T_0} (1 + |T|^3) e^{-D_u(T)} dT &\ll_\epsilon e^{-c_\epsilon K^2}. \end{aligned} \quad (73)$$

For $T_0 \leq |T| \leq \eta$, the variance bound (64) and the damping estimate (68) give

$$\begin{aligned} \sqrt{\lambda V(u)} \int_{T_0 \leq |T| \leq \eta} (1 + |T|^3) e^{-D_u(T)} dT \\ \ll_\epsilon \sqrt{\lambda} \exp\left(\frac{B+1}{2} \delta + 4 \log(2 + \delta) - c_\epsilon \lambda Q e^{B\delta}\right) = o_\epsilon(1). \end{aligned} \quad (74)$$

The exponent on the right of (74) decreases in $\delta \geq \epsilon/2$, since its derivative is $(B+1)/2 + 4/(2+\delta) - c_\epsilon \lambda B Q e^{B\delta} < 0$ for sufficiently large λ . At $\delta = \epsilon/2$, that exponent is $-c_\epsilon \lambda + O_\epsilon(1)$. Thus the middle-frequency contribution tends to zero uniformly even as $u \rightarrow \infty$.

Finally, for $|T| \geq \eta$, (69) supplies the positive-shell damping and a linear gamma tail, so

$$\sqrt{\lambda V(u)} \int_{|T| \geq \eta} (1 + |T|^3) e^{-D_u(T)} dT \ll_\epsilon e^{-c_\epsilon \lambda}, \quad (75)$$

uniformly in δ : as in (74), the damping $-c_\epsilon \lambda Q e^{B\delta}$ absorbs the growth of $\sqrt{V(u)}$. Equations (73)–(75) prove the saddle formula on every $u \geq U$. $\square$

Corollary 4.9. *For every fixed $0 < \epsilon < \epsilon_0$, there is d_ϵ such that, for every integer $d \geq d_\epsilon$,*

$$\begin{aligned} f_+(r) &> 0 \quad (r \geq e^{v(u_*)}), \\ f_-(r) &< 0 \quad (r \geq e^{v(u_0)}), \\ f_0(r) &> 0 \quad (r \geq e^{v(u_0)}). \end{aligned} \quad (76)$$

Proof. For $u > -1$, the prefactor of $I_{\lambda, P_j}(u)$ in (71) is positive. Hence (70) identifies the sign of $f_j(e^{v(u)})$ with that of $P_j(iu)$. Equations (56) and (63) give $v'(u) = V(u) > 0$ on $[u_*, \infty)$, while (44) and the positive shell w_B give $v(u) \rightarrow \infty$. Thus $[u_*, \infty)$ parametrizes every radius $r \geq e^{v(u_*)}$, and $[u_0, \infty)$ parametrizes every radius $r \geq e^{v(u_0)}$. The signs in (39) now give (76). $\square$

4.5. Positivity and the sharp upper bound. Corollary 4.9 proves the required signs outside the saddle radii. To finish the construction, we must also show $f_+(r) > 0$ for $0 \leq r \leq r_* = e^{v(u_*)}$. Shifting the Mellin contour below $O(\log \lambda)$ gamma poles expresses $f_+(r)/f_+(0)$ as a truncated exponential series plus a uniformly negligible remainder.

Lemma 4.10. *Fix $0 < \epsilon < \epsilon_0$, let $\lambda = d/2$, set $r_* = e^{v(u_*)}$, and write $h'_1 = \int_0^\infty w(a) a \sinh a da$. As $d \rightarrow \infty$,*

$$\sup_{0 \leq r \leq r_*} \left| e^{\pi e^{2h'_1} r^2} \frac{f_+(r)}{f_+(0)} - 1 \right| \rightarrow 0.$$

In particular, $f_+(r) > 0$ on $[0, r_]$ for all sufficiently large d .*

Proof. First identify the range on which the truncated exponential series must approximate e^{-y} . Put

$$y = \pi e^{2h'_1} r^2, \quad H(u) = \int_0^\infty w(a) a \sinh(ua) da, \quad \eta_* = 1 + u_* = \frac{\log \lambda}{4\lambda}.$$

The normalized contour height u_* tends to -1 , the normalized height of the first gamma pole, while the gamma shape parameter $\lambda(1 + u_*)/2 = (\log \lambda)/8$ still diverges. This choice makes (70) applicable and keeps $y(r_*)$ logarithmic. Indeed, H is odd and, for fixed ϵ , has bounded derivative near -1 . Consequently

$$H(u_*) + h'_1 = H(-1 + \eta_*) - H(-1) = O_\epsilon(\eta_*).$$

By (44),

$$y(r_*) = \exp\left(\psi\left(\frac{\lambda\eta_*}{2}\right) + 2H(u_*) + 2h'_1\right).$$

The standard digamma asymptotic [DLMF, § 5.11] and $\lambda\eta_*/2 = (\log \lambda)/8$ therefore give

$$0 \leq y \leq y(r_*) = \frac{1}{8} \log \lambda + O_\epsilon(1). \quad (77)$$

Since e^{-y} can be as small as a negative power of λ on (77), the approximation error must be controlled relative to e^{-y} .

To recover enough exponential-series coefficients, set

$$N = \lceil \log \lambda \rceil, \quad p = N + \frac{1}{2}, \quad \kappa = 1 + \frac{2p}{\lambda}.$$

Since $p = N + 1/2$, the contour $t = s - i(\lambda + 2p)$ lies strictly between consecutive gamma poles. Furthermore, $p \asymp \log \lambda$ makes the exponential-series tail negligible on (77). To justify shifting to this contour, we verify that the multiplier $e^{\lambda h_\epsilon(t/\lambda)}$ preserves a uniform exponential decay margin. Indeed, $pA/\lambda = o_\epsilon(1)$, so the negative-shell taper satisfies

$$b(a)e^{2pa/\lambda} \leq 1 - c\epsilon \quad (a_0 \leq a \leq A)$$

for an absolute $c > 0$ and all sufficiently large λ . For every intermediate contour height $0 \leq q \leq \kappa$, the positive shell contributes nonpositively to $\operatorname{Re} h_\epsilon(s/\lambda - iq) - h_\epsilon(iq)$, and the negative shell gives

$$\begin{aligned} \lambda[\operatorname{Re} h_\epsilon(s/\lambda - iq) - h_\epsilon(iq)] &\leq \frac{(1 - c\epsilon)\lambda}{2} \int_0^\infty \frac{1 - \cos(as/\lambda)}{a^2} da \\ &= (1 - c\epsilon) \frac{\pi|s|}{4}. \end{aligned}$$

Standard gamma asymptotics [DLMF, § 5.11] supply the complementary factor $e^{-\pi|s|/4}$. Hence the integrand decays like $e^{-c_\epsilon|s|}$ on the vertical sides, permitting the shift in (38) to $t = s - i(\lambda + 2p)$. This shift crosses exactly the poles $t = -i(\lambda + 2n)$, $0 \leq n \leq N$.

Applying the residue formula (41) and the origin value (42) gives

$$\frac{f_+(r)}{f_+(0)} = \sum_{n=0}^N \frac{(-y)^n}{n!} A_{\lambda,n} + \mathcal{R}_{\lambda,p}(r), \quad (78)$$

$$A_{\lambda,n} = e^{\lambda[h_\epsilon(i(1+2n/\lambda)) - h_\epsilon(i)] - 2nh'_1} \frac{P_+(-i(1+2n/\lambda))}{\beta},$$

$$|A_{\lambda,n} - 1| \ll_\epsilon \frac{n(1+n)}{\lambda} \quad (0 \leq n \leq N). \quad (79)$$

Indeed, Taylor expansion at $u = 1$ gives

$$\lambda \left[h_\epsilon\left(i\left(1 + \frac{2n}{\lambda}\right)\right) - h_\epsilon(i) \right] = 2nh'_1 + O_\epsilon\left(\frac{n^2}{\lambda}\right), \quad \frac{P_+(-i(1+2n/\lambda))}{\beta} = 1 + O_\epsilon\left(\frac{n}{\lambda}\right),$$

uniformly for $n \leq N$; here $N^2/\lambda = o(1)$.

For $r > 0$, the remainder $\mathcal{R}_{\lambda,p}(r)$ in (78) is the integral over $t = s - i(\lambda + 2p)$:

$$\begin{aligned} \mathcal{R}_{\lambda,p}(r) &= \frac{\pi^{\lambda/2+p} r^{2p}}{2\pi f_+(0)} \int_{\mathbb{R}} \pi^{is/2} \Gamma(-p - is/2) \\ &\quad \cdot e^{\lambda h_\epsilon(s/\lambda - i\kappa)} P_+(s/\lambda - i\kappa) r^{is} ds. \end{aligned}$$

The gamma reflection and product estimates [DLMF, §§ 5.5, 5.8], together with $p \in \mathbb{Z} + \frac{1}{2}$, give the standard bound

$$|\Gamma(-p - is/2)| \ll \frac{e^{-\pi|s|/4}}{\Gamma(1+p)},$$

$$\lambda[\operatorname{Re} h_\epsilon(s/\lambda - i\kappa) - h_\epsilon(i\kappa)] \leq (1 - c\epsilon) \frac{\pi|s|}{4}.$$

Since $P_+(s/\lambda - i\kappa)/\beta = O_\epsilon(1 + |s|^3)$, integration in s now gives

$$|\mathcal{R}_{\lambda,p}(r)| \ll_\epsilon \frac{y^p}{\Gamma(1+p)}. \quad (80)$$

Here we used

$$\lambda[h_\epsilon(i\kappa) - h_\epsilon(i)] = 2ph'_1 + O_\epsilon(p^2/\lambda), \quad (\pi r^2)^p e^{2ph'_1} = y^p,$$

and absorbed $p^2/\lambda = o(1)$. The estimate extends to $r = 0$ by continuity.

To compare (78) with e^{-y} , we bound the coefficient errors $A_{\lambda,n} - 1$, the contour remainder $\mathcal{R}_{\lambda,p}(r)$, and the omitted terms $\sum_{n>N} (-y)^n/n!$. Equations (79), (80), and (77) give

$$e^y \sum_{n=0}^N \frac{y^n}{n!} |A_{\lambda,n} - 1| \ll_\epsilon \frac{(1+y)^2 e^{2y}}{\lambda} \ll_\epsilon \frac{(\log \lambda)^2}{\lambda^{3/4}},$$

$$e^y |\mathcal{R}_{\lambda,p}(r)| \ll_\epsilon \frac{e^y y^p}{\Gamma(1+p)},$$

$$e^y \sum_{n>N} \frac{y^n}{n!} \ll \frac{e^y y^{N+1}}{(N+1)!}.$$
(81)

The first estimate follows from $\sum_{n \geq 0} n(1+n)y^n/n! = (y^2 + 2y)e^y$. For the two tails, put $L = \log \lambda$. Since $y \leq L/8 + O_\epsilon(1)$ and $p = L + O(1)$, Stirling's formula [DLMF, § 5.11] gives

$$\log \frac{e^y y^p}{\Gamma(1+p)} \leq \left(\frac{1}{8} + 1 - \log 8 \right) L + O_\epsilon(\log L).$$

The same estimate holds with p replaced by $N+1$. Since $1/8 + 1 - \log 8 < 0$, both tails are smaller than the coefficient error in (81). Therefore (78) yields, uniformly on $0 \leq r \leq r_*$,

$$\frac{f_+(r)}{f_+(0)} = e^{-y} \left(1 + O_\epsilon \left(\frac{(\log \lambda)^2}{\lambda^{3/4}} \right) \right) > 0 \quad (0 \leq r \leq r_*). \quad \square \quad (82)$$

Proof of Theorem 4.1. Set $R_{\epsilon,d} = e^{v(u_0)}$. The Fourier identities and origin values follow from (40) and (42). Corollary 4.9 gives the required exterior signs, while (82) supplies positivity of f_+ on the remaining interval. Thus

$$\begin{aligned} \widehat{f}_- &= f_+ > 0, & f_-(0) &= f_+(0) > 0, \\ \widehat{f}_0 &= f_0, & f_0(0) &= 0, & f_-(r) &< 0 < f_0(r) \quad (r \geq R_{\epsilon,d}). \end{aligned} \quad (83)$$

For fixed ϵ , the saddle equation (44) and the digamma asymptotic [DLMF, § 5.11] give

$$\lim_{d \rightarrow \infty} \frac{R_{\epsilon,d}}{\sqrt{d}} = \sqrt{\frac{1+u_0}{4\pi}} \exp \left(\int_0^\infty w(a) a \sinh(u_0 a) da \right). \quad (84)$$

The two shell contributions in Lemma 4.2 give

$$\int_0^\infty w(a) a \sinh(u_0 a) da = -\frac{1}{2} \log \frac{\pi}{2} + O(\epsilon).$$

Since $u_0 \rightarrow 1$, (33) and (84) give

$$\lim_{\epsilon \downarrow 0} \lim_{d \rightarrow \infty} \frac{R_{\epsilon,d}}{\sqrt{d}} = \frac{1}{\pi}. \quad \square \quad (85)$$

Proof of Theorem 1.1. By (83), the dilation $F_{\epsilon,d}(x) = f_-(R_{\epsilon,d}x)$ is admissible and satisfies $F_{\epsilon,d}(0)/\widehat{F}_{\epsilon,d}(0) = R_{\epsilon,d}^d$ by Fourier scaling. Inserting $F_{\epsilon,d}$ into (3) gives

$$\text{LP}_d \leq \frac{v_d}{2^d} R_{\epsilon,d}^d. \quad (86)$$

The universal lower bound (28), the admissible upper bound (86), and (85), together with $v_d^{1/d} \sqrt{d} \rightarrow \sqrt{2\pi e}$, imply

$$\sqrt{\frac{e}{2\pi}} \leq \liminf_{d \rightarrow \infty} \text{LP}_d^{1/d} \leq \limsup_{d \rightarrow \infty} \text{LP}_d^{1/d} \leq \frac{\sqrt{2\pi e}}{2\pi} = \sqrt{\frac{e}{2\pi}}. \quad \square$$

Proof of Theorem 1.2. The lower bound is Proposition 3.7. For the upper bound, define $g_{\epsilon,d,-} = f_+ - f_-$ and $g_{\epsilon,d,+} = f_0$. Equation (40) and Fourier inversion give

$$\widehat{g}_{\epsilon,d,-} = -g_{\epsilon,d,-}, \quad \widehat{g}_{\epsilon,d,+} = g_{\epsilon,d,+}.$$

By (83), both $g_{\epsilon,d,-}$ and $g_{\epsilon,d,+}$ vanish at the origin and are strictly positive outside $R_{\epsilon,d}$; in particular, neither is zero. The definition (6) therefore gives $A_\zeta(d) \leq R_{\epsilon,d}$ for both signs. Combining this bound with (85) gives

$$\frac{1}{\pi} \leq \liminf_{d \rightarrow \infty} \frac{A_\zeta(d)}{\sqrt{d}} \leq \limsup_{d \rightarrow \infty} \frac{A_\zeta(d)}{\sqrt{d}} \leq \lim_{\epsilon \downarrow 0} \lim_{d \rightarrow \infty} \frac{R_{\epsilon,d}}{\sqrt{d}} = \frac{1}{\pi}. \quad \square$$

APPENDIX A. COMPARISON OF THE SIGN-UNCERTAINTY CONSTANTS

For an anti-self-Fourier radial function g , the central Mellin moment $M_g(d/2)$ vanishes. Integrating the radial tail of g therefore produces a self-Fourier function with a strictly smaller last-sign radius.

Proposition A.1. *Let $d \geq 1$, put $\lambda = d/2$, and suppose that $0 \neq g \in L_{\text{rad}}^1(\mathbb{R}^d; \mathbb{R})$ satisfies $\widehat{g} = -g$ and $g(0) = 0$. Define $T_d g(0) = 0$ and, for $x \neq 0$, set*

$$(T_d g)(x) = \frac{\lambda}{2} \int_1^\infty t^{\lambda-1} g(tx) dt, \quad (87)$$

where g denotes its continuous Fourier-inversion representative. Then $T_d g$ is nonzero, continuous, and integrable, with

$$\widehat{T_d g} = T_d g, \quad (T_d g)(0) = 0, \quad \|T_d g\|_1 \leq \frac{1}{2} \|g\|_1.$$

If $r(g) < \infty$, then $r(T_d g) < r(g)$; if g is Schwartz, then so is $T_d g$. Consequently, the constants in (6) satisfy

$$A_+(d) < A_-(d) \quad (d \geq 1).$$

Proof. Since $g = -\widehat{g} \in L^1$, Fourier inversion makes g bounded and continuous. In particular, $\int_{\mathbb{R}^d} |g(x)| |x|^{-\lambda} dx < \infty$, by splitting the integral at $|x| = 1$. The Mellin–Fourier identity in (9) was established for Schwartz functions, so we first verify its central consequence directly for the integrable eigenfunction g . Set $J(t) = \int_{\mathbb{R}^d} g(x) e^{-\pi t |x|^2} dx$. Gaussian duality and Tonelli’s theorem give

$$J(t) = -t^{-\lambda} J(1/t), \quad \int_0^\infty t^{\lambda/2-1} |J(t)| dt \leq \frac{\Gamma(\lambda/2)}{\pi^{\lambda/2}} \int_{\mathbb{R}^d} |g(x)| |x|^{-\lambda} dx < \infty.$$

Hence the substitution $t \mapsto 1/t$ makes $\int_0^\infty t^{\lambda/2-1} J(t) dt$ equal to its negative. Gaussian integration and polar coordinates therefore give the central Mellin cancellation

$$M_g(\lambda) = \int_0^\infty g(r) r^{\lambda-1} dr = 0. \quad (88)$$

The integral defining $T_d g(x)$ converges absolutely for $x \neq 0$, since g is integrable and $s^{\lambda-1} \ll_x s^{d-1}$ for $s \geq |x|$. Tonelli’s theorem and $d = 2\lambda$ yield $\|T_d g\|_1 \leq (\lambda/2) \|g\|_1 \int_1^\infty t^{-\lambda-1} dt = \|g\|_1/2$,

also justifying Fourier transformation under the integral. Fourier scaling and (88) give

$$\begin{aligned}\widehat{T_d g}(\xi) &= \frac{\lambda}{2} \int_1^\infty t^{\lambda-d-1} \widehat{g}(\xi/t) dt \\ &= -\frac{\lambda}{2} \int_0^1 s^{\lambda-1} g(s\xi) ds \\ &= \frac{\lambda}{2} \int_1^\infty s^{\lambda-1} g(s\xi) ds = T_d g(\xi).\end{aligned}\tag{89}$$

The small-scale representation in (89) extends continuously to $x = 0$, with value $-g(0)/2 = 0$. Differentiating the large-scale representation gives $(x \cdot \nabla + \lambda)T_d g = -\lambda g/2$, so $T_d g \neq 0$. If g is Schwartz, differentiating the small-scale representation gives smoothness at the origin, and differentiating the large-scale representation gives rapid decay at infinity; thus $T_d g$ is Schwartz.

If $R = r(g) < \infty$, then $R > 0$, since otherwise $g \geq 0$ and $\int g = \widehat{g}(0) = -g(0) = 0$ would force $g = 0$. Hence $g(s) \geq 0$ for $s \geq R$, and (5) and (87) give

$$(T_d g)(r) = \frac{\lambda}{2} r^{-\lambda} \int_r^\infty s^{\lambda-1} g(s) ds > 0 \quad (r \geq R).$$

Indeed, equality at any $r \geq R$ would make both g and $\widehat{g} = -g$ compactly supported, contradicting Fourier analyticity. In particular, $T_d g(R) > 0$, so continuity gives $r(T_d g) < R$. Applying this strict decrease to an attained radial negative-sign extremizer [CG19, Thm. 1.4] gives $A_+(d) < A_-(d)$. $\square$

REFERENCES

- [AJCHLT20] N. Afkhami-Jeddi, H. Cohn, T. Hartman, D. de Laat, and A. Tajdini, *High-dimensional sphere packing and the modular bootstrap*, J. High Energy Phys. **12** (2020), article 066, doi:10.1007/JHEP12(2020)066.
- [Ahl79] L. V. Ahlfors, *Complex Analysis*, 3rd ed., McGraw-Hill, New York, 1979.
- [BCK10] J. Bourgain, L. Clozel, and J.-P. Kahane, *Principe d’Heisenberg et fonctions positives*, Ann. Inst. Fourier (Grenoble) **60** (2010), 1215–1232, doi:10.5802/aif.2552.
- [Coh02] H. Cohn, *New upper bounds on sphere packings. II*, Geom. Topol. **6** (2002), 329–353, doi:10.2140/gt.2002.6.329.
- [CDG24] H. Cohn, D. Dong, and F. Gonçalves, *Sign uncertainty principles and low-degree polynomials*, Proc. Amer. Math. Soc. Ser. B **11** (2024), 224–228, doi:10.1090/bproc/219.
- [CE03] H. Cohn and N. Elkies, *New upper bounds on sphere packings. I*, Ann. of Math. (2) **157** (2003), 689–714, doi:10.4007/annals.2003.157.689.
- [CG19] H. Cohn and F. Gonçalves, *An optimal uncertainty principle in twelve dimensions via modular forms*, Invent. Math. **217** (2019), 799–831, doi:10.1007/s00222-019-00875-4.
- [CKMRV17] H. Cohn, A. Kumar, S. D. Miller, D. Radchenko, and M. S. Viazovska, *The sphere packing problem in dimension 24*, Ann. of Math. (2) **185** (2017), 1017–1033, doi:10.4007/annals.2017.185.3.8.
- [CKMRV22] H. Cohn, A. Kumar, S. D. Miller, D. Radchenko, and M. S. Viazovska, *Universal optimality of the E_8 and Leech lattices and interpolation formulas*, Ann. of Math. (2) **196** (2022), 983–1082, doi:10.4007/annals.2022.196.3.3.
- [CZ14] H. Cohn and Y. Zhao, *Sphere packing bounds via spherical codes*, Duke Math. J. **163** (2014), 1965–2002, doi:10.1215/00127094-2738857.
- [Edw25] R. Edwin, *Fourier inequalities and sign uncertainty*, preprint, 2025, arXiv:2505.15994.
- [GOSR21] F. Gonçalves, D. Oliveira e Silva, and J. P. G. Ramos, *On regularity and mass concentration phenomena for the sign uncertainty principle*, J. Geom. Anal. **31** (2021), 6080–6101, doi:10.1007/s12220-020-00519-7.
- [GOSS17] F. Gonçalves, D. Oliveira e Silva, and S. Steinerberger, *Hermite polynomials, linear flows on the torus, and an uncertainty principle for roots*, J. Math. Anal. Appl. **451** (2017), 678–711, doi:10.1016/j.jmaa.2017.02.030.
- [Gor00] D. V. Gorbachev, *Extremal problem for entire functions of exponential spherical type, connected with the Levenshtein bound on the sphere packing density in $\mathbb{R}^n$* , Izv. Tula State Univ. Ser. Math. Mech. Inform. **6** (2000), 71–78; in Russian.
- [HMR19] T. Hartman, D. Mazáč, and L. Rastelli, *Sphere packing and quantum gravity*, J. High Energy Phys. **12** (2019), article 048, doi:10.1007/JHEP12(2019)048.
- [KL78] G. A. Kabatianskii and V. I. Levenshtein, *On bounds for packings on a sphere and in space*, Problems Inform. Transmission **14** (1978), 1–17.
- [Lev79] V. I. Levenshtein, *On bounds for packings in n -dimensional Euclidean space*, Soviet Math. Dokl. **20** (1979), 417–421.

- [DLMF] F. W. J. Olver et al., eds., *NIST Digital Library of Mathematical Functions*, Release 1.2.7, National Institute of Standards and Technology, 2026, dlmf.nist.gov.
- [PK01] R. B. Paris and D. Kaminski, *Asymptotics and Mellin–Barnes Integrals*, Encyclopedia of Mathematics and its Applications, vol. 85, Cambridge University Press, Cambridge, 2001.
- [RV19] D. Radchenko and M. S. Viazovska, *Fourier interpolation on the real line*, Publ. Math. Inst. Hautes Études Sci. **129** (2019), 51–81, [doi:10.1007/s10240-018-0101-z](https://doi.org/10.1007/s10240-018-0101-z).
- [SZ24] N. T. Sardari and M. Zargar, *New upper bounds for spherical codes and packings*, Math. Ann. **389** (2024), 3653–3703, [doi:10.1007/s00208-023-02738-z](https://doi.org/10.1007/s00208-023-02738-z).
- [SW71] E. M. Stein and G. Weiss, *Introduction to Fourier Analysis on Euclidean Spaces*, Princeton Mathematical Series, vol. 32, Princeton University Press, Princeton, 1971.
- [Via17] M. S. Viazovska, *The sphere packing problem in dimension 8*, Ann. of Math. (2) **185** (2017), 991–1015, [doi:10.4007/annals.2017.185.3.7](https://doi.org/10.4007/annals.2017.185.3.7).
- [Z24] M. Zargar, *Stiefel manifolds and upper bounds for spherical codes and packings*, preprint, 2024, [arXiv:2407.10697](https://arxiv.org/abs/2407.10697).

CHAPTER 2

Improved Bounds for Binary and Spherical Codes

ABSTRACT. We construct two-point linear-programming certificates that improve the high-dimensional bounds for unrestricted binary and spherical codes. For every fixed relative distance $0 < \delta < 1/2$, our binary-code bound strictly improves the optimized McEliece–Rodemich–Rumsey–Welch exponent. For every fixed maximum inner product $0 < s < 1$, our spherical-code bound strictly improves the optimized Kabatianskii–Levenshtein exponent, including its spherical-cap optimization. These are the first improvements to the respective general high-dimensional exponents since 1977 and 1978. In the limit $s \rightarrow 1$, the spherical construction also recovers the optimal sphere-packing exponent of the Euclidean Cohn–Elkies linear program, obtained independently in a companion paper.

CONTENTS

1. Introduction
2. The first binary bound
3. The optimized binary bound
4. Representation graphs for spherical codes
5. Spherical harmonics and stabilizer representations
6. The spherical transition graph
7. Asymptotic spherical-code bounds
8. Sphere packings
- Appendix A. Orthogonal representations and asymptotic dimensions
- Appendix B. The spherical-to-Euclidean limit
- References

1. INTRODUCTION

A binary code $C \subseteq \{0, 1\}^n$ has Hamming distance $d_H(x, y)$, the number of coordinates in which two words differ; its minimum distance is the smallest such value for distinct codewords. Let $A_2(n, d)$ denote the largest code size with minimum distance at least d . A spherical code $C \subseteq \mathbb{S}^{n-1}$ has maximum inner product at most s when $\langle x, y \rangle \leq s$ for distinct points; let $A(n, s)$ denote its largest size. For fixed relative distance $0 < \delta < 1/2$ and maximum inner product $0 < s < 1$, the corresponding asymptotic rates are

$$R_2(\delta) = \limsup_{n \rightarrow \infty} \frac{1}{n} \log_2 A_2(n, \lceil \delta n \rceil), \quad R_{\text{sph}}(s) = \limsup_{n \rightarrow \infty} \frac{1}{n} \log_2 A(n, s).$$

The best previous general bounds on these rates are due to McEliece, Rodemich, Rumsey, and Welch (MRRW) for unrestricted binary codes and to Kabatianskii and Levenshtein for spherical codes [MRRW77, KL78]. Both arise from Delsarte's two-point programs [Del72, Del73, DGS77]. A Delsarte certificate is a polynomial $F(t) = \sum_{j=0}^M f_j P_j(t)$ in the Krawtchouk, Hahn, or Gegenbauer basis associated with the code space, with $P_0 = 1$. If $f_0 > 0$, $f_j \geq 0$ for $1 \leq j \leq M$, and $F(t) \leq 0$ at every normalized inner product allowed by the distance constraint, then $|C| \leq F(1)/f_0$. Our constructions verify positivity directly through projection Gram factorizations, without requiring explicit orthogonal-polynomial formulas or a converse characterization of positive-definite kernels.

In the classical spectral construction, each retained harmonic space contributes a single vector associated with a code point [BN06]. We instead associate a d_E -dimensional subspace with each point x , inside a common D -dimensional ambient space. The subspaces move with the points: a symmetry carrying x to y carries the subspace at x to the subspace at y . If P_x denotes the corresponding orthogonal projection, the overlap $\text{tr}(P_x P_y)$ is still a scalar function of the distance. Write s for the largest permitted normalized inner product and Λ for the largest eigenvalue of the associated weighted transition matrix. If $\Lambda > s$, the projection bound (50) gives

$$|C| \leq \frac{1 - s}{\Lambda - s} \frac{D}{d_E}.$$

Thus an exponentially large projection rank improves the rate provided Λ remains bounded away from the threshold s . For binary codes, the transition matrices are tridiagonal; for spherical codes, the full construction uses weighted graphs with several degree indices.

Bachoc and Vallentin also exploit higher harmonics of the rotations fixing a point [BV08, §3]. Their fixed base point yields matrix-valued, three-point semidefinite programs. Here the subspace moves with each code point, giving scalar two-point certificates whose bounds contain its dimension through D/d_E .

1.1. Binary codes: statement of the result. We first recall the classical binary-code bounds and then state the two constructions needed to improve the fully optimized second MRRW bound. The whole-cube construction improves the first MRRW bound, while the constant-weight construction handles parameter ranges where optimization over layers gives a stronger classical bound.

With $0 \log_2 0 = 0$, define the binary entropy function H_2 and an auxiliary function g on $[0, 1]$ by

$$H_2(u) = -u \log_2 u - (1 - u) \log_2 (1 - u), \quad g(v) = H_2 \left(\frac{1 - \sqrt{1 - v}}{2} \right).$$

The first MRRW bound is [MRRW77]

$$M_1(\delta) = H_2 \left(\frac{1}{2} - \sqrt{\delta(1 - \delta)} \right). \quad (1)$$

In particular, $R_2(\delta) \leq M_1(\delta)$. To state the whole-cube improvement, introduce normalized degree parameters $0 \leq b < a \leq 1/2$. In the finite construction, $a = L/n$ records the largest retained Fourier degree, meaning the number of coordinates in a Fourier monomial, and $b = k/n$

records the degree of the Boolean harmonic subspace attached to each code point. The spaces are constructed in §2. Set

$$\Gamma_H(a, b) = \frac{2(a-b)(1-a-b)}{\sqrt{a(1-a)}}.$$

The whole-cube construction gives the tridiagonal matrix (20), whose limiting largest eigenvalue is $\Gamma_H(a, b)$ by Lemma 2.2. Its resulting exponent is

$$\kappa_H(\delta) = \inf_{\substack{0 \leq b < a \leq 1/2 \\ \Gamma_H(a, b) > 1-2\delta}} (H_2(a) - H_2(b)). \quad (2)$$

Here $H_2(a)$ and $H_2(b)$ are the dimension exponents of the retained ambient space and attached harmonic subspace, respectively. Theorem 2.3 proves that, for every $0 < \delta < 1/2$,

$$R_2(\delta) \leq \kappa_H(\delta) < M_1(\delta).$$

The second MRRW bound optimizes the classical construction over constant-weight layers [MRRW77]. Its scalar parameter τ reparametrizes the polynomial degree: on the classical spectral boundary, a largest retained degree un corresponds to $\tau = 2\sqrt{u(1-u)}$. The objective is

$$F_\delta(\tau) = 1 + g(\tau^2) - g(\tau^2 + 2\delta\tau + 2\delta), \quad 0 \leq \tau \leq 1 - 2\delta.$$

The endpoints are

$$F_\delta(1 - 2\delta) = M_1(\delta), \quad F_\delta(0) = 1 - g(2\delta),$$

and the second MRRW bound is the full optimized exponent

$$M_2(\delta) = \min_{0 \leq \tau \leq 1-2\delta} F_\delta(\tau). \quad (3)$$

Thus $R_2(\delta) \leq M_2(\delta) \leq \min\{M_1(\delta), F_\delta(0)\}$. The minimum over all τ can occur in the interior and be strictly smaller than both endpoint values. Therefore, improving M_1 alone need not improve M_2 , so we refine the constant-weight construction as well.

A weight- w layer of the binary cube consists of all binary words having exactly w coordinates equal to 1. Fixing one such word partitions the coordinates into its w -element support and its $(n-w)$ -element complement. The construction attaches the tensor product of Boolean harmonic spaces of degrees p and q on these two coordinate sets and retains layer degrees up to L . Write $\alpha = w/n$, $\beta = p/n$, $\gamma = q/n$, and $u = L/n$ for the normalized layer weight, harmonic degrees, and largest retained layer degree, respectively. §3 constructs these spaces and an associated tridiagonal matrix. The parameter ranges are

$$\frac{\delta}{2} < \alpha < \frac{1}{2}, \quad 0 \leq \beta < \frac{\alpha}{2}, \quad 0 \leq \gamma < \frac{1-\alpha}{2}, \quad (4)$$

and

$$\beta + \gamma < u < \min\{\alpha, \alpha - \beta + \gamma, 1 - \alpha + \beta - \gamma\}. \quad (5)$$

The following function determines which parameter choices yield a code bound; its origin is explained in Lemma 3.5. Introduce the affine coordinates

$$z = 1 - 2u, \quad m = 1 - 2\alpha, \quad \zeta = 1 - 2\beta - 2\gamma, \quad \xi = 1 - 2\alpha + 2\beta - 2\gamma.$$

Define

$$\Lambda_{\alpha, \beta, \gamma}(u) = \frac{(\zeta\xi - mz^2)^2}{z^2(1-m^2)(1-z^2)} + \frac{(z^2 - \xi^2)(\zeta^2 - z^2)}{z^2(1-m^2)\sqrt{1-z^2}}. \quad (6)$$

Two weight- αn words at Hamming distance δn have normalized layer-distance coordinate $1 - \delta/(2\alpha(1-\alpha))$; see (28). The construction applies when the preceding function exceeds this coordinate:

$$\Lambda_{\alpha, \beta, \gamma}(u) > 1 - \frac{\delta}{2\alpha(1-\alpha)}. \quad (7)$$

Let $\mathcal{D}_\delta$ be the set of all $(\alpha, \beta, \gamma, u)$ in the ranges (4)–(5) that also satisfy (7). The resulting unrestricted-code exponent from constant-weight (CW) layers is

$$\kappa_{\text{CW}}(\delta) = \inf_{(\alpha, \beta, \gamma, u) \in \mathcal{D}_\delta} \left\{ 1 - H_2(\alpha) + H_2(u) - \alpha H_2\left(\frac{\beta}{\alpha}\right) - (1-\alpha) H_2\left(\frac{\gamma}{1-\alpha}\right) \right\}. \quad (8)$$

The four terms have a direct dimension interpretation. Passing from the whole cube to a weight- αn layer costs $1 - H_2(\alpha)$, since that layer has $2^{(H_2(\alpha)+o(1))n}$ words. The retained ambient space has dimension $2^{(H_2(u)+o(1))n}$, while the subspace attached to a code point has dimension

$$2^{(\alpha H_2(\beta/\alpha)+(1-\alpha)H_2(\gamma/(1-\alpha))+o(1))n}.$$

Dividing the ambient dimension by this subspace dimension produces the two entropy subtractions in (8). Finally, set

$$\kappa_{\text{bin}}(\delta) = \min\{\kappa_{\text{H}}(\delta), \kappa_{\text{CW}}(\delta)\}.$$

Theorem 1.1. *For every fixed $0 < \delta < 1/2$,*

$$R_2(\delta) \leq \kappa_{\text{bin}}(\delta) < M_2(\delta).$$

The classical MRRW constructions are boundary subfamilies of the enlarged variational problems: $b = 0$ recovers M_1 , and $\beta = \gamma = 0$ recovers F_δ . [Theorem 2.3](#) and [Proposition 3.7](#) show that introducing positive harmonic degrees strictly improves the first MRRW bound and the classical constant-weight objective at every interior optimizing layer.

If the endpoint $\tau = 1 - 2\delta$ minimizes (3), then $M_2 = M_1$ and the whole-cube construction gives the strict improvement. Otherwise, the constant-weight construction strictly improves a minimizing layer. [Theorem 3.8](#) combines these cases to prove [Theorem 1.1](#).

1.2. Spherical codes: statement of the result. Fix a maximum inner product $0 < s < 1$. We first recall the classical spherical-code bound and its spherical-cap optimization, then illustrate our improvement using harmonics on the directions orthogonal to a code point. Finally, we state the full hierarchy; its sphere-packing consequence follows in the next subsection.

For $u \geq 0$, define the spherical harmonic dimension exponent by

$$\text{H}_{\text{sph}}(u) = (1+u) \log_2(1+u) - u \log_2 u, \quad \text{H}_{\text{sph}}(0) = 0. \quad (9)$$

Indeed, the space of degree- $un + o(n)$ spherical harmonics has dimension $2^{(\text{H}_{\text{sph}}(u)+o(1))n}$; the exact formula is recorded in (59). Applying the Kabatianskii–Levenshtein bound directly to the whole sphere gives

$$A(n, s) \leq 2^{(\text{H}_{\text{sph}}(a_0(s))+o(1))n}, \quad a_0(s) = \frac{1}{2} \left((1-s^2)^{-1/2} - 1 \right).$$

A stronger classical bound first restricts the code to a sufficiently populated spherical cap and projects that cap onto a lower-dimensional sphere. Sidelnikov’s spherical-slice inequality replaces the original inner-product threshold s by a smaller threshold $0 \leq t \leq s$, at an exponential cost $\frac{1}{2} \log_2((1-t)/(1-s))$ [Sid74]. Combining this spherical-cap reduction with the Kabatianskii–Levenshtein construction gives the optimized classical exponent [KL78, Thm. 4]

$$B_{\text{KL}}(s) = \inf_{0 \leq t \leq s} \left\{ \text{H}_{\text{sph}}(a_0(t)) + \frac{1}{2} \log_2 \frac{1-t}{1-s} \right\}. \quad (10)$$

Thus $R_{\text{sph}}(s) \leq B_{\text{KL}}(s)$.

For $x \in \mathbb{S}^{n-1}$, let $E_{k,x} = \mathcal{H}_k(x^\perp)$ be the space of degree- k spherical harmonics on the unit sphere in $x^\perp$. For $k = 0$, this is the classical space of constants; for $k = 1$, it consists of linear functions on $x^\perp$. Every ambient harmonic space $\mathcal{H}_i(\mathbb{R}^n)$ with $i \geq k$ contains a naturally associated copy of $E_{k,x}$; see §5.2 for the construction.

Retain these copies for $k \leq i \leq L$, and write $a = L/n$, $b = k/n$, with $0 \leq b < a$. Coordinate multiplication connects consecutive ambient harmonic spaces, giving a tridiagonal matrix indexed by $i = k, \dots, L$. Define

$$\Gamma_{\text{row}}(a, b) = \frac{(a-b)(1+a+b)}{(1+2a)\sqrt{a(1+a)}}, \quad \Phi_{\text{row}}(a, b) = \text{H}_{\text{sph}}(a) - \text{H}_{\text{sph}}(b). \quad (11)$$

This construction gives, whenever $2\Gamma_{\text{row}}(a, b) > s$,

$$A(n, s) \leq 2^{(\Phi_{\text{row}}(a,b)+o(1))n}.$$

Since $\dim E_{k,x} = 2^{(\mathbf{H}_{\text{sph}}(b) + o(1))n}$, the projection rank subtracts $\mathbf{H}_{\text{sph}}(b)$ from the ambient exponent $\mathbf{H}_{\text{sph}}(a)$. Optimizing over $b > 0$ strictly improves both the direct Kabatianskii–Levenshtein bound and, after spherical-cap reduction, its optimized version.

For the full hierarchy, fix an integer $r \geq 0$ and choose vectors $\mathbf{a} = (a_1, \dots, a_{r+1})$ and $\mathbf{b} = (b_1, \dots, b_r)$ satisfying

$$a_1 > b_1 > a_2 > \dots > b_r > a_{r+1} \geq 0.$$

For $r = 0$, the chain means simply $a_1 \geq 0$. The relevant orthogonal-group representations are indexed by Young diagrams, and a_ℓ and b_m are their row lengths divided by n . The hierarchy level r counts the rows of the subspace associated with each code point; the ambient representation has one additional row. Their representation-theoretic meaning and the resulting weighted adjacency matrices are developed in §§5.3 and 7. Level $r = 0$ recovers the classical construction; when $r = 1$, setting $a_2 = 0$ gives the harmonics on $x^\perp$ in (11). Introduce the quadratic coordinates and scalar function

$$\mathcal{A}(u) = u(1+u), \quad q(u) = \frac{\sqrt{u(1+u)}}{1+2u}, \quad x_\ell = \mathcal{A}(a_\ell), \quad y_m = \mathcal{A}(b_m).$$

For $1 \leq \ell, j \leq r+1$ and $1 \leq m \leq r$, define R_ℓ , $\Gamma_r = \Gamma_r(\mathbf{a}, \mathbf{b})$, and $\Phi_r = \Phi_r(\mathbf{a}, \mathbf{b})$ by

$$R_\ell = \frac{\prod_{m=1}^r (x_\ell - y_m)}{\prod_{j \neq \ell} (x_\ell - x_j)}, \quad \Gamma_r = \sum_{\ell=1}^{r+1} R_\ell q(a_\ell), \quad \Phi_r = \sum_{\ell=1}^{r+1} \mathbf{H}_{\text{sph}}(a_\ell) - \sum_{m=1}^r \mathbf{H}_{\text{sph}}(b_m).$$

The displayed strict inequalities make every R_ℓ positive, and Lagrange interpolation gives $\sum_\ell R_\ell = 1$; see (81). The spectral calculation and dimension estimate giving $2\Gamma_r$ and Φ_r appear in Lemmas 7.2 and A.1. Empty products and sums take their usual values, so the definitions include $r = 0$.

Theorem 1.2. Fix $r \in \mathbb{Z}_{\geq 0}$, $0 < s < 1$, and nonnegative vectors $\mathbf{a} = (a_1, \dots, a_{r+1})$ and $\mathbf{b} = (b_1, \dots, b_r)$ satisfying $a_\ell > b_\ell > a_{\ell+1}$ for $1 \leq \ell \leq r$. Then

$$2\Gamma_r(\mathbf{a}, \mathbf{b}) > s \implies A(n, s) \leq 2^{(\Phi_r(\mathbf{a}, \mathbf{b}) + o(1))n}.$$

Here $o(1) \rightarrow 0$ as $n \rightarrow \infty$ with $r, s, \mathbf{a}, \mathbf{b}$ fixed. Consequently,

$$R_{\text{sph}}(s) \leq \inf_{r \in \mathbb{Z}_{\geq 0}} \inf_{a_1 > b_1 > \dots > b_r > a_{r+1} \geq 0} \inf_{2\Gamma_r(\mathbf{a}, \mathbf{b}) \geq s} \Phi_r(\mathbf{a}, \mathbf{b}).$$

The non-strict constraint $2\Gamma_r \geq s$ follows by approximating its equality boundary with parameters satisfying $2\Gamma_r > s$. Write $\kappa_r(s)$ for the inner infimum in Theorem 1.2, taken over $\mathbf{a}, \mathbf{b}$ at the fixed hierarchy level r , and set $\kappa_r(0) = 0$. Applying the spherical-cap reduction from (10) gives

$$\bar{\kappa}_r(s) = \inf_{0 \leq t \leq s} \left\{ \kappa_r(t) + \frac{1}{2} \log_2 \frac{1-t}{1-s} \right\}.$$

Consequently,

$$R_{\text{sph}}(s) \leq \inf_{r \geq 0} \bar{\kappa}_r(s).$$

Level 0 recovers the classical Kabatianskii–Levenshtein construction [KL78], so $\bar{\kappa}_0 = B_{\text{KL}}$ by (10). The one-row construction (11) is obtained from level 1 by setting $a_2 = 0$; write $\bar{\kappa}_{\text{row}}$ for its spherical-cap-optimized exponent. The full strict hierarchy, including this intermediate one-row improvement, is proved in Corollary 7.6.

Corollary 1.3. For every $0 < s < 1$ and every $r \geq 0$,

$$\kappa_{r+1}(s) < \kappa_r(s), \quad \bar{\kappa}_{r+1}(s) < \bar{\kappa}_r(s).$$

Moreover,

$$R_{\text{sph}}(s) \leq \inf_{r \geq 0} \bar{\kappa}_r(s) < \bar{\kappa}_1(s) < \bar{\kappa}_{\text{row}}(s) < \bar{\kappa}_0(s) = B_{\text{KL}}(s).$$

1.3. Sphere-packing consequence. Let Δ_n denote the maximal packing density in $\mathbb{R}^n$. For $-1 < s < 1$, projection from the upper hemisphere gives [Sid74, KL78, CZ14]

$$\Delta_n \leq \left(\frac{1-s}{2} \right)^{n/2} A(n+1, s).$$

Consequently, if $A(n, s) \leq 2^{(B(s)+o(1))n}$, then

$$\Delta_n \leq 2^{(B(s)-\frac{1}{2}\log_2 \frac{2}{1-s}+o(1))n}.$$

For spherical codes at a fixed threshold s , passing to a spherical cap can improve the exponent by replacing s with some $t \leq s$. The packing transfer behaves differently: the cost of passing to that cap exactly offsets the change in its geometric packing factor. Indeed, at hierarchy level r ,

$$\frac{1}{2} \log_2 \frac{2}{1-s} - \frac{1}{2} \log_2 \frac{1-t}{1-s} - \kappa_r(t) = \frac{1}{2} \log_2 \frac{2}{1-t} - \kappa_r(t).$$

Thus transferring a cap-improved spherical bound at threshold s gives exactly the packing exponent obtained by transferring the direct spherical bound at threshold t ; see (93). Optimizing over the hierarchy and the threshold therefore requires no additional cap optimization.

Corollary 1.4. As $n \rightarrow \infty$,

$$\Delta_n \leq 2^{-(\lambda_*+o(1))n}, \quad \lambda_* = \frac{1}{2} \log_2 \frac{2\pi}{e}.$$

This improves the classical Kabatianskii–Levenshtein sphere-packing exponent [KL78, Lev79] and attains the Cohn–Elkies rate conjectured in [AJCHLT20, Conjs. 3.1–3.2]. The exact hierarchy limit and the order of the dimension, hierarchy-depth, and small-angle limits are proved in Theorem 8.3.

Remark 1.5. Gorbachev and, independently, Cohn and Elkies developed the Euclidean Fourier-analytic linear program for sphere packing [Gor00, CE03]; optimizing its auxiliary functions is called the Cohn–Elkies linear program. The companion paper in Chapter 1 gives an independent, self-contained determination of its optimal exponent by proving a lower bound for every auxiliary function satisfying the Cohn–Elkies Fourier-positivity and sign conditions and constructing matching Euclidean witnesses. If LP_n denotes the optimal density bound of that program, its result is

$$\lim_{n \rightarrow \infty} \text{LP}_n^{1/n} = \sqrt{\frac{e}{2\pi}}.$$

Cohn and Zhao’s upper-hemisphere construction also converts our spherical certificates into feasible Euclidean auxiliary functions [CZ14, Thm. 3.4 and subsequent discussion]. The subsequent discussion covers $s > 1/2$, as required in the limit $s \uparrow 1$. §B further identifies the main weight in the companion paper’s Euclidean construction with the small-angle limit of our spherical certificates.

Remark 1.6. The spherical and binary constructions differ in the multiplicities that can arise when an ambient representation is restricted to a point stabilizer. This distinction is expressed by Gelfand pairs and strong Gelfand pairs [Kra76, CST08]. For a compact group G and a point stabilizer H , the pair (G, H) is a Gelfand pair when $\dim_{\mathbb{C}} \text{Hom}_H(\mathbf{1}, V) \leq 1$ for every irreducible complex G -representation V . Thus each ambient representation contains at most one stabilizer-fixed line, as in the classical constructions. The pair is a strong Gelfand pair when $\dim_{\mathbb{C}} \text{Hom}_H(E, V) \leq 1$ for every irreducible complex H -representation E . This stronger condition permits nontrivial stabilizer representations while retaining scalar transition coefficients.

The spherical pair $(SO(n), SO(n-1))$ is strong by the multiplicity-free branching rule (72) [GW09, Thms. 8.1.3–8.1.4]. For the Hamming cube and a weight- w layer, the respective pairs are

$$(B_n, S_n), \quad (S_n, S_w \times S_{n-w}), \quad B_n = \{\pm 1\}^n \rtimes S_n.$$

Both are Gelfand pairs but need not be strong. For example, the Littlewood–Richardson coefficient $c_{(2,1),(2,1)}^{(3,2,1)} = 2$ produces multiplicity two in both restrictions. Nevertheless, the Fourier

spaces in §2.1 and the two-row Johnson spaces in Lemma 3.1 contain the selected stabilizer types with multiplicity one. Consequently, both binary arguments remain scalar. More general binary stabilizer types can require matrix-valued transitions, but they are unnecessary for the all-distance improvement proved here.

Remark 1.7. Our bounds concern unrestricted codes with fixed relative distance or fixed maximum inner product in dimensions tending to infinity. Previous spherical constructions in the same high-dimensional, fixed-angle setting improve the direct Kabatianskii–Levenshtein bound for some angles. Relative to the classical bound after spherical-cap optimization, however, those results improve the bound by constant multiplicative factors rather than its exponential rate [SZ24, Z24]. Other results address restricted code families, different asymptotic regimes, or specific dimensions. The Cohn–Elkies bound is sharp in dimensions 8 and 24, where the E_8 and Leech lattice packings are optimal, respectively [Via17, CKMRV17]. The n -dimensional kissing number is $\tau_n = A(n, 1/2)$. Already hierarchy level 2 gives $\tau_n \leq 2^{(0.39661+o(1))n}$, compared with the optimized classical exponent $0.400944\dots$. Its values are known exactly in dimensions 3 [SvdW53], 4 [Mus08], 8 [Lev79], and 24 [OS79]. Semidefinite programs improve further finite-dimensional binary and spherical bounds [Sch05, BV08]. MacWilliams’s weight-transform identities and subsequent spectral and higher-order methods apply to the restricted class of linear codes [Mac63, FT05, CJJ22, LL23, CJJLL26], while another binary-code improvement concerns the regime $d = n/2 - \Theta(\sqrt{n})$, rather than fixed relative distance [PMP23]. Constructions of Delsarte dual certificates, limitations of particular spectral methods, and lower bounds on spherical, Hamming, and constant-weight linear-programming optima provide additional context but do not exclude the certificates constructed here [Sam01, Sam04, NS05, CD25, Sam25].

2. THE FIRST BINARY BOUND

We first improve the first MRRW bound at every relative distance by working directly on the Hamming cube. The construction uses Fourier levels, elementary Boolean addition and deletion maps, and a symmetric tridiagonal matrix indexed by consecutive Fourier degrees. Its new feature is that every degree carries a copy of the same higher-dimensional harmonic space. In this section, the superscript $\square$ identifies whole-cube representation spaces and dimensions. Later sections use J and S for the Johnson-layer and spherical constructions, respectively.

2.1. Fourier levels and Boolean harmonics. The classical first MRRW bound comes from a symmetric tridiagonal matrix whose indices are Fourier degrees and whose nonzero off-diagonal entries connect consecutive degrees. To improve that bound, we attach a larger Boolean harmonic space to every selected degree. Identify the binary alphabet with $\{\pm 1\}$, and write $X_n = \{\pm 1\}^n$. The normalized inner product records Hamming distance:

$$\ell_x = \frac{x}{\sqrt{n}}, \quad t(x, y) = \langle \ell_x, \ell_y \rangle = 1 - \frac{2d_H(x, y)}{n}.$$

Give real-valued functions on X_n the uniform-probability inner product. The Fourier characters $\chi_S(x) = \prod_{r \in S} x_r$, indexed by $S \subseteq [n]$, form an orthonormal basis; write $e_S = \chi_S$ for its basis vectors. Their degree- i subspace is

$$V_i^\square = \text{span}\{e_S : |S| = i\}, \quad D_i^\square = \dim V_i^\square = \binom{n}{i}.$$

At a word x , the classical spectral construction [BN06, §3.1] uses the unit vector

$$v_{i,x} = \binom{n}{i}^{-1/2} \sum_{|S|=i} \chi_S(x) e_S.$$

It is unchanged by every symmetry of the cube fixing x . The inner products $\langle v_{i,x}, v_{i,y} \rangle$ are normalized Krawtchouk polynomials [Del73, MRRW77], but their explicit formulas will not be needed: we use the Fourier spaces and their coordinate transitions (17)–(18) instead. Multiplication by $t(x, \cdot)$ connects only consecutive Fourier levels. On the lines spanned by $v_{i,x}$, its

off-diagonal coefficient between degrees i and $i + 1$ is $\sqrt{(i+1)(n-i)}/n$. Thus the classical construction is governed by the symmetric tridiagonal matrix with these neighboring-degree entries.

To replace each fixed line by a larger subspace, introduce the Boolean addition and deletion operators [Sri11, Fei12]:

$$Ue_S = \sum_{r \notin S} e_{S \cup \{r\}}, \quad De_S = \sum_{r \in S} e_{S \setminus \{r\}}.$$

These formulas define U and D on the Fourier basis and extend linearly to $\bigoplus_{i=0}^n V_i^\square$. Since adding an element to S is equivalent to deleting the same element from the resulting set, their matrix coefficients satisfy

$$\langle Ue_S, e_T \rangle = \langle e_S, De_T \rangle.$$

Consequently, $U^* = D$. Counting additions and deletions in the two possible orders also gives

$$(DU - UD)|_{V_i^\square} = (n - 2i)\text{id}_{V_i^\square}. \quad (12)$$

For $0 \leq k \leq n/2$, define the degree- k Boolean harmonic space, with $V_{-1}^\square = 0$, by

$$E_k = \ker(D : V_k^\square \rightarrow V_{k-1}^\square).$$

Thus a coefficient vector on the k -element subsets belongs to E_k when its coefficients sum to zero over all subsets containing any specified $(k-1)$ -element subset. In particular, E_0 is the constant line, and

$$E_1 = \left\{ \sum_{r=1}^n a_r e_{\{r\}} : \sum_{r=1}^n a_r = 0 \right\}$$

has dimension $n - 1$. More generally, for $1 \leq k \leq n/2$, (12) implies

$$\|Uf\|^2 = \|Df\|^2 + (n - 2k + 2)\|f\|^2 \quad (f \in V_{k-1}^\square).$$

Thus $U : V_{k-1}^\square \rightarrow V_k^\square$ is injective, and its adjoint $D : V_k^\square \rightarrow V_{k-1}^\square$ is surjective. Rank-nullity therefore gives

$$d_k^\square = \dim E_k = \binom{n}{k} - \binom{n}{k-1},$$

where $\binom{n}{-1} = 0$. For $h \in E_k$, induction using (12) yields the following identities, valid for $1 \leq r \leq n - 2k$ and $0 \leq r \leq n - 2k$, respectively:

$$DU^r h = r(n - 2k - r + 1)U^{r-1}h, \quad \|U^r h\|^2 = r! \frac{(n - 2k)!}{(n - 2k - r)!} \|h\|^2. \quad (13)$$

The norm formula in (13) follows by repeatedly applying $U^* = D$. Consequently, for every $k \leq i \leq n - k$, the normalized addition map

$$\varphi_i h = \frac{U^{i-k} h}{\sqrt{(i-k)!(n-2k)!/(n-i-k)!}} \quad (h \in E_k) \quad (14)$$

has $\|\varphi_i h\| = \|h\|$ and therefore embeds E_k isometrically into $V_i^\square$. If $R_x e_S = \chi_S(x) e_S$, then $\varphi_{i,x} = R_x \varphi_i$ transports this copy from the all-ones word to x . For $k = 0$, its image is precisely the classical fixed line $\mathbb{R}v_{i,x}$. For $k > 0$, every retained Fourier level instead carries a copy of the same $d_k^\square$ -dimensional space; when k is proportional to n , that dimension grows exponentially.

The coordinate transitions have an elementary explicit form. For neighboring levels, define isometric addition and deletion maps by

$$C_{i,i+1} e_T = \frac{1}{\sqrt{i+1}} \sum_{r \in T} e_r \otimes e_{T \setminus \{r\}}, \quad |T| = i + 1, \quad (15)$$

$$C_{i+1,i} e_T = \frac{1}{\sqrt{n-i}} \sum_{r \notin T} e_r \otimes e_{T \cup \{r\}}, \quad |T| = i. \quad (16)$$

Here e_r is the r th standard basis vector of $\mathbb{R}^n$, $C_{i,i+1} : V_{i+1}^\square \rightarrow \mathbb{R}^n \otimes V_i^\square$, and $C_{i+1,i} : V_i^\square \rightarrow \mathbb{R}^n \otimes V_{i+1}^\square$. Each map is an isometry: the displayed summands are distinct orthonormal basis

tensors, and inputs to the two maps have $i+1$ and $n-i$ summands, respectively. At a common output level $V_i^\square$, the deletion map $C_{i,i+1}$ is supported on tensors $e_r \otimes e_S$ with $r \notin S$, whereas the addition map $C_{i,i-1}$ is supported on tensors with $r \in S$. Their supports are disjoint, so their ranges are orthogonal. Set

$$Q_i = (i - k + 1)(n - i - k).$$

For $f \in V_i^\square$ and $g \in V_{i+1}^\square$, the defining basis formulas give

$$C_{i,i+1}^*(\ell_x \otimes R_x f) = \frac{R_x U f}{\sqrt{n(i+1)}}, \quad C_{i+1,i}^*(\ell_x \otimes R_x g) = \frac{R_x D g}{\sqrt{n(n-i)}}.$$

By (13) and (14), the normalized Boolean harmonic embeddings satisfy

$$U\varphi_i h = \sqrt{Q_i} \varphi_{i+1} h, \quad D\varphi_{i+1} h = \sqrt{Q_i} \varphi_i h.$$

Combining these identities gives

$$C_{i,i+1}^*(\ell_x \otimes \varphi_{i,x} h) = \sqrt{\frac{Q_i}{n(i+1)}} \varphi_{i+1,x} h, \quad (17)$$

$$C_{i+1,i}^*(\ell_x \otimes \varphi_{i+1,x} h) = \sqrt{\frac{Q_i}{n(n-i)}} \varphi_{i,x} h. \quad (18)$$

Thus the squared forward and reverse transition coefficients are

$$p_{i,i+1} = \frac{Q_i}{n(i+1)}, \quad p_{i+1,i} = \frac{Q_i}{n(n-i)}.$$

For $k = 0$, all outgoing coefficients at each Fourier degree sum to one on the full path. For $k > 0$, they generally do not, because coordinate multiplication also reaches representations outside the selected path. The coefficients of the two orientations of a given edge are likewise generally unequal, but they satisfy the dimension-weighted balance identity

$$\binom{n}{i} p_{i,i+1} = \binom{n}{i+1} p_{i+1,i}. \quad (19)$$

This identity will convert the directed transitions into a real symmetric matrix, whose largest eigenvalue controls the code bound.

2.2. The finite whole-cube bound. We now symmetrize the directed Fourier-degree transitions and combine their largest eigenvector with the Boolean coordinate maps to construct a scalar positive-definite kernel. This gives a bound for codes of fixed finite length. Fix $0 \leq k < L \leq n - k$, and select the Fourier levels $V_k^\square, \dots, V_L^\square$. Their symmetric transition matrix $J_H = J_H(n, k, L)$ has zero diagonal and off-diagonal entries

$$(J_H)_{i,i+1} = (J_H)_{i+1,i} = c_{i,H}^{(k)} := \frac{(i - k + 1)(n - i - k)}{n\sqrt{(i+1)(n-i)}}. \quad (20)$$

Indeed, (19) says that the directed transition matrix is diagonally similar to a symmetric matrix. The resulting symmetric edge weight is the geometric mean of the coefficients in its two directions:

$$\sqrt{p_{i,i+1} p_{i+1,i}} = \frac{Q_i}{n\sqrt{(i+1)(n-i)}}.$$

For $k = 0$, J_H is the classical Krawtchouk matrix underlying the first MRRW bound [MRRW77]; its tridiagonal spectral formulation appears in [BN06, §3.1]. For positive k , J_H is still indexed by consecutive Fourier degrees and connects only neighboring degrees, but every degree carries a copy of the $d_k^\square$ -dimensional harmonic space. The resulting weighted path is shown in Figure 1.

Theorem 2.1. *Let n, d, k, L be integers with $1 \leq d \leq n$ and $0 \leq k < L \leq n - k$. Set $s = 1 - 2d/n$ and $\lambda = \lambda_{\max}(J_H(n, k, L))$. If $\lambda > s$, then*

$$A_2(n, d) \leq \left(\frac{1-s}{d_k^\square(\lambda-s)} \right) \sum_{i=k}^L \binom{n}{i}. \quad (21)$$




FIGURE 1. The whole-cube representation path at fixed Boolean harmonic degree k . Each Fourier level contains the isometric copy $\varphi_{i,x}E_k$, and neighboring levels have symmetric edge weight $c_{i,H}^{(k)}$ from (20). The classical path corresponds to the one-dimensional fiber E_0 .

Proof. We first construct a positive-definite projection kernel from the Boolean coordinate maps and the Perron eigenvector of J_H . Let $I = \{k, \dots, L\}$ index the vertices of $J_H(n, k, L)$, and put

$$\mathcal{V} = \bigoplus_{i \in I} V_i^\square, \quad D_{\text{amb}} = \dim \mathcal{V} = \sum_{i \in I} D_i^\square.$$

Since the off-diagonal entries in (20) are strictly positive, its largest eigenvalue $\lambda = \lambda_{\max}(J_H) > 0$ has a strictly positive unit eigenvector $v = (v_i)_{i \in I}$. Write

$$w_i = \sqrt{D_i^\square} v_i, \quad Z = \sum_{i \in I} w_i.$$

The dimension balance (19) and $J_H v = \lambda v$ imply

$$\sum_{\substack{i \in I \\ |i-j|=1}} p_{i,j} w_i = \lambda w_j \quad (j \in I). \quad (22)$$

For each j , dividing by λw_j therefore turns the incoming weights into a probability distribution on its retained neighbors. The square roots of these normalized weights are the coefficients in the coordinate isometry below.

For each word $x \in X_n$, combine its copies of the Boolean harmonic space E_k into the isometric embedding

$$\Psi_x : E_k \longrightarrow \mathcal{V}, \quad \Psi_x h = \bigoplus_{i \in I} \sqrt{\frac{w_i}{Z}} \varphi_{i,x} h.$$

Let $P_x = \Psi_x \Psi_x^*$ be the orthogonal projection onto its image. Thus every P_x has rank $d_k^\square$, and its scalar overlap with another projection is

$$K(x, y) = \text{tr}(P_x P_y) = \|P_x P_y\|_{\text{HS}}^2 \geq 0, \quad \|A\|_{\text{HS}}^2 = \text{tr}(A^* A).$$

As a Hilbert–Schmidt Gram kernel, K is positive definite. Coordinate permutations and sign changes transport P_x to the projection associated with the transformed word, so $K(x, y)$ depends only on $d_H(x, y)$.

For $f = \bigoplus_{j \in I} f_j \in \mathcal{V}$, the concrete coordinate maps (15) and (16) define an operator between the corresponding direct sums:

$$\begin{aligned} \mathcal{B} : \mathcal{V} &\longrightarrow \mathbb{R}^n \otimes \mathcal{V} = \bigoplus_{i \in I} (\mathbb{R}^n \otimes V_i^\square), \\ \mathcal{B} f &= \bigoplus_{i \in I} \sum_{\substack{j \in I \\ |i-j|=1}} \sqrt{\frac{p_{i,j} w_i}{\lambda w_j}} C_{i,j} f_j. \end{aligned}$$

For each i , the images of $C_{i,i-1}$ and $C_{i,i+1}$ in $\mathbb{R}^n \otimes V_i^\square$ are orthogonal: their respective basis tensors $e_r \otimes e_S$ satisfy $r \in S$ and $r \notin S$. Therefore (22) gives

$$\|\mathcal{B} f\|^2 = \sum_{j \in I} \frac{\|f_j\|^2}{\lambda w_j} \sum_{\substack{i \in I \\ |i-j|=1}} p_{i,j} w_i = \|f\|^2.$$

Thus $\mathcal{B}$ is an isometry and $\mathcal{B}^*\mathcal{B} = \text{id}_{\mathcal{V}}$. For a harmonic vector $h \in E_k$, the explicit contractions (17) and (18) likewise give

$$\mathcal{B}^*(\ell_x \otimes \Psi_x h) = \bigoplus_{j \in I} \frac{\varphi_{j,x} h}{\sqrt{\lambda w_j Z}} \sum_{\substack{i \in I \\ |i-j|=1}} p_{i,j} w_i = \bigoplus_{j \in I} \sqrt{\frac{\lambda w_j}{Z}} \varphi_{j,x} h.$$

Consequently,

$$\mathcal{B}^*(\ell_x \otimes \Psi_x h) = \sqrt{\lambda} \Psi_x h. \quad (23)$$

To expose the positivity needed for the code bound, set

$$L_x = (\ell_x \otimes \text{id})P_x, \quad G_x = \mathcal{B}P_x, \quad \Theta_x = L_x - \sqrt{\lambda} G_x.$$

These are maps from $\mathcal{V}$ to $\mathbb{R}^n \otimes \mathcal{V}$. Their Hilbert–Schmidt pairings are

$$\begin{aligned} \langle L_x, L_y \rangle_{\text{HS}} &= \text{tr} [P_x (\ell_x \otimes \text{id})^* (\ell_y \otimes \text{id}) P_y] = t(x, y) K(x, y), \\ \langle G_x, G_y \rangle_{\text{HS}} &= \text{tr} (P_x \mathcal{B}^* \mathcal{B} P_y) = \text{tr} (P_x P_y) = K(x, y). \end{aligned}$$

Since $\text{im } P_x = \text{im } \Psi_x$, (23) implies

$$\mathcal{B}^*(\ell_x \otimes \text{id})P_x = \sqrt{\lambda} P_x. \quad (24)$$

Taking the adjoint of (24), and then applying the same identity at y , yields both mixed pairings:

$$\begin{aligned} \langle L_x, G_y \rangle_{\text{HS}} &= \text{tr} [P_x (\ell_x \otimes \text{id})^* \mathcal{B} P_y] = \sqrt{\lambda} K(x, y), \\ \langle G_x, L_y \rangle_{\text{HS}} &= \text{tr} [P_x \mathcal{B}^* (\ell_y \otimes \text{id}) P_y] = \sqrt{\lambda} K(x, y). \end{aligned}$$

Expanding $\langle L_x - \sqrt{\lambda} G_x, L_y - \sqrt{\lambda} G_y \rangle_{\text{HS}}$ now gives

$$\langle \Theta_x, \Theta_y \rangle_{\text{HS}} = (t(x, y) - \lambda) K(x, y). \quad (25)$$

In particular, for every $s < \lambda$,

$$(t(x, y) - s) K(x, y) = (\lambda - s) K(x, y) + \langle \Theta_x, \Theta_y \rangle_{\text{HS}}.$$

Both summands are positive-definite Gram kernels, so $(t - s)K$ is a scalar two-point Delsarte certificate.

Let $C \subseteq X_n$ have minimum Hamming distance at least d , and put $N_C = |C|$. The projection kernel constructed above satisfies $K(x, y) \geq 0$. Hence $(t(x, y) - s)K(x, y) \leq 0$ for distinct $x, y \in C$, whereas its diagonal value is $(1 - s)d_k^{\square}$. Summing (25) over C^2 gives

$$(\lambda - s) \left\| \sum_{x \in C} P_x \right\|_{\text{HS}}^2 \leq N_C (1 - s) d_k^{\square}.$$

Since every P_x has rank $d_k^{\square}$ and $D_{\text{amb}} = \sum_{i=k}^L \binom{n}{i}$, trace Cauchy–Schwarz gives

$$\left\| \sum_{x \in C} P_x \right\|_{\text{HS}}^2 \geq \frac{N_C^2 (d_k^{\square})^2}{\sum_{i=k}^L \binom{n}{i}}.$$

Combining these estimates proves (21). $\square$

2.3. Asymptotic improvement over the first MRRW bound. To compare the finite bound with M_1 , take both the harmonic degree and the largest retained Fourier degree proportional to n . The largest eigenvalue of the explicitly defined tridiagonal matrix $J_H(n, k, L)$ determines whether the spectral condition $\lambda_{\max}(J_H(n, k, L)) > 1 - 2d/n$ in Theorem 2.1 holds, while the dimension of E_k gives the exponential saving. Recall the binary entropy H_2 , the asymptotic rate R_2 , the first MRRW exponent M_1 , and the whole-cube exponent κ_H from the introduction. The spectral calculations use the Rayleigh–Ritz principle: for every real symmetric matrix J ,

$$\lambda_{\max}(J) = \max_{a \neq 0} \frac{\langle a, Ja \rangle}{\langle a, a \rangle}. \quad (26)$$

Lemma 2.2. Suppose $0 \leq b < a \leq 1/2$, and let $k_n < L_n \leq \lfloor n/2 \rfloor$ satisfy $k_n/n \rightarrow b$ and $L_n/n \rightarrow a$. Then

$$\lim_{n \rightarrow \infty} \lambda_{\max}(J_H(n, k_n, L_n)) = \Gamma_H(a, b).$$

Proof. For a degree i with $i/n \rightarrow x > 0$, the neighboring-degree entry of $J_H(n, k_n, L_n)$ satisfies

$$c_{i,H}^{(k_n)} \rightarrow g_b(x) := \frac{(x-b)(1-x-b)}{\sqrt{x(1-x)}} = \sqrt{x(1-x)} - \frac{b(1-b)}{\sqrt{x(1-x)}}.$$

The finite entries are increasing through the selected degrees. Indeed, put $Y_i = (i+1)(n-i)$ and $T_k = k_n(n+1-k_n)$. Then

$$nc_{i,H}^{(k_n)} = \frac{Y_i - T_k}{\sqrt{Y_i}}.$$

Both Y_i and the function $Y \mapsto (Y - T_k)/\sqrt{Y}$ increase for $k_n \leq i < L_n \leq n/2$. Consequently, the largest row sum gives

$$\lambda_{\max}(J_H(n, k_n, L_n)) \leq 2c_{L_n-1,H}^{(k_n)} \rightarrow 2g_b(a).$$

Choose integers $m_n \rightarrow \infty$ with $m_n = o(n)$, and define a vector supported on the last m_n Fourier degrees by

$$f_{L_n-m_n+r} = \sin \frac{\pi r}{m_n+1} \quad (1 \leq r \leq m_n).$$

Although the entries vary across the full Fourier-degree path, they converge uniformly to $g_b(a)$ on these last $m_n = o(n)$ degrees. The restricted matrix therefore becomes a constant-weight nearest-neighbor adjacency operator. Its first discrete sine vector and (26) give

$$\lambda_{\max}(J_H(n, k_n, L_n)) \geq 2g_b(a) \cos \frac{\pi}{m_n+1} + o(1).$$

Taking $n \rightarrow \infty$ and recalling $\Gamma_H(a, b) = 2g_b(a)$ proves the claim. $\square$

Theorem 2.3. For every fixed $0 < \delta < 1/2$,

$$R_2(\delta) \leq \kappa_H(\delta) < M_1(\delta).$$

Proof. Fix $0 \leq b < a \leq 1/2$ with $\Gamma_H(a, b) > 1 - 2\delta$, as in (2), and take $k = \lfloor bn \rfloor$, $L = \lfloor an \rfloor$, and $d_n = \lceil \delta n \rceil$. The finite distance threshold is $s_n = 1 - 2d_n/n \rightarrow 1 - 2\delta$. By Lemma 2.2, $\lambda_{\max}(J_H(n, k, L)) \rightarrow \Gamma_H(a, b)$. Stirling's formula gives

$$\frac{1}{n} \log_2 \sum_{i=k}^L \binom{n}{i} = H_2(a) + o(1), \quad \frac{1}{n} \log_2 d_k^\square = H_2(b) + o(1).$$

Strict feasibility in (2) gives $\lambda_{\max}(J_H(n, k, L)) - s_n > 0$, bounded away from zero for large n , so the prefactor in (21) stays bounded. Taking logarithmic rates and then the infimum over (a, b) gives $R_2(\delta) \leq \kappa_H(\delta)$.

At the classical fixed-line spectral boundary, put $a_0 = \frac{1}{2} - \sqrt{\delta(1-\delta)}$ and $b = 0$. Then

$$\Gamma_H(a_0, 0) = 2\sqrt{a_0(1-a_0)} = 1 - 2\delta, \quad H_2(a_0) - H_2(0) = M_1(\delta).$$

At this boundary point,

$$\partial_a \Gamma_H > 0, \quad -\frac{\partial_b \Gamma_H}{\partial_a \Gamma_H} = \frac{2}{1-2a_0}.$$

Thus every $c > 2/(1-2a_0)$ gives strictly feasible parameter pairs $a = a_0 + cb$ for sufficiently small $b > 0$. The ambient entropy $H_2(a)$ increases by $O_\delta(b)$, whereas $H_2(b) = b \log_2(1/b) + O(b)$. The harmonic-space entropy therefore dominates the spectral cost as $b \downarrow 0$, proving $\kappa_H(\delta) < M_1(\delta)$. $\square$

3. THE OPTIMIZED BINARY BOUND

The first MRRW exponent is not always the smallest classical bound. To improve the optimized second MRRW exponent M_2 , we adapt the whole-cube projection argument of [Theorem 2.1](#) to a constant-weight layer and then transfer the resulting estimate back to unrestricted binary codes. Johnson degrees replace Fourier degrees, and two Boolean harmonic spaces, on the support and its complement, replace the single whole-cube harmonic space.

3.1. Constant-weight layers and their harmonic spaces. To improve the constant-weight part of the second MRRW bound, we first reduce an unrestricted code to words with a common weight. Identify such words with their supports, and write $X_w = \binom{[n]}{w}$ for the weight- w layer. Let $A_J(n, w, d)$ be the largest size of a code in X_w with minimum Hamming distance at least d . Averaging the intersection of an unrestricted binary code with all translates of the weight- w layer gives the Bassalygo–Elias inequality [[Bas65](#), [MRRW77](#)]:

$$A_2(n, d) \leq \frac{2^n}{\binom{n}{w}} A_J(n, w, d). \quad (27)$$

When $w/n \rightarrow \alpha$, its prefactor has exponential rate $1 - H_2(\alpha)$. It therefore suffices to construct a constant-weight bound that improves the classical Hahn-polynomial construction after accounting for this reduction.

A base word $x \in X_w$ partitions the coordinates into x and x^c , of sizes w and $N = n - w$, respectively. Complementation preserves Hamming distance, so we assume $1 \leq w < n/2$; the middle layer follows by taking $w/n \rightarrow 1/2$.

Coordinate permutations preserve the layer, and those fixing x separately permute its support and complement. The distance between two words in the layer is even: half their Hamming distance is the number of coordinates removed from one support to obtain the other. For $x, y \in X_w$, put

$$r(x, y) = w - |x \cap y| = \frac{1}{2} d_H(x, y), \quad t(x, y) = 1 - \frac{nr(x, y)}{wN}. \quad (28)$$

The unit vectors $\ell_x = \sqrt{n/(wN)}(\mathbf{1}_x - (w/n)\mathbf{1})$ satisfy $\langle \ell_x, \ell_y \rangle = t(x, y)$.

Give functions on X_w the uniform-probability inner product. The standard Johnson harmonic decomposition, indexed by degrees $0 \leq j \leq w$, is the constant-weight analogue of the Fourier decomposition [[Del73](#)]:

$$\mathbb{R}^{X_w} = \bigoplus_{j=0}^w V_j^J, \quad D_j^J = \dim V_j^J = \binom{n}{j} - \binom{n}{j-1}. \quad (29)$$

As a representation of the coordinate permutation group S_n , the space V_j^J is the Specht module indexed by the two-row partition $(n-j, j)$ [[Del73](#), [Sri11](#)]. In particular, these spaces are pairwise inequivalent and irreducible. In the classical construction, the permutations fixing x fix a single line in V_j^J , whose overlap kernel is a Hahn polynomial in $r(x, y)$. Neither its explicit polynomial formula nor the adjacency eigenvalues will be needed: the argument uses the Johnson harmonic spaces and the associated Hahn recurrence coefficients (35). We replace the fixed line by a product of Boolean harmonic spaces, one on the support and one on its complement, extending the degree- k Boolean harmonics of §2.1. For a finite set A , identify $\mathbb{R}^{\binom{A}{p}}$ with real functions on its p -element subsets. For $0 \leq p \leq |A|/2$, let $\mathcal{E}_p(A) = \ker \partial$, where $\partial : \mathbb{R}^{\binom{A}{p}} \rightarrow \mathbb{R}^{\binom{A}{p-1}}$ is the Boolean lowering map $(\partial f)(T) = \sum_{a \in A \setminus T} f(T \cup \{a\})$, with $\mathbb{R}^{\binom{A}{-1}} = 0$. Thus f is harmonic precisely when its values sum to zero over the p -subsets extending each $(p-1)$ -subset. Its harmonic dimension is

$$d_p(|A|) := \dim \mathcal{E}_p(A) = \binom{|A|}{p} - \binom{|A|}{p-1}. \quad (30)$$

For $0 \leq p \leq w/2$ and $0 \leq q \leq N/2$, define the attached harmonic space and its dimension by

$$E_x^{p,q} = \mathcal{E}_p(x) \otimes \mathcal{E}_q(x^c), \quad d_{p,q}^J = d_p(w)d_q(N).$$

Thus p and q are harmonic degrees on the support and its complement, not dimensions. These Boolean harmonic spaces are the Specht modules indexed by $(w - p, p)$ and $(N - q, q)$, respectively [Sri11]. The two-row Littlewood–Richardson restriction rule for $S_n \downarrow S_w \times S_N$ determines exactly which Johnson degrees contain their tensor product $E_x^{p,q}$, and shows that each such degree contains it with multiplicity one [GW09, §9.3.5].

Lemma 3.1. *Let n, w, p, q be integers with $1 \leq w < n/2$, $0 \leq p \leq w/2$, and $0 \leq q \leq (n - w)/2$. Put $N = n - w$, and fix $x \in X_w$. The Johnson space V_j^J contains a copy of $E_x^{p,q}$ preserved by the permutations fixing x precisely when*

$$j_- = p + q \leq j \leq j_+ := \min\{w, w - p + q, N + p - q\}. \quad (31)$$

This copy occurs once, and its isometric embedding $\varphi_{j,x} : E_x^{p,q} \rightarrow V_j^J$, commuting with those permutations, is unique up to sign.

3.2. The associated Hahn tridiagonal matrix. To adapt the projection argument of Theorem 2.1 to the layer X_w , we need the coordinate transitions between the copies of $E_x^{p,q}$ in successive Johnson degrees. These transitions are governed by an associated Hahn recurrence. For fixed x , let M_x denote multiplication by $y \mapsto t(x, y)$. Its action on the Johnson harmonic spaces satisfies the degree-one product rule [BCV23, (90)]:

$$M_x V_j^J \subseteq V_{j-1}^J \oplus V_j^J \oplus V_{j+1}^J. \quad (32)$$

Because M_x commutes with all permutations fixing x , Lemma 3.1 reduces its action on the copies of $E_x^{p,q}$ to a scalar three-term recurrence. Its coefficients are conveniently expressed using six standard Hahn-recurrence parameters:

$$j_1 = \frac{w}{2} - p, \quad j_2 = \frac{N}{2} - q, \quad j = \frac{n}{2} - j, \quad m_0 = \frac{n}{2} - w, \quad \Sigma = j_1 + j_2, \quad \Delta_j = j_2 - j_1. \quad (33)$$

Here j_1, j_2 depend on the two Boolean harmonic degrees, j depends on the Johnson degree, and m_0 depends only on the layer weight. Write

$$\begin{aligned} \mu_j^{p,q} &= \frac{m_0}{2} \frac{j_2(j_2 + 1) - j_1(j_1 + 1)}{j(j + 1)}, \\ \nu_j^{p,q} &= \frac{\sqrt{(j^2 - m_0^2)(j^2 - \Delta_j^2)((\Sigma + 1)^2 - j^2)}}{2j\sqrt{(2j - 1)(2j + 1)}}. \end{aligned} \quad (34)$$

Only $j_- \leq j < j_+$ require $\nu_j^{p,q}$. Define the diagonal and off-diagonal recurrence coefficients by

$$b_j^{p,q} = \frac{n\mu_j^{p,q} - m_0^2}{wN}, \quad c_j^{p,q} = \frac{n\nu_j^{p,q}}{wN}. \quad (35)$$

The numbers $b_j^{p,q}$ and $c_j^{p,q}$ are the associated Hahn recurrence coefficients [BCV23, (91)–(92)], with $(j_1, j_2, j, m) = (j_2, j_1, j, m_0)$ and with the operator normalized to act by multiplication by $t(x, \cdot)$. Their identification with Hahn polynomials is classical [Koo81, §4]. Setting $p = q = 0$ recovers the ordinary Hahn recurrence on the stabilizer-fixed lines, while positive p or q introduces a higher-dimensional harmonic space.

We now fix the embeddings in Lemma 3.1 unambiguously. Suppose $j_- \leq j_+$; otherwise no Johnson harmonic space contains $E_x^{p,q}$. Choose a reference support x_0 and an isometric stabilizer-equivariant embedding $\varphi_{j_-,x_0} : E_{x_0}^{p,q} \rightarrow V_{j_-}^J$. Write Π_j^J for orthogonal projection onto V_j^J . For each successive degree, choose the sign of its isometric embedding so that, for every $Y \in E_{x_0}^{p,q}$,

$$\Pi_{j+1}^J(t(x_0, \cdot)\varphi_{j,x_0}Y) = c_j^{p,q}\varphi_{j+1,x_0}Y, \quad j_- \leq j < j_+. \quad (36)$$

The associated Hahn coefficients $c_j^{p,q}$ are positive throughout this range, so this condition determines each successive sign. For every support $x = gx_0$, transport all the resulting embeddings by $\varphi_{j,x} = g\varphi_{j,x_0}g^{-1}$. This definition is independent of the permutation g : two choices differ by a permutation fixing x_0 , and every φ_{j,x_0} intertwines their actions.

Proposition 3.2. *Let n, w, p, q be integers with $1 \leq w < n/2$, $0 \leq p \leq w/2$, and $0 \leq q \leq (n-w)/2$. Put $N = n - w$, and set $j_- = p + q$ and $j_+ = \min\{w, w - p + q, N + p - q\}$. Suppose $j_- \leq j_+$. For every $x \in X_w$, $Y \in E_x^{p,q}$, and $j_- \leq j \leq j_+$, the embeddings fixed in (36) satisfy*

$$t(x, \cdot) \varphi_{j,x} Y = c_{j-1}^{p,q} \varphi_{j-1,x} Y + b_j^{p,q} \varphi_{j,x} Y + c_j^{p,q} \varphi_{j+1,x} Y, \quad (37)$$

where the degree- $(j-1)$ term is omitted when $j = j_-$, and the degree- $(j+1)$ term is omitted when $j = j_+$. Every retained off-diagonal coefficient satisfies $c_j^{p,q} > 0$ for $j_- \leq j < j_+$.

Proof. Let $H_x = S_x \times S_{x^c} \leq S_n$ be the stabilizer of the support x ; its two factors permute x and x^c , respectively. Multiplication by $t(x, \cdot)$ and each projection Π_i^J commute with H_x . Consequently,

$$\Pi_i^J M_x \varphi_{j,x} : E_x^{p,q} \longrightarrow V_i^J$$

is H_x -equivariant. By Lemma 3.1, its image is zero unless $j_- \leq i \leq j_+$, and otherwise it lies in the unique copy $\varphi_{i,x} E_x^{p,q}$. The Johnson product rule (32) further restricts i to $j-1, j, j+1$. Since $E_x^{p,q}$ is the outer tensor product of two absolutely irreducible real Specht modules, every H_x -equivariant endomorphism of $E_x^{p,q}$ is a real scalar. Therefore each surviving component is a scalar multiple of $\varphi_{i,x}$.

The associated Hahn recurrence and coefficient formulas [BCV23, (90)–(92)], with the parameters and coordinate normalization in (33)–(35), give the diagonal scalar $b_j^{p,q}$ and the absolute value $c_j^{p,q}$ of the degree- $(j+1)$ scalar. For $j_- \leq j < j_+$,

$$0 < \max\{m_0, |\Delta_J|\} < j \leq \Sigma.$$

Indeed, $j \geq j_- = p + q$ gives $j \leq \Sigma$, while

$$\frac{n}{2} - j_+ = \max\{m_0, |\Delta_J|\}$$

and $j < j_+$ give the strict lower inequality. Therefore all three radicand factors in (34) are positive, as is its denominator, proving $c_j^{p,q} > 0$. The orientation (36) fixes the degree- $(j+1)$ component at x_0 . Since $t(gx_0, gy) = t(x_0, y)$, transporting the embeddings gives $c_j^{p,q} \varphi_{j+1,x}$ for every x .

If $j > j_-$, self-adjointness of M_x and the isometry of the embeddings give, for $Y, Z \in E_x^{p,q}$,

$$\langle \varphi_{j-1,x} Z, M_x \varphi_{j,x} Y \rangle = \langle M_x \varphi_{j-1,x} Z, \varphi_{j,x} Y \rangle = c_{j-1}^{p,q} \langle Z, Y \rangle.$$

Thus the degree- $(j-1)$ component is $c_{j-1}^{p,q} \varphi_{j-1,x} Y$, proving (37). At the lower endpoint, the formal preceding index $j_- - 1$ gives $j = \Sigma + 1$, making $(\Sigma + 1)^2 - j^2 = 0$. At the upper endpoint, $j = \max\{m_0, |\Delta_J|\}$; hence $j^2 - m_0^2 = 0$ if $j_+ = w$, and $j^2 - \Delta_J^2 = 0$ if $j_+ = w - p + q$ or $j_+ = N + p - q$. These are precisely the missing transitions beyond the branching range. When $j_- = j_+$, both off-diagonal terms are absent. $\square$

Write $b_j^0 = b_j^{0,0}$ and $c_j^0 = c_j^{0,0}$ for the classical fixed-line coefficients. The fixed-line diagonal satisfies $b_0^0 = 0$ and, for $1 \leq j \leq w$,

$$b_j^0 = \frac{m_0^2 j (n - j + 1)}{w N j (j + 1)} > 0,$$

and $c_j^0 > 0$ for $0 \leq j < w$.

The associated Hahn coefficients describe multiplication by the scalar distance coordinate $t(x, \cdot)$. The projection construction instead requires isometric maps between Johnson spaces and their tensor products with the coordinate space $W = \mathbf{1}^\perp \subseteq \mathbb{R}^n$. Their normalization is determined by the fixed-line case of (37).

For Johnson degrees i, j in (31) with $|i - j| \leq 1$, set

$$r_{i,j}^{p,q} = \begin{cases} b_i^{p,q}, & i = j, \\ c_{\min\{i,j\}}^{p,q}, & |i - j| = 1, \end{cases} \quad r_{i,j}^0 = \begin{cases} b_i^0, & i = j, \\ c_{\min\{i,j\}}^0, & |i - j| = 1. \end{cases}$$

Thus $r_{i,j}^{p,q}$ is the scalar coefficient connecting the copies of $E_x^{p,q}$ in V_i^J and V_j^J , whereas $r_{i,j}^0$ is the corresponding coefficient for the classical fixed lines. Except when $i = j = 0$, the fixed-line coefficient is strictly positive.

Lemma 3.3. *Let i, j be Johnson degrees in (31) with $|i - j| \leq 1$ and $(i, j) \neq (0, 0)$. There is an S_n -equivariant isometry*

$$C_{i,j} : V_j^J \longrightarrow W \otimes V_i^J$$

whose contraction onto the selected harmonic spaces is

$$C_{i,j}^*(\ell_x \otimes \varphi_{i,x} Y) = \sqrt{p_{i,j}} \varphi_{j,x} Y, \quad p_{i,j} = \frac{(r_{i,j}^{p,q})^2}{r_{i,j}^0} \sqrt{\frac{D_j^J}{D_i^J}}, \quad (38)$$

for every $x \in X_w$ and $Y \in E_x^{p,q}$. Consequently,

$$D_i^J p_{i,j} = D_j^J p_{j,i}, \quad \sqrt{p_{i,j} p_{j,i}} = \frac{(r_{i,j}^{p,q})^2}{r_{i,j}^0}. \quad (39)$$

Proof. Write σ for uniform probability measure on X_w , and let K_i and K_j be the reproducing kernels of V_i^J and V_j^J . Permutations act transitively on X_w , so these kernels have constant diagonal values $K_i(x, x) = D_i^J$ and $K_j(x, x) = D_j^J$. Hence

$$z_{i,x} = \frac{K_i(x, \cdot)}{\sqrt{D_i^J}}, \quad z_{j,x} = \frac{K_j(x, \cdot)}{\sqrt{D_j^J}}$$

are the unit stabilizer-fixed vectors positive at x . Their ordinary Hahn recurrence, the $p = q = 0$ case of (37), says that

$$\Pi_j^J(t(x, \cdot) z_{i,x}) = r_{i,j}^0 z_{j,x}. \quad (40)$$

Define the coordinate-multiplication map by

$$\mathcal{M}_{i,j} : W \otimes V_i^J \longrightarrow V_j^J, \quad \mathcal{M}_{i,j}(v \otimes f) = \Pi_j^J(\langle v, \ell \rangle f).$$

It commutes with coordinate permutations, so the positive operator $\mathcal{M}_{i,j} \mathcal{M}_{i,j}^*$ on the irreducible Johnson space V_j^J is a scalar multiple of the identity. To identify that scalar, expand the Hilbert–Schmidt norm in orthonormal bases. The coordinate identity following (28), together with permutation invariance and the fixed-line recurrence (40), gives

$$\begin{aligned} \|\mathcal{M}_{i,j}\|_{\text{HS}}^2 &= \iint_{X_w^2} t(y, z) K_i(y, z) K_j(y, z) d\sigma(y) d\sigma(z) \\ &= \int_{X_w} t(x, y) K_i(x, y) K_j(x, y) d\sigma(y) \\ &= r_{i,j}^0 \sqrt{D_i^J D_j^J}. \end{aligned}$$

Taking the trace on the D_j^J -dimensional output space now yields

$$\mathcal{M}_{i,j} \mathcal{M}_{i,j}^* = m_{i,j} \text{id}_{V_j^J}, \quad m_{i,j} = r_{i,j}^0 \sqrt{\frac{D_i^J}{D_j^J}}.$$

Let $\varepsilon_{i,j} = 1$ when $r_{i,j}^{p,q} \geq 0$ and $\varepsilon_{i,j} = -1$ otherwise. Then

$$C_{i,j} = \frac{\varepsilon_{i,j}}{\sqrt{m_{i,j}}} \mathcal{M}_{i,j}^*$$

is an equivariant isometry. By (37), its contraction on $E_x^{p,q}$ is

$$C_{i,j}^*(\ell_x \otimes \varphi_{i,x} Y) = \frac{|r_{i,j}^{p,q}|}{\sqrt{m_{i,j}}} \varphi_{j,x} Y.$$

Squaring this coefficient gives (38). Since $r_{i,j}^{p,q} = r_{j,i}^{p,q}$ and $r_{i,j}^0 = r_{j,i}^0$, applying the same formula in the reverse direction proves (39). $\square$

By Lemma 3.3, the symmetric matrix entry associated with a Johnson transition is the square of its associated Hahn coefficient divided by the corresponding ordinary Hahn coefficient. For $j_- < L \leq j_+$, retain the consecutive Johnson degrees $j_-, \dots, L$. Their coordinate-transition matrix is tridiagonal: its off-diagonal entries connect neighboring Johnson degrees, and a diagonal entry is allowed at each degree. Write $\hat{J} = \hat{J}(n, w, p, q, L)$ for the symmetric matrix

$$\hat{J}_{j,j} = \hat{b}_j^{p,q} := \begin{cases} (b_j^{p,q})^2 / b_j^0, & j \geq 1, \\ 0, & j = 0, \end{cases} \quad (41)$$

$$\hat{J}_{j,j+1} = \hat{J}_{j+1,j} = \hat{c}_j^{p,q} := \frac{(c_j^{p,q})^2}{c_j^0}. \quad (42)$$

With the Johnson dimensions D_j^J from (29), the squared directed contraction coefficients in Lemma 3.3 are

$$p_{j,j+1} = \hat{c}_j^{p,q} \sqrt{\frac{D_{j+1}^J}{D_j^J}}, \quad p_{j+1,j} = \hat{c}_j^{p,q} \sqrt{\frac{D_j^J}{D_{j+1}^J}}, \quad p_{j,j} = \hat{b}_j^{p,q}.$$

For $p = q = 0$, the diagonal and off-diagonal coefficients sum to one over the complete Johnson-degree path. For higher harmonic types, or after truncating that path, the retained coefficients need not have this normalization. Consequently, the Johnson dimensions satisfy

$$D_j^J p_{j,j+1} = D_{j+1}^J p_{j+1,j}. \quad (43)$$

The geometric means of the forward and reverse coefficients are exactly the off-diagonal entries of $\hat{J}$. Since $j_- < L$, the Johnson matrix $\hat{J}$ has a positive largest eigenvalue $\lambda = \lambda_{\max}(\hat{J})$, so $\lambda > s$ implies $\lambda > \max\{s, 0\}$.

Theorem 3.4. *Let n, w, p, q, L, d be integers with $1 \leq w < n/2$, $0 \leq p \leq w/2$, $0 \leq q \leq (n-w)/2$, and*

$$p + q < L \leq \min\{w, w - p + q, n - w + p - q\}.$$

Suppose d is even and $2 \leq d \leq 2w$. Put

$$N = n - w, \quad s = 1 - \frac{nd}{2wN}, \quad \lambda = \lambda_{\max}(\hat{J}(n, w, p, q, L)),$$

where $\hat{J}$ is given by (41)–(42). If $\lambda > s$, then

$$A_J(n, w, d) \leq \frac{1-s}{\lambda-s} \frac{\sum_{j=p+q}^L \left(\binom{n}{j} - \binom{n}{j-1} \right)}{\left(\binom{w}{p} - \binom{w}{p-1} \right) \left(\binom{N}{q} - \binom{N}{q-1} \right)}. \quad (44)$$

Proof. Let $I = \{p+q, \dots, L\}$, let $v = (v_j)_{j \in I}$ be the positive unit Perron eigenvector of the Johnson matrix $\hat{J}$ in (41)–(42), and put

$$\omega_j = \sqrt{D_j^J} v_j, \quad Z = \sum_{j \in I} \omega_j.$$

The Johnson dimension balance (43) gives

$$\sum_{\substack{i \in I \\ |i-j| \leq 1}} p_{i,j} \omega_i = \lambda \omega_j. \quad (45)$$

As in the whole-cube identity (22), dividing by $\lambda \omega_j$ normalizes the incoming transition weights to sum to one. Combine the copies of $E_x^{p,q}$ in the selected Johnson spaces by setting

$$\Psi_x Y = \bigoplus_{j \in I} \sqrt{\frac{\omega_j}{Z}} \varphi_{j,x} Y, \quad P_x = \Psi_x \Psi_x^*.$$

These projections have rank $d_{p,q}^J$ on the ambient space $\mathcal{V} = \bigoplus_{j \in I} V_j^J$.

Write $W = \mathbf{1}^\perp \subseteq \mathbb{R}^n$, and let $C_{i,j} : V_j^J \rightarrow W \otimes V_i^J$ be the normalized Johnson coordinate isometries constructed in [Lemma 3.3](#) for $(i,j) \neq (0,0)$; set $C_{0,0} = 0$, consistently with $p_{0,0} = 0$. For fixed i , the nonzero images for distinct j are orthogonal: they are invariant irreducible subspaces of $W \otimes V_i^J$ carrying pairwise inequivalent permutation representations V_j^J . Their contraction coefficients on $E_x^{p,q}$ are $\sqrt{p_{i,j}}$, by [Lemma 3.3](#). Define $\mathcal{B} : \mathcal{V} \rightarrow W \otimes \mathcal{V}$ by

$$\mathcal{B}f = \bigoplus_{i \in I} \sum_{\substack{j \in I \\ |i-j| \leq 1}} \sqrt{\frac{p_{i,j} \omega_i}{\lambda \omega_j}} C_{i,j} f_j, \quad f = \bigoplus_{j \in I} f_j \in \mathcal{V}.$$

Orthogonality and the Perron identity (45) show that $\mathcal{B}$ is an isometry and $\mathcal{B}^*(\ell_x \otimes \Psi_x Y) = \sqrt{\lambda} \Psi_x Y$. Consequently, if $\Theta_x = (\ell_x \otimes \text{id})P_x - \sqrt{\lambda} \mathcal{B}P_x$, direct expansion gives

$$\langle \Theta_x, \Theta_y \rangle_{\text{HS}} = (t(x, y) - \lambda) \text{tr}(P_x P_y).$$

For a constant-weight code C of size N_C , the off-diagonal terms of $(t(x, y) - s) \text{tr}(P_x P_y)$ are nonpositive, and the diagonal terms equal $(1 - s)d_{p,q}^J$. Summation and trace Cauchy-Schwarz therefore give

$$(\lambda - s) \frac{N_C^2 (d_{p,q}^J)^2}{\sum_{j=p+q}^L D_j^J} \leq N_C (1 - s) d_{p,q}^J.$$

Substituting $D_j^J = \binom{n}{j} - \binom{n}{j-1}$ and $d_{p,q}^J = d_p(w) d_q(N)$ proves (44). The overlap is a scalar function of Johnson distance because coordinate permutations are transitive on pairs of supports having the same intersection size. $\square$

3.3. The asymptotic constant-weight bound. The finite constant-weight certificate in [Theorem 3.4](#) bounds a code by the ratio between the total dimensions of the retained Johnson spaces and the dimension of the attached Boolean harmonic space. We now pass to high-dimensional limits, estimate its largest eigenvalue, and include the Bassalygo–Elias cost of returning to the whole cube.

Fix $0 < \delta < 1/2$, and recall the normalized parameters $\alpha = w/n$, $\beta = p/n$, $\gamma = q/n$, and $u = L/n$ from the introduction. They represent the relative layer weight, the two normalized Boolean harmonic degrees, and the largest retained Johnson degree, respectively. Their ranges are (4)–(5), and (6) gives an asymptotic lower bound for the largest eigenvalue of the Johnson recurrence matrix $\hat{J}$.

Lemma 3.5. *Suppose $(\alpha, \beta, \gamma, u)$ satisfies (4) and (5), and choose integers with*

$$\frac{w}{n} \rightarrow \alpha, \quad \frac{p}{n} \rightarrow \beta, \quad \frac{q}{n} \rightarrow \gamma, \quad \frac{L}{n} \rightarrow u.$$

For the Johnson tridiagonal matrix $\hat{J}(n, w, p, q, L)$, one has

$$\liminf_{n \rightarrow \infty} \lambda_{\max}(\hat{J}(n, w, p, q, L)) \geq \Lambda_{\alpha, \beta, \gamma}(u). \quad (46)$$

Proof. Let z, m, ζ, ξ be the normalized coordinates preceding (6). Uniformly for $|j/n - u| \rightarrow 0$, substitution in (35) gives

$$\begin{aligned} b_j^{p,q} &\rightarrow \frac{m(\zeta\xi - mz^2)}{z^2(1 - m^2)}, & b_j^0 &\rightarrow \frac{m^2(1 - z^2)}{z^2(1 - m^2)}, \\ c_j^{p,q} &\rightarrow \frac{\sqrt{(z^2 - m^2)(z^2 - \xi^2)(\zeta^2 - z^2)}}{2z^2(1 - m^2)}, & c_j^0 &\rightarrow \frac{(z^2 - m^2)\sqrt{1 - z^2}}{2z^2(1 - m^2)}. \end{aligned}$$

The strict inequalities in (4) and (5) keep all denominators bounded away from zero and place the terminal degree strictly inside the Johnson-degree range (31), so $j_- < L < j_+$ for large n . Choose $m_n \rightarrow \infty$ with $m_n = o(n)$. Then $L - m_n \geq j_-$ and $L < j_+$ for large n , so the discrete sine vector

$$f_{L-m_n+r} = \sin \frac{\pi r}{m_n + 1} \quad (1 \leq r \leq m_n)$$

is supported on Johnson degrees $j_- \leq j \leq j_+$.

By (41) and (42), the diagonal entries on this terminal degree interval converge uniformly to

$$B = \frac{(\zeta\xi - mz^2)^2}{z^2(1 - m^2)(1 - z^2)},$$

while the neighboring-degree entries converge uniformly to

$$C = \frac{(z^2 - \xi^2)(\zeta^2 - z^2)}{2z^2(1 - m^2)\sqrt{1 - z^2}}.$$

The Rayleigh–Ritz principle (26) applied to the displayed sine vector gives

$$\lambda_{\max}(\hat{J}) \geq B + 2C \cos \frac{\pi}{m_n + 1} + o(1).$$

Since $B + 2C = \Lambda_{\alpha,\beta,\gamma}(u)$, taking $n \rightarrow \infty$ proves (46). $\square$

Theorem 3.6. *Fix $0 < \delta < 1/2$. If $(\alpha, \beta, \gamma, u) \in \mathcal{D}_\delta$, then*

$$R_2(\delta) \leq 1 - H_2(\alpha) + H_2(u) - \alpha H_2\left(\frac{\beta}{\alpha}\right) - (1 - \alpha)H_2\left(\frac{\gamma}{1 - \alpha}\right). \quad (47)$$

Proof. Choose integers with $w/n \rightarrow \alpha$, $p/n \rightarrow \beta$, $q/n \rightarrow \gamma$, and $L/n \rightarrow u$, and put $d_n = 2\lceil \delta n \rceil / 2$. This is the largest even integer not exceeding $\lceil \delta n \rceil$, so $d_n/n \rightarrow \delta$ and $A_2(n, \lceil \delta n \rceil) \leq A_2(n, d_n)$. Since $\alpha > \delta/2 > 0$, the finite-theorem hypotheses $2 \leq d_n \leq 2w$ hold for all sufficiently large n . By Lemma 3.5, the largest eigenvalue of $\hat{J}(n, w, p, q, L)$ has limiting lower bound $\Lambda_{\alpha,\beta,\gamma}(u)$. The dimensions of the retained Johnson spaces telescope to

$$\sum_{j=p+q}^L \left(\binom{n}{j} - \binom{n}{j-1} \right) = \binom{n}{L} - \binom{n}{p+q-1} = 2^{nH_2(u) + o(n)}.$$

Stirling’s formula and (30) likewise give

$$\frac{1}{n} \log_2 d_{p,q}^J = \alpha H_2\left(\frac{\beta}{\alpha}\right) + (1 - \alpha)H_2\left(\frac{\gamma}{1 - \alpha}\right) + o(1).$$

The finite threshold $1 - nd_n/(2w(n - w))$ converges to the right-hand side of (7). Thus strict feasibility and (46) bound the prefactor in (44). Taking the logarithmic rate of this finite Johnson bound and adding the Bassalygo–Elias cost $1 - H_2(\alpha)$ from (27) proves (47). $\square$

Taking the infimum in (47) over $\mathcal{D}_\delta$ recovers the constant-weight exponent (8) and gives

$$R_2(\delta) \leq \kappa_{\text{CW}}(\delta).$$

3.4. Strict comparison with the classical bound. We now show that the two binary constructions together strictly improve the fully optimized classical bound at every relative distance. Recall the MRRW objective F_δ , its endpoint values $M_1(\delta)$ and $F_\delta(0)$, and its optimized value $M_2(\delta)$ from (1)–(3). In particular, $M_2(\delta) \leq M_1(\delta)$. We first identify the classical constant-weight objective inside our construction by taking the one-dimensional fixed-line choice $p = q = 0$. A higher-dimensional choice then improves every interior minimizer, while the whole-cube certificate handles the endpoint with value M_1 .

We first identify the fixed-line specialization $\beta = \gamma = 0$ with the fully optimized classical objective (3). Write $A = \alpha(1 - \alpha)$ and $U = u(1 - u)$. Then

$$1 - \Lambda_{\alpha,0,0}(u) = \frac{A - U}{A(1 + 2\sqrt{U})}. \quad (48)$$

On the spectral boundary, $\tau = 2\sqrt{U}$ gives

$$4A = \tau^2 + 2\delta\tau + 2\delta, \quad 1 - H_2(\alpha) + H_2(u) = 1 + g(\tau^2) - g(\tau^2 + 2\delta\tau + 2\delta).$$

Equivalently, $u = (1 - \sqrt{1 - \tau^2})/2$ and $\alpha = (1 - \sqrt{1 - \tau^2 - 2\delta\tau - 2\delta})/2$. Hence $0 < \tau < 1 - 2\delta$ corresponds exactly to $0 < u < \alpha < 1/2$; moreover, $\alpha > \delta/2$, so the shell-domain conditions hold. Every boundary point is approached by strictly feasible $u' > u$, while $\tau = 0$ and $\tau = 1 - 2\delta$ are obtained as $u \downarrow 0$ and $\alpha \uparrow 1/2$, respectively. Consequently,

$$\kappa_{\text{CW}}(\delta) \leq M_2(\delta).$$

Proposition 3.7. Fix $0 < \delta < 1/2$, and suppose $0 < u < \alpha < 1/2$ and

$$\Lambda_{\alpha,0,0}(u) = 1 - \frac{\delta}{2\alpha(1-\alpha)}.$$

Then there exist $\beta, \gamma > 0$ and u' with $(\alpha, \beta, \gamma, u') \in \mathcal{D}_\delta$ whose rate in (47) is strictly smaller than $1 - H_2(\alpha) + H_2(u)$.

Proof. Put $\beta = \gamma = \varepsilon$. By (48),

$$\partial_u \Lambda_{\alpha,0,0}(u) = (1-2u) \left[\frac{1}{A(1+2\sqrt{U})} + \frac{A-U}{A\sqrt{U}(1+2\sqrt{U})^2} \right] > 0.$$

Here $A = \alpha(1-\alpha)$ and $U = u(1-u)$. Smoothness at the interior point shows that replacing $(\beta, \gamma) = (0, 0)$ by $(\varepsilon, \varepsilon)$ changes the spectral value by $O(\varepsilon)$. Thus, for a sufficiently large fixed C , the choice $u_\varepsilon = u + C\varepsilon$ satisfies (4)–(7) strictly. The ambient entropy increases by $H_2(u_\varepsilon) - H_2(u) = O(\varepsilon)$, whereas the two harmonic-space dimensions contribute

$$\alpha H_2\left(\frac{\varepsilon}{\alpha}\right) + (1-\alpha) H_2\left(\frac{\varepsilon}{1-\alpha}\right) = 2\varepsilon \log_2 \frac{1}{\varepsilon} + O(\varepsilon).$$

This gain dominates $O(\varepsilon)$, proving the strict improvement. $\square$

Theorem 3.8. For every fixed $0 < \delta < 1/2$,

$$\kappa_{\text{bin}}(\delta) = \min\{\kappa_{\text{H}}(\delta), \kappa_{\text{CW}}(\delta)\} < M_2(\delta).$$

Proof. Suppose first that $M_2(\delta) < M_1(\delta)$. The endpoint $\tau = 1 - 2\delta$ cannot minimize F_δ , since its value is $M_1(\delta)$. At the other endpoint $\tau = 0$, $g(\tau^2) = O(\tau^2 \log(1/\tau)) = o(\tau)$, so $F'_\delta(0+) = -2\delta g'(2\delta) < 0$. Hence every minimizer is interior, and Proposition 3.7 gives $\kappa_{\text{CW}}(\delta) < M_2(\delta)$. If $M_2(\delta) = M_1(\delta)$, Theorem 2.3 instead gives $\kappa_{\text{H}}(\delta) < M_2(\delta)$. In either case, $\kappa_{\text{bin}}(\delta) < M_2(\delta)$. $\square$

4. REPRESENTATION GRAPHS FOR SPHERICAL CODES

The whole-cube construction in §2 uses the tridiagonal matrix J_{H} in (20), indexed by Fourier degree. The constant-weight construction in §3 uses a tridiagonal matrix indexed by Johnson degree, with off-diagonal entries (42). In each case, coordinate multiplication connects neighboring degrees. The simplest spherical construction similarly uses a tridiagonal matrix indexed by harmonic degree, coming from the Gegenbauer recurrence (65). Stronger spherical constructions, however, require ambient spaces indexed by several integers, with coordinate transitions in several independent directions.

This section extends the elementary projection argument to an arbitrary finite connected weighted graph. The vertices represent rotation-invariant ambient spaces containing a common moving subspace, and the edges record the action of the distance coordinate. The concrete spherical ambient and stabilizer spaces are introduced in §5; their explicit coordinate-transition weights are computed in §6.

4.1. The general projection bound. We first isolate the projection argument that underlies both the concrete whole-cube calculation and the more general spherical construction. The only inputs are moving orthogonal projections and an isometry compatible with the distance coordinate.

Proposition 4.1. Let X be a set, and associate a unit vector ℓ_x in a Euclidean space W with every $x \in X$. Write $t(x, y) = \langle \ell_x, \ell_y \rangle$, let $\mathcal{V}$ be a D -dimensional Hilbert space, and let P_x be a rank- d orthogonal projection on $\mathcal{V}$ for every $x \in X$, where $1 \leq d \leq D$. Suppose an isometry $\mathcal{B} : \mathcal{V} \rightarrow W \otimes \mathcal{V}$ satisfies, for some $\Lambda \geq 0$, every $x \in X$, and every $u \in \text{im } P_x$,

$$\mathcal{B}^*(\ell_x \otimes u) = \sqrt{\Lambda} u. \tag{49}$$

If $\Lambda > s$, every code $C \subseteq X$ satisfying $t(x, y) \leq s$ at distinct code points obeys

$$|C| \leq \frac{1-s}{\Lambda-s} \frac{D}{d}. \tag{50}$$

Proof. It suffices to consider finite codes: an infinite code would contain finite subcodes of arbitrarily large cardinality. The projection overlap satisfies

$$K(x, y) = \text{tr}(P_x P_y) = \|P_x P_y\|_{\text{HS}}^2 \geq 0.$$

Define the residual map

$$\Theta_x = (\ell_x \otimes \text{id})P_x - \sqrt{\Lambda} \mathcal{B} P_x.$$

The same Hilbert–Schmidt calculation as in (25), using $\mathcal{B}^* \mathcal{B} = \text{id}_{\mathcal{V}}$ and (49), gives

$$\begin{aligned} \langle \Theta_x, \Theta_y \rangle_{\text{HS}} &= (t(x, y) - \Lambda) K(x, y), \\ (t(x, y) - s) K(x, y) &= (\Lambda - s) K(x, y) + \langle \Theta_x, \Theta_y \rangle_{\text{HS}}. \end{aligned} \quad (51)$$

The left side is nonpositive at distinct code points and equals $(1-s)d$ on the diagonal. Summing over C^2 , discarding the nonnegative Gram term, and applying trace Cauchy–Schwarz give

$$(\Lambda - s) \frac{|C|^2 d^2}{D} \leq (\Lambda - s) \left\| \sum_{x \in C} P_x \right\|_{\text{HS}}^2 \leq |C|(1-s)d.$$

Rearrangement proves (50). $\square$

We now apply Proposition 4.1 to moving subspaces on the sphere. The rotation group and the subgroup fixing a base point determine which subspaces can be assembled into a projection.

Let $n \geq 3$, let $X = \mathbb{S}^{n-1}$, and let $G = SO(n)$ be its rotation group. Give $W = \mathbb{R}^n$ the usual rotation action. For $x \in X$, the unit distance-coordinate vector is $\ell_x = x$, so $t(x, y) = \langle x, y \rangle$. Fix a base point o , and write H for the rotations fixing o . The action of H on $o^\perp$ identifies H with $SO(n-1)$.

We call an H -representation a stabilizer representation because $H = \text{Stab}_{SO(n)}(o)$ is the subgroup fixing the base point o . All representations below are finite-dimensional real Hilbert spaces with orthogonal group actions. For example, H acts trivially on the fixed line $\mathbb{R}o$, acts by ordinary rotations on $o^\perp$, and acts on the degree- k harmonic polynomials $\mathcal{H}_k(o^\perp)$ by rotating their input vectors. These examples have dimensions 1, $n-1$, and $\dim \mathcal{H}_k(o^\perp)$, respectively.

Choose a d -dimensional irreducible stabilizer representation E , and let Ω be a finite collection of pairwise inequivalent irreducible ambient G -representations V_λ . We require the restriction of each V_λ to H to contain exactly one copy of E :

$$\dim_{\mathbb{R}} \text{Hom}_H(E, V_\lambda) = 1, \quad D_\lambda = \dim V_\lambda.$$

The one-dimensional intertwiner space determines a unique copy of E in each V_λ . Identify E with its copy in one fixed ambient representation, and choose isometric embeddings $\varphi_{\lambda, o} : E \rightarrow V_\lambda$ at the base point. Composing any H -equivariant endomorphism of E with $\varphi_{\lambda, o}$ produces another intertwiner. Hence the one-dimensional intertwiner hypothesis also gives $\text{End}_H(E) = \mathbb{R} \text{id}_E$. If $x = go$, transport the stabilizer space and all its embeddings simultaneously:

$$E_x = gE, \quad \varphi_{\lambda, x}(gY) = g\varphi_{\lambda, o}Y \quad (Y \in E).$$

These definitions are independent of the rotation carrying o to x : two such rotations differ by an element of H , and every $\varphi_{\lambda, o}$ commutes with H . In particular, a single rotation transports all embeddings, so their relative signs are preserved.

4.2. Coordinate transitions and their balanced graph. We next define the spherical transition graph and its weights. The vertices are the ambient spaces V_λ . If V_ν occurs inside $W \otimes V_\lambda$, its inclusion gives a possible coordinate transition from λ to ν . Choose an isometric map commuting with rotations:

$$C_{\lambda, \nu} : V_\nu \longrightarrow W \otimes V_\lambda.$$

For fixed λ , the images associated with distinct ν are orthogonal because they are inequivalent irreducible subrepresentations of $W \otimes V_\lambda$. Contracting with the coordinate vector ℓ_x gives a map between the two copies of E_x . Since these copies occur with multiplicity one, $\text{End}_H(E) = \mathbb{R} \text{id}_E$, and the map commutes with the rotations fixing x , it acts by a real scalar $c_{\lambda, \nu}$:

$$C_{\lambda, \nu}^*(\ell_x \otimes \varphi_{\lambda, x} Y) = c_{\lambda, \nu} \varphi_{\nu, x} Y. \quad (52)$$

Choose the sign of each coordinate map so that $c_{\lambda,\nu} \geq 0$, and write $p_{\lambda,\nu} = c_{\lambda,\nu}^2$. The graph retains an edge only when this coefficient is nonzero. Set $p_{\lambda,\nu} = 0$ when no such inclusion exists or its contraction coefficient vanishes. Thus occurrence in the tensor product allows a transition, but the chosen stabilizer representation determines whether its edge is present. Opposite directions generally have different squared coefficients. We require the following dimension-weighted reciprocity:

$$D_\lambda p_{\lambda,\nu} = D_\nu p_{\nu,\lambda}. \quad (53)$$

Thus multiplying each directed weight by the dimension of the source representation makes the two directions agree. For the spherical transitions, this identity is verified explicitly in (76) and §A.2. The undirected weighted adjacency matrix J therefore has symmetric edge weights

$$J_{\lambda,\nu} = p_{\lambda,\nu} \sqrt{\frac{D_\lambda}{D_\nu}} = \sqrt{p_{\lambda,\nu} p_{\nu,\lambda}}. \quad (54)$$

Indeed, if $\mathbf{T} = (p_{\lambda,\nu})$ and $\mathbf{D} = \text{diag}(D_\lambda)$, dimension balance (53) gives the symmetric similarity transform

$$J = \mathbf{D}^{1/2} \mathbf{T} \mathbf{D}^{-1/2}. \quad (55)$$

Assume that the finite graph on Ω is connected and has positive edge weights. The Perron–Frobenius theorem supplies a strictly positive eigenvector of the weighted adjacency matrix J for the largest eigenvalue $\Lambda = \lambda_{\max}(J)$.

Theorem 4.2. *Let $n \geq 3$, let $G = SO(n)$, and let $H = \text{Stab}_G(o) \simeq SO(n-1)$ for a base point $o \in \mathbb{S}^{n-1}$. Let E be a d -dimensional irreducible real H -representation, and let Ω consist of finitely many pairwise inequivalent irreducible real G -representations V_λ whose coordinate-transition graph is connected. Suppose $\dim_{\mathbb{R}} \text{Hom}_H(E, V_\lambda) = 1$ for each $\lambda \in \Omega$, and suppose the coordinate-transition coefficients satisfy (52) and (53). Let J be the weighted adjacency matrix in (54), and put $\Lambda = \lambda_{\max}(J)$ and $D_\lambda = \dim V_\lambda$. If $-1 < s < 1$ and $\Lambda > \max\{s, 0\}$, then*

$$A(n, s) \leq \frac{1-s}{d(\Lambda-s)} \sum_{\lambda \in \Omega} D_\lambda. \quad (56)$$

Proof. Let v be the positive unit Perron eigenvector of J in (54), and put $w_\lambda = \sqrt{D_\lambda} v_\lambda$ and $Z = \sum_\lambda w_\lambda$. The dimension-weighted reciprocity (53) and the eigenvector identity $Jv = \Lambda v$ give

$$\sum_{\lambda \in \Omega} p_{\lambda,\nu} w_\lambda = \sqrt{D_\nu} \sum_{\lambda \in \Omega} J_{\nu,\lambda} v_\lambda = \Lambda w_\nu. \quad (57)$$

Consequently, $p_{\lambda,\nu} w_\lambda / (\Lambda w_\nu)$ sums to one over $\lambda \in \Omega$ for each fixed ν . These Perron-normalized weights form the probability distributions used in the coordinate isometry below. Thus

$$\Psi_x Y = \bigoplus_{\lambda \in \Omega} \sqrt{\frac{w_\lambda}{Z}} \varphi_{\lambda,x} Y$$

is an isometric embedding into $\mathcal{V} = \bigoplus_{\lambda \in \Omega} V_\lambda$. Set $P_x = \Psi_x \Psi_x^*$. For $h = \bigoplus_{\nu \in \Omega} h_\nu \in \mathcal{V}$, define $\mathcal{B} : \mathcal{V} \rightarrow W \otimes \mathcal{V}$ by

$$\mathcal{B}h = \bigoplus_{\lambda \in \Omega} \sum_{\substack{\nu \in \Omega \\ p_{\lambda,\nu} > 0}} \sqrt{\frac{p_{\lambda,\nu} w_\lambda}{\Lambda w_\nu}} C_{\lambda,\nu} h_\nu. \quad (58)$$

Orthogonality of the images of $C_{\lambda,\nu}$ and (57) give

$$\|\mathcal{B}h\|^2 = \sum_{\nu \in \Omega} \frac{\|h_\nu\|^2}{\Lambda w_\nu} \sum_{\lambda \in \Omega} p_{\lambda,\nu} w_\lambda = \|h\|^2,$$

so $\mathcal{B}$ is an isometry. Using (52) and then (57) gives

$$\mathcal{B}^*(\ell_x \otimes \Psi_x Y) = \bigoplus_{\nu \in \Omega} \frac{\varphi_{\nu,x} Y}{\sqrt{\Lambda w_\nu Z}} \sum_{\lambda \in \Omega} p_{\lambda,\nu} w_\lambda = \bigoplus_{\nu \in \Omega} \sqrt{\frac{\Lambda w_\nu}{Z}} \varphi_{\nu,x} Y = \sqrt{\Lambda} \Psi_x Y.$$

Proposition 4.1, with $D = \sum_{\lambda} D_{\lambda}$, proves (56). Finally, define $F(x, y) = (\langle x, y \rangle - s) \operatorname{tr}(P_x P_y)$. Rotational invariance makes $\operatorname{tr}(P_x P_y)$ a function of $\langle x, y \rangle$ alone. The residual Gram decomposition (51) makes F positive definite, and $F(x, y) \leq 0$ when $\langle x, y \rangle \leq s$. If σ is uniform probability measure on the sphere, its constant coefficient satisfies

$$f_0 = \iint F(x, y) d\sigma(x) d\sigma(y) \geq (\Lambda - s) \left\| \int P_x d\sigma(x) \right\|_{\text{HS}}^2 \geq (\Lambda - s) \frac{d^2}{\sum_{\lambda \in \Omega} D_{\lambda}} > 0.$$

Thus F satisfies every requirement of a two-point Delsarte certificate. $\square$

The full coordinate decomposition can also contain edges leading to ambient spaces outside the selected vertex set Ω . A finite truncation discards those edges, so the retained transition weights need not sum to one. This causes no difficulty: (57) normalizes the Perron-reweighted coefficients in (58), making that operator an isometry. Thus **Theorem 4.2** applies to arbitrary finite connected truncations, including multidimensional rectangular boxes.

5. SPHERICAL HARMONICS AND STABILIZER REPRESENTATIONS

To apply the moving-subspace bound of **Theorem 4.2** to spherical codes, we need a family of $SO(n)$ -representations, each containing the same representation of the point stabilizer $SO(n-1)$. We must also identify the coordinate transitions between the ambient representations and the dimensions of their moving subspaces. This section describes the relevant representations and works out the one-dimensional construction. The general transition weights are computed in §6, and their spectral and dimension asymptotics are analyzed in §7.

Throughout the spherical arguments, assume $n \geq 3$. The multi-row construction will impose the stronger stable-range condition $n \geq 2r + 4$. This restriction does not affect the asymptotic bounds: every hierarchy level r is fixed before $n \rightarrow \infty$.

The group $SO(n)$ acts transitively on $\mathbb{S}^{n-1}$. For a point x , the rotations fixing x act on its orthogonal complement

$$x^{\perp} = \{u \in \mathbb{R}^n : \langle u, x \rangle = 0\},$$

and every rotation of $x^{\perp}$ extends uniquely to a rotation fixing x . Hence the point stabilizer of x is naturally identified with $SO(n-1)$. We call a representation of this point-fixing subgroup a stabilizer representation. The fixed line $\mathbb{R}x$, on which every such rotation acts trivially, is the one-dimensional example; the tangent space $x^{\perp}$, on which the same rotations act in the usual way, is a higher-dimensional example. A rotation carrying x to another point y also carries the tangent space $x^{\perp}$ to $y^{\perp}$. Thus the stabilizer representations used below move naturally with the code point.

We first recall the classical construction, which retains the one-dimensional stabilizer-fixed line in each selected space of spherical harmonics. We next replace this line by a space of degree- k harmonics on $x^{\perp}$, obtaining a weighted one-dimensional path. Finally, more general stabilizer representations yield multidimensional transition graphs and the full hierarchy.

5.1. Spherical harmonics and the classical bound. Classical spherical linear programming uses harmonic spaces, their stabilizer-fixed lines, and the associated Gegenbauer kernels. The resulting tridiagonal spectral construction uses one fixed line in each harmonic degree, just as the classical whole-cube construction uses one fixed line in each Fourier degree in §2.1.

Equip $\mathbb{S}^{n-1}$ with rotation-invariant probability measure. A polynomial on $\mathbb{R}^n$ is homogeneous of degree m when scaling its argument by c scales its value by c^m , and it is harmonic when its Euclidean Laplacian vanishes. Let $V_m^{\text{S}} = \mathcal{H}_m(\mathbb{R}^n)$ denote the restrictions of these homogeneous harmonic polynomials to $\mathbb{S}^{n-1}$. Rotations act on V_m^{S} by changing the input of a polynomial. These mutually orthogonal irreducible spaces are the spherical analogues of Fourier levels [GW09, §§8.1.1–8.1.2]. For example, V_0^{S} consists of constant functions, and V_1^{S} consists of the linear functions $u \mapsto \langle v, u \rangle$. The integer m records polynomial degree, not the dimension of the harmonic space. That dimension is

$$D_m^{\text{S}} = \dim V_m^{\text{S}} = \frac{2m + n - 2}{n - 2} \binom{m + n - 3}{m}. \quad (59)$$

Fix $x \in \mathbb{S}^{n-1}$. A function invariant under all rotations fixing x depends only on its radial coordinate $t = \langle x, u \rangle$. Inside each V_m^S , the invariant functions form a one-dimensional subspace, called the stabilizer-fixed line. Its generator normalized by $P_m(1) = 1$ is $P_m(t) = C_m^{(n-2)/2}(t)/C_m^{(n-2)/2}(1)$, where C_m^η denotes the degree- m Gegenbauer polynomial. For an orthonormal basis $Y_{m,1}, \dots, Y_{m,D_m^S}$ of V_m^S , define its reproducing kernel by

$$K_m^S(u, v) = \sum_{a=1}^{D_m^S} Y_{m,a}(u) Y_{m,a}(v).$$

This kernel is independent of the chosen orthonormal basis. The spherical addition formula identifies it explicitly [BV08, (4)]:

$$K_m^S(u, v) = D_m^S P_m(\langle u, v \rangle). \quad (60)$$

The defining orthonormal-basis sum is positive definite. Therefore, so is $P_m(\langle u, v \rangle)$, and non-negative Gegenbauer coefficients suffice for positive definiteness. (The converse is Schoenberg's theorem [S42] and is not needed here.) Thus, if $f = \sum_{m=0}^M f_m P_m$ is a polynomial, $f_0 > 0$, $f_m \geq 0$, and $f(t) \leq 0$ for $-1 \leq t \leq s$, the spherical linear-programming bound is [DGS77]

$$A(n, s) \leq \frac{f(1)}{f_0}.$$

Just as in the Fourier-path construction of §2.1, the trivial stabilizer representation in V_m^S is therefore the line spanned by $P_m(\langle x, \cdot \rangle)$. The Gegenbauer recurrence (64) shows that multiplication by $u \mapsto \langle x, u \rangle$ connects this line only to the fixed lines of degrees $m-1$ and $m+1$. Thus consecutive harmonic degrees form a weighted path, whose tridiagonal matrix gives the classical spectral certificate [BN06, §3.3]. The explicit zonal polynomials identify the classical fixed lines. The addition formula (60) also normalizes the coordinate maps in Lemma 5.1; apart from this normalization, the new construction uses the harmonic spaces and normalized recurrence (65)–(66). Lemma 5.2 below calculates the largest-eigenvalue limit of this path; its case $k=0$ recovers the direct whole-sphere Kabatianskii–Levenshtein spectral bound [KL78], before the additional spherical-cap optimization.

5.2. Harmonics on the tangent sphere. The classical construction associates one stabilizer-fixed vector with each selected harmonic degree. We replace that vector with a common space E_x of harmonic polynomials on $x^\perp$, embedded inside several ambient harmonic spaces V_i^S . Consecutive ambient degrees are connected by a symmetric tridiagonal matrix, as in the Fourier-degree construction of §2. The projection bound Proposition 4.1 then divides by $\dim E_x$ rather than by the classical rank 1. For positive harmonic degree, assume $n \geq 4$, so the stabilizer representations used below have real-scalar endomorphisms. We now describe E_x , its embeddings $\varphi_{i,x}$, the weighted adjacency matrix $J_{k,L}$, and the resulting spherical-code bound.

For a base point x , the unit sphere in its tangent space is

$$\mathbb{S}(x^\perp) = x^\perp \cap \mathbb{S}^{n-1}.$$

Equip this sphere with its rotation-invariant probability measure. The stabilizer $SO(n-1)$ acts on it by rotating its tangent directions. For $k \geq 0$, let

$$E_x = \mathcal{H}_k(x^\perp), \quad d_k^S = \dim E_x = \binom{n+k-2}{k} - \binom{n+k-4}{k-2}. \quad (61)$$

Here and below, a binomial coefficient with negative lower index is zero. Thus E_x consists of degree- k homogeneous harmonic polynomials on $x^\perp$, restricted to $\mathbb{S}(x^\perp)$. Equivalently, its elements are symmetric traceless k -tensors on $x^\perp$: a symmetric tensor T represents the polynomial $\xi \mapsto T(\xi, \dots, \xi)$, and harmonicity means that contracting any two tensor indices gives zero. In particular, k is the polynomial degree or tensor order; it is not the dimension d_k^S , which is generally much larger. For $k=0$, E_x is the classical one-dimensional space of constants; for $k=1$, it is naturally identified with $x^\perp$; and for $k=2$, it consists of traceless quadratic forms on $x^\perp$.

The tangent harmonic spaces $E_x = \mathcal{H}_k(x^\perp)$ also appear in the matrix-valued kernels of Bachoc and Vallentin [BV08, (9), Thms. 3.1–3.2]. There the stabilizer fixes one distinguished base point, producing kernels in three inner products. Here E_x moves with the code point, and the overlap of the corresponding moving projections remains a scalar two-point kernel.

To see which ambient harmonic spaces contain E_x , restrict their rotation action from $SO(n)$ to the subgroup fixing x . The multiplicity-free orthogonal-group branching rule gives [GW09, Thms. 8.1.3–8.1.4]

$$V_i^S \downarrow_{SO(n-1)} \simeq \bigoplus_{k=0}^i \mathcal{H}_k(x^\perp). \quad (62)$$

Every summand occurs with multiplicity one, meaning that V_i^S contains exactly one subspace of each indicated harmonic type. In particular, it contains one copy of E_x if $i \geq k$, and none if $i < k$. Write $\varphi_{i,x} : E_x \rightarrow V_i^S$ for the isometric embedding of this unique copy.

Every $u \in \mathbb{S}^{n-1}$ away from the poles $u = \pm x$ has the unique decomposition

$$u = tx + \sqrt{1-t^2} \xi, \quad t = \langle x, u \rangle, \quad \xi \in \mathbb{S}(x^\perp).$$

In these coordinates, every element of the copy of E_x inside V_i^S separates into a tangential harmonic Y of degree k and a radial polynomial of degree $i-k$:

$$(\varphi_{i,x} Y)(u) = p_{i-k}(t)(1-t^2)^{k/2} Y(\xi), \quad (63)$$

where $Y \in E_x$ and the radial polynomial p_{i-k} is a scalar multiple of $C_{i-k}^{k+(n-2)/2}$ [BV08, (8), proof of Thm. 3.2]. Choose this scalar multiple to have positive leading coefficient and to make $\varphi_{i,x}$ an isometry. These choices fix the signs of all the embeddings, and the expression extends continuously over the poles. Multiplication by the base-point coordinate $t = \langle x, u \rangle$ preserves the tangential harmonic Y and acts only on its radial polynomial. Set $\eta = k + (n-2)/2$. The standard Gegenbauer identity [BN06, §3.3]

$$2(j+\eta)tC_j^\eta(t) = (j+1)C_{j+1}^\eta(t) + (j+2\eta-1)C_{j-1}^\eta(t) \quad (64)$$

with $j = i-k$, where $C_{-1}^\eta = 0$, shows that only degrees $i+1$ and $i-1$ occur. Applying the isometric normalization in (63) gives

$$t\varphi_{i,x} = \alpha_i^{(k)} \varphi_{i+1,x} + \beta_i^{(k)} \varphi_{i-1,x}, \quad (65)$$

with

$$(\alpha_i^{(k)})^2 = \frac{(i-k+1)(i+k+n-2)}{(2i+n-2)(2i+n)}, \quad (\beta_i^{(k)})^2 = \frac{(i-k)(i+k+n-3)}{(2i+n-4)(2i+n-2)}. \quad (66)$$

The recurrence (65) describes scalar multiplication by $u \mapsto \langle x, u \rangle$. To apply the projection bound Theorem 4.2, we must instead normalize the vector-valued coordinate maps between neighboring harmonic degrees. The required normalization is determined by the classical fixed-line coefficient $\alpha_i^{(0)}$ in (66).

Lemma 5.1. *Let $i \geq k$, abbreviate $D_j = D_j^S$, and let Π_j be the orthogonal projection onto V_j^S . Define the upward and downward coordinate-multiplication maps by*

$$\begin{aligned} M_i^\uparrow : \mathbb{R}^n \otimes V_i^S &\longrightarrow V_{i+1}^S, & M_i^\uparrow(v \otimes f) &= \Pi_{i+1}[\langle v, \cdot \rangle f], \\ M_i^\downarrow : \mathbb{R}^n \otimes V_{i+1}^S &\longrightarrow V_i^S, & M_i^\downarrow(v \otimes g) &= \Pi_i[\langle v, \cdot \rangle g]. \end{aligned}$$

These maps satisfy

$$M_i^\uparrow (M_i^\uparrow)^* = m_i^\uparrow \text{id}_{V_{i+1}^S}, \quad M_i^\downarrow (M_i^\downarrow)^* = m_i^\downarrow \text{id}_{V_i^S},$$

where the normalization constants are

$$m_i^\uparrow = \alpha_i^{(0)} \sqrt{\frac{D_i}{D_{i+1}}} = \frac{i+1}{2i+n}, \quad m_i^\downarrow = \alpha_i^{(0)} \sqrt{\frac{D_{i+1}}{D_i}} = \frac{i+n-2}{2i+n-2}. \quad (67)$$

Consequently, the coordinate inclusions

$$C_i^\uparrow = \frac{(M_i^\uparrow)^*}{\sqrt{m_i^\uparrow}} : V_{i+1}^S \longrightarrow \mathbb{R}^n \otimes V_i^S,$$

$$C_i^\downarrow = \frac{(M_i^\downarrow)^*}{\sqrt{m_i^\downarrow}} : V_i^S \longrightarrow \mathbb{R}^n \otimes V_{i+1}^S$$

are isometries, and their contractions onto the tangent-harmonic copies satisfy

$$(C_i^\uparrow)^*(x \otimes \varphi_{i,x} Y) = \frac{\alpha_i^{(k)}}{\sqrt{m_i^\uparrow}} \varphi_{i+1,x} Y,$$

$$(C_i^\downarrow)^*(x \otimes \varphi_{i+1,x} Y) = \frac{\alpha_i^{(k)}}{\sqrt{m_i^\downarrow}} \varphi_{i,x} Y.$$

The resulting symmetric edge weight is

$$J_{i,i+1} = \frac{(\alpha_i^{(k)})^2}{\alpha_i^{(0)}}. \quad (68)$$

Proof. The addition formula (60) identifies the reproducing kernel $K_j = K_j^S$ of V_j^S as $K_j(u, v) = D_j P_j(\langle u, v \rangle)$. Thus its unit stabilizer-fixed vector at x is $z_{j,x} = K_j(x, \cdot) / \sqrt{D_j}$. Write $t_x(u) = \langle x, u \rangle$. Expanding the Hilbert–Schmidt norm in orthonormal bases and using rotation invariance gives

$$\|M_i^\uparrow\|_{\text{HS}}^2 = \iint_{(\mathbb{S}^{n-1})^2} \langle u, v \rangle K_i(u, v) K_{i+1}(u, v) d\sigma(u) d\sigma(v)$$

$$= \sqrt{D_i D_{i+1}} \langle t_x z_{i,x}, z_{i+1,x} \rangle = \alpha_i^{(0)} \sqrt{D_i D_{i+1}}.$$

The final equality is the recurrence (65) for $k = 0$. For every standard basis vector $e_r \in \mathbb{R}^n$, self-adjointness of coordinate multiplication gives

$$\langle M_i^\uparrow(e_r \otimes f), g \rangle = \langle f, M_i^\downarrow(e_r \otimes g) \rangle.$$

Summing over r shows that the two coordinate maps have the same Hilbert–Schmidt norm. Moreover, $M_i^\uparrow(M_i^\uparrow)^*$ and $M_i^\downarrow(M_i^\downarrow)^*$ commute with rotations. Since their target harmonic spaces are irreducible, they are scalar operators. Taking traces identifies their respective scalar factors as $\alpha_i^{(0)} \sqrt{D_i/D_{i+1}}$ and $\alpha_i^{(0)} \sqrt{D_{i+1}/D_i}$. Substituting (59) and (66) with $k = 0$ yields (67). The stated contractions follow directly from (65); in the downward direction, $\beta_{i+1}^{(k)} = \alpha_i^{(k)}$. Their squared coefficients and dimension-weighted reciprocity satisfy

$$p_{i,i+1} = \frac{(\alpha_i^{(k)})^2}{m_i^\uparrow}, \quad p_{i+1,i} = \frac{(\alpha_i^{(k)})^2}{m_i^\downarrow}, \quad D_i p_{i,i+1} = D_{i+1} p_{i+1,i}. \quad (69)$$

Since $m_i^\uparrow m_i^\downarrow = (\alpha_i^{(0)})^2$, the symmetric edge weight $\sqrt{p_{i,i+1} p_{i+1,i}}$ is (68). $\square$

Fix $k < L$, and retain the unique copy of E_x in each of $V_k^S, \dots, V_L^S$. The resulting transition graph is the one-dimensional path

$$V_k^S \longleftrightarrow V_{k+1}^S \longleftrightarrow \dots \longleftrightarrow V_L^S. \quad (70)$$

The vertices record ambient harmonic degree; they do not record dimensions of subspaces. Substituting (66) into (68) gives the symmetric weight of the edge $V_i^S \leftrightarrow V_{i+1}^S$:

$$c_i^{(k)} = \frac{(i - k + 1)(i + k + n - 2)}{\sqrt{(i + 1)(i + n - 2)(2i + n - 2)(2i + n)}}. \quad (71)$$

Let $J_{k,L}$ be the weighted adjacency matrix of the harmonic-degree path: its diagonal entries vanish and its $(i, i + 1)$ entry is $c_i^{(k)}$. Write $\lambda_{k,L} = \lambda_{\max}(J_{k,L})$. The positive Perron eigenvector

of $J_{k,L}$, as in [Theorem 4.2](#), specifies how to combine the copies of E_x into a single rank- d_k^S projection inside $V_k^S \oplus \cdots \oplus V_L^S$. [Theorem 4.2](#) therefore gives

$$A(n, s) \leq \frac{1-s}{d_k^S(\lambda_{k,L} - s)} \sum_{i=k}^L D_i^S \quad \text{if } \lambda_{k,L} > s.$$

Here $\sum_{i=k}^L D_i^S$ is the dimension of the ambient direct sum, while d_k^S , given in (61), is the rank of the moving projection. For $k=0$, this rank is 1, recovering the classical fixed-line construction. When k grows proportionally to n , d_k^S grows exponentially and its logarithm is subtracted from the code exponent. This gain over the classical fixed-line construction comes from a scalar two-point Delsarte certificate, even though the moving subspaces have dimension greater than one.

The next lemma identifies the spectral limit of this one-dimensional harmonic-degree path, including the classical Kabatianskii–Levenshtein case $k=0$.

Lemma 5.2. *Fix $0 \leq b < a$, and suppose $k/n \rightarrow b$ and $L/n \rightarrow a$. Then*

$$\lambda_{k,L} \rightarrow 2\Gamma_{\text{row}}(a, b) = \frac{2(a-b)(1+a+b)}{(1+2a)\sqrt{a(1+a)}}.$$

Proof. Uniformly for i/n in a compact subinterval of $(0, \infty)$, the edge weight (71) converges to

$$\Gamma_{\text{row}}(u, b) = \frac{u(u+1) - b(b+1)}{(1+2u)\sqrt{u(u+1)}}.$$

This function increases on $u \geq b$: its first summand $\sqrt{u(u+1)}/(1+2u)$ increases, while $b(b+1)/((1+2u)\sqrt{u(u+1)})$ decreases. If $b > 0$, this convergence is uniform over the entire harmonic-degree path, so its largest edge weight converges to $\Gamma_{\text{row}}(a, b)$. If $b = 0$, the same conclusion follows by first restricting to $i \geq \varepsilon n$. The omitted edges satisfy

$$0 \leq c_i^{(k)} \leq c_i^{(0)} = \sqrt{\frac{(i+1)(i+n-2)}{(2i+n-2)(2i+n)}} \leq \sqrt{\frac{i+1}{n-2}},$$

because $(i-k+1)(i+k+n-2) = (i+1)(i+n-2) - k(k+n-3)$. Letting $\varepsilon \downarrow 0$ proves the claim at $b=0$ as well. The maximum row sum therefore gives $\limsup \lambda_{k,L} \leq 2\Gamma_{\text{row}}(a, b)$. For $m_n \rightarrow \infty$ with $m_n = o(n)$, the edge weights on the last m_n vertices converge uniformly to $\Gamma_{\text{row}}(a, b)$. Applying (26) to the first discrete sine vector on that terminal path gives the matching lower bound. $\square$

When $b=0$, the limiting eigenvalue is $2q(a)$, and the classical feasibility boundary $2q(a) = s$ gives $a = a_0(s)$ from the introduction and the direct whole-sphere exponent $H_{\text{sph}}(a_0(s))$. Applying the additional spherical-cap optimization gives instead the classical exponent $B_{\text{KL}}(s)$ in (10). Positive b gives the tangent-harmonic spectral quantity in (11).

The harmonic-degree path (70) uses only ambient spaces of scalar spherical harmonics. Scalar multiplication by the base-point coordinate connects just the neighboring degrees $i-1$ and $i+1$, as in (65). The moving-subspace construction also allows other ambient representations, however: its rotation-equivariant coordinate maps take values in $\mathbb{R}^n \otimes V_i^S$, followed by contraction with the base-point coordinate as in (52). This vector-valued tensor product contains one more irreducible representation besides the two neighboring harmonic spaces.

For integers $i \geq j \geq 0$, let $V_{(i,j)}$ denote the irreducible orthogonal-group tensor representation whose Young diagram has two rows of lengths i and j ; its precise highest-weight description is given in §5.3. A zero second row gives $V_{(i,0)} = V_i^S$. For $i \geq 1$ and $n \geq 6$, tensoring with the standard representation $\mathbb{R}^n$ gives [Kra18, (B.2)]

$$\mathbb{R}^n \otimes V_{(i,0)} \simeq V_{(i+1,0)} \oplus V_{(i-1,0)} \oplus V_{(i,1)}.$$

The first two summands correspond to the neighboring harmonic degrees already present in (65). The third, $V_{(i,1)}$, has row lengths i and 1; it is an ambient tensor representation rather than a space of scalar functions on the sphere. For $i \geq k \geq 1$, the restriction of $V_{(i,1)}$ still contains the

stabilizer representation E_x . Allowing all the additional spaces $V_{(i,j)}$ containing E_x gives the vertex set

$$\{(i, j) : i \geq k \geq j \geq 0\}.$$

The graph on vertices (i, j) now has edges changing either i or j ; the harmonic-degree path is the boundary $j = 0$. The positive Perron eigenvector of the two-dimensional weighted adjacency matrix again assembles a moving projection of rank d_k^S . Since $SO(n)$ is transitive on ordered pairs with a prescribed inner product, the overlap of two such projections depends only on $\langle x, y \rangle$. Consequently, the enlarged graph still produces a scalar two-point Delsarte certificate.

5.3. The multi-row hierarchy. The degree- k tangent-harmonic subspace consists of symmetric traceless k -tensors on $x^\perp$. This subspace supports both the one-dimensional harmonic-degree path and its enlargement to the two-dimensional graph on (i, j) . Higher-dimensional transition graphs arise by allowing more general orthogonal-group tensor representations.

Fix a hierarchy level $r \geq 0$ and assume $n \geq 2r + 4$. Since r remains fixed as $n \rightarrow \infty$, this dimension restriction is automatic. The irreducible orthogonal-group tensor representations used here correspond to Young diagrams: arrays of boxes whose row lengths form a weakly decreasing sequence of nonnegative integers. This sequence, completed by zero entries as necessary, is the dominant highest weight of the representation. For example, the one-row diagram (k) corresponds to symmetric traceless k -tensors; additional rows describe other tensor symmetries.

We use ambient $SO(n)$ -representations corresponding to Young diagrams with at most $r + 1$ rows and stabilizer $SO(n - 1)$ -representations corresponding to diagrams with at most r rows. Under $n \geq 2r + 4$, their remaining highest-weight coordinates vanish, so none reaches the final available orthogonal-group weight coordinate. This avoids the exceptional splitting that can occur outside the stable range. The hierarchy level counts these rows, not the dimension of a representation or the number of graph vertices.

Fix an irreducible $SO(n - 1)$ -representation E_μ with highest weight

$$\mu = (\mu_1, \dots, \mu_r, 0, \dots), \quad \mu_1 \geq \dots \geq \mu_r \geq 0.$$

Thus r bounds the number of nonzero stabilizer rows. An ambient representation V_λ of $SO(n)$ has highest weight

$$\lambda = (\lambda_1, \dots, \lambda_{r+1}, 0, \dots), \quad \lambda_1 \geq \dots \geq \lambda_{r+1} \geq 0.$$

The orthogonal-group branching rule [GW09, Thms. 8.1.3–8.1.4] says that restricting V_λ from $SO(n)$ to $SO(n - 1)$ contains E_μ exactly when every stabilizer row length lies between the two adjacent ambient row lengths:

$$\lambda_1 \geq \mu_1 \geq \lambda_2 \geq \dots \geq \mu_r \geq \lambda_{r+1} \geq 0. \tag{72}$$

The harmonic decomposition (62) is its specialization to one-row ambient representations. Every such copy occurs with multiplicity one. Thus, after choosing μ , the interlacing inequalities identify the ambient representations that contain the selected stabilizer type. These ambient weights λ will be the vertices of the transition graph. The stabilizer weight μ remains fixed throughout that graph: only λ varies from vertex to vertex. When taking the high-dimensional limit, the initial choice of μ may depend on n , but it is still fixed within each individual graph.

The standard $SO(n)$ representation $\mathbb{R}^n$ has highest weight $(1, 0, \dots, 0)$. Tensoring V_λ with $\mathbb{R}^n$ adds or removes one box from its Young diagram, and hence increases or decreases a single row length by one [Kra18, (B.2)]. In odd dimensions, a potential unchanged-weight term is absent under the zero-tail assumption, as explained in §A.2. The transition graph therefore joins λ to $\lambda \pm e_\ell$, where e_ℓ is the ℓ th coordinate vector, provided the resulting row lengths remain weakly decreasing and satisfy (72) for the same fixed μ . Thus the graph is the nearest-neighbor lattice graph on the ambient weights inside the interlacing region. When $r = 0$, it is the classical one-dimensional path of spherical harmonic degrees. When $\mu = (k)$, it is the two-dimensional region $i \geq k \geq j \geq 0$. Generally, r stabilizer rows permit $r + 1$ ambient coordinates and hence an $(r + 1)$ -dimensional transition graph.

The representation-theoretic inputs are the orthogonal-group branching condition (72) and Weyl's dimension formula (111) [GW09, §7.1.2 and §§8.1.1–8.1.2], together with the explicit

coordinate-multiplication coefficients in Kravchuk's conventions [Kra18, §§2.1–2.3 and App. B]. Kravchuk's $\text{Spin}(n)$ formulas apply unchanged to integral tensor representations of $SO(n)$. §A.2, specifically (109) and (110), reduces the odd- and even-dimensional formulas to the common transition weights used below. §A.3 records the dimensions of V_λ and E_μ in (111) and their uniform exponential rates in Lemma A.1.

6. THE SPHERICAL TRANSITION GRAPH

§5.3 described which ambient representations V_λ contain a chosen stabilizer representation E_μ and which pairs of ambient representations are connected by coordinate multiplication. To obtain a spherical-code bound, we must determine the strength of each coordinate transition. We first give the directed squared coefficients, then combine opposite directions into a symmetric weighted adjacency matrix J_Ω . The largest eigenvalue of J_Ω and the dimensions of V_λ and E_μ determine the finite-dimensional projection bound in Theorem 4.2.

Fix $r \geq 0$ and $n \geq 2r + 4$. Choose a stabilizer representation E_μ with highest weight $\mu = (\mu_1, \dots, \mu_r, 0, \dots)$. The highest weights $\lambda = (\lambda_1, \dots, \lambda_{r+1}, 0, \dots)$ of the ambient representations containing E_μ satisfy (72), and each ambient representation contains exactly one copy of E_μ .

Coordinate multiplication can increase or decrease one row length λ_ℓ . The resulting coefficients take a uniform form in odd and even dimensions after adding the standard orthogonal-group offsets to the row lengths. For $1 \leq \ell \leq r + 1$ and $1 \leq m \leq r$, set

$$\hat{\lambda}_\ell = \lambda_\ell + \frac{n}{2} - \ell, \quad \hat{\mu}_m = \mu_m + \frac{n-1}{2} - m, \quad \rho_{n,r} = \frac{n}{2} - r - 1. \quad (73)$$

Write e_ℓ for the ℓ th coordinate vector. The scalar $c_{\ell,\pm}(\lambda; \mu)$ describes multiplication by the base-point coordinate between the copies of E_μ in V_λ and $V_{\lambda \pm e_\ell}$, in the sense of (52). Define its squared magnitude by $p_{\ell,\pm}(\lambda; \mu) = |c_{\ell,\pm}(\lambda; \mu)|^2$. If the target row lengths are not weakly decreasing or fail (72), there is no target representation containing E_μ , and its transition coefficient is zero. On the full graph of ambient weights containing the fixed stabilizer representation E_μ , these squared coefficients sum to one at each vertex, as in (75). They may therefore be viewed as transition probabilities on that full graph, but this probabilistic interpretation is not needed for the code bound. Indeed, the construction retains a finite vertex set Ω and omits every edge leaving Ω , so the remaining coefficients generally sum to less than one. What the construction uses is the dimension-weighted reciprocity (76), which extends (69) and allows the finite directed transition matrix to be symmetrized. The positive Perron eigenvector of the resulting symmetric weighted adjacency matrix is then reweighted as in (57) to produce the moving projections.

Proposition 6.1. *Fix $r \geq 0$, $n \geq 2r + 4$, and ambient and stabilizer highest weights λ, μ satisfying (72). If $1 \leq \ell \leq r + 1$ and $\lambda \pm e_\ell$ is a dominant weight still satisfying (72), its directed squared coordinate coefficient is*

$$p_{\ell,\pm}(\lambda; \mu) = \frac{(\hat{\lambda}_\ell \pm \rho_{n,r}) \prod_{m=1}^r ((\hat{\lambda}_\ell \pm \frac{1}{2})^2 - \hat{\mu}_m^2)}{2\hat{\lambda}_\ell \prod_{\substack{q=1 \\ q \neq \ell}}^{r+1} (\hat{\lambda}_\ell^2 - \hat{\lambda}_q^2)}. \quad (74)$$

With the zero convention above, the full transition graph satisfies

$$\sum_{\ell=1}^{r+1} (p_{\ell,+}(\lambda; \mu) + p_{\ell,-}(\lambda; \mu)) = 1, \quad (75)$$

and, for every graph edge $\lambda \leftrightarrow \lambda + e_\ell$,

$$\dim V_\lambda p_{\ell,+}(\lambda; \mu) = \dim V_{\lambda+e_\ell} p_{\ell,-}(\lambda + e_\ell; \mu). \quad (76)$$

Proof. The coefficients come from coordinate multiplication between the multiplicity-free stabilizer copies. Kravchuk gives separate odd- and even-dimensional expressions for these coefficients [Kra18, (B.12)–(B.13)]. §A.2 identifies their normalization and records the corresponding

squared coefficients in (109) and (110). In both expressions, factors belonging to zero highest-weight coordinates cancel, leaving the common formula (74). The same appendix also explains why the apparent odd-dimensional self-loop is absent.

For $Y \in E_\mu$, the base-point coordinate ℓ_x is fixed by the stabilizer. Therefore $\ell_x \otimes \varphi_{\lambda,x} Y$ has stabilizer type E_μ , and its orthogonal projections onto the ambient irreducible summands have squared norms $p_{\ell,\pm}(\lambda; \mu) \|Y\|^2$. Orthogonality of these summands and $\|\ell_x\| = 1$ give

$$\|Y\|^2 = \|\ell_x \otimes \varphi_{\lambda,x} Y\|^2 = \sum_{\ell=1}^{r+1} (p_{\ell,+}(\lambda; \mu) + p_{\ell,-}(\lambda; \mu)) \|Y\|^2,$$

proving (75) on the full graph. Finally, applying the exact Weyl dimension formula (111) to the two endpoints of an existing edge gives (76); the calculation appears at the end of §A.2. $\square$

Now choose a finite connected set Ω of ambient weights satisfying (72), and retain only edges whose two endpoints belong to Ω . Although the two directed weights on an edge generally differ, (76) says that each agrees after multiplication by the dimension of its source representation. Thus (55) converts the directed coefficient matrix into a symmetric weighted adjacency matrix J_Ω . The edge weight of J_Ω is the geometric mean of the forward and reverse squared coefficients:

$$J_{\lambda, \lambda + e_\ell} = \sqrt{p_{\ell,+}(\lambda; \mu) p_{\ell,-}(\lambda + e_\ell; \mu)}, \quad (77)$$

with every other entry zero. Write $d_\mu = \dim E_\mu$ for the rank of the moving projection, $D_\lambda = \dim V_\lambda$ for each ambient representation dimension, and $\Lambda_\Omega = \lambda_{\max}(J_\Omega)$ for the largest eigenvalue of the finite weighted adjacency matrix. Theorem 4.2, applied with $G = SO(n)$, $H = SO(n-1)$, and $W = \mathbb{R}^n$, produces a scalar positive-definite kernel depending only on the inner product and gives the following bound.

Theorem 6.2. *Fix $-1 < s < 1$, $r \geq 0$, $n \geq 2r + 4$, a stabilizer representation E_μ , and a finite connected set Ω of ambient weights satisfying (72). If $\Lambda_\Omega > \max\{s, 0\}$, then*

$$A(n, s) \leq \frac{1-s}{d_\mu(\Lambda_\Omega - s)} \sum_{\lambda \in \Omega} D_\lambda. \quad (78)$$

6.1. The two-row graph. To make the general graph concrete, return to the degree- k tangent harmonics $E_\mu = \mathcal{H}_k(x^\perp)$, whose stabilizer weight is the single row $\mu = (k)$. An ambient representation may now have two rows $\lambda = (i, j)$. By (72), it contains E_μ exactly when $i \geq k \geq j \geq 0$. Thus the vertices form a two-dimensional region: the one-dimensional harmonic-degree path is the boundary $j = 0$, and $j > 0$ supplies the ambient representations omitted from that boundary. The resulting lattice strip is shown in Figure 2.

The two possible coordinate directions join (i, j) to $(i \pm 1, j)$ or to $(i, j \pm 1)$, respectively, and (74) gives the four directed squared coefficients

$$\begin{aligned} p_{i,+} &= \frac{(i-k+1)(i+k+n-2)(i+n-3)}{(i-j+1)(i+j+n-3)(2i+n-2)}, & p_{i,-} &= \frac{(i-k)(i+k+n-3)(i+1)}{(i-j+1)(i+j+n-3)(2i+n-2)}, \\ p_{j,+} &= \frac{(k-j)(j+k+n-3)(j+n-4)}{(i-j+1)(i+j+n-3)(2j+n-4)}, & p_{j,-} &= \frac{j(k-j+1)(j+k+n-4)}{(i-j+1)(i+j+n-3)(2j+n-4)}. \end{aligned} \quad (79)$$

The factors $i-k$, $k-j$, and j in (79) make the corresponding transitions vanish at the boundaries $i = k$, $j = k$, and $j = 0$, respectively. At $j = 0$, $p_{i,+}$ and $p_{i,-}$ describe the one-dimensional harmonic-degree path, whereas

$$p_{j,+}(i, 0) = \frac{k(k+n-3)}{(i+1)(i+n-3)}$$

is positive precisely when $k > 0$. This edge from $(i, 0)$ to $(i, 1)$ is absent from the classical fixed-line construction. Retaining vertices with $j > 0$ incorporates such additional edges into $\lambda_{\max}(J_\Omega)$ while leaving the stabilizer dimension d_k^S unchanged.




FIGURE 2. The spherical two-row representation graph for a fixed stabilizer weight $\mu = (k)$. Ambient vertices (i, j) satisfy $i \geq k \geq j \geq 0$; horizontal and vertical edges change i and j , respectively. Their directed squared weights are given in (79), and (77) converts opposite directions into symmetric edge weights. The blue boundary is the one-row subgraph. The dashed box marks a lattice region near an upper corner. For any large fixed box size, its edge weights become asymptotically constant as $n \rightarrow \infty$, and a product of discrete sine waves bounds the largest eigenvalue from below. Letting the box size increase gives Lemma 7.2.

The ambient dimension needed in (78) is obtained by specializing the exact Weyl formula (111) in §A.3 to $\lambda = (i, j)$:

$$D_{i,j} = \frac{(2i+n-2)(2j+n-4)(i-j+1)(i+j+n-3)}{(n-2)(n-4)(i+1)(i+n-3)} \binom{i+n-3}{i} \binom{j+n-5}{j}.$$

The identity $D_{i,j} p_{i,+}(i, j) = D_{i+1,j} p_{i,-}(i+1, j)$ and its j -direction analogue give (76); restricting to $j = 0$ recovers (71).

7. ASYMPTOTIC SPHERICAL-CODE BOUNDS

We now evaluate the finite spherical-code bound (78) as the dimension tends to infinity. The spectral constraint involves $\Lambda_\Omega = \lambda_{\max}(J_\Omega)$, while the exponent in the bound is determined by

$$\frac{\sum_{\lambda \in \Omega} \dim V_\lambda}{\dim E_\mu}.$$

We keep the number r of stabilizer rows fixed, let the row lengths of E_μ and V_λ grow proportionally to n , and compute both the spectral constraint and the dimension-ratio exponent from the limiting coordinates.

Fix $r \geq 0$ before taking $n \rightarrow \infty$. Suppose the ambient and stabilizer row lengths satisfy $\lambda_\ell/n \rightarrow a_\ell$ and $\mu_m/n \rightarrow b_m$, respectively. We require strict separation between successive ambient and stabilizer rows, while allowing the final ambient row to vanish:

$$a_1 > b_1 > a_2 > \cdots > b_r > a_{r+1} \geq 0. \quad (80)$$

For $r = 0$, this condition means simply $a_1 \geq 0$. Two scalar functions describe the limiting edge weights of J_Ω . The quadratic change of variables $\mathcal{A}(u) = u(1+u)$ converts shifted highest-weight differences into differences of nonnegative real numbers. The function $q(u)$ is the limiting symmetric edge weight in the classical level-zero harmonic path at degree $un + o(n)$. Put

$$\mathcal{A}(u) = u(1+u), \quad q(u) = \frac{\sqrt{u(1+u)}}{1+2u}.$$

Write $x_\ell = \mathcal{A}(a_\ell)$ and $y_m = \mathcal{A}(b_m)$. Since $\mathcal{A}$ is strictly increasing, the interlacing inequalities (80) give $x_1 > y_1 > x_2 > \cdots > y_r > x_{r+1} \geq 0$. The residues R_ℓ associated with the interlacing

nodes x_ℓ, y_m , the resulting weighted edge strength Γ_r , and the ambient-to-stabilizer dimension exponent Φ_r are

$$\begin{aligned} R_\ell(\mathbf{a}, \mathbf{b}) &= \frac{\prod_{m=1}^r (x_\ell - y_m)}{\prod_{\substack{q=1 \\ q \neq \ell}}^{r+1} (x_\ell - x_q)}, & \Gamma_r(\mathbf{a}, \mathbf{b}) &= \sum_{\ell=1}^{r+1} R_\ell(\mathbf{a}, \mathbf{b}) q(a_\ell), \\ \Phi_r(\mathbf{a}, \mathbf{b}) &= \sum_{\ell=1}^{r+1} \mathsf{H}_{\text{sph}}(a_\ell) - \sum_{m=1}^r \mathsf{H}_{\text{sph}}(b_m). \end{aligned} \tag{81}$$

The R_ℓ are the residues of $\prod_m (z - y_m) / \prod_\ell (z - x_\ell)$. The strict interlacing in (80) makes every residue positive, and comparison at $z = \infty$ gives $\sum_{\ell=1}^{r+1} R_\ell = 1$. The next subsection constructs finite weighted adjacency matrices J_{Ω_n} satisfying $\lambda_{\max}(J_{\Omega_n}) \geq 2\Gamma_r + o(1)$. The same subsection shows that the normalized logarithm of the dimension ratio in (78) converges to Φ_r .

7.1. Spectral and dimension asymptotics. The spectral estimate is local, even though transition weights vary across the full ambient-weight graph. Keep the stabilizer weight μ fixed and consider a lattice box of side m_n near an ambient weight λ_n^* , where $m_n \rightarrow \infty$, $m_n = o(n)$, and $\lambda_n^*/n \rightarrow \mathbf{a}$. Every vertex of this box has the same limiting normalized ambient weight $\mathbf{a}$. Consequently, the edge weights freeze to constants depending only on their coordinate direction:

$$w_\ell = R_\ell(\mathbf{a}, \mathbf{b}) q(a_\ell), \quad (J_{\text{loc}} f)(v) = \sum_\ell w_\ell (f(v + e_\ell) + f(v - e_\ell)).$$

Here f is zero outside the box. Thus $2 \sum_\ell w_\ell I - J_{\text{loc}}$ is the usual weighted Dirichlet lattice Laplacian, whose first eigenvector is a product of discrete sine vectors.

We first calculate the limiting edge weights in Lemma 7.1; the local sine test function then yields the eigenvalue lower bound in Lemma 7.2. Finally, Lemma A.1 in §A.3 determines the exponential ratio of ambient to stabilizer dimensions.

Lemma 7.1. *Fix $r \geq 0$. Suppose (80) holds, $\lambda_\ell/n \rightarrow a_\ell$, and $\mu_m/n \rightarrow b_m$. For every existing coordinate transition,*

$$p_{\ell, \pm}(\lambda; \mu) \longrightarrow \frac{R_\ell(2a_\ell + 1 \pm 1)}{2(1 + 2a_\ell)}. \tag{82}$$

Consequently,

$$J_{\lambda, \lambda + e_\ell} \longrightarrow R_\ell q(a_\ell). \tag{83}$$

Both limits hold uniformly whenever

$$\max_\ell \left| \frac{\lambda_\ell}{n} - a_\ell \right| + \max_{1 \leq m \leq r} \left| \frac{\mu_m}{n} - b_m \right| = o(1).$$

The second maximum is understood as zero when $r = 0$.

Proof. Substitute (73) into (74) and use $(u + \frac{1}{2})^2 - (v + \frac{1}{2})^2 = \mathcal{A}(u) - \mathcal{A}(v)$. The quadratic and linear factors give (82); the balanced geometric mean of the forward and reverse coefficients gives (83). The strict interlacing in (80) keeps all limiting denominator differences nonzero, so the rational formulas converge uniformly on shrinking neighborhoods of the prescribed row lengths. $\square$

Choose integral stabilizer row lengths $\mu_m = b_m n + o(n)$ and upper ambient row lengths $N_{\ell, n} = a_\ell n + o(n)$, taking $N_{r+1, n} = 0$ if $a_{r+1} = 0$. Let Ω_n consist of the ambient weights satisfying the interlacing inequalities (72) and

$$\mu_\ell \leq \lambda_\ell \leq N_{\ell, n} \quad (1 \leq \ell \leq r), \quad 0 \leq \lambda_{r+1} \leq N_{r+1, n}.$$

Its upper corner is $\lambda_n^* = (N_{1, n}, \dots, N_{r+1, n})$.

Lemma 7.2. *The weighted adjacency matrix of the finite ambient-weight set Ω_n satisfies*

$$\liminf_{n \rightarrow \infty} \lambda_{\max}(J_{\Omega_n}) \geq 2\Gamma_r(\mathbf{a}, \mathbf{b}). \tag{84}$$

Proof. Choose $m_n \rightarrow \infty$ with $m_n = o(n)$. Strict interlacing allows every positive coordinate of λ_n^* to decrease independently by $0, 1, \dots, m_n - 1$ while remaining inside Ω_n ; a zero last coordinate stays fixed. The resulting vertices form a rectangular box with m_n vertices in each active coordinate direction. Since each side length is $o(n)$, [Lemma 7.1](#) shows that all edges in the ℓ th coordinate direction have weight $R_\ell q(a_\ell) + o(1)$, uniformly throughout the box.

On a one-dimensional path with m_n vertices and constant edge weight w , the first sine vector has Rayleigh quotient $2w \cos(\pi/(m_n + 1))$. Taking the product of these sine vectors over the coordinate directions therefore gives

$$2 \sum_{\ell=1}^{r+1} R_\ell q(a_\ell) \cos \frac{\pi}{m_n + 1} + o(1).$$

The Rayleigh–Ritz principle (26) now proves (84). If $a_{r+1} = 0$, the rectangular box has no direction corresponding to the last ambient coordinate, but the omitted contribution $R_{r+1}q(0)$ also vanishes. Hence the same lower bound holds when $a_{r+1} = 0$. $\square$

The dimension calculation is separate from the spectral estimate. At the upper corner $\lambda = \lambda_n^*$, [Lemma A.1](#) gives the exponential growth rates of the ambient representation and the stabilizer representation:

$$\frac{1}{n} \log_2 \dim V_\lambda = \sum_{\ell=1}^{r+1} H_{\text{sph}}(a_\ell) + o(1), \quad \frac{1}{n} \log_2 \dim E_\mu = \sum_{m=1}^r H_{\text{sph}}(b_m) + o(1). \quad (85)$$

Fix $A > \max_\ell a_\ell$. For all sufficiently large n , every $\lambda \in \Omega_n$ satisfies $\lambda_\ell \leq N_{\ell,n} \leq An$. Since H_{sph} is increasing, [Lemma A.1](#) therefore gives uniformly on Ω_n

$$\frac{1}{n} \log_2 \dim V_\lambda = \sum_{\ell=1}^{r+1} H_{\text{sph}}\left(\frac{\lambda_\ell}{n}\right) + O_{r,A}\left(\frac{\log(n+2)}{n}\right) \leq \sum_{\ell=1}^{r+1} H_{\text{sph}}\left(\frac{N_{\ell,n}}{n}\right) + O_{r,A}\left(\frac{\log(n+2)}{n}\right).$$

Thus the upper corner λ_n^* maximizes the ambient dimension exponent, with uniform error $O_{r,A}(\log(n+2)/n)$. There are only $O(n^{r+1})$ ambient weights in Ω_n , so $\sum_{\lambda \in \Omega_n} \dim V_\lambda$ differs from the largest ambient dimension by a subexponential factor:

$$\frac{1}{n} \log_2 \sum_{\lambda \in \Omega_n} \dim V_\lambda = \sum_{\ell=1}^{r+1} H_{\text{sph}}(a_\ell) + o(1). \quad (86)$$

Proof of Theorem 1.2. First suppose

$$2\Gamma_r(\mathbf{a}, \mathbf{b}) > s,$$

so [Lemma 7.2](#) implies that $\Lambda_{\Omega_n} - s$ stays uniformly positive. Hence the prefactor in (78) is bounded, while (85) and (86) give dimension-ratio exponent Φ_r . Consequently,

$$A(n, s) \leq 2^{(\Phi_r(\mathbf{a}, \mathbf{b}) + o(1))n}.$$

The estimate extends to the feasibility boundary $2\Gamma_r = s > 0$. For $c > 0$, define the scaling map $S_c(u) = \mathcal{A}^{-1}(c\mathcal{A}(u))$. Replacing each a_ℓ, b_m by $S_c(a_\ell), S_c(b_m)$ multiplies all quadratic coordinates by c , preserving interlacing and the residues R_ℓ . If $c > 1$, strict monotonicity of q increases Γ_r . Taking $c \downarrow 1$ therefore approximates each tuple satisfying $2\Gamma_r = s$ by strictly interlacing tuples with $2\Gamma_r > s$ whose dimension exponents converge to Φ_r . Taking the infimum over the tuples satisfying (80) and $2\Gamma_r \geq s$, and then over all finite hierarchy levels, proves the second assertion of [Theorem 1.2](#). $\square$

Define

$$\kappa_r(s) = \inf_{\substack{a_1 > b_1 > \dots > b_r > a_{r+1} \geq 0 \\ 2\Gamma_r(\mathbf{a}, \mathbf{b}) \geq s}} \Phi_r(\mathbf{a}, \mathbf{b}), \quad \kappa_\infty(s) = \inf_{r \geq 0} \kappa_r(s).$$

The quantity $\kappa_r(s)$ is the best exponent obtained directly on the whole sphere at hierarchy level r . By [Theorem 1.2](#), approximating tuples with $2\Gamma_r = s$ by tuples with $2\Gamma_r > s$ bounds the whole-sphere code exponent by $\kappa_r(s)$. For any error tolerance, choose a finite level approximating

$\kappa_\infty(s)$ and keep that level fixed as $n \rightarrow \infty$. Thus

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log_2 A(n, s) \leq \kappa_\infty(s).$$

In particular, the hierarchy depth is fixed before the dimension limit. We can further improve a fixed-angle bound by the spherical-cap reduction recalled in (10): restrict to a sufficiently populated spherical cap and project to a lower-dimensional sphere with inner-product threshold $0 \leq t \leq s$. The spherical-slice inequality of Sidelnikov and Kabatianskii–Levenshtein [Sid74, KL78], recalled in (10), charges the exponential cost $\frac{1}{2} \log_2((1-t)/(1-s))$ for restricting to the cap and projecting to the smaller sphere. Applying this inequality at each fixed hierarchy level gives

$$\bar{\kappa}_r(s) = \inf_{0 \leq t \leq s} \left\{ \kappa_r(t) + \frac{1}{2} \log_2 \frac{1-t}{1-s} \right\}, \quad \bar{\kappa}_\infty(s) = \inf_{r \geq 0} \bar{\kappa}_r(s), \quad (87)$$

where $\kappa_r(0) = 0$. Consequently,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log_2 A(n, s) \leq \bar{\kappa}_\infty(s).$$

Thus κ_r is the exponent obtained directly on the whole sphere, and $\bar{\kappa}_r$ includes the additional optimization over spherical caps.

7.2. Strict improvement at every hierarchy level. The direct and spherical-cap-optimized exponents $\kappa_r(s)$ and $\bar{\kappa}_r(s)$ are nonincreasing in r , since a level- r tuple is recovered as a limit of level- $(r+1)$ tuples. We prove that both decrease strictly at every $0 < s < 1$. The first comparisons isolate the intermediate harmonic-degree path (70): replacing the classical fixed line by the moving stabilizer space $\mathcal{H}_k(x^\perp)$ gives one strict improvement, and allowing a second ambient highest-weight coordinate gives another. Extending this argument to all hierarchy levels requires controlling minimizing sequences whose coordinates escape to infinity.

For $r = 0$, $\Gamma_0(a) = q(a)$ and $\Phi_0(a) = \mathbf{H}_{\text{sph}}(a)$; the spectral boundary $2q(a_0(s)) = s$ recovers the direct whole-sphere Kabatianskii–Levenshtein exponent $\mathbf{H}_{\text{sph}}(a_0(s))$. The intermediate one-row construction is the restriction $a_2 = 0$ of level 1. By (11), its spectral quantity is $\Gamma_1((a, 0), b) = \Gamma_{\text{row}}(a, b)$, and its dimension exponent is $\mathbf{H}_{\text{sph}}(a) - \mathbf{H}_{\text{sph}}(b)$. The optimized one-row rate is

$$\kappa_{\text{row}}(s) = \inf_{\substack{a > b > 0 \\ 2\Gamma_{\text{row}}(a, b) \geq s}} (\mathbf{H}_{\text{sph}}(a) - \mathbf{H}_{\text{sph}}(b)), \quad (88)$$

and its cap-optimized exponent $\bar{\kappa}_{\text{row}}$ is obtained by replacing κ_r with κ_{row} in (87).

Two perturbations drive the strict improvements. Opening a zero terminal ambient coordinate increases the spectral quantity and decreases the dimension exponent at the same level; once that coordinate is positive, appending a stabilizer row gives the same improvements one level higher.

Proposition 7.3. *Fix $0 < s < 1$ and $r \geq 0$, and let $\mathbf{a} = (a_1, \dots, a_{r+1})$ and $\mathbf{b} = (b_1, \dots, b_r)$ satisfy*

$$a_1 > b_1 > a_2 > \dots > b_r > a_{r+1} \geq 0, \quad 2\Gamma_r(\mathbf{a}, \mathbf{b}) \geq s.$$

For $r = 0$, the interlacing condition reduces to $a_1 \geq 0$.

- (1) *If $r \geq 1$ and $a_{r+1} = 0$, there are strictly interlacing vectors $\mathbf{a}' \in \mathbb{R}^{r+1}$ and $\mathbf{b}' \in \mathbb{R}^r$ with $a'_{r+1} > 0$ such that*

$$2\Gamma_r(\mathbf{a}', \mathbf{b}') > 2\Gamma_r(\mathbf{a}, \mathbf{b}) \geq s, \quad \Phi_r(\mathbf{a}', \mathbf{b}') < \Phi_r(\mathbf{a}, \mathbf{b}).$$

- (2) *If $a_{r+1} > 0$, there are strictly interlacing vectors $\mathbf{a}' \in \mathbb{R}^{r+2}$ and $\mathbf{b}' \in \mathbb{R}^{r+1}$ with $a'_{r+2} = 0 < b'_{r+1} < a'_{r+1}$ such that*

$$2\Gamma_{r+1}(\mathbf{a}', \mathbf{b}') > 2\Gamma_r(\mathbf{a}, \mathbf{b}) \geq s, \quad \Phi_{r+1}(\mathbf{a}', \mathbf{b}') < \Phi_r(\mathbf{a}, \mathbf{b}).$$

Proof. Write $\Gamma = \Gamma_r(\mathbf{a}, \mathbf{b})$, $\Phi = \Phi_r(\mathbf{a}, \mathbf{b})$, and work in the quadratic coordinates x_ℓ, y_m . For the first assertion, choose $t > 0$ small and open the zero terminal coordinate by setting

$$\tilde{\mathbf{a}} = (a_1, \dots, a_r, \mathcal{A}^{-1}(t^2)), \quad \Gamma_{\text{open}} = \Gamma_r(\tilde{\mathbf{a}}, \mathbf{b}), \quad \Phi_{\text{open}} = \Phi_r(\tilde{\mathbf{a}}, \mathbf{b}).$$

Expanding (81) gives

$$\Gamma_{\text{open}} = \Gamma + \frac{\prod_m y_m}{\prod_{\ell \leq r} x_\ell} t + O(t^2), \quad \Phi_{\text{open}} = \Phi + O(t^2 \log(1/t)).$$

Replacing the zero terminal ambient coordinate by $\tilde{a}_{r+1} = \mathcal{A}^{-1}(t^2)$ therefore increases the spectral quantity to first order, while its direct contribution to the dimension exponent is smaller. To turn this spare spectral gain into a strict reduction of the exponent, rescale all quadratic coordinates simultaneously using S_c . Its entropy derivative is

$$\left. \frac{d}{d \log c} \mathbf{H}_{\text{sph}}(S_c(u)) \right|_{c=1} = \psi(u) := \frac{u(1+u)}{1+2u} \log_2 \frac{1+u}{u}, \quad \psi(0) := 0.$$

With $v = 1 + 2u$, strict monotonicity follows from

$$4(\log 2) \frac{d\psi}{dv} = (1 + v^{-2}) \log \frac{v+1}{v-1} - \frac{2}{v} > 0.$$

Hence the strict interlacing in (80) gives $\sum_\ell \psi(a_\ell) - \sum_m \psi(b_m) > 0$. For sufficiently small fixed $\eta > 0$, put $c = 1 - \eta t$, $a'_\ell = S_c(\tilde{a}_\ell)$, and $b'_m = S_c(b_m)$. This common contraction preserves a positive spectral gain of order t while decreasing the dimension exponent by order t , proving the first assertion.

For the second assertion, choose $0 < \varepsilon < a_{r+1}$, and append a new stabilizer row and a zero ambient row:

$$\tilde{\mathbf{a}} = (a_1, \dots, a_{r+1}, 0), \quad \tilde{\mathbf{b}} = (b_1, \dots, b_r, \varepsilon), \quad e = \mathcal{A}(\varepsilon).$$

Each existing interpolation weight becomes $R_\ell^{\text{new}} = R_\ell(1 - e/x_\ell)$, while the new zero node contributes nothing to the spectral quantity. Write $\Gamma_{\text{app}} = \Gamma_{r+1}(\tilde{\mathbf{a}}, \tilde{\mathbf{b}})$ and $\Phi_{\text{app}} = \Phi_{r+1}(\tilde{\mathbf{a}}, \tilde{\mathbf{b}})$. Then

$$\Gamma_{\text{app}} = \Gamma - e \sum_\ell \frac{R_\ell q(a_\ell)}{x_\ell}, \quad \Phi_{\text{app}} = \Phi - \mathbf{H}_{\text{sph}}(\varepsilon).$$

Put

$$D_\Gamma = \sum_\ell \frac{R_\ell q(a_\ell)}{2(1 + 4x_\ell)} > 0, \quad L = \sum_\ell \frac{R_\ell q(a_\ell)}{x_\ell}.$$

Under a common expansion of all quadratic coordinates, the logarithmic derivative of the original spectral quantity Γ is D_Γ . For fixed $\theta > 0$, choosing $\log c = (L/D_\Gamma + \theta)e$ and setting $a'_\ell = S_c(\tilde{a}_\ell)$, $b'_m = S_c(b_m)$ more than restores the spectral loss eL , at dimension-exponent cost $O(e)$. Since $e = O(\varepsilon)$ and $\mathbf{H}_{\text{sph}}(\varepsilon) = \varepsilon \log_2(1/\varepsilon) + O(\varepsilon)$, the expanded level- $(r+1)$ tuple has spectral quantity larger than Γ and dimension exponent smaller than Φ . $\square$

At level 1, these perturbations separate the classical construction, its one-row refinement, and the complete level-one construction; the same calculation supplies a boundary estimate for the general argument.

Lemma 7.4. *For every fixed $0 < s < 1$,*

$$\kappa_1(s) < \kappa_{\text{row}}(s) < \kappa_0(s), \quad \bar{\kappa}_1(s) < \bar{\kappa}_{\text{row}}(s) < \bar{\kappa}_0(s) = B_{\text{KL}}(s).$$

Proof. Since increasing b decreases the one-row objective $\mathbf{H}_{\text{sph}}(a) - \mathbf{H}_{\text{sph}}(b)$ in (88), the optimum over b for fixed $a > a_0(s)$ lies on the spectral boundary $b = B_s(a)$, where

$$B_s(a) = \frac{\sqrt{1 + 4a(1+a)} - 2s(1+2a)\sqrt{a(1+a)} - 1}{2}.$$

At $a = a_0(s) + \delta$, $B_s(a) = \frac{1}{2}\sqrt{1 - s^2} \delta + O_s(\delta^2)$. Since $\mathbf{H}_{\text{sph}}(u) = u \log_2(1/u) + O(u)$, the boundary objective is smaller than $\mathbf{H}_{\text{sph}}(a_0(s))$ for sufficiently small $\delta > 0$. Thus $\kappa_{\text{row}}(s) < \kappa_0(s)$.

As $a \rightarrow \infty$, the boundary objective tends to $-\frac{1}{2} \log_2(1-s)$, which is larger than $H_{\text{sph}}(a_0(s))$. Indeed, set $t = s/(1 + \sqrt{1-s^2}) \in (0, 1)$. Since $s = 2t/(1+t^2)$ and $a_0(s) = t^2/(1-t^2)$, the difference, multiplied by $\log 2$, is

$$\log(1+t) + \frac{1}{2} \log(1+t^2) - \frac{2t^2}{1-t^2} \log \frac{1}{t}.$$

The last term is smaller than t , since $\log(1/t) < (t^{-1} - t)/2$. The first two terms are larger than t : their difference from t vanishes at zero and has derivative $t^2(1-t)/((1+t)(1+t^2)) > 0$. Thus the one-variable minimization over a attains its infimum at a finite interior point.

Applying the first part of [Proposition 7.3](#) to the minimizing level-one tuple $((a, 0), (B_s(a)))$ gives $\kappa_1(s) < \kappa_{\text{row}}(s)$.

To justify optimization over spherical caps, put $G(t, a) = H_{\text{sph}}(a) - H_{\text{sph}}(B_t(a))$ for $a \geq a_0(t)$. This function is continuous, including at $a = a_0(t)$, where $G(t, a_0(t)) = \kappa_0(t)$. The explicit formula for $B_t(a)$ gives $B_t(a)/a \rightarrow \sqrt{1-t}$, locally uniformly for $0 < t < 1$. Hence

$$G(t, a) \longrightarrow -\frac{1}{2} \log_2(1-t) \quad (a \rightarrow \infty),$$

with the same local uniformity. The strict improvement over $H_{\text{sph}}(a_0(t))$ and the larger limiting objective as $a \rightarrow \infty$ therefore restrict all minimizing a to a common bounded interval $[a_0(t), A]$ for t in a sufficiently small neighborhood of any fixed $t_0 \in (0, 1)$. Parameterizing that interval by $a = a_0(t) + u(A - a_0(t))$, $0 \leq u \leq 1$, expresses $\kappa_{\text{row}}(t)$ as the minimum of a continuous function on a fixed compact interval. Thus κ_{row} is continuous on $(0, 1)$. Since

$$0 \leq \kappa_{\text{row}}(t) \leq \kappa_0(t) = O(t^2 \log(1/t)) \quad (t \downarrow 0),$$

it extends continuously to $t = 0$ with value zero. The classical and one-row spherical-cap objectives are consequently continuous on the compact interval $[0, s]$ and attain their minima. Neither minimum occurs at $t = 0$, since the slice cost decreases by $t/(2 \log 2) + O(t^2)$, whereas both code exponents increase by at most $O(t^2 \log(1/t))$. Applying $\kappa_{\text{row}}(t) < \kappa_0(t)$ and $\kappa_1(t) < \kappa_{\text{row}}(t)$ at the minimizing thresholds for the classical and one-row spherical-cap objectives gives the two cap-optimized inequalities. $\square$

The first-level comparison above uses an attained one-row optimizer. At higher levels, [Proposition 7.3](#) improves each fixed tuple, but this does not by itself give a strict inequality between optimized exponents: an ambient coordinate and its interlacing stabilizer coordinate can diverge together while their entropy difference remains bounded. The next lemma identifies precisely these noncompact limits.

Lemma 7.5. *Fix $r \geq 0$, and consider a sequence of strictly interlacing level- r tuples with bounded exponents Φ_r . After passing to a subsequence, there are $0 \leq j \leq r$, a strictly interlacing level- j tuple $(\mathbf{a}, \mathbf{b})$, and $0 < c \leq 1$ such that*

$$\Phi_r \longrightarrow \Phi_j(\mathbf{a}, \mathbf{b}) - \log_2 c, \quad 1 - 2\Gamma_r \longrightarrow c^2(1 - 2\Gamma_j(\mathbf{a}, \mathbf{b})).$$

If $c < 1$, then $j < r$. Conversely, this limiting datum is approximable at level $j+1$ when $c < 1$ and at level j when $c = 1$.

Proof. Put $f(x) = H_{\text{sph}}(\mathcal{A}^{-1}(x))$. This increasing function satisfies $f(x) = \frac{1}{2} \log_2 x + \log_2 e + o(1)$ as $x \rightarrow \infty$. In the quadratic coordinates from (81), the exponent is a sum of nonnegative terms:

$$\Phi_r = f(x_{r+1}) + \sum_{i=1}^r (f(x_i) - f(y_i)).$$

Thus x_{r+1} stays bounded, and every divergent x_i has a divergent paired y_i with y_i/x_i bounded away from zero. After extracting a subsequence, the divergent pairs are $i = 1, \dots, k$, and

$$\frac{y_i}{x_i} \longrightarrow \alpha_i \in (0, 1], \quad c^2 = \prod_{i=1}^k \alpha_i.$$

Take limits of the remaining finite coordinates and cancel coincident numerator and denominator factors in the rational Stieltjes transform

$$m(z) = \frac{\prod_{i=1}^r (z - y_i)}{\prod_{i=1}^{r+1} (z - x_i)} = \sum_{i=1}^{r+1} \frac{R_i}{z - x_i}.$$

Weak interlacing leaves a strictly interlacing level- j tuple with $j \leq r - k$, and $m(z) \rightarrow c^2 m_j(z)$ for $z < 0$. Consequently, on the compactified half-line, the residue measures converge to

$$\sum_i R_i \delta_{x_i} \longrightarrow c^2 \sum_i R_i^{(j)} \delta_{x_i^{(j)}} + (1 - c^2) \delta_\infty.$$

Since $q(\mathcal{A}^{-1}(x)) \rightarrow 1/2$ as $x \rightarrow \infty$, integrating this function gives the asserted limit of Γ_r . The asymptotic formula for f gives the corresponding limit of Φ_r . If $c < 1$, at least one pair diverges, so $j < r$.

Conversely, for $c < 1$, prepend the quadratic coordinates $x = M$ and $y = c^2 M$ to the level- j tuple and let $M \rightarrow \infty$. If $c = 1$, no additional pair is required. Further levels can be added without changing the limit by appending a stabilizer coordinate tending to zero and a zero ambient coordinate; a zero terminal coordinate can first be opened by an arbitrarily small amount. At a positive spectral feasibility boundary, slightly decrease c when $c < 1$; when $c = 1$, apply the common scaling S_d with $d \downarrow 1$. Either perturbation makes the approximations strictly feasible without changing their limiting exponent. $\square$

The compactification parameter c records precisely the spherical-cap cost. Indeed, if $t = 2\Gamma_j(\mathbf{a}, \mathbf{b})$ and $s = 1 - c^2(1 - t)$, then $-\log_2 c = \frac{1}{2} \log_2((1 - t)/(1 - s))$, exactly the cost in (87). The local perturbations can therefore be applied to the finite residual tuple to prove strict improvement of the optimized exponents at every level.

Corollary 7.6. *For every $r \geq 0$ and $0 < s < 1$,*

$$\kappa_{r+1}(s) < \kappa_r(s), \quad \bar{\kappa}_{r+1}(s) < \bar{\kappa}_r(s).$$

At level 1, moreover, $\kappa_1(s) < \kappa_{\text{row}}(s) < \kappa_0(s)$ and $\bar{\kappa}_1(s) < \bar{\kappa}_{\text{row}}(s) < \bar{\kappa}_0(s) = B_{\text{KL}}(s)$.

Proof. Appending vanishing rows shows that $0 \leq \kappa_r(t) \leq \kappa_0(t)$. A minimizing sequence for $\kappa_r(s)$ therefore has bounded exponent. Apply Lemma 7.5 to obtain a level- j tuple and $0 < c \leq 1$ with

$$\kappa_r(s) = \Phi_j - \log_2 c, \quad c^2(1 - 2\Gamma_j) \leq 1 - s.$$

If its terminal ambient coordinate vanishes and $j \geq 1$, the first part of Proposition 7.3, applied at the threshold $2\Gamma_j > 0$, decreases Φ_j and increases Γ_j . Retaining c , the resulting datum is approximable at level at most r , contradicting minimality. If $j = 0$ and the ambient coordinate vanishes, then $\Gamma_j = 0$, so feasibility gives

$$\Phi_j - \log_2 c \geq -\frac{1}{2} \log_2(1 - s) > \kappa_0(s) \geq \kappa_r(s),$$

where the strict middle inequality was proved in the proof of Lemma 7.4. Hence the terminal ambient coordinate is positive.

The second part of Proposition 7.3 now produces a level- $(j + 1)$ tuple with strictly smaller exponent and strictly larger spectral quantity. Retain the factor c . If $c = 1$, this datum is approximable at level $j + 1 \leq r + 1$; if $c < 1$, it is approximable at level $j + 2 \leq r + 1$. Its strict spectral margin persists under approximation, proving $\kappa_{r+1}(s) < \kappa_r(s)$.

The same compactification shows that κ_r is lower semicontinuous on $(0, 1)$: if $s_n \rightarrow s$, nearly minimizing tuples have uniformly bounded exponents, and their limiting datum is feasible at s and approximable at level at most r . At zero, $0 \leq \kappa_r(t) \leq \kappa_0(t) \rightarrow 0$. Therefore the infimum defining $\bar{\kappa}_r(s)$ is attained on $[0, s]$. It is not attained at zero: as $t \downarrow 0$,

$$\kappa_r(t) \leq \kappa_0(t) = O(t^2 \log(1/t)), \quad \frac{1}{2} \log_2(1 - t) = -\frac{t}{2 \log 2} + O(t^2).$$

Evaluating the first strict inequality at the positive minimizing threshold proves $\bar{\kappa}_{r+1}(s) < \bar{\kappa}_r(s)$. The first-level comparisons are Lemma 7.4. $\square$

8. SPHERE PACKINGS

We now convert the spherical hierarchy into an upper bound for the maximal Euclidean sphere-packing density Δ_n . The upper-hemisphere comparison (89) bounds Δ_n by the product of a spherical-code size and an explicit spherical-cap volume factor. Unlike the coding problem, where the separation threshold s is fixed, the packing comparison allows us to optimize over s . The resulting packing exponent balances the spherical-code exponent against the spherical-cap volume exponent.

8.1. From spherical codes to sphere packings. We combine the upper-hemisphere comparison (89) with the finite-level spherical-code bounds and show that spherical-cap optimization does not further improve the angle-optimized packing exponent.

For $-1 < s < 1$, Sidelnikov's upper-hemisphere inequality [Sid74], in the normalization of [CZ14, (2.1)], states that

$$\Delta_n \leq \left(\frac{1-s}{2}\right)^{n/2} A(n+1, s). \quad (89)$$

The factor $((1-s)/2)^{n/2}$ is the geometric cost of the upper-hemisphere comparison. Thus any spherical-code exponent $A(n, s) \leq 2^{(B(s)+o(1))n}$ gives

$$\Delta_n \leq 2^{(B(s) - \frac{1}{2} \log_2 \frac{2}{1-s} + o(1))n}. \quad (90)$$

At hierarchy level r , define the positive sphere-packing decay exponent γ_r obtained by optimizing (90):

$$\begin{aligned} \gamma_r &= \sup_{0 < s < 1} \left(\frac{1}{2} \log_2 \frac{2}{1-s} - \kappa_r(s) \right) \\ &= \sup_{a_1 > b_1 > \dots > b_r > a_{r+1} \geq 0} \left(\frac{1}{2} \log_2 \frac{2}{1 - 2\Gamma_r(\mathbf{a}, \mathbf{b})} - \Phi_r(\mathbf{a}, \mathbf{b}) \right). \end{aligned} \quad (91)$$

For a fixed parameter tuple $(\mathbf{a}, \mathbf{b})$ satisfying (80), the packing objective $\frac{1}{2} \log_2(2/(1-s)) - \Phi_r(\mathbf{a}, \mathbf{b})$ increases with s up to the feasibility boundary $s = 2\Gamma_r(\mathbf{a}, \mathbf{b})$. Substitution at that boundary gives the second formula in (91).

For each fixed s , we may first choose a spherical-cap threshold $0 \leq t \leq s$ and use the cap-optimized spherical-code exponent $\bar{\kappa}_\infty(s)$ defined in (87). Substituting the resulting estimate $A(n+1, s) \leq 2^{(\bar{\kappa}_\infty(s)+o(1))n}$ into the upper-hemisphere inequality (89), and then optimizing the resulting packing decay exponent over s , gives

$$\gamma_\infty = \sup_{0 < s < 1} \left(\frac{1}{2} \log_2 \frac{2}{1-s} - \bar{\kappa}_\infty(s) \right). \quad (92)$$

The spherical-cap reduction does not improve the angle-optimized packing bound. Indeed, changing the inner-product threshold from s to t incurs the spherical-slice exponent $\frac{1}{2} \log_2((1-t)/(1-s))$, which cancels the change in the upper-hemisphere volume exponent:

$$\frac{1}{2} \log_2 \frac{2}{1-s} - \frac{1}{2} \log_2 \frac{1-t}{1-s} = \frac{1}{2} \log_2 \frac{2}{1-t}.$$

Consequently, substituting (87) into (92) gives

$$\begin{aligned} \gamma_\infty &= \sup_{0 < t < 1} \left(\frac{1}{2} \log_2 \frac{2}{1-t} - \kappa_\infty(t) \right) \\ &= \sup_{r \geq 0} \gamma_r. \end{aligned} \quad (93)$$

Every level- r tuple satisfying (80) can be approximated at level $r+1$ by appending $b_{r+1} = \varepsilon$ and $a_{r+2} = 0$; if $a_{r+1} = 0$, first replace it by a sufficiently small positive number. Taking $\varepsilon \downarrow 0$ gives $\gamma_{r+1} \geq \gamma_r$. Therefore $\gamma_\infty = \sup_r \gamma_r = \lim_{r \rightarrow \infty} \gamma_r$. Thus spherical-cap reduction improves fixed-angle code bounds but does not improve the fully angle-optimized packing exponent γ_∞ . In particular,

$$\Delta_n \leq 2^{-(\gamma_\infty + o(1))n}.$$

8.2. A relative-entropy upper bound. A relative-entropy identity, stated precisely in (99), bounds every finite-level packing exponent by a common limiting threshold. The candidate threshold is

$$\lambda_* = \frac{1}{2} \log_2 \frac{2\pi}{e}. \quad (94)$$

Throughout the following argument, $\log$ is natural and D_{KL} denotes relative entropy.

To bound every finite hierarchy level by λ_* , we encode the interlacing quadratic coordinates $x_\ell = \mathcal{A}(a_\ell)$ and $y_m = \mathcal{A}(b_m)$ by the union of $[0, x_{r+1}]$ and the intervals $[y_m, x_m]$. The Stieltjes transform K of the indicator of that union simultaneously records the spectral quantity Γ_r and the dimension exponent Φ_r . Tilting an explicit reference probability measure by e^{-K} identifies the gap between λ_* and the packing objective with a nonnegative relative entropy.

Proposition 8.1. *Fix $r \geq 0$ and a parameter tuple satisfying (80), and set $x_\ell = \mathcal{A}(a_\ell)$, $y_m = \mathcal{A}(b_m)$, and $c = 1/4$. Since the positive residues R_ℓ sum to one and $0 \leq 2q(a_\ell) < 1$, put*

$$Z = 1 - 2\Gamma_r(\mathbf{a}, \mathbf{b}) > 0.$$

Define the interlacing indicator ξ , the associated Stieltjes transform K , two reference measures ρ, ν , and the tilted measure ρ_K , with all measures supported on $0 < t < c$:

$$\xi(u) = \mathbf{1}_{[0, x_{r+1}]}(u) + \sum_{m=1}^r \mathbf{1}_{[y_m, x_m]}(u), \quad (95)$$

$$K(t) = \int_0^\infty \frac{\xi(u)}{t+u} du, \quad (96)$$

$$\rho(dt) = \frac{dt}{\pi \sqrt{t(c-t)}}, \quad \nu(dt) = \frac{dt}{\sqrt{c-t}}, \quad (97)$$

$$\rho_K(dt) = Z^{-1} e^{-K(t)} \rho(dt). \quad (98)$$

The measures ρ , ν , and ρ_K all have total mass one, and

$$\lambda_* - \left(\frac{1}{2} \log_2 \frac{2}{1 - 2\Gamma_r(\mathbf{a}, \mathbf{b})} - \Phi_r(\mathbf{a}, \mathbf{b}) \right) = \frac{D_{\text{KL}}(\nu \parallel \rho_K)}{2 \log 2}. \quad (99)$$

In particular, $\gamma_r \leq \lambda_$ for every $r \geq 0$.*

Proof. The interpolation weights R_ℓ from (81) form a probability distribution on the poles x_ℓ . The rational Stieltjes transform $m(z)$ of that distribution has the equivalent product and partial-fraction representations

$$m(z) = \frac{\prod_{m=1}^r (z - y_m)}{\prod_{\ell=1}^{r+1} (z - x_\ell)} = \sum_{\ell=1}^{r+1} \frac{R_\ell}{z - x_\ell}.$$

Integrating (95) on each interlacing interval and then taking partial fractions gives

$$\begin{aligned} K(t) &= \log \frac{t + x_{r+1}}{t} + \sum_{m=1}^r \log \frac{t + x_m}{t + y_m}, \\ e^{-K(t)} &= \frac{t}{t + x_{r+1}} \prod_{m=1}^r \frac{t + y_m}{t + x_m} = t \sum_{\ell=1}^{r+1} \frac{R_\ell}{t + x_\ell}. \end{aligned} \quad (100)$$

Put $h(x) = (\log 2) \mathbf{H}_{\text{sph}}(\mathcal{A}^{-1}(x))$. Integrating h' on $[0, x_{r+1}]$ and the interlacing intervals $[y_m, x_m]$ gives

$$(\log 2) \Phi_r = \int_0^\infty h'(u) \xi(u) du. \quad (101)$$

Both reference measures ρ and ν have mass one because $c = 1/4$. Differentiating $h(x) = (\log 2) \mathbf{H}_{\text{sph}}(\mathcal{A}^{-1}(x))$ gives

$$h'(x) = \frac{1}{1 + 2a} \log \frac{1 + a}{a}, \quad a = \mathcal{A}^{-1}(x).$$

The substitutions $t = c \sin^2 \theta$ and $t = c(1 - v^2)$, respectively, give

$$2q(\mathcal{A}^{-1}(x)) = \sqrt{\frac{x}{x+c}} = \int_0^c \frac{x}{x+t} \rho(dt), \quad 2h'(x) = \int_0^c \frac{\nu(dt)}{x+t}.$$

By (100), (101), and Tonelli's theorem,

$$Z = \int_0^c e^{-K(t)} \rho(dt), \quad 2(\log 2)\Phi_r = \int_0^c K(t) \nu(dt). \quad (102)$$

The first identity in (102) normalizes ρ_K . Dividing the two reference densities and applying the same substitution $t = c(1 - v^2)$ gives

$$\frac{d\nu}{d\rho}(t) = \pi\sqrt{t}, \quad \int_0^c \log t \nu(dt) = -2.$$

Consequently,

$$D_{\text{KL}}(\nu \parallel \rho_K) = \int_0^c \left(\log \frac{d\nu}{d\rho}(t) + K(t) + \log Z \right) \nu(dt) = \log \pi - 1 + 2(\log 2)\Phi_r + \log Z.$$

Rearrangement and (94) give (99). Nonnegativity of relative entropy and (91) give $\gamma_r \leq \lambda_*$. $\square$

Corollary 8.2. *For every $r \geq 0$,*

$$\gamma_r < \gamma_{r+1} < \lambda_*.$$

Proof. Set $c = 1/4$, and for $0 \leq x < \infty$ define

$$I(x) = \int_0^c \frac{t}{t+x} \rho(dt), \quad f_x(t) = \frac{t}{(t+x)I(x)}, \quad \rho_x = f_x \rho.$$

Include the endpoints by setting $f_0 = 1$ and $f_\infty(t) = 2t/c$. The second identity in (100) expresses ρ_K as a probability mixture of at most $r+1$ measures ρ_x . Conversely, a mixture with distinct finite nodes x_i and positive weights α_i is obtained by choosing

$$R_i = \frac{\alpha_i/I(x_i)}{\sum_j \alpha_j/I(x_j)}.$$

The zeros of $\sum_i R_i/(z - x_i)$ strictly interlace its poles, giving a level- r tuple after approximation or the addition of vanishing rows. Nodes at infinity are obtained by approximation.

The bounds

$$\frac{c}{4(c/2+x)} \leq I(x) \leq \frac{c}{c+x}, \quad \frac{t}{c} \leq f_x(t) \leq 4$$

show that relative entropy is continuous on the compact space of mixtures of at most $r+1$ nodes in $[0, \infty]$. Indeed, $|\log t|$ is integrable against ν . Consequently, Proposition 8.1 gives, with $x_i \in [0, \infty]$,

$$\gamma_r = \lambda_* - \frac{1}{2 \log 2} \min_{\substack{g = \sum_{i=1}^{r+1} \alpha_i f_{x_i} \\ \alpha_i \geq 0, \sum_i \alpha_i = 1}} D_{\text{KL}}(\nu \parallel g\rho).$$

Write $h = d\nu/d\rho = \pi\sqrt{t}$. The probability measure $\eta(dx) = I(x)x^{-1/2} dx$ on $(0, \infty)$ satisfies

$$\int_0^\infty f_x(t) \eta(dx) = \int_0^\infty \frac{t}{t+x} \frac{dx}{\sqrt{x}} = \pi\sqrt{t} = h(t).$$

No finite mixture g equals h : if it contains a positive atom at zero, then $g(0) > 0$; otherwise, $g(t) = O(t)$ as $t \downarrow 0$, whereas $h(t) = \pi\sqrt{t}$. Thus the minimum relative entropy is positive, giving $\gamma_r < \lambda_*$.

Let g minimize the relative entropy at level r , and put $A(x) = \int_0^c f_x(t)g(t)^{-1} \nu(dt)$. Then

$$\int_0^\infty A(x) \eta(dx) = \int_0^c \frac{h(t)^2}{g(t)} \rho(dt) > 1,$$

where strictness follows from Cauchy–Schwarz and $g \neq h$. Hence $A(x) > 1$ for some finite $x > 0$ distinct from the existing mixture nodes. Since $f_x/g \leq c/(xI(x))$, differentiation under the integral is justified. For $g_\varepsilon = (1 - \varepsilon)g + \varepsilon f_x$,

$$\left. \frac{d}{d\varepsilon} D_{\text{KL}}(\nu \| g_\varepsilon \rho) \right|_{\varepsilon=0} = 1 - A(x) < 0.$$

This level- $(r + 1)$ mixture has smaller relative entropy, proving $\gamma_{r+1} > \gamma_r$. $\square$

8.3. A sharp Chebyshev construction. Corollary 8.2 shows that the finite-level exponents increase strictly while remaining below λ_* . To show that their limit equals this threshold, we form interlacing parameter tuples from Chebyshev roots and critical points. The resulting sphere-packing bound agrees with the full-space companion in Chapter 1, but is obtained here entirely from spherical certificates.

Theorem 8.3. *The spherical hierarchy attains the threshold*

$$\lim_{r \rightarrow \infty} \gamma_r = \gamma_\infty = \lambda_* = \frac{1}{2} \log_2 \frac{2\pi}{e}.$$

Consequently,

$$\Delta_n \leq 2^{-(\lambda_* + o(1))n}. \quad (103)$$

Proof. By Corollary 8.2, every finite-level exponent satisfies $\gamma_r < \lambda_*$. To obtain matching exponents from below, fix $R > 0$, set $N = r + 1$, and choose the roots and critical points of a Chebyshev polynomial shifted from $[-1, 1]$ to $[0, R]$. The roots and critical points, for $1 \leq \ell \leq N$ and $1 \leq m < N$, respectively, are

$$x_\ell = \frac{R}{2} \left(1 + \cos \frac{(2\ell - 1)\pi}{2N} \right), \quad y_m = \frac{R}{2} \left(1 + \cos \frac{m\pi}{N} \right). \quad (104)$$

The roots x_ℓ and the critical points y_m strictly interlace. If Q_N is the monic shifted Chebyshev polynomial with roots $x_1, \dots, x_N$, then Q'_N has roots $y_1, \dots, y_{N-1}$. The logarithmic derivative of Q_N therefore gives

$$\frac{\prod_{m=1}^{N-1} (z - y_m)}{\prod_{\ell=1}^N (z - x_\ell)} = \frac{Q'_N(z)}{NQ_N(z)} = \frac{1}{N} \sum_{\ell=1}^N \frac{1}{z - x_\ell}.$$

Comparing this partial-fraction expansion with the residues in (81) shows that every interpolation weight is $R_\ell = 1/N$. Put $a_\ell = \mathcal{A}^{-1}(x_\ell)$ and $b_m = \mathcal{A}^{-1}(y_m)$. For $t > 0$, the interlacing indicator ξ_N and its Stieltjes transform K_N are

$$\xi_N(u) = \mathbf{1}_{[0, x_N]}(u) + \sum_{m=1}^{N-1} \mathbf{1}_{[y_m, x_m]}(u), \quad K_N(t) = \int_0^R \frac{\xi_N(u)}{t + u} du.$$

By (102),

$$\begin{aligned} Z_N &= 1 - 2\Gamma_{N-1}(\mathbf{a}, \mathbf{b}) = 1 - \frac{1}{N} \sum_{\ell=1}^N \sqrt{\frac{x_\ell}{x_\ell + 1/4}} \\ &= \int_0^{1/4} e^{-K_N(t)} \frac{dt}{\pi \sqrt{t(1/4 - t)}}. \end{aligned}$$

Under $u = \frac{R}{2}(1 + \cos \theta)$, the critical points and the endpoints $\theta = 0, \pi$ divide $(0, \pi)$ into the N equal intervals $((m - 1)\pi/N, m\pi/N)$. Each interval contains one root at its midpoint, and ξ_N selects exactly half of the interval. The change of variables contributes the Jacobian $\frac{R}{2} \sin \theta$. Averaging against continuous test functions and then using their density in $L^1(0, R)$ gives weak-* convergence in $L^\infty(0, R)$. Explicitly,

$$\lim_{N \rightarrow \infty} \int_0^R f(u) \xi_N(u) du = \frac{1}{2} \int_0^R f(u) du \quad \text{for every } f \in L^1(0, R).$$

Since $h'(u) = O(\log(1/u))$ as $u \downarrow 0$, both h' and $u \mapsto (t+u)^{-1}$ belong to $L^1(0, R)$ for every fixed $t > 0$. Therefore, (96) and (101) give

$$\begin{aligned} K_N(t) &\longrightarrow \frac{1}{2} \log \frac{t+R}{t}, \\ (\log 2) \Phi_{N-1} &\longrightarrow \frac{1}{2} h(R). \end{aligned} \tag{105}$$

Since $0 \leq e^{-K_N} \leq 1$, dominated convergence gives

$$Z_N \longrightarrow Z_R = \int_0^{1/4} \sqrt{\frac{t}{t+R}} \rho(dt) = \frac{2}{\pi} \arcsin \frac{1}{\sqrt{4R+1}}. \tag{106}$$

Taking $N \rightarrow \infty$ with R fixed gives

$$\liminf_{r \rightarrow \infty} \gamma_r \geq \frac{\log(2/Z_R) - h(R)}{2 \log 2}. \tag{107}$$

As $R \rightarrow \infty$,

$$Z_R = \frac{1 + o(1)}{\pi \sqrt{R}}, \quad h(R) = \frac{1}{2} \log R + 1 + o(1),$$

so the right side of (107) tends to λ_* . Together with the fixed-level upper bound and (93), this proves the threshold identity.

For $\varepsilon > 0$, choose successively R , a finite r , and $s < 2\Gamma_r$ whose packing exponent exceeds $\lambda_* - \varepsilon$. Keep these parameters fixed. Then Theorem 1.2 and (89) give

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log_2 \Delta_n \leq -\lambda_* + \varepsilon.$$

Taking $n \rightarrow \infty$ before $\varepsilon \downarrow 0$ proves (103). $\square$

APPENDIX A. ORTHOGONAL REPRESENTATIONS AND ASYMPTOTIC DIMENSIONS

The spherical arguments in §§6 and 7 require two representation-theoretic inputs. The coordinate-transition formulas give Proposition 6.1 and the finite spherical-code bound in Theorem 6.2. The uniform representation-dimension estimates give (85) and (86), which enter the proof of the main spherical bound in Theorem 1.2.

The spherical transition graph has vertices given by ambient $SO(n)$ -representations containing a fixed representation of the point stabilizer $SO(n-1)$, and its edges come from multiplication by a coordinate. The ambient and stabilizer spaces are described in §5.3. §A.1 identifies the possible vertices and edges, and §A.2 derives the common transition formula (74), its full-graph normalization (75), and its dimension-weighted balance (76). These are the three conclusions of Proposition 6.1. Finally, §A.3 gives the exact Weyl dimensions and proves the uniform estimates in Lemma A.1, which determine the dimension ratio in (78). Throughout, the number of nonzero highest-weight coordinates stays fixed as n increases.

For the standard orthogonal-group conventions and formulas, see Kravchuk [Kra18, §§2.1–2.3 and App. B]. There, (2.18)–(2.21) give multiplicity-free restriction; (B.2) and its boundary correction give the tensor product; (B.3)–(B.4) specify shifted weights; (B.5)–(B.6) give Weyl dimensions; and (B.11)–(B.13) give coordinate-multiplication coefficients. Source equation numbers below refer to this citation.

A.1. Stable tensor conventions. The vertices of the spherical graph are determined by the orthogonal-group restriction rule (72); its possible edges are determined by tensoring with the standard representation. We first record the labeling and tensor conventions needed for both inputs. As in §5.3, the ambient rotation group is $SO(n)$, and the stabilizer of a point of the sphere is $SO(n-1)$. We call representations of $SO(n)$ ambient representations and representations of its point-fixing subgroup stabilizer representations.

Fix $r \geq 0$ and $n \geq 2r + 4$, and retain ambient representations with at most $r + 1$ nonzero rows and stabilizer representations with at most r nonzero rows. The remaining highest-weight coordinates, called the zero tails, all vanish. Because at least one such coordinate remains, these

representations avoid the exceptional splitting that can occur when a Young diagram reaches the final orthogonal-group weight coordinate. In this stable range, their irreducible tensor representations are indexed by dominant integral highest weights: weakly decreasing sequences of nonnegative integers whose nonzero entries are the row lengths of the corresponding Young diagram. Write the ambient and stabilizer weights, respectively, as

$$\lambda = (\lambda_1, \dots, \lambda_{r+1}, 0, \dots), \quad \mu = (\mu_1, \dots, \mu_r, 0, \dots).$$

The restriction of the ambient representation V_λ to $SO(n-1)$ contains the stabilizer representation E_μ exactly when their row lengths interlace as in (72):

$$\lambda_1 \geq \mu_1 \geq \lambda_2 \geq \dots \geq \mu_r \geq \lambda_{r+1} \geq 0.$$

Each such stabilizer representation occurs with multiplicity one. The rotation-equivariant coordinate maps in (52) are obtained by tensoring with the standard $SO(n)$ representation $\mathbb{R}^n$, of highest weight $(1, 0, \dots, 0)$. The tensor-product formula (B.2) therefore gives the possible coordinate transitions $\lambda \leftrightarrow \lambda \pm e_\ell$, provided the target weight remains dominant and still interlaces μ . Although the odd-dimensional tensor-product formula initially appears to contain an additional same-weight term, its zero-tail boundary correction removes that term. Hence these transitions are precisely the edges of the graph in §6.

A.2. Cancellation of the parity-dependent factors. §A.1 identifies the possible directed edges $\lambda \rightarrow \lambda \pm e_\ell$. Their squared coordinate coefficients are the quantities $p_{\ell, \pm} = |c_{\ell, \pm}|^2$ in Proposition 6.1, normalized by the unit base-point coordinate and the isometric embeddings in (52). Kravchuk's (B.12) and (B.13) give these coefficients separately for odd and even ambient dimensions. Factors belonging to zero highest-weight coordinates cancel in both expressions, yielding the common formula (74). Parseval then gives (75), and the Weyl dimension formula gives (76). As explained in §5.3, the $\text{Spin}(n)$ formulas apply unchanged to integral $SO(n)$ tensor representations.

In Kravchuk's notation, the highest weights $\mathbf{m}_n$, $\mathbf{m}_{n-1}$, and $\mathbf{m}_n(\pm\ell)$ correspond to λ , μ , and $\lambda \pm e_\ell$, respectively. Kravchuk's coordinate-multiplication formulas use shifted weights, obtained by adding the conventional dimension-dependent offsets to the highest-weight coordinates. In the notation of (73), put

$$L_\ell = \lambda_\ell + \frac{n}{2} - \ell, \quad M_m = \mu_m + \frac{n-1}{2} - m, \quad \rho_{n,r} = \frac{n}{2} - r - 1. \quad (108)$$

The variables $x_{d,j}$ in (B.3)–(B.4) satisfy

$x_{n,\ell}$	$n = 2N + 1$	$n = 2N$
$x_{n-1,m}$	$L_\ell - \frac{1}{2}$	L_ℓ
	M_m	$M_m - \frac{1}{2}$.

To interpret Kravchuk's coordinate coefficients, the multiplicity-free branching formulas (2.18)–(2.21), applied successively along $SO(n) \supset SO(n-1) \supset \dots \supset SO(2)$ give the orthonormal Gelfand–Tsetlin basis: its basis vectors are indexed by a sequence of interlacing highest weights, one for each group in the chain. Choose $SO(n-1)$ to fix the coordinate line $\mathbb{R}e_1$. The symbol $\bullet$ in (B.11)–(B.13) selects this fixed one-dimensional summand in $\mathbb{R}^n \downarrow_{SO(n-1)} = \mathbb{R}e_1 \oplus e_1^\perp$. The squared coefficient for this fixed summand is exactly the directed transition weight $p_{\ell, \pm}$ associated with (52); no additional coordinate normalization is required.

First suppose the ambient dimension is odd, $n = 2N + 1$. Write $L = L_\ell$ for the shifted active coordinate in (108). Squaring the odd-dimensional coordinate coefficient (B.12) gives

$$p_{\ell, \pm} = \frac{\prod_{m=1}^N ((L \pm \frac{1}{2})^2 - M_m^2)}{2L(L \pm \frac{1}{2}) \prod_{\substack{q=1 \\ q \neq \ell}}^N (L^2 - L_q^2)}. \quad (109)$$

Put $A = N - r - 1$, so $\rho_{n,r} = A + \frac{1}{2}$. For $r + 1 \leq m \leq N$ and $r + 2 \leq q \leq N$, the zero tail gives

$$M_m = N - m, \quad L_q = N + \frac{1}{2} - q.$$

The trailing numerator and denominator factors cancel as

$$\frac{\prod_{m=r+1}^N ((L \pm \frac{1}{2})^2 - M_m^2)}{(L \pm \frac{1}{2}) \prod_{q=r+2}^N (L^2 - L_q^2)} = L \pm \left(A + \frac{1}{2}\right) = L \pm \rho_{n,r}.$$

Empty products are 1. Substitution into (109) leaves only numerator indices $m \leq r$ and denominator indices $q \leq r + 1$.

The odd-dimensional Pieri formula in (B.2) appears to contain a copy of V_λ itself. However, its boundary correction applies because $N \geq r + 2$ and the final highest-weight coordinate λ_N is zero. Decreasing that zero coordinate gives a reflected weight representing the same $SO(2N + 1)$ representation; the correction subtracts the repeated term. Consequently,

$$\text{Hom}_{SO(2N+1)}(V_\lambda, \mathbb{R}^{2N+1} \otimes V_\lambda) = 0.$$

Hence coordinate multiplication has no transition from V_λ back to itself. On the zero-tail boundary $\lambda_N = 0$, the coefficient in (B.11) has the indeterminate form $0/0$, so the remaining transitions must be determined from the corrected tensor decomposition.

Now suppose the ambient dimension is even, $n = 2N$, and again write $L = L_\ell$. Squaring the even-dimensional coordinate coefficient (B.13) gives

$$p_{\ell,\pm} = \frac{\prod_{m=1}^{N-1} ((L \pm \frac{1}{2})^2 - M_m^2)}{2 \prod_{\substack{q=1 \\ q \neq \ell}}^N (L^2 - L_q^2)}. \quad (110)$$

Put $A = N - r - 2$, so $\rho_{n,r} = A + 1$. For $r + 1 \leq m \leq N - 1$ and $r + 2 \leq q \leq N$, the zero tail gives

$$M_m = N - \frac{1}{2} - m, \quad L_q = N - q.$$

The numerator factors for M_m and the denominator factors for L_q cancel according to

$$\frac{\prod_{m=r+1}^{N-1} ((L \pm \frac{1}{2})^2 - M_m^2)}{\prod_{q=r+2}^N (L^2 - L_q^2)} = \frac{L \pm (A + 1)}{L} = \frac{L \pm \rho_{n,r}}{L}.$$

Substitution into (110) shows that both parities give

$$p_{\ell,\pm} = \frac{(L_\ell \pm \rho_{n,r}) \prod_{m=1}^r ((L_\ell \pm \frac{1}{2})^2 - M_m^2)}{2L_\ell \prod_{\substack{q=1 \\ q \neq \ell}}^{r+1} (L_\ell^2 - L_q^2)}.$$

The common odd- and even-dimensional expression is (74). If the row lengths of $\lambda \pm e_\ell$ are not weakly decreasing or fail (72), there is no target representation containing E_μ , and the transition probability is zero. Applying Parseval to the orthogonal irreducible projections of $\ell_x \otimes \varphi_{\lambda,x} Y$, whose squared norm is $\|Y\|^2$, gives the full-graph normalization (75). On every existing edge, (111) gives

$$\frac{D_n(\lambda + e_\ell)}{D_n(\lambda)} = \frac{p_{\ell,+}(\lambda; \mu)}{p_{\ell,-}(\lambda + e_\ell; \mu)}.$$

Equivalently, $D_n(\lambda) p_{\ell,+}(\lambda; \mu) = D_n(\lambda + e_\ell) p_{\ell,-}(\lambda + e_\ell; \mu)$. Thus weighting each transition by the dimension of its source makes the forward and reverse directions equal. This is (76), and hence verifies the reciprocity assumption (53) used to construct the symmetric spherical graph.

A.3. Weyl dimensions. The transition calculation gives the spectral constraint in [Proposition 6.1](#), but the spherical-code bound (78) also contains the ratio between the total ambient dimension and the dimension of the common stabilizer space. We record exact formulas for both dimensions and establish the uniform exponential estimates stated in [Lemma A.1](#).

Write $D_n(\lambda) = \dim V_\lambda$. For a dominant integral highest weight λ with at most $r+1$ nonzero rows, Weyl's dimension formula [[GW09](#), §7.1.2] simplifies to

$$D_n(\lambda) = \prod_{i=1}^{r+1} \frac{2\lambda_i + n - 2i}{n - 2i} \frac{(\lambda_i + n - i - r - 2)!}{(n - i - r - 2)!} \frac{(r + 1 - i)!}{(\lambda_i + r + 1 - i)!} \times \prod_{1 \leq i < j \leq r+1} \frac{(\lambda_i - \lambda_j + j - i)(\lambda_i + \lambda_j + n - i - j)}{(j - i)(n - i - j)}. \quad (111)$$

Formula (111) is the zero-tail specialization of (B.5)–(B.6). If $r \geq 1$, substituting $(n, \lambda, r) \mapsto (n - 1, \mu, r - 1)$ gives the stabilizer dimension $\dim E_\mu$. If $r = 0$, the stabilizer representation is trivial and has dimension 1, and we write $D_{n-1}(0) = 1$. The corresponding exponential rates are expressed using the spherical entropy H_{sph} defined in (9).

Lemma A.1. *Fix $r \geq 0$ and $A < \infty$, and assume $n \geq 2r + 4$. Uniformly over dominant integral ambient weights $\lambda = (\lambda_1, \dots, \lambda_{r+1}, 0, \dots)$ and stabilizer weights $\mu = (\mu_1, \dots, \mu_r, 0, \dots)$ with $0 \leq \lambda_i, \mu_j \leq An$,*

$$\begin{aligned} \frac{1}{n} \log_2 D_n(\lambda) &= \sum_{i=1}^{r+1} H_{\text{sph}}\left(\frac{\lambda_i}{n}\right) + O_{r,A}\left(\frac{\log(n+2)}{n}\right), \\ \frac{1}{n} \log_2 D_{n-1}(\mu) &= \sum_{j=1}^r H_{\text{sph}}\left(\frac{\mu_j}{n}\right) + O_{r,A}\left(\frac{\log(n+2)}{n}\right). \end{aligned}$$

In particular, if $\lambda_i/n \rightarrow a_i$ and $\mu_j/n \rightarrow b_j$, the dimension exponents converge to $\sum_i H_{\text{sph}}(a_i)$ and $\sum_j H_{\text{sph}}(b_j)$, respectively.

Proof. Apply uniform Stirling to the finitely many factorials in (111). Dominance bounds every active difference between 1 and $O_{r,A}(n)$, so all remaining factors contribute $O_{r,A}(\log(n+2))$, including on chamber walls and at zero rows. For $r \geq 1$, the stabilizer estimate follows identically in dimension $n - 1$; for $r = 0$, it is the identity $\log_2 \dim E_0 = 0$. $\square$

APPENDIX B. THE SPHERICAL-TO-EUCLIDEAN LIMIT

The companion paper in Chapter 1 analyzes the Cohn–Elkies sphere-packing linear program directly in Euclidean space, proving both a universal bound for its Fourier-positive auxiliary functions satisfying the required normalization and sign conditions and a matching construction. Neither argument requires the spherical hierarchy developed here. Nevertheless, the two approaches share more than their final exponent: several objects in the Euclidean construction arise naturally when the inner-product threshold s of our spherical certificates tends to 1.

Cohn and Zhao's upper-hemisphere construction gives the correspondence between spherical certificates and Euclidean Cohn–Elkies auxiliary functions [[CZ14](#), Thm. 3.4 and subsequent discussion]. We recall the resulting change of scale, then identify the probability measure governing the optimal spherical hierarchy with the limiting measure in the companion's Euclidean argument. The limiting spherical potential also recovers the principal weight in its Mellin-transform construction. The companion uses a second, smaller weight to control its Fourier estimates; that additional correction has no spherical counterpart here.

B.1. The upper-hemisphere correspondence. Let $g_s(t)$ be a positive-definite spherical certificate on $\mathbb{S}^d$, nonpositive for $t \leq s$, and let $g_{s,0} > 0$ be its constant Gegenbauer coefficient. Set

$$L_s = \sqrt{\frac{2}{1-s}}, \quad \pi(u) = (u, \sqrt{1-|u|^2}) \quad (|u| < 1).$$

For a nonzero nonnegative radial cutoff supported in the radius- L_s ball, $\eta \in C_c^\infty(B_{L_s}^d)$, the Cohn–Zhao upper-hemisphere construction [CZ14, Thm. 3.4 and subsequent discussion] gives the Euclidean function

$$f_s(x) = \int_{\mathbb{R}^d} \eta(z + 2x) \eta(z) g_s(\langle \pi((z + 2x)/L_s), \pi(z/L_s) \rangle) dz.$$

The integrand is zero unless both arguments of η belong to its support. Rotational invariance makes f_s a radial Schwartz function, and positive definiteness of g_s gives $\hat{f}_s \geq 0$. If $|x| \geq 1$, the spherical inner product in the integrand is at most $1 - 2|x|^2/L_s^2 \leq s$, so $f_s(x) \leq 0$. Thus f_s is a Euclidean Cohn–Elkies auxiliary function. Centering a Gram factorization of g_s gives

$$f_s(0) = g_s(1) \|\eta\|_2^2, \quad \hat{f}_s(0) \geq 2^{-d} g_{s,0} \|\eta\|_1^2.$$

Taking cutoffs approaching the indicator of $B_{L_s}^d$ therefore gives [CZ14, Thm. 3.4 and subsequent discussion]

$$\text{LP}_d \leq \left(\frac{1-s}{2} \right)^{d/2} \frac{g_s(1)}{g_{s,0}}.$$

In particular, $L_s \rightarrow \infty$ as $s \uparrow 1$, giving the Euclidean scaling used below.

B.2. The limiting hyperbolic measure. The spherical and Euclidean optimality arguments each single out a probability measure. To compare them, we first describe how the spherical measure arises from the interlacing Chebyshev construction in the proof of Theorem 8.3. Fix $T > 0$, and let $x_1, \dots, x_N$ and $y_1, \dots, y_{N-1}$ be, respectively, the roots and critical points of the shifted Chebyshev polynomial on $[0, T]$, as in (104). These two sets of nodes interlace. The indicator ξ_N selects the alternating intervals between them, and its Stieltjes transform K_N converts those intervals into the exponential tilt of the reference probability measure $\rho(dt) = dt/(\pi\sqrt{t(1/4-t)})$ on $(0, 1/4)$:

$$\begin{aligned} \xi_N(v) &= \mathbf{1}_{[0, x_N]}(v) + \sum_{m=1}^{N-1} \mathbf{1}_{[y_m, x_m]}(v), \\ K_N(t) &= \int_0^T \frac{\xi_N(v)}{t+v} dv, \quad Z_N = \int_0^{1/4} e^{-K_N(t)} \rho(dt). \end{aligned}$$

Here $t \in (0, 1/4)$ is the spectral variable on which ρ is supported, Z_N normalizes the tilted measure in (98), and the hierarchy level is $N - 1$. Keeping T fixed and letting $N \rightarrow \infty$, (105) and (106) give

$$K_T(t) = \frac{1}{2} \log \frac{t+T}{t}, \quad Z_T = \frac{2}{\pi} \arcsin \frac{1}{\sqrt{4T+1}}.$$

The limiting feasible inner-product threshold is $s_T = 1 - Z_T$, so

$$1 - s_T \sim \frac{1}{\pi\sqrt{T}}, \quad L_T = \sqrt{\frac{2}{1-s_T}} \sim \sqrt{2\pi} T^{1/4}.$$

In particular, increasing T realizes the small-angle limit $s_T \uparrow 1$.

The reference probability measures ρ and ν were introduced in (97). The relative entropy of ν from the tilted measure controls the gap to the optimal packing exponent in (99). To identify ν with the Euclidean limiting measure, change from the bounded spectral variable $t \in (0, 1/4)$ to the unbounded hyperbolic coordinate $a > 0$ defined by

$$t = \frac{1}{4} \operatorname{sech}^2 a. \quad (112)$$

Under this change of variables, the reference and target measures in (97) are

$$\rho(dt) = \frac{2}{\pi} \operatorname{sech} a da, \quad \nu(dt) = \operatorname{sech}^2 a da. \quad (113)$$

The tilted probability measure ρ_{K_T} from (98) then has density

$$\rho_{K_T}(da) = \frac{2}{\pi Z_T} \frac{\operatorname{sech} a}{\sqrt{1 + 4T \cosh^2 a}} da \rightarrow \operatorname{sech}^2 a da.$$

The convergence holds in total variation by dominated convergence and $\pi Z_T \sqrt{T} \rightarrow 1$.

The Euclidean lower-bound argument in Chapter 1, Section 3.1 obtains its corresponding probability measure from the Poisson kernel of a complex strip. If $u \in \mathbb{R}$ denotes the rescaled Mellin frequency, the limiting measure is

$$p(u) du = \frac{\pi}{4} \operatorname{sech}^2\left(\frac{\pi u}{2}\right) du, \quad u \in \mathbb{R}.$$

The pushforward of $p(u) du$ under $a = \pi|u|/2$ is exactly $\operatorname{sech}^2 a da$. Thus the target probability measure in the spherical relative-entropy bound and the measure in the Euclidean lower bound coincide in hyperbolic coordinates. The hyperbolic densities in (113) also give

$$\frac{d\nu}{d\rho}(a) = \frac{\pi}{2} \operatorname{sech} a, \quad D_{\text{KL}}(\nu \parallel \rho) = \log \pi - 1,$$

since $\int_0^\infty \operatorname{sech}^2 a \log \cosh a da = 1 - \log 2$.

B.3. The limiting Mellin weight. The common probability measure identifies the spherical and Euclidean optimality arguments at the distributional level. We now give a stronger correspondence: the limiting spherical potential determines the principal weight in the Euclidean construction in Chapter 1, Section 4.1.

For a radial profile $g(r)$, its Mellin transform is $M_g(\zeta) = \int_0^\infty g(r) r^{\zeta-1} dr$. In dimension d , put $\lambda = d/2$ and $X_g(\tau) = M_g(\lambda - i\tau)$. On this symmetry line, the Mellin–Fourier identity is given in Chapter 1, Section 2.2:

$$X_{\widehat{g}}(\tau) = m_\lambda(\tau) X_g(-\tau), \quad m_\lambda(\tau) = \pi^{i\tau} \frac{\Gamma((\lambda - i\tau)/2)}{\Gamma((\lambda + i\tau)/2)}.$$

The common envelope in the companion construction is

$$E_\lambda(\tau) = \pi^{i\tau/2} \Gamma\left(\frac{\lambda - i\tau}{2}\right) \exp(\lambda h_\epsilon(\tau/\lambda)).$$

If h_ϵ is even, then $m_\lambda(\tau) E_\lambda(-\tau) = E_\lambda(\tau)$, so the modification by h_ϵ preserves Fourier reflection; see Chapter 1, Section 4.1. It is assembled from oscillations $\cos(az)$, where $z = \tau/\lambda$ and $a > 0$ is an oscillation scale, not a spatial radius. The oscillation weight occupies two disjoint intervals of scales: its principal negative interval determines the sharp packing exponent, whereas a smaller positive interval controls the global Fourier estimates. Only the negative interval has a counterpart in the limiting spherical potential.

Under the hyperbolic change of variables (112), the potential K_T diverges by an additive constant as $T \rightarrow \infty$. Subtract this constant:

$$\widetilde{K}_T(a) = K_T\left(\frac{1}{4} \operatorname{sech}^2 a\right) - \frac{1}{2} \log(4T) = \frac{1}{2} \log\left(\cosh^2 a + \frac{1}{4T}\right).$$

It follows that

$$\widetilde{K}_T(a) \longrightarrow K_\infty(a) = \log \cosh a, \quad e^{-K_\infty(a)} = \operatorname{sech} a. \quad (114)$$

More precisely, the Euclidean construction modifies the logarithm of its Mellin profile by the even function

$$h_\epsilon(z) = \int_0^\infty w_\epsilon(a) (\cos(az) - 1) da.$$

Here the rescaled frequency z may be complex. Fix sufficiently small $\epsilon > 0$, and put

$$I_\epsilon = [\epsilon^2, \log(1/\epsilon)], \quad b_\epsilon(a) = 1 - 2\epsilon(1+a).$$

The principal negative part of the oscillation weight is

$$w_{\epsilon,-}(a) = -\frac{b_\epsilon(a) e^{-2a}}{2a^2 \cosh a} \mathbf{1}_{I_\epsilon}(a).$$

The companion supplements this negative part with an exponentially smaller positive correction on a distant scale interval. That correction controls its global Fourier estimates without affecting the limiting radius; see Chapter 1, Sections 4.1–4.2.

To recover the negative weight $w_{\epsilon,-}$ from the spherical potential, define the finite-cutoff weight

$$w_{\epsilon,T}(a) = -\frac{b_\epsilon(a)e^{-2a}}{2a^2}e^{-\tilde{K}_T(a)}\mathbf{1}_{I_\epsilon}(a).$$

For each fixed $\epsilon > 0$, (114) gives $w_{\epsilon,T} \rightarrow w_{\epsilon,-}$ uniformly on I_ϵ . Thus the principal weight of the full-space construction is an exponential tilt of the limiting spherical Stieltjes potential.

As $\epsilon \downarrow 0$, the intervals I_ϵ increase to $(0, \infty)$, and $b_\epsilon(a) \rightarrow 1$ at each fixed $a > 0$. Thus the principal negative weight converges pointwise to $w_*(a) = -e^{-2a-K_\infty(a)}/(2a^2)$, and its limiting Mellin modification is

$$h_*(z) = \frac{1}{2} \int_0^\infty \frac{e^{-2a-K_\infty(a)}}{a^2} (1 - \cos(az)) da.$$

The integral converges locally uniformly on $\{z \in \mathbb{C} : |\operatorname{Im} z| < 3\}$. At the limiting radius-determining parameter $u = 1$, the contribution of w_* to the logarithm of the radius is

$$\int_0^\infty w_*(a) a \sinh a da = -\frac{1}{2} \int_0^\infty \frac{e^{-2a} \tanh a}{a} da. \quad (115)$$

The expansion

$$e^{-2a} \tanh a = \sum_{j=0}^\infty e^{-(4j+2)a} (1 - e^{-2a})^2$$

has nonnegative summands, so Tonelli's theorem permits termwise integration. For $u, v > 0$, the elementary integral identity

$$\int_0^\infty \frac{e^{-ua} - e^{-va}}{a} da = \log \frac{v}{u}$$

then gives Wallis's product:

$$\int_0^\infty \frac{e^{-2a} \tanh a}{a} da = \log \prod_{j=0}^\infty \frac{(2j+2)^2}{(2j+1)(2j+3)} = \log \frac{\pi}{2}.$$

The companion's Fourier-pair theorem in Chapter 1, Theorem 4.1 constructs radial Schwartz functions f_-, f_+ and an exterior radius $R_{\epsilon,d}$ such that $\hat{f}_- = f_+ > 0$, $f_-(0) = f_+(0) > 0$, and $f_-(x) < 0$ whenever $|x| \geq R_{\epsilon,d}$. In its radius formula, the principal negative weight contributes (115), whereas the positive correction vanishes in the limit. Consequently,

$$\lim_{\epsilon \downarrow 0} \lim_{d \rightarrow \infty} \frac{R_{\epsilon,d}}{\sqrt{d}} = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \log \frac{\pi}{2}\right) = \frac{1}{\pi}.$$

If v_d denotes the volume of the unit ball in $\mathbb{R}^d$, then $v_d^{1/d} \sim \sqrt{2\pi e/d}$. Thus we recover the same Cohn–Elkies density rate as (103):

$$\lim_{\epsilon \downarrow 0} \lim_{d \rightarrow \infty} v_d^{1/d} \frac{R_{\epsilon,d}}{2} = \sqrt{\frac{e}{2\pi}} = 2^{-\lambda*}.$$

Thus the spherical construction recovers both the target probability measure and the principal oscillation weight in the companion's optimal full-space construction. This identifies limiting objects, without asserting that finite-dimensional spherical certificates converge to a particular Euclidean auxiliary function. The order of limits also matters: the dimension first tends to infinity at fixed hierarchy level, the hierarchy level then increases at fixed T , and finally $T \rightarrow \infty$ forces $s_T \uparrow 1$. If the last two limits are taken jointly, choosing $N/\sqrt{T} \rightarrow \infty$ retains every fixed neighborhood of the zero endpoint of the interlacing interval, since the smallest shifted Chebyshev node satisfies

$$x_N = \frac{T}{2} \left(1 - \cos \frac{\pi}{2N}\right) \sim \frac{\pi^2 T}{16N^2} \rightarrow 0.$$

REFERENCES

- [AJCHLT20] N. Afkhami-Jeddi, H. Cohn, T. Hartman, D. de Laat, and A. Tajdini, *High-dimensional sphere packing and the modular bootstrap*, J. High Energy Phys. **12** (2020), article 066, doi:10.1007/JHEP12(2020)066.
- [BV08] C. Bachoc and F. Vallentin, *New upper bounds for kissing numbers from semidefinite programming*, J. Amer. Math. Soc. **21** (2008), 909–924, doi:10.1090/S0894-0347-07-00589-9.
- [BN06] A. Barg and D. Nogin, *Spectral approach to linear programming bounds on codes*, Problems Inform. Transmission **42** (2006), 77–89, doi:10.1134/S0032946006020025.
- [Bas65] L. A. Bassalygo, *New upper bounds for error correcting codes*, Problems Inform. Transmission **1** (1965), no. 4, 32–35.
- [BCV23] P.-A. Bernard, N. Crampé, and L. Vinet, *Entanglement of free fermions on Johnson graphs*, J. Math. Phys. **64** (2023), no. 6, article 061903, doi:10.1063/5.0099879.
- [CST08] T. Ceccherini-Silberstein, F. Scarabotti, and F. Tolli, *Harmonic Analysis on Finite Groups: Representation Theory, Gelfand Pairs and Markov Chains*, Cambridge Studies in Advanced Mathematics, vol. 108, Cambridge University Press, Cambridge, 2008, doi:10.1017/CBO9780511619823.
- [CD25] A. Chailloux and T. Debris-Alazard, *New solutions to Delsarte’s dual linear programs*, IEEE Trans. Inform. Theory **71** (2025), no. 1, 297–316, doi:10.1109/TIT.2024.3476974.
- [CE03] H. Cohn and N. Elkies, *New upper bounds on sphere packings. I*, Ann. of Math. (2) **157** (2003), 689–714, doi:10.4007/annals.2003.157.689.
- [CKMRV17] H. Cohn, A. Kumar, S. D. Miller, D. Radchenko, and M. S. Viazovska, *The sphere packing problem in dimension 24*, Ann. of Math. (2) **185** (2017), 1017–1033, doi:10.4007/annals.2017.185.3.8.
- [CZ14] H. Cohn and Y. Zhao, *Sphere packing bounds via spherical codes*, Duke Math. J. **163** (2014), 1965–2002, doi:10.1215/00127094-2738857.
- [CJJ22] L. N. Coregliano, F. G. Jeronimo, and C. Jones, *A complete linear programming hierarchy for linear codes*, in *13th Innovations in Theoretical Computer Science Conference*, LIPIcs **215** (2022), article 51, 51:1–51:22, doi:10.4230/LIPIcs.ITCS.2022.51.
- [CJJLL26] L. N. Coregliano, F. G. Jeronimo, C. Jones, N. Linial, and E. Loyfer, *Higher-order Delsarte dual LPs: lifting, constructions and completeness*, in *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, LIPIcs **362** (2026), article 44, 44:1–44:22, doi:10.4230/LIPIcs.ITCS.2026.44.
- [Del72] P. Delsarte, *Bounds for unrestricted codes, by linear programming*, Philips Res. Rep. **27** (1972), 272–289.
- [Del73] P. Delsarte, *An algebraic approach to the association schemes of coding theory*, Philips Res. Rep. Suppl. **10** (1973), vi+97 pp.
- [DGS77] P. Delsarte, J. M. Goethals, and J. J. Seidel, *Spherical codes and designs*, Geom. Dedicata **6** (1977), 363–388, doi:10.1007/BF03187604.
- [Fei12] P. Feinsilver, *Representations of $\mathfrak{sl}(2)$ in the Boolean lattice, and the Hamming and Johnson schemes*, Infin. Dimens. Anal. Quantum Probab. Relat. Top. **15** (2012), no. 3, article 1250019, doi:10.1142/S0219025712500191.
- [FT05] J. Friedman and J.-P. Tillich, *Generalized Alon–Boppana theorems and error-correcting codes*, SIAM J. Discrete Math. **19** (2005), 700–718, doi:10.1137/S0895480102408353.
- [GW09] R. Goodman and N. R. Wallach, *Symmetry, Representations, and Invariants*, Graduate Texts in Mathematics, vol. 255, Springer, New York, 2009, doi:10.1007/978-0-387-79852-3.
- [Gor00] D. V. Gorbachev, *Extremal problem for entire functions of exponential spherical type, connected with the Levenshtein bound on the sphere packing density in $\mathbb{R}^n$* , Izv. Tula State Univ. Ser. Math. Mech. Inform. **6** (2000), 71–78; in Russian.
- [KL78] G. A. Kabatianskii and V. I. Levenshtein, *On bounds for packings on a sphere and in space*, Problems Inform. Transmission **14** (1978), 1–17.
- [Koo81] T. H. Koornwinder, *Clebsch–Gordan coefficients for $SU(2)$ and Hahn polynomials*, Nieuw Arch. Wisk. (3) **29** (1981), 140–155.
- [Kra76] M. Krämer, *Multiplicity free subgroups of compact connected Lie groups*, Arch. Math. (Basel) **27** (1976), 28–36, doi:10.1007/BF01224637.
- [Kra18] P. Kravchuk, *Casimir recursion relations for general conformal blocks*, J. High Energy Phys. **2018**, no. 2, article 011, doi:10.1007/JHEP02(2018)011.
- [Lev79] V. I. Levenshtein, *On bounds for packings in n -dimensional Euclidean space*, Soviet Math. Dokl. **20** (1979), 417–421.
- [LL23] E. Loyfer and N. Linial, *New LP-based upper bounds in the rate-vs.-distance problem for binary linear codes*, IEEE Trans. Inform. Theory **69** (2023), no. 5, 2886–2899, doi:10.1109/TIT.2023.3236660.
- [Mac63] J. MacWilliams, *A theorem on the distribution of weights in a systematic code*, Bell Syst. Tech. J. **42** (1963), 79–94, doi:10.1002/j.1538-7305.1963.tb04003.x.
- [MRRW77] R. J. McEliece, E. R. Rodemich, H. C. Rumsey, Jr., and L. R. Welch, *New upper bounds on the rate of a code via the Delsarte–MacWilliams inequalities*, IEEE Trans. Inform. Theory **23** (1977), 157–166, doi:10.1109/TIT.1977.1055688.

- [Mus08] O. R. Musin, *The kissing number in four dimensions*, Ann. of Math. (2) **168** (2008), 1–32, [doi:10.4007/annals.2008.168.1](https://doi.org/10.4007/annals.2008.168.1).
- [NS05] M. Navon and A. Samorodnitsky, *On Delsarte’s linear programming bounds for binary codes*, in *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005)*, 2005, 327–336, [doi:10.1109/SFCS.2005.55](https://doi.org/10.1109/SFCS.2005.55).
- [OS79] A. M. Odlyzko and N. J. A. Sloane, *New bounds on the number of unit spheres that can touch a unit sphere in n dimensions*, J. Combin. Theory Ser. A **26** (1979), 210–214, [doi:10.1016/0097-3165\(79\)90074-8](https://doi.org/10.1016/0097-3165(79)90074-8).
- [PMP23] J. C.-J. Pang, H. MahdaviFar, and S. S. Pradhan, *New bounds on the size of binary codes with large minimum distance*, IEEE J. Sel. Areas Inform. Theory **4** (2023), 219–231, [doi:10.1109/JSAIT.2023.3295836](https://doi.org/10.1109/JSAIT.2023.3295836).
- [Sam01] A. Samorodnitsky, *On the optimum of Delsarte’s linear program*, J. Combin. Theory Ser. A **96** (2001), 261–287, [doi:10.1006/jcta.2001.3176](https://doi.org/10.1006/jcta.2001.3176).
- [Sam04] A. Samorodnitsky, *On linear programming bounds for spherical codes and designs*, Discrete Comput. Geom. **31** (2004), 385–394, [doi:10.1007/s00454-003-2858-0](https://doi.org/10.1007/s00454-003-2858-0).
- [Sam25] A. Samorodnitsky, *On the difficulty to beat the first linear programming bound for binary codes*, IEEE Trans. Inform. Theory **71** (2025), no. 4, 2383–2388, [doi:10.1109/TIT.2024.3504268](https://doi.org/10.1109/TIT.2024.3504268).
- [SZ24] N. T. Sardari and M. Zargar, *New upper bounds for spherical codes and packings*, Math. Ann. **389** (2024), 3653–3703, [doi:10.1007/s00208-023-02738-z](https://doi.org/10.1007/s00208-023-02738-z).
- [S42] I. J. Schoenberg, *Positive definite functions on spheres*, Duke Math. J. **9** (1942), 96–108, [doi:10.1215/S0012-7094-42-00908-6](https://doi.org/10.1215/S0012-7094-42-00908-6).
- [Sch05] A. Schrijver, *New code upper bounds from the Terwilliger algebra and semidefinite programming*, IEEE Trans. Inform. Theory **51** (2005), 2859–2866, [doi:10.1109/TIT.2005.851748](https://doi.org/10.1109/TIT.2005.851748).
- [SvdW53] K. Schütte and B. L. van der Waerden, *Das Problem der dreizehn Kugeln*, Math. Ann. **125** (1953), 325–334, [doi:10.1007/BF01343127](https://doi.org/10.1007/BF01343127).
- [Sid74] V. M. Sidel’nikov, *New bounds for densest packing of spheres in n -dimensional Euclidean space*, Math. USSR-Sb. **24** (1974), no. 1, 147–157, [doi:10.1070/SM1974v024n01ABEH001911](https://doi.org/10.1070/SM1974v024n01ABEH001911).
- [Sri11] M. K. Srinivasan, *Symmetric chains, Gelfand–Tsetlin chains, and the Terwilliger algebra of the binary Hamming scheme*, J. Algebraic Combin. **34** (2011), 301–322, [doi:10.1007/s10801-010-0272-2](https://doi.org/10.1007/s10801-010-0272-2).
- [Via17] M. S. Viazovska, *The sphere packing problem in dimension 8*, Ann. of Math. (2) **185** (2017), 991–1015, [doi:10.4007/annals.2017.185.3.7](https://doi.org/10.4007/annals.2017.185.3.7).
- [Z24] M. Zargar, *Stiefel manifolds and upper bounds for spherical codes and packings*, 2024, [arXiv: 2407.10697](https://arxiv.org/abs/2407.10697).

CHAPTER 3

A Counterexample to the Soficity Conjecture

ABSTRACT. A countable group is sofic if finite pieces of its multiplication table can be faithfully approximated by permutations of finite sets. We prove that the unit group $L_{\mathbb{F}_2}(1, 2)^\times$ of the binary Leavitt algebra is not sofic, disproving the soficity conjecture. The proof starts from Kun’s expander decomposition for property-(T) groups and the Kun–Thom centralizer obstruction. A general expander-matching argument recovers a single expanding component from a union of expanders. Elementary groups over the Leavitt algebra then force Thompson’s group V to be locally embeddable into finite groups, a contradiction.

CONTENTS

1. Introduction
 2. An expander-matching criterion
 3. The binary Leavitt configuration
- References

1. INTRODUCTION

A countable group is *sofic* if every finite portion of its multiplication table can be approximated by permutations of a finite set: multiplication must hold at almost every point, and every nonidentity element must move almost every point. Gromov introduced this approximation property in his work on symbolic dynamics [Gro99]; Weiss subsequently named sofic groups and asked whether a nonsofic group exists [Wei00, p. 351]. The question became known as the *soficity conjecture*: is every countable group sofic [Pes08, Open Question 3.8]?

Let $\mathbb{F}_2$ be the field with two elements. The noncommutative binary Leavitt algebra [Lea62] is

$$R = L_{\mathbb{F}_2}(1, 2) = \mathbb{F}_2\langle s_0, s_1, t_0, t_1 \mid t_i s_j = \delta_{ij}, s_0 t_0 + s_1 t_1 = 1 \rangle. \quad (1)$$

For example, $t_0 s_0 = 1$, whereas $s_0 t_0$ is a proper idempotent. Write $R^\times$ for its group of invertible elements: $x \in R^\times$ means that some $y \in R$ satisfies $xy = yx = 1$. Since R is finitely generated over the finite field $\mathbb{F}_2$, both R and $R^\times$ are countable.

Theorem 1.1. *The unit group $L_{\mathbb{F}_2}(1, 2)^\times$ is not sofic.*

Related work

Earlier work gave several conditional routes to nonsofic groups. Bowen and Burton showed that flexible permutation stability of $\text{PSL}_d(\mathbb{Z})$ for some $d \geq 5$ would produce such a group [BB20]. Gohla and Thom showed that suitable central extensions of p -adic lattices would be nonsofic under a permutation-stability hypothesis; Chapman, Dikstein, and Lubotzky gave an algebraic and combinatorial treatment of this implication [GT24, CDL24]. Theorem 1.1 requires no unproved stability hypothesis.

The Aldous–Lyons conjecture asked whether every unimodular random rooted network is a local weak limit of finite networks [AL07, Question 10.1]. It was disproved in companion papers by Bowen, Chapman, Lubotzky, and Vidick and by Bowen, Chapman, and Vidick [BCLV24, BCV25]. Their tour-de-force argument combines subgroup tests with a refinement of the compression techniques behind the breakthrough $\text{MIP}^* = \text{RE}$ of Ji, Natarajan, Vidick, Wright, and Yuen, which also disproved Connes’s embedding conjecture [JNVWY20]. The second companion paper also gives another proof that Connes’s embedding conjecture is false [BCV25]. To explain the relationship with soficity, let F_d be a free group of finite rank $d \geq 2$. An *invariant random subgroup* of F_d is a conjugation-invariant probability distribution on its subgroups. It is *co-sofic* if it is a weak limit of the stabilizer distributions associated with uniformly chosen points of finite F_d -sets. The Aldous–Lyons conjecture was equivalent to the assertion that, for every finite rank $d \geq 2$, every invariant random subgroup of F_d is co-sofic [Gel18, Section 6].

The soficity conjecture is precisely the restriction of this assertion to invariant random subgroups concentrated on a single normal subgroup. Indeed, if $N \triangleleft F_d$ and δ_N denotes the point mass at N , then

$$F_d/N \text{ is sofic} \iff \delta_N \text{ is a co-sofic invariant random subgroup of } F_d;$$

see [Gel18, Section 6, Proposition 6.1]. Every finitely generated group is a quotient of some F_d , and soficity is detected on finitely generated subgroups. Therefore, the soficity conjecture asks exactly whether every such δ_N is co-sofic.

Our proof constructs a finitely generated nonsofic subgroup $G \leq R^\times$. For any surjection $F_d \twoheadrightarrow G$, its kernel N therefore gives a non-co-sofic point mass δ_N . Equivalently, the generator-marked Cayley diagram of G is a deterministic unimodular network that is not a local weak limit of finite networks with the same markings. This makes no assertion about the underlying unmarked Cayley graph. In contrast, a general non-co-sofic invariant random subgroup need not be supported on a normal subgroup and therefore need not produce a nonsofic quotient.

The original motivation for soficity was Gottschalk’s surjunctivity conjecture. A group H is *surjunctive* if, for every finite alphabet A , every injective H -equivariant continuous map $A^H \rightarrow A^H$ is surjective. Gromov and Weiss proved that every sofic group is surjunctive [Gro99, Wei00]. Bowen and Chapman constructed a surjunctive non-co-sofic invariant random subgroup [BC25], but their example is not concentrated on a normal subgroup and therefore does not produce a surjunctive nonsofic group. We do not know whether $R^\times$ is surjunctive. A positive answer

would produce a surjunctive nonsofic group, while a negative answer would disprove Gottschalk's conjecture. Since surjunctivity passes to subgroups [Wei00, Lemma 1.1], a positive answer would also establish surjunctivity of the copy $V \leq R^\times$ constructed below.

A group is *hyperlinear* if finite pieces of its multiplication table admit asymptotically faithful models by finite-dimensional unitary matrices in normalized Hilbert–Schmidt distance. The conjecture that every discrete group is hyperlinear, known as Connes's embedding conjecture for groups, remains a major open problem [Pes08, Open Question 3.9]: no nonhyperlinear discrete group is known. For permutations σ, τ of n points, their permutation matrices satisfy $\|P_\sigma - P_\tau\|_{2,n}^2 = 2d_H(\sigma, \tau)$, where $\|\cdot\|_{2,n}$ is the normalized Hilbert–Schmidt norm and d_H is normalized Hamming distance. Thus every sofic group is hyperlinear [Pes08, Theorem 3.3], but Theorem 1.1 does not determine whether $L_{\mathbb{F}_2}(1, 2)^\times$ is hyperlinear. Hyperlinearity is equivalent to Connes embeddability of the group von Neumann algebra [Pes08, Theorem 8.5]. Although $\text{MIP}^* = \text{RE}$ disproved Connes's embedding conjecture for general von Neumann algebras [JNVWY20], it does not produce a nonembeddable group von Neumann algebra.

The corresponding approximation problem for the unnormalized Frobenius norm has a different answer. De Chiffre, Glebsky, Lubotzky, and Thom [DCGLT20] constructed central extensions of arithmetic lattices that are not approximable by finite-dimensional unitaries in this norm. Because hyperlinearity uses the normalized Hilbert–Schmidt norm, their examples do not produce a nonhyperlinear group.

There are also conditional approaches to nonhyperlinear groups. Dogon showed that flexible Hilbert–Schmidt stability of certain property-(T) groups would make nonsplit central extensions nonhyperlinear; his examples include extensions of $\text{Sp}_{2g}(\mathbb{Z})$ with $g \geq 2$ [Dog23]. Dogon and Vigdorovich showed that if $\text{SL}_2(\mathbb{Z}[1/p])$ is flexibly Hilbert–Schmidt stable for some prime p , then a finite central extension of that group is nonhyperlinear [DV26]. Both stability hypotheses remain open.

Two further conjectures hold for sofic groups. Lück's determinant conjecture asserts that the modified Fuglede–Kadison determinant of every matrix over $\mathbb{Z}[H]$ is at least one [ES05]. The algebraic eigenvalue conjecture asserts that every eigenvalue of such a matrix acting on $\ell^2(H)^n$ is an algebraic integer [Tho08]. Neither is settled for $R^\times$. Manzoor constructed an invariant random subgroup satisfying the determinant conjecture that is not co-hyperlinear [Man25], but this does not determine the conjecture for $R^\times$. The Kervaire–Laudenbach conjecture asks whether every one-variable equation with coefficients in a group and nonzero total exponent of the variable has a solution in some larger group. It holds for every hyperlinear group [NT22, Theorem 1.3]. Consequently, it would hold for $R^\times$ if that group were hyperlinear, whereas its failure would exhibit a nonhyperlinear group.

Kaplansky's direct-finiteness conjecture asks whether $ab = 1$ in a group algebra always implies $ba = 1$. The superficially similar relation $t_0 s_0 = 1 \neq s_0 t_0$ does not disprove this conjecture: neither s_0 nor t_0 belongs to $R^\times$, so their relation holds in R , not in $\mathbb{F}_2[R^\times]$. For every group H , the rings $\mathbb{Z}[H]$ and $k[H]$ for fields k of characteristic zero are already stably finite [BFF24, Theorem 3.4 and Corollary 3.5]. Thus only positive-characteristic fields remain open for $R^\times$. If $R^\times$ were surjunctive, then $k[R^\times]$ would be stably finite for every field [BFF24, Corollary 3.25].

Binary self-similarity

The defining relations of R can be written as rectangular-matrix identities:

$$S = \begin{bmatrix} s_0 & s_1 \end{bmatrix}, \quad T = \begin{bmatrix} t_0 \\ t_1 \end{bmatrix}, \quad ST = 1, \quad TS = I_2.$$

Thus $S : R^2 \rightarrow R$ and $T : R \rightarrow R^2$ are inverse right R -module maps, and

$$R \cong R^2 \quad \text{as right } R\text{-modules.}$$

This is an isomorphism of modules, not of unital rings.

The same relations connect the Leavitt algebra with Thompson's group. The Cuntz algebra $\mathcal{O}_2$ is generated by two isometries and their adjoints subject to the same formal relations as s_i, t_i . It is the universal C^* -completion of the complex Leavitt $*$ -algebra $L_{\mathbb{C}}(1, 2)$; our algebra R has the same relations in characteristic two. Replacing initial binary words realizes Thompson's group V

both in the unitary group of $\mathcal{O}_2$ [HO17, Section 4] and in the unit group of the binary Leavitt algebra over any field [BS16, Section 2.2].

For a finite binary word a , let $[a]$ denote the cylinder consisting of infinite binary strings beginning with a . A *complete prefix code* is a finite set of pairwise prefix-incomparable words whose cylinders partition all infinite binary strings. Repeatedly splitting a cylinder into its two children produces complete prefix codes with any prescribed number $n \geq 1$ of words. After ordering its n words, every such code E determines a unital ring isomorphism

$$\Theta_E : M_n(R) \xrightarrow{\sim} R. \quad (2)$$

Write $\text{EL}_n(R)$ for the subgroup of $\text{GL}_n(R)$ generated by the elementary matrices $I_n + rE_{ij}$, where $i \neq j$ and $r \in R$. Its image

$$\text{EL}_E(R) := \Theta_E(\text{EL}_n(R)) \leq R^\times$$

is the same elementary group realized inside $R^\times$ through the binary prefix decomposition. Since R is a finitely generated ring, the Ershov–Jaikin–Zapirain theorem gives these elementary groups property (T) when $n \geq 3$ [EJ10, Theorem 1.1]. For the specific nine-word code D constructed in Equation (14), put

$$G := \text{EL}_D(R) = \Theta_D(\text{EL}_9(R)) \leq R^\times.$$

Hence $G \cong \text{EL}_9(R)$. The number nine is useful for the later prefix construction; the matrix isomorphism (2) holds in every rank. Because soficity passes to subgroups, it suffices to prove that G is not sofic. In fact, the stronger identity $G = R^\times$ is independently known [KT26, Corollary 4.4], but the proof does not require it.

Proof overview

A family of finite graphs of uniformly bounded degree is an *expander family* if there exists $\gamma > 0$ such that, in every graph, each vertex set U containing at most half the vertices has at least $\gamma|U|$ edges to its complement. Given permutations modeling a finite generating set, their *generator graph* joins each point of the model set to its image under each generator. Kazhdan’s property (T) is a uniform spectral-gap condition on unitary representations. Kun’s theorem says that the generator graph of a sofic approximation to a property-(T) group becomes a disjoint union of uniformly expanding graphs after changing a proportion of edges that tends to zero along the approximation [Kun19, Theorem 1]. The expansion constant is uniform, but the number of components can grow without bound: taking increasingly many disjoint copies preserves a sofic approximation. The Kun–Thom obstruction requires a stronger hypothesis: if K has property (T), J is finitely generated, and a sofic approximation of $K \times J$ has a *single* uniformly expanding K -generator graph on the whole approximation set, then J is *locally embeddable into finite groups*, or LEF [KT19, Theorem 1.1]. This means that every finite portion of the multiplication table of J embeds exactly into a finite group.

The single-expander hypothesis in the Kun–Thom theorem cannot be replaced by an arbitrary union of expanders: a commuting group may move between components without being LEF. For example, set

$$\Lambda = \text{SL}_3(\mathbb{Z}), \quad B = \text{BS}(2, 3) = \langle a, b \mid ab^2a^{-1} = b^3 \rangle.$$

The group Λ has property (T) and is residually finite, hence sofic. The group B is finitely presented and sofic but not residually finite [Pes08, Example 4.6]; hence it is not LEF [VG97, Theorem 2.2]. Nevertheless, $\Lambda \times B$ is sofic [ES06, Theorem 1(1)]. Thus the expanding Λ -components alone cannot force their commuting group B to be LEF.

Proposition 2.3 bridges this gap without assuming that Kun’s decomposition has only one component. It is a general group-theoretic criterion; the Leavitt algebra enters later, when we construct an example satisfying its hypotheses. Suppose that G and its subgroup $\Gamma \leq G$ both have property (T), and that

$$G = \langle \Gamma, t_1, \dots, t_m \rangle, \quad t_i \Gamma t_i^{-1} \leq \Gamma \quad (1 \leq i \leq m).$$

Suppose further that a finitely generated subgroup $J \leq G$ satisfies

$$[\Gamma, J] = 1, \quad \Gamma \cap J = \{1\}, \quad t_1 J t_1^{-1} \leq \Gamma.$$

Thus $\Gamma \times J$ embeds in G , and the same conjugation moves both commuting factors inside Γ :

$$t_1(\Gamma \times J)t_1^{-1} = (t_1\Gamma t_1^{-1}) \times (t_1Jt_1^{-1}) \leq \Gamma.$$

This additional nesting is absent from the direct-product example above. Under these assumptions, [Proposition 2.3](#) proves that soficity of G would force J to be LEF.

To prove this implication, suppose that G has a sofic approximation on a finite set Y . Apply Kun's theorem separately to the generator graphs for Γ and G , obtaining two generally different partitions of Y into expanding components. Call their respective parts Γ -components and G -components. For each i , let p_i be the permutation approximating t_i . Since $t_i\Gamma t_i^{-1} \leq \Gamma$, expansion implies that, apart from components containing $o(|Y|)$ vertices altogether, each transported component $p_i(C)$ lies almost entirely inside some Γ -component D . Different transported components may, however, initially select the same D .

For $z \in Y$, let $C(z)$ be its Γ -component and define $M(z) = |C(z)|$. If A is a G -component, let m_A be a median of the values $M(z)$ for $z \in A$, and define

$$f(z) = \frac{M(z)}{M(z) + m_A}, \quad z \in A.$$

The inclusions $t_i\Gamma t_i^{-1} \leq \Gamma$ imply that f is almost nondecreasing along each p_i , while the Γ -generators almost preserve f . Since every generator acts by a permutation, the total increase of f equals its total decrease. Expansion of the G -components therefore forces f to remain close to its median $1/2$ outside $o(|Y|)$ vertices. It follows that a transported component $p_i(C)$ and its target D have approximately the same size. Since $p_i(C)$ already lies almost entirely in D , it occupies more than half of D . Distinct transported components are disjoint and therefore cannot select the same target.

For $i = 1$, this injective matching makes every fixed Γ -word preserve the transported components $p_1(C)$ at almost every vertex. Since $t_1Jt_1^{-1} \leq \Gamma$, each J -generator is the conjugate by t_1^{-1} of an element of Γ . Conjugating the component-preservation statement by p_1^{-1} therefore shows that the J -generators preserve the original Γ -components at almost every vertex. The Γ -generators already preserve these components. We can consequently select one original expanding component on which the required multiplication, commutation, and distinctness tests all hold outside a negligible set. Restrict the Γ - and J -generator actions to that component, complete their partial permutations, discard a negligible set to restore uniform expansion, and complete the permutations once more. The result is a sofic approximation of $\Gamma \times J$ on a single expander, so the Kun–Thom theorem implies that J is LEF.

It remains to realize the hypotheses of [Proposition 2.3](#) inside the concrete group $G = \text{EL}_D(R)$. In [Section 3](#), we verify that G has property (T), construct a property-(T) subgroup $\Gamma \leq G$, two units $u, v \in G$, and a subgroup $J \leq G$, and prove that

$$\begin{aligned} G &= \langle \Gamma, u, v \rangle, & u\Gamma u^{-1}, v\Gamma v^{-1} &\leq \Gamma, \\ \Gamma \times J &\leq G, & uJu^{-1} &\leq \Gamma, & J &\cong V. \end{aligned}$$

Here Γ acts inside the cylinder $[0]$, whereas $J \cong V$ acts inside the disjoint cylinder $[1000]$; conjugation by u moves J into the cylinder $[0001] \subseteq [0]$. Thompson's group V is finitely presented, infinite, and simple [[CFP96](#)], so it is not LEF [[VG97](#), Theorem 2.2]. If G were sofic, [Proposition 2.3](#) applied with $t_1 = u$ and $t_2 = v$ would nevertheless force J to be LEF. This contradiction proves that G , and hence $R^\times$, is not sofic.

2. AN EXPANDER-MATCHING CRITERION

We first isolate the part of the proof that does not depend on the binary Leavitt algebra. Its purpose is to turn a sofic approximation consisting of many expanding components into an approximation of a commuting direct product on one expanding component.

For a nonempty finite set Y , write

$$d_H(p, q) = \frac{|\{z \in Y : pz \neq qz\}|}{|Y|} \quad (p, q \in \text{Sym}(Y)).$$

A *sofic approximation* of a countable group H consists of maps $p_n : H \rightarrow \text{Sym}(Y_n)$ satisfying

$$p_n(1) = 1, \quad d_H(p_n(gh), p_n(g)p_n(h)) \rightarrow 0, \quad d_H(p_n(g), 1) \rightarrow 1 \quad (g \neq 1).$$

Here the multiplication condition holds for every $g, h \in H$. By taking disjoint copies, we may assume that $|Y_n| \rightarrow \infty$.

A group J is *locally embeddable into finite groups*, or *LEF*, if, for every finite $F \subseteq J$, there are a finite group B and an injection $\phi : F \rightarrow B$ such that

$$\phi(xy) = \phi(x)\phi(y) \quad (x, y, xy \in F).$$

Thus the multiplication table of F embeds exactly, although the finite group B may depend on F .

For a graph L and $U \subseteq V(L)$, let $\partial_L U$ be the set of edges joining U to $V(L) \setminus U$. If C is a connected component of L , we also write $\partial_C U$ for the boundary of $U \subseteq C$ in the induced graph on C . Parallel edges are counted with multiplicity, while loops contribute nothing. Edge symmetric differences are likewise counted with multiplicity. We call C a γ -*expander* if

$$|\partial_C U| \geq \gamma \min\{|U|, |C \setminus U|\} \quad (U \subseteq C). \quad (3)$$

Throughout this section, a finite symmetric generating set $S = S^{-1} \subseteq H$ includes the identity. Normalize inverse labels to be inverse permutations, with the identity label acting exactly as the identity. The corresponding S -labelled *generator graph* has vertex set Y_n and an arc from z to $p_n(s)z$ for every $s \in S$. Pair inverse arcs to obtain an undirected multigraph; fixed points, including the identity label, give loops. Parallel edges retain their multiplicities. All such graphs have uniformly bounded degree. Adding identity loops does not change their edge boundaries, but makes the associated random walk lazy, as in Kun's property-(T) argument.

The two external inputs differ in a crucial respect: Kun's theorem produces possibly many expanders; the Kun–Thom theorem requires one expander on the entire approximation set.

Theorem 2.1 (Kun's expander decomposition). *Let H be an infinite finitely generated group with property (T), let $S = S^{-1} \subseteq H$ be a finite generating set containing 1, and let $p_n : H \rightarrow \text{Sym}(Y_n)$ be a sofic approximation with $|Y_n| \rightarrow \infty$. Write X_n for its S -generator multigraph. There are a constant $\gamma > 0$ and unlabelled, uniformly bounded-degree multigraphs L_n on the same vertex sets such that*

$$|E(X_n) \triangle E(L_n)| = o(|Y_n|)$$

and every connected component of L_n is a γ -expander.

This is the expander-decomposition conclusion of [Kun19, Theorem 1], stated with the identity-inclusive generator and multigraph conventions used in its property-(T) argument. If a graph convention suppresses identity loops, adjoin the same identity loop at every vertex of both graphs; this changes neither their boundaries nor their edge symmetric difference. The reference graphs are unlabelled: their edges need not retain the original generator labels, and decompositions for different generating families need not be compatible.

Theorem 2.2 (Kun–Thom's expander-centralizer theorem). *Let K be a group with property (T), let J be a finitely generated group, and fix a finite symmetric generating set $S_K \subseteq K$ of K containing 1. Suppose that $K \times J$ has a sofic approximation $p_n : K \times J \rightarrow \text{Sym}(Y_n)$. Write $X_{K,n}$ for its S_K -generator multigraph, and suppose that, for some $\lambda > 0$,*

$$|\partial_{X_{K,n}} U| \geq \lambda \min\{|U|, |Y_n \setminus U|\} \quad (U \subseteq Y_n)$$

for every n . Then J is LEF.

We use the permutation-multigraph version of [KT19, Theorem 1.1 and Section 4]; its boundary counts repeated edges with multiplicity.

Proposition 2.3. *Let $\Gamma \leq G$ be infinite finitely generated groups with property (T). Suppose that*

$$G = \langle \Gamma, t_1, \dots, t_m \rangle, \quad t_i \Gamma t_i^{-1} \leq \Gamma \quad (1 \leq i \leq m), \quad m \geq 1.$$

If a finitely generated subgroup $J \leq G$ satisfies

$$[\Gamma, J] = 1, \quad \Gamma \cap J = \{1\}, \quad t_1 J t_1^{-1} \leq \Gamma,$$

then soficity of G implies that J is LEF.

Proof. Assume that G is sofic, and fix one approximation $p_n : G \rightarrow \text{Sym}(Y_n)$. By disjoint amplification, arrange that

$$N = |Y_n| \longrightarrow \infty.$$

All $o(N)$ estimates refer to this same approximation. For readability, write $p_g = p_n(g)$ and suppress n from the graphs and partitions.

The hypotheses identify ΓJ with $\Gamma \times J$ and imply $t_1(\Gamma J)t_1^{-1} \leq \Gamma$. Our goal is to extract a sofic approximation of $\Gamma \times J$ whose Γ -generator graph is one expander. We first compare the expander decompositions for Γ and G , then show that J almost preserves the Γ -components. Finally, we select one such component and repair the restricted generator actions.

Step 1: Locate a dominant original component inside each transported component.

Choose a finite symmetric generating set S_Γ for Γ containing 1. Adjoin the distinct nonidentity elements among $t_i^{\pm 1}$ to obtain a symmetric generating set S_G for G . For every inverse pair added in this way, choose one of the original t_i as its positive representative. Thus each positive nonidentity G -generator either belongs to Γ or is one of the given compressing elements. Normalize inverse labels to be inverse permutations and nonidentity involutive labels to be exact involutions, allowing their $o(N)$ fixed points; this changes only $o(N)$ vertices [ES06, Lemma 2.1].

Apply Theorem 2.1 first to the restricted sofic approximation of Γ , using S_Γ , and then to the sofic approximation of G , using S_G . Denote the resulting reference graphs, component partitions, and expansion constants by

reference graph	connected components	expansion constant
L_Γ	$\mathcal{Q}$	$\gamma_\Gamma > 0$
L_G	$\mathcal{A}$	$\gamma_G > 0$.

We call members of $\mathcal{Q}$ *original components* and members of $\mathcal{A}$ *ambient components*. Both partitions concern the same vertex set, but neither is assumed to refine the other. Each reference graph differs from its corresponding generator graph in $o(N)$ edges. Since L_Γ has no edges between original components, an S_Γ -generator crosses between them on only $o(N)$ vertices. If $w = s_1 \cdots s_r$ is a fixed Γ -word, write $p_w = p_{s_1} \cdots p_{s_r}$. The same crossing estimate then holds for p_w .

Set $\tau_i = p_{t_i}$. Transport the vertices and edges of L_Γ through τ_i , obtaining

$$I_i = \tau_i L_\Gamma, \quad \mathcal{P}_i = \{\tau_i C : C \in \mathcal{Q}\}.$$

Each member of $\mathcal{P}_i$ is a γ_Γ -expanding connected component of I_i . Apart from $o(N)$ edited edges, the edges of I_i follow the permutations $\tau_i p_s \tau_i^{-1}$, $s \in S_\Gamma$. Approximate multiplicativity identifies each such permutation, outside $o(N)$ vertices, with a fixed Γ -word representing $t_i s t_i^{-1}$. Since these words almost preserve the original components, we obtain

$$\#\{I_i\text{-edges joining different members of } \mathcal{Q}\} = o(N) \quad (1 \leq i \leq m). \quad (4)$$

For $P \in \mathcal{P}_i$, choose $Q(P) \in \mathcal{Q}$ maximizing $|P \cap Q(P)|$, and define its unmatched mass by

$$L(P) = |P| - |P \cap Q(P)|.$$

Partition P into the sets $P \cap Q$, $Q \in \mathcal{Q}$. If the largest part has more than half the vertices, apply expansion to all other parts, whose total size is $L(P)$. Otherwise, every part has at most half the vertices, so apply expansion to all of them. Each edge between different parts is counted at most twice, giving

$$\gamma_\Gamma L(P) \leq 2\#\{I_i[P]\text{-edges joining different members of } \mathcal{Q}\}.$$

Together with (4), this gives $\sum_{P \in \mathcal{P}_i} L(P) = o(N)$. Choose $\eta_n \downarrow 0$ sufficiently slowly that, simultaneously for every i ,

$$\sum_{\substack{P \in \mathcal{P}_i \\ L(P) > \eta_n |P|}} |P| = o(N). \quad (5)$$

For each i , all but $o(N)$ vertices belong to components of $\mathcal{P}_i$ whose unmatched mass is at most η_n times their size. It remains to rule out several transported components selecting the same original component.

Step 2: Use a bounded median normalization to compare component sizes.

For each vertex z , let $C(z) \in \mathcal{Q}$ be its original component and set

$$M(z) = |C(z)|.$$

Suppose that $P = \tau_i C$ satisfies $L(P) \leq \eta_n |P|$. For every $z \in C$ with $\tau_i z \in P \cap Q(P)$,

$$M(\tau_i z) = |Q(P)| \geq (1 - \eta_n) |P| = (1 - \eta_n) M(z).$$

By (5), the discarded components and the unmatched portions of the remaining components occupy $o(N)$ vertices. Consequently,

$$M(\tau_i z) \geq (1 - \eta_n) M(z) \quad \text{outside } o(N) \text{ vertices} \quad (1 \leq i \leq m). \quad (6)$$

For each $s \in S_\Gamma$, we likewise have $M(p_s z) = M(z)$ outside $o(N)$ vertices.

The values of M need not be uniformly bounded, so these exceptional sets cannot control its total variation. For each ambient component $A \in \mathcal{A}$, let $m_A > 0$ be a median of the vertex-indexed multiset

$$(M(z) : z \in A).$$

Thus at least half of these values are at most m_A , and at least half are at least m_A . In particular, an original component C contributes $|C \cap A|$ copies of $|C|$. Define

$$f(z) = \frac{M(z)}{M(z) + m_A}, \quad z \in A.$$

Then $0 < f < 1$, and $1/2$ is a median of f on every A .

Every G -generator crosses between components of L_G on only $o(N)$ vertices. On each remaining generator arc, both endpoints use the same normalizing median m_A . For $s \in S_\Gamma$, this gives $f(p_s z) = f(z)$ outside $o(N)$ vertices. For a positive compressing generator t_i , use (6) and

$$\frac{(1 - \eta)x}{(1 - \eta)x + a} \geq \frac{x}{x + a} - \eta \quad (x, a > 0, 0 \leq \eta < 1).$$

Thus every positive G -generator s satisfies

$$f(p_s z) \geq f(z) - \eta_n \quad \text{outside } o(N) \text{ vertices}.$$

Because p_s is a permutation,

$$\sum_{z \in Y_n} (f(p_s z) - f(z)) = 0.$$

The total decrease of f is at most $\eta_n N + o(N) = o(N)$: the inequality above controls the nonexceptional vertices, and $0 < f < 1$ controls the exceptional ones. Since the displayed sum is zero, the total increase equals the total decrease. Hence

$$\sum_{z \in Y_n} |f(p_s z) - f(z)| = o(N).$$

The same estimate holds for inverse generators. Summing over S_G , and noting that changing one edge affects the total variation by at most one, therefore gives

$$\sum_{\{z, w\} \in E(L_G)} |f(z) - f(w)| = o(N). \quad (7)$$

Expansion now forces f to concentrate near its median $1/2$. Fix $A \in \mathcal{A}$. The finite coarea identity says that

$$\int_0^1 |\partial_A \{z \in A : f(z) > t\}| dt = \sum_{\{z, w\} \in E(L_G[A])} |f(z) - f(w)|.$$

Indeed, an edge $\{z, w\}$ crosses the displayed level set precisely for levels between $f(z)$ and $f(w)$. For $0 < t < 1/2$, the set $\{z \in A : f(z) \leq t\}$ contains at most $|A|/2$ vertices; for $1/2 < t < 1$, the same holds for $\{z \in A : f(z) > t\}$. Apply (3) to these sublevel and superlevel sets, use that complementary sets have the same boundary, and split the coarea integral at $1/2$. This yields

$$\gamma_G \sum_{z \in A} \left| f(z) - \frac{1}{2} \right| \leq \sum_{\{z, w\} \in E(L_G[A])} |f(z) - f(w)|.$$

Summing over A and using (7), we obtain

$$\sum_{z \in Y_n} \left| f(z) - \frac{1}{2} \right| = o(N).$$

Choose $\delta_n \downarrow 0$ sufficiently slowly that

$$E_n = \left\{ z : \left| f(z) - \frac{1}{2} \right| > \delta_n \right\} \quad \text{satisfies} \quad |E_n| = o(N).$$

Since

$$M(z) = m_A \frac{f(z)}{1 - f(z)} \quad (z \in A),$$

any $z, w \in A \setminus E_n$ satisfy

$$\rho_n^{-1} \leq \frac{M(w)}{M(z)} \leq \rho_n, \quad \rho_n = \left(\frac{1 + 2\delta_n}{1 - 2\delta_n} \right)^2 \rightarrow 1. \quad (8)$$

We have therefore shown that the sizes of original components meeting the same set $A \setminus E_n$ differ by a factor tending uniformly to one.

Step 3: Match the first transported partition almost bijectively.

All t_i were needed to control the full G -generator graph. To show that J preserves the original components, however, we need only the first transport. Set

$$\tau = \tau_1, \quad \mathcal{P} = \mathcal{P}_1.$$

Retain a component $P = \tau C \in \mathcal{P}$ precisely when

$$L(P) \leq \eta_n |P|$$

and there exists $z \in C$ such that

$$\tau z \in P \cap Q(P), \quad z, \tau z \notin E_n, \quad z, \tau z \text{ belong to the same } A \in \mathcal{A}. \quad (9)$$

The last condition means that both vertices have the same median m_A .

The vertices for which τ crosses between ambient components form a set

$$H_n = \{z : z \text{ and } \tau z \text{ belong to different members of } \mathcal{A}\}$$

of size $o(N)$. Indeed, t_1 is a fixed word in S_G , and each generator crosses $\mathcal{A}$ on only $o(N)$ vertices. If P satisfies the loss bound but has no vertex as in (9), then

$$\tau^{-1}(P \cap Q(P)) \subseteq H_n \cup E_n \cup \tau^{-1}(E_n).$$

The sets $\tau^{-1}(P \cap Q(P))$ are disjoint as P varies, and each has at least $(1 - \eta_n)|P|$ vertices when the loss bound holds. Thus the components satisfying the loss bound but having no witness contribute only $o(N)$ vertices. The components failing the loss bound contribute another $o(N)$ vertices by (5). Therefore all nonretained components together contain $o(N)$ vertices.

For a retained P , choose z as in (9). Then

$$M(z) = |C| = |P|, \quad M(\tau z) = |Q(P)|.$$

The size comparison (8) gives

$$\rho_n^{-1}|P| \leq |Q(P)| \leq \rho_n|P|.$$

Since at most $\eta_n|P|$ vertices of P lie outside $Q(P)$, it follows that

$$|P \triangle Q(P)| \leq (\rho_n - 1 + 2\eta_n)|P| = o(|P|).$$

In particular, for all sufficiently large n ,

$$|P \cap Q(P)| > \frac{1}{2}|Q(P)|.$$

Distinct transported components are disjoint, so they cannot both occupy a strict majority of the same original component. Therefore $P \mapsto Q(P)$ is injective on retained components. The nonretained components and unmatched portions of retained components together contain $o(N)$ vertices, so the matched intersections cover $N - o(N)$ vertices.

Fix a Γ -word w . Discard vertices outside the matched intersections, vertices whose p_w -images lie outside those intersections, and vertices for which p_w crosses between original components. Each

discarded set has size $o(N)$. For every remaining vertex z , the vertices z and $p_w z$ lie in matched intersections belonging to the same original component. Injectivity of the matching therefore puts them in the same transported component. Thus

$$\#\{z : z \text{ and } p_w z \text{ lie in different members of } \mathcal{P}\} = o(N). \quad (10)$$

For each $j \in J$, choose a fixed Γ -word w_j representing $t_1 j t_1^{-1}$, and put

$$q_j = \tau^{-1} p_{w_j} \tau.$$

Approximate multiplicativity shows that q_j agrees with p_j outside $o(N)$ vertices. Moreover, $\tau q_j z = p_{w_j} \tau z$, and the transported component containing τz is $\tau C(z)$. Hence z and $q_j z$ belong to the same original component exactly when τz and $p_{w_j} \tau z$ belong to the same transported component. Applying (10) to w_j therefore gives

$$\#\{z : C(q_j z) \neq C(z)\} = o(N). \quad (11)$$

Thus, for every fixed generator of either Γ or J , its approximating permutation preserves the original-component partition outside $o(N)$ vertices. It remains to select one original component and turn the restricted actions into an expanding approximation of $\Gamma \times J$.

Step 4: Select one component carrying all required word tests.

Choose finite symmetric generating sets

$$T_\Gamma = S_\Gamma, \quad T_J \subseteq J \setminus \{1\}, \quad T = T_\Gamma \cup T_J.$$

Thus $1 \in T_\Gamma \subseteq T$; its generator permutation is the identity. The set T_J may be empty when J is trivial. For $s \in T_\Gamma$, define

$$q_s = p_s.$$

For $j \in T_J$, use the permutation q_j constructed in Step 3. Normalize inverse labels and make involutive J -labels exact by replacing all cycles of length greater than two with fixed points. These changes affect $o(N)$ vertices and can alter any fixed word test at only $o(N)$ starting vertices. Since $[\Gamma, J] = 1$ and $\Gamma \cap J = \{1\}$, multiplication identifies $\Gamma \times J$ with its subgroup of G , and the permutations q_t satisfy its fixed equality and distinctness tests outside $o(N)$ vertices.

For $z \in Y_n$, recall that $C(z)$ is its original component. Define the set of generator exits by

$$C_n = \{z : C(q_t z) \neq C(z) \text{ for some } t \in T\}.$$

Every Γ -generator preserves the original components outside $o(N)$ vertices, while (11) applies to every J -generator. Hence $|C_n| = o(N)$. Let Δ_n consist of vertices incident to the edge-multiset difference between the Γ -generator graph and L_Γ . Then $|\Delta_n| = o(N)$.

For $\ell \geq 0$, let W_ℓ be the finite set of *formal* words in the alphabet T of length at most ℓ , including the empty word. Different formal words may represent the same group element. For $w \in W_\ell$, write q_w for the corresponding product of permutations and $\bar{w} \in \Gamma \times J$ for the represented group element. Words representing the same element should have the same endpoint, whereas words representing distinct elements should have distinct endpoints. The vertices where either requirement fails form

$$\begin{aligned} F_{n,\ell} = & \bigcup_{\substack{w, w' \in W_\ell \\ \bar{w} = \bar{w}'}} \{z : q_w z \neq q_{w'} z\} \\ & \cup \bigcup_{\substack{w, w' \in W_\ell \\ \bar{w} \neq \bar{w}'}} \{z : q_w z = q_{w'} z\}. \end{aligned}$$

For $\ell \geq 1$, collect all exceptional vertices in the explicitly defined set

$$B_{n,\ell} = \Delta_n \cup F_{n,\ell} \cup \bigcup_{w \in W_{\ell-1}} q_w^{-1}(C_n).$$

The last term records word paths for which some generator step crosses between original components. For each fixed ℓ , we have $|B_{n,\ell}| = o(N)$. Choose $\ell_n \rightarrow \infty$ sufficiently slowly that

$$e_n = \frac{|B_{n,\ell_n}|}{N} \rightarrow 0, \quad B_n = B_{n,\ell_n}.$$

Averaging over all original components with weights $|Q|$ yields a component $Q_n \in \mathcal{Q}$ satisfying

$$|Q_n \cap B_n| \leq e_n |Q_n| \quad (12)$$

for every n . For each element of the word-metric ball in Γ , with generating set T_Γ and radius $\lfloor \ell_n/2 \rfloor$, choose a formal word of that length or shorter representing it. At any $z \in Q_n \setminus B_n$, the exit tests keep every such word path inside Q_n , and the distinctness tests give different endpoints for different elements. Evaluating the chosen words at z therefore injects the entire word-metric ball into Q_n . Since Γ is infinite, we have

$$|Q_n| \rightarrow \infty.$$

Step 5: Complete the partial actions and remove the additive expansion error.

We first restrict the generator actions to the selected component and complete them to permutations. The resulting generator graph is close to an expander but may fail to expand on very small sets. Removing a small exceptional set and completing the restricted permutations a second time will produce one uniformly expanding generator graph.

Set $P = Q_n$, now an original component rather than a transported component. For one representative of every nonidentity inverse pair in T , restrict q_t to its internal arcs:

$$A_t = \{z \in P : q_t z \in P\}, \quad q_t : A_t \rightarrow q_t(A_t).$$

The omitted sets $P \setminus A_t$ and $P \setminus q_t(A_t)$ have the same cardinality. Choose a bijection between them to extend this partial bijection to a permutation $\hat{q}_t \in \text{Sym}(P)$. For an involution, the missing domain equals the missing range, and we complete it by fixed points. Use $\hat{q}_t^{-1}$ for the inverse label and $\hat{q}_1 = \text{id}_P$ for the identity label. Because

$$P \setminus A_t \subseteq C_n \cap P \subseteq B_n \cap P,$$

(12) shows that only $O(e_n |P|)$ values are changed. Every word path of length at most ℓ_n starting in $P \setminus B_n$ remains in P , so its equality, commutation, and distinctness tests are preserved. To obtain maps on the whole group, choose once and for all a formal representative word w_g for every $g \in \Gamma \times J$, with w_1 empty and $w_t = t$ for every nonidentity $t \in T$, and set

$$\hat{p}_n(g) = \hat{q}_{w_g} \in \text{Sym}(Q_n).$$

For fixed g, h , the equality tests compare w_{gh} with the concatenation $w_g w_h$, and the distinctness tests compare w_g with the empty word when $g \neq 1$. Since $\ell_n \rightarrow \infty$, these tests eventually apply, and their exceptional proportions tend to zero. Therefore $\hat{p}_n$ is a sofic approximation of $\Gamma \times J$ on Q_n .

Let I_0 be the T_Γ -generator multigraph of $\hat{p}_n$. Since $\Delta_n \cup C_n \subseteq B_n$,

$$|E(I_0) \triangle E(L_\Gamma[P])| = O_{|S_\Gamma|}(e_n |P|).$$

Here $L_\Gamma[P]$ is precisely the original γ_Γ -expanding component on P . Consequently, for some $a_n \rightarrow 0$,

$$|\partial_{I_0} X| \geq \gamma_\Gamma \min\{|X|, |P \setminus X|\} - a_n |P| \quad (X \subseteq P). \quad (13)$$

The additive error does not control small sets, so the current graph need not yet be an expander.

Fix $0 < \lambda < \gamma_\Gamma$. If there is a nonempty $B_{\text{cut}} \subseteq P$ with

$$|B_{\text{cut}}| \leq \frac{|P|}{2}, \quad |\partial_{I_0} B_{\text{cut}}| < \lambda |B_{\text{cut}}|,$$

choose such a set of maximum cardinality; otherwise, set $B_{\text{cut}} = \emptyset$. Equation (13) implies

$$|B_{\text{cut}}| \leq \frac{a_n}{\gamma_\Gamma - \lambda} |P| = o(|P|).$$

Put $Z = P \setminus B_{\text{cut}}$. We claim that, for all sufficiently large n , the induced graph $I_0[Z]$ is a λ -expander.

Let $X \subseteq Z$ satisfy $0 < |X| \leq |Z|/2$. If $|X| + |B_{\text{cut}}| \leq |P|/2$, maximality of B_{cut} gives

$$\lambda(|X| + |B_{\text{cut}}|) \leq |\partial_{I_0}(X \cup B_{\text{cut}})| \leq |\partial_{I_0[Z]} X| + |\partial_{I_0} B_{\text{cut}}|.$$

It follows that $|\partial_{I_0[Z]}X| \geq \lambda|X|$. Otherwise,

$$|X| > \frac{|P|}{2} - |B_{\text{cut}}|.$$

If d bounds the degrees of I_0 , (13) yields

$$|\partial_{I_0[Z]}X| \geq \gamma_\Gamma|X| - a_n|P| - d|B_{\text{cut}}| \geq \lambda|X|$$

for all sufficiently large n , because $|B_{\text{cut}}| = o(|P|)$ and $|X| = (\frac{1}{2} - o(1))|P|$. If B_{cut} is empty, the same conclusion is immediate from its defining maximality condition. This proves the claim.

The permutations $\hat{q}_t$ still act on P , not on Z . Restrict each one to

$$\hat{q}_t : \{z \in Z : \hat{q}_t z \in Z\} \longrightarrow \hat{q}_t(\{z \in Z : \hat{q}_t z \in Z\}).$$

The omitted domain and range inside Z again have equal cardinality, so extend this partial bijection to a permutation of Z . Again use fixed points for an involution and inverse permutations for inverse labels. The resulting $\tilde{q}_t \in \text{Sym}(Z)$ differs from the restriction of $\hat{q}_t$ on at most

$$|B_{\text{cut}}| = o(|P|) = o(|Z|)$$

vertices per generator. Every internal Γ -edge of $I_0[Z]$ is retained. Thus the Γ -generator multigraph of $(\tilde{q}_t)_{t \in T}$ contains $I_0[Z]$ as a spanning submultigraph. The additional completion edges cannot decrease any edge boundary, so this generator graph is a λ -expander on the whole of Z .

For each fixed word w , compare its path under the permutations $\hat{q}_t$ on P with its path under $\tilde{q}_t$ on Z . These paths can first differ only when a prefix encounters a repaired generator value. Since each prefix acts by a permutation,

$$|\{z \in Z : \tilde{q}_w z \neq \hat{q}_w z\}| \leq |w| |B_{\text{cut}}|.$$

Thus every fixed equality, distinctness, and commutation test still fails on only $o(|Z|)$ vertices. As

$$|Z| = (1 - o(1))|P| \longrightarrow \infty,$$

the maps $g \mapsto \tilde{q}_{w_g}$ give a sofic approximation of $\Gamma \times J$ whose Γ -generator graph is one uniform expander. Theorem 2.2 implies that J is LEF. $\square$

3. THE BINARY LEAVITT CONFIGURATION

We apply Proposition 2.3 inside the unit group of the noncommutative binary Leavitt algebra R from (1). Its binary prefix structure supplies the contracting conjugations and the commuting copy of Thompson's group, while the Ershov–Jaikin–Zapirain theorem supplies property (T) for the elementary groups used in the construction. We construct these groups and then verify the hypotheses of Proposition 2.3.

3.1. Prefix codes and elementary groups. Let W be the $\mathbb{F}_2$ -vector space with one basis vector $\mathbf{e}_\omega$ for every infinite binary string $\omega \in \{0, 1\}^\mathbb{N}$. The formulas

$$s_i \mathbf{e}_\omega = \mathbf{e}_{i\omega}, \quad t_i \mathbf{e}_{j\omega} = \delta_{ij} \mathbf{e}_\omega$$

satisfy the defining relations of R . Thus s_i prefixes a string by i , whereas t_i deletes that prefix when present and otherwise gives zero. The operators s_0^k are distinct, so R is infinite.

For a finite binary word $a = i_1 \cdots i_r$, set

$$s_a = s_{i_1} \cdots s_{i_r}, \quad t_a = t_{i_r} \cdots t_{i_1}, \quad e_a = s_a t_a.$$

For the empty word, both products are 1. The *cylinder* $[a]$ consists of the infinite binary strings beginning with a . A *prefix code* $E = (a_1, \dots, a_q)$ is a list of words none of which is a prefix of another. It is *complete* if the cylinders $[a_i]$ partition all infinite binary strings. Write $e_E = \sum_i e_{a_i}$. The prefix-code identities

$$t_{a_i} s_{a_j} = \delta_{ij}, \quad (s_{a_i} t_{a_j})(s_{a_k} t_{a_l}) = \delta_{jk} s_{a_i} t_{a_l}$$

follow from the incomparability of its words. Moreover, the Leavitt relation gives

$$e_a = s_a(s_0 t_0 + s_1 t_1) t_a = e_{a0} + e_{a1}.$$

Every finite complete prefix code is obtained from the one-word code consisting of the empty word by repeatedly replacing a word a with its children $a0, a1$. The displayed identity therefore shows that $e_E = 1$ whenever E is complete. The prefix-code identities also give a ring isomorphism

$$\Theta_E : M_q(R) \xrightarrow{\sim} e_E R e_E, \quad (r_{ij}) \mapsto \sum_{i,j} s_{a_i} r_{ij} t_{a_j}.$$

Its inverse sends $x \in e_E R e_E$ to $(t_{a_i} x s_{a_j})_{i,j}$. If E is complete, this identifies $M_q(R)$ with R . Otherwise, a unit h in the corner extends to a unit of R as $h + 1 - e_E$, acting identically outside the cylinders in E .

In particular, define the *elementary prefix group*

$$\text{EL}_E(R) = \langle 1 + s_{a_i} r t_{a_j} : i \neq j, r \in R \rangle \leq R^\times.$$

Indeed, the corner-unit extension satisfies

$$(1 - e_E) + \Theta_E(I_q + r E_{ij}) = 1 + s_{a_i} r t_{a_j} \quad (i \neq j),$$

so $\text{EL}_E(R)$ is the copy of $\text{EL}_q(R)$ acting identically outside the cylinders of E . We call its displayed generators *elementary E -roots*.

Take the three prefix blocks

$$\begin{aligned} \alpha &= (000, 001, 01), \\ \beta &= (1000, 1001, 101), \\ \nu &= (1100, 1101, 111), \\ D &= (\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3, \nu_1, \nu_2, \nu_3). \end{aligned} \tag{14}$$

Their cylinders partition $[0]$, $[10]$, and $[11]$, respectively. Therefore D is a complete nine-leaf code, whereas α covers only $[0]$. Set

$$G = \text{EL}_D(R) \cong \text{EL}_9(R), \quad \Gamma = \text{EL}_\alpha(R) \cong \text{EL}_3(R).$$

Since $e_\alpha = e_0$, these identifications give

$$\Gamma = \{1 - e_0 + \Theta_\alpha(A) : A \in \text{EL}_3(R)\} = \{\Theta_D(\text{diag}(A, I_6)) : A \in \text{EL}_3(R)\}.$$

Thus the corner action is extended by the identity on the six complementary leaves, and every elementary α -root is also a D -root. Consequently,

$$\Gamma \leq G \leq R^\times.$$

Both groups are infinite: for distinct code leaves, the map

$$r \mapsto 1 + s_{a_i} r t_{a_j}$$

embeds the infinite additive group of R into their corresponding elementary root subgroup.

The ring R is generated by $\{1, s_0, s_1, t_0, t_1\}$. For $x_{ij}(r) = I_q + r E_{ij}$, the identities

$$x_{ij}(r + r') = x_{ij}(r)x_{ij}(r'), \quad [x_{ij}(r), x_{jk}(r')] = x_{ik}(rr') \quad (i, j, k \text{ distinct})$$

show that the elementary roots with these five coefficients generate $\text{EL}_q(R)$ whenever $q \geq 3$. Thus G and Γ are finitely generated. Both have property (T) by the Ershov–Jaikin–Zapirain theorem [EJ10, Theorem 1.1].

3.2. Copies of Thompson’s group. To define Thompson’s group V [CFP96], take two complete prefix codes

$$E = (a_1, \dots, a_q), \quad E' = (b_1, \dots, b_q),$$

together with a bijection $a_i \mapsto b_i$. Every infinite binary string has a unique expression $a_i \omega$, where $a_i \in E$ and ω is its remaining infinite tail. The corresponding *prefix replacement* acts by

$$g(a_i \omega) = b_i \omega.$$

Thus it replaces the initial word a_i by b_i and leaves the infinite tail unchanged. Since E' is also complete, this map is a bijection; its inverse replaces each b_i by a_i . Thompson’s group V consists of all such prefix replacements.

For example, the complete codes

$$E = (0, 10, 11), \quad E' = (10, 0, 11)$$

give the replacement

$$g(0\omega) = 10\omega, \quad g(10\omega) = 0\omega, \quad g(11\omega) = 11\omega.$$

It exchanges the cylinders $[0]$ and $[10]$ while fixing $[11]$.

We next realize V inside G , together with smaller copies inside Γ . In the standard Leavitt-algebra realization [BS16, Section 2.2], a prefix-replacement table $g : E \rightarrow E'$ gives the element

$$U_g = \sum_{a \in E} s_{g(a)} t_a, \quad U_{g^{-1}} = \sum_{a \in E} s_a t_{g(a)}.$$

Because both codes are complete, the prefix-code identities give

$$U_g U_{g^{-1}} = U_{g^{-1}} U_g = 1.$$

To see that this construction depends only on the represented prefix replacement, observe that

$$s_b t_a = s_{b0} t_{a0} + s_{b1} t_{a1}.$$

Thus replacing one table entry $a \mapsto b$ by the two entries $a0 \mapsto b0$ and $a1 \mapsto b1$ does not change U_g . Any two tables representing the same prefix replacement have a common refinement. Likewise, given two replacements g and h , refine the range code of g and the domain code of h until they agree. The prefix-code identities then give

$$U_h U_g = U_{h \circ g}, \quad U_{\text{id}} = 1.$$

Hence $g \mapsto U_g$ is a well-defined homomorphism. Moreover, U_g sends $\mathbf{e}_{a\omega}$ to $\mathbf{e}_{g(a)\omega}$. Distinct prefix replacements act differently on some basis vector, so their units are distinct. We therefore obtain a faithful copy $V \leq R^\times$. For a binary word l , write $V_l \cong V$ for the copy acting by prefix replacements inside $[l]$ and fixing its complement. Its tables can be chosen to leave every complementary cylinder unchanged, so

$$g - 1 \in e_l R e_l \quad (g \in V_l). \quad (15)$$

Lemma 3.1. *We have $V \leq G$. If l extends one of the three words α_i , then $V_l \leq \Gamma$.*

Proof. For incomparable words ρ, σ , the unit

$$\tau_{\rho, \sigma} = 1 + e_\rho + e_\sigma + s_\rho t_\sigma + s_\sigma t_\rho$$

exchanges the cylinders $[\rho]$ and $[\sigma]$ and fixes their complement. These cylinder swaps generate V [BQ17, Theorem 1.1].

Fix any prefix code $E = (a_1, \dots, a_q)$ with $q \geq 2$. Suppose first that $\rho = a_i w$ and $\sigma = a_j w'$ extend different leaves of this code. Then

$$P = s_\rho t_\sigma = s_{a_i} (s_w t_{w'}) t_{a_j}, \quad Q = s_\sigma t_\rho = s_{a_j} (s_{w'} t_w) t_{a_i}$$

make $1 + P$ and $1 + Q$ elementary E -roots. Since the coefficient field has characteristic two,

$$\tau_{\rho, \sigma} = (1 + P)(1 + Q)(1 + P) \in \text{EL}_E(R).$$

If instead both words extend the same leaf a_i , choose a different leaf a_j . The transposition identity

$$\tau_{\rho, \sigma} = \tau_{\rho, a_j} \tau_{\sigma, a_j} \tau_{\rho, a_j}$$

again puts $\tau_{\rho, \sigma}$ in $\text{EL}_E(R)$.

For the complete code D , choose k large enough that every ρw and σw , $w \in \{0, 1\}^k$, extends a leaf of D . Then

$$\tau_{\rho, \sigma} = \prod_{w \in \{0, 1\}^k} \tau_{\rho w, \sigma w} \in \text{EL}_D(R).$$

Hence $V \leq G$. If l extends α_i , the same-leaf argument with $E = \alpha$ puts every cylinder swap supported on $[l]$ in Γ , proving $V_l \leq \Gamma$. $\square$

3.3. Two contractions and a commuting obstruction. We construct two prefix replacements that send Γ into the same subgroup and together recover G . The three prefixes $\zeta = (100, 101, 11)$ partition $[1]$. Define $u, v \in V \leq G$ by the tables

$$\begin{array}{c|ccc} & \alpha_i & \beta_i & \nu_i \\ \hline u & \alpha_i 0 & \alpha_i 1 & \zeta_i \\ v & \alpha_i 0 & \zeta_i & \alpha_i 1 \end{array} \quad (1 \leq i \leq 3). \quad (16)$$

Each column lists a source prefix, and an entry b below a source prefix a means that the replacement sends $a\omega$ to $b\omega$. The notation $\alpha_i 0$ and $\alpha_i 1$ means appending the indicated bit. In either row, the six target cylinders $[\alpha_i 0], [\alpha_i 1]$ partition $[0]$, and the three cylinders $[\zeta_i]$ partition $[1]$. Thus both rows are complete prefix replacements.

Both u and v send $[\alpha_i]$ to $[\alpha_i 0]$. Their other prefix images are

$$u([\beta_i]) = [\alpha_i 1], \quad u([\nu_i]) = [\zeta_i], \quad v([\beta_i]) = [\zeta_i], \quad v([\nu_i]) = [\alpha_i 1].$$

Define the obstruction group on the first leaf of the β -block:

$$J = V_{\beta_1} = V_{1000} \leq G.$$

The group Γ acts only on $[0]$, whereas J acts only on the disjoint cylinder $[1000]$. Figure 1 displays these two supports inside the nine-leaf code.




FIGURE 1. The nine-leaf prefix code. The blue, teal, and ochre blocks partition $[0]$, $[10]$, and $[11]$, respectively. The group Γ acts only on the blue block, while J acts only on the violet leaf β_1 in the teal block. Their supports are disjoint.

To verify the direct-product structure algebraically, the incomparability of 0 and 1000 gives

$$e_0 e_{1000} = e_{1000} e_0 = 0.$$

Every $g \in \Gamma$ satisfies $g - 1 \in e_0 R e_0$, while (15) gives $j - 1 \in e_{1000} R e_{1000}$ for every $j \in J$. Orthogonality therefore implies

$$(g - 1)(j - 1) = (j - 1)(g - 1) = 0.$$

Thus g and j commute. Moreover, if $g = j$, then their common difference from 1 belongs to both orthogonal corners and is therefore zero. Hence

$$[\Gamma, J] = 1, \quad \Gamma \cap J = \{1\}, \quad \Gamma \times J \leq G.$$

Thompson's group V is finitely presented, infinite, and simple [CFP96]. Every finitely presented LEF group is residually finite [VG97, Theorem 2.2], whereas an infinite simple group has no nontrivial finite quotient. Thus $J \cong V$ is finitely generated but not LEF.

Proposition 3.2. *The prefix replacements in (16) satisfy*

$$u\Gamma u^{-1} = v\Gamma v^{-1} = \text{EL}_{(\alpha_1 0, \alpha_2 0, \alpha_3 0)}(R) \leq \Gamma, \quad uJu^{-1} = V_{0001} \leq \Gamma \quad (17)$$

and

$$G = \langle \Gamma, u, v \rangle. \quad (18)$$

Proof. For $g = u, v$, conjugating an elementary α -root gives

$$g(1 + s_{\alpha_i} r t_{\alpha_j})g^{-1} = 1 + s_{\alpha_i 0} r t_{\alpha_j 0}, \quad i \neq j, \quad r \in R.$$

Thus both conjugates equal $\text{EL}_{(\alpha_1 0, \alpha_2 0, \alpha_3 0)}(R)$. This elementary prefix group lies in Γ , since

$$s_{\alpha_i 0} r t_{\alpha_j 0} = s_{\alpha_i} (s_0 r t_0) t_{\alpha_j}.$$

Moreover, u sends $\beta_1 = 1000$ to $\alpha_1 1 = 0001$, so

$$u J u^{-1} = V_{0001} \leq \Gamma$$

by Lemma 3.1. This proves (17).

To recover G , partition its six α - and β -leaves into three two-leaf blocks

$$C_i = \{\alpha_i, \beta_i\} \quad (1 \leq i \leq 3).$$

Set

$$X_i = s_{\alpha_i} t_0 + s_{\beta_i} t_1, \quad Y_j = s_0 t_{\alpha_j} + s_1 t_{\beta_j}.$$

The inverse prefix table for u gives

$$u^{-1}(1 + s_{\alpha_i} r t_{\alpha_j})u = 1 + X_i r Y_j \quad (i \neq j).$$

Write $\ell_{i,0} = \alpha_i$ and $\ell_{i,1} = \beta_i$. Given a matrix $B = (b_{pq}) \in M_2(R)$, take

$$r = \sum_{p,q \in \{0,1\}} s_p b_{pq} t_q.$$

Since $t_p r s_q = b_{pq}$, we obtain

$$X_i r Y_j = \sum_{p,q \in \{0,1\}} s_{\ell_{i,p}} b_{pq} t_{\ell_{j,q}}.$$

Taking a matrix with one nonzero entry produces every elementary root whose source and target lie in different blocks C_i and C_j . For any three distinct code leaves a, b, c , the elementary commutator identity

$$[1 + s_a r t_c, 1 + s_c t_b] = 1 + s_a r t_b$$

also supplies the roots between two leaves in the same block: choose c in any different block, so that both roots on the left join distinct blocks. Consequently,

$$\text{EL}_{(\alpha,\beta)}(R) \leq \langle \Gamma, u \rangle.$$

The same argument with v and the three blocks $\{\alpha_i, \nu_i\}$ gives

$$\text{EL}_{(\alpha,\nu)}(R) \leq \langle \Gamma, v \rangle.$$

It remains to connect a β -leaf to a ν -leaf. For any such leaves a, b , choose an α -leaf c . Both factors on the left-hand side of

$$[1 + s_a r t_c, 1 + s_c t_b] = 1 + s_a r t_b$$

belong to the six-leaf groups just obtained. Interchanging a and b also gives the reverse elementary root. Thus $\langle \Gamma, u, v \rangle$ contains every elementary D -root, proving (18). $\square$

Proof of Theorem 1.1. Proposition 3.2 shows that $\Gamma \leq G$, $J = V_{1000}$, $t_1 = u$, and $t_2 = v$ satisfy the hypotheses of Proposition 2.3. If G were sofic, Proposition 2.3 would make J LEF, which it is not. Therefore G is not sofic. Since $G \leq R^\times$ and soficity passes to subgroups, $R^\times$ is not sofic either. $\square$

Acknowledgments. We thank Henry Bradford, Michael Chapman, Alon Dogon, and Francesco Fournier-Facio for helpful comments on the manuscript.

REFERENCES

- [AL07] D. Aldous and R. Lyons, *Processes on unimodular random networks*, Electron. J. Probab. **12** (2007), 1454–1508, [arXiv:math/0603062](#).
- [BQ17] C. Bleak and M. Quick, *The infinite simple group V of Richard J. Thompson: presentations by permutations*, Groups Geom. Dyn. **11** (2017), 1401–1436, [arXiv:1511.02123](#).
- [BB20] L. Bowen and P. Burton, *Flexible stability and nonsocficity*, Trans. Amer. Math. Soc. **373** (2020), 4469–4481, [doi:10.1090/tran/8047](#).
- [BC25] L. Bowen and M. Chapman, *Surjunctivity does not characterize cosocficity of invariant random subgroups*, preprint, 2025, [arXiv:2511.06586](#).
- [BCLV24] L. Bowen, M. Chapman, A. Lubotzky, and T. Vidick, *The Aldous–Lyons conjecture I: Subgroup tests*, preprint, 2024, [arXiv:2408.00110](#).
- [BCV25] L. Bowen, M. Chapman, and T. Vidick, *The Aldous–Lyons conjecture II: Undecidability*, preprint, 2025, [arXiv:2501.00173](#).
- [BFF24] H. Bradford and F. Fournier-Facio, *Hopfian wreath products and the stable finiteness conjecture*, Math. Z. **308** (2024), article no. 58, [doi:10.1007/s00209-024-03589-3](#).
- [BS16] N. Brownlowe and A. P. W. Sørensen, $L_{2,\mathbb{Z}} \otimes L_{2,\mathbb{Z}}$ does not embed in $L_{2,\mathbb{Z}}$, J. Algebra **456** (2016), 1–22, [doi:10.1016/j.jalgebra.2016.01.040](#).
- [CFP96] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Introductory notes on Richard Thompson’s groups*, Enseign. Math. (2) **42** (1996), no. 3–4, 215–256, [available online](#).
- [CDL24] M. Chapman, Y. Dikstein, and A. Lubotzky, *Conditional non-socficity of p -adic Deligne extensions: On a theorem of Gohla and Thom*, preprint, 2024, [arXiv:2410.02913](#).
- [DCGLT20] M. De Chiffre, L. Glebsky, A. Lubotzky, and A. Thom, *Stability, cohomology vanishing, and nonapproximable groups*, Forum Math. Sigma **8** (2020), article no. e18, [doi:10.1017/fms.2020.5](#).
- [Dog23] A. Dogon, *Flexible Hilbert–Schmidt stability versus hyperlinearity for property (T) groups*, Math. Z. **305** (2023), article no. 58, [doi:10.1007/s00209-023-03387-3](#).
- [DV26] A. Dogon and I. Vigdorovich, *Hyperlinearity, stability and asymptotic spectral gap of higher rank lattices*, preprint, 2026, [arXiv:2506.20843v2](#).
- [ES05] G. Elek and E. Szabó, *Hyperlinearity, essentially free actions and L^2 -invariants: The sofic property*, Math. Ann. **332** (2005), 421–441, [arXiv:math/0408400](#).
- [ES06] G. Elek and E. Szabó, *On sofic groups*, J. Group Theory **9** (2006), 161–171, [arXiv:math/0305352](#).
- [EJ10] M. Ershov and A. Jaikin-Zapirain, *Property (T) for noncommutative universal lattices*, Invent. Math. **179** (2010), 303–347, [arXiv:0809.4095](#).
- [Gel18] T. Gelander, *A view on invariant random subgroups and lattices*, in *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018*, vol. II, World Scientific, 2018, 1339–1362, [arXiv:1807.06979](#).
- [GT24] L. Gohla and A. Thom, *High-dimensional expansion and soficity of groups*, preprint, 2024, [arXiv:2403.09582](#).
- [Gro99] M. Gromov, *Endomorphisms of symbolic algebraic varieties*, J. Eur. Math. Soc. **1** (1999), 109–197.
- [HO17] U. Haagerup and K. K. Olesen, *Non-inner amenability of the Thompson groups T and V* , J. Funct. Anal. **272** (2017), 4838–4852, [doi:10.1016/j.jfa.2017.02.003](#).
- [JNVWY20] Z. Ji, A. Natarajan, T. Vidick, J. Wright, and H. Yuen, $\text{MIP}^* = \text{RE}$, preprint, 2020, [arXiv:2001.04383](#).
- [KT26] H. V. Khanh and V. H. Thanh, *Matrix generators for the unit groups of $L_K(1, d)$* , preprint, 2026, [arXiv:2607.10351](#).
- [Kun19] G. Kun, *On sofic approximations of property (T) groups*, preprint, 2019, [arXiv:1606.04471v5](#).
- [KT19] G. Kun and A. Thom, *Inapproximability of actions and Kazhdan’s property (T)*, preprint, 2019, [arXiv:1901.03963](#).
- [Lea62] W. G. Leavitt, *The module type of a ring*, Trans. Amer. Math. Soc. **103** (1962), 113–130, [doi:10.1090/S0002-9947-1962-0132764-X](#).
- [Man25] A. Manzoor, *Invariant random subgroups, soficity, and Lück’s determinant conjecture*, preprint, 2025, [arXiv:2508.15154](#).
- [NT22] M. Nitsche and A. Thom, *Universal solvability of group equations*, J. Group Theory **25** (2022), 1–10, [arXiv:1811.07737](#).
- [Pes08] V. G. Pestov, *Hyperlinear and sofic groups: A brief guide*, Bull. Symbolic Logic **14** (2008), 449–480, [arXiv:0804.3968](#).
- [Tho08] A. Thom, *Sofic groups and diophantine approximation*, Comm. Pure Appl. Math. **61** (2008), 1155–1171, [arXiv:math/0701294](#).
- [VG97] A. M. Vershik and E. I. Gordon, *Groups that are locally embeddable in the class of finite groups*, Algebra i Analiz **9** (1997), no. 1, 71–97; English transl., St. Petersburg Math. J. **9** (1998), 49–67, [Math-Net.Ru aa751](#).
- [Wei00] B. Weiss, *Sofic groups and dynamical systems*, Sankhyā Ser. A **62** (2000), 350–359.

CHAPTER 4

A Counterexample to Connes's Rigidity Conjecture

ABSTRACT. Connes conjectured that the group von Neumann algebra of an ICC group with Kazhdan's property (T) determines the group up to isomorphism. We disprove the conjecture by constructing a countably infinite family of pairwise nonisomorphic, mutually commensurable, finitely generated ICC property- (T) groups with isomorphic group von Neumann algebras. The construction exploits the fact that $L^\infty(\hat{A}, m_{\hat{A}})$ depends on the underlying probability space, not on the group law: binary carry produces different compact abelian group structures with the same Haar measure and group action. Our examples also answer Popa's finite-to-one question in the negative and attain his countability bound.

CONTENTS

1. Introduction
 2. Group factors, property (T) , and Fourier duality
 3. A torsion-free property- (T) acting group
 4. The binary-carry construction
 5. ICC, property (T) , and the basic counterexample
 6. An infinite fiber of the group-factor functor
- References

1. INTRODUCTION

For a countable discrete group G , write $L(G)$ for its group von Neumann algebra. A group is *ICC* if it is infinite and every nonidentity conjugacy class is infinite; in this case, $L(G)$ is a II_1 factor. Connes's rigidity conjecture asks whether the group factor of an ICC property- (T) group determines the group. It grew out of his foundational rigidity theorem [Con80], goes back to his Kingston proceedings article [Con82], and appears explicitly as Problem 1 in his 1994 monograph [Con94, Chapter 5, Appendix B, Problem 1, p. 551], where the group factor is denoted by $R(G)$.

Conjecture 1.1 (Connes). *Let G and H be countable ICC groups with property (T) . If $L(G) \cong L(H)$, then $G \cong H$.*

Theorem 1.2 (Main result). *There exist finitely generated ICC groups with property (T) ,*

$$\Lambda, \Gamma_0, \Gamma_1, \Gamma_2, \dots,$$

that are pairwise nonisomorphic and satisfy

$$L(\Gamma_n) \cong L(\Lambda) \quad (n \geq 0).$$

Moreover, for every $n \geq 0$, the group Γ_n contains a subgroup isomorphic to Γ_0 of index 2^{4n} . In particular, the groups $(\Gamma_n)_{n \geq 0}$ are mutually commensurable.

Consequences and context. As recorded in Corollary 5.10, the pair Λ, Γ_0 already disproves Connes's conjecture. To establish this two-group result, it suffices to verify the ICC property and property (T) for Λ : both then transfer to Γ_0 through the common group factor. The full family in Theorem 1.2 additionally requires an intrinsic invariant to distinguish the groups Γ_n , and gives a stronger negative result: even the finite-to-one weakening fails. In his Madrid ICM address, Popa proved that the group-factor functor on ICC property- (T) groups is at most countable-to-one [Pop07, Section 4, pp. 457–458]; an alternative proof appears in [IPV13, Proposition 3.5]. He subsequently asked whether its fibers must in fact be finite [Pop13, p. 9]. Since $L(\Lambda)$ has countably infinitely many pairwise nonisomorphic ICC property- (T) group realizations, the answer is negative and Popa's countability bound is sharp.

The amenable case gives the opposite extreme: by Connes's celebrated classification theorem for injective factors [Con76], every amenable ICC group has the same group factor, the hyperfinite II_1 factor. Property (T) [Kaz67] was expected to prevent this collapse.

The conjecture also belongs to the broader theory of W^* -superrigidity. A countable group G is W^* -superrigid if $L(G) \cong L(H)$ implies $G \cong H$ for every countable group H . Furman's orbit-equivalence rigidity theorem provided an earlier analogue for probability-measure-preserving actions of higher-rank lattices [Fur99]. Popa introduced and developed deformation/rigidity theory, beginning with foundational work on Bernoulli shifts and rigid Cartan inclusions [Pop06a, Pop06b]. For suitable malleable actions of rigid groups, his strong-rigidity theorems [Pop06c, Pop06d] recover the acting group and its probability-space action from the crossed product: a group-measure-space analogue of Connes's conjecture. Ioana subsequently proved W^* -superrigidity for Bernoulli actions of ICC property- (T) groups [Ioa11]; this determines the action from its crossed product and does not assert rigidity of the bare group factor. For accounts of the theory, see Popa's ICM address [Pop07] and the later surveys [Vae10, Ioa18]. The first W^* -superrigid groups were constructed in [IPV13]; the first examples with property (T) followed in [CIOS23]. Thus some ICC property- (T) groups are indeed determined by their group factors; what Theorem 1.2 disproves is that property (T) alone guarantees this conclusion, even up to finite ambiguity.

Infinite fibers were already known outside the property- (T) setting. Indeed, as shown in [IPV13, Theorem 1.2], for every nontrivial finite abelian group H_0 and every $n \geq 3$, there are infinitely many pairwise nonisomorphic groups H satisfying

$$L(H) \cong L(H_0 \wr \text{PSL}_n(\mathbb{Z})).$$

The ordinary wreath product on the right does not have property (T). The new feature of Theorem 1.2 is an infinite fiber entirely within the ICC property-(T) class, precisely the class to which the countability bound applies.

Proof outline. For a countable abelian K -module A , Fourier transform gives

$$L(A \rtimes K) \cong L^\infty(\widehat{A}, m_{\widehat{A}}) \rtimes K.$$

The crossed product remembers the Haar probability space and its K -action, but need not remember the compact group law on $\widehat{A}$. Our strategy is therefore to put different K -invariant compact abelian group structures on a single probability space, and then to recover their differences from the dual discrete groups.

Section 3 constructs a torsion-free ICC property-(T) group

$$K = \ker \left(\mathrm{SL}_4(\mathbb{Z}[t]) \xrightarrow{g(t) \mapsto g(0) \bmod 3} \mathrm{SL}_4(\mathbb{F}_3) \right)$$

surjecting onto $\mathrm{SL}_4(\mathbb{F}_2[t])$. For $V = \mathbb{F}_2[t]^4$, the divided-square module

$$B = \mathrm{span}_{\mathbb{F}_2} \{v \otimes v : v \in V\}, \quad D = V \oplus B,$$

supplies the common compact K -space $\widehat{D} = X \times Y$, where $X = V^*$ and $Y = B^*$.

The elementary model for the construction is the four-point probability space $\mathbb{F}_2^2$. Coordinatewise addition gives the Klein four-group, whereas

$$(x, y) \diamond (x', y') = (x + x', y + y' + xx')$$

gives $\mathbb{Z}/4\mathbb{Z}$; both have the same uniform Haar measure. Section 4 globalizes this binary carry to a compact group C_0 with the same measured K -action as $X \times Y$. Writing

$$\Lambda = D \rtimes K, \quad \Gamma_0 = \widehat{C_0} \rtimes K,$$

Fourier transform identifies $L(\Gamma_0) \cong L(\Lambda)$, while order-four torsion distinguishes the two groups.

Section 5 proves that Λ is ICC and has property (T). The ICC assertion follows from an infinite-orbit calculation. For property (T), transitivity on primitive vectors and a quadratic Boolean support bound give a uniform spectral estimate for the semidirect product. Both properties then transfer to Γ_0 through the common group factor.

Finally, Section 6 shifts the carry by n coefficient positions in each of the four directions, producing compact groups C_n with the same measured K -action. Their duals give

$$\Gamma_n = \widehat{C_n} \rtimes K, \quad L(\Gamma_n) \cong L(\Lambda).$$

Pontryagin duality embeds Γ_0 in Γ_n with index 2^{4n} . The finite-orbit part of the intrinsic quotient $E_n[2]/2E_n$, where $E_n = \widehat{C_n}$, also has order 2^{4n} , recovering n from the abstract group and proving that the resulting family is pairwise nonisomorphic.

2. GROUP FACTORS, PROPERTY (T), AND FOURIER DUALITY

2.1. Group von Neumann algebras. Let G be a countable discrete group, and write $\mathcal{U}(\mathcal{H})$ for the unitary group of a Hilbert space $\mathcal{H}$. The left regular representation of G is

$$\lambda_G : G \longrightarrow \mathcal{U}(\ell^2(G)), \quad \lambda_G(g)\delta_h = \delta_{gh}.$$

The *group von Neumann algebra* of G and its canonical trace are

$$L(G) = \lambda_G(G)'', \quad \tau_G(x) = \langle x\delta_1, \delta_1 \rangle.$$

For a finite sum $x = \sum_g a_g \lambda_G(g)$, one has $\tau_G(x) = a_1$. A von Neumann algebra is a *factor* if its center is $\mathbb{C}1$; an infinite-dimensional factor admitting a faithful normal tracial state is a II_1 factor. The group G is *ICC* if every nonidentity element has an infinite conjugacy class. For infinite G , the group-factor criterion asserts that $L(G)$ is a II_1 factor if and only if G is ICC [MN43].

2.2. Property (T). Let G be a countable discrete group and let $\pi : G \rightarrow \mathcal{U}(\mathcal{H})$ be a unitary representation. It has *almost invariant unit vectors* if, for every finite $F \subseteq G$ and $\varepsilon > 0$, there is a unit vector $\xi \in \mathcal{H}$ such that

$$\max_{g \in F} \|\pi(g)\xi - \xi\| < \varepsilon.$$

The group G has *property (T)* if every unitary representation with almost invariant unit vectors has a nonzero invariant vector. Property (T) passes to quotients and is preserved under passage to finite-index subgroups and finite-index extensions. Every countable discrete property-(T) group is finitely generated; see [BHV08].

For a subgroup $A \leq G$, the pair (G, A) has *relative property (T)* if every unitary representation of G with almost invariant unit vectors has a nonzero A -invariant vector. If $A \triangleleft G$, relative property (T) for (G, A) , together with property (T) for G/A , implies property (T) for G .

Lemma 2.1. *Let G and H be countable discrete groups such that $L(G) \cong L(H)$. If H is ICC and has property (T), then G is ICC and has property (T).*

Proof. Since H is ICC, $L(H)$, and therefore $L(G)$, is a II_1 factor. The group-factor criterion shows that G is ICC. By the Connes–Jones characterization [CJ85], an ICC group has property (T) if and only if its group factor does. Applying this equivalence first to H and then to G proves the assertion. $\square$

2.3. Pontryagin duality and crossed products. The algebra $L^\infty(\widehat{A}, m_{\widehat{A}})$ depends only on the Haar probability space of $\widehat{A}$, not on its compact group law. Once the measure and the K -action are fixed, so is the crossed product.

Put $\mathbb{T} = \{z \in \mathbb{C} : |z| = 1\}$. For a locally compact abelian group A , write

$$\widehat{A} = \text{Hom}_{\text{cont}}(A, \mathbb{T})$$

for its Pontryagin dual, equipped with the compact-open topology. When A is countable and discrete, $\widehat{A}$ is compact; denote its normalized Haar measure by $m_{\widehat{A}}$ and its trivial character by 1. If a countable discrete group K acts on A by automorphisms, the induced actions on characters and functions are

$$(k \cdot \chi)(a) = \chi(k^{-1} \cdot a), \quad \alpha_k(f)(\chi) = f(k^{-1} \cdot \chi).$$

The semidirect product $A \rtimes K$ has multiplication $(a, k)(b, h) = (a + k \cdot b, kh)$. The corresponding crossed product $L^\infty(\widehat{A}, m_{\widehat{A}}) \rtimes K$ is generated, in its standard representation, by $L^\infty(\widehat{A}, m_{\widehat{A}})$ and unitaries $(v_k)_{k \in K}$ satisfying

$$v_k f v_k^* = \alpha_k(f), \quad v_k v_l = v_{kl}.$$

Under

$$\ell^2(A \rtimes K) \cong \ell^2(A) \otimes \ell^2(K),$$

Fourier transform in the A -coordinate identifies the group unitaries of A with multiplication by $\chi \mapsto \chi(a)$, and those of K with the unitaries implementing the dual action. Hence

$$L(A \rtimes K) \cong L^\infty(\widehat{A}, m_{\widehat{A}}) \rtimes K. \tag{2.1}$$

In particular, let A' be another countable discrete abelian K -module. A measurable bijection modulo null sets $\theta : \widehat{A} \rightarrow \widehat{A}'$ is measure-preserving and K -equivariant when

$$\theta_* m_{\widehat{A}} = m_{\widehat{A}'}, \quad \theta(k \cdot \chi) = k \cdot \theta(\chi) \quad \text{for every } k \in K \text{ and almost every } \chi.$$

Then $f \mapsto f \circ \theta^{-1}$, together with the identity on the implementing unitaries, gives

$$L(A \rtimes K) \cong L(A' \rtimes K).$$

No compatibility between the group laws on $\widehat{A}$ and $\widehat{A}'$ is required; this is the mechanism underlying our construction.

3. A TORSION-FREE PROPERTY-(T) ACTING GROUP

We need a torsion-free ICC property-(T) group that surjects onto $\mathrm{SL}_4(\mathbb{F}_2[t])$. A congruence condition at the different prime 3 will provide torsion-freeness without losing the characteristic-2 action. Define

$$K = \ker\left(\mathrm{SL}_4(\mathbb{Z}[t]) \xrightarrow{g(t) \mapsto g(0) \bmod 3} \mathrm{SL}_4(\mathbb{F}_3)\right).$$

Set

$$R = \mathbb{F}_2[t], \quad Q = \mathrm{SL}_4(R).$$

The groups K and Q are countable.

For a commutative ring S , let

$$\mathrm{EL}_n(S) = \langle I + fE_{ij} : f \in S, i \neq j \rangle \leq \mathrm{SL}_n(S),$$

where E_{ij} is the (i, j) -matrix unit. Thus $\mathrm{EL}_n(S)$ is the *elementary linear group*. Since $\mathbb{Z}$ is a regular ring of Krull dimension 1 and $\mathrm{SK}_1(\mathbb{Z}) = 0$, the elementary generation theorem for polynomial rings [Sus77, Corollary 6.6] gives

$$\mathrm{EL}_n(\mathbb{Z}[t]) = \mathrm{SL}_n(\mathbb{Z}[t]) \quad (n \geq 3).$$

Property (T) for elementary linear groups over finitely generated unital rings [EJ10, Theorem 1.1] therefore gives property (T) for $\mathrm{SL}_4(\mathbb{Z}[t])$.

Proposition 3.1. *The group K has property (T).*

Proof. Evaluation at $t = 0$ followed by reduction modulo 3 has finite image. Thus K has finite index in $\mathrm{SL}_4(\mathbb{Z}[t])$ and inherits property (T). $\square$

We use the following level-3 congruence lemma [Min87].

Lemma 3.2. *For every $d \geq 2$, the group*

$$\ker(\mathrm{SL}_d(\mathbb{Z}) \longrightarrow \mathrm{SL}_d(\mathbb{F}_3))$$

is torsion-free.

Proof. If the displayed group contains a nonidentity torsion element, it contains one of prime order p . Write this element as

$$A = I + 3^r B,$$

where $r \geq 1$ is maximal, so $B \not\equiv 0 \pmod{3}$. If $p \neq 3$, then

$$I = A^p \equiv I + p3^r B \pmod{3^{r+1}},$$

contradicting $B \not\equiv 0 \pmod{3}$. If $p = 3$, then

$$0 = \frac{A^3 - I}{3^{r+1}} = B + 3^r B^2 + 3^{2r-1} B^3.$$

Reduction modulo 3 again gives $B \equiv 0 \pmod{3}$, a contradiction. $\square$

Lemma 3.3. *The group K is torsion-free and ICC. Reduction modulo 2 gives a surjection*

$$\pi_2 : K \rightarrow Q.$$

Proof. For torsion-freeness, suppose $g(t) \in K$ satisfies $g(t)^m = I$ for some $m \geq 1$. Its constant term $g(0)$ is a torsion element of $\ker(\mathrm{SL}_4(\mathbb{Z}) \rightarrow \mathrm{SL}_4(\mathbb{F}_3))$, so Lemma 3.2 gives $g(0) = I$. If $g \neq I$, let $r \geq 1$ be its first nonzero t -adic degree:

$$g(t) = I + t^r A(t), \quad A(0) \neq 0.$$

Since $\mathbb{Z}$ is an integral domain,

$$g(t)^m \equiv I + mt^r A(0) \not\equiv I \pmod{t^{r+1}},$$

contradicting $g^m = I$.

For the ICC property, let

$$u_{ij}(f) = I + 3fE_{ij} \in K \quad (i \neq j, f \in \mathbb{Z}[t]).$$

A matrix commuting with every E_{ij} is scalar, and every scalar element of $\mathrm{SL}_4(\mathbb{Z}[t])$ has finite order. Thus, for a nonidentity $g \in K$, torsion-freeness provides $i \neq j$ such that

$$[E_{ij}, g] = E_{ij}g - gE_{ij}$$

is nonzero. If $u_{ij}(f)$ and $u_{ij}(h)$ produce the same conjugate of g , then g commutes with $u_{ij}(f-h)$, and hence

$$3(f-h)[E_{ij}, g] = 0.$$

As $\mathbb{Z}[t]$ is an integral domain and $[E_{ij}, g] \neq 0$, we must have $f = h$. Varying f therefore produces infinitely many conjugates.

Finally, reduction modulo 2 gives

$$u_{ij}(f) \mapsto I + \bar{f}E_{ij}.$$

Every elementary matrix over R arises in this way. Since R is Euclidean, these matrices generate $Q = \mathrm{SL}_4(R)$, proving that π_2 is onto. $\square$

Remark 3.4. Although the module and orbit calculations factor through Q , this group cannot replace K :

$$u = I + E_{12} + E_{23} \in Q, \quad u^2 = I + E_{13} \neq I, \quad u^4 = I,$$

so Q has an element of order 4. Torsion-freeness of K , by contrast, ensures that Λ has no element of order 4, since D has exponent 2; this distinguishes Λ from Γ_0 in Proposition 4.3. It also identifies E_n as the torsion subgroup of Γ_n , the intrinsic starting point for recovering n in Proposition 6.9.

4. THE BINARY-CARRY CONSTRUCTION

The distinction between a measured space and its group law already appears on the four-point set $\mathbb{F}_2 \times \mathbb{F}_2$. Consider

$$\begin{aligned} (x, y) + (x', y') &= (x + x', y + y'), \\ (x, y) \diamond (x', y') &= (x + x', y + y' + xx'), \end{aligned}$$

with all coordinates on the right computed in $\mathbb{F}_2$. The first law gives the Klein four-group; for the second,

$$(x, y) \mapsto \tilde{x} + 2\tilde{y} \pmod{4}$$

is an isomorphism with $\mathbb{Z}/4\mathbb{Z}$, where tildes denote representatives in $\{0, 1\}$. Both laws nevertheless have the same uniform Haar probability measure. Their difference is the binary carry xx' .

We globalize this carry K -equivariantly. The resulting compact abelian group C_0 admits a K -equivariant, measure-preserving homeomorphism from $X \times Y = \widehat{D}$, but its group law is not the coordinatewise one. Its Pontryagin dual E_0 will define Γ_0 .

4.1. The common linear and quadratic modules. Let $V = R^4$, let $e_1, \dots, e_4$ be its standard basis, and put

$$e = e_1, \quad b(v) = v \otimes_{\mathbb{F}_2} v.$$

Define

$$B = \mathrm{span}_{\mathbb{F}_2} \{b(v) : v \in V\} \subseteq V \otimes_{\mathbb{F}_2} V, \quad D = V \oplus B.$$

The tensor product is over $\mathbb{F}_2$, not over R .

Let K act on V through $\pi_2 : K \rightarrow \mathrm{SL}_4(R)$. Its diagonal action on $V \otimes_{\mathbb{F}_2} V$ preserves B , since

$$k \cdot b(v) = b(k \cdot v),$$

and therefore acts on D . These are countable discrete $\mathbb{F}_2$ -vector spaces, so D has exponent 2 and

$$0 \longrightarrow V \xrightarrow{v \mapsto (v, 0)} D \xrightarrow{(v, w) \mapsto w} B \longrightarrow 0$$

is a split K -equivariant exact sequence.

The polarization of b is bilinear:

$$b(u + v) - b(u) - b(v) = u \otimes v + v \otimes u.$$

For an ordered $\mathbb{F}_2$ -basis $(u_i)_{i \geq 1}$ of V , a basis of B is

$$u_i \otimes u_i, \quad u_i \otimes u_j + u_j \otimes u_i \quad (i < j).$$

Indeed, if $v = \sum_i x_i u_i$, then

$$b(v) = \sum_i x_i (u_i \otimes u_i) + \sum_{i < j} x_i x_j (u_i \otimes u_j + u_j \otimes u_i).$$

Conversely,

$$u_i \otimes u_j + u_j \otimes u_i = b(u_i + u_j) + b(u_i) + b(u_j),$$

and linear independence follows from the tensor-product basis $(u_i \otimes u_j)_{i,j}$. These are precisely the tensors fixed by interchanging the two factors; thus

$$B = (V \otimes_{\mathbb{F}_2} V)^{\mathfrak{S}_2} = \Gamma_{\mathbb{F}_2}^2(V),$$

the divided square of V . Crucially, this is the invariant subspace, not the ordinary symmetric-square quotient

$$\text{Sym}_{\mathbb{F}_2}^2(V) = (V \otimes_{\mathbb{F}_2} V)_{\mathfrak{S}_2}.$$

Indeed, the canonical map from invariants to coinvariants kills every off-diagonal basis element $u_i \otimes u_j + u_j \otimes u_i$.

For $q \in B^* = \text{Hom}_{\mathbb{F}_2}(B, \mathbb{F}_2)$, the displayed expansion gives

$$q(b(v)) = \sum_i x_i q(u_i \otimes u_i) + \sum_{i < j} x_i x_j q(u_i \otimes u_j + u_j \otimes u_i).$$

Thus B^* parametrizes the Boolean quadratic functions on V with zero constant term: diagonal coordinates supply the linear terms, and off-diagonal coordinates the quadratic terms. In particular, for every finite $J \subseteq \mathbb{N}$, the restriction of $v \mapsto q(b(v))$ to $\text{span}_{\mathbb{F}_2}\{u_j : j \in J\}$ is a Boolean polynomial of degree at most 2.

Via

$$(\ell, q) \longmapsto [(v, w) \mapsto (-1)^{\ell(v) + q(w)}],$$

Pontryagin duality identifies

$$\widehat{D} = X \times Y, \quad X = V^* = \text{Hom}_{\mathbb{F}_2}(V, \mathbb{F}_2), \quad Y = B^* = \text{Hom}_{\mathbb{F}_2}(B, \mathbb{F}_2).$$

With their pointwise-convergence topologies, X and Y are compact products of copies of $\mathbb{F}_2$. Thus $\widehat{D} = X \times Y$ has coordinatewise addition and Haar probability measure

$$m = m_X \otimes m_Y.$$

The dual actions are

$$(k \cdot \ell)(v) = \ell(k^{-1} \cdot v), \quad (k \cdot q)(w) = q(k^{-1} \cdot w).$$

4.2. The carry group. For $a \in \mathbb{F}_2$, let $\tilde{a} \in \{0, 1\} \subset \mathbb{Z}$ denote its standard representative. For $(\ell, q) \in X \times Y$, define $F_{\ell, q} : V \rightarrow \mathbb{Z}/4\mathbb{Z}$ by

$$F_{\ell, q}(v) = \widetilde{\ell(v)} + 2\widetilde{q(b(v))} \pmod{4},$$

and set

$$C_0 = \{F_{\ell, q} : (\ell, q) \in X \times Y\} \subseteq (\mathbb{Z}/4\mathbb{Z})^V.$$

Equip $(\mathbb{Z}/4\mathbb{Z})^V$ with the product topology and C_0 with the subspace topology. The K -action is

$$(k \cdot F)(v) = F(k^{-1} \cdot v).$$

It preserves C_0 , since $k \cdot F_{\ell, q} = F_{k \cdot \ell, k \cdot q}$.

For $s, s' \in \mathbb{F}_2$, the standard lifts satisfy the binary-carry identity

$$\tilde{s} + \tilde{s}' = \widetilde{s + s'} + 2\tilde{s} \tilde{s}' \pmod{4}.$$

For $\ell, \ell' \in X$, set

$$r_{\ell, \ell'} = (\ell \otimes \ell')|_B \in Y.$$

Here $(\ell \otimes \ell')(u \otimes v) = \ell(u)\ell'(v)$. Since

$$r_{\ell, \ell'}(b(v)) = \ell(v)\ell'(v),$$

and the elements $b(v)$ span B , the map $(\ell, \ell') \mapsto r_{\ell, \ell'}$ is symmetric, bilinear, and normalized: $r_{\ell, 0} = r_{0, \ell} = 0$. Bilinearity gives the cocycle identity

$$r_{\ell, \ell'} + r_{\ell + \ell', \ell''} = r_{\ell', \ell''} + r_{\ell, \ell' + \ell''}.$$

Evaluation at any $w \in B$ uses only finitely many coordinates of ℓ and ℓ' , so r is continuous for the pointwise topologies. It is therefore a continuous, symmetric, normalized, K -equivariant Y -valued 2-cocycle on X . Applying the binary-carry identity at each $v \in V$ gives

$$F_{\ell, q} + F_{\ell', q'} = F_{\ell + \ell', q + q' + r_{\ell, \ell'}}. \quad (4.2)$$

Thus C_0 contains 0 and is closed under addition. Since the ambient group has exponent 4, every $F \in C_0$ has inverse $-F = 3F \in C_0$; hence C_0 is a subgroup.

Reduction modulo 2 recovers ℓ from $F_{\ell, q}$, and the remaining values $q(b(v))$ recover q , since the $b(v)$ span B . Hence $(\ell, q) \mapsto F_{\ell, q}$ is a bijection under which pointwise addition becomes

$$(\ell, q) \star_0 (\ell', q') = (\ell + \ell', q + q' + r_{\ell, \ell'}).$$

Proposition 4.1. *The set C_0 is a compact subgroup of $(\mathbb{Z}/4\mathbb{Z})^V$. The map*

$$\Phi_0 : X \times Y \longrightarrow C_0, \quad \Phi_0(\ell, q) = F_{\ell, q},$$

is a K -equivariant homeomorphism that identifies the Haar measure of C_0 with m . It is a topological group isomorphism for $\star_0$, but not for coordinatewise addition on $X \times Y = \widehat{D}$.

Proof. Each coordinate of Φ_0 is continuous, and the preceding recovery argument shows that Φ_0 is bijective. It is therefore a homeomorphism from the compact space $X \times Y$ onto its image in the Hausdorff group $(\mathbb{Z}/4\mathbb{Z})^V$; in particular, C_0 is compact and closed. The identity $b(k \cdot v) = k \cdot b(v)$ gives K -equivariance.

Under Φ_0 , translation by $F_{\ell', q'}$ is the triangular map

$$(\ell, q) \longmapsto (\ell + \ell', q + q' + r_{\ell, \ell'}).$$

This map preserves m : its first coordinate is translation in X , and each fiber map is translation in Y . Consequently, $(\Phi_0)_*m$ is the Haar measure of C_0 .

Equation (4.2) identifies Φ_0 as a topological group isomorphism for $\star_0$. It cannot be a homomorphism for coordinatewise addition: if $\ell(e) = 1$, then $r_{\ell, \ell}(b(e)) = 1$ and $F_{\ell, 0}(e) = 1$. Thus $F_{\ell, 0}$ has order 4, whereas $X \times Y$ with coordinatewise addition has exponent 2. $\square$

4.3. Dualizing the carry extension. The binary carry is invisible to the measured K -space but survives Pontryagin duality as a nonsplit extension. Let $E_0 = \widehat{C_0}$, written additively and equipped with the dual K -action, so that biduality identifies C_0 canonically with $\widehat{E_0}$. For $v \in V$, define

$$\varepsilon_v^{(0)} \in E_0, \quad \varepsilon_v^{(0)}(F) = i^{F(v)},$$

where $i = \sqrt{-1}$. Since V is countable, C_0 is compact metrizable and E_0 is countable. Moreover, $4\varepsilon_e^{(0)} = 0$, and any $\ell \in X$ with $\ell(e) = 1$ gives $\varepsilon_e^{(0)}(F_{\ell, 0}) = i$. Thus $\varepsilon_e^{(0)}$ has order 4, exhibiting the obstruction to splitting.

Lemma 4.2. *There are K -equivariant exact sequences*

$$\begin{aligned} 0 \longrightarrow Y \longrightarrow \widehat{D} = X \times Y \longrightarrow X \longrightarrow 0, \\ 0 \longrightarrow Y \xrightarrow{j_0} C_0 \xrightarrow{\rho_0} X \longrightarrow 0, \end{aligned}$$

whose Pontryagin duals are

$$\begin{aligned} 0 \longrightarrow V \longrightarrow D = V \oplus B \longrightarrow B \longrightarrow 0, \\ 0 \longrightarrow V \xrightarrow{\iota_0} E_0 \xrightarrow{\sigma_0} B \longrightarrow 0. \end{aligned}$$

The sequences involving $\widehat{D}$ and D split K -equivariantly. The sequence involving E_0 does not split even as an extension of abelian groups; equivalently, the map ρ_0 admits no continuous homomorphic section. The maps for the carry extension satisfy

$$j_0(q) = F_{0,q}, \quad \rho_0(F_{\ell,q}) = \ell, \quad \iota_0(v)(F_{\ell,q}) = (-1)^{\ell(v)}.$$

The map σ_0 is characterized by

$$\eta(F_{0,q}) = (-1)^{q(\sigma_0(\eta))} \quad (\eta \in E_0, q \in Y),$$

and

$$2\varepsilon_v^{(0)} = \iota_0(v), \quad \sigma_0(\varepsilon_v^{(0)}) = b(v).$$

Proof. The sequences for $\widehat{D}$ and D split by their direct-sum decompositions. For the carry sequence, (4.2) shows that j_0 and ρ_0 are continuous homomorphisms with

$$\ker \rho_0 = \{F_{0,q} : q \in Y\} = j_0(Y).$$

Since $\rho_0(F_{\ell,0}) = \ell$, the sequence is exact; all its maps are K -equivariant. Pontryagin duality gives the stated discrete extension, with σ_0 obtained by restricting a character to $j_0(Y)$ and using $\widehat{Y} = B$. For evaluation characters,

$$\varepsilon_v^{(0)}(F_{0,q}) = i^{2q(b(v))} = (-1)^{q(b(v))},$$

so $\sigma_0(\varepsilon_v^{(0)}) = b(v)$. Since the $b(v)$ span B , these characters also give surjectivity of σ_0 .

If $\sigma_0(\eta) = 0$, then η factors through $C_0/j_0(Y) \cong X$; since $\widehat{X} = V$, it equals $\iota_0(v)$ for a unique $v \in V$. Finally,

$$(\varepsilon_v^{(0)}(F_{\ell,q}))^2 = (-1)^{F_{\ell,q}(v)} = (-1)^{\ell(v)},$$

which gives $2\varepsilon_v^{(0)} = \iota_0(v)$.

If the discrete sequence split, then $E_0 \cong V \oplus B$ would have exponent 2, contradicting the order-four element $\varepsilon_e^{(0)}$. By Pontryagin duality, ρ_0 therefore admits no continuous homomorphic section. $\square$

Define

$$\Gamma_0 = E_0 \rtimes K, \quad \Lambda = D \rtimes K.$$

The common measured K -action identifies the group factors; the order-four carry distinguishes the groups.

Proposition 4.3. *There is an isomorphism*

$$L(\Gamma_0) \cong L(\Lambda),$$

but $\Gamma_0 \not\cong \Lambda$.

Proof. Equation (2.1) and Proposition 4.1 give

$$\begin{aligned} L(\Gamma_0) &\cong L^\infty(C_0, m) \rtimes K \\ &\cong L^\infty(X \times Y, m) \rtimes K \cong L(\Lambda). \end{aligned}$$

Since K is torsion-free, every finite-order element of $\Lambda = D \rtimes K$ lies in D , which has exponent 2. But $\varepsilon_e^{(0)} \in E_0 \subset \Gamma_0$ has order 4. Hence $\Gamma_0 \not\cong \Lambda$. $\square$

5. ICC, PROPERTY (T), AND THE BASIC COUNTEREXAMPLE

The goal of this section is the two-group counterexample of Corollary 5.10. Proposition 4.3 already shows that Γ_0 and Λ are nonisomorphic but have isomorphic group von Neumann algebras. By Lemma 2.1, it therefore suffices to prove that Λ is ICC and has property (T). These two properties reduce, respectively, to infinite orbits on D and a uniform spectral estimate on $\widehat{D} = X \times Y$.

5.1. The ICC property.

Lemma 5.1. *Let an ICC group H act by automorphisms on an abelian group A . If every nonzero element of A has an infinite H -orbit, then $A \rtimes H$ is ICC.*

Proof. If $h \neq 1$, the conjugacy class of (a, h) maps onto the infinite conjugacy class of h in H . If $h = 1$ and $a \neq 0$, then

$$(0, k)(a, 1)(0, k)^{-1} = (k \cdot a, 1) \quad (k \in H).$$

The latter conjugacy class is infinite because the H -orbit of a is infinite. $\square$

Since $D = V \oplus B$, Lemma 5.1 reduces the ICC assertion to the following orbit calculation, which will also distinguish the groups in Section 6.

Lemma 5.2. *Every nonzero element of $V \oplus B$ has an infinite Q -orbit.*

Proof. Elementary transvections give infinite orbits on V ; on its tensor square, specialization in one tensor variable distinguishes the same transvections.

For $v \in V \setminus \{0\}$, choose j with $v_j \neq 0$ and $i \neq j$. The transvections

$$g_n = I + t^n E_{ij} \in Q \quad (n \geq 1), \quad g_n v = v + t^n v_j e_i$$

give pairwise distinct vectors.

For $B \subseteq V \otimes_{\mathbb{F}_2} V$, introduce separate variables for the tensor factors:

$$S = \mathbb{F}_2[t_1, t_2], \quad V \otimes_{\mathbb{F}_2} V \cong (R \otimes_{\mathbb{F}_2} R)^{16} \cong S^{16}, \quad t_1 = t \otimes 1, \quad t_2 = 1 \otimes t,$$

so $g \in Q$ acts as $g(t_1) \otimes g(t_2)$. Given $w \in S^{16} \setminus \{0\}$, remove its largest common power of t_2 :

$$w = t_2^a w_0, \quad w_0(t_1, 0) \neq 0.$$

Write

$$w_0(t_1, 0) = \sum_{j=1}^4 e_j \otimes w_j, \quad w_j \in \mathbb{F}_2[t_1]^4.$$

Choose j with $w_j \neq 0$ and $i \neq j$, and set

$$z = (E_{ij} \otimes I)w_0(t_1, 0) \neq 0, \quad h_r = I + t^r E_{ij} \in Q \quad (r \geq 1).$$

If $h_r w = h_s w$, cancel t_2^a and specialize $t_2 = 0$. Since $h_r(0) = I$ in the second tensor factor, this gives

$$(t_1^r + t_1^s)z = 0.$$

As $\mathbb{F}_2[t_1]^{16}$ is torsion-free and $z \neq 0$, we obtain $r = s$. Every nonzero element of B therefore has an infinite orbit. Finally, for $(v, w) \in V \oplus B$, use the equivariant projection to V if $v \neq 0$, and the preceding argument if $v = 0$ and $w \neq 0$. $\square$

Proposition 5.3. *The group Λ is ICC.*

Proof. Lemma 5.2 and the surjection $K \twoheadrightarrow Q$ show that every nonzero element of D has an infinite K -orbit. Since K is ICC by Lemma 3.3, Lemma 5.1 shows that $\Lambda = D \rtimes K$ is ICC. $\square$

5.2. A spectral criterion for property (T). The extension

$$1 \longrightarrow D \longrightarrow \Lambda \longrightarrow K \longrightarrow 1$$

has a property-(T) quotient by Proposition 3.1. By the extension principle in Subsection 2.2, it therefore suffices to prove relative property (T) for (Λ, D) . The following spectral criterion combines these two steps: an invariant spectral measure satisfying a uniform detection estimate must have an atom at the trivial character.

If $\pi : A \rtimes H \rightarrow \mathcal{U}(\mathcal{H})$ is a unitary representation with A abelian, the joint spectral theorem gives a unique projection-valued measure P on $\hat{A}$ such that

$$\pi(a) = \int_{\hat{A}} \chi(a) dP(\chi) \quad (a \in A).$$

Its covariance relation is

$$\pi(h)P(U)\pi(h)^{-1} = P(h \cdot U)$$

for every Borel set $U \subseteq \widehat{A}$ and $h \in H$.

An H -fixed unit vector therefore induces an H -invariant scalar spectral probability measure, and $P(\{1\})$ projects onto the A -fixed subspace. The criterion is a spectral form of the relative-property- (T) criterion in [CT11].

Lemma 5.4. *Let a countable property- (T) group H act by automorphisms on a countable abelian group A . Suppose there are a finite set $J \subseteq A$ and $c > 0$ such that every H -invariant Borel probability measure μ on $\widehat{A}$ satisfies*

$$\sum_{a \in J} \int_{\widehat{A}} |\chi(a) - 1|^2 d\mu(\chi) \geq c(1 - \mu(\{1\})). \quad (5.1)$$

Then $A \rtimes H$ has property (T) .

Proof. Suppose $A \rtimes H$ does not have property (T) . Choose a unitary representation π with asymptotically invariant unit vectors (ξ_n) but no nonzero invariant vector. For a Kazhdan pair (F, κ) of H , let P_H project onto $\mathcal{H}^H$. The Kazhdan estimate on $(\mathcal{H}^H)^\perp$ gives

$$\|\xi_n - P_H \xi_n\| \leq \kappa^{-1} \max_{h \in F} \|\pi(h)\xi_n - \xi_n\| \rightarrow 0.$$

Hence, after discarding finitely many terms,

$$\eta_n = \frac{P_H \xi_n}{\|P_H \xi_n\|}$$

are H -invariant and satisfy $\|\eta_n - \xi_n\| \rightarrow 0$.

Let P be the spectral measure of $\pi|_A$ and define

$$\mu_n(U) = \langle P(U)\eta_n, \eta_n \rangle.$$

Covariance makes μ_n H -invariant. For every $a \in J$,

$$\|\pi(a)\eta_n - \eta_n\| \leq 2\|\eta_n - \xi_n\| + \|\pi(a)\xi_n - \xi_n\| \rightarrow 0,$$

and hence

$$\int_{\widehat{A}} |\chi(a) - 1|^2 d\mu_n(\chi) = \|\pi(a)\eta_n - \eta_n\|^2 \rightarrow 0.$$

As J is finite, (5.1) forces $\mu_n(\{1\}) \rightarrow 1$. Thus $P(\{1\})\eta_n \neq 0$ for large n . This vector is A -fixed by definition of $P(\{1\})$, and H -fixed by covariance, since the trivial character is H -fixed. It is therefore $A \rtimes H$ -invariant, a contradiction. $\square$

5.3. A quadratic Boolean estimate. Let $\mathbb{F}_2$ be the field with two elements. For $m \geq 2$, every function $p : \mathbb{F}_2^m \rightarrow \mathbb{F}_2$ has a unique multilinear polynomial representative; let $\deg p$ denote its degree and $\text{supp}(p) = \{x \in \mathbb{F}_2^m : p(x) \neq 0\}$.

Lemma 5.5. *Let $m \geq 2$. If a nonzero polynomial function $p : \mathbb{F}_2^m \rightarrow \mathbb{F}_2$ has degree at most 2, then*

$$|\text{supp}(p)| \geq 2^{m-2}.$$

Equivalently, p is nonzero on at least one quarter of the Boolean cube.

Proof. The assertion is immediate for $m = 2$. For $m \geq 3$, write

$$p(x', x_m) = p_0(x') + x_m p_1(x'),$$

where $\deg p_0 \leq 2$ and $\deg p_1 \leq 1$. If $p_1 = 0$, then $p_0 \neq 0$; induction gives

$$|\text{supp}(p)| = 2|\text{supp}(p_0)| \geq 2 \cdot 2^{(m-1)-2} = 2^{m-2}.$$

Suppose that $p_1 \neq 0$. A nonzero affine-linear function on $\mathbb{F}_2^{m-1}$ is nonzero on at least 2^{m-2} points: it is either the constant function 1, or its two fibers have equal size. For every $x' \in \text{supp}(p_1)$, exactly one of $p(x', 0)$ and $p(x', 1)$ is nonzero. Therefore

$$|\text{supp}(p)| \geq |\text{supp}(p_1)| \geq 2^{m-2}. \quad \square$$

5.4. A uniform evaluation estimate. The spectral estimate comes from one K -orbit: transitivity makes detection at e equivalent to detection at any primitive vector, while Lemma 5.5 ensures that every nontrivial character is detected by many such vectors.

A vector $v = (v_1, \dots, v_4) \in R^4$ is *primitive* if its coordinates generate the unit ideal. Let

$$\mathcal{P} = \{v \in R^4 : (v_1, v_2, v_3, v_4) = R\}.$$

For $N \geq 1$, put

$$R_N = \text{span}_{\mathbb{F}_2}\{1, t, \dots, t^{N-1}\}, \quad V_N = R_N^4.$$

Lemma 5.6. *The action of K on $\mathcal{P}$ is transitive, and*

$$|\mathcal{P} \cap V_N| = 7 \cdot 2^{4N-3} + 1.$$

Equivalently,

$$\frac{|\mathcal{P} \cap V_N|}{|V_N|} = \frac{7}{8} + 2^{-4N}.$$

Proof. Since $\pi_2 : K \twoheadrightarrow \text{SL}_4(R)$ is onto, the K -orbits on V are the $\text{SL}_4(R)$ -orbits. As R is Euclidean, elementary row reduction takes any primitive vector to $e = (1, 0, 0, 0)$, proving transitivity.

Write $P_N = |\mathcal{P} \cap V_N|$, with $P_0 = 0$. Partitioning $V_N \setminus \{0\}$ by the monic greatest common divisor of the four coordinates gives

$$2^{4N} - 1 = \sum_{d=0}^{N-1} 2^d P_{N-d}.$$

Subtract twice the corresponding identity for $N-1$:

$$P_N = (2^{4N} - 1) - 2(2^{4N-4} - 1) = 7 \cdot 2^{4N-3} + 1. \quad \square$$

For $z = (\ell, q) \in X \times Y$, define

$$z[v] = (\ell(v), q(b(v))) \in \mathbb{F}_2^2.$$

We say that v *detects* z if $z[v] \neq (0, 0)$. The two bits are evaluated separately by $(v, 0)$ and $(0, b(v))$ in $D = V \oplus B$; adding them in $\mathbb{F}_2$ would allow cancellation.

Lemma 5.7. *If μ is a K -invariant probability measure on $X \times Y$, then*

$$\mu\{z : z[e] \neq (0, 0)\} \geq \frac{1}{7}(1 - \mu(\{(0, 0)\})). \quad (5.3)$$

Proof. By K -invariance, detection has the same probability at every primitive vector. Indeed, set

$$p_\mu = \mu\{z : z[e] \neq (0, 0)\}.$$

Then

$$(k \cdot z)[k \cdot v] = z[v] \implies \mu\{z : z[v] \neq (0, 0)\} = p_\mu \quad (v \in \mathcal{P}).$$

Suppose $z = (\ell, q)$ is detected by some $v \in V_N$. At least one of

$$v \mapsto \ell(v), \quad v \mapsto q(b(v))$$

is a nonzero Boolean polynomial on $V_N \cong \mathbb{F}_2^{4N}$. The first is linear and the second has degree at most 2, since $b(v) = v \otimes v$ and q is linear. Lemma 5.5 therefore gives at least 2^{4N-2} detecting vectors. By Lemma 5.6, exactly

$$|V_N| - |\mathcal{P} \cap V_N| = 2^{4N-3} - 1$$

of them can be nonprimitive, so at least $2^{4N-3} + 1$ primitive vectors detect z . Write

$$U_N = \{z : z[v] \neq (0, 0) \text{ for some } v \in V_N\}.$$

Integrating the pointwise count and using $|\mathcal{P} \cap V_N| = 7 \cdot 2^{4N-3} + 1$ yields

$$p_\mu |\mathcal{P} \cap V_N| \geq (2^{4N-3} + 1) \mu(U_N) \geq \frac{1}{7} |\mathcal{P} \cap V_N| \mu(U_N).$$

Since $V = \bigcup_N V_N$ and the $b(v)$ span B , the sets U_N increase to $(X \times Y) \setminus \{(0, 0)\}$. Divide by $|\mathcal{P} \cap V_N|$ and let $N \rightarrow \infty$ to obtain (5.3). $\square$

Remark 5.8. The bound $1/7 = (1/4 - 1/8)/(7/8)$ comes from restricting detection to primitive vectors: a nonzero Boolean polynomial of degree at most 2 detects at least $1/4$ of the cube, while the nonprimitive vectors have asymptotic density $1/8$.

More generally, in rank $j \geq 3$, the nonprimitive vectors have asymptotic density 2^{1-j} . The same argument gives the bound

$$c_j = \frac{1/4 - 2^{1-j}}{1 - 2^{1-j}} = \frac{2^{j-3} - 1}{2^{j-1} - 1}.$$

It is positive exactly when $j \geq 4$, with $c_4 = 1/7$. Thus 4 is the smallest rank for which this argument yields a spectral gap.

Proposition 5.9. *The group Λ has property (T).*

Proof. Let μ be a K -invariant Borel probability measure on $\widehat{D} = X \times Y$, and choose

$$d_1 = (e, 0), \quad d_2 = (0, b(e)).$$

Writing $\mathbf{1}_S$ for the indicator of a set S , a character $z = (\ell, q) \in \widehat{D} = X \times Y$ satisfies

$$|z(d_1) - 1|^2 + |z(d_2) - 1|^2 \geq 4 \mathbf{1}_{\{z[e] \neq (0,0)\}}.$$

After integration, Lemma 5.7 gives

$$\sum_{j=1}^2 \int_{\widehat{D}} |\chi(d_j) - 1|^2 d\mu(\chi) \geq \frac{4}{7} (1 - \mu(\{1\})).$$

Thus (5.1) holds with $(A, H, J, c) = (D, K, \{d_1, d_2\}, 4/7)$, proving property (T) for Λ . $\square$

These rigidity results and the common group factor now give the promised two-group counterexample.

5.5. The basic counterexample.

Corollary 5.10 (Two-group counterexample). *The groups Λ and Γ_0 are nonisomorphic ICC property-(T) groups with isomorphic group von Neumann algebras.*

Proof. Proposition 4.3 shows that the groups are nonisomorphic and that their group von Neumann algebras are isomorphic. Propositions 5.3 and 5.9 show that Λ is ICC and has property (T). Applying Lemma 2.1 to the factor isomorphism gives the same properties for Γ_0 . $\square$

6. AN INFINITE FIBER OF THE GROUP-FACTOR FUNCTOR

To extend Corollary 5.10 to a countably infinite fiber, we keep the measured K -space $X \times Y$ fixed and apply the same binary carry after shifting each of the four coordinate sequences by n . The common measured action identifies the resulting group factors; an intrinsic finite-orbit invariant recovers the $4n$ carry-free coordinates. The groups Γ_n will also be pairwise commensurable.

6.1. Pulling back the carry. For $n \geq 0$, define

$$T_n : X \longrightarrow X, \quad (T_n \ell)(v) = \ell(t^n v).$$

The R -linearity of the K -action makes T_n continuous and K -equivariant. It is surjective: for $\lambda \in X$, extend the functional $t^n v \mapsto \lambda(v)$ on $t^n V$ to V . Its kernel is

$$Z_n = \{\ell \in X : \ell|_{t^n V} = 0\} \cong \text{Hom}_{\mathbb{F}_2}(V/t^n V, \mathbb{F}_2), \quad |Z_n| = 2^{4n}.$$

For $\ell, \ell' \in X$, put

$$c_n(\ell, \ell') = r_{T_n \ell, T_n \ell'} \in Y.$$

Equip $X \times Y$ with the operation

$$(\ell, q) \star_n (\ell', q') = (\ell + \ell', q + q' + c_n(\ell, \ell')). \quad (6.3)$$

To see the original four-point carry inside this formula, write

$$x_{i,j} = \ell(t^j e_i), \quad y_{i,j} = q(b(t^j e_i)) \quad (1 \leq i \leq 4, j \geq 0).$$

Since

$$c_n(\ell, \ell')(b(t^j e_i)) = x_{i,j+n} x'_{i,j+n},$$

the diagonal coordinate pair $(x_{i,j+n}, y_{i,j})$ adds by

$$(x, y) \diamond (x', y') = (x + x', y + y' + xx').$$

This is exactly the $\mathbb{Z}/4\mathbb{Z}$ carry from Section 4. The first n coordinates $x_{i,0}, \dots, x_{i,n-1}$ in each of the four directions never enter the carry.

For $n = 0$, Proposition 4.1 identifies this coordinate model with the function-space group C_0 of Section 4 via $\Phi_0(\ell, q) = F_{\ell,q}$. We use this as a standing identification. For every $n \geq 0$, write C_n for $X \times Y$ equipped with $\star_n$.

Lemma 6.1. *For every $n \geq 0$, the operation $\star_n$ makes C_n a compact abelian group with Haar probability measure m . The group K acts by continuous automorphisms through the coordinate action*

$$k \cdot (\ell, q) = (k \cdot \ell, k \cdot q),$$

which is independent of n .

Proof. The pulled-back carry c_n is continuous, symmetric, bilinear, and K -equivariant. Its cocycle identity makes (6.3) associative; symmetry makes it commutative. The identity and inverse are $(0, 0)$ and $(\ell, q + c_n(\ell, \ell))$, respectively. Thus C_n is a compact abelian group and the coordinate K -action is by continuous automorphisms.

Translation by (ℓ', q') has the form

$$(\ell, q) \mapsto (\ell + \ell', q + q' + c_n(\ell, \ell')).$$

It translates X and then translates each Y -fiber. Hence it preserves m , which is therefore the Haar measure of every C_n . $\square$

At this point the measured K -space, and therefore the resulting group factor, no longer depends on n . We must now recover the shifted carry from the compact group structure. Two maps isolate its effects: ρ_n describes the carry extension, while p_n identifies the $4n$ carry-free coordinates.

Define

$$\begin{aligned} \rho_n : C_n &\longrightarrow X, & \rho_n(\ell, q) &= \ell, \\ p_n : C_n &\longrightarrow C_0, & p_n(\ell, q) &= (T_n \ell, q). \end{aligned}$$

The first map exhibits C_n as an extension of X by Y . Since $c_n(\ell, \ell') = r_{T_n \ell, T_n \ell'}$, the second is a surjective, continuous K -equivariant homomorphism with

$$\ker p_n = Z_n \times \{0\}.$$

Thus p_n forgets exactly the $4n$ carry-free bits. Its dual will give the finite-index inclusion in Theorem 1.2.

A continuous linear section of T_n gives a noncanonical isomorphism of compact abelian groups

$$C_n \cong C_0 \times (\mathbb{Z}/2\mathbb{Z})^{4n}.$$

For $n > 0$, this splitting is not K -equivariant: the carry-free coordinates cannot be separated from the tail as K -modules. Subsection 6.5 shows that their K -action recovers n .

To recover the group structure forgotten by the common measure, set

$$E_n = \widehat{C_n}, \quad \Gamma_n = E_n \rtimes K.$$

Under the standing identification via Φ_0 , E_0 and Γ_0 agree with the groups defined in Section 4.

To describe the extension dual to ρ_n , define, for $v \in V$,

$$\iota_n(v)(\ell, q) = (-1)^{\ell(v)}.$$

For $\eta \in E_n$, define $\sigma_n(\eta) \in B$ by

$$\eta(0, q) = (-1)^{q(\sigma_n(\eta))} \quad (q \in Y).$$

This uniquely determines $\sigma_n(\eta)$, since restriction to $\{0\} \times Y$ is a character of Y and $\hat{Y} = B$.

The next lemma makes the shifted carry visible on the dual side: every E_n retains the original order-four elements, and their doubles record the shift t^n .

Lemma 6.2. *For every $n \geq 0$, the group E_n is countable and fits into the K -equivariant exact sequence*

$$0 \longrightarrow V \xrightarrow{\iota_n} E_n \xrightarrow{\sigma_n} B \longrightarrow 0. \quad (6.7)$$

For $v \in V$, define

$$\varepsilon_v^{(n)} = \varepsilon_v^{(0)} \circ p_n \in E_n.$$

Then

$$\sigma_n(\varepsilon_v^{(n)}) = b(v), \quad 2\varepsilon_v^{(n)} = \iota_n(t^n v). \quad (6.8)$$

In particular, $\varepsilon_e^{(n)}$ has order 4.

Proof. Pontryagin duality applied to the exact sequence of compact abelian K -groups

$$0 \longrightarrow Y \xrightarrow{q \mapsto (0, q)} C_n \xrightarrow{\rho_n} X \longrightarrow 0$$

gives (6.7); metrizability of C_n gives countability of E_n . The two formulas in (6.8) follow from

$$\begin{aligned} \varepsilon_v^{(n)}(0, q) &= \varepsilon_v^{(0)}(F_{0, q}) = (-1)^{q(b(v))}, \\ (\varepsilon_v^{(n)}(\ell, q))^2 &= (-1)^{(T_n \ell)(v)} = (-1)^{\ell(t^n v)}. \end{aligned}$$

Since $t^n e \neq 0$ and ι_n is injective, $\varepsilon_e^{(n)}$ has order 4. □

The essential point is the factor t^n in (6.8): the lift of $b(v)$ doubles to $\iota_n(t^n v)$. Thus n records where the carry begins. Subsection 6.5 will recover the $4n$ carry-free coordinates from the finite-orbit part of $E_n[2]/2E_n$.

6.2. The common group factor. The compact group law depends on n , but the measured K -space does not. Fourier transform therefore gives the same crossed product throughout the family.

Proposition 6.3. *For every $n \geq 0$, there is an isomorphism*

$$L(\Gamma_n) \cong L(\Lambda).$$

Proof. By Lemma 6.1, the identity $(C_n, m) \rightarrow (X \times Y, m)$ is a K -equivariant isomorphism of probability spaces. Equation (2.1) therefore gives

$$\begin{aligned} L(\Gamma_n) &\cong L^\infty(C_n, m) \rtimes K \\ &\cong L^\infty(X \times Y, m) \rtimes K \cong L(\Lambda). \end{aligned} \quad \square$$

6.3. ICC and property (T). Both rigidity properties transfer along the common group factor.

Proposition 6.4. *For every $n \geq 0$, the group Γ_n is ICC and has property (T).*

Proof. Propositions 5.3 and 5.9 show that Λ is ICC and has property (T). Proposition 6.3 identifies $L(\Gamma_n)$ with $L(\Lambda)$. The conclusion now follows from Lemma 2.1. □

6.4. Finite-index embeddings and commensurability.

Proposition 6.5 (Finite-index embeddings). *For every $n \geq 0$, the group Γ_n contains a subgroup isomorphic to Γ_0 of index 2^{4n} . Consequently, the groups $(\Gamma_n)_{n \geq 0}$ are mutually commensurable.*

Proof. Dualizing

$$0 \longrightarrow Z_n \times \{0\} \longrightarrow C_n \xrightarrow{p_n} C_0 \longrightarrow 0$$

gives a K -equivariant exact sequence

$$0 \longrightarrow E_0 \xrightarrow{p_n^*} E_n \xrightarrow{\text{res}_{Z_n}} \widehat{Z}_n \longrightarrow 0, \quad p_n^*(\eta) = \eta \circ p_n.$$

Here $\text{res}_{Z_n}(\eta)$ is the character $\ell \mapsto \eta(\ell, 0)$ of Z_n . Hence $p_n^*(E_0) \rtimes K \cong \Gamma_0$ is a subgroup of Γ_n of index

$$|\widehat{Z}_n| = |Z_n| = 2^{4n}.$$

Since every Γ_n contains a finite-index copy of Γ_0 , the groups are mutually commensurable. $\square$

6.5. Pairwise nonisomorphism. The parameter n is detected by the finite orbits on an intrinsic 2-torsion quotient. By (6.3), doubling in C_n is

$$2(\ell, q) = (0, c_n(\ell, \ell)).$$

Hence C_n and E_n have exponent dividing 4. Since K is torsion-free, the torsion elements of $\Gamma_n = E_n \rtimes K$ form the characteristic abelian subgroup

$$\text{Tor}(\Gamma_n) = E_n.$$

For an abelian group A , write

$$A[2] = \{a \in A : 2a = 0\}, \quad 2A = \{2a : a \in A\}.$$

Since $4E_n = 0$, conjugation induces an intrinsic action of Γ_n/E_n on $E_n[2]/2E_n$. Define

$$i(\Gamma_n) = |\{a \in E_n[2]/2E_n : |(\Gamma_n/E_n) \cdot a| < \infty\}|.$$

Because $E_n = \text{Tor}(\Gamma_n)$, this is an invariant of the abstract group. We will prove

$$i(\Gamma_n) = |V/t^n V| = 2^{4n},$$

so its first two values are 1 and 16. The calculation proceeds through the doubling map on E_n , with ι_n and σ_n as in (6.7).

Lemma 6.6. *There is a unique K -equivariant surjective $\mathbb{F}_2$ -linear map*

$$d : B \longrightarrow V, \quad d(b(v)) = v.$$

Proof. Fix an ordered $\mathbb{F}_2$ -basis $(u_i)_{i \geq 1}$ of V . Using the corresponding basis for B from Subsection 4.1, define d by

$$u_i \otimes u_i \longmapsto u_i, \quad u_i \otimes u_j + u_j \otimes u_i \longmapsto 0 \quad (i < j).$$

The expansion of $b(v)$ gives $d(b(v)) = v$. Since the elements $b(v)$ span B , this identity implies surjectivity, uniqueness, and K -equivariance: $d(k \cdot b(v)) = k \cdot v = k \cdot d(b(v))$. $\square$

Lemma 6.7. *For every $n \geq 0$ and $\eta \in E_n$,*

$$2\eta = \iota_n(t^n d(\sigma_n(\eta))). \tag{6.10}$$

Consequently,

$$2E_n = \iota_n(t^n V) \quad \text{and} \quad E_n[2] = \sigma_n^{-1}(\ker d). \tag{6.11}$$

Proof. Both sides of (6.10) are homomorphisms in η . By (6.8), they agree on $\varepsilon_v^{(n)}$; they both vanish on $\iota_n(V)$. These elements generate E_n , since their images under σ_n span B and $\ker \sigma_n = \iota_n(V)$. This proves the identity.

Surjectivity of d and σ_n gives $2E_n = \iota_n(t^n V)$. Since multiplication by t^n on V and ι_n are injective, (6.10) also gives $E_n[2] = \sigma_n^{-1}(\ker d)$. $\square$

Now set

$$A_n = E_n[2]/2E_n.$$

For its finite-orbit subgroup, write

$$A_n^{\text{fin}} = \{a \in A_n : K \cdot a \text{ is finite}\}.$$

Lemma 6.8. *For every $n \geq 0$, there is a K -equivariant exact sequence*

$$0 \longrightarrow V/t^n V \longrightarrow A_n \longrightarrow \ker d \longrightarrow 0. \quad (6.13)$$

The maps are

$$v + t^n V \longmapsto \iota_n(v) + 2E_n, \quad \eta + 2E_n \longmapsto \sigma_n(\eta).$$

Under the resulting identification of $V/t^n V$ with its image in A_n ,

$$A_n^{\text{fin}} = V/t^n V, \quad |A_n^{\text{fin}}| = 2^{4n}.$$

Proof. By (6.11), restriction of σ_n gives the exact sequence

$$0 \longrightarrow V \xrightarrow{\iota_n} E_n[2] \xrightarrow{\sigma_n} \ker d \longrightarrow 0.$$

Quotienting by $2E_n = \iota_n(t^n V)$ gives (6.13).

Every element of $V/t^n V$ has finite orbit. Conversely, if $a \in A_n$ has nonzero image in $\ker d$, that image has infinite K -orbit by Lemma 5.2 and $K \twoheadrightarrow Q$. Equivariance forces a to have infinite orbit as well. Thus $A_n^{\text{fin}} = V/t^n V$, which has order 2^{4n} because $V/t^n V \cong (R/(t^n))^4$. $\square$

Proposition 6.9. *The groups $(\Gamma_n)_{n \geq 0}$ are pairwise nonisomorphic.*

Proof. An isomorphism $\Gamma_n \cong \Gamma_m$ preserves the characteristic torsion subgroup and hence induces an isomorphism $A_n \cong A_m$ intertwining the quotient-group actions. It therefore identifies their finite-orbit subgroups. Lemma 6.8 gives

$$2^{4n} = |A_n^{\text{fin}}| = |A_m^{\text{fin}}| = 2^{4m},$$

so $n = m$. $\square$

Proof of Theorem 1.2. Propositions 6.3 and 6.9 show that the Γ_n are pairwise nonisomorphic and have group factors isomorphic to $L(\Lambda)$. By Lemma 6.2, every Γ_n contains an element of order 4, whereas Λ has none by Proposition 4.3. Thus $\Lambda \not\cong \Gamma_n$ for every n .

Propositions 5.3, 5.9, and 6.4 give the ICC and property-(T) assertions. Since the groups are countable, property (T) also implies finite generation. Finally, Proposition 6.5 realizes Γ_0 as a subgroup of Γ_n of index 2^{4n} , so the Γ_n are mutually commensurable. $\square$

Acknowledgments. We thank Sorin Popa for valuable comments on the historical context, the countability theorem, and the finite-to-one question, and François Charles and Cyril Houdayer for careful readings and helpful suggestions.

During the preparation of this manuscript, we learned of independent and concurrent work by Shuoxing Zhou also establishing a counterexample to Connes's rigidity conjecture, developed in part with the assistance of GPT-5.6 Sol.

REFERENCES

- [BHV08] B. Bekka, P. de la Harpe, and A. Valette, *Kazhdan's Property (T)*, New Mathematical Monographs **11**, Cambridge University Press, 2008.
- [CIOS23] I. Chifan, A. Ioana, D. Osin, and B. Sun, *Wreath-like products of groups and their von Neumann algebras I: W^* -superrigidity*, Ann. of Math. (2) **198** (2023), no. 3, 1261–1303.
- [Con76] A. Connes, *Classification of injective factors. Cases II_1 , II_∞ , III_λ , $\lambda \neq 1$* , Ann. of Math. (2) **104** (1976), no. 1, 73–115.
- [Con80] A. Connes, *A factor of type II_1 with countable fundamental group*, J. Operator Theory **4** (1980), no. 1, 151–153.
- [Con82] A. Connes, *Classification des facteurs*, in *Operator Algebras and Applications, Part 2 (Kingston, Ont., 1980)*, Proc. Sympos. Pure Math. **38**, American Mathematical Society, Providence, R.I., 1982, 43–109, doi:10.1090/pspum/038.2/679497.
- [Con94] A. Connes, *Noncommutative Geometry*, Academic Press, San Diego, 1994.

- [CJ85] A. Connes and V. F. R. Jones, *Property T for von Neumann algebras*, Bull. London Math. Soc. **17** (1985), no. 1, 57–62.
- [CT11] Y. Cornuier and R. Tessera, *A characterization of relative Kazhdan property T for semidirect products with abelian groups*, Ergodic Theory Dynam. Systems **31** (2011), no. 3, 793–805.
- [EJ10] M. Ershov and A. Jaikin-Zapirain, *Property (T) for noncommutative universal lattices*, Invent. Math. **179** (2010), no. 2, 303–347.
- [Fur99] A. Furman, *Orbit equivalence rigidity*, Ann. of Math. (2) **150** (1999), no. 3, 1083–1108, [Annals of Mathematics](#).
- [Ioa11] A. Ioana, *W^* -superrigidity for Bernoulli actions of property (T) groups*, J. Amer. Math. Soc. **24** (2011), no. 4, 1175–1226, [doi:10.1090/S0894-0347-2011-00706-6](#).
- [Ioa18] A. Ioana, *Rigidity for von Neumann algebras*, in *Proceedings of the International Congress of Mathematicians 2018*, Vol. III, World Scientific, 2018, 1639–1672.
- [IPV13] A. Ioana, S. Popa, and S. Vaes, *A class of superrigid group von Neumann algebras*, Ann. of Math. (2) **178** (2013), no. 1, 231–286.
- [Kaz67] D. A. Kazhdan, *Connection of the dual space of a group with the structure of its closed subgroups*, Funct. Anal. Appl. **1** (1967), no. 1, 63–65.
- [Min87] H. Minkowski, *Ueber den arithmetischen Begriff der Aequivalenz und über die endlichen Gruppen linearer ganzzahliger Substitutionen*, J. Reine Angew. Math. **100** (1887), 449–458.
- [MN43] F. J. Murray and J. von Neumann, *On rings of operators. IV*, Ann. of Math. (2) **44** (1943), 716–808.
- [Pop06a] S. Popa, *Some rigidity results for non-commutative Bernoulli shifts*, J. Funct. Anal. **230** (2006), no. 2, 273–328.
- [Pop06b] S. Popa, *On a class of type II_1 factors with Betti numbers invariants*, Ann. of Math. (2) **163** (2006), no. 3, 809–899.
- [Pop06c] S. Popa, *Strong rigidity of II_1 factors arising from malleable actions of w -rigid groups, I*, Invent. Math. **165** (2006), no. 2, 369–408.
- [Pop06d] S. Popa, *Strong rigidity of II_1 factors arising from malleable actions of w -rigid groups, II*, Invent. Math. **165** (2006), no. 2, 409–451.
- [Pop07] S. Popa, *Deformation and rigidity for group actions and von Neumann algebras*, in *Proceedings of the International Congress of Mathematicians (Madrid, 2006)*, Vol. I, European Mathematical Society, Zürich, 2007, 445–477, [doi:10.4171/022-1/18](#).
- [Pop13] S. Popa, *Some open problems in W^* -rigidity*, problem list, Paris, June 2013, <https://www.math.ucla.edu/~popa/ProblemsJune2013.pdf>.
- [Sus77] A. A. Suslin, *On the structure of the special linear group over polynomial rings*, Math. USSR-Izv. **11** (1977), no. 2, 221–238.
- [Vae10] S. Vaes, *Rigidity for von Neumann algebras and their invariants*, in *Proceedings of the International Congress of Mathematicians (Hyderabad, India, 2010)*, Vol. III, Hindustan Book Agency, 2010, 1624–1650, [arXiv:1008.3610](#).

CHAPTER 5

Circuit and Formula Lower Bounds for the Permanent

ABSTRACT. We study the exact symbolic computation of the $n \times n$ permanent over $\mathbb{C}$ by arithmetic circuits and formulas. We prove two lower bounds.

- Division-free circuits with unrestricted reuse of intermediate values require $\Omega(n^2 \log \log n)$ arithmetic gates.
- Arithmetic formulas require $\Omega(n^4 / \log n)$ variable-labeled leaves, even when division is allowed provided all denominators are nonzero rational functions.

The proof of the circuit lower bound constructs an affine specialization whose gradient vanishes on a sufficiently low-dimensional set, then combines simultaneous differentiation with a geometric degree bound. The formula lower bound proof identifies many algebraically independent coefficients associated with a short matching of matrix entries and adds the resulting requirements over entry-disjoint matchings. We also explain how both arguments use properties of the permanent and do not directly imply a similar result for the determinant.

CONTENTS

1. Introduction
 2. Overview of the circuit lower bound
 3. Overview of the formula lower bounds
 4. A geometric circuit bound
 5. Critical loci of minor sums
 6. An affine specialization of the permanent
 7. Circuit parameters and the lower bound
 8. Coefficient transcendence degree
 9. Independent coefficients from a matching
 10. Formula lower bounds with and without division
 11. Comparison with the determinant
 12. Related work
- References

1. INTRODUCTION

Fix an integer $n \geq 1$, and let $X = (x_{ij})_{i,j \in [n]}$ be an $n \times n$ matrix of algebraically independent variables over $\mathbb{C}$, where $[j] = \{1, \dots, j\}$. The *permanent* of X is the polynomial

$$\text{per}_n(X) = \sum_{\sigma \in S_n} \prod_{i=1}^n x_{i,\sigma(i)}.$$

Here S_n is the set of permutations of $[n]$. Thus a monomial of per_n chooses exactly one entry from each row and each column; equivalently, it records a perfect matching in the complete bipartite graph on the row and column indices. The closely related *determinant* is

$$\det_n(X) = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_{i=1}^n x_{i,\sigma(i)},$$

which differs from the permanent only through the signs of its monomials. Nevertheless, the determinant has polynomial-size arithmetic circuits [Ber84], whereas the permanent is complete for Valiant’s class VNP [Val79a]. Understanding whether the permanent admits polynomial-sized algebraic circuits is the central problem in algebraic complexity theory [BCS97, SY10].

Throughout, “compute” means compute the polynomial exactly as a formal expression over $\mathbb{C}$, not merely evaluate it on a particular input matrix. An *arithmetic circuit* is a finite directed acyclic graph with input gates labeled by individual variables or arbitrary complex constants, binary internal gates labeled by $+$, $-$, or $\times$, and one designated output gate. Intermediate values may be reused arbitrarily: a gate can feed its output to any number of later gates. The circuit is division-free, and its *size* is the number of arithmetic gates; input gates are not counted. We write $C(f)$ for the minimum size of a circuit whose output is the polynomial f . There are no restrictions on depth, intermediate degrees, fan-out, or cancellations.

An *arithmetic formula* has the same kinds of inputs and binary arithmetic gates, but its underlying computation is a rooted tree. Consequently, two uses of the same intermediate polynomial must be computed by separate subformulas. For a formula Φ , let $L(\Phi)$ be its total number of leaves, let $L_{\text{var}}(\Phi)$ count only the leaves labeled by variables, with repeated occurrences counted separately, and let $G(\Phi)$ be its number of internal arithmetic gates. A *formula with division* additionally permits $\div$ gates and is interpreted in the rational-function field $\mathbb{C}(x_{ij} : i, j \in [n])$. Such a formula is *valid* if every denominator is a nonzero element of this field. Its final output is required to equal the polynomial being computed; in particular, intermediate rational functions need not themselves be polynomials.

The permanent depends on all $N = n^2$ matrix variables, giving the elementary $\Omega(N)$ circuit lower bound. The classical division-free formula lower bounds for the permanent and determinant are $\Omega(n^3)$ [KS14, §3.2, Thm. 5]; for the determinant, Kalorkoti’s cubic lower bound holds even when division is allowed [Kal85]. Our first result improves the unrestricted, division-free circuit lower bound specifically for the permanent.

Theorem 1.1. *For every $n \geq 2^{16}$,*

$$(1.1) \quad C(\text{per}_n) \geq \frac{n^2}{144} (\log_2 \log_2 n - 3).$$

In particular, $C(\text{per}_n) = \Omega(n^2 \log \log n)$ and $C(\text{per}_n)/n^2 \rightarrow \infty$.

The next result concerns the more restrictive formula model. Its primary measure is the number of variable occurrences, which also lower-bounds the total number of leaves and vertices.

Theorem 1.2. *If $n \geq 32$ and a division-free arithmetic formula Φ computes per_n , then*

$$L_{\text{var}}(\Phi) \geq \frac{n^4}{128 \log_2 n}, \quad G(\Phi) \geq \frac{n^4}{256 \log_2 n}.$$

The variable-leaf bound also holds for the total numbers of leaves and vertices.

The same asymptotic variable-occurrence bound remains true when a formula may use valid divisions; the explicit constants change.

Theorem 1.3. *If $n \geq 32$ and a valid arithmetic formula Φ with division computes per_n , then*

$$L_{\text{var}}(\Phi) \geq \frac{n^4}{192 \log_2 n}, \quad G(\Phi) \geq \frac{n^4}{384 \log_2 n}.$$

Again, the variable-leaf bound holds for the total numbers of leaves and vertices.

In terms of the $N = n^2$ input variables, the circuit bound is $\Omega(N \log \log N)$, and both formula bounds are $\Omega(N^2 / \log N)$. The circuit and formula results use different arguments: circuits permit efficient simultaneous computation of first derivatives because intermediate values can be reused, whereas the formula bounds charge algebraically independent coefficients to distinct occurrences of selected input variables. The permanent-specific inputs required by these arguments are not available for the determinant; §11 identifies both obstructions.

2. OVERVIEW OF THE CIRCUIT LOWER BOUND

Our circuit lower bound combines a geometric measure of polynomial complexity with a permanent-specific construction. The measure applies to a homogeneous polynomial whose gradient vanishes on a sufficiently small set; the construction realizes such a polynomial by fixing and identifying inputs of per_n . We describe both components, illustrate the necessary cancellation on a four-by-four permanent, and give complete proofs in §4–§7.

Stage 1: Many solutions force many multiplication gates. Consider the homogeneous polynomial

$$H(z_1, \dots, z_k) = \frac{1}{d} \sum_{i=1}^k z_i^d, \quad \nabla H(z) = (z_1^{d-1}, \dots, z_k^{d-1}), \quad d \geq 2.$$

For a target whose coordinates are all nonzero, the gradient has $(d-1)^k$ distinct preimages. If affine operations are free, a circuit using q multiplication gates describes the same fiber by q quadratic gate equations and k affine output equations. Bézout’s inequality bounds its number of isolated solutions by 2^q . Reverse-mode differentiation computes the gradient of a size- L circuit with at most $3L$ multiplications [BS83]. Consequently,

$$(d-1)^k \leq 2^{3L}, \quad L \geq \frac{k \log_2(d-1)}{3}.$$

Section 4 supplies the full gate-equation and reverse-mode differentiation arguments.

The advantage of this argument is that it does not require the coordinates of the gradient to be individual powers. The same solution count holds for any homogeneous map $F : \mathbb{C}^k \rightarrow \mathbb{C}^k$ of degree $d-1$ whose only zero is the origin: its generic fiber consists of $(d-1)^k$ points. This extension, proved in Lem. 4.1, is what allows the degree argument to apply to a polynomial obtained from the permanent.

Stage 2: A small critical locus produces the required gradient map. Let $P \in \mathbb{C}[x_1, \dots, x_m]$ be a homogeneous polynomial of degree d . Its *critical locus* is

$$\text{Crit}(P) = \{x \in \mathbb{C}^m : \partial P / \partial x_1 = \dots = \partial P / \partial x_m = 0\}.$$

Its dimension measures the number of freely varying coordinates in the common zero set; its codimension is $m - \dim \text{Crit}(P)$. In particular, a critical locus of codimension at least k is small enough for a generic k -dimensional linear subspace to meet it only at the origin.

Choose a linear injection $W : \mathbb{C}^k \rightarrow \mathbb{C}^m$ whose image is such a subspace. Then $\nabla P(Wu) \neq 0$ whenever $u \neq 0$. A second generic linear map $A : \mathbb{C}^m \rightarrow \mathbb{C}^k$ preserves this nonvanishing on the entire family of gradient directions. Therefore

$$F(u) = A \nabla P(Wu)$$

is a homogeneous degree- $(d-1)$ map from $\mathbb{C}^k$ to itself with $F^{-1}(0) = \{0\}$. Lem. 4.2 justifies the two generic choices. Linear maps cost no multiplications, so combining slicing, reverse-mode differentiation, and the preceding Bézout bound gives

$$L \geq \frac{k \log_2(d-1)}{3}$$

for any circuit computing P . Proposition 4.3 records the same implication when P is a nonzero scalar multiple of an affine specialization of per_n : replacing the permanent's inputs by affine linear forms transfers its circuit to P without adding multiplication gates.

To obtain a superquadratic lower bound, it therefore suffices to construct a degree- d specialization with d growing with n and with $k = \Theta(n^2)$. A direct application to the full permanent does not supply the needed codimension. The remaining stages construct a different specialization whose critical locus can be controlled.

Stage 3: Control the critical locus of a sum of subpermanents. For a $t \times s$ matrix $X = (x_{ia})$ and $3 \leq d \leq \min\{t, s\}$, define

$$M_{t,s,d}(X) = \sum_{\substack{I \subseteq [t], J \subseteq [s] \\ |I|=|J|=d}} \text{per}(X_{I,J}).$$

Thus $M_{t,s,d}$ sums the weights of all size- d matchings between t row vertices and s column vertices. It is homogeneous of degree d in ts variables.

For each nonempty subset $B \subseteq [t]$, introduce the power sum

$$p_B(X) = \sum_{a=1}^s \prod_{i \in B} x_{ia}.$$

There are $2^t - 1$ such quantities, shared by every column. To differentiate $M_{t,s,d}$ with respect to x_{ia} , first insist that row i is matched to column a . The remaining $d - 1$ rows must be matched injectively to columns other than a . Inclusion–exclusion on partitions expresses this injectivity requirement as a polynomial in the $p_B(X)$ and the entries of column a .

Section 5 uses these derivative equations together with the first $d - 2$ elementary symmetric functions of each column. Once the $2^t - 1$ shared power sums and these $s(d - 2)$ additional quantities are specified, the critical-point equations leave only finitely many possibilities for the column entries. Algebraically, those equations reduce every $(d - 1)$ st power of a column coordinate to expressions of lower column degree. Repeated reduction makes the coordinate algebra finite over the ring generated by the specified quantities. This yields the dimension bound

$$\dim \text{Crit}(M_{t,s,d}) \leq 2^t - 1 + s(d - 2).$$

If b copies occupy disjoint variable blocks and receive nonzero scalar weights, their critical loci form a Cartesian product. Corollary 5.2 therefore bounds the resulting dimension by

$$b(2^t - 1 + s(d - 2)).$$

Stage 4: Realize the desired block polynomial inside the permanent. Partition $r = bt$ rows into blocks $R_1, \dots, R_b$ of size t , and let X be an $r \times s$ matrix of variables. We want the polynomial

$$P(X) = \sum_{h=1}^b \lambda_h M_{t,s,d}(X_{R_h,[s]}), \quad \lambda_h \neq 0,$$

so that the preceding critical-locus estimate applies. The challenge is to realize this specific sum as one permanent, rather than as a separate sum of circuits.

A four-by-four cancellation example. The required block-selective cancellation is already visible in

$$\text{per} \begin{pmatrix} u & v & 1 & 1 \\ w & z & 1 & 1 \\ p & q & 2 & -2 \\ r & s & 2 & -2 \end{pmatrix} = -8(uz + vw) + 2(ps + qr).$$

Expand according to the two rows assigned to the first two columns. If both chosen rows belong to the upper block, their variable contribution is $\text{per} \begin{pmatrix} u & v \\ w & z \end{pmatrix} = uz + vw$, while the complementary constant block has permanent

$$\text{per} \begin{pmatrix} 2 & -2 \\ 2 & -2 \end{pmatrix} = -8.$$

If both chosen rows belong to the lower block, the variable contribution is $ps + qr$, and the complementary constant block has permanent

$$\text{per} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} = 2.$$

If one chosen row comes from each block, the complementary constant block likewise contains one row of each type, and its permanent vanishes:

$$\text{per} \begin{pmatrix} 1 & 1 \\ 2 & -2 \end{pmatrix} = -2 + 2 = 0.$$

Thus every mixed-block contribution cancels, leaving precisely the two within-block permanents. The example has degree $d = 2$, so $\log_2(d - 1) = 0$ and it cannot itself produce a useful circuit lower bound. Its purpose is to exhibit the cancellation that the general construction must reproduce at a larger degree.

The general specialization. Lemma 6.1 constructs a constant $r \times (r - d)$ matrix U for which

$$B(X) = \begin{pmatrix} X & U \\ \mathbf{1} & 0 \end{pmatrix}$$

is square of size $r + s - d$; the bottom-left all-ones block has $s - d$ rows. Every nonzero term of $\text{per}(B(X))$ chooses exactly d entries from X . The contribution of a chosen set of d variable rows is controlled by a complementary permanent of U .

The columns of U are chosen using d th roots of unity. In commuting square-zero variables $z_1, \dots, z_r$, write $y_h = \sum_{i \in R_h} z_i$. For $h \geq 2$, set $Y = y_1 + \dots + y_{h-1}$, and let ζ be a primitive d th root of unity. The corresponding root-of-unity product has the form

$$\prod_{j=0}^{d-1} (Y + 2\zeta^j y_h) = Y^d + (-1)^{d+1} 2^d y_h^d.$$

Combining these factors across the blocks forces every contributing set of d variable rows to lie entirely inside a single R_h ; choices meeting several blocks cancel. All surviving block weights are nonzero. Thus $\text{per}(B(X))$ is a nonzero scalar multiple of $P(X)$. Section 6 proves the construction and computes its nonzero block weights.

Stage 5: Choose the degree and count the surviving variables. Section 7 takes

$$d = \left\lfloor \frac{\log_2 n}{4} \right\rfloor, \quad t = 4d, \quad b = \left\lfloor \frac{n}{2t} \right\rfloor, \quad r = bt, \quad s = n - r + d.$$

Then $B(X)$ has size n and X contains $m = rs = \Theta(n^2)$ independent variables. The choice $t \leq \log_2 n$ keeps 2^t at most n , while $d = t/4$ keeps the columnwise contribution $bs(d - 2)$ below $m/4$. The full critical-locus bound is therefore smaller than $m/2$, permitting a slice of dimension $k = \lfloor m/4 \rfloor = \Theta(n^2)$. Stage 2 now gives

$$\text{C}(\text{per}_n) \geq \frac{k \log_2(d - 1)}{3} = \Omega(n^2 \log \log n).$$

The quantitative bookkeeping in Proposition 7.1 gives the explicit constant in Theorem 1.1.

3. OVERVIEW OF THE FORMULA LOWER BOUNDS

The formula bounds exploit the fact that a formula is a tree: unlike a circuit, it cannot reuse a computation in several places. Our strategy identifies many small, disjoint groups of matrix entries whose influence on the permanent is individually rich. Each group forces many occurrences of its variables in any formula, and disjointness allows us to add these requirements. We first describe the measure of influence, then the matching construction that makes the measure large, and finally the packing and the extension to division.

Stage 1: Measuring information in a group of variables. Let Y be a nonempty set of variables of a polynomial $f \in \mathbb{C}[Y, Z]$, where Z denotes all remaining variables. Expanding in the variables of Y gives

$$f = \sum_{\alpha \in \mathbb{N}^Y} c_\alpha(Z) Y^\alpha, \quad c_\alpha(Z) \in \mathbb{C}[Z].$$

We measure how much independent information these coefficients contain by

$$\text{td}_Y(f) = \text{trdeg}_{\mathbb{C}} \mathbb{C}(c_\alpha(Z) : \alpha \in \mathbb{N}^Y).$$

Here polynomials $c_1, \dots, c_r \in \mathbb{C}[Z]$ are *algebraically independent* if no nonzero polynomial $H \in \mathbb{C}[T_1, \dots, T_r]$ satisfies $H(c_1, \dots, c_r) = 0$. For example, independent variables z_1, z_2 are algebraically independent, whereas z_1, z_1^2 are dependent because $T_2 - T_1^2$ vanishes after substituting these two polynomials. The transcendence degree above is the maximum number of algebraically independent coefficient polynomials. It counts independent parameters, not merely the number of distinct or linearly independent coefficients.

For a formula Φ , let $t_Y(\Phi)$ be the number of leaves labeled by variables from Y , counting repeated occurrences. Section 8 proves that every division-free formula computing f satisfies

$$(3.1) \quad \text{td}_Y(f) \leq 4t_Y(\Phi) - 2 \leq 4t_Y(\Phi), \quad t_Y(\Phi) \geq 1.$$

The reason is that any subformula containing no Y -variable computes a polynomial in Z . Along a path where only one child contains a Y -variable, the other child therefore modifies the marked computation by an affine map $u \mapsto Au + B$, with $A, B \in \mathbb{C}[Z]$. An entire such path remains a single affine map. Only a gate with two Y -containing children combines genuinely different marked computations. The tree connecting $t_Y(\Phi)$ marked leaves has at most $t_Y(\Phi) - 1$ such branching gates, and each gate introduces at most four coefficient parameters; the final affine map contributes two more. Consequently, all coefficients of f belong to a field generated by at most $4t_Y(\Phi) - 2$ quantities. This step turns algebraic independence into a direct lower bound on variable occurrences.

Stage 2: Obtaining quadratically many independent coefficients from a logarithmic-size matching. A matching Y in the $n \times n$ variable matrix X is a set of entries with no repeated row or column. Set

$$\ell = \lceil \log_2 n \rceil, \quad k = 2\ell, \quad m = n - k,$$

and choose a matching of k entries. For $n \geq 32$, we have $m \geq n/2$, so $m^2 = \Omega(n^2)$ even though $|Y| = O(\log n)$. After permuting rows and columns, suppose that the marked variables are $Y = \{y_i = x_{ii} : i \in [k]\}$. Since the permanent is multilinear, its expansion takes the particularly simple form

$$\text{per}_n(X) = \sum_{S \subseteq [k]} c_S(X \setminus Y) \prod_{i \in S} y_i.$$

The coefficient c_S sums the permutation terms using exactly the marked diagonal entries indexed by S . Equivalently, delete their rows and columns and set every other marked entry to zero.

The key assertion, proved in Section 9, is

$$(3.2) \quad \text{td}_Y(\text{per}_n) \geq m^2.$$

To see what must be shown, divide the k marked indices into two sets E and F of size ℓ . For each pair $0 \leq \alpha, \beta < m$, let $P(\alpha) \subseteq E$ and $Q(\beta) \subseteq F$ encode the binary digits of α and β , respectively, and set

$$S(\alpha, \beta) = [k] \setminus (P(\alpha) \cup Q(\beta)).$$

Because $m \leq 2^\ell$, these choices specify m^2 distinct coefficients $c_{S(\alpha, \beta)}$. Separate the marked and unmarked row and column indices by writing

$$X = \begin{pmatrix} D & U \\ V & W \end{pmatrix}, \quad W = (w_{ab})_{a, b \in [m]}.$$

Here D uses the k internal indices, W uses the m external indices, and U, V connect the two groups. We establish algebraic independence by showing that the square Jacobian

$$\left(\frac{\partial c_{S(\alpha, \beta)}}{\partial w_{ab}} \right)_{(\alpha, \beta), (a, b)}$$

is invertible at one explicit assignment of the unmarked variables. The characteristic-zero Jacobian criterion then proves (3.2).

The assignment uses distinct complex numbers $p_1, \dots, p_m$ and, independently, distinct complex numbers $q_1, \dots, q_m$. Set the internal off-diagonal entries to zero and every w_{ab} to one. If $e_u \in E$ and $f_v \in F$ correspond to binary positions $u, v \in \{0, \dots, \ell - 1\}$, assign

$$U_{e_u, b} = p_b^{2^u}, \quad U_{f_v, b} = 1, \quad V_{a, e_u} = 1, \quad V_{a, f_v} = q_a^{2^v}$$

to the internal-to-external and external-to-internal entries, respectively. The derivative with respect to w_{ab} fixes the external matching edge (a, b) . Because all other external entries are one, the remaining weighted choices separate into an external column choice depending only on p_b and an external row choice depending only on q_a . More precisely, the evaluated Jacobian has entries

$$L_{\alpha, \beta} g_\alpha(p_b) h_\beta(q_a), \quad L_{\alpha, \beta} \neq 0,$$

where g_α and h_β are univariate polynomials of exact degrees α and β , respectively. Thus the p -evaluation and q -evaluation matrices are Vandermonde matrices multiplied by invertible triangular change-of-basis matrices. Their Kronecker product, and hence the whole Jacobian, is invertible. The 5×5 example at the beginning of Section 9 illustrates this row-column separation before the general construction.

Stage 3: Packing matchings and summing their costs. One matching now forces at least $m^2/4$ occurrences of its k variables by (3.1) and (3.2). To obtain the full formula bound, Section 10 partitions many of the n^2 matrix entries into pairwise entry-disjoint matchings of size k . Index rows and columns by $\{0, \dots, n - 1\}$ and, for $0 \leq \tau < n$ and $0 \leq j < \lfloor n/k \rfloor$, define

$$Y_{\tau, j} = \{x_{jk+r, (jk+r+\tau) \bmod n} : 0 \leq r < k\}.$$

Each $Y_{\tau, j}$ is a matching. The row determines j , while the column-minus-row difference modulo n determines τ , so no matrix entry belongs to two of these matchings. Their number is

$$\nu = n \left\lfloor \frac{n}{k} \right\rfloor \geq \frac{n^2}{2k}.$$

Since their variable sets are disjoint, each variable-labeled leaf is charged at most once. Therefore

$$L_{\text{var}}(\Phi) \geq \sum_{\tau, j} t_{Y_{\tau, j}}(\Phi) \geq \frac{\nu m^2}{4} = \Omega\left(\frac{n^4}{\log n}\right).$$

Keeping the explicit inequalities $m \geq n/2$ and $k \leq 4 \log_2 n$ gives the constant in Theorem 1.2. The passage to internal-gate bounds uses only the binary-tree identity $G(\Phi) = L(\Phi) - 1$.

Stage 4: Allowing valid division. The matching construction and its Jacobian concern the permanent itself, so they remain unchanged when the formula contains division gates. Only the occurrence-charging statement requires a new argument. Put $R = \mathbb{C}(Z)$, the field of rational functions in the unmarked variables. Along a path with one marked child, an unmarked child now computes an element of R , and all four arithmetic operations transform the marked value by a fractional linear map

$$u \mapsto \frac{au + b}{cu + d}, \quad a, b, c, d \in R.$$

Validity ensures that the actual denominator is nonzero as a rational function. After dividing the representing matrix by one of its nonzero entries, the map is described by at most three parameters; the matrix is not required to be invertible. A gate combining two marked subformulas therefore contributes at most six parameters, and the final map contributes at most three. This places the output in $K(Y)$ for a subfield $K \subseteq R$ generated over $\mathbb{C}$ by at most $6t_Y(\Phi) - 3$ elements.

There is one additional subtlety: membership in $K(Y)$ alone does not immediately say that the coefficients of the output belong to K . However, the output is also the polynomial $\text{per}_n \in R[Y]$, and the elementary field-intersection identity

$$K(Y) \cap R[Y] = K[Y]$$

recovers precisely that conclusion. Consequently, Section 10 establishes

$$\text{td}_Y(\text{per}_n) \leq 6t_Y(\Phi) - 3 < 6t_Y(\Phi).$$

Applying this inequality to the same ν disjoint matchings gives

$$L_{\text{var}}(\Phi) \geq \frac{\nu m^2}{6} = \Omega\left(\frac{n^4}{\log n}\right),$$

with the explicit constant of Theorem 1.3. The Jacobian specialization is applied only to polynomial coefficients of the permanent: a formula denominator need not remain nonzero at that numerical assignment.

4. A GEOMETRIC CIRCUIT BOUND

Our goal in this section is to convert a geometric property of a homogeneous polynomial into an arithmetic circuit lower bound. For a polynomial $P \in \mathbb{C}[x_1, \dots, x_m]$, its gradient and critical locus are

$$\nabla P(x) = \left(\frac{\partial P}{\partial x_1}(x), \dots, \frac{\partial P}{\partial x_m}(x) \right), \quad \text{Crit}(P) = \{x \in \mathbb{C}^m : \nabla P(x) = 0\}.$$

The critical locus is an affine algebraic set: it consists of the common solutions to finitely many polynomial equations. Its *dimension* is, informally, the largest number of algebraically independent parameters that can vary on one of its components; its *codimension* in $\mathbb{C}^m$ is $m - \dim \text{Crit}(P)$. We will show that when P has degree d and its critical locus has codimension at least k , computing P requires at least $k \log_2(d-1)/3$ multiplication gates. The bound then transfers to the permanent whenever P is an affine specialization of it.

The argument has three components. First, we bound the complexity of a homogeneous square polynomial map whose only zero is the origin. Second, we obtain such a map from the gradient of P by restricting its input to a suitable k -dimensional linear subspace and taking k linear combinations of its outputs. Finally, we use reverse-mode differentiation to compute this restricted gradient from a circuit for P .

Throughout the section, we use the *multiplicative-complexity model*: affine combinations of inputs and previously computed results are free, while each binary multiplication costs one. Since an ordinary size- L arithmetic circuit contains at most L multiplication gates, its multiplicative complexity is at most L . Treating affine operations as free will also let us perform linear changes of variables and linear combinations of outputs without affecting our bounds.

The first lemma formalizes the gradient warm-up in §2. Its hypothesis that the origin is the only common zero forces a generic fiber of a homogeneous degree- e polynomial map to have e^k solutions. Conversely, a circuit using q multiplications can describe the same solutions with only q quadratic equations, which have at most 2^q isolated solutions.

Lemma 4.1. *Let $F = (F_1, \dots, F_k) : \mathbb{C}^k \rightarrow \mathbb{C}^k$ have homogeneous coordinates of a common degree $e \geq 1$. If $F^{-1}(0) = \{0\}$ and a circuit in this model computes all F_i using q multiplications, then $e^k \leq 2^q$.*

Proof. We first count the solutions of $F(u) = a$ for a suitably chosen target $a = (a_1, \dots, a_k)$. Introducing one additional variable z , homogenize these equations to obtain

$$(4.1) \quad F_i(u) - a_i z^e = 0, \quad i \in [k],$$

in projective space $\mathbb{P}^k$. A point of $\mathbb{P}^k$ is a nonzero vector $(u_1, \dots, u_k, z)$ considered up to multiplication by a nonzero scalar. Points with $z \neq 0$ can be normalized to $z = 1$ and are exactly the usual affine solutions of $F(u) = a$. The remaining points, with $z = 0$, are the potential solutions “at infinity.”

There are no such solutions at infinity: if $[u_1 : \cdots : u_k : 0]$ satisfied (4.1), then $u \neq 0$ and $F(u) = 0$, contradicting the hypothesis. Moreover, every positive-dimensional projective algebraic set intersects every projective hyperplane. Thus a positive-dimensional common zero set of (4.1) would intersect the hyperplane $z = 0$, which we have just ruled out. The intersection is therefore zero-dimensional. Projective Bézout now says that the total number of its points, counted with their algebraic multiplicities, is the product of the k defining degrees:

$$\underbrace{e \cdot e \cdots e}_{k \text{ factors}} = e^k.$$

Since there are no points at infinity, this is also the number of affine solutions counted with multiplicity.

To obtain e^k *distinct* solutions, rather than merely e^k solutions counted with multiplicity, we verify that a generic target a has a reduced fiber. Put

$$R = \mathbb{C}[u_1, \dots, u_k], \quad B = \mathbb{C}[F_1, \dots, F_k] \subseteq R.$$

Because the origin is the only common zero of the F_i , the Nullstellensatz gives

$$(u_1, \dots, u_k)^h \subseteq (F_1, \dots, F_k)$$

for some integer h . In particular, every monomial of total degree at least h belongs to the ideal generated by the F_i . Since that ideal is homogeneous, a monomial of degree $j \geq h$ can be written as a sum of terms $F_i G_i$, with each nonzero G_i homogeneous of degree $j - e < j$. Repeating this reduction expresses every element of R as a B -linear combination of the finitely many monomials of total degree less than h . Thus R is a finite B -module.

A finite ring extension preserves Krull dimension, so $\dim B = \dim R = k$. Since B is generated by the k elements $F_1, \dots, F_k$, this dimension equality also shows that these elements are algebraically independent. Consequently, $B \cong \mathbb{C}[y_1, \dots, y_k]$, and the original map $F : \mathbb{C}^k \rightarrow \mathbb{C}^k$ is a finite dominant morphism: finiteness gives finite fibers, and dominance means that the image is dense in the target. Characteristic zero rules out inseparability, so generic smoothness supplies a nonempty open set of targets a whose fibers are reduced, meaning that every solution has multiplicity one [Har77, Ch. III, Cor. 10.7]. Choose such an a . Its fiber consists of exactly e^k distinct points by the preceding Bézout count.

We next describe the same fiber using the multiplication gates of the circuit. Order those gates topologically and introduce a new variable v_j for the output of the j th gate. Because additions, subtractions, and scalar operations are free, its two inputs are affine functions $A_j(u, v_{<j})$ and $B_j(u, v_{<j})$ of the original inputs and the preceding multiplication outputs. Its gate equation is therefore

$$v_j = A_j(u, v_{<j})B_j(u, v_{<j}), \quad j \in [q].$$

Similarly, each circuit output is an affine function $H_i(u, v)$, so requiring that the output equal a adds the equations

$$H_i(u, v) = a_i, \quad i \in [k].$$

We have obtained q equations of degree at most two and k equations of degree at most one in the $q + k$ unknowns (u, v) . For every u , the gate equations determine the v_j successively and uniquely. Thus solutions of the complete system correspond bijectively to the e^k points of the selected fiber. The affine Bézout inequality [Hei83] bounds the number of isolated solutions by the product of the equation degrees, which is at most 2^q . Hence $e^k \leq 2^q$. $\square$

We now explain how to produce the kind of square polynomial map required by Lem. 4.1 from a gradient. The gradient of a degree- d homogeneous polynomial has degree $d - 1$, but it has m outputs and may vanish on a positive-dimensional critical locus. The next lemma removes these obstacles in two steps: restrict the input to a k -dimensional subspace that avoids nonzero critical points, and then project the m gradient coordinates down to k without introducing a new zero.

Lemma 4.2. *Let $P \in \mathbb{C}[x_1, \dots, x_m]$ be homogeneous of degree $d \geq 2$, and let $\delta \geq 0$ satisfy $\dim \text{Crit}(P) \leq \delta$. If $1 \leq k < m$ and $k + \delta \leq m$, there are linear maps $W : \mathbb{C}^k \rightarrow \mathbb{C}^m$ and*

$A : \mathbb{C}^m \rightarrow \mathbb{C}^k$, with W injective, such that the homogeneous map $F(u) = A\nabla P(Wu)$ has degree $d - 1$ and $F^{-1}(0) = \{0\}$.

Proof. Since P is homogeneous and $d \geq 2$, all coordinates of ∇P are homogeneous of degree $d - 1 \geq 1$. In particular, $\text{Crit}(P)$ is a cone: if x is critical, so is every scalar multiple of x . Its nonzero points can consequently be viewed as directions in projective space:

$$\mathbb{P} \text{Crit}(P) = \{[x] \in \mathbb{P}^{m-1} : x \in \text{Crit}(P) \setminus \{0\}\}.$$

Removing the scalar parameter decreases dimension by one, so this projectivized critical locus has dimension at most $\delta - 1$. If $\text{Crit}(P) = \{0\}$, its projectivization is empty and the avoidance condition below is automatic.

We use the standard projective avoidance principle: if an algebraic subset of $\mathbb{P}^N$ has dimension at most r , then a generic projective l -plane misses that subset whenever $l + r < N$. One way to see the dimension count is to consider pairs consisting of a point in the subset and an l -plane containing that point. For a fixed point, the condition that an l -plane contain it has codimension $N - l$ in the space of all l -planes. Allowing the point to vary introduces at most r parameters, so the planes meeting the subset still form a proper exceptional set whenever $r < N - l$ [Har77, Ch. I].

Apply this principle first in $\mathbb{P}^{m-1}$ with $l = k - 1$ and $r = \delta - 1$. The hypothesis $k + \delta \leq m$ gives

$$(k - 1) + (\delta - 1) < m - 1,$$

so some projective $(k - 1)$ -plane avoids $\mathbb{P} \text{Crit}(P)$. Such a plane is the projectivization of a k -dimensional linear subspace of $\mathbb{C}^m$. Choose an injective linear map $W : \mathbb{C}^k \rightarrow \mathbb{C}^m$ whose image is that subspace. Avoidance means exactly that

$$\nabla P(Wu) \neq 0 \quad \text{for every } u \neq 0.$$

At this point the gradient has no nonzero zero on the restricted input space, but it still has m output coordinates. Homogeneity and the displayed nonvanishing let us associate to every nonzero u the projective direction of its gradient:

$$[u] \mapsto [\nabla P(Wu)].$$

This is a well-defined morphism $\mathbb{P}^{k-1} \rightarrow \mathbb{P}^{m-1}$: replacing u by cu multiplies the gradient by c^{d-1} and therefore does not change its projective direction. Its image is a projective algebraic set of dimension at most $k - 1$.

Apply the avoidance principle again, now to this image and to projective planes of dimension $m - k - 1$. The required inequality is

$$(m - k - 1) + (k - 1) = m - 2 < m - 1.$$

Hence there is an $(m - k)$ -dimensional linear subspace $K \subseteq \mathbb{C}^m$ whose projectivization does not meet any gradient direction. Choose a rank- k linear map $A : \mathbb{C}^m \rightarrow \mathbb{C}^k$ with kernel K . If $A\nabla P(Wu) = 0$ for some $u \neq 0$, then the nonzero vector $\nabla P(Wu)$ belongs to K , contradicting the choice of K . Therefore

$$F(u) = A\nabla P(Wu) \implies F^{-1}(0) = \{0\}.$$

Its coordinates remain homogeneous of degree $d - 1$, as required. $\square$

Combining the preceding lemmas now gives a circuit criterion in terms of the codimension of the critical locus. The only additional ingredient is that all first derivatives of a circuit output can be computed together with only a constant-factor increase in the number of multiplication gates.

Proposition 4.3. *Suppose a homogeneous polynomial P of degree $d \geq 2$ in m variables is a nonzero scalar multiple of an affine specialization of per_n . If $\dim \text{Crit}(P) \leq m - k$ for some $1 \leq k < m$, then every arithmetic circuit of size L computing per_n satisfies*

$$(4.2) \quad k \log_2(d - 1) \leq 3L.$$

Proof. Start with a size- L circuit for per_n . Since P is a nonzero scalar multiple of an affine specialization, substituting affine linear forms for the input entries and applying the scalar factor produces a circuit for P . These operations are free in the multiplicative-complexity model, so the resulting circuit has at most L multiplication gates.

The Baur–Strassen differentiation theorem [BS83, Thm. 2] computes all m first derivatives of P simultaneously with at most three times as many multiplication gates. To see why the factor is three in this model, evaluate the original circuit in topological order, retaining the output of each multiplication gate. In the reverse traversal, the adjoint $\bar{v}$ of an intermediate value v records the derivative of the final output with respect to v . If a forward gate computes $v = ab$, the chain rule gives the updates

$$\bar{a} \leftarrow \bar{a} + \bar{v}b, \quad \bar{b} \leftarrow \bar{b} + \bar{v}a.$$

The forward product uses one multiplication, and these two updates use at most two more; additions and scalar operations remain free. Thus all of ∇P can be computed with at most $3L$ multiplications, regardless of the number m of derivatives.

Apply Lem. 4.2 with $\delta = m - k$. Its hypothesis is satisfied because $\dim \text{Crit}(P) \leq m - k$ and $k + (m - k) = m$. We obtain linear maps W and A for which

$$F(u) = A\nabla P(Wu)$$

is a homogeneous degree- $(d - 1)$ map $\mathbb{C}^k \rightarrow \mathbb{C}^k$ whose only zero is the origin. Precomposing the gradient circuit with W and taking the output combinations prescribed by A cost no multiplications. Therefore F is computable with at most $3L$ multiplication gates, and Lem. 4.1 gives

$$(d - 1)^k \leq 2^{3L}.$$

Taking base-two logarithms proves (4.2). $\square$

5. CRITICAL LOCI OF MINOR SUMS

The circuit bound from Prop. 4.3 becomes useful when we find a degree- d polynomial in many variables whose critical locus has comparatively small dimension. This section supplies the necessary estimate for a polynomial that sums all size- d matchings between t rows and s columns. The following section will realize several disjoint copies of this polynomial as an affine specialization of the permanent.

Let $X = (x_{ia})$ be a $t \times s$ matrix of variables, and suppose $1 \leq d \leq \min\{t, s\}$. Define

$$(5.1) \quad M_{t,s,d}(X) = \sum_{\substack{I \subseteq [t] \\ |I|=d}} \sum_{\substack{J \subseteq [s] \\ |J|=d}} \text{per}(X_{I,J}).$$

Every term of $\text{per}(X_{I,J})$ chooses a bijection from I to J . Equivalently, $M_{t,s,d}$ is the sum of the weights $\prod_{j \in I} x_{j,\phi(j)}$ over all d -element row sets I and all injections $\phi : I \hookrightarrow [s]$. Thus its monomials correspond precisely to matchings of d row–column pairs.

To describe the critical-point equations, for every nonempty $B \subseteq [t]$ introduce

$$p_B(X) = \sum_{a=1}^s \prod_{j \in B} x_{ja}.$$

For example, $p_{\{j\}}$ is the sum of row j , while $p_{\{j,j'\}}$ records the sum of products obtained by assigning both rows to the same column. There are $\rho = 2^t - 1$ such quantities. We will first treat them as freely chosen parameters and then impose their actual values $p_B(X)$. This enlargement can only increase the dimension being estimated, but makes the critical-point equations uniform from column to column.

The basic idea is especially transparent for $d = 2$. Fix a column a , write $u_j = x_{ja}$, and abbreviate $p_j = p_{\{j\}}(X)$. A matching counted by $\partial M_{t,s,2}/\partial x_{ia}$ already uses the pair (i, a) ; its other pair may use any row $j \neq i$ and any column other than a . Therefore

$$\frac{\partial M_{t,s,2}}{\partial x_{ia}} = \sum_{j \neq i} (p_j - u_j).$$

At a critical point, put $v_j = p_j - u_j$. Subtracting the equations for two different values of i shows that all v_j are equal. Each equation then reads $(t-1)v_j = 0$. Since $t \geq 2$ and the ground field has characteristic zero, every v_j vanishes. Thus, once the row sums p_j have been fixed, the entire column $(u_1, \dots, u_t)$ is determined.

For $d \geq 3$, the power sums alone do not determine a column. We will additionally retain its first $d-2$ elementary symmetric functions and prove that these data leave only finitely many possibilities. Counting the $2^t - 1$ shared power-sum parameters and the $d-2$ additional parameters for each of the s columns gives the desired bound.

Proposition 5.1. *If $3 \leq d \leq \min\{t, s\}$, then*

$$(5.2) \quad \dim \text{Crit}(M_{t,s,d}) \leq 2^t - 1 + s(d-2).$$

Proof. Put $q = d-1$, so $q \geq 2$ and $q < t$. The argument has three stages. First, we rewrite every first derivative using the power-sum parameters and one column of X . Next, we identify the highest-degree part of the resulting equations. Finally, we use that highest-degree part to show that, once $q-1 = d-2$ symmetric functions per column are fixed, all the remaining column coordinates admit only finitely many possibilities.

Fix a row i and a column a . Differentiating a matching monomial with respect to x_{ia} keeps precisely those matchings that use the pair (i, a) . Removing that pair leaves a q -element set of other rows, injected into the columns other than a . Consequently,

$$\frac{\partial M_{t,s,d}}{\partial x_{ia}}(X) = \sum_{\substack{I \subseteq [t] \setminus \{i\} \\ |I|=q}} \sum_{\phi: I \hookrightarrow [s] \setminus \{a\}} \prod_{j \in I} x_{j, \phi(j)}.$$

The obstacle is the injectivity requirement: independently summing over a column for each row would also count assignments in which several rows choose the same column. Inclusion–exclusion over these collisions gives a convenient expression in terms of the p_B .

For a finite set I , let $\Pi(I)$ be its set of partitions into nonempty blocks. For $\pi \in \Pi(I)$, define

$$\mu(\pi) = \prod_{B \in \pi} (-1)^{|B|-1} (|B|-1)!.$$

Here $\mu(\pi)$ is the usual Möbius coefficient for the partition lattice. Temporarily regard all p_B as independent parameters, and for a column vector $u = (u_1, \dots, u_t)$ define

$$(5.3) \quad Q_i(u; p) = \sum_{\substack{I \subseteq [t] \setminus \{i\} \\ |I|=q}} \sum_{\pi \in \Pi(I)} \mu(\pi) \prod_{B \in \pi} \left(p_B - \prod_{j \in B} u_j \right).$$

To verify the inclusion–exclusion explicitly, specialize to $u = X_{\bullet a} = (x_{1a}, \dots, x_{ta})$ and $p = p(X)$. For any nonempty block B ,

$$p_B(X) - \prod_{j \in B} x_{ja} = \sum_{c \neq a} \prod_{j \in B} x_{jc}.$$

The product of these expressions over $B \in \pi$ counts all maps $\phi: I \rightarrow [s] \setminus \{a\}$ that are constant on every block of π , with the matching weight $\prod_{j \in I} x_{j, \phi(j)}$. For a fixed map ϕ , its total coefficient in the partition sum is therefore

$$\sum_{\substack{\pi \in \Pi(I) \\ \pi \text{ refines the fiber partition of } \phi}} \mu(\pi).$$

This coefficient factors over the nonempty fibers C of ϕ . The factor for a fiber C is

$$\sum_{\pi \in \Pi(C)} \prod_{B \in \pi} (-1)^{|B|-1} (|B|-1)! = \sum_{\sigma \in S_C} \text{sgn}(\sigma).$$

Indeed, a permutation with cycle blocks B contributes the displayed sign, and there are $(|B|-1)!$ cycles on each fixed block B . The sum of permutation signs is 1 when $|C| = 1$ and 0 when

$|C| \geq 2$. Hence precisely the maps with singleton fibers—that is, the injections—survive. Summing over I proves

$$(5.4) \quad \frac{\partial M_{t,s,d}}{\partial x_{ia}}(X) = Q_i(X_{\bullet a}; p(X)).$$

We next identify the part of Q_i having highest degree in the column variables u , while treating every p_B as a coefficient of degree zero. For a fixed q -element set I , the only way to obtain column degree q is to take $-\prod_{j \in B} u_j$ from every block factor in (5.3). Thus the coefficient of $\prod_{j \in I} u_j$ in degree q is

$$(-1)^q \sum_{\pi \in \Pi(I)} \prod_{B \in \pi} (|B| - 1)! = (-1)^q q!.$$

For the equality, specify the cycle blocks of a permutation of I ; each prescribed block B can be arranged into a cycle in $(|B| - 1)!$ ways. Summing over partitions therefore counts all $q!$ permutations. Put $\alpha = (-1)^q q!$, and write $e_j(u)$ for the j th elementary symmetric polynomial in $u_1, \dots, u_t$, with $e_0(u) = 1$. Summing the degree- q monomials over all $I \subseteq [t] \setminus \{i\}$ gives

$$(5.5) \quad Q_i(u; p) = \alpha e_q(u_1, \dots, \hat{u}_i, \dots, u_t) + R_i(u; p), \quad \deg_u R_i \leq q - 1.$$

Here the hat means that coordinate u_i is omitted. In particular, although the lower-degree remainder can depend on the shared parameters p_B , the highest-degree part has a simple symmetric form independent of those parameters.

We now convert this structure into the promised parameter count. Introduce independent coordinates p_B and a separate column vector $u^{(a)} = (u_1^{(a)}, \dots, u_t^{(a)})$ for every $a \in [s]$. The equations $Q_i(u^{(a)}; p) = 0$ define an affine set whose coordinate ring is

$$\mathcal{A} = \mathbb{C}[p_B, u_i^{(a)}] / (Q_i(u^{(a)}; p) : 1 \leq i \leq t, 1 \leq a \leq s).$$

This is the incidence construction: its points record one common collection of power-sum parameters together with s columns satisfying the corresponding critical-point equations. A true critical point of $M_{t,s,d}$ gives such a point by taking $u^{(a)} = X_{\bullet a}$ and $p_B = p_B(X)$. We allow additional incidence points for which the p_B need not equal these actual power sums, since an upper bound for the larger set also bounds the critical locus.

In addition to the ρ power-sum parameters, retain the first $q - 1$ elementary symmetric functions of each column. To keep these parameters formally independent, introduce the polynomial ring

$$\mathcal{B} = \mathbb{C}[p_B, z_{a,j} : \emptyset \neq B \subseteq [t], 1 \leq a \leq s, 1 \leq j < q]$$

and map it into $\mathcal{A}$ by sending each p_B to its namesake and each $z_{a,j}$ to $e_j(u^{(a)})$. Equivalently, this map makes $\mathcal{A}$ into a $\mathcal{B}$ -algebra. The map need not be injective: any relations among the actual parameter values can only decrease the final dimension.

Our remaining task is to prove that $\mathcal{A}$ is generated by finitely many elements as a module over $\mathcal{B}$. Fix one column, suppress its index, and put $\kappa = t - q > 0$. Two elementary symmetric-polynomial identities give

$$e_q(u_1, \dots, \hat{u}_i, \dots, u_t) = \sum_{j=0}^q (-u_i)^{q-j} e_j(u), \quad \sum_{i=1}^t e_q(u_1, \dots, \hat{u}_i, \dots, u_t) = \kappa e_q(u).$$

For the first identity, expand $\prod_{h \neq i} (1 + u_h z) = \prod_h (1 + u_h z) / (1 + u_i z)$ and compare coefficients of z^q . For the second, each squarefree monomial of degree q occurs once for each of the $t - q = \kappa$ omitted coordinates outside its support.

The incidence equations and (5.5) say that $\alpha e_q(u_1, \dots, \hat{u}_i, \dots, u_t) + R_i(u; p) = 0$ for every i . Summing these equations and using the second elementary symmetric identity gives

$$\kappa \alpha e_q(u) + \sum_{h=1}^t R_h(u; p) = 0.$$

This eliminates the only elementary symmetric function $e_q(u)$ not already recorded among the base parameters $z_{a,j}$. Insert the first symmetric identity into the i th incidence equation, replace

$e_j(u)$ by $z_{a,j}$ for $1 \leq j < q$, and eliminate $e_q(u)$ using the summed equation. Since $\alpha(-1)^q = q!$, the resulting equality in $\mathcal{A}$ is

$$(5.6) \quad \kappa q! u_i^q = -\kappa \alpha \sum_{j=1}^{q-1} (-u_i)^{q-j} z_{a,j} - \kappa R_i(u; p) + \sum_{h=1}^t R_h(u; p).$$

Treating the p_B and $z_{a,j}$ as base coefficients, every term on the right has degree at most $q-1$ in the coordinates of the chosen column. Because $\kappa > 0$ and the ground field has characteristic zero, the leading coefficient $\kappa q!$ is invertible. Therefore (5.6) replaces u_i^q by a $\mathcal{B}$ -linear combination of monomials of strictly smaller total column degree.

Apply this replacement whenever any coordinate exponent is at least q . Each replacement decreases total degree in the corresponding column, so the process terminates. Repeating it for all s columns shows that the finitely many monomials

$$\prod_{a=1}^s \prod_{i=1}^t (u_i^{(a)})^{r_{a,i}}, \quad 0 \leq r_{a,i} < q,$$

generate $\mathcal{A}$ as a $\mathcal{B}$ -module. In particular, fixing all the parameters p_B and $z_{a,j}$ leaves a finite-dimensional coordinate algebra for the possible columns; this is the algebraic form of the claimed finite-choice property.

A finite algebra cannot have larger Krull dimension than its parameter ring: if the map $\mathcal{B} \rightarrow \mathcal{A}$ has kernel K , then $\mathcal{A}$ is finite over $\mathcal{B}/K$, and the standard dimension theorem for finite algebras gives $\dim \mathcal{A} = \dim(\mathcal{B}/K) \leq \dim \mathcal{B}$ [Eis95]. Since $\mathcal{B}$ is a polynomial ring in $\rho + s(q-1)$ independent parameters, we obtain

$$\dim \mathcal{A} \leq \dim \mathcal{B} = \rho + s(q-1);$$

Finally, to return from the enlarged incidence set to the actual critical locus, (5.4) induces a surjection from $\mathcal{A}$ onto $\mathbb{C}[X]/(\partial M_{t,s,d}/\partial x_{ia} : i, a)$ by sending p_B to $p_B(X)$ and $u_i^{(a)}$ to x_{ia} . The map is surjective because the images of the $u_i^{(a)}$ are all the matrix variables. Passing to a quotient cannot increase Krull dimension, and the dimension of an affine common-zero set equals the Krull dimension of its coordinate ring. Hence

$$\dim \text{Crit}(M_{t,s,d}) \leq \dim \mathcal{A} \leq \rho + s(q-1) = 2^t - 1 + s(d-2),$$

which proves (5.2). □

The circuit application requires several copies of this polynomial on disjoint sets of variables. Because their critical-point equations involve separate variable blocks, the individual dimension bounds add.

Corollary 5.2. *Suppose $3 \leq d \leq \min\{t, s\}$, let $X^{(1)}, \dots, X^{(b)}$ be disjoint $t \times s$ variable blocks, and let $\lambda_1, \dots, \lambda_b \in \mathbb{C}^\times$. Then*

$$P(X) = \sum_{h=1}^b \lambda_h M_{t,s,d}(X^{(h)}) \implies \dim \text{Crit}(P) \leq b(2^t - 1 + s(d-2)).$$

Proof. For a variable in block $X^{(h)}$, the corresponding derivative of P is λ_h times the derivative of $M_{t,s,d}(X^{(h)})$. Since $\lambda_h \neq 0$, this derivative vanishes exactly when the corresponding derivative of the block polynomial vanishes. Different blocks share no variables, and therefore

$$\text{Crit}(P) = \prod_{h=1}^b \text{Crit}(M_{t,s,d}(X^{(h)})).$$

Dimensions add under products of complex affine algebraic sets. Applying Prop. 5.1 to each factor gives the stated bound. □

6. AN AFFINE SPECIALIZATION OF THE PERMANENT

Corollary 5.2 bounds the critical locus of a sum of minor-sum polynomials supported on disjoint row blocks. To apply Prop. 4.3, however, we must first realize such a sum as a specialization of one permanent. Our immediate objective is therefore to construct a matrix of constants whose permanent minors cancel whenever the omitted rows come from several different blocks, but remain nonzero when all the omitted rows come from one block. Appending these constant columns to a matrix of variables will then produce exactly the blockwise minor sums from the preceding section.

Let $b, t, s \geq 1$, partition $[r]$ into b disjoint blocks $R_1, \dots, R_b$ of size t , where $r = bt$, and suppose that $1 \leq d \leq \min\{t, s\}$. For $T \subseteq [r]$, let $\mathbf{1}_T \in \mathbb{C}^r$ denote the vector that is 1 on T and 0 elsewhere. Choose a primitive d th root of unity ζ , and set

$$\theta = (-1)^{d+1} 2^d.$$

The following two families of constant columns form a matrix $U \in \mathbb{C}^{r \times (r-d)}$:

- for each $h \in [b]$, take $t - d$ copies of $\mathbf{1}_{R_h}$;
- for each $h = 2, \dots, b$ and $j = 0, \dots, d - 1$, take $\mathbf{1}_{R_1 \cup \dots \cup R_{h-1}} + 2\zeta^j \mathbf{1}_{R_h}$.

The first family contributes $b(t - d)$ columns and the second contributes $(b - 1)d$ columns, for a total of $b(t - d) + (b - 1)d = bt - d = r - d$. If $t = d$, the first family is empty; if $b = 1$, the second family is empty.

Lemma 6.1. *There are $\lambda_1, \dots, \lambda_b \in \mathbb{C}^\times$ such that $\sum_h \lambda_h M_{t,s,d}(X_{R_h,[s]})$ is, up to a nonzero scalar, an affine specialization of per_{r+s-d} . More precisely, for the block matrix*

$$B(X) = \begin{pmatrix} X & U \\ \mathbf{1} & 0 \end{pmatrix}, \quad X = (x_{ia})_{i \in [r], a \in [s]},$$

where the lower-left block is the $(s - d) \times s$ all-ones matrix and the lower-right block is the $(s - d) \times (r - d)$ zero matrix, one has

$$(6.1) \quad \text{per}(B(X)) = c \sum_{h=1}^b \lambda_h M_{t,s,d}(X_{R_h,[s]}), \quad c = (s - d)!(t - d)!(t!)^{b-1} \neq 0.$$

Proof. First observe why minors of U determine the desired specialization. The matrix $B(X)$ has $r + s - d$ rows and columns. Because its lower-right block is zero, each of its $s - d$ lower rows must be matched to one of the first s columns in every nonzero permanent term. This leaves exactly d of those columns to be matched to upper rows. Let $J \subseteq [s]$ be those d columns and $I \subseteq [r]$ the d upper rows matched to them. The remaining $r - d$ upper rows must be matched to the $r - d$ columns of U . Finally, the remaining lower-left submatrix is an $(s - d) \times (s - d)$ all-ones matrix, whose permanent is $(s - d)!$. Summing over I and J consequently gives

$$(6.2) \quad \text{per}(B(X)) = (s - d)! \sum_{\substack{I \subseteq [r], J \subseteq [s] \\ |I|=|J|=d}} \text{per}(U_{I^c,[r-d]}) \text{per}(X_{I,J}).$$

It remains to show that the constant coefficient $\text{per}(U_{I^c,[r-d]})$ vanishes unless I is contained in one row block, and to compute its value in the remaining cases.

To keep track of these minors, introduce the commutative square-zero algebra

$$\mathcal{S} = \mathbb{C}[z_1, \dots, z_r] / (z_1^2, \dots, z_r^2).$$

Its monomials $z_K = \prod_{i \in K} z_i$, indexed by subsets $K \subseteq [r]$, form a vector-space basis: a term containing any z_i twice vanishes. Consequently, if A is an $r \times q$ matrix and $|K| = q$, expansion of its column forms gives

$$(6.3) \quad [z_K] \prod_{j=1}^q \left(\sum_{i=1}^r A_{ij} z_i \right) = \text{per}(A_{K,[q]}).$$

Indeed, each surviving contribution chooses a distinct row for each column, and the contributions producing z_K are precisely the bijections between the q columns and the rows in K . When $q = 0$,

both the empty product and the permanent of the empty matrix are 1. Thus the minors of U are encoded by a single product of its column forms.

For each block, put

$$y_h = \sum_{i \in R_h} z_i.$$

Because the variables commute and have square zero,

$$(6.4) \quad y_h^j = j! \sum_{\substack{K \subseteq R_h \\ |K|=j}} z_K \quad (0 \leq j \leq t), \quad y_h^{t+1} = 0.$$

In particular, $y_h^t = t! \prod_{i \in R_h} z_i$: the power y_h^t saturates the entire block, so multiplying it by any additional z_i with $i \in R_h$ gives zero.

The column form of $\mathbf{1}_{R_h}$ is y_h . If $Y_{h-1} = y_1 + \cdots + y_{h-1}$, the column form of $\mathbf{1}_{R_1 \cup \cdots \cup R_{h-1}} + 2\zeta^j \mathbf{1}_{R_h}$ is $Y_{h-1} + 2\zeta^j y_h$. The ordinary factorization $\prod_{j=0}^{d-1} (A + \zeta^j B) = A^d - (-B)^d$, applied in the commutative algebra $\mathcal{S}$, therefore gives

$$\prod_{j=0}^{d-1} (Y_{h-1} + 2\zeta^j y_h) = Y_{h-1}^d + \theta y_h^d.$$

All mixed powers of Y_{h-1} and y_h have canceled. Multiplying over both families of columns now yields

$$(6.5) \quad \prod_{j=1}^{r-d} \left(\sum_{i=1}^r U_{ij} z_i \right) = \left(\prod_{h=1}^b y_h^{t-d} \right) \prod_{h=2}^b (Y_{h-1}^d + \theta y_h^d).$$

For example, when $b = 2$, the product of all column forms of U is

$$y_1^{t-d} y_2^{t-d} (y_1^d + \theta y_2^d) = y_1^t y_2^{t-d} + \theta y_1^{t-d} y_2^t.$$

In the first term all rows of R_1 appear and exactly d rows of R_2 are omitted; in the second, the roles of the blocks are reversed. In particular, no term can omit rows from both blocks. Taking $t = s = d = 2$ gives $\theta = -4$ and recovers the 4×4 cancellation in §2.

We next verify that the same block-selection property persists for any number of blocks. For $1 \leq h \leq b$, let

$$C_h = \left(\prod_{g=1}^h y_g^{t-d} \right) \prod_{g=2}^h (Y_{g-1}^d + \theta y_g^d).$$

We claim that there are coefficients $\lambda_i^{(h)}$ such that

$$C_h = \sum_{i=1}^h \lambda_i^{(h)} y_i^{t-d} \prod_{\substack{g \in [h] \\ g \neq i}} y_g^t.$$

For $h = 1$, this is the identity $C_1 = y_1^{t-d}$, with $\lambda_1^{(1)} = 1$. Suppose the identity holds for h blocks. By definition,

$$C_{h+1} = C_h y_{h+1}^{t-d} (Y_h^d + \theta y_{h+1}^d).$$

The term containing θy_{h+1}^d saturates the new block and multiplies each previous coefficient by θ . In the term containing Y_h^d , every old block other than R_i is already saturated in the i th summand. Therefore all mixed terms in $Y_h^d = (y_1 + \cdots + y_h)^d$ vanish after multiplication by that summand, and the only surviving term is y_i^d . Hence

$$\left(\sum_{i=1}^h \lambda_i^{(h)} y_i^{t-d} \prod_{\substack{g \in [h] \\ g \neq i}} y_g^t \right) Y_h^d = \left(\sum_{i=1}^h \lambda_i^{(h)} \right) \prod_{g=1}^h y_g^t.$$

It follows that the coefficients satisfy

$$\lambda_i^{(h+1)} = \theta \lambda_i^{(h)} \quad (1 \leq i \leq h), \quad \lambda_{h+1}^{(h+1)} = \sum_{i=1}^h \lambda_i^{(h)}.$$

If $S_h = \sum_{i=1}^h \lambda_i^{(h)}$, then $S_1 = 1$ and $S_{h+1} = (1 + \theta)S_h$, so $S_h = (1 + \theta)^{h-1}$. Consequently, at $h = b$ the coefficients are

$$\lambda_1 = \theta^{b-1}, \quad \lambda_h = \theta^{b-h}(1 + \theta)^{h-2} \quad (2 \leq h \leq b).$$

They are all nonzero: $\theta \neq 0$ and $|\theta| = 2^d > 1$ imply $1 + \theta \neq 0$. Combining the induction with (6.5) proves

$$(6.6) \quad \prod_{j=1}^{r-d} \left(\sum_{i=1}^r U_{ij} z_i \right) = \sum_{h=1}^b \lambda_h y_h^{t-d} \prod_{g \neq h} y_g^t,$$

the promised expression in which exactly one row block is unsaturated.

Finally, let $I \subseteq [r]$ have size d . By (6.3), the coefficient of z_{I^c} on the left of (6.6) is $\text{per}(U_{I^c, [r-d]})$. On the right, the h th summand saturates every block except R_h . Thus it contributes to z_{I^c} only if $I \subseteq R_h$. When this occurs, (6.4) gives a factor $(t-d)!$ from y_h^{t-d} and a factor $t!$ from each of the other $b-1$ blocks. Therefore

$$\text{per}(U_{I^c, [r-d]}) = \begin{cases} \lambda_h (t-d)!(t!)^{b-1}, & I \subseteq R_h, \\ 0, & I \text{ meets more than one block.} \end{cases}$$

Substituting this identity into (6.2) and grouping the surviving terms by their block R_h gives

$$\text{per}(B(X)) = (s-d)!(t-d)!(t!)^{b-1} \sum_{h=1}^b \lambda_h \sum_{\substack{I \subseteq R_h, J \subseteq [s] \\ |I|=|J|=d}} \text{per}(X_{I,J}).$$

The inner sum is $M_{t,s,d}(X_{R_h, [s]})$, and the factorial prefactor is nonzero over $\mathbb{C}$. This proves (6.1) and realizes the blockwise polynomial needed for Cor. 5.2 as a specialization of one permanent. $\square$

7. CIRCUIT PARAMETERS AND THE LOWER BOUND

We now assemble the geometric gradient bound, the critical-locus estimate, and the block-selective specialization to prove Thm. 1.1. There are two quantitative requirements. First, the specialization must retain $m = \Theta(n^2)$ genuinely variable matrix entries. Second, its critical locus must have dimension bounded away from m , so that a linear slice of dimension $k = \Theta(m)$ avoids every nonzero critical point. Once these requirements hold, Prop. 4.3 supplies a lower bound proportional to $k \log_2(d-1)$.

The critical-locus bound from Cor. 5.2 has two costs per row block: $2^t - 1$ global power-sum parameters and $s(d-2)$ column parameters. Setting $t = 4d$ makes the second cost less than one quarter of the ts variables in each block. We will separately require $2^t - 1 \leq ts/4$, so the first cost also consumes at most one quarter. Choosing d on the order of $\log n$ will make these requirements compatible while preserving the logarithmic gain in the gradient bound.

Fix $d \geq 3$, and introduce the parameters

$$(7.1) \quad t = 4d, \quad b = \left\lfloor \frac{n}{2t} \right\rfloor, \quad r = bt, \quad s = n - r + d, \quad m = rs, \quad k = \left\lfloor \frac{m}{4} \right\rfloor.$$

Here t is the number of rows in each block, b is the number of blocks, and $r = bt$ is the total number of variable rows. Taking $b = \lfloor n/(2t) \rfloor$ places r near $n/2$, leaving roughly $n/2$ variable columns. More precisely, the choice $s = n - r + d$ makes the square matrix $B(X)$ of Lem. 6.1 have exactly

$$r + s - d = n$$

rows and columns. Its variable upper-left block is an $r \times s$ matrix, so $m = rs$ counts its variable entries. Finally, $k = \lfloor m/4 \rfloor$ is the dimension of the gradient slice to which Prop. 4.3 will be applied.

If $n \geq 6t$, then $b \geq 1$ and $s \geq n/2$. In particular, $d \leq t$ and $d \leq s$, so all hypotheses needed to construct the degree- d specialization in Lem. 6.1 hold. The next proposition isolates the remaining numerical condition before we choose the degree as a function of n .

Proposition 7.1. *Let n, d be positive integers, and define t, b, r, s, m, k by (7.1). Suppose*

$$(7.2) \quad d \geq 3, \quad n \geq 6t, \quad 4(2^t - 1) \leq ts.$$

Then every arithmetic circuit of size L computing per_n satisfies

$$(7.3) \quad L \geq \frac{n^2 \log_2(d-1)}{144}.$$

Proof. Let $\rho = 2^t - 1$, and let

$$P(X) = \sum_{h=1}^b \lambda_h M_{t,s,d}(X_{R_h,[s]})$$

be the weighted block polynomial from (6.1). Lem. 6.1 makes P a nonzero scalar multiple of an affine specialization of per_n , so it has precisely the form required by Prop. 4.3. Since all λ_h are nonzero and the row blocks involve disjoint variables, Cor. 5.2 gives

$$\dim \text{Crit}(P) \leq D := b(\rho + s(d-2)).$$

We estimate the two summands of D separately. The scale condition $4\rho \leq ts$ gives

$$b\rho \leq \frac{bts}{4} = \frac{m}{4}.$$

For the column-parameter contribution, the choice $t = 4d$ gives

$$bs(d-2) < bsd = \frac{bts}{4} = \frac{m}{4}.$$

Consequently,

$$\dim \text{Crit}(P) \leq D < \frac{m}{2}.$$

Since $k = \lfloor m/4 \rfloor \leq m/4$, we have $k + D < m$, and in particular $k + D \leq m$. Thus a k -dimensional linear slice can avoid the nonzero critical locus, as required in Lem. 4.2. Applying Prop. 4.3 to this specialization gives

$$(7.4) \quad k \log_2(d-1) \leq 3L.$$

It remains to express the slice dimension k in terms of the original matrix size n . From the definition of b ,

$$\frac{n}{2} - t \leq r = t \left\lfloor \frac{n}{2t} \right\rfloor \leq \frac{n}{2}.$$

The hypothesis $n \geq 6t$ implies $n/2 - t \geq n/3$. Therefore

$$\frac{n}{3} \leq r \leq \frac{n}{2}, \quad s = n - r + d \geq \frac{n}{2}, \quad m = rs \geq \frac{n^2}{6}.$$

In particular $m \geq 8$, and the floor in the definition of k loses at most a factor of two:

$$k = \left\lfloor \frac{m}{4} \right\rfloor \geq \frac{m}{8} \geq \frac{n^2}{48}.$$

Equivalently, $n^2 \leq 48k$. Substituting this estimate into (7.4) yields

$$L \geq \frac{k \log_2(d-1)}{3} \geq \frac{n^2 \log_2(d-1)}{144},$$

as claimed. □

Proof of Thm. 1.1. The fixed-degree bound becomes strongest when d grows while 2^{4d} remains small enough to satisfy the scale condition. For $n \geq 2^{16}$, put

$$\ell = \log_2 n, \quad d = \left\lfloor \frac{\ell}{4} \right\rfloor, \quad t = 4d.$$

Then $\ell \geq 16$, so $d \geq 4$ and $t \leq \ell$. The inequality $n \geq 6t$ follows from $n \geq 2^{16}$ and $t \leq \log_2 n$. As in the proposition, $r \leq n/2$ and $s \geq n/2$. Moreover, $2^t \leq 2^\ell = n$, while $d \geq 4$ gives $ds \geq n$. Therefore

$$2^t - 1 < n \leq ds, \quad 4(2^t - 1) \leq 4ds = ts.$$

Thus all three hypotheses of (7.2) hold.

Finally, $\ell \geq 16$ ensures

$$d - 1 = \left\lfloor \frac{\ell}{4} \right\rfloor - 1 \geq \frac{\ell}{4} - 2 \geq \frac{\ell}{8}.$$

Taking logarithms gives

$$\log_2(d - 1) \geq \log_2 \ell - 3 = \log_2 \log_2 n - 3.$$

Substitution into (7.3) proves (1.1). The factor $\log_2 \log_2 n - 3$ grows without bound, so the resulting circuit lower bound is superquadratic. $\square$

8. COEFFICIENT TRANSCENDENCE DEGREE

We now turn from circuits to formulas. The first task is to define a complexity measure that records how much independent information a polynomial carries in the coefficients associated with a selected set of variables. We then show that, for any division-free formula, this information is controlled by the number of occurrences of those variables. The next section will exhibit a matching of $O(\log n)$ matrix entries for which the permanent has $\Omega(n^2)$ independent coefficients. The formula lower bound will follow by applying the present section to many entry-disjoint matchings and adding their occurrence requirements.

Let $\mathcal{X}$ be a finite set of variables, let $\emptyset \neq Y \subseteq \mathcal{X}$, and put $Z = \mathcal{X} \setminus Y$. We call Y the selected variables and Z the remaining variables. To expose the information associated with Y , regard a polynomial in all the variables as a polynomial in Y whose coefficients are themselves polynomials in Z . Thus every $f \in \mathbb{C}[\mathcal{X}] = \mathbb{C}[Z][Y]$ has a unique expansion

$$f = \sum_{\alpha \in \mathbb{N}^Y} f_\alpha Y^\alpha, \quad f_\alpha \in \mathbb{C}[Z], \quad Y^\alpha = \prod_{y \in Y} y^{\alpha(y)}.$$

Although the index set $\mathbb{N}^Y$ is infinite, only finitely many f_α are nonzero. The multi-index 0 denotes the all-zero exponent vector, so f_0 is obtained by setting every variable of Y to zero.

Counting distinct coefficient polynomials would overestimate the information they contain: several coefficients can satisfy polynomial relations. Instead, we count the maximum number that are algebraically independent. Recall that polynomials $c_1, \dots, c_r \in \mathbb{C}[Z]$ are algebraically independent over $\mathbb{C}$ if

$$Q(c_1, \dots, c_r) \neq 0 \quad \text{for every nonzero } Q \in \mathbb{C}[T_1, \dots, T_r].$$

For example, if $Y = \{y\}$ and $f = yz_1 + z_2$, the two coefficients z_1, z_2 are algebraically independent. In contrast, for $f = yz^2 + z$, the coefficients $c_0 = z$ and $c_1 = z^2$ satisfy $c_1 - c_0^2 = 0$, and therefore carry only one algebraically independent parameter.

Define its coefficient transcendence degree by

$$\text{td}_Y(f) = \text{trdeg}_{\mathbb{C}} \mathbb{C}(f_\alpha : \alpha \in \mathbb{N}^Y).$$

Here $\mathbb{C}(Z)$ denotes the rational-function field in the remaining variables, and $\mathbb{C}(f_\alpha : \alpha \in \mathbb{N}^Y)$ is the smallest subfield containing $\mathbb{C}$ and all the coefficient polynomials. In particular, the field includes the constant-in- Y coefficient f_0 . Its transcendence degree is precisely the largest number of algebraically independent coefficients in the expansion. Passing to the field generated by the coefficients does not assume that they are rational rather than polynomial: it simply lets us measure how many independent parameters generate all of them. In the examples above, the coefficient transcendence degrees are two and one, respectively.

For a formula Φ , let $t_Y(\Phi)$ denote the number of leaves labeled by variables in Y , counted with multiplicity. Thus two leaves carrying the same variable contribute two occurrences. Our goal is to bound $\text{td}_Y(f)$ in terms of $t_Y(\Phi)$, independently of the number or complexity of the subformulas involving only Z . The coefficient-transcendence method is standard [KS14, §3.2.1, Lem. 7]; the following proof supplies the explicit constant needed here.

Lemma 8.1. *If a division-free formula Φ computes $f \in \mathbb{C}[\mathcal{X}]$ and $t = t_Y(\Phi) \geq 1$, then*

$$\text{td}_Y(f) \leq 4t - 2 \leq 4t.$$

Proof. Mark a vertex of the formula exactly when its subformula has a descendant leaf labeled by a variable in Y . The marked vertices form the portion of the formula connecting the t selected leaves to the root. A marked gate can have either one or two marked children. The two cases play different roles: a gate with one marked child merely changes an affine wrapper around the selected-variable computation, while a gate with two marked children combines two such computations and can introduce new parameters.

First suppose a marked gate has exactly one marked child, whose output is u . The other child contains no selected-variable leaf and hence computes some polynomial $h \in \mathbb{C}[Z]$. According to the gate operation and the order of its inputs, the output is one of

$$u + h, \quad u - h, \quad h - u, \quad hu.$$

Each expression has the form $Au + B$ with $A, B \in \mathbb{C}[Z]$. Moreover, a whole path of such gates can be absorbed into a single affine map, since

$$A_2(A_1u + B_1) + B_2 = (A_2A_1)u + (A_2B_1 + B_2).$$

Thus an arbitrarily long path with only one marked child at each gate does not require a new independent parameter at every gate.

We formalize this observation by induction. For every marked subformula with $s \geq 1$ selected-variable leaves, we claim that its output admits a representation

$$(8.1) \quad Ag + B, \quad A, B \in \mathbb{C}[Z], \quad g \in \mathbb{C}(\Gamma)[Y], \quad |\Gamma| \leq 4s - 4,$$

for some finite $\Gamma \subseteq \mathbb{C}(Z)$. The outer coefficients A, B are not yet counted in Γ ; they will be charged when the corresponding marked computation meets another one, or at the root. At a marked leaf labeled by $y \in Y$, choose

$$g = y, \quad A = 1, \quad B = 0, \quad \Gamma = \emptyset.$$

This gives the claim for $s = 1$.

If a marked gate has only one marked child, compose its affine map with the outer affine map already supplied for that child. The resulting slope and intercept still belong to $\mathbb{C}[Z]$, the inner polynomial g is unchanged, and no new element is added to Γ .

It remains to consider a gate with two marked children. Suppose they have $s_1, s_2 \geq 1$ selected-variable leaves and, by induction, their respective outputs are

$$A_1g_1 + B_1 \quad \text{and} \quad A_2g_2 + B_2, \quad g_i \in \mathbb{C}(\Gamma_i)[Y], \quad |\Gamma_i| \leq 4s_i - 4.$$

Adjoin the four outer coefficients and set

$$\Gamma = \Gamma_1 \cup \Gamma_2 \cup \{A_1, B_1, A_2, B_2\} \subseteq \mathbb{C}(Z).$$

Both child outputs now belong to $\mathbb{C}(\Gamma)[Y]$. Applying the gate operation $+$, $-$, or $\times$ therefore produces another polynomial $g \in \mathbb{C}(\Gamma)[Y]$. Represent the resulting output with the identity outer map, namely $A = 1$ and $B = 0$. Since $s = s_1 + s_2$, the parameter count is

$$|\Gamma| \leq (4s_1 - 4) + (4s_2 - 4) + 4 = 4s - 4.$$

This completes the induction. Equivalently, the marked part of a binary formula has $t-1$ genuine branching points; each can be charged at most four parameters, regardless of the lengths of the intervening one-child paths.

Apply (8.1) to the root, which is marked because $t \geq 1$. Its output is $f = Ag + B$ with $|\Gamma| \leq 4t - 4$. Adjoin the final slope and intercept by taking

$$\Gamma' = \Gamma \cup \{A, B\} \subseteq \mathbb{C}(Z).$$

Then $|\Gamma'| \leq 4t-2$ and $f \in \mathbb{C}(\Gamma')[Y]$, so every coefficient f_α belongs to $\mathbb{C}(\Gamma')$. A field generated by r elements has transcendence degree at most r , whether or not those generators are algebraically independent. Therefore

$$\text{td}_Y(f) \leq \text{trdeg}_{\mathbb{C}} \mathbb{C}(\Gamma') \leq |\Gamma'| \leq 4t-2,$$

as claimed. $\square$

To apply this upper bound to the permanent, we also need a practical way to certify that a large family of its coefficients is algebraically independent. The characteristic-zero Jacobian criterion supplies exactly that test. Given coefficient polynomials $c_1, \dots, c_r \in \mathbb{C}[Z]$, form the matrix

$$J(c_1, \dots, c_r) = \left(\frac{\partial c_i}{\partial z} \right)_{i \in [r], z \in Z}.$$

Its rank is computed over the rational-function field $\mathbb{C}(Z)$. The Jacobian criterion states that

$$\text{trdeg}_{\mathbb{C}} \mathbb{C}(c_1, \dots, c_r) = \text{rank}_{\mathbb{C}(Z)} J(c_1, \dots, c_r)$$

in characteristic zero [BMS11, Thm. 6]. Consequently, to prove that r coefficients are algebraically independent, it is enough to identify r remaining variables and show that the corresponding square Jacobian minor is a nonzero polynomial. It is enough in turn to find one specialization of those and the other remaining variables where this minor evaluates to a nonzero complex number. The next section constructs precisely such a specialization for $r = \Omega(n^2)$ coefficients associated with one short matching.

9. INDEPENDENT COEFFICIENTS FROM A MATCHING

A matching is a set of matrix entries with no repeated row or column. Our goal in this section is to prove that, for a matching Y of only $O(\log n)$ entries, the expansion of per_n in the variables of Y contains $\Omega(n^2)$ algebraically independent coefficients. Combined with the coefficient bound from Lem. 8.1, this will force every division-free formula for the permanent to contain $\Omega(n^2)$ occurrences of the variables in each such matching; an analogous coefficient bound later treats formulas with division. The remaining step, packing many entry-disjoint matchings, is deferred to §10.

The independence proof has two conceptual ingredients. First, split the marked matching into two equal parts, one encoding external column choices and the other encoding external row choices. Second, specialize the unmarked entries so that the Jacobian of selected coefficients separates into independent row and column evaluation matrices. We begin with a small instance that displays this separation without the binary-index bookkeeping needed in general.

9.1. A five-by-five warm-up. Choose distinct $p_1, p_2, p_3 \in \mathbb{C}$ and, independently, distinct $q_1, q_2, q_3 \in \mathbb{C}$, and consider

$$B = \begin{pmatrix} y_e & 0 & p_1 & p_2 & p_3 \\ 0 & y_f & 1 & 1 & 1 \\ 1 & q_1 & w_{11} & w_{12} & w_{13} \\ 1 & q_2 & w_{21} & w_{22} & w_{23} \\ 1 & q_3 & w_{31} & w_{32} & w_{33} \end{pmatrix}.$$

The matching here consists of the two diagonal variables y_e, y_f . Call their rows and columns *internal*, and call the remaining three rows and columns *external*. Since the permanent is multilinear, its expansion in the marked variables is

$$\text{per}(B) = c_{ef} y_e y_f + c_f y_f + c_e y_e + c_{\emptyset}.$$

Thus $c_{ef}, c_f, c_e, c_{\emptyset}$ collect the permutation terms using both marked diagonal entries, only y_f , only y_e , or neither, respectively. They are polynomials in the external block entries w_{ab} .

Differentiating a coefficient with respect to w_{ab} keeps precisely the permutation terms that match external row a to external column b . Set

$$G_b = \sum_{j \neq b} p_j, \quad H_a = \sum_{i \neq a} q_i,$$

and, only after taking the derivative, evaluate all w_{ij} at 1. Counting the remaining matchings gives

$$\frac{\partial c_{ef}}{\partial w_{ab}} = 2, \quad \frac{\partial c_f}{\partial w_{ab}} = 2G_b, \quad \frac{\partial c_e}{\partial w_{ab}} = 2H_a, \quad \frac{\partial c_{\emptyset}}{\partial w_{ab}} = G_b H_a.$$

Each count can be seen directly:

- For c_{ef} , both internal rows and columns are already matched. After the derivative fixes (a, b) , the other two external rows can match the other two external columns in $2! = 2$ ways.
- For c_f , the f -row uses its marked diagonal entry. The e -row must use an external column $j \neq b$, with weight p_j , giving G_b . The internal e -column can receive either external row other than a , giving the additional factor 2.
- For c_e , the roles are reversed: the f -row can use either remaining external column, while the external row $i \neq a$ matched to the f -column contributes q_i . The result is $2H_a$.
- For $c_{\emptyset}$, neither internal diagonal is used. The external column used by the e -row contributes G_b , and the external row matched to the f -column contributes H_a . These choices are independent, and all remaining matches are forced.

Restricting to the four external variables $w_{11}, w_{12}, w_{21}, w_{22}$ yields the Jacobian

$$\text{diag}(2, 2, 2, 1) \begin{pmatrix} 1 & 1 & 1 & 1 \\ G_1 & G_2 & G_1 & G_2 \\ H_1 & H_1 & H_2 & H_2 \\ G_1 H_1 & G_2 H_1 & G_1 H_2 & G_2 H_2 \end{pmatrix}.$$

The second factor is the Kronecker product $V_H \otimes V_G$ of the two-point Vandermonde matrices $V_H = \begin{pmatrix} 1 & 1 \\ H_1 & H_2 \end{pmatrix}$ and $V_G = \begin{pmatrix} 1 & 1 \\ G_1 & G_2 \end{pmatrix}$. The row choice and column choice are independent, so these two evaluation factors separate. Since $G_2 - G_1 = p_1 - p_2$ and $H_2 - H_1 = q_1 - q_2$, their Kronecker product has determinant $(p_1 - p_2)^2 (q_1 - q_2)^2 \neq 0$. The diagonal prefactor is also invertible, so the four coefficients are algebraically independent. The characteristic-zero Jacobian criterion recalled in §8 justifies this final implication: a nonzero Jacobian minor at even one specialization certifies independence before specialization.

The example identifies the general plan. We will replace the single e -index and f -index by ℓ indices of each type. Subsets of these indices encode external column and row choices in binary, and the resulting two evaluation matrices become $m \times m$ Vandermonde-type matrices. Their Kronecker product then certifies the independence of m^2 coefficients.

9.2. The parameters and the chosen coefficient family. Fix $n \geq 32$ and introduce the formula parameters

$$(9.1) \quad \ell = \lceil \log_2 n \rceil, \quad k = 2\ell, \quad m = n - k.$$

These parameters satisfy $k \leq n/2$, $m \geq k + 1$, and $m \leq 2^\ell$. Indeed, $n > 4\ell$ when $\ell = 5$, and when $\ell \geq 6$ one has $n \geq 2^{\ell-1} + 1 \geq 4\ell + 1$. Here k is the number of marked entries and m is the number of unmarked rows and columns left after those entries are placed on the diagonal. The inequality $m \leq 2^\ell$ means that ℓ bits suffice to index each of the m external rows or columns; the inequality $m \geq k + 1$ ensures that there are enough external indices for every injection that appears below.

Lemma 9.1. *If Y is a matching of k entries of X , then $\text{td}_Y(\text{per}_n) \geq m^2$.*

Proof. We first identify m^2 coefficients indexed by a pair $(\alpha, \beta) \in \{0, \dots, m-1\}^2$. We then specialize the unmarked entries, count the derivatives of those coefficients with respect to the m^2 external-block entries, and show that the resulting Jacobian is invertible. The counting separates external column choices from external row choices, exactly as in the 5×5 example.

Step 1: Coefficients of marked diagonal entries. By permuting rows and columns, suppose $Y = \{y_1, \dots, y_k\}$, where $y_i = x_{ii}$. For $S \subseteq [k]$, let c_S denote the coefficient of $\prod_{i \in S} y_i$. Multilinearity

gives

$$(9.2) \quad c_S = \left(\prod_{i \in S} \frac{\partial}{\partial y_i} \right) \text{per}_n(X) \Big|_{y_1 = \dots = y_k = 0}.$$

Indeed, differentiating once with respect to each y_i for $i \in S$ selects the permutation terms that use those marked diagonal entries. Setting all marked variables to zero then discards terms using any additional marked entry. Equivalently, c_S is the permanent of the matrix obtained by deleting the rows and columns indexed by S and setting the remaining marked diagonal entries to zero.

Step 2: Separate internal and external choices. Partition the marked indices as $[k] = E \sqcup F$, where $E = \{e_0, \dots, e_{\ell-1}\}$ and $F = \{f_0, \dots, f_{\ell-1}\}$. Call these k rows and columns internal, index the m external rows and columns by $[m]$, and write $X = \begin{pmatrix} D & U \\ V & W \end{pmatrix}$, where D is $k \times k$, U is $k \times m$, V is $m \times k$, and $W = (w_{ab})$ is $m \times m$. Thus U matches internal rows to external columns, V matches external rows to internal columns, and W matches external rows to external columns.

Choose distinct $p_1, \dots, p_m \in \mathbb{C}$ and, independently, distinct $q_1, \dots, q_m$. For $i \neq j$ in $[k]$, $0 \leq u, v < \ell$, and $a, b \in [m]$, specialize the unmarked variables by

$$(9.3) \quad D_{ij} = 0, \quad w_{ab} = 1,$$

$$U_{e_u, b} = p_b^{2^u}, \quad U_{f_v, b} = 1,$$

$$(9.4) \quad V_{a, e_u} = 1, \quad V_{a, f_v} = q_a^{2^v}.$$

Denote this specialization by ξ ; derivatives are taken before evaluation at ξ . The e -indices carry powers of p_b when an internal row chooses external column b , whereas the f -indices carry powers of q_a when external row a chooses an internal column. All other internal–external edges have weight 1. The powers 2^u and 2^v ensure that subsets of E and F encode integers without colliding.

For $0 \leq \alpha, \beta < m$, let $P(\alpha) \subseteq E$ and $Q(\beta) \subseteq F$ record the 1-digits of the ℓ -digit binary expansions of α and β , with digit 0 least significant. Set $T(\alpha, \beta) = P(\alpha) \cup Q(\beta)$ and $S(\alpha, \beta) = [k] \setminus T(\alpha, \beta)$. Since $m \leq 2^\ell$, these definitions yield m^2 coefficients $c_{S(\alpha, \beta)}$ of distinct marked monomials. The complement in the definition of S is important: the indices in S have already been matched through their marked diagonal entries, while the indices in T remain present in the coefficient permanent and must be matched through the unmarked blocks.

Our immediate goal is to show that these m^2 coefficients have an invertible Jacobian with respect to the m^2 variables w_{ab} . Consider the square matrix

$$J_{(\alpha, \beta), (a, b)} = \frac{\partial c_{S(\alpha, \beta)}}{\partial w_{ab}} \Big|_{\xi}.$$

Once $\det J \neq 0$ has been established, the Jacobian criterion immediately gives the conclusion of the lemma. We first derive a combinatorial expression for each entry of J .

Step 3: Count the surviving permutation terms. Fix α, β, a, b , and abbreviate $P = P(\alpha)$, $Q = Q(\beta)$, $T = P \cup Q$, and $r = |T|$. Because $r \leq k \leq m - 1$, differentiation fixes external row a to external column b and leaves precisely the internal rows and columns indexed by T . There are $m - 1$ external rows and $m - 1$ external columns still available. At ξ , the internal block on T is zero: its off-diagonal entries are zero by (9.3), and its marked diagonal entries were set to zero in (9.2). Therefore no remaining internal row can match an internal column. Instead, every internal row must match an external column, and every internal column must receive an external row.

The first set of choices is an injection $\rho : T \hookrightarrow [m] \setminus \{b\}$, assigning a distinct external column to each internal row. The second is an injection $\theta : T \hookrightarrow [m] \setminus \{a\}$, assigning a distinct external row to each internal column. They are independent because ρ selects columns and θ selects rows. The $m - 1 - r$ external rows and columns not used by these injections can then be matched in $(m - 1 - r)!$ ways, each of weight 1, since W specializes to the all-ones matrix. Hence

$$(9.5) \quad J_{(\alpha, \beta), (a, b)} = (m - 1 - r)! \sum_{\rho : T \hookrightarrow [m] \setminus \{b\}} \prod_{i \in T} U_{i, \rho(i)} \sum_{\theta : T \hookrightarrow [m] \setminus \{a\}} \prod_{j \in T} V_{\theta(j), j}.$$

This is the general counterpart of the four elementary matching counts in the warm-up. To see the row–column separation in a form that proves invertibility, we now express each injection sum as the evaluation of a univariate polynomial.

Step 4: Encode forbidden-column choices by univariate polynomials. For $P \subseteq E$, let $D_P = \sum_{e_u \in P} 2^u$ and $H_P = \sum_{\rho: P \hookrightarrow [m]} \prod_{e_u \in P} p_{\rho(e_u)}^{2^u}$. Here H_P is the total weight of injections into all m external columns; in particular, $H_\emptyset = 1$. Define $g_P \in \mathbb{C}[z]$ recursively by

$$(9.6) \quad g_\emptyset(z) = 1, \quad g_P(z) = H_P - \sum_{e_u \in P} z^{2^u} g_{P \setminus \{e_u\}}(z) \quad \text{for } P \neq \emptyset.$$

To interpret this recursion, fix b . An injection from P into $[m]$ either avoids b or maps exactly one index $e_u \in P$ to b . In the latter case, the remaining indices form an injection from $P \setminus \{e_u\}$ that avoids b , and the chosen index contributes $p_b^{2^u}$. Consequently, subtracting all injections that use b from the unrestricted sum H_P gives precisely the recursion (9.6) evaluated at $z = p_b$. Induction on $|P|$ therefore yields

$$(9.7) \quad g_P(p_b) = \sum_{\rho: P \hookrightarrow [m] \setminus \{b\}} \prod_{e_u \in P} p_{\rho(e_u)}^{2^u}, \quad \deg g_P = D_P, \quad [z^{D_P}]g_P(z) = (-1)^{|P|}|P|!.$$

For completeness, the degree and leading coefficient follow from the same induction. For $P \neq \emptyset$, H_P is constant, whereas each term $z^{2^u} g_{P \setminus \{e_u\}}(z)$ in (9.6) has degree D_P and leading coefficient $(-1)^{|P|-1}(|P|-1)!$. The preceding minus sign and the $|P|$ choices of e_u give the claimed coefficient $(-1)^{|P|}|P|!$. In particular, $g_{\{e_0\}}(z) = \sum_b p_b - z$, recovering the first factor in the example. These leading coefficients are nonzero because the ground field has characteristic zero.

The row-side construction is identical. For $Q \subseteq F$, put

$$D_Q = \sum_{f_v \in Q} 2^v, \quad K_Q = \sum_{\theta: Q \hookrightarrow [m]} \prod_{f_v \in Q} q_{\theta(f_v)}^{2^v},$$

and recursively define

$$h_\emptyset(z) = 1, \quad h_Q(z) = K_Q - \sum_{f_v \in Q} z^{2^v} h_{Q \setminus \{f_v\}}(z) \quad \text{for } Q \neq \emptyset.$$

Partitioning row injections according to whether they use the forbidden row a proves

$$(9.8) \quad h_Q(q_a) = \sum_{\theta: Q \hookrightarrow [m] \setminus \{a\}} \prod_{f_v \in Q} q_{\theta(f_v)}^{2^v}, \quad \deg h_Q = D_Q, \quad [z^{D_Q}]h_Q(z) = (-1)^{|Q|}|Q|!.$$

At this point the nontrivial weighted choices have been isolated: $g_P(p_b)$ records the P -indices that choose external columns, and $h_Q(q_a)$ records the Q -indices that choose external rows. It remains to account for indices whose edge weights are 1 and then prove that the resulting evaluation matrices have full rank.

Step 5: Factor the Jacobian entries. Write $(M)_s = M(M-1)\cdots(M-s+1)$ and $(M)_0 = 1$. In the first sum of (9.5), choose the images of P first. Their total weight is $g_P(p_b)$ by (9.7). Because every index in Q has U -weight one, each such injection has exactly $(m-1-|P|)_{|Q|}$ extensions to the labeled indices in Q : there are $m-1-|P|$ available external columns for the first index, one fewer for the next, and so on. Similarly, in the second sum first choose the images of Q . Their total weight is $h_Q(q_a)$, and each choice has $(m-1-|Q|)_{|P|}$ extensions to the indices in P , all with V -weight one.

Substituting these two counts into (9.5) separates the column parameter p_b from the row parameter q_a :

$$(9.9) \quad J_{(\alpha, \beta), (a, b)} = L_{P, Q} g_P(p_b) h_Q(q_a),$$

$$(9.10) \quad L_{P, Q} = (m-1-|P|-|Q|)!(m-1-|P|)_{|Q|}(m-1-|Q|)_{|P|} \neq 0.$$

All factors are nonzero because $|P| + |Q| \leq m-1$; in particular, every falling factorial counts injections that actually exist. The scalar $L_{P, Q}$ depends on the selected coefficient, but not on the external row a or column b .

Step 6: Apply two Vandermonde matrices. Since $D_{P(\alpha)} = \alpha$ and $D_{Q(\beta)} = \beta$, the evaluation matrices $A_p = (g_{P(\alpha)}(p_b))_{b,\alpha}$ and $A_q = (h_{Q(\beta)}(q_a))_{a,\beta}$ are invertible. Indeed, (9.7) shows that the polynomials

$$g_{P(0)}(z), g_{P(1)}(z), \dots, g_{P(m-1)}(z)$$

have respective degrees $0, 1, \dots, m-1$ and nonzero leading coefficients. Their coefficient matrix in the monomial basis $1, z, \dots, z^{m-1}$ is therefore triangular with nonzero diagonal. Evaluating at the distinct points $p_1, \dots, p_m$ multiplies this coefficient matrix by the Vandermonde matrix $(p_b^j)_{b \in [m], 0 \leq j < m}$, which is invertible because the p_b are distinct. Thus A_p is invertible. The same argument, using (9.8) and the distinct q_a , proves that A_q is invertible.

After permuting columns, the matrix $(g_{P(\alpha)}(p_b)h_{Q(\beta)}(q_a))_{(\alpha,\beta),(a,b)}$ is $A_p^\top \otimes A_q^\top$ and is therefore invertible. By (9.9), J differs from it only by the nonzero row scalars (9.10). Hence $\det J \neq 0$. The m^2 selected coefficients are consequently algebraically independent by the Jacobian criterion, giving $\text{td}_Y(\text{per}_n) \geq m^2$. $\square$

10. FORMULA LOWER BOUNDS WITH AND WITHOUT DIVISION

We have established the two local ingredients needed for a formula lower bound. For a matching Y of $k = O(\log n)$ matrix entries, Lem. 9.1 produces $m^2 = \Omega(n^2)$ algebraically independent Y -coefficients of per_n . If the formula has no division, Lem. 8.1 charges those coefficients to occurrences of the variables in Y . Our first task is to select many matchings for which the resulting occurrence charges do not overlap. We then extend the charging argument to formulas with division by showing that rational operations along a path with only one marked child at each gate contribute only a constant number of field parameters.

10.1. Packing matchings and the division-free bound. Throughout this section, $k = 2\lceil \log_2 n \rceil$ and $m = n - k$ are the parameters from (9.1). The useful notion of disjointness concerns matrix *entries*, not their rows and columns: two matchings may use the same row or the same column, as long as they never contain the same variable x_{ij} . This weaker condition is exactly what prevents a variable-labeled leaf from being charged to two different matchings.

Proof of Thm. 1.2. Index the rows and columns by $\{0, \dots, n-1\}$. For a cyclic offset $0 \leq \tau < n$ and a row-block index $0 \leq j < \lfloor n/k \rfloor$, define

$$Y_{\tau,j} = \{x_{jk+r, (jk+r+\tau) \bmod n} : 0 \leq r < k\}.$$

Each $Y_{\tau,j}$ is a matching: its k row indices are distinct, and adding the fixed offset τ modulo n sends them to k distinct column indices. Moreover, any entry in one of these matchings determines its parameters uniquely. Its row index i determines $j = \lfloor i/k \rfloor$, and its column-minus-row difference modulo n determines τ . Thus the matchings are pairwise entry-disjoint even though different matchings can share rows and columns.

There are n choices of τ and $\lfloor n/k \rfloor$ choices of j , so their total number is

$$(10.1) \quad \nu = n \left\lfloor \frac{n}{k} \right\rfloor \geq \frac{n^2}{2k},$$

where the inequality follows from $k \leq n/2$ and $\lfloor x \rfloor \geq x/2$ for $x \geq 2$.

Every matrix variable must occur at least once in Φ : otherwise its output would be independent of that variable, whereas $\partial \text{per}_n / \partial x_{ij}$ is the nonzero permanent of the complementary $(n-1) \times (n-1)$ submatrix. In particular, $t_Y(\Phi) \geq 1$ for every selected matching Y , so Lem. 8.1 applies. Together with Lem. 9.1, it gives the per-matching requirement

$$m^2 \leq \text{td}_Y(\text{per}_n) \leq 4t_Y(\Phi).$$

In words, a matching containing only $k = O(\log n)$ distinct variables must account for at least $m^2/4 = \Omega(n^2)$ variable-labeled leaves, counted with repetition.

Because the matchings are entry-disjoint, a leaf labeled by x_{ij} contributes to at most one of the numbers $t_Y(\Phi)$. Consequently, $\sum_Y t_Y(\Phi) \leq L_{\text{var}}(\Phi)$. Summing the preceding local

requirement, and then using (10.1), $m \geq n/2$, and $k \leq 4 \log_2 n$, gives

$$\begin{aligned} L_{\text{var}}(\Phi) &\geq \sum_Y t_Y(\Phi) \geq \frac{\nu m^2}{4} \geq \frac{n^2 m^2}{8k} \\ &\geq \frac{n^4}{32k} \geq \frac{n^4}{128 \log_2 n}. \end{aligned}$$

The total leaf and vertex counts are at least the number $L_{\text{var}}(\Phi)$ of variable-labeled leaves. Finally, a rooted binary formula with $L(\Phi)$ leaves has exactly $G(\Phi) = L(\Phi) - 1$ internal gates. Since the permanent is nonconstant and $n \geq 32$, $L(\Phi) \geq 2$, and hence $G(\Phi) \geq L(\Phi)/2 \geq L_{\text{var}}(\Phi)/2$. This gives the stated gate bound as well. $\square$

10.2. Allowing division: recovering a coefficient field. To extend the argument, fix a nonempty set Y of marked variables, let Z contain all remaining variables, and put $R = \mathbb{C}(Z)$. A formula with division is evaluated symbolically in the rational-function field

$$\mathbb{C}(Y, Z) = R(Y).$$

Validity means that the denominator at each division gate is a nonzero element of this field; it does *not* require the denominator to be nonzero after every numerical specialization of Z . In particular, the point used in (9.3)–(9.4) to prove algebraic independence might make some denominator of the formula vanish. This causes no problem: that point is applied only to the polynomial coefficients of per_n , not to the rational formula or any of its gates.

The new difficulty is that a division formula may represent its output using rational expressions in the marked variables. Merely placing the output in a small rational-function field $K(Y)$ does not immediately say that its polynomial coefficients belong to K . The next elementary observation supplies precisely that missing implication.

Lemma 10.1. *For fields $K \subseteq E$ and a finite set Y of indeterminates, $K(Y) \cap E[Y] = K[Y]$ inside $E(Y)$.*

Proof. The inclusion $K[Y] \subseteq K(Y) \cap E[Y]$ is immediate. Conversely, suppose $f \in K(Y) \cap E[Y]$. Membership in $K(Y)$ gives polynomials $P, Q \in K[Y]$ with $Q \neq 0$ and

$$Qf = P \quad \text{in } E[Y].$$

To compare the coefficients in E with those in K , regard E as a K -vector space. Extend a basis of its one-dimensional subspace $K \subseteq E$ to a basis of E , and let $\pi : E \rightarrow K$ be the associated K -linear projection. Thus π fixes every element of K . Apply π coefficientwise to polynomials in $E[Y]$, writing π_* for the resulting map.

Although π need not preserve products of arbitrary elements of E , its K -linearity ensures that it commutes with multiplication by the polynomial $Q \in K[Y]$. Therefore

$$Q\pi_*(f) = \pi_*(Qf) = \pi_*(P) = P.$$

Subtracting from $Qf = P$ gives $Q(f - \pi_*(f)) = 0$. Since $E[Y]$ is an integral domain and $Q \neq 0$, we conclude that $f = \pi_*(f) \in K[Y]$, as required. $\square$

It remains to find a small field K containing the rational parameters that genuinely influence the marked variables. As in the division-free argument, mark each vertex whose subformula contains a Y -labeled leaf. A maximal path of gates with only one marked child can contain many unmarked computations, but all of its influence is captured by one fractional linear transformation. The essential count is therefore charged at gates where two marked subformulas meet.

Lemma 10.2. *Suppose that a valid formula with division computes a polynomial $f \in \mathbb{C}[Y, Z]$ and contains $t \geq 1$ leaves labeled by variables in Y . Then $\text{td}_Y(f) \leq 6t - 3 < 6t$.*

Proof. Put $R = \mathbb{C}(Z)$, and mark precisely those vertices having at least one Y -labeled descendant. At a gate with one marked child, the other child has no marked leaves and therefore

computes a rational function $h \in R$. If the marked child computes $u \in R(Y)$, the possible outputs are

$$u + h, \quad u - h, \quad h - u, \quad hu, \quad u/h, \quad h/u,$$

where validity requires $h \neq 0$ for u/h and $u \neq 0$ for h/u . Every such operation has the fractional linear form

$$u \mapsto \frac{au + b}{cu + d}, \quad M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in R^{2 \times 2}.$$

For example, the transformations $u + h$, hu , u/h , and h/u are represented, respectively, by

$$\begin{pmatrix} 1 & h \\ 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} h & 0 \\ 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 \\ 0 & h \end{pmatrix}, \quad \begin{pmatrix} 0 & h \\ 1 & 0 \end{pmatrix}.$$

It is important not to require these matrices to be invertible. When $h = 0$, multiplication by h is a valid constant-zero operation represented by the singular, but nonzero, second matrix.

Composition of fractional linear transformations corresponds to multiplication of their representing matrices. Moreover, although singular matrices are allowed, the composite matrix for a valid marked path cannot become the zero matrix. To see this, suppose that the current marked value is

$$u_j = \frac{a_j u_0 + b_j}{c_j u_0 + d_j}, \quad c_j u_0 + d_j \neq 0,$$

and the next unary marked gate applies $u \mapsto (au + b)/(cu + d)$. The denominator represented by the composite matrix is

$$c(a_j u_0 + b_j) + d(c_j u_0 + d_j) = (c_j u_0 + d_j)(cu_j + d) \neq 0 \quad \text{in } R(Y).$$

The first factor is nonzero by the inductive representation, and the second is nonzero because the gate is valid. In particular, the composite matrix has a nonzero entry. Dividing all four entries by any one nonzero entry leaves the represented rational function unchanged and normalizes that entry to 1. Such a normalized matrix is therefore described by at most three elements of R , even when it is singular.

We now prove the following more precise invariant by induction on a marked subformula with $s \geq 1$ marked leaves: its output can be written as

$$\frac{ag + b}{cg + d}, \quad g \in \mathbb{C}(\Gamma)(Y), \quad \Gamma \subseteq R, \quad |\Gamma| \leq 6s - 6,$$

for some matrix over R , with $cg + d \neq 0$ in $R(Y)$. The outer matrix is deliberately not counted in Γ until the subformula meets another marked subformula or reaches the root.

For a marked leaf labeled by $y \in Y$, take $g = y$, $\Gamma = \emptyset$, and the identity matrix. If the root gate of the subformula has only one marked child, compose its fractional linear transformation with the child's outer matrix. The preceding validity argument shows that the resulting denominator remains nonzero, and no new element needs to be added to Γ .

Suppose instead that the gate has two marked children with $s_1, s_2 \geq 1$ marked leaves. Normalize each child's outer matrix and adjoin its at most three nontrivial entries to the union of the two existing parameter sets. If Γ denotes the resulting set, both child outputs belong to the single field $\mathbb{C}(\Gamma)(Y)$. Apply the gate operation to these two outputs, use the resulting rational function as the new inner function g , and take the identity as the new outer matrix. When the operation is division, the divisor is nonzero in $R(Y)$ by validity, hence also nonzero in the subfield $\mathbb{C}(\Gamma)(Y)$; thus the new inner function is well-defined. Finally, because $s = s_1 + s_2$, the parameter count is at most

$$(6s_1 - 6) + (6s_2 - 6) + 6 = 6(s_1 + s_2) - 6.$$

This completes the induction.

At the formula root, the invariant gives at most $6t - 6$ parameters for the inner function. Normalize the final outer matrix and adjoin its at most three remaining entries. We obtain a set $\Gamma' \subseteq R$ satisfying

$$|\Gamma'| \leq (6t - 6) + 3 = 6t - 3$$

such that the complete formula output belongs to $K(Y)$, where $K = \mathbb{C}(\Gamma') \subseteq R$. On the other hand, the hypothesis that the output is a polynomial gives $f \in \mathbb{C}[Y, Z] \subseteq R[Y]$. Applying Lem. 10.1 with $E = R$ therefore yields

$$f \in K(Y) \cap R[Y] = K[Y].$$

Thus every coefficient of f as a polynomial in Y belongs to the same field K . Since a field generated by $|\Gamma'|$ elements has transcendence degree at most $|\Gamma'|$,

$$\text{td}_Y(f) \leq \text{trdeg}_{\mathbb{C}} K \leq |\Gamma'| \leq 6t - 3 < 6t.$$

This is the desired division-formula analogue of Lem. 8.1. $\square$

10.3. Completing the lower bound with division. The coefficient-independence construction is a property of the permanent itself and therefore remains unchanged when the formula is allowed to use division. Only its translation into a lower bound on marked occurrences changes, from a factor of 4 to a factor of 6.

Proof of Thm. 1.3. Use the same parameters (9.1) and the same ν pairwise entry-disjoint matchings constructed in (10.1). If some matching Y contributes no variable-labeled leaf, every gate of the formula belongs to the rational-function field $\mathbb{C}(X \setminus Y)$. Its output would then be independent of all variables in Y , contradicting its equality to per_n . Therefore $t_Y(\Phi) \geq 1$ for every selected matching, as required to apply Lem. 10.2.

For each matching, Lem. 9.1 and Lem. 10.2 now give

$$m^2 \leq \text{td}_Y(\text{per}_n) \leq 6t_Y(\Phi) - 3 < 6t_Y(\Phi).$$

As before, entry-disjointness guarantees that no variable-labeled leaf is charged twice. Summing over all ν matchings, and using $m \geq n/2$ and $k \leq 4 \log_2 n$, yields

$$\begin{aligned} L_{\text{var}}(\Phi) &\geq \sum_Y t_Y(\Phi) \geq \frac{\nu m^2}{6} \geq \frac{n^2 m^2}{12k} \\ &\geq \frac{n^4}{48k} \geq \frac{n^4}{192 \log_2 n}. \end{aligned}$$

Finally, allowing division changes the allowed labels of internal gates but does not change the fact that the formula is a rooted binary tree. Consequently, $L(\Phi) \geq L_{\text{var}}(\Phi)$ and $G(\Phi) = L(\Phi) - 1 \geq L(\Phi)/2$, giving the remaining leaf, vertex, and gate bounds exactly as in the division-free case. $\square$

11. COMPARISON WITH THE DETERMINANT

For the generic determinant, the known unrestricted formula bounds are $\Omega(n^3)$ and $n^{O(\log n)}$. The lower bound holds even with division [Kal85], while the upper bound follows by unfolding the polynomial-size, $O(\log^2 n)$ -depth circuits of [Ber84]. It is therefore natural to ask whether either of our permanent arguments also improves the corresponding determinant bound. We show that the answer is negative for two separate reasons. On the formula side, the coefficients associated with a small set of determinant entries satisfy many algebraic relations. On the circuit side, every homogeneous specialization of the determinant has a critical locus that is too large for our geometric certificate to give a superquadratic bound.

11.1. Coefficient independence and a cubic formula ceiling. The formula argument requires many algebraically independent coefficients for small, entry-disjoint sets of variables. Recall that $\text{td}_Y(f)$ counts the maximum number of algebraically independent coefficients obtained by expanding f in the variables of Y ; those coefficients are polynomials in the remaining variables. For the permanent, a matching of only $O(\log n)$ entries contributes $\Omega(n^2)$ independent coefficients. We first explain why this crucial feature fails for the determinant.

Let Y be a matching of $k < n$ entries of a generic $n \times n$ matrix X . Permuting rows and columns moves its marked entries to the first k diagonal positions. Separate the corresponding rows and columns from the others, and write

$$X = \begin{pmatrix} \text{diag}(Y) + D_0 & U \\ V & W \end{pmatrix}, \quad \delta = \det W, \quad S = D_0 - UW^{-1}V.$$

Here $\text{diag}(Y)$ is the $k \times k$ diagonal matrix of marked variables, D_0 contains the unmarked entries of the same block and has zero diagonal, and U, V, W also contain only unmarked variables. The complementary block W is a generic $(n - k) \times (n - k)$ matrix, so its determinant δ is a nonzero polynomial. Consequently, W is invertible over the rational-function field of the unmarked variables: its entries need not be invertible as polynomials, but $W^{-1} = \text{adj}(W)/\delta$ is a well-defined matrix of rational functions. The Schur complement identity then gives

$$\det X = \delta \det(\text{diag}(Y) + S).$$

Although the right-hand side uses rational functions to describe the coefficients, the identity holds in the rational-function field, and its left-hand side is still the original polynomial.

To read off an individual coefficient, let $T \subseteq [k]$. In the determinant expansion, selecting y_i for every $i \in T$ fixes those rows and columns to their diagonal positions. The remaining rows and columns therefore contribute the principal minor on $[k] \setminus T$:

$$\left[\prod_{i \in T} y_i \right] \det X = \delta \det S_{[k] \setminus T, [k] \setminus T},$$

where the determinant of an empty matrix is 1. Thus every Y -coefficient belongs to the field generated by the single element δ and the k^2 entries of S . A field generated by $k^2 + 1$ elements has transcendence degree at most $k^2 + 1$, and hence

$$\text{td}_Y(\det X) \leq k^2 + 1.$$

Consequently, matchings of size $O(\log n)$ yield at most $O(\log^2 n)$ independent coefficients for the determinant, whereas §9 produces $\Omega(n^2)$ independent coefficients for the permanent. The determinant's Schur complement compresses a potentially exponential family of coefficients into only $O(k^2)$ parameters.

One might try to recover a stronger formula lower bound by choosing marked entries that are not a matching. The same obstruction still applies. Suppose Y is any nonempty set of k marked entries, where $1 \leq k < n$. Its occupied row set and occupied column set each have size at most k ; enlarge the smaller set, if necessary, so both have a common size $s \leq k$. After permuting rows and columns, all marked entries lie in the upper-left $s \times s$ block, which we write as $B(Y) + D_0$. Here $B(Y)$ contains the marked variables and vanishes in the other positions, while D_0 contains the unmarked entries. Because $s < n$, the complementary block W is again generically invertible. The Schur complement expresses every Y -coefficient in the field generated by $\delta = \det W$ and the s^2 entries of $D_0 - UW^{-1}V$. Therefore $\text{td}_Y(\det X) \leq s^2 + 1 \leq nk$. For the last inequality, $s \leq k < n$ implies $s^2 + 1 \leq k^2 + 1 \leq nk$. If $k \geq n$, we instead use the trivial bound $\text{td}_Y(\det X) \leq n^2 \leq nk$: there are only n^2 matrix variables in total.

We have therefore proved $\text{td}_Y(\det X) \leq n|Y|$ for every nonempty block Y . If $Y_1, \dots, Y_q$ are pairwise entry-disjoint, then their total size is at most n^2 , and summing the individual bounds gives

$$(11.1) \quad \sum_{j=1}^q \text{td}_{Y_j}(\det X) \leq n \sum_{j=1}^q |Y_j| \leq n^3.$$

Thus summing blockwise coefficient transcendence degrees cannot give a supercubic determinant formula lower bound, regardless of how the entry-disjoint blocks are chosen.

11.2. Full matchings recover the cubic determinant bound. The cubic ceiling is not merely an upper limit on what this method might prove: using full matchings, the same method actually recovers a cubic lower bound. Let Y consist of all n diagonal entries, and write

$$X = \text{diag}(Y) + A, \quad a_{ii} = 0.$$

Expanding as above, the Y -coefficients are exactly the principal minors of the zero-diagonal matrix A ; a principal minor indexed by $I \subseteq [n]$ is $\det A_{I,I}$.

The first question is how many of these principal minors can be algebraically independent. If D is an invertible diagonal matrix, then

$$(DAD^{-1})_{I,I} = D_{I,I}A_{I,I}D_{I,I}^{-1},$$

so diagonal conjugation leaves every principal minor unchanged. On the generic open set where $a_{1i} \neq 0$ for $i > 1$, choose $\lambda_i = a_{1i}$ and $D = \text{diag}(1, \lambda_2, \dots, \lambda_n)$. The matrix $B = DAD^{-1}$ satisfies $b_{1i} = 1$ for every $i > 1$. Conversely, $A = D^{-1}BD$, so the original off-diagonal entries are birationally equivalent to the $n - 1$ nonzero parameters λ_i together with the remaining normalized entries of B . The latter entries are therefore algebraically independent, and their number is

$$n(n - 1) - (n - 1) = (n - 1)^2$$

because the normalization fixes precisely the $n - 1$ first-row entries. In particular, all principal minors belong to the rational-function field generated by these $(n - 1)^2$ normalized entries. That field has transcendence degree $(n - 1)^2$, giving the corresponding upper bound.

For the matching lower bound, it is enough to recover all normalized entries up to algebraic ambiguity from the principal minors. Continue to write A for the normalized matrix, and set $t_i = a_{i1}$ for $i > 1$. For distinct $i, j > 1$, its principal minors of sizes two and three are

$$\begin{aligned} \det A_{\{1,i\},\{1,i\}} &= -t_i, \\ \det A_{\{i,j\},\{i,j\}} &= -a_{ij}a_{ji}, \\ \det A_{\{1,i,j\},\{1,i,j\}} &= a_{ij}t_j + a_{ji}t_i. \end{aligned}$$

Thus the field generated by the principal minors contains t_i , the product $p_{ij} = a_{ij}a_{ji}$, and the linear combination $q_{ij} = a_{ij}t_j + a_{ji}t_i$. Multiplying the last identity by a_{ij} gives

$$t_j a_{ij}^2 - q_{ij} a_{ij} + p_{ij} t_i = 0.$$

Its leading coefficient t_j is a nonzero rational function, so a_{ij} is algebraic of degree at most two over the field of principal minors; the same argument applies to a_{ji} . Every normalized entry is consequently algebraic over that field. Algebraic extensions preserve transcendence degree, so the upper bound established by normalization is attained:

$$\text{td}_Y(\det X) = (n - 1)^2.$$

Now let Φ be a formula computing $\det X$, and partition all n^2 entries into the n cyclic full matchings

$$Y_\tau = \{x_{i,(i+\tau) \bmod n} : 0 \leq i < n\}, \quad 0 \leq \tau < n,$$

where rows and columns are indexed by $\{0, \dots, n - 1\}$. Each matching has coefficient transcendence degree $(n - 1)^2$ by the preceding argument, since row and column permutations only relabel the entries. If the formula is division-free, Lem. 8.1 gives $(n - 1)^2 \leq 4t_{Y_\tau}(\Phi) - 2$. If it permits valid division, Lem. 10.2 instead gives $(n - 1)^2 \leq 6t_{Y_\tau}(\Phi) - 3$. Summing over the n disjoint matchings therefore yields

$$L_{\text{var}}(\Phi) \geq \frac{n((n - 1)^2 + 2)}{4}, \quad L_{\text{var}}(\Phi) \geq \frac{n((n - 1)^2 + 3)}{6}$$

without division and with valid division. Thus the cubic ceiling in (11.1) is attained up to a constant factor.

11.3. Critical loci obstruct the geometric circuit argument. The circuit proof needs an affine specialization with many active variables, reasonably high degree, and a critical locus of large codimension. We first inspect the determinant analogue of the specialization used for the permanent, and then show that the obstruction persists for every homogeneous determinant specialization.

Replacing the permanent by the determinant in the example from §2 gives

$$\det \begin{pmatrix} u & v & 1 & 1 \\ w & z & 1 & 1 \\ p & q & 2 & -2 \\ r & s & 2 & -2 \end{pmatrix} = 4((u-w)(q-s) - (v-z)(p-r)),$$

a single 2×2 determinant rather than a sum of independent block permanents. The cancellations have compressed all the variable dependence into a much smaller matrix.

More generally, suppose $2 \leq d \leq \min\{r, s\}$, and let X, U, V have respective sizes $r \times s$, $r \times (r-d)$, and $(s-d) \times s$, with U, V constant. The block matrix

$$\begin{pmatrix} X & U \\ V & 0 \end{pmatrix}$$

is square of size $r + s - d$. If U fails to have full column rank, its last $r-d$ columns are linearly dependent. If V fails to have full row rank, its last $s-d$ rows are linearly dependent. In either case the determinant is identically zero.

Suppose therefore that U and V have their full respective ranks. Choose a full-row-rank $d \times r$ matrix A whose rows span the left kernel of U , so $AU = 0$. Likewise, choose a full-column-rank $s \times d$ matrix C whose columns span the kernel of V , so $VC = 0$. Constant changes of basis separating these kernels from complementary subspaces give

$$\det \begin{pmatrix} X & U \\ V & 0 \end{pmatrix} = c \det(AXC),$$

where $c \neq 0$ depends only on the constant basis changes. In other words, the original determinant retains only the induced map from the d -dimensional kernel of V to the d -dimensional quotient by the image of U .

The linear map $X \mapsto AXC$ is surjective onto the space of $d \times d$ matrices. Indeed, full row rank provides a right inverse A' with $AA' = I_d$, while full column rank provides a left inverse C' with $C'C = I_d$; for any $d \times d$ matrix H , taking $X = A'HC'$ gives $AXC = H$. The first derivatives of a determinant are its $(d-1) \times (d-1)$ cofactors. They vanish simultaneously precisely when the matrix has rank at most $d-2$. The space of $d \times d$ matrices of rank at most r has dimension $r(2d-r)$: a rank- r factorization has two $d \times r$ factors, with r^2 parameters accounting for their common change of basis. Taking $r = d-2$, its codimension in the d^2 -dimensional matrix space is

$$d^2 - (d-2)(d+2) = (d - (d-2))^2 = 4.$$

If $\mathcal{L}(X) = AXC$, the chain rule gives $d(\det \circ \mathcal{L})_X = \mathcal{L}^*(d \det_{\mathcal{L}(X)})$. Surjectivity of $\mathcal{L}$ makes its dual $\mathcal{L}^*$ injective, so the critical locus is exactly the inverse image of the smaller determinant's critical locus. Moreover, choosing a complementary subspace to $\ker \mathcal{L}$ identifies the domain with $\ker \mathcal{L} \oplus \mathbb{C}^{d \times d}$. Under this identification, an inverse image is a product with $\ker \mathcal{L}$, which preserves codimension. The specialized determinant therefore also has critical-locus codimension 4. Thus this block construction cannot provide the $\Theta(n^2)$ critical-locus codimension that drove the permanent circuit bound.

Importantly, changing the specialization does not resolve the problem. Let $n \geq 5$, and obtain P by substituting affine linear forms into $\det_n$. Suppose that the resulting polynomial is homogeneous of degree $d > 2$. For a homogeneous polynomial of degree d , Euler's identity says

$$dP(x) = \sum_i x_i \frac{\partial P}{\partial x_i}(x).$$

Because the ground field has characteristic zero, every critical point of P also satisfies $P(x) = 0$. Consequently, $\text{Crit}(P)$ is exactly the singular locus of the hypersurface defined by $P = 0$.

Recall that the determinantal complexity $\text{dc}(P)$ is the smallest size of a matrix of affine linear forms having determinant P . Since P itself is an affine specialization of $\det_n$, the original specialization provides such a representation of size n ; therefore $\text{dc}(P) \leq n$. If $\text{codim Crit}(P) > 4$, then [ABV17, Thm. 1.2] implies $\text{codim Crit}(P) + 1 \leq \text{dc}(P) \leq n$: the singular locus of a

polynomial with a small determinantal representation cannot have arbitrarily large codimension. Otherwise, $\text{codim Crit}(P) \leq 4 \leq n - 1$. In either case,

$$\text{codim Crit}(P) \leq n - 1.$$

Finally, the slicing parameter k in Lem. 4.2 cannot exceed the critical-locus codimension: its hypothesis is precisely $\dim \text{Crit}(P) \leq m - k$, where m is the number of variables of P . Therefore $k \leq n - 1$. Also $d \leq n$, since substituting affine linear forms into a degree- n determinant cannot increase its degree. The largest circuit lower bound obtainable from (4.2) is consequently

$$\frac{k \log_2(d - 1)}{3} \leq \frac{(n - 1) \log_2(n - 1)}{3}.$$

For $d = 2$, the certificate is zero because $\log_2(d - 1) = 0$. Even for larger degrees, it is only $O(n \log n)$. This is below the elementary $n^2 - 1$ gate lower bound: the determinant depends on all n^2 individual input variables, and combining n^2 distinct inputs into one output with binary arithmetic gates requires at least $n^2 - 1$ gates. Thus this method, in its present form, cannot establish a superquadratic determinant lower bound.

12. RELATED WORK

The two lower bounds interact with several bodies of work that use similar words for substantially different computational models and complexity measures. We first distinguish those models and recall their established bounds. We then place the geometric circuit proof, the coefficient-independence formula proof, and the restricted-depth literature in their respective contexts. In particular, a lower bound for representing a polynomial as a determinant, a lower bound for general arithmetic circuits, and a lower bound for formulas are logically different statements.

12.1. Historical bounds and computational models. The modern complexity-theoretic comparison between permanent and determinant begins with two distinct results of Valiant. His algebraic completeness theorem makes the permanent complete for VNP over fields of characteristic different from two [Val79a]. Here completeness concerns families of polynomials and arithmetic projections: variables of one polynomial are replaced by variables or field constants to obtain another polynomial. His separate counting-complexity theorem proves that evaluating the permanent of a zero-one matrix is #P-complete [Val79b]. The two results motivate related questions, but one concerns symbolic arithmetic computation and the other a Boolean counting problem. In characteristic two, the determinant and permanent coincide, so even the distinction between the two polynomials depends on the ground field.

Several surveys describe this broader landscape. Bürgisser–Clausen–Shokrollahi [BCS97], Shpilka–Yehudayoff [SY10], and Kayal–Saptharishi [KS14] introduce algebraic complexity and arithmetic lower-bound methods. Bürgisser surveys algebraic completeness classes [Bür24]; Chen–Kayal–Wigderson survey partial-derivative techniques [CKW11]; and Bürgisser, Landsberg, Manivel, and Weyman [BLMW11], together with Bläser and Ikenmeyer [BI25], discuss geometric complexity theory.

It is important to distinguish three models. An unrestricted arithmetic circuit may reuse any previously computed intermediate value. A formula is a circuit whose underlying graph is a tree, so every reuse must instead be recomputed. A determinantal representation of f is an expression $f = \det(A)$ for a matrix A of affine linear forms; its minimum matrix dimension is the *determinantal complexity* $\text{dc}(f)$. The determinant is complete, under polynomial-size projections, for algebraic branching programs, or equivalently weakly skew circuits. It is not known to be complete for unrestricted polynomial-size arithmetic circuits [MP08]. Consequently, a superpolynomial lower bound on $\text{dc}(\text{per}_n)$ separates the permanent from the branching-program/weakly-skew model but does not by itself establish $\text{VP} \neq \text{VNP}$ for general circuits.

There is nevertheless a weaker connection to the general-circuit model. Valiant’s determinant simulation, together with circuit balancing, transforms a polynomial-size general circuit into a

quasipolynomial-size determinantal representation [Val79a, VSB83, BI25]. A *superquasipolynomial* lower bound on determinantal complexity would therefore rule out polynomial-size general circuits, whereas a merely superpolynomial determinantal bound would not. The distinction matters for Thm. 1.1: its lower bound concerns the number of gates in an unrestricted division-free circuit directly, not the dimension of a determinantal representation.

Write $N = n^2$ for the number of entries of an $n \times n$ matrix. Before the present result, the general-circuit lower bound known specifically for either $\det_n$ or per_n was only $\Omega(n^2) = \Omega(N)$. The classical $\Omega(N \log d)$ lower bounds of Strassen and Baur–Strassen apply to other explicit degree- d polynomial families, such as sums of powers; they do not assert the same bound for either matrix polynomial [Str73a, BS83, KS14]. Thm. 1.1 instead establishes the permanent-specific bound $\Omega(N \log \log N)$ for division-free circuits. It neither proves such a bound for the determinant nor treats circuits with division.

For upper bounds, the determinant has polynomial-size division-free circuits [Ber84]. Balancing and unfolding suitable circuits gives determinant formulas of size $n^{O(\log n)}$ [Hya79, VSB83]. For the permanent, Ryser’s inclusion–exclusion identity yields exponential-size circuits and, if intermediate computations cannot be shared, formulas of size $O(n^2 2^n)$ [Rys63].

For unrestricted division-free formulas, the previous lower bounds for both matrix polynomials were $\Omega(n^3) = \Omega(N^{3/2})$ [Kal85, KS14]. Kalorkoti’s determinant lower bound even permits rational formulas with valid division [Kal85]; his published result should not be interpreted as the corresponding division-formula theorem for the permanent. Thms. 1.2 and 1.3 raise the permanent-specific bounds to $\Omega(n^4 / \log n) = \Omega(N^2 / \log N)$ without division and with valid division, respectively. This comparison is specific to the permanent: for a different explicit polynomial, Chatterjee–Kumar–She–Volk prove an $\Omega(N^2)$ unrestricted formula lower bound for a suitable elementary symmetric polynomial [CKSV22]. Raz’s $n^{\Omega(\log n)}$ bounds for the determinant and permanent are numerically stronger but concern syntactically multilinear formulas, a restricted model rather than the general formulas considered here [Raz09].

12.2. Determinantal representations and geometric lower bounds. The permanent-versus-determinant problem is often stated in terms of determinantal representations: how large must a matrix of affine linear forms be if its determinant is per_n ? Early work of von zur Gathen and Cai studied versions of this projection problem [vzG87, Cai90]. Mignon–Ressayre proved

$$\text{dc}(\text{per}_n) \geq \frac{n^2}{2}$$

over characteristic zero by comparing Hessians at suitable points [MR04]. Cai–Chen–Li extend a quadratic bound to fields of characteristic different from two [CCL10]; in the opposite direction, Grenet gives determinantal representations of size $2^n - 1$ [Gre12].

There are also geometric variants in which the permanent is allowed to arise as a limit of determinantal representations. Landsberg–Manivel–Ressayre obtain quadratic lower bounds for this *border determinantal complexity* using dual varieties [LMR13]. Mulmuley–Sohoni place determinantal representation and orbit-closure questions in the representation-theoretic framework of geometric complexity theory [MS01]. All these results measure the dimension of a determinant representation, or its border analogue. They are not lower bounds on the number of gates in an arbitrary arithmetic circuit, and therefore should not be confused with prior superquadratic general-circuit bounds.

The geometric quantity relevant to the present paper is more elementary to describe. If P is a polynomial in m variables, its critical locus $\text{Crit}(P)$ is the set where all first partial derivatives vanish; its codimension in the ambient affine space $\mathbb{C}^m$ is $m - \dim \text{Crit}(P)$. The singular locus of the hypersurface $P = 0$ additionally requires $P = 0$; its codimension here is likewise measured in $\mathbb{C}^m$, not inside the hypersurface. For a homogeneous polynomial of positive degree over $\mathbb{C}$, Euler’s identity implies that the derivative equations already force $P = 0$, so these two loci agree.

These loci have several antecedents in algebraic complexity. Alper–Bogart–Velasco lower-bound determinantal complexity using the codimension of the singular locus and determine

the exact determinantal complexity of the 3×3 permanent [ABV17]. Chatterjee–Kumar–She–Volk use the common zero set of first derivatives for quadratic branching-program and formula lower bounds [CKSV22], while Gesmundo–Ghosal–Ikenmeyer–Lysikov relate singular-locus codimension to homogeneous branching-program complexity [GGIL22].

Our use of the same geometric object is different. We first obtain a homogeneous polynomial P by replacing entries of the permanent matrix with affine linear forms. We then prove that $\text{Crit}(P)$ has large codimension, choose a linear subspace avoiding its nonzero points, and apply a degree bound to the restricted gradient map. The resulting lower bound concerns the original unrestricted circuit directly; at no point do we identify circuit size with determinantal complexity.

The singular-locus theorem of Alper–Bogart–Velasco also makes the determinant obstruction transparent. Suppose $n \geq 5$, substitute affine linear forms into $\det_n$, and assume the resulting polynomial P is homogeneous of degree at least three. The specialization supplies an $n \times n$ determinantal representation, so $\text{dc}(P) \leq n$. If $\text{codim Crit}(P) > 4$, the singular-locus theorem in [ABV17] gives $\text{codim Crit}(P) + 1 \leq \text{dc}(P) \leq n$. If the codimension is at most 4, the assumption $n \geq 5$ already bounds it by $n - 1$. Consequently, in either case,

$$\text{codim Crit}(P) \leq n - 1.$$

Thus a determinant specialization cannot supply the $\Theta(n^2)$ critical-locus codimension needed by the present circuit argument. Section 11 explains the implication, including the low-codimension case, in detail.

A related literature studies the polynomial ideals generated by subpermanents. Laubacher and Swanson investigate permanental ideals [LS00]. Efremenko, Landsberg, Schenck, and Weyman study minimal free resolutions of such ideals motivated by geometric complexity theory [ELSW18a]. Boralevi, Carlini, Michałek, and Ventura give codimension bounds for permanental varieties [BCM25]. In particular, their results imply that, for $n \geq 6$, the critical locus of the full permanent satisfies

$$6 \leq \text{codim Crit}(\text{per}_n) \leq 2n.$$

The upper bound shows why applying our gradient argument directly to per_n cannot work: its critical-locus codimension is only $O(n)$, rather than the $\Theta(n^2)$ required in Prop. 4.3. The specially constructed affine specialization is therefore essential. There is also an important distinction between the ideals themselves. The cited permanental ideal papers usually generate an ideal from a *collection* of individual subpermanents, whereas §5 studies the first derivatives of a single *sum* of subpermanents. These families of equations need not define the same geometric object.

12.3. Differentiation, degree bounds, and the affine specialization. Proposition 4.3 combines two classical ideas: efficient simultaneous differentiation and a degree bound for polynomial maps. Strassen used Bézout-type degree arguments to lower-bound the simultaneous computation of powers and elementary symmetric functions [Str73a]. Baur–Strassen then showed that a straight-line program computing one polynomial f can be augmented to compute f and *all* its first partial derivatives with only constant-factor overhead [BS83]. The history and many applications of partial-derivative techniques are surveyed by Chen–Kayal–Wigderson [CKW11]; the Nisan–Wigderson partial-derivative method is another important, model-dependent development [NW97].

The special polynomial

$$f(x_1, \dots, x_k) = d^{-1} \sum_{i=1}^k x_i^d$$

makes the degree argument particularly transparent. Its gradient is $(x_1^{d-1}, \dots, x_k^{d-1})$, and setting these outputs equal to generic nonzero constants gives $(d-1)^k$ distinct solutions. Computing the same outputs with q multiplications produces a system containing q quadratic multiplication-gate equations. Bézout’s inequality bounds the number of isolated solutions of that system by 2^q . Consequently $(d-1)^k \leq 2^q$, yielding $q = \Omega(k \log d)$. This argument works because the chosen gradient has an isolated common zero and a large generic fiber. Neither fact is an

automatic consequence of the degree and number of variables of an arbitrary polynomial, so the classical $\Omega(k \log d)$ conclusion cannot simply be substituted for a permanent or determinant lower bound.

The exact differentiation overhead also depends on the cost model. Proposition 4.3 counts nonscalar multiplications and treats affine operations as free. For an original product gate, one forward multiplication and at most two reverse-mode multiplications suffice, giving the factor three used there. The full Baur–Strassen theorem also accounts for additions, scalar operations, and divisions; it does not justify the unqualified statement that a circuit with s total gates always has a gradient circuit with at most $3s$ total gates [BS83]. Furthermore, reverse-mode differentiation reuses intermediate values, an operation available to circuits but not to formulas. Ramya–Shastri exhibit multi-output formula and planar-circuit separations demonstrating that the analogous constant-overhead claim fails in those restricted models [RS26]. This is why the formula argument needs its separate coefficient-transcendence measure.

The remaining ingredients of the circuit proof likewise have classical precedents. A multiplication gate can be represented by a quadratic equation for its output; affine Bézout bounds the number of isolated common solutions; a generic linear slice avoids a sufficiently low-dimensional critical locus; and the degree of the resulting finite gradient map counts a generic fiber [Str73a, BS83, Hei83, KS14]. The challenge specific to the permanent is to construct a specialization for which these standard tools apply simultaneously: the specialized polynomial must have degree $d \asymp \log n$, retain $\Theta(n^2)$ variables, and have critical-locus codimension $\Theta(n^2)$.

The first ingredient in that construction is the minor-sum polynomial $M_{t,s,d}$ of (5.1). Its monomials encode size- d matchings between the rows and columns of a $t \times s$ matrix; summing the corresponding products is equivalent to summing all $d \times d$ permanental minors. To analyze its derivatives, §5 rewrites sums over injective assignments of rows to columns in terms of the power sums p_B . This is Möbius inversion on the partition lattice: each partition records which rows were assigned the same column, and its Möbius coefficient corrects the resulting overcount. The general incidence-algebra framework originates with Rota [Rot64]; the specific permanental expansion is the rectangular Binet–Minc identity [Min79]. Forbes also uses this identity for characteristic-independent set-multilinearization [For24].

The second ingredient is realizing an appropriate sum of these minor-sum polynomials as a single specialization of a larger permanent. Friedland–Levy already realize the sum of all $d \times d$ permanental minors as a larger permanent by adjoining all-ones and zero blocks [FL06]. Lemma 6.1 refines that completion: it selects constant columns using roots of unity so that the d upper rows not assigned to those constant columns all belong to a single row block. The surviving terms therefore split into the separate block polynomials whose critical loci can be bounded independently.

The square-zero algebra used to verify this cancellation has its own precedents. Feinsilver–McSorley use commuting square-zero “zeon” variables whose induced matrix coefficients are permanents [FM11]; Butera–Pernici use commuting Grassmann-even nilpotent variables to encode sums of permanental minors and bipartite matchings [BP15]. Signed and Fourier/root-of-unity coefficient filters likewise occur in Ryser’s and Glynn’s permanent identities and in Aaronson–Hance’s generalized Glynn estimator [Rys63, Gly10, AH14]. These precedents supply the underlying square-zero and cancellation ideas, but not the particular simultaneous block-selective specialization of Lem. 6.1. For the determinant, complementary maximal minors instead satisfy Plücker relations, obstructing the same block-supported pattern. The construction therefore uses a structural difference between permanents and determinants rather than merely replacing unsigned terms with signed ones.

12.4. Formula methods, coefficient independence, and division. The formula argument adapts a block-counting principle introduced in a different computational setting. Nečiporuk partitioned the variables of a Boolean function and lower-bounded formula size by adding the information required by the individual blocks [Nec66]. Kalorkoti translated this general idea to rational arithmetic formulas using algebraic independence of coefficient families [Kal85]. The exposition in [KS14, §3.2] describes the corresponding arithmetic measure explicitly.

Concretely, choose a block Y of variables of a polynomial f and expand f as a polynomial in Y . Its coefficients are polynomials in the remaining variables. The coefficient transcendence degree $\text{td}_Y(f)$ counts the maximum number of those coefficients that are algebraically independent: no nonzero polynomial relation with complex coefficients holds among the chosen coefficient polynomials. A formula with few occurrences of variables from Y can introduce only few independent coefficients. If the variable blocks are disjoint, their occurrence requirements can be added without counting any formula leaf twice. Applying this principle to n entry-disjoint full diagonals of a matrix gives the classical $\Omega(n^3)$ division-free formula bounds for both the permanent and determinant.

An especially close antecedent for the choice of blocks is work of Hrubeš–Joglekar [HJ25]. Their proof also selects a matching of $\Theta(\log n)$ matrix entries, exploits a specialization of the permanent on that matching, packs $\Theta(n^2/\log n)$ pairwise entry-disjoint matchings, and adds the resulting occurrence requirements. In the model of read-bounded determinantal representations, they obtain an $\Omega(n^{5/2}/\log n)$ lower bound on variable-containing matrix entries and rule out read- $o(\sqrt{n}/\log n)$ representations. The present argument uses the same matching-and-packing architecture in a different model. It proves that each short matching already produces $\Omega(n^2)$ algebraically independent coefficient polynomials, and charges this quantity directly to leaves of an unrestricted arithmetic formula.

The new independence witness can be understood through the Jacobian criterion. Over characteristic zero, coefficient polynomials are algebraically independent when an appropriate Jacobian matrix has full rank. Lem. 9.1 specializes the unmarked matrix entries so that this Jacobian factors into two Vandermonde evaluation matrices, one associated with an external row and the other with an external column. Their Kronecker product has full rank, certifying $\Omega(n^2)$ independent coefficients for a matching of only $O(\log n)$ entries. Algebraic independence and the Jacobian criterion also play central roles in identity testing and restricted-model lower bounds [BMS11]. In particular, Agrawal–Saha–Saptharishi–Saxena use Jacobian and transcendence-degree methods to lower-bound all immanants, including the permanent and determinant, in bounded-occurrence and bounded-transcendence-depth models [ASSS16]. Boralevi–Carlini–Michalek–Ventura also study algebraic independence of selected subpermanents [BCMV25], but their coefficient families and geometric objectives differ from those in Lem. 9.1.

This block-based method has a nearly matching intrinsic ceiling. For an N -variate multilinear polynomial and a block Y , its coefficient contribution satisfies

$$\text{td}_Y(f) \leq \min\{2^{|Y|}, N - |Y|\}.$$

Indeed, multilinearity gives at most $2^{|Y|}$ coefficient polynomials, while those coefficients involve only the $N - |Y|$ remaining variables. Summing this contribution over a partition of the variables gives at most $O(N^2/\log N)$ in the usual Nečiporuk–Kalorkoti framework; see [CKSV22]. Hence the permanent bound $\Omega(N^2/\log N)$ essentially saturates this particular framework. The $\Omega(N^2)$ lower bound for an elementary symmetric polynomial in [CKSV22] uses different methods.

The determinant illustrates why choosing very short matchings is specific to the permanent. The Schur-complement calculation in §11 shows that a matching Y of k entries has $\text{td}_Y(\det X) \leq k^2 + 1$. For $k = \Theta(\log n)$, this gives only $O(\log^2 n)$ independent coefficients, rather than the $\Omega(n^2)$ supplied by Lem. 9.1 for the permanent. Relations among the relevant principal minors are studied explicitly by Lin–Sturmfels [LinS09]. The same section proves that even arbitrary entry-disjoint blocks cannot make this coefficient-summing method yield a supercubic determinant formula bound.

Finally, allowing division requires care about the computational model. Strassen’s division-elimination theorem often replaces a circuit with division by a division-free circuit computing the same polynomial, but the cost depends on its degree [Str73b]. This transformation does not preserve the sharp formula-size or variable-occurrence estimates needed here. Kalorkoti instead analyzes rational formulas directly [Kal85], an approach followed by Lem. 10.2. Along a formula path containing only one marked child at each gate, the marked value is transformed by a fractional linear map $u \mapsto (au + b)/(cu + d)$. Such maps require only a bounded number of

parameters from the field of unmarked variables; at branching gates, the parameters contributed by the two marked children are added. A field-intersection argument then shows that when the final output is a polynomial, its coefficients belong to the resulting parameter field. Formula validity requires denominators to be nonzero as rational functions; it does *not* require those denominators to remain nonzero at the later Jacobian specialization, which is applied only to the polynomial coefficients.

12.5. Lower bounds in restricted arithmetic models. Many stronger-looking lower bounds impose structural restrictions on the computation rather than on the target polynomial. In a monotone computation over a semiring, subtraction and cancellation are unavailable. Jerrum–Snir obtain exponential lower bounds for monotone computations of the permanent [JS82]; these arguments do not apply here because our circuits and formulas allow arbitrary complex constants and cancellation.

At every multiplication gate of a syntactically multilinear computation, the two input subcomputations use disjoint sets of variables. Consequently every intermediate polynomial is multilinear, a stronger requirement than multilinearity of the final output alone. For this restricted formula model, Raz proves $n^{\Omega(\log n)}$ lower bounds for both the determinant and permanent [Raz09]. Raz–Yehudayoff prove exponential lower bounds for constant-depth multilinear circuits [RY09]. Although the determinant and permanent are themselves multilinear, a general circuit or formula computing one of them may pass through nonmultilinear intermediate polynomials; the restricted bounds therefore do not transfer to the unrestricted models of this paper.

Another important restriction is circuit depth. A depth-three $\Sigma\Pi\Sigma$ circuit expresses a polynomial as a sum of products of affine linear forms. Over characteristic zero, Shpilka–Wigderson use partial derivatives on low-codimension affine subspaces to obtain nearly quadratic depth-three determinant bounds in the number $N = n^2$ of matrix variables [SW01]. Over fixed finite fields, Grigoriev–Karpinski obtain exponential depth-three lower bounds for the determinant; Grigoriev–Razborov and the exposition in [KS14, §7] give corresponding results and variants for the permanent [GK98, GR00]. Those arguments exploit finite-field evaluation and are not characteristic-zero general-circuit arguments. In characteristic two, the determinant and permanent again coincide.

The distinction between homogeneous and nonhomogeneous circuits is also essential. In a homogeneous circuit, intermediate gates respect the degree grading; a nonhomogeneous circuit may create terms of different degrees and cancel them later. For homogeneous depth-three circuits, exponential lower bounds hold for both the determinant and permanent [NW97]. Nevertheless, over $\mathbb{Q}$ the determinant has nonhomogeneous depth-three circuits of size $\exp(O(\sqrt{n} \log n))$ [GKKS16]. Thus a homogeneous lower bound does not automatically extend to a general circuit of the same depth.

Several complexity measures yield additional restricted-model bounds. Nisan–Wigderson use the dimension of a space of partial derivatives [NW97]; Jacobian and algebraic-independence methods apply to bounded-occurrence and related models [ASSS16]; and shifted partial derivatives yield bounds in homogeneous, low-depth, or bounded-bottom-fan-in settings [GKKS14]. In particular, such methods give $\exp(\Omega(\sqrt{n}))$ lower bounds for homogeneous depth-four circuits of bottom fan-in $O(\sqrt{n})$ computing either the $n \times n$ determinant or the $n \times n$ permanent.

The fact that many of these measures treat the two polynomials alike is itself supported by barrier results. Efremenko–Landsberg–Schenck–Weyman show that shifted partial derivatives cannot separate the padded permanent from the determinant in the relevant orbit-closure regime [ELSW18b]; Gesmundo–Landsberg identify related limitations for unpadded derivative spaces [GL19]. In a different geometric complexity theory setting, Bürgisser–Ikenmeyer–Panova prove a barrier for representation-theoretic “occurrence obstructions” [BIP19]. These results help explain why a large derivative-based measure often proves a lower bound for *both* matrix polynomials in a restricted model, rather than a permanent-versus-determinant separation. With a different symmetry restriction, Dawar–Wiltschko do obtain an exponential permanent lower bound while retaining polynomial-size determinant circuits [DW25]. This is a genuine separation inside that restricted symmetric model, not a lower bound for unrestricted circuits.

Depth-reduction theorems explain why shallow circuits nevertheless attract sustained attention. Agrawal–Vinay reduce general arithmetic circuits to depth four [AV08], and Gupta–Kamath–Kayal–Saptharishi obtain a “chasm at depth three” over characteristic zero [GKKS16]. These simulations increase circuit size superpolynomially at the relevant degrees, so an arbitrary lower bound for the resulting shallow model does not immediately yield a comparable unrestricted-circuit lower bound.

Recent constant-depth lower bounds extend beyond the early homogeneous and quadratic depth-three results. Limaye–Srinivasan–Tavenas prove superpolynomial lower bounds for *general, possibly nonhomogeneous* constant-depth circuits computing iterated matrix multiplication over characteristic zero or sufficiently large characteristic [LST25]; Forbes extends the conclusion to every field [For24]. Iterated matrix multiplication is a polynomial-size projection of both the determinant and permanent: its source–sink path polynomial can be represented by a cycle-cover matrix with self-loops, with the determinant differing only by a uniform sign. Since substituting variables and constants preserves constant depth, a small constant-depth circuit for either matrix polynomial would produce one for the projected iterated matrix multiplication polynomial. Thus these results also give superpolynomial constant-depth lower bounds for both matrix families.

More precisely, over characteristic zero, the depth-three consequence is at least $n^{\Omega(\sqrt{\log n})}$ for each $n \times n$ matrix polynomial [LST25]. Over every field, one obtains in particular the depth-three bound

$$\exp\left(\Omega\left(\frac{(\log n)^{3/2}}{\sqrt{\log \log n}}\right)\right)$$

[For24]. Bhargav–Dutta–Saxena improve the constant-depth parameters and identify a barrier for the associated measure [BDS24]. None of these results conflicts with polynomial-size *unrestricted-depth* determinant circuits, and none distinguishes the permanent from the determinant in the constant-depth model under discussion.

REFERENCES

- [AH14] S. Aaronson and T. Hance, *Generalizing and derandomizing Gurvits’s approximation algorithm for the permanent*, Quantum Inf. Comput. **14** (2014), nos. 7–8, 541–559, doi:10.26421/QIC14.7-8-1.
- [ASSS16] M. Agrawal, C. Saha, R. Saptharishi, and N. Saxena, *Jacobian hits circuits: Hitting sets, lower bounds for depth- D occur- k formulas and depth-3 transcendence degree- k circuits*, SIAM J. Comput. **45** (2016), no. 4, 1533–1562, doi:10.1137/130910725.
- [AV08] M. Agrawal and V. Vinay, *Arithmetic circuits: A chasm at depth four*, in *49th Annual IEEE Symposium on Foundations of Computer Science*, IEEE, 2008, 67–75, doi:10.1109/FOCS.2008.32.
- [ABV17] J. Alper, T. Bogart, and M. Velasco, *A lower bound for the determinantal complexity of a hypersurface*, Found. Comput. Math. **17** (2017), no. 3, 829–836, doi:10.1007/s10208-015-9300-x.
- [BS83] W. Baur and V. Strassen, *The complexity of partial derivatives*, Theoret. Comput. Sci. **22** (1983), no. 3, 317–330, doi:10.1016/0304-3975(83)90110-X.
- [BMS11] M. Beecken, J. Mittmann, and N. Saxena, *Algebraic independence and blackbox identity testing*, in *Automata, Languages and Programming, Part II*, L. Aceto, M. Henzinger, and J. Sgall, eds., Lecture Notes in Comput. Sci. **6756**, Springer, 2011, 137–148, doi:10.1007/978-3-642-22012-8_10.
- [Ber84] S. J. Berkowitz, *On computing the determinant in small parallel time using a small number of processors*, Inform. Process. Lett. **18** (1984), no. 3, 147–150, doi:10.1016/0020-0190(84)90018-8.
- [BDS24] C. S. Bhargav, S. Dutta, and N. Saxena, *Improved lower bound, and proof barrier, for constant depth algebraic circuits*, ACM Trans. Comput. Theory **16** (2024), no. 4, Art. 23, 1–22, doi:10.1145/3689957.
- [BI25] M. Bläser and C. Ikenmeyer, *Introduction to geometric complexity theory*, Theory Comput. Libr., Grad. Surv. **10** (2025), 1–166, doi:10.4086/toc.gs.2025.010.
- [BCM25] A. Boralevi, E. Carlini, M. Michalek, and E. Ventura, *On the codimension of permanent varieties*, Adv. Math. **461** (2025), Art. 110079, doi:10.1016/j.aim.2024.110079.
- [Bür24] P. Bürgisser, *Completeness classes in algebraic complexity theory*, 2024, arXiv:2406.06217.
- [BCS97] P. Bürgisser, M. Clausen, and M. A. Shokrollahi, *Algebraic Complexity Theory*, Grundlehren Math. Wiss. **315**, Springer, Berlin, 1997, doi:10.1007/978-3-662-03338-8.
- [BIP19] P. Bürgisser, C. Ikenmeyer, and G. Panova, *No occurrence obstructions in geometric complexity theory*, J. Amer. Math. Soc. **32** (2019), no. 1, 163–193, doi:10.1090/jams/908.
- [BLMW11] P. Bürgisser, J. M. Landsberg, L. Manivel, and J. Weyman, *An overview of mathematical issues arising in the geometric complexity theory approach to VP $\neq$ VNP*, SIAM J. Comput. **40** (2011), no. 4, 1179–1209, doi:10.1137/090765328.

- [BP15] P. Butera and M. Pernici, *Sums of permanent minors using Grassmann algebra*, Int. J. Graph Theory Appl. **1** (2015), no. 2, 83–96, [arXiv:1406.5337](#).
- [Cai90] J.-Y. Cai, *A note on the determinant and permanent problem*, Inform. Comput. **84** (1990), no. 1, 119–127, [doi:10.1016/0890-5401\(90\)90036-H](#).
- [CCL10] J.-Y. Cai, X. Chen, and D. Li, *Quadratic lower bound for permanent vs. determinant in any characteristic*, Comput. Complexity **19** (2010), no. 1, 37–56, [doi:10.1007/s00037-009-0284-2](#).
- [CKSV22] P. Chatterjee, M. Kumar, A. She, and B. L. Volk, *Quadratic lower bounds for algebraic branching programs and formulas*, Comput. Complexity **31** (2022), no. 2, Art. 8, [doi:10.1007/s00037-022-00223-8](#).
- [CKW11] X. Chen, N. Kayal, and A. Wigderson, *Partial derivatives in arithmetic complexity and beyond*, Found. Trends Theor. Comput. Sci. **6** (2011), nos. 1–2, 1–138, [doi:10.1561/04000000043](#).
- [DW25] A. Dawar and G. Wilsenach, *Symmetric arithmetic circuits*, Theory Comput. **21** (2025), no. 14, 1–32, [doi:10.4086/toc.2025.v021a014](#).
- [ELSW18a] K. Efremenko, J. M. Landsberg, H. Schenck, and J. Weyman, *On minimal free resolutions of sub-permanents and other ideals arising in complexity theory*, J. Algebra **503** (2018), 8–20, [doi:10.1016/j.jalgebra.2018.01.021](#).
- [ELSW18b] K. Efremenko, J. M. Landsberg, H. Schenck, and J. Weyman, *The method of shifted partial derivatives cannot separate the permanent from the determinant*, Math. Comp. **87** (2018), 2037–2045, [doi:10.1090/mcom/3284](#).
- [Eis95] D. Eisenbud, *Commutative Algebra with a View Toward Algebraic Geometry*, Grad. Texts in Math. **150**, Springer, New York, 1995, [doi:10.1007/978-1-4612-5350-1](#).
- [FM11] P. Feinsilver and J. McSorley, *Zeons, permanents, the Johnson scheme, and generalized derangements*, Int. J. Combin. **2011** (2011), Art. 539030, [doi:10.1155/2011/539030](#).
- [For24] M. A. Forbes, *Low-depth algebraic circuit lower bounds over any field*, in *39th Computational Complexity Conference*, Leibniz Int. Proc. Inform. **300**, Schloss Dagstuhl, 2024, 31:1–31:16, [doi:10.4230/LIPIcs.CCC.2024.31](#).
- [FL06] S. Friedland and D. Levy, *A polynomial-time approximation algorithm for the number of k -matchings in bipartite graphs*, in *Mathematical Papers in Honour of Eduardo Marques de Sá*, Textos Mat. Sér. B **39**, Univ. Coimbra, Coimbra, 2006, 61–67, [arXiv:cs/0607135](#).
- [GGIL22] F. Gesmundo, P. Ghosal, C. Ikenmeyer, and V. Lysikov, *Degree-restricted strength decompositions and algebraic branching programs*, in *42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, Leibniz Int. Proc. Inform. **250**, Schloss Dagstuhl, 2022, 20:1–20:15, [doi:10.4230/LIPIcs.FSTTCS.2022.20](#).
- [GL19] F. Gesmundo and J. M. Landsberg, *Explicit polynomial sequences with maximal spaces of partial derivatives and a question of K. Mulmuley*, Theory Comput. **15** (2019), no. 3, 1–24, [doi:10.4086/toc.2019.v015a003](#).
- [Gly10] D. G. Glynn, *The permanent of a square matrix*, European J. Combin. **31** (2010), no. 7, 1887–1891, [doi:10.1016/j.ejc.2010.01.010](#).
- [Gre12] B. Grenet, *An upper bound for the permanent versus determinant problem*, manuscript, 2012, [author’s manuscript](#).
- [GK98] D. Grigoriev and M. Karpinski, *An exponential lower bound for depth 3 arithmetic circuits*, in *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, ACM, 1998, 577–582, [doi:10.1145/276698.276872](#).
- [GR00] D. Grigoriev and A. A. Razborov, *Exponential lower bounds for depth 3 arithmetic circuits in algebras of functions over finite fields*, Appl. Algebra Engrg. Comm. Comput. **10** (2000), no. 6, 465–487, [doi:10.1007/s002009900021](#).
- [GKKS14] A. Gupta, P. Kamath, N. Kayal, and R. Saptharishi, *Approaching the chasm at depth four*, J. ACM **61** (2014), no. 6, Art. 33, 1–16, [doi:10.1145/2629541](#).
- [GKKS16] A. Gupta, P. Kamath, N. Kayal, and R. Saptharishi, *Arithmetic circuits: A chasm at depth 3*, SIAM J. Comput. **45** (2016), no. 3, 1064–1079, [doi:10.1137/140957123](#).
- [Har77] R. Hartshorne, *Algebraic Geometry*, Grad. Texts in Math. **52**, Springer, New York, 1977, [doi:10.1007/978-1-4757-3849-0](#).
- [Hei83] J. Heintz, *Definability and fast quantifier elimination in algebraically closed fields*, Theoret. Comput. Sci. **24** (1983), no. 3, 239–277, [doi:10.1016/0304-3975\(83\)90002-6](#); corrigendum, **39** (1985), no. 2–3, 343, [doi:10.1016/0304-3975\(85\)90150-1](#).
- [HJ25] P. Hrubeš and P. S. Joglekar, *On read- k projections of the determinant*, in *42nd International Symposium on Theoretical Aspects of Computer Science*, Leibniz Int. Proc. Inform. **327**, Schloss Dagstuhl, 2025, 53:1–53:7, [doi:10.4230/LIPIcs.STACS.2025.53](#).
- [Hya79] L. Hyafil, *On the parallel evaluation of multivariate polynomials*, SIAM J. Comput. **8** (1979), no. 2, 120–123, [doi:10.1137/0208010](#).
- [JS82] M. Jerrum and M. Snir, *Some exact complexity results for straight-line computations over semirings*, J. ACM **29** (1982), no. 3, 874–897, [doi:10.1145/322326.322341](#).
- [Kal85] K. A. Kalorkoti, *A lower bound for the formula size of rational functions*, SIAM J. Comput. **14** (1985), no. 3, 678–687, [doi:10.1137/0214050](#).

- [KS14] N. Kayal and R. Saptharishi, *A selection of lower bounds for arithmetic circuits*, in *Perspectives in Computational Complexity*, M. Agrawal and V. Arvind, eds., Progr. Comput. Sci. Appl. Logic **26**, Birkhäuser, Cham, 2014, 77–115, doi:10.1007/978-3-319-05446-9_5.
- [LMR13] J. M. Landsberg, L. Manivel, and N. Ressayre, *Hypersurfaces with degenerate duals and the geometric complexity theory program*, Comment. Math. Helv. **88** (2013), no. 2, 469–484, doi:10.4171/CMH/292.
- [LS00] R. Laubenbacher and I. Swanson, *Permanental ideals*, J. Symbolic Comput. **30** (2000), no. 2, 195–205, doi:10.1006/jSCO.2000.0363.
- [LST25] N. Limaye, S. Srinivasan, and S. Tavenas, *Superpolynomial lower bounds against low-depth algebraic circuits*, J. ACM **72** (2025), no. 4, Art. 26, 1–35, doi:10.1145/3734215.
- [LinS09] S. Lin and B. Sturmfels, *Polynomial relations among principal minors of a 4×4 -matrix*, J. Algebra **322** (2009), no. 11, 4121–4131, doi:10.1016/j.jalgebra.2009.06.026.
- [MP08] G. Malod and N. Portier, *Characterizing Valiant’s algebraic complexity classes*, J. Complexity **24** (2008), no. 1, 16–38, doi:10.1016/j.jCO.2006.09.006.
- [MR04] T. Mignon and N. Ressayre, *A quadratic bound for the determinant and permanent problem*, Int. Math. Res. Not. **2004** (2004), no. 79, 4241–4253, doi:10.1155/S1073792804142566.
- [Min79] H. Minc, *Evaluation of permanents*, Proc. Edinburgh Math. Soc. (2) **22** (1979), no. 1, 27–32, doi:10.1017/S0013091500027760.
- [MS01] K. D. Mulmuley and M. Sohoni, *Geometric complexity theory I: An approach to the P vs. NP and related problems*, SIAM J. Comput. **31** (2001), no. 2, 496–526, doi:10.1137/S009753970038715X.
- [Nec66] É. I. Nechiporuk, *On a Boolean function*, Dokl. Akad. Nauk SSSR **169** (1966), no. 4, 765–766; English transl., Soviet Math. Dokl. **7** (1966), 999–1000, Math-Net.Ru:dan32449.
- [NW97] N. Nisan and A. Wigderson, *Lower bounds on arithmetic circuits via partial derivatives*, Comput. Complexity **6** (1996/97), no. 3, 217–234, doi:10.1007/BF01294256.
- [RS26] C. Ramya and P. Shastri, *Lower bounds for planar arithmetic circuits*, ACM Trans. Comput. Theory **18** (2026), no. 1, Art. 9, 1–23, doi:10.1145/3778858; conference version, *15th Innovations in Theoretical Computer Science Conference*, Leibniz Int. Proc. Inform. **287**, Schloss Dagstuhl, 2024, 91:1–91:22, doi:10.4230/LIPIcs.ITCS.2024.91.
- [Raz09] R. Raz, *Multi-linear formulas for permanent and determinant are of super-polynomial size*, J. ACM **56** (2009), no. 2, Art. 8, 1–17, doi:10.1145/1502793.1502797.
- [RY09] R. Raz and A. Yehudayoff, *Lower bounds and separations for constant depth multilinear circuits*, Comput. Complexity **18** (2009), no. 2, 171–207, doi:10.1007/s00037-009-0270-8.
- [Rot64] G.-C. Rota, *On the foundations of combinatorial theory I: Theory of Möbius functions*, Z. Wahrscheinlichkeitstheorie Verw. Gebiete **2** (1964), 340–368, doi:10.1007/BF00531932.
- [Rys63] H. J. Ryser, *Combinatorial Mathematics*, Carus Math. Monogr. **14**, Mathematical Association of America, 1963.
- [SW01] A. Shpilka and A. Wigderson, *Depth-3 arithmetic circuits over fields of characteristic zero*, Comput. Complexity **10** (2001), no. 1, 1–27, doi:10.1007/PL00001609.
- [SY10] A. Shpilka and A. Yehudayoff, *Arithmetic circuits: A survey of recent results and open questions*, Found. Trends Theor. Comput. Sci. **5** (2010), no. 3–4, 207–388, doi:10.1561/04000000039.
- [Str73a] V. Strassen, *Die Berechnungskomplexität von elementarsymmetrischen Funktionen und von Interpolationskoeffizienten*, Numer. Math. **20** (1973), 238–251, doi:10.1007/BF01436566.
- [Str73b] V. Strassen, *Vermeidung von Divisionen*, J. Reine Angew. Math. **264** (1973), 184–202, EuDML:151394.
- [Val79a] L. G. Valiant, *Completeness classes in algebra*, in *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, ACM, 1979, 249–261, doi:10.1145/800135.804419.
- [Val79b] L. G. Valiant, *The complexity of computing the permanent*, Theoret. Comput. Sci. **8** (1979), no. 2, 189–201, doi:10.1016/0304-3975(79)90044-6.
- [VSB83] L. G. Valiant, S. Skyum, S. Berkowitz, and C. Rackoff, *Fast parallel computation of polynomials using few processors*, SIAM J. Comput. **12** (1983), no. 4, 641–644, doi:10.1137/0212043.
- [vzG87] J. von zur Gathen, *Permanent and determinant*, Linear Algebra Appl. **96** (1987), 87–100, doi:10.1016/0024-3795(87)90337-5.

CHAPTER 6

Exponential Parallel Repetition for All Two-Player Entangled Games

ABSTRACT. In a two-player entangled game G , a referee sends a question to each of two noncommunicating players, receives their answers, and decides whether they win. The players may share an arbitrary entangled quantum state, and the supremum of their winning probabilities is the entangled value $\omega^*(G)$. In the repeated game $G^{\otimes n}$, the referee plays n independent copies and accepts only if the players win every copy.

Raz's celebrated classical parallel repetition theorem (STOC 1995) shows that, for classical players, the repeated value decreases exponentially whenever the original value is less than one. Whether the same holds for arbitrary entangled games has been a longstanding open problem. We resolve this quantum analogue affirmatively: for every finite two-player, one-round game G with $\omega^*(G) = 1 - \varepsilon < 1$ and answer alphabets A, B ,

$$\omega^*(G^{\otimes n}) \leq \exp\left(-c_{\text{qs}} \frac{\varepsilon^{13}}{\varepsilon + \log(|A||B|)} n\right), \quad n \geq 1,$$

for a universal constant $c_{\text{qs}} > 0$.

Previously, Yuen (ICALP 2016) proved polynomial decay for arbitrary entangled games, and Bavarian, Vidick, and Yuen (STOC 2017) proved exponential decay for anchored games obtained by modifying the original game. Our proof builds on Yuen's conditioning and dependency-breaking framework. Its main new ingredient is a postselection-stable quantum sampleability estimate that avoids an inverse dependence on the probability of the conditioning event.

CONTENTS

1. Introduction
 2. Preliminaries
 3. Proof of the parallel repetition theorem
 4. Proof of the postselected sampleability lemma
 - A. Deferred auxiliary proofs
- References

1 Introduction

Parallel repetition is a basic method for reducing the error of multiprover interactive proofs. Starting from a game that dishonest players cannot always win, the verifier plays many independent copies and accepts only if every copy is won. The fundamental question is whether requiring simultaneous success makes the players' winning probability exponentially small.

In more detail, a finite two-player, one-round game is specified by $G = (X, Y, A, B, \mu, V)$. A referee samples a pair of questions $(x, y) \sim \mu$, sends x to Alice and y to Bob, and receives answers $a \in A$ and $b \in B$. The players win if $V(x, y, a, b) = 1$. They may agree on a strategy and share randomness before the game, but they cannot communicate after receiving their questions. The maximum winning probability of such a strategy is the classical value $\omega(G)$.

In the parallel repetition $G^{\otimes n}$, the referee independently samples n question pairs and accepts precisely when the players win all n coordinates. The coordinates are independent in the referee's experiment, but Alice's answer to a coordinate can depend on all her questions, and similarly for Bob. Consequently, the winning probability of an arbitrary repeated-game strategy need not factor across coordinates. Nevertheless, Raz's parallel repetition theorem shows that every finite classical game with $\omega(G) < 1$ satisfies $\omega(G^{\otimes n}) \leq \exp(-c_G n)$ for some $c_G > 0$ depending on the game [Raz98]. Holenstein later gave an information-theoretic proof and improved the quantitative dependence on the soundness gap [Hol09].

In the quantum setting, an entangled game has exactly the same questions, answers, and acceptance rule, but the players are additionally allowed to share a bipartite quantum state before receiving their questions. Alice measures her part of this state according to x , Bob measures his part according to y , and their classical outcomes determine their answers. Their optimal winning probability, taken over arbitrary finite-dimensional shared states and local measurements, is the entangled value $\omega^*(G)$. In $G^{\otimes n}$, each player may perform one joint measurement depending on their entire question tuple. Playing the coordinates independently shows only that

$$\omega^*(G^{\otimes n}) \geq \omega^*(G)^n;$$

the difficulty is to prove a matching exponential *upper* bound for unrestricted entangled strategies.

The general quantum analogue of Raz's theorem was explicitly noted as open by 2004 [CHTW04, footnote 2]. In its basic qualitative form, the question is:

(Quantum Parallel Repetition Conjecture) For every finite two-player entangled game G with $\omega^(G) < 1$, is there a constant $c_G > 0$ such that*

$$\omega^*(G^{\otimes n}) \leq \exp(-c_G \cdot n) \quad \text{for every } n \geq 1?$$

Exponential repetition was subsequently proved for several important classes of games, but the question for arbitrary games remained open. The strongest general theorem, due to Yuen, established polynomial rather than exponential decay [Yue16].

1.1 Our result

We give an affirmative answer to the quantum parallel repetition question for every finite two-player, one-round entangled game.

Theorem 1.1 (Quantum parallel repetition theorem for two-player games). *There exists a universal constant $c_{\text{qs}} > 0$ such that the following holds. Let $G = (X, Y, A, B, \mu, V)$ be any finite two-player, one-round entangled game with nonempty answer alphabets, and put*

$$\varepsilon = 1 - \omega^*(G) > 0, \quad \ell = \log(|A||B|).$$

Then, for every integer $n \geq 1$,

$$\omega^*(G^{\otimes n}) \leq \exp\left(-c_{\text{qs}} \frac{\varepsilon^{13}}{\varepsilon + \log(|A||B|)} n\right) \quad (n \geq 1).$$

The exponent 13 is not claimed to be optimal. It arises from the quantitative loss in a standard quantum correlated-sampling lemma [DSV15, Lemma 17]. The central point is that the decay is exponential for *every* finite entangled game.

1.2 Previous work

Classical parallel repetition. Raz proved exponential parallel repetition for all finite classical two-player, one-round games [Raz98]. Holenstein subsequently introduced an information-theoretic approach based on correlated sampling and obtained the bound [Hol09, Theorem 2.5]

$$\omega(G^{\otimes n}) \leq \exp\left(-\Omega\left(\frac{\varepsilon^3 n}{1 + \log(|A||B|)}\right)\right), \quad \varepsilon = 1 - \omega(G).$$

The dependence on the soundness gap can be improved for certain structured games [Rao11], while Raz’s counterexample shows that the strongest possible linear-gap bound does not hold for all games [Raz11].

Entangled games with additional structure. Perfect parallel repetition holds for quantum XOR games [CSUU08], and exponential repetition has been established for entangled unique games [KRT10], projection games [DSV15], and free games, where the players’ questions are independently distributed [JPY14, CS15]. Other bounds apply to games with strictly positive Cartesian question support but can depend on the minimum question probability [CS14]. These results do not resolve the question for arbitrary predicates and correlated question distributions.

General games. For unrestricted entangled games, Yuen proved [Yue16, Theorem 1] that if $\omega^*(G) = 1 - \varepsilon$, then, for $n \geq 2$,

$$\omega^*(G^{\otimes n}) \leq \frac{c s_G \log n}{\varepsilon^{17} n^{1/4}},$$

where c is universal and s_G is the bit-length of the answer alphabet. This theorem already applies to arbitrary question distributions and unchanged games; the remaining gap is that its decay is polynomial rather than exponential.

Bavarian, Vidick, and Yuen obtained exponential amplification by transforming the original game into an anchored game [BVY17]. Earlier work of Kempe and Vidick also obtained general-game amplification by introducing additional consistency or dummy questions [KV11]. These results are useful for hardness amplification, but a repetition theorem for a modified game does not imply one for the original game.

Our proof follows the general conditioning and sampleability framework of Yuen [Yue16, Section 3.2], which itself builds on Holenstein’s classical correlated sampling [Hol09, Lemma 5.2 and Corollary 5.3] and the quantum correlated-sampling lemma of Dinur, Steurer, and Vidick [DSV15, Lemma 17]. The new ingredient is a quantum sampleability estimate that remains effective even when the conditioning event has exponentially small probability.

2 Preliminaries

We first fix probability and operator notation, then formally define finite one-round entangled games and their standard parallel repetition. Standard finite-dimensional quantum background appears in [Wat18, Chapters 2 and 3] and [NC10].

2.1 Notation and probability conventions

All logarithms are natural. For $n \geq 1$, write $[n] = \{1, \dots, n\}$. If $S \subseteq [n]$ and x^n is a tuple, write $x_S = (x_j)_{j \in S}$, and similarly for y_S, a_S, b_S . Empty products equal one.

Named probability distributions are written in calligraphic font. In particular, $\mathcal{P}$ denotes an original experiment, $\mathcal{Q}$ its conditioned distribution, and $\mathcal{J}_A, \mathcal{J}_B$ locally generated distributions. The referee's question distribution retains the conventional symbol μ ; Roman E_x^a, F_y^b denote measurement operators, not probability distributions.

For probability distributions $\mathcal{D}, \mathcal{E}$ on a finite set, relative entropy and total variation are

$$D(\mathcal{D} \parallel \mathcal{E}) = \sum_u \mathcal{D}(u) \log \frac{\mathcal{D}(u)}{\mathcal{E}(u)},$$

$$d_{\text{TV}}(\mathcal{D}, \mathcal{E}) = \frac{1}{2} \sum_u |\mathcal{D}(u) - \mathcal{E}(u)|.$$

The conventions $0 \log 0 = 0$ and $D(\mathcal{D} \parallel \mathcal{E}) = +\infty$ when $\text{supp } \mathcal{D} \not\subseteq \text{supp } \mathcal{E}$ are understood. With natural logarithms, Pinsker's inequality reads

$$d_{\text{TV}}(\mathcal{D}, \mathcal{E}) \leq \sqrt{\frac{1}{2} D(\mathcal{D} \parallel \mathcal{E})}.$$

Probabilities and expectations under a specified distribution are written $\mathbb{P}$ and $\mathbb{E}$. Conditioning is used only when the conditioned event has strictly positive probability.

For a finite-dimensional Hilbert space $\mathcal{H}$, $|\psi\rangle$ denotes a vector, $F^\dagger$ the adjoint of an operator, and $F \succeq 0$ positive semidefiniteness. A positive contraction satisfies $0 \preceq F \preceq I$. A positive-operator-valued measure (POVM) is a finite family of positive semidefinite operators summing to the identity.

2.2 Finite one-round two-player entangled games

The referee samples one possibly correlated question pair and sends one question to each player. Entanglement can correlate the answers, but the players cannot communicate after receiving their questions.

Definition 2.1 (Finite one-round two-player entangled game). A game is a tuple

$$G = (X, Y, A, B, \mu, V), \quad V : X \times Y \times A \times B \longrightarrow \{0, 1\},$$

where the question and answer sets are finite, both answer alphabets A, B are nonempty, and μ is an arbitrary probability distribution on $X \times Y$. A finite-dimensional tensor-product strategy consists of finite local Hilbert spaces $\mathcal{H}_A, \mathcal{H}_B$, a unit vector $\psi \in \mathcal{H}_A \otimes \mathcal{H}_B$, and local POVMs

$$E_x^a \succeq 0, \quad \sum_{a \in A} E_x^a = I_{\mathcal{H}_A}, \quad F_y^b \succeq 0, \quad \sum_{b \in B} F_y^b = I_{\mathcal{H}_B}.$$

Its winning probability is

$$\omega^*(G; \psi, E, F) = \sum_{x,y} \mu(x, y) \sum_{a,b} V(x, y, a, b) \langle \psi | E_x^a \otimes F_y^b | \psi \rangle.$$

The entangled value is the supremum

$$\omega^*(G) = \sup_{\substack{\mathcal{H}_A, \mathcal{H}_B \text{ finite} \\ \psi, E, F}} \omega^*(G; \psi, E, F).$$

No uniform bound on the local dimensions and no attaining strategy are assumed.

Define the question marginals

$$\mu_X(x) = \sum_y \mu(x, y), \quad \mu_Y(y) = \sum_x \mu(x, y).$$

A question with zero marginal probability can be deleted without changing either game value. On the resulting positive-marginal question sets, write

$$\mu(x | y) = \frac{\mu(x, y)}{\mu_Y(y)}, \quad \mu(y | x) = \frac{\mu(x, y)}{\mu_X(x)}.$$

These expressions are evaluated only where their denominators are positive. The actual question-pair support is

$$E_\mu = \{(x, y) \in X \times Y : \mu(x, y) > 0\};$$

it need not equal $X \times Y$ and need not be connected.

The tensor product expresses locality, not independence of the players' answers. In particular, writing $\rho_{AB} = |\psi\rangle\langle\psi|$, the actual Born probabilities are

$$\mathbb{P}(a, b | x, y) = \text{Tr}(\rho_{AB}(E_x^a \otimes F_y^b)).$$

The state is prepared before either question is received.

2.3 Standard parallel repetition

Definition 2.2 (Standard parallel repetition). For $n \geq 1$, the repeated game $G^{\otimes n}$ samples

$$(x^n, y^n) \sim \mu^{\otimes n},$$

sends x^n to Alice and y^n to Bob, permits arbitrary finite-dimensional joint local POVMs

$$\{E_{x^n}^{a^n} : a^n \in A^n\}, \quad \{F_{y^n}^{b^n} : b^n \in B^n\},$$

and accepts according to

$$V^{\otimes n}(x^n, y^n, a^n, b^n) = \prod_{i=1}^n V(x_i, y_i, a_i, b_i).$$

Write $v_n = \omega^*(G^{\otimes n})$. Thus a strategy wins the repeated game exactly when it wins every original coordinate.

For $S \subseteq [n]$, write

$$\mu^S(x_S, y_S) = \prod_{j \in S} \mu(x_j, y_j), \quad \mu_X^S(x_S) = \prod_{j \in S} \mu_X(x_j), \quad \mu_Y^S(y_S) = \prod_{j \in S} \mu_Y(y_j).$$

These products express independence between distinct game coordinates, not independence between the two questions of one coordinate.

Although the sampled question pairs are independent between coordinates, Alice's i th answer may depend on all of x^n , and Bob's i th answer may depend on all of y^n . In particular,

$$\omega^*(G^{\otimes n}) \geq \omega^*(G)^n$$

is a lower bound obtained from independent strategies, not an upper bound on unrestricted joint strategies. For a repeated strategy, write $\rho_{AB}^{(n)} = |\psi^{(n)}\rangle\langle\psi^{(n)}|$ for its preshared density operator.

The resulting precise game definitions are those used in Theorem 1.1, which was stated in the introduction.

3 Proof of the parallel repetition theorem

We prove Theorem 1.1 by contradiction. Suppose that a strategy for $G^{\otimes n}$ wins all n coordinates simultaneously with probability exceeding the claimed exponential bound. Lemma 3.1 selects a small core D such that, conditional on winning every coordinate in D , a uniformly chosen remaining coordinate wins with high probability. Lemma 3.2 represents this conditioned coordinate by an ideal shared state and local measurements.

This ideal experiment is not yet a legal single-game strategy: the conditioned question distribution may differ from μ , its history is not supplied to the players, and its state depends jointly on their separate questions. The central Lemma 3.3 produces locally generated histories and locally describable states that approximate the ideal experiment. Lemmas 3.4 and 3.5 then use classical and quantum correlated sampling to synchronize those histories and prepare the state, and Lemma 3.6 converts them into an actual strategy for G . Finally, Section 3.5 chooses the parameters so that this strategy wins with probability greater than $\omega^*(G)$, giving the contradiction.

This section mainly follows the conditioning-and-rounding template of Yuen [Yue16, Section 3.2]. The new ingredient is Lemma 3.3, whose proof is deferred to Section 4.

3.1 The contradiction and the conditioning lemma

Fix an actual finite-dimensional strategy for $G^{\otimes n}$. Its shared state is $|\psi\rangle$, its local measurements are $\{E_{x^n}^{a^n}\}_{a^n \in A^n}$ and $\{F_{y^n}^{b^n}\}_{b^n \in B^n}$. Write (X^n, Y^n, A^n, B^n) for the random question-and-answer word generated by this strategy; its distribution is

$$\mathcal{P}(x^n, y^n, a^n, b^n) = \mu^{\otimes n}(x^n, y^n) \langle \psi | E_{x^n}^{a^n} \otimes F_{y^n}^{b^n} | \psi \rangle.$$

Write W_i for the event that coordinate i is won and set

$$W_D = \bigcap_{j \in D} W_j, \quad \vartheta = \mathcal{P}(W_{[n]}).$$

The letter ϑ always denotes the all-coordinate success probability; the letter r below denotes a classical history.

Lemma 3.1 (Quantitative greedy conditioning). *Suppose $\vartheta > 0$, let $0 < \delta < 1$, and assume*

$$\frac{\log(1/\vartheta)}{\delta} < n.$$

There exists $D \subsetneq [n]$ such that

$$\begin{aligned} |D| &\leq \frac{\log(1/\vartheta)}{\delta}, \\ \mathcal{P}(W_D) &\geq \vartheta, \\ \frac{1}{n - |D|} \sum_{i \in [n] \setminus D} \mathcal{P}(W_i \mid W_D) &\geq 1 - \delta. \end{aligned}$$

The proof appears in Appendix A.2. Recall the conditioning parameters from the proof overview: p is the probability of winning every coordinate in D , q is the average conditional winning probability on a remaining coordinate, and η is the information cost per remaining coordinate. For the set D supplied by Lemma 3.1, these are

$$\begin{aligned} m &= n - |D|, & p &= \mathcal{P}(W_D), \\ q &= \frac{1}{m} \sum_{i \in [n] \setminus D} \mathcal{P}(W_i \mid W_D), & \eta &= \frac{\log(1/p) + |D| \log(|A||B|)}{m}. \end{aligned} \tag{1}$$

We will explicitly choose the parameters in Section 3.5, where the bounds on $|D|$ and p from Lemma 3.1 will give $0 \leq \eta \leq 1$; see (17). We will also define a universal rounding constant $B_{\text{qs}} \geq 1$ in (11). With these choices, the operational goal is the following implication:

$$0 \leq \eta \leq 1 \implies \omega^*(G) \geq q - B_{\text{qs}} \eta^{1/12}.$$

Lemma 3.6 proves this implication by constructing a genuine single-game strategy. Its left-hand side is computed in the conditioned repeated-game experiment; its right-hand side concerns an actual single-game strategy receiving fresh questions from μ . Thus it suffices to make q close to one and η small; the probability p itself need not approach one.

3.2 The ideal conditioned single-game strategy

Fix $D \subsetneq [n]$ with $p = \mathcal{P}(W_D) > 0$. To extract one instance of G , we keep one coordinate $i \notin D$ *live*: its questions X_i, Y_i are withheld from the history and will serve as the referee's single-game questions. The history reveals both core questions X_D, Y_D , so the core winning event W_D can be checked once the core answers are recorded. At every other coordinate it reveals at least one of the two questions. This breaks the correlation between the remaining unrevealed Alice and Bob question strings, while randomized reveal orders later make the live question a uniformly chosen martingale increment. Section 4 constructs those orders and their distribution explicitly.

To help construct their single-game strategy, Alice and Bob use finite-valued public randomness $\lambda \in \Lambda$, independent of the game questions, to sample a uniformly random live coordinate $i \in [n] \setminus D$ and revealed-coordinate sets $C_X, C_Y \subseteq [n]$ for their respective questions. These sets satisfy

$$D \subseteq C_X \cap C_Y, \quad C_X \cup C_Y = [n] \setminus \{i\}. \quad (2)$$

The first condition records both questions on D ; the second withholds both live questions X_i, Y_i and reveals at least one question at every other coordinate. In particular, this second condition is what makes the question factorization below possible.

The conditioned history. Let Z denote the random pair of answer words produced on D , taking values in $A^D \times B^D$. Define the revealed-question history and full history by

$$T = (\lambda, X_{C_X}, Y_{C_Y}), \quad Z \in A^D \times B^D, \quad R = (T, Z).$$

Because T contains λ , conditioning on $T = t$ automatically fixes both reveal sets C_X, C_Y . Continue to denote the augmented distribution on $\Lambda \times X^n \times Y^n \times A^n \times B^n$ by $\mathcal{P}$, and write

$$\mathcal{Q} = \mathcal{P}(\cdot \mid W_D).$$

The history R is an analytical variable, not information supplied to either player. Since R contains λ , a realized history r determines its live coordinate i ; we nevertheless display i explicitly to identify the extracted coordinate. Define the posterior tuple distribution by

$$\mathcal{Q}(i, r, x, y) = \mathcal{P}(R = r, X_i = x, Y_i = y \mid W_D).$$

Here (i, r, x, y) denotes realized values, while R, X_i, Y_i denote random variables; the four tuple entries need not be independent.

Why the unrevealed questions factor. We now see why the coverage property in (2) matters. Once (T, X_i, Y_i) is fixed, the still-random Alice questions occur only on $[n] \setminus (C_X \cup \{i\})$, while the still-random Bob questions occur only on $[n] \setminus (C_Y \cup \{i\})$. Their intersection is empty precisely because $C_X \cup C_Y = [n] \setminus \{i\}$. Therefore no unrevealed pair $(X_j, Y_j) \sim \mu$ remains jointly

sampled. Independence of the original question pairs across coordinates then gives the prior factorization

$$\begin{aligned}\mathcal{P}(X^n = x^n, Y^n = y^n \mid T = t, X_i = x, Y_i = y) \\ = \mathcal{P}(X^n = x^n \mid T = t, X_i = x) \mathcal{P}(Y^n = y^n \mid T = t, Y_i = y).\end{aligned}\tag{3}$$

This factorization holds under the prior $\mathcal{P}$, not under the conditioned distribution $\mathcal{Q}$.

Local effects and the ideal conditioned state. Fix a realized history $r = (t, z)$, where t is a realized value of T and $z = (a_D, b_D) \in A^D \times B^D$ is a realized value of Z . The marginal repeated-game POVM effects for these recorded answers are

$$\begin{aligned}E_{x^n}^{a_D} &= \sum_{\substack{\tilde{a}^n \in A^n \\ \tilde{a}_D = a_D}} E_{x^n}^{\tilde{a}^n}, \\ F_{y^n}^{b_D} &= \sum_{\substack{\tilde{b}^n \in B^n \\ \tilde{b}_D = b_D}} F_{y^n}^{\tilde{b}^n}.\end{aligned}$$

Thus $E_{x^n}^{a_D}$ is the effect of Alice producing the particular answer word a_D on D , regardless of her answers elsewhere; $F_{y^n}^{b_D}$ has the analogous meaning for Bob. Average these effects over each player's unrevealed questions to obtain the effective local effects

$$\begin{aligned}H_{r,x} &= \mathbb{E}[E_{X^n}^{a_D} \mid T = t, X_i = x], \\ K_{r,y} &= \mathbb{E}[F_{Y^n}^{b_D} \mid T = t, Y_i = y].\end{aligned}$$

Thus $H_{r,x}$ and $K_{r,y}$ are the local effective effects for the fixed answer words a_D, b_D . The question factorization (3) gives the exact probability of observing those answer words:

$$\begin{aligned}p_r(x, y) &:= \mathcal{P}(Z = z \mid T = t, X_i = x, Y_i = y) \\ &= \langle \psi \mid H_{r,x} \otimes K_{r,y} \mid \psi \rangle.\end{aligned}\tag{4}$$

A POVM effect specifies an outcome probability but does not uniquely determine the operator representing its conditional state-update branch. To specify the branch, choose one common cross operator, or purification,

$$\Gamma(F) : \mathcal{H} \longrightarrow \mathcal{H} \otimes \mathcal{A}, \quad 0 \preceq F \preceq I,$$

satisfying

$$\Gamma(F)^\dagger \Gamma(F) = F.\tag{5}$$

Equation (5) says that the cross operator preserves the exact Born probability of its effect. Definition 4.2 constructs a common finite-dimensional resolvent purification, and Lemma 4.3 supplies its operator-entropy estimate.

The resulting conditioned branch and its normalized state are

$$\phi_{r,x,y} = (\Gamma(H_{r,x}) \otimes \Gamma(K_{r,y})) \mid \psi \rangle, \quad \Psi_{r,x,y} = \frac{\phi_{r,x,y}}{\|\phi_{r,x,y}\|}.$$

By (5) and (4),

$$\|\phi_{r,x,y}\|^2 = p_r(x, y).$$

Therefore $\Psi_{r,x,y}$ is well defined on every branch of positive $\mathcal{Q}$ -probability. This state is only *ideal*: its definition uses the full history r and both live questions (x, y) , whereas Alice and Bob separately know only their own questions. Lemma 3.3 will replace it by nearby states described using separate local information.

There are two separate questions about the ideal branch. First, if Alice and Bob were handed $\Psi_{r,x,y}$, could they reproduce the original strategy's answers at coordinate i ? Second, can they approximately obtain this state from their own questions? The next lemma answers only the first question. The substantially harder Lemma 3.3 answers the second.

Lemma 3.2 (The ideal state exactly simulates the remaining coordinate). *Fix (i, r, x, y) such that $\mathcal{Q}(i, r, x, y) > 0$. If Alice and Bob share the ideal state $\Psi_{r,x,y}$, there are local POVMs $\{\tilde{M}_{r,x}^a\}_{a \in A}$ and $\{\tilde{N}_{r,y}^b\}_{b \in B}$, determined by (r, x) and (r, y) , respectively, such that*

$$\langle \Psi_{r,x,y} | \tilde{M}_{r,x}^a \otimes \tilde{N}_{r,y}^b | \Psi_{r,x,y} \rangle = \mathcal{Q}(A_i = a, B_i = b \mid R = r, X_i = x, Y_i = y).$$

The proof and the corresponding answer-refinement operators appear in Appendix A.1.

Why the ideal state wins with probability q . By Lemma 3.2, measuring $\Psi_{r,x,y}$ produces exactly the original strategy's answers at coordinate i , conditioned on the history and live questions. Its winning probability on this branch is therefore

$$\mathcal{Q}(W_i \mid R = r, X_i = x, Y_i = y).$$

Average over $(i, R, X_i, Y_i) \sim \mathcal{Q}$. Since i is uniform in $[n] \setminus D$ and independent of the original experiment, the law of total expectation gives

$$\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}}[\mathcal{Q}(W_i \mid R = r, X_i = x, Y_i = y)] = \frac{1}{m} \sum_{j \in [n] \setminus D} \mathcal{P}(W_j \mid W_D) = q.$$

This is an ideal *analytical* success probability: it does not assume that the players can sample $\mathcal{Q}$, obtain the history, or prepare the question-dependent state.

Sampling histories from separate questions. The ideal experiment is not yet a legal single-game strategy: under $\mathcal{Q}$, its questions may be biased by conditioning, its history is not given to either player, and $\Psi_{r,x,y}$ depends jointly on both questions. Instead, consider the following actual experiment.

The referee draws $(x, y) \sim \mu$, and the players use shared randomness to choose a uniform $i \in [n] \setminus D$. Alice sees x but not y , so she samples a history from the conditional distribution of R given her own question:

$$r_A \sim \mathcal{Q}(R \mid i, X_i = x).$$

Similarly, Bob sees y but not x , and samples

$$r_B \sim \mathcal{Q}(R \mid i, Y_i = y).$$

These histories need not agree automatically. They will be coupled by classical correlated sampling using the players' shared randomness.

Let $\mathcal{J}_A$ denote the distribution generated when Alice samples a history using her own question, and let $\mathcal{J}_B$ denote the analogous distribution generated by Bob. Their respective tuple distributions are

$$\begin{aligned} \mathcal{J}_A(i, r, x, y) &= \frac{1}{m} \mu(x, y) \mathcal{Q}(r \mid i, X_i = x), \\ \mathcal{J}_B(i, r, x, y) &= \frac{1}{m} \mu(x, y) \mathcal{Q}(r \mid i, Y_i = y). \end{aligned} \tag{6}$$

Both retain the original question distribution μ .

Locally describable candidate states. Even when $r_A = r_B = r$, the target $\Psi_{r,x,y}$ still depends on both questions. To obtain states that Alice and Bob can describe separately, average the corresponding effects over the other player's unknown question:

$$\bar{H}_{r,y} = \sum_{x'} \mu(x' | y) H_{r,x'}, \quad \bar{K}_{r,x} = \sum_{y'} \mu(y' | x) K_{r,y'}.$$

Here $\bar{H}_{r,y}$ is Bob's estimate of Alice's effect, using his question y , while $\bar{K}_{r,x}$ is Alice's estimate of Bob's effect, using her question x . Collect the exact branch and its two locally describable alternatives:

$$\begin{aligned} \phi_{r,x,y} &= (\Gamma(H_{r,x}) \otimes \Gamma(K_{r,y})) |\psi\rangle, \\ \phi_{r,x}^A &= (\Gamma(H_{r,x}) \otimes \Gamma(\bar{K}_{r,x})) |\psi\rangle, \\ \phi_{r,y}^B &= (\Gamma(\bar{H}_{r,y}) \otimes \Gamma(K_{r,y})) |\psi\rangle. \end{aligned} \tag{7}$$

Let

$$\Psi_{r,x}^A = \frac{\phi_{r,x}^A}{\|\phi_{r,x}^A\|}, \quad \Psi_{r,y}^B = \frac{\phi_{r,y}^B}{\|\phi_{r,y}^B\|}.$$

Alice can specify the entire bipartite state $\Psi_{r,x}^A$ from (r, x) , while Bob can specify $\Psi_{r,y}^B$ from (r, y) .

The attempt to mimic the ideal experiment therefore requires two guarantees: the locally generated histories must agree and remain close in distribution to $\mathcal{Q}$, and both locally described states must approximate $\Psi_{r,x,y}$. The next lemma provides exactly these guarantees.

3.3 The central state-alignment and history-sampleability lemma

The following lemma contains the new technical estimates. Its first conclusion says that the ideal conditioned state admits two nearby descriptions, one available to each player. Its second says that both players can approximately sample the same conditioned history despite receiving fresh questions from μ . The complete proof is deferred to Section 4.

Recall that $(i, r, x, y) \sim \mathcal{Q}$ denotes a tuple drawn from the conditioned distribution; although r already determines i , the coordinate is displayed explicitly.

Lemma 3.3 (Postselected state alignment and history sampleability). *For every finite repeated-game strategy and every $D \subsetneq [n]$ with $p = \mathcal{P}(W_D) > 0$, use the parameters, history, states, and locally generated distributions defined above. Then*

$$\begin{aligned} \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,x}^A \right\|^2 &\leq 8\eta, \\ \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,y}^B \right\|^2 &\leq 8\eta, \end{aligned} \tag{8}$$

Moreover, with

$$\kappa = \sqrt{\frac{3}{2}\eta},$$

one has $d_{\text{TV}}(\mathcal{Q}, \mathcal{J}_A), d_{\text{TV}}(\mathcal{Q}, \mathcal{J}_B) \leq \kappa$. There exists finite shared randomness such that, on $(x, y) \sim \mu$, Alice and Bob sample a common uniform $i \in [n] \setminus D$ and histories r_A, r_B with respective tuple distributions $\mathcal{J}_A, \mathcal{J}_B$, satisfying

$$\mathbb{P}(r_A \neq r_B) \leq 4\kappa. \tag{9}$$

The state inequalities are averages under the posterior distribution: for example, the first states that when the repeated experiment is conditioned on W_D , the ideal state is typically close to the state described by Alice. They are not pointwise claims about every history. The total-variation bounds compare the ideal posterior with histories that Alice or Bob can actually generate using their own question. Finite classical correlated sampling then gives (9).

3.4 Constructing the actual single-game strategy

On fresh referee questions $(x, y) \sim \mu$, the desired strategy has three operational steps. First, Alice and Bob use shared randomness to produce approximately matching histories r_A, r_B . Second, they use their local descriptions $\Psi_{r_A, x}^A$ and $\Psi_{r_B, y}^B$ to approximately prepare the ideal state $\Psi_{r, x, y}$. Finally, they apply the local coordinate measurements from Lemma 3.2. We first state the two standard sampling tools and then give the full strategy and its error analysis.

Classical sampling is the familiar Holenstein construction [Hol09, Lemma 5.2 and Corollary 5.3]; a direct proof using finite shared randomness appears in Appendix A.7. The quantum input is the state-preparation theorem of Dinur, Steurer, and Vidick [DSV15, Lemma 17], which applies to arbitrary bipartite states.

Lemma 3.4 (Exact finite classical correlated sampling). *Let $\mathcal{F}$ be a finite family of probability distributions on a finite set $\mathcal{R}$. There exists one finite shared random variable and, for each $\mathcal{D} \in \mathcal{F}$, an output $R_{\mathcal{D}}$ having exactly distribution $\mathcal{D}$, such that*

$$\mathbb{P}(R_{\mathcal{D}} \neq R_{\mathcal{E}}) \leq 2 \, \text{d}_{\text{TV}}(\mathcal{D}, \mathcal{E}) \quad (\mathcal{D}, \mathcal{E} \in \mathcal{F}). \quad (10)$$

Apply the lemma to the finite family

$$\{\mathcal{Q}(R \mid i, X_i = x), \quad \mathcal{Q}(R \mid i, Y_i = y)\}_{i, x, y}.$$

After sharing a uniform i , Alice uses her actual question x and Bob his actual question y . Averaging (10) against $\mu(x, y)$ gives

$$\begin{aligned} \mathbb{P}(r_A \neq r_B) &\leq 2 \, \text{d}_{\text{TV}}(\mathcal{J}_A, \mathcal{J}_B) \\ &\leq 2(\text{d}_{\text{TV}}(\mathcal{J}_A, \mathcal{Q}) + \text{d}_{\text{TV}}(\mathcal{Q}, \mathcal{J}_B)) \leq 4\kappa, \end{aligned}$$

which proves the finite-shared-randomness guarantee in Lemma 3.3.

Lemma 3.5 (Quantum correlated sampling). *There exists a universal $K_{\text{qs}} \geq 1$ with the following property. Given $d \geq 1$ and $0 < \alpha \leq 1$, there exist a finite $d' \geq 1$ and assignments of local unitaries*

$$\sigma_A \mapsto U(\sigma_A), \quad \sigma_B \mapsto V(\sigma_B)$$

on $\mathbb{C}^{dd'}$, indexed by unit vectors in $\mathbb{C}^d \otimes \mathbb{C}^{d'}$, such that, simultaneously for every pair σ_A, σ_B of such unit vectors,

$$\|(U(\sigma_A) \otimes V(\sigma_B)) |E_{dd'}\rangle - |\sigma_A\rangle |E_{d'}\rangle\| \leq K_{\text{qs}} \max\{\alpha^{1/12}, \|\sigma_A - \sigma_B\|^{1/6}\}.$$

Here the finite embezzlement state is

$$|E_b\rangle = \left(\sum_{j=1}^b \frac{1}{j} \right)^{-1/2} \sum_{j=1}^b \frac{1}{\sqrt{j}} |j\rangle |j\rangle.$$

The initial state $|E_{dd'}\rangle$ and catalyst dimension d' depend on d, α , not on σ_A, σ_B .

This is precisely the quantum correlated-sampling lemma of Dinur, Steurer, and Vidick [DSV15, Lemma 17].

Rounding into a legal single-game strategy. Set the following universal constants:

$$U_* = K_{\text{qs}}(1 + 2^{1/6}) + 2, \quad B_{\text{qs}} = (5 + 2U_*)\sqrt{\frac{3}{2}} + 2K_{\text{qs}}32^{1/12} + 2\sqrt{8}. \quad (11)$$

Lemma 3.6 (Postselection-stable finite-dimensional rounding). *In the setting of Lemma 3.3, suppose $0 \leq \eta \leq 1$. For every $0 < \alpha \leq 1$, there is an actual finite-dimensional tensor-product strategy S_α for the unchanged game G . On questions $(x, y) \sim \mu$, the strategy first uses classical correlated sampling to generate approximately matching histories, then uses quantum correlated sampling to prepare the state associated with that history, and finally applies the coordinate measurements of Lemma 3.2. Its winning probability satisfies*

$$\text{win}(S_\alpha) \geq q - B_{\text{qs}}\eta^{1/12} - 2K_{\text{qs}}\alpha^{1/12}. \quad (12)$$

The questions of this strategy have distribution μ , and its shared state is prepared before either question is received. In particular,

$$\omega^*(G) \geq q - B_{\text{qs}}\eta^{1/12}. \quad (13)$$

Proof. First, synchronize the history. Classical correlated sampling provides finite shared randomness h , including a uniform coordinate i , from which Alice computes $r_A = r_A(h, x)$ and Bob computes $r_B = r_B(h, y)$. By Lemma 3.3,

$$d_{\text{TV}}(\mathcal{Q}, \mathcal{J}_A), d_{\text{TV}}(\mathcal{Q}, \mathcal{J}_B) \leq \kappa, \quad \mathbb{P}(r_A \neq r_B) \leq 4\kappa, \quad \kappa \leq \sqrt{\frac{3}{2}\eta}. \quad (14)$$

Next, prepare the ideal state. Pad all candidate states to a common finite dimension d , and apply quantum correlated sampling once, obtaining d' and local unitaries U, V . If $\nu(h)$ denotes the distribution of the classical shared randomness, the question-independent shared state is

$$|\Omega\rangle = \left(\sum_h \sqrt{\nu(h)} |h\rangle_A |h\rangle_B \right) \otimes |E_{dd'}\rangle. \quad (15)$$

For analysis, fix a branch on which $r_A = r_B = r$. Alice knows (r, x) , Bob knows (r, y) , and they apply $U(\Psi_{r,x}^A)$ and $V(\Psi_{r,y}^B)$, respectively. Quantum correlated sampling requires these candidate states to be close. Since both approximate the same ideal state,

$$\begin{aligned} \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x}^A - \Psi_{r,y}^B \right\|^2 &\leq 2\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x}^A - \Psi_{r,x,y} \right\|^2 \\ &\quad + 2\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,y}^B \right\|^2 \leq 32\eta. \end{aligned}$$

Quantum correlated sampling therefore approximately prepares $\Psi_{r,x}^A$, which is itself close to $\Psi_{r,x,y}$. The resulting average state-preparation error is $O(\alpha^{1/12} + \eta^{1/12})$.

Finally, measure the extracted coordinate. Alice and Bob apply the measurements from Lemma 3.2. On the ideal state their average winning probability is q ; closeness of the prepared state changes this probability by $O(\alpha^{1/12} + \eta^{1/12})$. Mismatching histories and the change from $\mathcal{Q}$ to the actual question distribution cost $O(\kappa) = O(\sqrt{\eta})$. The precise calculation in Appendix A.8 gives (12). Taking the supremum over the finite-dimensional strategies S_α and letting $\alpha \rightarrow 0^+$ yields (13). $\square$

3.5 The final contradiction and exponential bound

We now prove Theorem 1.1, including its explicit distribution-uniform rate and the prefactor 1. Set

$$c_{\text{qs}} = \frac{1}{8(4B_{\text{qs}})^{12}} > 0. \quad (16)$$

Proof of Theorem 1.1. First suppose the game satisfies $\omega^*(G) = 0$. From any finite-dimensional strategy for $G^{\otimes n}$, the players can construct a single-game strategy by choosing a coordinate and presampling all other question pairs from μ using finite shared randomness. They insert their

actual referee questions in the chosen coordinate, run the original repeated local measurements, and return that coordinate's answers. Every all-coordinate win is a win of the resulting single-game strategy. Hence

$$\omega^*(G^{\otimes n}) \leq \omega^*(G) = 0 \quad (n \geq 1),$$

and the assertion holds.

Suppose, therefore, that $0 < \varepsilon < 1$, and write

$$\ell = \log(|A||B|), \quad \gamma = c_{\text{qs}} \frac{\varepsilon^{13}}{\varepsilon + \ell}, \quad \delta = \frac{\varepsilon}{4}.$$

Fix an arbitrary $n \geq 1$. If the desired conclusion fails at this n , then

$$\omega^*(G^{\otimes n}) > e^{-\gamma n}.$$

The definition of the supremum therefore gives an actual finite-dimensional repeated strategy with

$$\vartheta = \mathcal{P}(W_{[n]}) > e^{-\gamma n}.$$

There is no assumption that the repeated value is attained. Since $B_{\text{qs}} \geq 1$,

$$\gamma = \frac{\varepsilon}{8(\varepsilon + \ell)} \left(\frac{\varepsilon}{4B_{\text{qs}}} \right)^{12} \leq \frac{\varepsilon}{8} = \frac{\delta}{2}.$$

Thus the hypothesis of Lemma 3.1 is satisfied, and it supplies $D \subsetneq [n]$ with

$$|D| < \frac{\gamma n}{\delta} \leq \frac{n}{2}, \quad m > \frac{n}{2}, \quad p \geq \vartheta, \quad q \geq 1 - \frac{\varepsilon}{4}.$$

In particular,

$$\log(1/p) < \gamma n, \quad |D|\ell < \frac{4\gamma n\ell}{\varepsilon}.$$

Using $m > n/2$ and (16), we get

$$\begin{aligned} \eta &= \frac{\log(1/p) + |D|\ell}{m} \\ &< 2\gamma \left(1 + \frac{4\ell}{\varepsilon} \right) \\ &= \frac{\varepsilon^{12}}{4(4B_{\text{qs}})^{12}} \frac{\varepsilon + 4\ell}{\varepsilon + \ell} \\ &\leq \left(\frac{\varepsilon}{4B_{\text{qs}}} \right)^{12} \leq 1. \end{aligned} \tag{17}$$

Choose the strictly positive finite-catalyst accuracy

$$\alpha = \left(\frac{\varepsilon}{16K_{\text{qs}}} \right)^{12} \in (0, 1].$$

Lemma 3.6 now produces an actual single-game finite-dimensional strategy with

$$\begin{aligned} \text{win}(S_\alpha) &\geq q - B_{\text{qs}}\eta^{1/12} - 2K_{\text{qs}}\alpha^{1/12} \\ &> 1 - \frac{\varepsilon}{4} - \frac{\varepsilon}{4} - \frac{\varepsilon}{8} = 1 - \frac{5\varepsilon}{8} > 1 - \varepsilon = \omega^*(G), \end{aligned}$$

contradicting the definition of $\omega^*(G)$. Therefore $\omega^*(G^{\otimes n}) \leq e^{-\gamma n}$ for this arbitrary n . As $n \geq 1$ was arbitrary,

$$\omega^*(G^{\otimes n}) \leq \exp\left(-\frac{1}{8(4B_{\text{qs}})^{12}} \frac{\varepsilon^{13}}{\varepsilon + \log(|A||B|)} n\right) \quad (n \geq 1).$$

Every question conditional above was used only on positive μ -support. Neither a minimum positive question probability nor a connected-support or component-dependent reduction enters the proof. Both the classical flag and the catalyst in (15) are finite for each fixed strategy and each $\alpha > 0$. The rate is consequently independent of the question distribution, its support, the question-alphabet sizes, and the entanglement dimension. $\square$

4 Proof of the postselected sampleability lemma

This section proves Lemma 3.3.

Reminder of Lemma 3.3 (Postselected state alignment and history sampleability). *For every finite repeated-game strategy and every $D \subsetneq [n]$ with $p = \mathcal{P}(W_D) > 0$, let the histories, branches, and distributions be defined by (1)–(7) and (6). Then*

$$\begin{aligned} \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,x}^A \right\|^2 &\leq 8\eta, \\ \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,y}^B \right\|^2 &\leq 8\eta, \end{aligned}$$

Furthermore, with

$$\kappa = \sqrt{\frac{3}{2}}\eta,$$

one has

$$\text{d}_{\text{TV}}(\mathcal{Q}, \mathcal{J}_A) \leq \kappa, \quad \text{d}_{\text{TV}}(\mathcal{Q}, \mathcal{J}_B) \leq \kappa.$$

There exists a finite shared random variable such that, on $(x, y) \sim \mu$, the players sample a common uniform $i \in [n] \setminus D$ and locally generated histories r_A, r_B with respective tuple marginals $\mathcal{J}_A, \mathcal{J}_B$ and

$$\mathbb{P}(r_A \neq r_B) \leq 4\kappa.$$

Recall notations. Fix the repeated-game strategy and $D \subsetneq [n]$ with $p = \mathcal{P}(W_D) > 0$. Throughout this section write

$$M = [n] \setminus D, \quad m = |M|, \quad \tau = \log(1/p), \quad s = |D| \log(|A||B|), \quad \eta = \frac{\tau + s}{m}.$$

Here W_D is the event that every core coordinate in D is won, and $\mathcal{Q} = \mathcal{P}(\cdot \mid W_D)$ is its posterior.

Proof roadmap. We begin by understanding what the state difference in (8) actually measures. Public randomness λ picks a live coordinate i and specifies which questions have already been revealed. The resulting history is

$$r = (t, z), \quad t = (\lambda, X_{C_X}, Y_{C_Y}), \quad z = (a_D, b_D).$$

Thus r records the core questions X_D, Y_D and answers a_D, b_D , but neither live question X_i, Y_i . Suppose now that $X_i = x$ and $Y_i = y$. The effects $H_{r,x}$ and $K_{r,y}$ describe Alice's and Bob's

recorded core answers a_D, b_D , respectively. Alice knows x , but not y , so from her point of view Bob's effect $K_{r,y}$ must instead be averaged:

$$\bar{K}_{r,x} = \mathbb{E}_{Y_i \sim \mu(\cdot|x)} K_{r,Y_i}.$$

Write $|\psi\rangle$ for the original shared state and Γ for a common purification. The ideal branch conditioned on both live questions (x, y) and the branch determined by Alice's information (r, x) are, respectively,

$$\phi_{r,x,y} = (\Gamma(H_{r,x}) \otimes \Gamma(K_{r,y})) |\psi\rangle, \quad \phi_{r,x}^A = (\Gamma(H_{r,x}) \otimes \Gamma(\bar{K}_{r,x})) |\psi\rangle.$$

The only difference is that the first branch sees Bob's actual question $Y_i = y$, while the second does not. The quantity $\phi_{r,x,y} - \phi_{r,x}^A$ is therefore precisely the change caused by revealing Y_i . Its normalized counterpart is the first state difference in (8).

A single reveal can have a large effect, so there is no reason for revealing Y_i in isolation to have a small cost. Instead, imagine keeping Alice's relevant questions fixed and gradually uncovering Bob's remaining questions, one at a time, in a random order π_X . Somewhere along the way we encounter the live coordinate i . Because the order is random, the number k_X of questions revealed before i is uniform. Let U be everything fixed before this process begins. After the first j questions have been uncovered, our current prediction for Bob's core-answer effect $F_{Y^n}^{b_D}$ is

$$G_j = \mathbb{E} \left[F_{Y^n}^{b_D} \middle| U, Y_{\pi_X^{\leq j}} \right].$$

As more questions are exposed, these predictions form a Doob martingale under the prior $\mathcal{P}$. Immediately before the live step k_X , the question Y_i is still unknown; immediately afterwards, it has been revealed. We therefore have

$$G_{k_X} = \bar{K}_{r,X_i}, \quad G_{k_X+1} = K_{r,Y_i}, \quad \phi_{r,X_i,Y_i} - \phi_{r,X_i}^A = \Gamma(H_{r,X_i}) \otimes (\Gamma(G_{k_X+1}) - \Gamma(G_{k_X})) |\psi\rangle.$$

This is the central reduction: the state difference we care about is exactly one randomly chosen increment of a much longer reveal process. A particular increment might be large, but a uniformly random one can be controlled by the total information budget.

The martingale above tracks effects. To control the resulting change in quantum states, we must also choose the purification Γ . One obvious choice would be $\Gamma(F) = F^{1/2}$, which already preserves Born probabilities, that is, $\Gamma(F)^\dagger \Gamma(F) = F$. Instead, we use the resolvent purification from Definition 4.2, which preserves the same Born probabilities. Its additional benefit is the entropy-control property in Lemma 4.3: with $H_1(v) = -v \log v$ and $H_1(0) = 0$,

$$\mathbb{E}_F \left[(\Gamma(F) - \Gamma(\bar{F}))^\dagger (\Gamma(F) - \Gamma(\bar{F})) \right] \preceq H_1(\bar{F}) - \mathbb{E}_F H_1(F), \quad \bar{F} = \mathbb{E}_F F.$$

The meaning is that the expected squared difference between a random effect F and its average $\bar{F}$, after applying Γ , is bounded by the decrease in operator entropy when $\bar{F}$ is refined to F . Apply this at every step of the reveal martingale $(G_j)_{j=0}^N$. The successive entropy decreases telescope, so their total, weighted by Alice's fixed effect H and the shared state, is at most the entropy of the initial branch probability:

$$\sum_{j=0}^{N-1} \mathbb{E} \left\| \Gamma(H) \otimes (\Gamma(G_{j+1}) - \Gamma(G_j)) |\psi\rangle \right\|^2 \leq H_1(p_0), \quad p_0 = \langle \psi | H \otimes G_0 | \psi \rangle.$$

Lemma 4.4 therefore bounds a uniformly random step k by its $1/N$ share of the total:

$$\mathbb{E} \left\| \Gamma(H) \otimes (\Gamma(G_{k+1}) - \Gamma(G_k)) |\psi\rangle \right\|^2 \leq \frac{H_1(p_0)}{N}.$$

For a fixed core-answer word z , the initial branch probability is $p_0 = p_z(U) = \mathcal{P}(Z = z \mid U)$. Lemma 4.5 shows that the total entropy of the accepted words is at most $p(\tau + s)$ on average. Intuitively, acceptance has total probability p , and recording its core-answer word contributes at most s additional units of entropy.

Lemma 4.6 applies the random-step estimate to the two live-question reveals. Write $\mathcal{P}^0$ for the prior distribution of public randomness and questions. The Alice- and Bob-reveal costs are

$$I_A = \mathbb{E}_{(i,t,x,y) \sim \mathcal{P}^0} \sum_{z: W_D(t,z)} \left\| \phi_{r,x,y} - \phi_{r,y}^B \right\|^2,$$

$$I_B = \mathbb{E}_{(i,t,x,y) \sim \mathcal{P}^0} \sum_{z: W_D(t,z)} \left\| \phi_{r,x,y} - \phi_{r,x}^A \right\|^2.$$

Here $r = (t, z)$, and the sums range over accepted core-answer words. The vectors $\phi_{r,x,y}, \phi_{r,x}^A, \phi_{r,y}^B$ are unnormalized branches, so I_A and I_B measure their unnormalized squared differences. The reverse description selects a block of N candidate live coordinates with size bias $2N/m$, while its uniform reveal step contributes $1/N$. These factors cancel, so Lemma 4.6 bounds both costs by

$$I_A, I_B \leq \frac{2}{m} p(\tau + s).$$

Lemma 4.7 then converts these unnormalized branch differences into the normalized-state differences that we actually want to bound:

$$\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,x}^A \right\|^2 \leq \frac{4}{p} I_B \leq 8\eta,$$

$$\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,y}^B \right\|^2 \leq \frac{4}{p} I_A \leq 8\eta.$$

This proves both desired state bounds in (8) without any inverse-postselection loss. What remains, the history-matching guarantee (9), is purely classical: Lemma 4.8 supplies the standard conditioning budget, and Lemma 4.9 combines it with the relative-entropy chain rule, Pinsker's inequality, and classical correlated sampling.

4.1 Setting up the revealed history and its martingale

The forward history. Choose i uniformly in M and partition $M \setminus \{i\}$ fairly into $L_X \sqcup L_Y$. Choose independent uniform random permutations $\pi_{X,-i}$ of L_X and $\pi_{Y,-i}$ of L_Y . The subscript $-i$ indicates that these permutations exclude the live coordinate i . Independently choose uniform cuts

$$k_X \in \{0, \dots, |L_X|\}, \quad k_Y \in \{0, \dots, |L_Y|\}.$$

This is a symmetric random-reveal variant of the dependency-breaking construction in Holenstein's classical embedding argument [Hol09, Section 3]. The random permutations and cuts k_X, k_Y allow the live question to be identified later with a uniformly random increment of a reveal martingale.

The public randomness and revealed-coordinate sets are

$$\lambda = (i, L_X, L_Y, \pi_{X,-i}, \pi_{Y,-i}, k_X, k_Y),$$

$$C_X = D \cup L_X \cup \pi_{Y,-i}^{\leq k_Y},$$

$$C_Y = D \cup L_Y \cup \pi_{X,-i}^{\leq k_X}.$$

Thus Alice's questions X_{L_X} and $X_{\pi_{Y,-i}^{\leq k_Y}}$ are revealed, as are Bob's questions Y_{L_Y} and $Y_{\pi_{X,-i}^{\leq k_X}}$. This is precisely the public randomness λ introduced in Section 3.2; in particular,

$$T = (\lambda, X_{C_X}, Y_{C_Y}), \quad C_X \cup C_Y = [n] \setminus \{i\}.$$

Write $\mathcal{P}^0$ for the marginal retaining the independent public randomness and questions but discarding all answers: $\mathcal{P}^0(\lambda, x^n, y^n) = \mathbb{P}(\lambda) \prod_{j=1}^n \mu(x_j, y_j)$. Its induced marginal on (i, T, X_i, Y_i) is denoted by $\mathcal{P}^0(i, t, x, y)$. Both core-question tuples X_D, Y_D are revealed, neither live question X_i, Y_i is revealed, and every other coordinate reveals at least one question. Therefore the unrevealed Alice and Bob index sets are disjoint, and the product prior $\mathcal{P}$ gives exactly (3). This is independence under $\mathcal{P}$, not under the posterior $\mathcal{Q}$.

The branches to be compared. For $z = (a_D, b_D)$ and $r = (t, z)$, recall

$$\begin{aligned} H_{r,x} &= \mathbb{E}[E_{X^n}^{a_D} \mid T = t, X_i = x], \\ K_{r,y} &= \mathbb{E}[F_{Y^n}^{b_D} \mid T = t, Y_i = y]. \end{aligned}$$

Prior conditional independence identifies the joint branch probability $p_r(x, y)$ with the notation already introduced in Section 3:

$$p_r(x, y) = \langle \psi \mid H_{r,x} \otimes K_{r,y} \mid \psi \rangle = \|\phi_{r,x,y}\|^2.$$

Averaging over the other player's unknown live question gives

$$\bar{H}_{r,y} = \sum_{x'} \mu(x' \mid y) H_{r,x'}, \quad \bar{K}_{r,x} = \sum_{y'} \mu(y' \mid x) K_{r,y'}.$$

Consequently,

$$\begin{aligned} \phi_{r,x,y} &= (\Gamma(H_{r,x}) \otimes \Gamma(K_{r,y})) \mid \psi \rangle, \\ \phi_{r,x}^A &= (\Gamma(H_{r,x}) \otimes \Gamma(\bar{K}_{r,x})) \mid \psi \rangle, \\ \phi_{r,y}^B &= (\Gamma(\bar{H}_{r,y}) \otimes \Gamma(K_{r,y})) \mid \psi \rangle. \end{aligned}$$

If $\mathcal{Q}(i, t, z, x, y) > 0$, then $\mu(x, y) > 0$ and $p_r(x, y) > 0$. Moreover,

$$\bar{H}_{r,y} \succeq \mu(x \mid y) H_{r,x}, \quad \bar{K}_{r,x} \succeq \mu(y \mid x) K_{r,y},$$

so

$$\|\phi_{r,x}^A\|^2 \geq \mu(y \mid x) p_r(x, y) > 0, \quad \|\phi_{r,y}^B\|^2 \geq \mu(x \mid y) p_r(x, y) > 0.$$

Thus all normalized states used on positive-posterior branches (i, t, z, x, y) exist. Since t records the core questions X_D, Y_D , the pair (t, z) determines whether W_D occurs, that is, whether every core coordinate in D is won. Summing the accepted branch probabilities therefore gives

$$\sum_{i,t,x,y,z} \mathcal{P}^0(i, t, x, y) \mathbf{1}_{W_D}(t, z) p_r(x, y) = p.$$

This exact branch mass p is the factor that will cancel when the unnormalized costs are converted into normalized-state distances.

The random live-coordinate martingale. Only three features of the reveal construction enter the proof: the live question is a uniform martingale step, its block is sampled with the appropriate size bias, and the neighboring effects are exactly the local effects we wish to compare.

Lemma 4.1 (Random live-coordinate martingales). *The forward history admits the following equivalent reverse descriptions.*

For the Alice-reveal direction, sample a partition $M = L_X \sqcup L_Y^+$ with

$$\mathbb{P}(L_X, L_Y^+) = 2^{-m} \frac{2|L_Y^+|}{m}.$$

A uniform order π_Y of L_Y^+ and an independent uniform cut $k_Y \in \{0, \dots, |L_Y^+| - 1\}$ determine $i = \pi_Y[k_Y + 1]$. With

$$U = (X_D, Y_D, X_{L_X}, Y_{L_Y}, Y_{\pi_{X,-i}^{\leq k_X}}), \quad (21)$$

the effects

$$\begin{aligned} F_j &= \mathbb{E} \left[E_{X^n}^{a_D} \mid U, X_{\pi_Y^{\leq j}} \right], & 0 \leq j \leq |L_Y^+|, \\ K &= \mathbb{E} \left[F_{Y^n}^{b_D} \mid U \right] \end{aligned} \quad (22)$$

form a prior martingale with fixed Bob effect K , and

$$F_{k_Y} = \bar{H}_{r,Y_i}, \quad F_{k_Y+1} = H_{r,X_i}, \quad K = K_{r,Y_i}. \quad (23)$$

For the Bob-reveal direction, sample $M = L_X^+ \sqcup L_Y$ with probability $2^{-m} 2^{|L_X^+|}/m$, take a uniform order π_X of L_X^+ , and let $i = \pi_X[k_X + 1]$ for a uniform cut k_X . With

$$U = (X_D, Y_D, X_{L_X^+}, Y_{L_Y}, X_{\pi_{Y,-i}^{\leq k_Y}}),$$

the effects

$$G_j = \mathbb{E} \left[F_{Y^n}^{b_D} \mid U, Y_{\pi_X^{\leq j}} \right], \quad 0 \leq j \leq |L_X^+|,$$

form a prior martingale with fixed Alice effect H_{r,X_i} , and

$$G_{k_X} = \bar{K}_{r,X_i}, \quad G_{k_X+1} = K_{r,Y_i}.$$

The live increments are therefore exactly the Alice- and Bob-reveal branch differences. Appendix A.3 verifies the reverse distribution and these martingale properties.

4.2 Resolvent purification and one-step entropy control

We now make precise the purification Γ used in the roadmap. The ordinary square-root choice $\Gamma(F) = F^{1/2}$ preserves Born probabilities; the resolvent purification below does the same while also yielding a useful one-step entropy estimate.

Definition 4.2 (Resolvent purification). Let $\mathcal{S}$ be the collection of all possible values of $H_{r,x}, K_{r,y}, \bar{H}_{r,y}, \bar{K}_{r,x}$ and the martingale effects F_j, G_j in Lemma 4.1. The question, answer, and public-randomness alphabets are finite, so $\mathcal{S}$ is finite. Let Σ be the union of the positive eigenvalues of the operators in $\mathcal{S}$. Set

$$g_\sigma(u) = \frac{\sigma}{\sigma + u}, \quad \mathcal{A}_0 = \text{span}\{g_\sigma : \sigma \in \Sigma\} \subseteq L^2((0, \infty), du), \quad \mathcal{A} = \mathcal{A}_0 \oplus \mathbb{C}|0\rangle.$$

The auxiliary space $\mathcal{A}$ is finite dimensional, with $\dim \mathcal{A} \leq |\Sigma| + 1$. For each effect F , define

$$[\Gamma(F)v](u) = F(F + uI)^{-1}v, \quad \Gamma(F) : \mathcal{H} \longrightarrow \mathcal{H} \otimes \mathcal{A}.$$

Verification of the Born-probability property. For every positive eigenvalue σ ,

$$\int_0^\infty \left(\frac{\sigma}{\sigma + u} \right)^2 du = \sigma,$$

so spectral calculus gives

$$\Gamma(F)^\dagger \Gamma(F) = F, \quad (24)$$

including for singular F .

The entropy-control property. Revealing one question replaces an averaged effect $\bar{F}$ by a more informative effect F . The key benefit of the resolvent purification is that the resulting squared movement is bounded by the corresponding decrease in operator entropy.

Lemma 4.3 (One-step purified variation is controlled by entropy). *Let F be a finitely supported random positive contraction taking values in $\mathcal{S}$, and suppose that $\bar{F} = \mathbb{E}_F F$ also belongs to $\mathcal{S}$. Recall*

$$H_1(v) = -v \log v \quad (0 < v \leq 1), \quad H_1(0) = 0,$$

and define $H_1(F)$ by functional calculus. Then

$$\mathbb{E}_F \left[(\Gamma(F) - \Gamma(\bar{F}))^\dagger (\Gamma(F) - \Gamma(\bar{F})) \right] \preceq H_1(\bar{F}) - \mathbb{E}_F H_1(F). \quad (25)$$

For positive definite effects, the entropy gap is the expected operator Bregman divergence introduced by Petz [Pet07]; closely related resolvent-based estimates were studied by Kim [Kim14]. A direct proof, also covering singular effects, appears in Appendix A.4.

4.3 A purified martingale has a small random increment

The one-step estimate becomes useful because its entropy terms telescope over an entire reveal martingale. This is the quantum counterpart of the classical relative-entropy chain rule.

Lemma 4.4 (Purified martingale increments). *Let $F_0, \dots, F_N$ be a finite positive-contraction martingale, where $N \geq 1$, F_0 is deterministic, and*

$$\mathbb{E}[F_{j+1} \mid F_0, \dots, F_j] = F_j.$$

Fix a positive contraction K on Bob's space and a shared unit vector $|\psi\rangle$. Let k be uniform in $\{0, \dots, N-1\}$, independently of the martingale, and write

$$p_0 = \langle \psi \mid F_0 \otimes K \mid \psi \rangle.$$

Recall

$$H_1(v) = -v \log v \quad (0 < v \leq 1), \quad H_1(0) = 0.$$

Then

$$\mathbb{E}_{k, \mathbf{F}} \left\| (\Gamma(F_{k+1}) - \Gamma(F_k)) \otimes \Gamma(K) \mid \psi \right\|^2 \leq \frac{H_1(p_0)}{N}. \quad (26)$$

The same conclusion holds conditionally when previously revealed information fixes F_0 and K .

Proof. Absorb Bob's fixed effect into the positive operator

$$\rho_K = \text{Tr}_B \left[(I \otimes K^{1/2}) \mid \psi \rangle \langle \psi \mid (I \otimes K^{1/2}) \right].$$

For every operator C on Alice's space,

$$\text{Tr}(\rho_K C) = \langle \psi \mid C \otimes K \mid \psi \rangle, \quad \text{Tr} \rho_K = \langle \psi \mid I \otimes K \mid \psi \rangle \leq 1. \quad (27)$$

In particular, $p_0 = \text{Tr}(\rho_K F_0)$.

Write

$$\Delta_j = \Gamma(F_{j+1}) - \Gamma(F_j), \quad \mathcal{F}_j = \sigma(F_0, \dots, F_j).$$

Conditional on $\mathcal{F}_j$, apply Lemma 4.3 with $F = F_{j+1}$ and $\bar{F} = F_j$. The martingale identity $\mathbb{E}[F_{j+1} \mid \mathcal{F}_j] = F_j$ gives

$$\mathbb{E}[\Delta_j^\dagger \Delta_j \mid \mathcal{F}_j] \preceq H_1(F_j) - \mathbb{E}[H_1(F_{j+1}) \mid \mathcal{F}_j].$$

Moreover, (24) and (27) give

$$\|\Delta_j \otimes \Gamma(K) |\psi\rangle\|^2 = \text{Tr}(\rho_K \Delta_j^\dagger \Delta_j).$$

Taking expectations, tracing against ρ_K , and summing over j makes the intermediate entropies cancel:

$$\sum_{j=0}^{N-1} \mathbb{E} \|\Delta_j \otimes \Gamma(K) |\psi\rangle\|^2 \leq \text{Tr}[\rho_K (H_1(F_0) - \mathbb{E} H_1(F_N))] \leq \text{Tr}(\rho_K H_1(F_0)). \quad (28)$$

For the last inequality use $H_1(F_N) \succeq 0$.

Diagonalize $F_0 = \sum_a \sigma_a |v_a\rangle \langle v_a|$ and put $w_a = \langle v_a | \rho_K | v_a \rangle$. Since $\sum_a w_a = \text{Tr} \rho_K \leq 1$, append an eigenvalue 0 with weight $1 - \text{Tr} \rho_K$. Concavity of H_1 then gives

$$\text{Tr}(\rho_K H_1(F_0)) = \sum_a w_a H_1(\sigma_a) \leq H_1\left(\sum_a w_a \sigma_a\right) = H_1(p_0).$$

Finally, the uniform cut k selects one of the N increments independently, so its expected squared movement is the sum in (28) divided by N . This proves (26). $\square$

4.4 From reveal costs to normalized state alignment

Our goal is the pair of normalized-state estimates in (8). On every positive-posterior branch, the exact posterior weight is

$$\mathcal{Q}(i, t, z, x, y) = \frac{\mathcal{P}^0(i, t, x, y) \mathbf{1}_{W_D}(t, z) \|\phi_{r,x,y}\|^2}{p}. \quad (29)$$

For nonzero vectors u, v ,

$$\left\| \frac{u}{\|u\|} - \frac{v}{\|v\|} \right\|^2 \leq \frac{4\|u - v\|^2}{\|u\|^2}.$$

Taking $u = \phi_{r,x,y}$ and $v = \phi_{r,y}^B$, the factor $\|u\|^2$ in (29) cancels the denominator above. Thus the desired normalized estimate reduces to the following unnormalized Alice-reveal cost:

$$\begin{aligned} I_A &= \mathbb{E}_{(i,t,x,y) \sim \mathcal{P}^0} \sum_{\substack{z: \\ \mathbf{1}_{W_D}(t,z)=1}} \left\| (\Gamma(H_{r,x}) - \Gamma(\bar{H}_{r,y})) \otimes \Gamma(K_{r,y}) |\psi\rangle \right\|^2 \\ &= \mathbb{E}_{(i,t,x,y) \sim \mathcal{P}^0} \sum_{\substack{z: \\ \mathbf{1}_{W_D}(t,z)=1}} \left\| \phi_{r,x,y} - \phi_{r,y}^B \right\|^2. \end{aligned}$$

It is an Alice-reveal cost because the operator that changes is Alice's effect $H_{r,x}$; it compares the ideal branch $\phi_{r,x,y}$ with Bob's surrogate $\phi_{r,y}^B$. It is unnormalized because neither branch vector is divided by its norm and the expectation is under the prior $\mathcal{P}^0$. Define the symmetric Bob-reveal cost by

$$I_B = \mathbb{E}_{(i,t,x,y) \sim \mathcal{P}^0} \sum_{\substack{z: \\ \mathbf{1}_{W_D}(t,z)=1}} \left\| \phi_{r,x,y} - \phi_{r,x}^A \right\|^2.$$

The two target estimates are therefore bounded by

$$\begin{aligned} \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,y}^B \right\|^2 &\leq \frac{4}{p} I_A, \\ \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,x}^A \right\|^2 &\leq \frac{4}{p} I_B. \end{aligned} \quad (30)$$

It remains to show that each unnormalized cost is at most $2p(\tau + s)/m$.

Lemma 4.5 (Accepted-word entropy budget). *For any background information U containing X_D, Y_D , let*

$$p_W(U) = \mathcal{P}(W_D | U), \quad p_z(U) = \mathcal{P}(Z = z | U).$$

The accepted answer words z satisfy

$$\sum_{\mathbf{1}_{W_D}(\vec{U}, z)=1} H_1(p_z(U)) \leq H_1(p_W(U)) + p_W(U)s, \quad (31)$$

and consequently

$$\mathbb{E}_{U \sim \mathcal{P}^0} \sum_{\mathbf{1}_{W_D}(\vec{U}, z)=1} H_1(p_z(U)) \leq p(\tau + s). \quad (32)$$

The proof is deferred to Appendix A.5.

Lemma 4.6 (Probability-weighted quantum alignment). *The two reveal costs satisfy*

$$I_A, I_B \leq \frac{2p(\tau + s)}{m}.$$

Proof. Use the Alice-reveal martingale already constructed in (22). Conditional on its background U , the accepted answer word z has prior probability

$$p_z(U) = \mathcal{P}(Z = z | U) = \langle \psi | F_0 \otimes K | \psi \rangle.$$

Its live cut k_Y is uniform in $\{0, \dots, |L_Y^+| - 1\}$, so Lemma 4.4, applied with $N = |L_Y^+|$ and $k = k_Y$, gives

$$\mathbb{E} \left[\left\| (\Gamma(F_{k_Y+1}) - \Gamma(F_{k_Y})) \otimes \Gamma(K) | \psi \right\|^2 \middle| U \right] \leq \frac{H_1(p_z(U))}{|L_Y^+|}.$$

By (23), the vector on the left is exactly $\phi_{r, X_i, Y_i} - \phi_{r, Y_i}^B$. Fix a partition $M = L_X \sqcup L_Y^+$, and write $\mathbb{E}_U$ for averaging over its remaining public randomness and background questions. Because U contains X_D, Y_D , accepted words are already determined before the random cut. Summing the preceding bound over those words gives

$$\mathbb{E}_U \sum_{z: W_D(U, z)} \frac{H_1(p_z(U))}{|L_Y^+|} \leq \frac{p(\tau + s)}{|L_Y^+|},$$

where (32) applies for each fixed partition because the public randomness is independent of the game. Finally, average over the size-biased partitions:

$$\begin{aligned} I_A &\leq \sum_{\substack{L_X \sqcup L_Y^+ = M \\ |L_Y^+| > 0}} \underbrace{2^{-m} \frac{2|L_Y^+|}{m}}_{\text{probability of the partition}} \frac{p(\tau + s)}{|L_Y^+|} \\ &= \frac{2p(\tau + s)}{m} \sum_{\substack{L_X \sqcup L_Y^+ = M \\ |L_Y^+| > 0}} 2^{-m} \leq \frac{2p(\tau + s)}{m}. \end{aligned}$$

The last sum is at most one because 2^{-m} is the probability of a fair partition and only the nonempty L_Y^+ blocks occur.

For I_B , use the symmetric Bob-reveal martingale $(G_j)_j$ constructed above. Alice's effect H_{r, X_i} remains fixed while Bob's questions $Y_{\pi_X^{\leq j}}$ are revealed, and its partition has size bias $2|L_X^+|/m$. The same martingale estimate and size-bias cancellation give $I_B \leq 2p(\tau + s)/m$. $\square$

Passing to normalized states. We now convert the unnormalized reveal-cost bounds into the normalized-state alignment required by Lemma 3.3.

Lemma 4.7 (Postselection-stable normalized alignment). *The ideal conditioned state and its locally describable alternatives satisfy*

$$\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,x}^A \right\|^2 \leq 8\eta, \quad \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x,y} - \Psi_{r,y}^B \right\|^2 \leq 8\eta.$$

Proof. On tuples of zero posterior probability, any undefined normalized state may be completed by a fixed arbitrary unit vector; such defaults do not contribute to a posterior expectation. Combining the reduction (30) with Lemma 4.6 gives both bounds:

$$\frac{4}{p}I_A, \frac{4}{p}I_B \leq \frac{8(\tau + s)}{m} = 8\eta.$$

Thus the p retained in the unnormalized entropy budget cancels the $1/p$ in the posterior, with no inverse-postselection loss. $\square$

4.5 Classical history generation and correlated sampling

The remaining task is classical: Alice and Bob must generate nearly matching histories r_A, r_B using their respective live questions X_i, Y_i . We collect the conditioning budget and its sampling consequence together to separate this standard argument from the quantum state alignment.

Lemma 4.8 (Conditioning and answer-history entropy). *Conditioning on W_D has exact relative-entropy cost*

$$D(\mathcal{P}(\cdot \mid W_D) \parallel \mathcal{P}) = \log(1/p) = \tau. \quad (33)$$

Recall $s = |D| \log(|A||B|)$. The answer-word alphabet $\mathcal{Z} = A^D \times B^D$ has $\log |\mathcal{Z}| = s$; for every revealed question tuple U and additional question block Y_C ,

$$D(\mathcal{Q}_{U,Y_C,Z} \parallel \mathcal{P}_{U,Y_C} \otimes \text{Unif}_{\mathcal{Z}}) \leq \tau + s. \quad (34)$$

Proof. The Radon–Nikodym derivative of $\mathcal{Q}$ with respect to $\mathcal{P}$ is $\mathbf{1}_{W_D}/p$, which gives (33). Data processing bounds the relative-entropy cost of every question marginal by τ . Conditionally on those questions, the divergence of Z from the uniform distribution on $\mathcal{Z}$ is at most $\log |\mathcal{Z}| = s$. The chain rule gives (34). $\square$

Lemma 4.9 (Locally generated histories approximate the posterior). *The locally generated history distributions satisfy*

$$D(\mathcal{Q} \parallel \mathcal{J}_A), D(\mathcal{Q} \parallel \mathcal{J}_B) \leq \frac{3\tau + 2s}{m}, \quad d_{\text{TV}}(\mathcal{Q}, \mathcal{J}_A), d_{\text{TV}}(\mathcal{Q}, \mathcal{J}_B) \leq \kappa.$$

Classical correlated sampling produces histories r_A, r_B with respective tuple distributions $\mathcal{J}_A, \mathcal{J}_B$ such that

$$\mathbb{P}(r_A \neq r_B) \leq 4\kappa. \quad (35)$$

The proof uses only the preceding conditioning budget, the classical relative-entropy chain rule, Pinsker’s inequality, and Lemma 3.4; it is deferred to Appendix A.6.

Conclusion. Lemma 4.7 proves the two state inequalities in Lemma 3.3, and Lemma 4.9 gives its total-variation and matching-history conclusions. Together they complete the proof.

A Deferred auxiliary proofs

A.1 Ideal-state extraction and local measurement refinements

For $r = (t, z)$, $z = (a_D, b_D)$, define the answer-refinement operators

$$\begin{aligned} H_{r,x}^a &= \mathbb{E}_{X^n \sim \mathcal{P}(\cdot | T=t, X_i=x)} [E_{X^n}^{a_D, a_i=a}], \\ K_{r,y}^b &= \mathbb{E}_{Y^n \sim \mathcal{P}(\cdot | T=t, Y_i=y)} [F_{Y^n}^{b_D, b_i=b}]. \end{aligned}$$

These are positive contractions satisfying

$$\sum_{a \in A} H_{r,x}^a = H_{r,x}, \quad \sum_{b \in B} K_{r,y}^b = K_{r,y}.$$

Lemma A.1 (Singular-safe local POVM refinement). *Let $0 \preceq F \preceq I$ and let $\{F^a : a \in A\}$ be positive operators satisfying $\sum_a F^a = F$. There is a genuine POVM $\{\mathbf{M}_F^a : a \in A\}$ on the finite purified space such that*

$$\Gamma(F)^\dagger \mathbf{M}_F^a \Gamma(F) = F^a \quad (a \in A). \quad (36)$$

Proof. Use the inverse square root $F^{[-1/2]}$ on $\text{supp}(F)$, extended by zero on $\ker F$, and put

$$U_F = \Gamma(F) F^{[-1/2]}.$$

By (5), U_F is an isometry on $\text{supp}(F)$. Choose a default $a_0 \in A$ and set

$$\mathbf{M}_F^a = U_F F^{[-1/2]} F^a F^{[-1/2]} U_F^\dagger + \mathbf{1}_{\{a=a_0\}} (I - U_F U_F^\dagger).$$

Since $0 \preceq F^a \preceq F$, each F^a is supported on $\text{supp}(F)$. The displayed operators are positive and sum to the identity on the entire finite purified space. Finally $U_F^\dagger \Gamma(F) = F^{1/2}$, and $\Gamma(F)$ has image in $\text{im } U_F$. These identities prove (36), including when F has a nontrivial kernel. $\square$

Proof of Lemma 3.2. Apply Lemma A.1 to $\{H_{r,x}^a\}_{a \in A}$ and $\{K_{r,y}^b\}_{b \in B}$, obtaining local POVMs $\{\widetilde{M}_{r,x}^a\}_{a \in A}$ and $\{\widetilde{N}_{r,y}^b\}_{b \in B}$. The refinement identities and $\|\phi_{r,x,y}\|^2 = p_r(x, y)$ give

$$\langle \Psi_{r,x,y} | \widetilde{M}_{r,x}^a \otimes \widetilde{N}_{r,y}^b | \Psi_{r,x,y} \rangle = \frac{\langle \psi | H_{r,x}^a \otimes K_{r,y}^b | \psi \rangle}{p_r(x, y)}.$$

Prior conditional independence identifies the numerator with the probability of the recorded history answers together with $(A_i, B_i) = (a, b)$; the denominator is the probability of the recorded history answers. Their ratio is therefore $\mathcal{Q}(A_i = a, B_i = b \mid R = r, X_i = x, Y_i = y)$, as claimed. $\square$

A.2 The conditioning lemma recalled and proved

Reminder of Lemma 3.1 (Quantitative greedy conditioning). *Suppose $\vartheta > 0$, let $0 < \delta < 1$, and assume*

$$\frac{\log(1/\vartheta)}{\delta} < n.$$

There exists $D \subsetneq [n]$ such that

$$\begin{aligned} |D| &\leq \frac{\log(1/\vartheta)}{\delta}, \\ \mathcal{P}(W_D) &\geq \vartheta, \\ \frac{1}{n - |D|} \sum_{i \in [n] \setminus D} \mathcal{P}(W_i \mid W_D) &\geq 1 - \delta. \end{aligned}$$

Proof of Lemma 3.1. Start with $D = \emptyset$. If

$$\frac{1}{n - |D|} \sum_{i \notin D} \mathbb{P}(\bar{W}_i \mid W_D) > \delta,$$

choose an $i \notin D$ whose individual conditional failure is larger than δ , and replace D by $D \cup \{i\}$. Every conditional probability is well defined, since

$$W_{[n]} \subseteq W_D, \quad \mathbb{P}(W_D) \geq \vartheta > 0.$$

Each addition strictly decreases the mass by a factor smaller than $1 - \delta$:

$$\mathbb{P}(W_{D \cup \{i\}}) = \mathbb{P}(W_D) \mathbb{P}(W_i \mid W_D) < (1 - \delta) \mathbb{P}(W_D).$$

Consequently, after $k \geq 1$ additions,

$$\vartheta \leq \mathbb{P}(W_D) < (1 - \delta)^k, \quad k < \frac{\log(1/\vartheta)}{-\log(1 - \delta)} \leq \frac{\log(1/\vartheta)}{\delta}.$$

The assumed strict inequality prevents all n coordinates from being added. Therefore the process stops at a proper set. The stopping condition is exactly the asserted average-success inequality; containment of the all-win event gives the claimed mass. If the process makes no additions, the same conclusions hold with $D = \emptyset$. $\square$

A.3 Random live-coordinate reveal martingales

Proof of Lemma 4.1. Consider first the Alice-reveal direction and write $N = |L_Y^+|$. In the forward experiment, choose i uniformly in M , partition the other $m - 1$ coordinates fairly, and choose the two orders and cuts uniformly. The probability of a particular outcome is

$$\frac{2^{1-m}}{m |L_X|! (|L_X| + 1) (N - 1)! N}.$$

In the reverse experiment, first choose $M = L_X \sqcup L_Y^+$ with probability $2^{-m} 2N/m$. Then choose a uniform order π_Y of L_Y^+ , a uniform cut $k_Y \in \{0, \dots, N - 1\}$, and set

$$i = \pi_Y[k_Y + 1], \quad L_Y = L_Y^+ \setminus \{i\}.$$

Deleting i from π_Y recovers $\pi_{Y,-i}$. Finally, choose $\pi_{X,-i}$ and k_X as in the forward experiment. The resulting outcome again has probability

$$\underbrace{2^{-m} \frac{2N}{m}}_{\text{size-biased partition}} \frac{1}{N!} \frac{1}{N} \frac{1}{|L_X|!} \frac{1}{|L_X| + 1} = \frac{2^{1-m}}{m |L_X|! (|L_X| + 1) (N - 1)! N}.$$

The reverse construction therefore preserves the exact history distribution while making the live cut k_Y uniform.

Fix the partition, its orders, k_X , and the background U from (21). Given U , the remaining questions X_j , $j \in L_Y^+$, are independent with distributions $\mu(\cdot \mid Y_j)$. Therefore the tower property gives

$$\mathbb{E}[F_{j+1} \mid U, X_{\pi_Y^{\leq j}}] = F_j,$$

so $(F_j)_{j=0}^N$ is a positive-contraction martingale. Bob's effect K is fixed because all questions on which it depends have already been included in U . At the live cut, averaging over X_i gives $F_{k_Y} = \bar{H}_{r,Y_i}$, revealing X_i gives $F_{k_Y+1} = H_{r,X_i}$, and the fixed effect is $K = K_{r,Y_i}$. This proves (23).

Interchanging Alice and Bob gives the second construction. Its partition has probability $2^{-m} 2|L_X^+|/m$; the uniform cut k_X selects the live question from L_X^+ , while Alice's effect remains fixed. The same conditional-independence and tower arguments give the martingale $(G_j)_j$ and the two stated effects at its live cut. $\square$

A.4 Resolvent purification and its entropy estimate

For positive definite effects, let $f(v) = v \log v$ and $H_1(v) = -v \log v$. The associated operator Bregman divergence is

$$D_f(F, \bar{F}) = f(F) - f(\bar{F}) - Df(\bar{F})[F - \bar{F}].$$

When $\bar{F} = \mathbb{E}_F F$, averaging cancels the derivative term and identifies the expected divergence with the entropy gap:

$$\mathbb{E}_F D_f(F, \bar{F}) = H_1(\bar{F}) - \mathbb{E}_F H_1(F).$$

Proof of Lemma 4.3. For $u > 0$, abbreviate

$$R_F = (F + uI)^{-1}, \quad R_{\bar{F}} = (\bar{F} + uI)^{-1}, \quad J = F - \bar{F}.$$

The pointwise difference between the two resolvent purifications is

$$L_u = F(F + uI)^{-1} - \bar{F}(\bar{F} + uI)^{-1} = uR_F J R_{\bar{F}} = uR_{\bar{F}} J R_F.$$

Since $uR_F^2 \preceq R_F$,

$$L_u^2 = u^2 R_{\bar{F}} J R_F^2 J R_{\bar{F}} \preceq u R_{\bar{F}} J R_F J R_{\bar{F}}.$$

The second-order resolvent identity gives

$$R_F - R_{\bar{F}} + R_{\bar{F}} J R_{\bar{F}} = R_{\bar{F}} J R_F J R_{\bar{F}}.$$

Taking expectations and using $\mathbb{E}_F J = 0$, we obtain

$$\mathbb{E}_F L_u^2 \preceq u(\mathbb{E}_F R_F - R_{\bar{F}}).$$

For every positive semidefinite C ,

$$\int_0^T u(C + uI)^{-1} du = TI - C \log(C + TI) + C \log C.$$

Integrate the preceding operator inequality, cancel the TI terms, and let $T \rightarrow \infty$. The remaining large- T terms vanish because $\log(F + TI) = (\log T)I + \log(I + F/T)$ and $\mathbb{E}_F F = \bar{F}$. Therefore

$$\int_0^\infty \mathbb{E}_F L_u^2 du \preceq \mathbb{E}_F [F \log F] - \bar{F} \log \bar{F} = H_1(\bar{F}) - \mathbb{E}_F H_1(F).$$

The integral on the left is precisely the left side of (25). □

A.5 Probability-weighted accepted-answer entropy

Proof of Lemma 4.5. Fix background information U containing X_D, Y_D . The accepted answer words have total probability $p_W(U)$ and their number is at most

$$|\mathcal{Z}| = |A|^{|D|} |B|^{|D|} = e^s.$$

Applying the log-sum inequality to their probabilities $p_z(U)$ gives

$$\sum_{\mathbf{1}_{W_D}(\vec{U}, z)=1} H_1(p_z(U)) \leq H_1(p_W(U)) + p_W(U)s,$$

which is (31). Since $\mathbb{E}_{U \sim \mathcal{P}^0} p_W(U) = p$ and H_1 is concave,

$$\mathbb{E}_{U \sim \mathcal{P}^0} H_1(p_W(U)) \leq H_1(p) = p\tau.$$

Averaging (31) therefore gives (32). □

A.6 Classical sampleability of the revealed history

Proof of Lemma 4.9. Recall the two locally generated tuple distributions

$$\mathcal{J}_A(i, r, x, y) = \frac{1}{m} \mu(x, y) \mathcal{Q}(r \mid i, X_i = x), \quad \mathcal{J}_B(i, r, x, y) = \frac{1}{m} \mu(x, y) \mathcal{Q}(r \mid i, Y_i = y).$$

The live coordinate i remains uniform under $\mathcal{Q}$, so the relative-entropy chain rule gives

$$\begin{aligned} D(\mathcal{Q} \parallel \mathcal{J}_A) &= \mathbb{E}_{i \sim \text{Unif}(M)} D(\mathcal{Q}_{X_i \mid i} \parallel \mu_X) \\ &\quad + \mathbb{E}_{(i, r, x, y) \sim \mathcal{Q}} D\left(\mathcal{Q}_{Y_i \mid i, X_i, T, Z} \parallel \mu(\cdot \mid X_i)\right). \end{aligned} \quad (40)$$

The product prior and (33) imply

$$\sum_{i \in M} D(\mathcal{Q}_{X_i} \parallel \mu_X) \leq D(\mathcal{Q}_{X_M} \parallel \mu_X^{\otimes m}) \leq \tau,$$

so the first term of (40) is at most τ/m .

For the second term use the Bob-reveal reverse construction from Section 4. Fix a partition $M = L_X^+ \sqcup L_Y$, its orders, the L_Y cut, and the background

$$U = (X_D, Y_D, X_{L_X^+}, Y_{L_Y}, X_{\pi_{Y, -i}^{\leq k_Y}}), \quad |L_X^+| > 0.$$

Equation (34) specializes to

$$D\left(\mathcal{Q}_{U, Y_{L_X^+}, Z} \parallel \mathcal{P}_{U, Y_{L_X^+}} \otimes \text{Unif}_Z\right) \leq \tau + s.$$

Under the reference distribution, conditionally on U , the variables Y_j , $j \in L_X^+$, are independent with distributions $\mu(\cdot \mid X_j)$. Apply the chain rule in the order π_X and discard its nonnegative initial term:

$$\sum_{j=1}^{|L_X^+|} \mathbb{E}_{\mathcal{Q}} D\left(\mathcal{Q}_{Y_{\pi_X[j]} \mid U, Z, Y_{\pi_X^{< j}}}\right) \parallel \mu(\cdot \mid X_{\pi_X[j]}) \leq \tau + s.$$

The live cut is uniform over these $|L_X^+|$ positions. Its conditioning data are precisely (i, T, X_i, Z) , while the reverse partition has weight $2|L_X^+|/m$. These factors cancel, giving

$$\mathbb{E}_{(i, r, x, y) \sim \mathcal{Q}} D\left(\mathcal{Q}_{Y_i \mid i, X_i, T, Z} \parallel \mu(\cdot \mid X_i)\right) \leq \frac{2(\tau + s)}{m}.$$

Consequently,

$$D(\mathcal{Q} \parallel \mathcal{J}_A) \leq \frac{3\tau + 2s}{m}.$$

Interchanging Alice and Bob gives the same bound for $\mathcal{J}_B$. Pinsker's inequality and $\eta = (\tau + s)/m$ imply

$$\text{d}_{\text{TV}}(\mathcal{Q}, \mathcal{J}_A), \text{d}_{\text{TV}}(\mathcal{Q}, \mathcal{J}_B) \leq \sqrt{\frac{3\tau + 2s}{2m}} \leq \sqrt{\frac{3}{2}} \eta = \kappa.$$

Finally, apply Lemma 3.4 to the finite family

$$\{\mathcal{Q}(R \mid i, X_i = x), \mathcal{Q}(R \mid i, Y_i = y)\}_{i, x, y},$$

including arbitrary fixed defaults for zero posterior marginals. Use public randomness to choose the uniform coordinate i and perform the correlated sampling. The resulting histories have exact tuple distributions $\mathcal{J}_A, \mathcal{J}_B$, and

$$\mathbb{P}(r_A \neq r_B) \leq 2 \text{d}_{\text{TV}}(\mathcal{J}_A, \mathcal{J}_B) \leq 2(\text{d}_{\text{TV}}(\mathcal{J}_A, \mathcal{Q}) + \text{d}_{\text{TV}}(\mathcal{Q}, \mathcal{J}_B)) \leq 4\kappa.$$

This proves (35). □

A.7 Proof of finite classical correlated sampling

Proof of Lemma 3.4. Temporarily generate shared independent proposals

$$(Z_j, U_j)_{j \geq 1}, \quad Z_j \sim \text{Unif}(\mathcal{R}), \quad U_j \sim \text{Unif}[0, 1].$$

The procedure indexed by $\mathcal{D}$ outputs the first Z_j for which $U_j \leq \mathcal{D}(Z_j)$. Each proposal is accepted with probability $1/|\mathcal{R}|$, and the output has exactly distribution $\mathcal{D}$. For $\mathcal{D}, \mathcal{E}$, inspect the first proposal accepted by either procedure. The probability that it is accepted by only one of them is

$$\frac{\sum_z |\mathcal{D}(z) - \mathcal{E}(z)|}{\sum_z \max\{\mathcal{D}(z), \mathcal{E}(z)\}} = \frac{2 \text{d}_{\text{TV}}(\mathcal{D}, \mathcal{E})}{1 + \text{d}_{\text{TV}}(\mathcal{D}, \mathcal{E})} \leq 2 \text{d}_{\text{TV}}(\mathcal{D}, \mathcal{E}).$$

When the first proposal is accepted by both, their outputs agree; thus this also bounds their eventual disagreement. Since $\mathcal{F}$ is finite, all procedures terminate simultaneously almost surely. Their joint output vector belongs to the finite set $\mathcal{R}^{\mathcal{F}}$. Use that vector itself as the finite shared random variable, and let each player read the component indexed by the locally specified distribution. This preserves the exact marginals and all the disagreement bounds without using an infinite randomness register. $\square$

A.8 Quantitative details for the single-game rounding

We finish the error calculation deferred from Lemma 3.6. On a matching valid branch (i, r, x, y) , Lemma 3.5 with

$$\sigma_A = \Psi_{r,x}^A, \quad \sigma_B = \Psi_{r,y}^B$$

bounds the distance from Alice's candidate state. A further triangle inequality bounds the distance from the ideal state by

$$\begin{aligned} e(i, r, x, y) &= K_{\text{qs}} \left(\alpha^{1/12} + \left\| \Psi_{r,x}^A - \Psi_{r,y}^B \right\|^{1/6} \right) \\ &\quad + \left\| \Psi_{r,x}^A - \Psi_{r,x,y} \right\|. \end{aligned}$$

Set $e = U_*$ outside the support of $\mathcal{Q}$. Since unit vectors have distance at most 2, one has $0 \leq e \leq U_*$.

The state alignment in Lemma 3.3 and Jensen's inequality give

$$\begin{aligned} \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x}^A - \Psi_{r,y}^B \right\|^{1/6} &\leq \left(\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x}^A - \Psi_{r,y}^B \right\|^2 \right)^{1/12} \\ &\leq (32\eta)^{1/12}. \end{aligned}$$

Similarly, $\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} \left\| \Psi_{r,x}^A - \Psi_{r,x,y} \right\| \leq \sqrt{8\eta}$. Consequently,

$$\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} e(i, r, x, y) \leq K_{\text{qs}}(\alpha^{1/12} + (32\eta)^{1/12}) + \sqrt{8\eta}.$$

Let $w(i, r, x, y)$ be the ideal branch winning probability, and set $w = 0$ outside the support of $\mathcal{Q}$. Then $0 \leq w \leq 1$, $\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} w(i, r, x, y) = q$, and (14) implies

$$\mathbb{E}_{(i,r,x,y) \sim \mathcal{J}_A} w(i, r, x, y) \geq q - \kappa, \quad \mathbb{E}_{(i,r,x,y) \sim \mathcal{J}_A} e(i, r, x, y) \leq \mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} e(i, r, x, y) + U_* \kappa.$$

Histories disagree with probability at most 4κ . On a matching branch, replacing the ideal state by one at distance e changes any measurement probability by at most $2e$. Therefore the strategy (15) satisfies

$$\begin{aligned} \text{win}(S_\alpha) &\geq q - 5\kappa - 2(\mathbb{E}_{(i,r,x,y) \sim \mathcal{Q}} e(i, r, x, y) + U_* \kappa) \\ &\geq q - (5 + 2U_*)\kappa - 2K_{\text{qs}}\alpha^{1/12} - 2K_{\text{qs}}(32\eta)^{1/12} - 2\sqrt{8\eta}. \end{aligned}$$

Finally, $0 \leq \eta \leq 1$ gives

$$\kappa \leq \sqrt{\frac{3}{2}} \eta^{1/12}, \quad \sqrt{\eta} \leq \eta^{1/12}.$$

Using the constants in (11) proves (12) with exactly the stated universal constant B_{qs} .

References

- [BVY17] M. Bavarian, T. Vidick, and H. Yuen. Hardness amplification for entangled games via anchoring. In *Proceedings of the 49th ACM Symposium on Theory of Computing*, pages 303–316, 2017. doi:10.1145/3055399.3055433. Revised full version, *Anchored parallel repetition for nonlocal games*: <https://arxiv.org/abs/1509.07466v2>.
- [CS14] A. Chailloux and G. Scarpa. Parallel repetition of entangled games with exponential decay via the superposed information cost. In *ICALP 2014, Lecture Notes in Computer Science* 8572, pages 296–307, 2014. doi:10.1007/978-3-662-43948-7_25; corrected full version: <https://arxiv.org/abs/1310.7787v3>.
- [CS15] A. Chailloux and G. Scarpa. Parallel repetition of free entangled games: simplification and improvements. *arXiv:1410.4397v2*, 2015 (first version 2014). <https://arxiv.org/abs/1410.4397v2>.
- [CHTW04] R. Cleve, P. Høyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th IEEE Conference on Computational Complexity*, pages 236–249, 2004. doi:10.1109/CCC.2004.1313847; <https://arxiv.org/abs/quant-ph/0404076v1>.
- [CSUU08] R. Cleve, W. Slofstra, F. Unger, and S. Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity* 17:282–299, 2008. doi:10.1007/s00037-008-0250-4.
- [DSV15] I. Dinur, D. Steurer, and T. Vidick. A parallel repetition theorem for entangled projection games. *Computational Complexity* 24(2):201–254, 2015. doi:10.1007/s00037-015-0098-3. Full arXiv version, including Lemma 17: [arXiv:1310.4113v2](https://arxiv.org/abs/1310.4113v2).
- [FMZ25] H. Fu, K. Mastel, and X. Zhang. Succinct perfect zero-knowledge for MIP*. *arXiv:2503.04517v2*, 2025. <https://arxiv.org/abs/2503.04517v2>.
- [Hol09] T. Holenstein. Parallel repetition: Simplification and the no-signaling case. *Theory of Computing* 5:141–172, 2009. doi:10.4086/toc.2009.v005a008.
- [JPY14] R. Jain, A. Pereszlényi, and P. Yao. A parallel repetition theorem for entangled two-player one-round games under product distributions. In *Proceedings of the 29th IEEE Conference on Computational Complexity*, pages 209–216, 2014. doi:10.1109/CCC.2014.29.
- [KRT10] J. Kempe, O. Regev, and B. Toner. Unique games with entangled provers are easy. *SIAM Journal on Computing* 39(7):3207–3229, 2010. doi:10.1137/090772885.
- [KV11] J. Kempe and T. Vidick. Parallel repetition of entangled games. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 353–362, 2011. doi:10.1145/1993636.1993684.
- [Kim14] I. H. Kim. Modulus of convexity for operator convex functions. *Journal of Mathematical Physics* 55:082201, 2014. doi:10.1063/1.4890292; <https://arxiv.org/abs/1310.0746v3>.
- [Lin25] J. Lin. $\text{MIP}^{\text{co}} = \text{coRE}$. *arXiv:2510.07162*, 2025. <https://arxiv.org/abs/2510.07162>.
- [NC10] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. 10th Anniversary Edition, Cambridge University Press, 2010.
- [Pet07] D. Petz. Bregman divergence as relative operator entropy. *Acta Mathematica Hungarica* 116:127–131, 2007. doi:10.1007/s10474-007-6014-9.
- [Rao11] A. Rao. Parallel repetition in projection games and a concentration bound. *SIAM Journal on Computing* 40(6):1871–1891, 2011. doi:10.1137/080734042.
- [Raz98] R. Raz. A parallel repetition theorem. *SIAM Journal on Computing* 27(3):763–803, 1998. doi:10.1137/S0097539795280895.
- [Raz11] R. Raz. A counterexample to strong parallel repetition. *SIAM Journal on Computing* 40(3):771–777, 2011. doi:10.1137/090747270.
- [Wat18] J. Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018. doi:10.1017/9781316848142.

- [Yue16] H. Yuen. A parallel repetition theorem for all entangled games. In *43rd International Colloquium on Automata, Languages, and Programming*, LIPIcs 55, article 77, 2016. doi:10.4230/LIPIcs.ICALP.2016.77. Full version: <https://arxiv.org/abs/1604.04340>.

CHAPTER 7

$n^{1/400}$ -Hardness of the Euclidean Closest Vector Problem

ABSTRACT. We study the hardness of approximating the Euclidean closest vector problem within a fixed polynomial in the lattice dimension. We present a deterministic polynomial-time many-one reduction from 3SAT to $\text{GapCVP}_{n^{1/400}}^{(2)}$, where n is the rank of a lattice given by an explicit square integer basis. The construction uses Reed–Solomon power-sum constraints over a characteristic-two field to encode assignments and their satisfying restrictions to individual clauses. It converts the resulting binary affine system into an integer lattice by coordinatewise reduction modulo two. Completeness follows from polynomial interpolation. For soundness, a low-weight binary solution gives small sets of selected field values. Reconstruction from their power sums produces separable root sets over a rational function field; additional power-sum constraints recover the Boolean clause assignments through valuations, and the clause equations identify a root compatible with every clause. The same construction yields approximation factors $n^{1/200}$ for binary nearest codeword and syndrome decoding and $n^{1/(200p)}$ for closest vector in every fixed rational ℓ_p norm, $p \geq 1$. The result is a direct reduction from 3SAT and therefore does not invoke the PCP theorem or assume the Projection Games Conjecture.

CONTENTS

1. Introduction
 2. Preliminaries
 3. Encoding Boolean assignments by Reed–Solomon codes
 4. Algebraic reconstruction of bounded polynomial-moment families
 5. Soundness of the nearest-codeword reduction
 6. Transfer to CVP and complexity
 7. Further consequences
- References

1 Introduction

Given a full-rank integer lattice and a target, the closest vector problem (CVP) asks for a lattice point of minimum Euclidean distance to the target. The computational complexity of CVP depends on how closely one seeks to approximate this minimum. Writing n for the lattice rank, we study approximation factors that grow as a fixed polynomial in n .

1.1 The closest vector problem and main result

To formulate the closest vector problem precisely, we first describe the lattice generated by an integer basis and the distance from a target to that lattice.

For a nonsingular integer matrix $B \in \mathbb{Z}^{n \times n}$, the lattice $L(B)$ generated by the columns of B is

$$L(B) = B\mathbb{Z}^n = \{Bz : z \in \mathbb{Z}^n\}.$$

The Euclidean distance from a target $t \in \mathbb{Q}^n$ to the lattice $L(B)$ is

$$\text{dist}_2(t, L(B)) = \min_{z \in \mathbb{Z}^n} \|t - Bz\|_2.$$

For an approximation factor $\gamma(n) \geq 1$, define the approximate closest vector problem $\text{GapCVP}_\gamma^{(2)}$ as follows. An instance consists of a nonsingular integer lattice basis $B \in \mathbb{Z}^{n \times n}$, a target $t \in \mathbb{Q}^n$, and a positive rational radius $r \in \mathbb{Q}_{>0}$, all encoded in binary. The task is to distinguish between the following two cases:

$$\begin{aligned} \text{YES : } & \text{dist}_2(t, L(B)) \leq r, \\ \text{NO : } & \text{dist}_2(t, L(B)) > \gamma(n)r. \end{aligned}$$

An algorithm must return the indicated answer whenever one of these conditions holds. If $r < \text{dist}_2(t, L(B)) \leq \gamma(n)r$, either answer is allowed.

CVP is a central computational problem in the geometry of numbers, with connections to computational number theory, integer optimization, and cryptography. The lattice-basis reduction algorithm of Lenstra, Lenstra, and Lovász led to applications in polynomial factorization [LLL82], and Babai used reduced bases to approximate nearest lattice points, with applications to integer programming [Bab86].

The importance of lattice problems extends to public-key and post-quantum cryptography. Ajtai established a foundational connection between worst-case lattice problems and average-case computational hardness [Ajt96]; Goldreich, Goldwasser, and Halevi proposed public-key encryption and signatures based on lattice-reduction problems [GGH97]; and Regev introduced learning with errors and related its hardness to worst-case lattice problems [Reg09]. The practical significance of this line of work is reflected in the standardization of lattice-based key encapsulation and digital signatures by the National Institute of Standards and Technology [NIST24a, NIST24b]. These cryptographic applications rely on appropriate structured or average-case assumptions, rather than directly on the worst-case NP-hardness of CVP. Nonetheless, they reinforce the importance of determining which approximation regimes for fundamental lattice problems remain computationally intractable.

Theorem 1 (Main theorem). *There is a deterministic polynomial-time mapping that assigns to each 3SAT formula φ a nonsingular matrix $B \in \mathbb{Z}^{n \times n}$, a target $t \in \mathbb{Z}^n$, and a radius $r \in \mathbb{Q}_{>0}$ such that*

$$\begin{aligned} \varphi \text{ satisfiable} & \implies \text{dist}_2(t, L(B)) \leq r, \\ \varphi \text{ unsatisfiable} & \implies \text{dist}_2(t, L(B)) > n^{1/400}r. \end{aligned}$$

Consequently, $\text{GapCVP}_{n^{1/400}}^{(2)}$ is NP-hard under deterministic polynomial-time many-one reductions.

The reduction uses no randomized step, gap-producing PCP, or Projection Games Conjecture.

Two related binary coding problems also play a central role. Write $\mathbb{F}_2$ for the field with two elements. A binary linear code $\mathcal{C} \leq \mathbb{F}_2^n$ is a linear subspace, and a generator matrix for $\mathcal{C}$ has rows spanning that subspace. For a binary word z , write $\text{wt}(z)$ for its *Hamming weight*, the number of nonzero coordinates. Given $\mathcal{C}$ and a received word $u \in \mathbb{F}_2^n$, the *nearest-codeword problem* asks for the minimum Hamming distance

$$d_H(u, \mathcal{C}) = \min_{c \in \mathcal{C}} \text{wt}(u - c).$$

Its promise version receives a generator for $\mathcal{C}$, the word u , and an integer radius $R > 0$. The equivalent *syndrome-decoding problem* receives a binary parity-check matrix H , a syndrome b , and the same radius, and asks for

$$\min\{\text{wt}(x) : x \in \mathbb{F}_2^n, Hx = b\}.$$

Here H is a parity-check matrix for $\mathcal{C} = \ker_{\mathbb{F}_2} H$. If $b = Hu$, the solutions of $Hx = b$ form the affine coset $u + \mathcal{C}$, so the two objectives agree. For approximation factor $\gamma(n)$, the YES threshold is R , and the strict NO threshold is $\gamma(n)R$; here n denotes the code block length.

The reduction also establishes approximation hardness within $n^{1/200}$ for nearest codeword and syndrome decoding. For every fixed rational $p \geq 1$, replacing the Euclidean norm with the ℓ_p norm gives a closest-vector factor $n^{1/(200p)}$; setting $p = 2$ recovers the main Euclidean exponent. Let N denote the input-size parameter and let q denote the size of the finite field used in the construction. The field size and lattice dimension satisfy

$$q = \Theta(N^{200}), \quad n \leq 40N^{401},$$

so the construction has large, but fixed-polynomial, output and bit complexity.

1.2 Proof overview

The reduction factors through binary nearest codeword:

$$3\text{SAT} \longrightarrow \text{binary nearest codeword} \longrightarrow \text{Euclidean CVP}.$$

The first arrow represents our main technical contribution; the second is the standard dimension-preserving reduction from binary nearest codeword to Euclidean CVP. We now explain the first arrow, the reduction from 3SAT to binary nearest codeword.

Let the input formula have m variables and clauses $C_1, \dots, C_\ell$. Choose a characteristic-two field $\mathbb{K} = \mathbb{F}_q$, distinct anchors $a_1, \dots, a_m \in \mathbb{K}$, and the evaluation set

$$P = \mathbb{K} \setminus \{a_1, \dots, a_m\}.$$

An assignment $\sigma \in \{0, 1\}^m$ determines the unique interpolation polynomial $Q_\sigma \in \mathbb{K}[X]$ satisfying

$$\deg Q_\sigma < m, \quad Q_\sigma(a_i) = \sigma_i.$$

For every $p \in P$ and $w \in \mathbb{K}$, introduce a binary coordinate $x_{0,p,w}$. The intended encoding of Q_σ is

$$x_{0,p,w} = 1 \iff w = Q_\sigma(p).$$

For a clause C , let I_C be the indices of its variables, and let $\mathcal{B}_C \subseteq \{0, 1\}^{I_C}$ be the set of assignments satisfying C . For every $\beta \in \mathcal{B}_C$, introduce another array of binary coordinates $x_{(C,\beta),p,w}$. If σ satisfies C , the intended assignment is

$$x_{(C,\beta),p,w} = 1 \iff \beta = \sigma|_{I_C} \text{ and } w = Q_\sigma(p).$$

Call these binary arrays *evaluation tables*: the global table has index $\tau = 0$, and the clause tables have indices $\tau = (C, \beta)$. Such an index τ is called the *table type*. For any type, define its support at p by

$$S_\tau(p) = \{w \in \mathbb{K} : x_{\tau,p,w} = 1\}.$$

For the intended assignment, the support of table 0 and each selected clause table is $\{Q_\sigma(p)\}$; all other supports are empty. Writing M for the total number of binary coordinates, the proof proceeds as follows.

- (i) Encoding and completeness (Section 3). Impose low-degree constraints on the ordinary and shifted moments

$$\mu_{\tau,j}(p) = \sum_{w \in S_\tau(p)} w^j, \quad \eta_{\tau,i,j}(p) = \sum_{w \in S_\tau(p)} \left(\frac{w - \beta_i}{p - a_i} \right)^j,$$

where the shifted moment is defined for a table indexed by $\tau = (C, \beta)$ and a variable i appearing in C . Since $p \notin \{a_1, \dots, a_m\}$, its denominator is nonzero. Require the moments to be low-degree polynomial evaluations, the support of table 0 to have odd size, and the tables for each clause to reproduce the indicators of table 0 modulo two. These conditions give a binary affine system $Hx = b$. For a satisfying assignment σ , the indicator assignment defined above gives a solution of weight

$$R = (\ell + 1)|P|,$$

as proved in Lemma 8.

- (ii) Algebraic reconstruction (Section 4). Lemma 10 treats a family of bounded sets whose power sums are low-degree polynomials. It reconstructs one separable polynomial over the rational function field $\mathbb{K}(X)$; its roots, in a finite separable extension, have exactly those prescribed power sums. This statement is proved independently of the reduction from 3SAT to binary nearest codeword.
- (iii) Nearest-codeword soundness (Section 5). A low-weight solution of $Hx = b$ yields bounded support sets $S_\tau(p)$ after discarding a small number of evaluation points separately for each type. For each table τ , Lemma 10 produces a root set $\mathcal{R}_\tau$; thus $\mathcal{R}_0$ corresponds to table 0, and $\mathcal{R}_{(C,\beta)}$ corresponds to the table indexed by (C, β) . Lemma 12 recovers the clause-assignment bit β_i from roots in $\mathcal{R}_{(C,\beta)}$. Lemma 13 associates each $\alpha \in \mathcal{R}_0$ with a satisfying assignment β for each clause C such that $\alpha \in \mathcal{R}_{(C,\beta)}$. These assignments agree on shared variables, giving Proposition 14: a sufficiently low-weight solution implies that the original formula is satisfiable.
- (iv) Transfer to CVP (Sections 2 and 6). Lemma 7 converts $Hx = b$ into an integer lattice Λ_H and a target u with

$$\text{dist}_2(u, \Lambda_H)^2 = \min_{\substack{x \in \mathbb{F}_2^M \\ Hx=b}} \text{wt}(x).$$

Consequently, the binary approximation factor $M^{1/200}$ becomes the Euclidean factor $n^{1/400}$, where the lattice rank is $n = M$. Section 6 combines completeness, soundness, and this standard reduction to prove Theorem 1.

- (v) Further consequences (Section 7). Corollaries 15 and 16 give the corresponding hardness results for binary decoding and for closest vector in each fixed finite rational norm.

1.3 Background and related work

Hardness of approximating closest vector. The complexity-theoretic study of CVP begins with the NP-hardness of its exact version, proved by van Emde Boas [vEB81]. Arora, Babai, Stern, and Sweedyk extended this line of work to approximation, establishing hardness within every fixed constant for problems involving lattices, codes, and systems of linear equations [ABSS97]. Dinur, Kindler, and Safra subsequently obtained an almost-polynomial hardness factor for CVP [DKS98], refined in the journal version by Dinur, Kindler, Raz, and Safra [DKRS03]. In particular, the latter work proves NP-hardness of Euclidean CVP within

$$n^{a/\log \log n} \quad \text{for some absolute } a > 0.$$

Although this factor grows faster than every fixed power of $\log n$, its exponent tends to zero. Consequently, it does not establish hardness within n^c for any fixed $c > 0$.

Complexity-theoretic limits and conditional results. The significance of strengthening the approximation factor becomes clear from upper bounds in a nearby regime. Goldreich and Goldwasser established interactive upper bounds for lattice approximation problems [GG00], and Aharonov and Regev proved that, for some absolute constant $C > 0$,

$$\text{GapCVP}_{C\sqrt{n}}^{(2)} \in \text{NP} \cap \text{coNP}.$$

Thus, NP-hardness at the square-root scale would imply $\text{NP} = \text{coNP}$ [AR05]. Between the exponent $1/400$ established here and this square-root barrier, it remains open to determine the largest $c \in [1/400, 1/2)$ for which $\text{GapCVP}_{n^c}^{(2)}$ is NP-hard.

Moshkovitz [Mos15, Section 5.1] and Mukhopadhyay [Muk22] obtain fixed-polynomial lattice inapproximability conditional on the Projection Games Conjecture. Such conditional conclusions differ from the unconditional, deterministic reduction from ordinary 3SAT in Theorem 1.

Coding problems and PCP-free reductions. The coding-theoretic problems used in our reduction have a related but distinct history. Berlekamp, McEliece, and van Tilborg proved the intractability of syndrome decoding [BMVT78], while Vardy established the hardness of computing the minimum distance of a linear code [Var97]. Syndrome decoding and nearest codeword concern an affine coset, whereas minimum distance concerns nonzero vectors in a linear subspace. Hardness for either formulation therefore does not automatically provide an approximation gap for the other.

Bhattachoplu, Guruswami, and Ren gave deterministic, PCP-free reductions establishing NP-hardness of approximating nearest codeword and minimum distance within any fixed constant factor over each fixed finite field [BGR25]. Almost-polynomial coding gaps in that work require the stronger assumption that NP has no quasipolynomial-time algorithms. The broader work of Bhattachoplu, Guruswami, Lee, and Ren also studies sparse vectors in real subspaces and lattices, with randomized reductions in those settings [BGLR25]. These results do not yield a deterministic n^δ -factor hardness result for Euclidean CVP with any fixed $\delta > 0$.

Algebraic reconstruction and lattice lifts. The reduction in this paper uses the polynomial-evaluation codes introduced by Reed and Solomon [RS60]. The associated reconstruction step is related to recovering algebraic information from unordered evaluations and to Reed–Solomon list recovery [Sud97, ALRS98, GS99, GR06]. The power-sum Hankel systems also have antecedents in Massey’s shift-register decoding and in subsequent key-equation methods [Mas69, RR00, ZGA11]. The structural conclusion needed here differs

from a conventional list-decoding output: it produces a complete set of algebraic roots, which need not be rational functions, rather than only a list of low-degree codewords.

Bennett and Peikert likewise combine Reed–Solomon codes with an integer lattice lift, but in a randomized hardness reduction for the shortest vector problem. That work also discusses the obstacles to derandomizing the reduction [BP23].

Fine-grained hardness. Recent work also examines how quickly approximate lattice problems can be solved. Aggarwal, Gupta, Morolia, and Zhang establish deterministic ETH-based lower bounds for constant-factor CVP in every finite ℓ_p norm [AGMZ26]. Huang, Ko, and Wang study related classical and quantum fine-grained reductions for constant-factor CVP and Max-Cut [HKW26]. These results address running-time lower bounds at constant approximation factors rather than unconditional NP-hardness at a fixed polynomial approximation factor.

2 Preliminaries

We recall finite-field evaluations, valuations, and the standard dimension-preserving reduction from binary decoding to CVP.

2.1 Finite-field evaluations

Let $\mathbb{K} = \mathbb{F}_q$, let $P \subseteq \mathbb{K}$, and let D be a nonnegative integer with $D < |P|$. Define

$$\text{RS}(P, D) = \left\{ (f(p))_{p \in P} : f \in \mathbb{K}[X], \deg f \leq D \right\}.$$

Since a nonzero polynomial of degree at most D has at most D roots, every codeword has a unique interpolating polynomial. An element of the evaluation code is a *codeword*, and a *parity check* is a linear equation satisfied by every codeword. A membership test is an explicit $\mathbb{K}$ -linear system: form the evaluation matrix of $1, X, \dots, X^D$, compute parity checks by Gaussian elimination, and require each parity check to vanish. After fixing a basis of $\mathbb{K}$ over $\mathbb{F}_2$, each such equation becomes a finite list of binary-linear equations.

2.2 Valuations

A valuation records orders of vanishing and poles at a specified point.

Rational functions and orders of vanishing. Let $\mathbb{K}$ be a field. Its rational function field in the indeterminate X is

$$F = \mathbb{K}(X) = \left\{ \frac{g(X)}{h(X)} : g, h \in \mathbb{K}[X], h \neq 0 \right\}.$$

Fix $a \in \mathbb{K}$ and write $\ell_a = X - a$. A rational function cannot always be evaluated at a : for example, $1/\ell_a$ has a pole there. Nevertheless, it has a well-defined order of vanishing at a .

For a nonzero polynomial $g \in \mathbb{K}[X]$, let $\text{ord}_{\ell_a}(g)$ be the largest integer $m \geq 0$ such that ℓ_a^m divides g . Equivalently,

$$g = \ell_a^m g_0, \quad g_0(a) \neq 0.$$

For nonzero rational functions, define

$$\text{ord}_{\ell_a} \left(\frac{g}{h} \right) = \text{ord}_{\ell_a}(g) - \text{ord}_{\ell_a}(h), \quad \text{ord}_{\ell_a}(0) = +\infty.$$

Unique factorization in $\mathbb{K}[X]$ shows that the definition does not depend on the chosen representation of g/h . Positive orders correspond to zeros, negative orders to poles, and order zero to functions that are nonzero and regular at a . In particular,

$$\text{ord}_{\ell_a}(\ell_a^3) = 3, \quad \text{ord}_{\ell_a}(1/\ell_a) = -1, \quad \text{ord}_{\ell_a}(c) = 0 \quad (c \in \mathbb{K}^*).$$

Valuations and their basic properties. An additive, rational-valued *valuation* on a field E is a function

$$v : E \longrightarrow \mathbb{Q} \cup \{+\infty\}$$

such that, for all $x, y \in E$,

$$\begin{aligned} v(x) = +\infty &\iff x = 0, \\ v(xy) &= v(x) + v(y), \\ v(x + y) &\geq \min\{v(x), v(y)\}. \end{aligned}$$

The third property is the *ultrametric inequality*. Multiplicativity gives

$$v(1) = v(-1) = 0, \quad v(-x) = v(x), \quad v(x^j) = jv(x) \quad (x \neq 0, j \geq 0).$$

Since $\ell_a = X - a$, the two notations $\text{ord}_{\ell_a} = \text{ord}_{X-a}$ denote the same integer-valued valuation on F .

The ultrametric inequality prevents a uniquely minimum-valuation summand from being canceled.

Lemma 2 (Minimum-valuation principle). *Let v be a valuation on E . If $v(x) \neq v(y)$, then*

$$v(x + y) = \min\{v(x), v(y)\}.$$

More generally, if $r \geq 2$, $x_1, \dots, x_r \in E^$, and exactly one x_j has minimum valuation, then*

$$v\left(\sum_{i=1}^r x_i\right) = v(x_j).$$

In particular, such a sum cannot be zero.

Proof. Suppose $v(x) < v(y)$. The ultrametric inequality gives $v(x+y) \geq v(x)$. If the inequality were strict, applying the same property to $x = (x+y) - y$ would give

$$v(x) \geq \min\{v(x+y), v(y)\} > v(x),$$

a contradiction. For the general statement, put $y = \sum_{i \neq j} x_i$. Repeated applications of the ultrametric inequality give

$$v(y) \geq \min_{i \neq j} v(x_i) > v(x_j),$$

so the two-term statement applies to $x_j + y$. □

Separable polynomials and splitting fields. Let $G(Y) \in F[Y]$ be a nonzero polynomial. A *splitting field* of G is an extension E/F in which G factors into linear polynomials and which is generated over F by those roots. The polynomial G is *separable* if its roots in a splitting field are distinct. Equivalently,

$$\gcd\left(G, \frac{dG}{dY}\right) = 1.$$

The derivative criterion is especially useful in positive characteristic. For example, if $\text{char } F = 2$, then

$$\frac{d}{dY}(Y^2 - c) = 0,$$

and $Y^2 - c$ has a repeated root in any splitting field.

Lemma 3 (Common separable splitting field). *Let $G_1, \dots, G_r \in F[Y]$ be separable polynomials. There is a single finite separable extension E/F in which all G_i split into linear factors.*

Proof. Take a splitting field E of the product $G_1 \cdots G_r$. It is generated by finitely many roots, so E/F is finite. Every generating root belongs to a separable polynomial G_i and is therefore separable over F . A finite extension generated by separable elements is separable, proving the claim. The product need not itself have distinct roots: different G_i may share a root. $\square$

Extending a valuation. Let E/F be a finite separable extension. An *extension* of ord_{ℓ_a} to E is a valuation v_a satisfying

$$v_a(f) = \text{ord}_{\ell_a}(f) \quad (f \in F^*), \quad v_a(\ell_a) = 1.$$

Such an extension need not be integer-valued. For example, if $u \in E$ satisfies $u^e = \ell_a$, then

$$e v_a(u) = v_a(u^e) = v_a(\ell_a) = 1, \quad v_a(u) = \frac{1}{e}.$$

The appearance of fractional orders is called *ramification*.

Lemma 4 (Existence and normalization of extended valuations). *Let $\mathbb{K}$ be a field, $F = \mathbb{K}(X)$, $a \in \mathbb{K}$, and E/F a finite separable extension. There exists a valuation*

$$v_a : E \longrightarrow \mathbb{Q} \cup \{+\infty\}$$

whose restriction to F is ord_{X-a} . Moreover, for some positive integer e , its values on E^ lie in $\frac{1}{e}\mathbb{Z}$.*

Proof. An element $z \in E$ is *integral over $\mathbb{K}[X]$* if it satisfies a monic polynomial with coefficients in $\mathbb{K}[X]$. The set of all such elements,

$$A = \{z \in E : z \text{ is integral over } \mathbb{K}[X]\},$$

is the *integral closure* of $\mathbb{K}[X]$ in E . Because E/F is finite and separable, A is a finite $\mathbb{K}[X]$ -module and a Dedekind domain. The lying-over property provides a nonzero prime ideal $\mathfrak{p} \subseteq A$ above (ℓ_a) , that is,

$$\mathfrak{p} \cap \mathbb{K}[X] = (\ell_a).$$

Localizing at $\mathfrak{p}$ means allowing denominators outside $\mathfrak{p}$:

$$A_{\mathfrak{p}} = \left\{ \frac{x}{s} : x \in A, s \in A \setminus \mathfrak{p} \right\}.$$

The localization of a Dedekind domain at a nonzero prime is a *discrete valuation ring*. In particular, its maximal ideal is generated by one element t , called a *uniformizer*, and every $z \in E^*$ has a unique expression

$$z = t^n w, \quad n \in \mathbb{Z}, \quad w \in A_{\mathfrak{p}}^*.$$

Here $A_{\mathfrak{p}}^*$ denotes the units of the local ring. Define $\text{ord}_{\mathfrak{p}}(z) = n$ and put

$$e = \text{ord}_{\mathfrak{p}}(\ell_a) > 0, \quad v_a(z) = \frac{\text{ord}_{\mathfrak{p}}(z)}{e}, \quad v_a(0) = +\infty.$$

The integer e is the *ramification index*. If $g \in \mathbb{K}[X] \setminus \{0\}$, write

$$g = \ell_a^m g_0, \quad g_0(a) \neq 0.$$

Since $\mathfrak{p} \cap \mathbb{K}[X] = (\ell_a)$, we have $g_0 \notin \mathfrak{p}$, so g_0 is a unit in $A_{\mathfrak{p}}$. It follows that

$$\text{ord}_{\mathfrak{p}}(g) = em, \quad v_a(g) = m = \text{ord}_{\ell_a}(g).$$

Applying the same identity to numerators and denominators shows that $v_a(f) = \text{ord}_{\ell_a}(f)$ for every $f \in F^*$. Finally, $\text{ord}_{\mathfrak{p}}(E^*) = \mathbb{Z}$, so $v_a(E^*) = \frac{1}{e}\mathbb{Z}$. $\square$

Valuation rings and congruence. For a valuation v_a on E , its *valuation ring* and maximal ideal are

$$\mathcal{O}_{v_a} = \{z \in E : v_a(z) \geq 0\}, \quad \mathfrak{m}_{v_a} = \{z \in E : v_a(z) > 0\}.$$

The quotient

$$\kappa(v_a) = \mathcal{O}_{v_a} / \mathfrak{m}_{v_a}$$

is the *residue field*. Reduction modulo $\mathfrak{m}_{v_a}$ generalizes evaluation at a : if $f \in \mathbb{K}[X]$, then

$$f(X) - f(a) \in \ell_a \mathbb{K}[X], \quad v_a(f(X) - f(a)) \geq 1,$$

so $f(X)$ and $f(a)$ have the same residue. Since $v_a(c) = 0$ for every $c \in \mathbb{K}^*$, distinct elements of $\mathbb{K}$ remain distinct in $\kappa(v_a)$.

Lemma 5 (Uniqueness of a constant residue). *Let v_a extend ord_{X-a} to E . For $\alpha \in E$ and $b, c \in \mathbb{K}$, if*

$$v_a(\alpha - b) > 0, \quad v_a(\alpha - c) > 0,$$

then $b = c$. The same conclusion holds if both lower bounds are replaced by 1.

Proof. If $b \neq c$, then $b - c \in \mathbb{K}^*$, and therefore $v_a(b - c) = 0$. On the other hand, the ultrametric inequality gives

$$v_a(b - c) = v_a((\alpha - c) - (\alpha - b)) \geq \min\{v_a(\alpha - c), v_a(\alpha - b)\} > 0,$$

a contradiction. □

2.3 From binary nearest codeword to CVP

A binary affine decoding instance specifies a matrix and a prescribed right-hand side, called its syndrome:

$$H \in \mathbb{F}_2^{r \times M}, \quad b \in \mathbb{F}_2^r,$$

and the affine system

$$Hx = b, \quad x \in \mathbb{F}_2^M,$$

where all matrix operations are over $\mathbb{F}_2$. The system is *consistent* if there exists at least one $x \in \mathbb{F}_2^M$ satisfying $Hx = b$; equivalently,

$$b \in \text{im}_{\mathbb{F}_2} H, \quad \text{or, equivalently,} \quad \text{rank}_{\mathbb{F}_2} H = \text{rank}_{\mathbb{F}_2} [H \mid b].$$

It is *inconsistent* if no such solution exists. Gaussian elimination tests consistency and, in the consistent case, computes a particular solution in polynomial time. For a consistent system, the associated *minimum-weight affine solution problem* asks for a solution having the fewest nonzero coordinates:

$$W(H, b) = \min_{\substack{x \in \mathbb{F}_2^M \\ Hx = b}} \text{wt}(x).$$

Here $\text{wt}(x)$, the *Hamming weight* of x , is the number of its nonzero coordinates. Minimizing this weight is the standard *binary syndrome-decoding problem* [BMVT78, ABSS97].

The associated binary *nearest-codeword problem* is obtained from the homogeneous code

$$\mathcal{C} = \ker_{\mathbb{F}_2} H,$$

and a particular solution $u \in \mathbb{F}_2^M$ of $Hu = b$ identifies the entire solution set:

$$\{x \in \mathbb{F}_2^M : Hx = b\} = u + \mathcal{C}.$$

Since subtraction and addition agree in characteristic two,

$$W(H, b) = \min_{c \in \mathcal{C}} \text{wt}(u - c) = d_H(u, \mathcal{C}),$$

where $d_H(u, \mathcal{C})$ is the Hamming distance from u to the nearest codeword of $\mathcal{C}$. Conversely, given a binary code $\mathcal{C}$, a parity-check matrix H for $\mathcal{C}$, and a received word u , set

$$b = Hu.$$

Then nearest-codeword decoding is exactly minimum-weight decoding of $Hx = b$. Gaussian elimination converts between generator and parity-check representations in polynomial time.

A closest-vector instance specifies a nonsingular integer basis $B \in \mathbb{Z}^{M \times M}$, the lattice

$$L(B) = B\mathbb{Z}^M,$$

and a target $u \in \mathbb{Z}^M$. Its Euclidean objective is

$$\text{dist}_2(u, L(B)) = \min_{z \in \mathbb{Z}^M} \|u - Bz\|_2.$$

The reduction from nearest codeword to closest vector is standard. Feige and Micciancio and Regev use the corresponding coding-to-lattice connection in their hardness results with preprocessing; Alekhnovich, Khot, Kindler, and Vishnoi explicitly apply the resulting square-root conversion of approximation factors [FM04, Reg04, AKKV11].

The homogeneous code and its integer lift. The homogeneous system associated with H defines the binary linear code already introduced above:

$$\mathcal{C} = \ker_{\mathbb{F}_2} H = \{c \in \mathbb{F}_2^M : Hc = 0\}.$$

The matrix H is a *parity-check matrix* for $\mathcal{C}$; its rows need not be linearly independent. Rank-nullity gives

$$k_0 = \dim_{\mathbb{F}_2} \mathcal{C} = M - \text{rank}_{\mathbb{F}_2} H.$$

Let

$$\rho : \mathbb{Z}^M \longrightarrow \mathbb{F}_2^M, \quad \rho(z) = z \bmod 2,$$

be coordinatewise reduction modulo two. The inverse image of the homogeneous code is

$$\Lambda_H = \rho^{-1}(\mathcal{C}) = \{z \in \mathbb{Z}^M : H(z \bmod 2) = 0\}.$$

This inverse-image lattice is classical [LS71]. Bennett and Peikert likewise combine the corresponding coding-to-lattice lift with Reed–Solomon codes in their randomized shortest-vector reduction [BP23]. We call Λ_H the *parity-lift lattice* of H .

For $c \in \mathcal{C}$, let $\tilde{c} \in \{0, 1\}^M$ denote its integer lift. Then

$$\Lambda_H = \bigcup_{c \in \mathcal{C}} (\tilde{c} + 2\mathbb{Z}^M), \quad 2\mathbb{Z}^M \subseteq \Lambda_H \subseteq \mathbb{Z}^M.$$

Because $\mathcal{C}$ is an additive subgroup of $\mathbb{F}_2^M$, its inverse image Λ_H is an additive subgroup of $\mathbb{Z}^M$. Since it contains $2\mathbb{Z}^M$, it is a full-rank integer lattice in $\mathbb{R}^M$.

An explicit lattice basis. A closest-vector instance requires an integer basis, not merely a membership condition. Gaussian elimination puts $\mathcal{C}$ into *systematic form*: after a permutation of the M coordinates, there are k_0 free information coordinates $f \in \mathbb{F}_2^{k_0}$, and the remaining coordinates are determined by a binary matrix

$$P_{\text{sys}} \in \mathbb{F}_2^{(M-k_0) \times k_0}.$$

In these coordinates,

$$\mathcal{C} = \{(f, P_{\text{sys}}f) : f \in \mathbb{F}_2^{k_0}\}.$$

Here P_{sys} is the *parity-coordinate matrix*: it gives the dependent parity coordinates as a function of the free information coordinates. Write

$$\tilde{P}_{\text{sys}} \in \{0, 1\}^{(M-k_0) \times k_0}$$

for its entrywise integer lift.

These free and parity coordinates give the following explicit lattice basis.

Lemma 6 (Explicit parity-lift lattice basis). *In systematic coordinate order, Λ_H has integer basis*

$$B_H = \begin{pmatrix} I_{k_0} & 0 \\ \tilde{P}_{\text{sys}} & 2I_{M-k_0} \end{pmatrix}. \quad (1)$$

In particular,

$$\Lambda_H = B_H \mathbb{Z}^M, \quad |\det B_H| = 2^{M-k_0} = 2^{\text{rank}_{\mathbb{F}_2} H}.$$

The basis and the coordinate permutation can be computed in polynomial time from H .

Proof. In systematic coordinates, write

$$z = (f, g), \quad f \in \mathbb{Z}^{k_0}, \quad g \in \mathbb{Z}^{M-k_0}.$$

The definition of Λ_H gives

$$\begin{aligned} z \in \Lambda_H &\iff g \equiv \tilde{P}_{\text{sys}} f \pmod{2} \\ &\iff g = \tilde{P}_{\text{sys}} f + 2w \quad \text{for some } w \in \mathbb{Z}^{M-k_0} \\ &\iff z = B_H \begin{pmatrix} f \\ w \end{pmatrix}. \end{aligned}$$

Thus $\Lambda_H = B_H \mathbb{Z}^M$. The determinant follows from the block-triangular form. Gaussian elimination computes P_{sys} and the permutation; the inverse permutation restores the basis to the original coordinate order. $\square$

The affine system becomes a lattice coset. When $Hx = b$ is consistent, Gaussian elimination provides a particular solution

$$u \in \{0, 1\}^M, \quad Hu = b \quad \text{over } \mathbb{F}_2.$$

Every binary solution differs from u by a codeword:

$$\{x \in \mathbb{F}_2^M : Hx = b\} = (u \bmod 2) + \mathcal{C}.$$

Likewise, the integer vectors having syndrome b form the lattice coset

$$\{z \in \mathbb{Z}^M : H(z \bmod 2) = b\} = u + \Lambda_H.$$

Thus the natural candidate for the closest-vector instance is the integer basis B_H and target u . The following lemma gives their exact distance relation.

Lemma 7 (Standard nearest-codeword-to-CVP reduction). *Let $H \in \mathbb{F}_2^{r \times M}$ and $b \in \mathbb{F}_2^r$, and suppose $Hx = b$ is consistent. In deterministic polynomial time, one can construct a nonsingular matrix*

$$B_H \in \mathbb{Z}^{M \times M}$$

and a target

$$u \in \{0, 1\}^M$$

such that

$$\text{dist}_2(u, L(B_H))^2 = W(H, b).$$

More generally, for every real $p \geq 1$, the same basis and target satisfy

$$\min_{\lambda \in L(B_H)} \|u - \lambda\|_p^p = W(H, b).$$

The lattice rank is exactly M .

Proof. Use Gaussian elimination over $\mathbb{F}_2$ to compute a particular solution u of $Hu = b$, a systematic form of $\ker_{\mathbb{F}_2} H$, and the basis B_H from Lemma 6. Restore the original coordinate order after constructing the basis. This takes polynomial time, and

$$L(B_H) = \Lambda_H.$$

Let $\lambda \in L(B_H)$ and put

$$x = (u - \lambda) \bmod 2.$$

Since $H(\lambda \bmod 2) = 0$, we have $Hx = b$. Each nonzero coordinate of x comes from an odd integer coordinate of $u - \lambda$, whose absolute value is at least one. Therefore

$$\|u - \lambda\|_p^p = \sum_{j=1}^M |u_j - \lambda_j|^p \geq \text{wt}(x) \geq W(H, b).$$

Conversely, let x be any binary solution of $Hx = b$ and let $\tilde{x} \in \{0, 1\}^M$ be its integer lift. Then

$$H((u - \tilde{x}) \bmod 2) = b - b = 0,$$

so $u - \tilde{x} \in \Lambda_H$. Choosing $\lambda = u - \tilde{x}$ gives

$$\|u - \lambda\|_p^p = \|\tilde{x}\|_p^p = \text{wt}(x).$$

Taking minima over binary solutions proves the ℓ_p identity. Substituting $p = 2$ gives

$$\text{dist}_2(u, L(B_H))^2 = W(H, b).$$

Since B_H is a nonsingular $M \times M$ integer matrix, the lattice rank is M . □

Consequently, a binary weight gap

$$W(H, b) \leq R \quad \text{versus} \quad W(H, b) > \Gamma R$$

becomes the closest-vector gap

$$\min_{\lambda \in \Lambda_H} \|u - \lambda\|_p \leq R^{1/p} \quad \text{versus} \quad \min_{\lambda \in \Lambda_H} \|u - \lambda\|_p > \Gamma^{1/p} R^{1/p}.$$

For $p = 2$, the approximation factor is $\sqrt{\Gamma}$. These are real-valued distance thresholds, which need not themselves be rational. The reduction also requires a particular solution u , and therefore applies only to consistent affine systems.

For example, take

$$H = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}, \quad b = 1.$$

The homogeneous code and its parity-lift lattice are

$$\mathcal{C} = \{000, 011, 101, 110\},$$

$$\Lambda_H = \{z \in \mathbb{Z}^3 : z_1 + z_2 + z_3 \equiv 0 \pmod{2}\}.$$

Here $k_0 = 2$, and a systematic-form matrix and lattice basis are

$$P_{\text{sys}} = \begin{pmatrix} 1 & 1 \end{pmatrix}, \quad B_H = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 1 & 2 \end{pmatrix}.$$

Choose $u = (1, 0, 0)$. The binary solutions of $Hx = 1$ are the odd-parity vectors, and their minimum weight is one. Correspondingly,

$$\text{dist}_2(u, \Lambda_H) = 1.$$

3 Encoding Boolean assignments by Reed–Solomon codes

Let φ be a 3SAT formula of encoding length s . The purpose of this section is to encode its Boolean assignments and clause constraints in a binary affine system whose structure is controlled by Reed–Solomon codes.

A single interpolation polynomial first records the Boolean assignment. Binary indicators then encode its values at the evaluation points, and clause-indexed copies of those indicators represent possible satisfying local assignments. Low-degree moment constraints tie these tables together without leaving the binary affine setting.

Delete tautological clauses and repeated occurrences of a literal. A formula containing an empty clause is trivially unsatisfiable, whereas one with no remaining clauses is trivially satisfiable; handle these cases separately. Otherwise delete unused variables. Write

$$v_1, \dots, v_m \quad \text{and} \quad C_1, \dots, C_\ell$$

for the remaining variables and clauses, where $m, \ell \geq 1$. Write $[m] := \{1, \dots, m\}$ for the set of variable indices.

3.1 Assignments as interpolation polynomials

Let

$$\mathbb{K} = \mathbb{F}_q$$

be a characteristic-two field large enough to contain distinct elements

$$a_1, \dots, a_m \in \mathbb{K}.$$

We call a_i the *anchor* of the Boolean variable v_i . Identify the Boolean values 0 and 1 with the corresponding elements of $\mathbb{K}$.

A Boolean assignment

$$\sigma = (\sigma_1, \dots, \sigma_m) \in \{0, 1\}^m$$

determines a unique interpolation polynomial

$$Q_\sigma \in \mathbb{K}[X], \quad \deg Q_\sigma \leq m - 1, \quad Q_\sigma(a_i) = \sigma_i \quad (1 \leq i \leq m).$$

Explicitly, Lagrange interpolation gives

$$Q_\sigma(X) = \sum_{i=1}^m \sigma_i \prod_{\substack{1 \leq h \leq m \\ h \neq i}} \frac{X - a_h}{a_i - a_h}.$$

All denominators are nonzero because the anchors are distinct. Conversely, any polynomial of degree at most $m - 1$ taking values in $\{0, 1\}$ at the anchors recovers a Boolean assignment by evaluation.

Let

$$P = \mathbb{K} \setminus \{a_1, \dots, a_m\}$$

be the set of non-anchor evaluation points. The vector

$$(Q_\sigma(p))_{p \in P} \in \text{RS}(P, m - 1)$$

is a Reed–Solomon encoding of the assignment polynomial. Excluding the anchors ensures that

$$p - a_i \neq 0 \quad (p \in P, 1 \leq i \leq m),$$

so these differences can be used as denominators. We write

$$d = m$$

for a convenient upper bound on the degree of an assignment polynomial.

Encode each polynomial evaluation by one binary coordinate for each field value: the coordinate for the actual value is one and all others are zero. This one-hot array is the *global evaluation table*.

3.2 Clause assignments and evaluation tables

For each clause, introduce one additional evaluation table for each of its satisfying local assignments.

For each clause C , let

$$I_C \subseteq [m]$$

be the indices of the variables occurring in C , and define

$$\mathcal{B}_C = \{\beta \in \{0, 1\}^{I_C} : \beta \text{ satisfies } C\}.$$

Thus $\mathcal{B}_C$ consists of all satisfying local assignments to the variables in C . Since a clause has at most three variables,

$$|I_C| \leq 3, \quad |\mathcal{B}_C| \leq 8.$$

If a global assignment σ satisfies C , then its restriction

$$\sigma|_{I_C} \in \mathcal{B}_C$$

selects one satisfying local assignment.

Index the global table and all clause-indexed copies by

$$\Theta = \{0\} \cup \{(C, \beta) : C \text{ a clause, } \beta \in \mathcal{B}_C\}.$$

We call an element of this index set a *table type*. Type 0 represents the global assignment polynomial. A type (C, β) represents a candidate local assignment for clause C and is called a *clause subtype*.

For a satisfying assignment σ , declare the global type 0 and, for each clause C , the clause subtype $(C, \sigma|_{I_C})$ to be *active*. Every other clause subtype is *inactive*. An active table records the one-hot evaluations of the assignment polynomial; an inactive table is identically zero.

For every table type, evaluation point, and candidate field value,

$$\tau \in \Theta, \quad p \in P, \quad w \in \mathbb{K},$$

introduce a binary indicator

$$x_{\tau,p,w} \in \mathbb{F}_2.$$

For a fixed type τ and point p , the coordinates

$$(x_{\tau,p,w})_{w \in \mathbb{K}}$$

form the *evaluation fiber* at p . For a satisfying assignment, a fiber of an active type is an *active fiber*; it has exactly one nonzero indicator, namely

$$x_{\tau,p,Q_\sigma(p)} = 1.$$

A general binary table can also have fibers containing several nonzero coordinates.

3.3 Reed–Solomon moment constraints

For each evaluation fiber, record the power sums of its selected field values.

Fix a moment budget $T \geq 0$ and assume

$$dT < |P|.$$

For $0 \leq j \leq T$, use the convention $w^0 = 1$, including when $w = 0$. Define the ordinary moments

$$\mu_{\tau,j}(p) = \sum_{w \in \mathbb{K}} x_{\tau,p,w} w^j.$$

If a fiber is supported only at $w = Q_\sigma(p)$, then

$$\mu_{\tau,j}(p) = Q_\sigma(p)^j.$$

As p varies, these are evaluations of the polynomial

$$Q_\sigma(X)^j, \quad \deg Q_\sigma(X)^j \leq (m-1)j \leq dj.$$

Therefore the ordinary moments should form a Reed–Solomon codeword of degree at most dj .

For a clause-indexed table, also form shifted power sums by subtracting each prescribed local Boolean value and dividing by the corresponding nonzero anchor difference.

For a clause type $\tau = (C, \beta)$ and an index $i \in I_C$, define the shifted moments

$$\eta_{(C,\beta),i,j}(p) = \sum_{w \in \mathbb{K}} x_{(C,\beta),p,w} \left(\frac{w - \beta_i}{p - a_i} \right)^j.$$

The denominator is nonzero because $p \in P$. If $\beta_i = \sigma_i$, then

$$Q_\sigma(a_i) = \beta_i,$$

and the factor theorem yields

$$\frac{Q_\sigma(X) - \beta_i}{X - a_i} \in \mathbb{K}[X].$$

When $m \geq 2$, this quotient has degree at most $m - 2$; when $m = 1$, it is zero. Thus an active clause fiber has shifted moments

$$\eta_{(C,\beta),i,j}(p) = \left(\frac{Q_\sigma(p) - \beta_i}{p - a_i} \right)^j,$$

which form a Reed–Solomon codeword of degree at most $(d - 1)j$.

Impose the following four groups of constraints:

(C1) For every $p \in P$,

$$\sum_{w \in \mathbb{K}} x_{0,p,w} = 1 \quad \text{in } \mathbb{F}_2.$$

Thus every global fiber has odd parity.

(C2) For every clause C , $p \in P$, and $w \in \mathbb{K}$,

$$x_{0,p,w} = \sum_{\beta \in \mathcal{B}_C} x_{(C,\beta),p,w} \quad \text{in } \mathbb{F}_2.$$

Thus the local fibers for C reproduce the global fiber modulo two.

(C3) For every $\tau \in \Theta$ and $0 \leq j \leq T$,

$$(\mu_{\tau,j}(p))_{p \in P} \in \text{RS}(P, dj).$$

(C4) For every $\tau = (C, \beta)$, $i \in I_C$, and $0 \leq j \leq T$,

$$(\eta_{\tau,i,j}(p))_{p \in P} \in \text{RS}(P, (d - 1)j).$$

In particular, condition (C4) uniquely determines the *shifted-moment interpolation polynomial*

$$\begin{aligned} \eta_{\tau,i,j}(X) &\in \mathbb{K}[X], \quad \deg \eta_{\tau,i,j} \leq (d - 1)j, \\ \eta_{\tau,i,j}(p) &= \sum_{w \in \mathbb{K}} x_{\tau,p,w} \left(\frac{w - \beta_i}{p - a_i} \right)^j \quad (p \in P). \end{aligned} \tag{2}$$

Uniqueness follows from $(d - 1)j < |P|$.

Although the moments involve powers of field elements, their unknowns are the binary indicators $x_{\tau,p,w}$; every field element multiplying an indicator is a known coefficient. Moreover, by Section 2, membership in a Reed–Solomon code is given by linear parity checks over $\mathbb{K}$. Expanding these checks in an $\mathbb{F}_2$ -basis of $\mathbb{K}$, all four families together form an explicit binary affine system

$$Hx = b \quad \text{over } \mathbb{F}_2. \tag{3}$$

3.4 Completeness of the encoding

The one-hot graph of a satisfying assignment, together with the clause copies indexed by its satisfying local restrictions, has the following completeness guarantee.

Lemma 8 (Satisfying assignments give low-weight binary solutions). *If φ is satisfiable, system (3) has a binary solution of Hamming weight*

$$R = (\ell + 1)|P|.$$

Proof. Let

$$\sigma = (\sigma_1, \dots, \sigma_m) \in \{0, 1\}^m$$

be a satisfying assignment and let Q_σ be its interpolation polynomial. For each $p \in P$, set

$$x_{0,p,Q_\sigma(p)} = 1$$

and set all other coordinates of the global fiber to zero. For each clause C , select its satisfying restriction

$$\beta = \sigma|_{I_C} \in \mathcal{B}_C.$$

Set

$$x_{(C,\beta),p,Q_\sigma(p)} = 1 \quad (p \in P)$$

and set all other clause-type coordinates to zero.

Each global fiber has exactly one nonzero coordinate, so condition (C1) holds. For each clause, exactly one subtype reproduces the global fiber, proving condition (C2). The ordinary moments of an active type are

$$Q_\sigma(p)^j,$$

and those of an inactive type are zero. Since

$$\deg Q_\sigma(X)^j \leq (m-1)j \leq dj,$$

condition (C3) follows. For an active subtype and $i \in I_C$, the shifted moments are the evaluations of

$$\left(\frac{Q_\sigma(X) - \sigma_i}{X - a_i} \right)^j.$$

The quotient has degree at most $m-2$ when $m \geq 2$ and is zero when $m = 1$. Its powers therefore satisfy condition (C4); at $j = 0$, the active shifted moment is the constant one. Finally, there are exactly $\ell + 1$ active coordinates at each evaluation point: one global coordinate and one coordinate for each clause. Hence

$$\text{wt}(x) = (\ell + 1)|P| = R. \quad \square$$

3.5 Parameter choices and encoding size

Fix the parameters and bound the size of the resulting binary affine system. Let

$$N = 100 + s + m + \ell, \quad d = m, \quad K = N^4, \quad T = N^{30}, \quad c = 1/400.$$

Choose the least integer e satisfying

$$2^e \geq N^{200},$$

and define

$$q = 2^e, \quad N^{200} \leq q < 2N^{200}, \quad \mathbb{K} = \mathbb{F}_q.$$

This field can be constructed deterministically. Enumerate the monic binary polynomials of degree

$$e = O(\log N)$$

and apply the standard polynomial-time irreducibility test until an irreducible polynomial is found. Even exhaustive enumeration costs at most

$$2^e = N^{O(1)}.$$

Since $q \geq N^{200} > m$, distinct anchors $a_1, \dots, a_m \in \mathbb{K}$ exist. Their complement satisfies

$$|P| = q - m, \quad dT \leq N^{31} < q - m = |P|,$$

so every Reed–Solomon degree bound above is strictly smaller than the number of evaluation points.

The type count is bounded by

$$|\Theta| = 1 + \sum_C |\mathcal{B}_C| \leq 1 + 8\ell.$$

There is one binary indicator for each triple (τ, p, w) , so the dimension of the affine system is

$$M = |\Theta| |P| q \leq (1 + 8\ell) q^2 \leq 40N^{401}. \quad (4)$$

Each $\mathbb{K}$ -linear Reed–Solomon parity check becomes

$$e = [\mathbb{K} : \mathbb{F}_2]$$

binary equations, and each code membership condition requires at most $|P|$ field-valued checks. Furthermore,

$$\sum_C |\mathcal{B}_C| |I_C| \leq 3(|\Theta| - 1).$$

Therefore, conditions (C1) through (C4) produce at most

$$|P| + \ell |P| q + 4|\Theta|(T + 1)|P|e = O(N^{401}) \quad (5)$$

binary equations. Here

$$|\Theta| \leq 9N, \quad |P| < q < 2N^{200}, \quad T = N^{30}, \quad e = O(\log N).$$

Together, (4) and (5) show that the dense binary encoding of H has at most

$$O(N^{802})$$

bits. In particular, the complete affine system is constructible in deterministic polynomial time.

4 Algebraic reconstruction of bounded polynomial-moment families

We isolate the following algebraic inverse problem: when do bounded, unordered sets whose power sums agree with low-degree polynomials admit a single global description? We show that these power sums are also the power sums of the roots of a single polynomial, independently of Boolean assignments, clauses, and binary affine systems.

4.1 Bounded moment families and algebraic root sets

Let $\mathbb{K}$ be a finite field and $P \subseteq \mathbb{K}$ a set of evaluation points. At each $p \in P$, we are given an unordered set $S(p) \subseteq \mathbb{K}$, forming the family

$$\mathcal{S} = (S(p))_{p \in P}.$$

The set $S(p)$ is called the *fiber at p* . There is no prescribed correspondence between the elements at different points, and the sizes of the sets may vary.

Each set must have at most K elements, and its power sums up to a moment budget T must agree with low-degree polynomials, even though its individual elements need not vary polynomially with p . To formulate this condition, write

$$m_j(p) := \sum_{w \in S(p)} w^j$$

for the j -th moment of $S(p)$, using the convention $w^0 = 1$ even when $w = 0$. In particular, $m_0(p)$ is $|S(p)|$ interpreted in $\mathbb{K}$; in characteristic two, it records the parity of the set size, rather than its integer value. The following definition makes the two regularity conditions precise.

Definition 9 (Bounded polynomial-moment family). Let $d, K \in \mathbb{Z}_{\geq 1}$, let $T \in \mathbb{Z}_{\geq 0}$ satisfy $dT < |P|$. A family $\mathcal{S} = (S(p))_{p \in P}$ is a (d, T, K) -bounded polynomial-moment family if it has the following two properties:

(M1) Uniformly bounded sets. Every pointwise set has at most K elements:

$$|S(p)| \leq K \quad (p \in P).$$

(M2) Low-degree polynomial moments. For every $0 \leq j \leq T$, there is a polynomial

$$\mu_j(X) \in \mathbb{K}[X], \quad \deg \mu_j \leq dj,$$

satisfying

$$\mu_j(p) = m_j(p) \quad (p \in P).$$

The polynomials $\mu_0, \dots, \mu_T$ are called the *moment polynomials* of the family.

Since $dj \leq dT < |P|$, the polynomial in condition (M2) is unique: two degree-at-most- dj polynomials agreeing on P must coincide. Condition (M1) holds at every point indexing the family and permits empty sets.

The global structure to be extracted. To see the kind of global structure we seek, first consider K low-degree polynomials

$$q_1(X), \dots, q_K(X) \in \mathbb{K}[X], \quad \deg q_t \leq d.$$

Suppose that their evaluations are pairwise distinct at every $p \in P$, and define

$$S(p) = \{q_1(p), \dots, q_K(p)\}.$$

Then the moment polynomials are simply

$$\mu_j(X) = \sum_{t=1}^K q_t(X)^j, \quad \deg \mu_j \leq dj.$$

In particular, evaluating μ_j at p gives the j -th moment of $S(p)$. The entire family is explained by the single polynomial

$$G(Y) = \prod_{t=1}^K (Y - q_t(X)) \in \mathbb{K}[X][Y].$$

Its roots are the global polynomials $q_1(X), \dots, q_K(X)$; at each p , evaluating its coefficients recovers the polynomial whose roots are the elements of $S(p)$.

Unlike this example, a bounded polynomial-moment family need not arise from globally labeled polynomials q_t . Nevertheless, its moment polynomials still yield an analogous global description. Set

$$F = \mathbb{K}(X), \quad h = \max_{p \in P} |S(p)| \leq K.$$

The desired global description consists of a monic separable polynomial

$$G(Y) \in F[Y], \quad \deg_Y G = h.$$

Its distinct roots should form, in a finite separable extension E/F , the root set

$$\mathcal{R} = \{\alpha \in E : G(\alpha) = 0\}, \quad |\mathcal{R}| = h,$$

whose power sums reproduce the original moment polynomials:

$$\sum_{\alpha \in \mathcal{R}} \alpha^j = \mu_j(X) \quad (0 \leq j \leq T).$$

The roots play the role of $q_1(X), \dots, q_K(X)$ in the example, but may be algebraic rather than polynomial or rational functions. Moreover, wherever $|S(p)| = h$, evaluating the coefficients of G at p recovers the monic polynomial whose roots are exactly the elements of $S(p)$.

Encoding a pointwise set without labeling it. Recovering a polynomial from its moment sequence follows a classical method of Massey [Mas69]. Here the moments additionally depend on the evaluation point.

The natural polynomial associated with a particular set $S(p)$ is

$$\prod_{w \in S(p)} (Y - w) \in \mathbb{K}[Y].$$

This product is monic, has degree $|S(p)|$, and vanishes precisely at the elements of $S(p)$. Its coefficients are symmetric functions of those elements: they do not change when the elements are reordered. Consequently, no ordering or labeling is required. If $S(p) = \emptyset$, the empty product is 1.

To explain how the moments determine this product, suppose that $h > 0$ and consider a point satisfying $|S(p)| = h$. Write

$$\prod_{w \in S(p)} (Y - w) = Y^h + \sum_{l=0}^{h-1} g_{p,l} Y^l, \quad g_{p,l} \in \mathbb{K}.$$

Since the product vanishes at every $w \in S(p)$, multiplying by w^i and summing gives

$$\begin{aligned} 0 &= \sum_{w \in S(p)} w^i \left(w^h + \sum_{l=0}^{h-1} g_{p,l} w^l \right) \\ &= m_{i+h}(p) + \sum_{l=0}^{h-1} g_{p,l} m_{i+l}(p) \quad (0 \leq i < h). \end{aligned}$$

Thus the coefficients $g_{p,l}$ satisfy a linear system whose entries are the moments. The coefficient matrix is a Hankel matrix: its entries depend only on the sum of their row and column indices. Here it is

$$(m_{i+l}(p))_{0 \leq i, l < h}.$$

Because $S(p)$ consists of h distinct elements, this matrix factors as a Vandermonde matrix times its transpose. The Vandermonde matrix lists successive powers of those distinct elements and is therefore invertible. Hence the moments determine the entire product polynomial, without identifying any individual element.

Constructing one global polynomial. The pointwise Hankel systems are linked by a simple observation: each entry $m_j(p)$ is the evaluation of the same moment polynomial $\mu_j(X)$. Consequently, all the pointwise Hankel matrices arise by specializing one polynomial matrix,

$$(\mu_{i+l}(X))_{0 \leq i, l < h}.$$

Solve one linear system with X left as an indeterminate:

$$\sum_{l=0}^{h-1} \mu_{i+l}(X) c_l(X) = -\mu_{i+h}(X) \quad (0 \leq i < h).$$

At any point where $|S(p)| = h$, the polynomial matrix specializes to the invertible pointwise Hankel matrix. Its determinant is therefore not identically zero. Hence the global system has a unique solution over the rational function field $F = \mathbb{K}(X)$. Division by the determinant may be necessary, which explains why its solutions need not belong to $\mathbb{K}[X]$.

The coefficient functions obtained from this one system define a single monic polynomial

$$G(Y) = Y^h + \sum_{l=0}^{h-1} c_l(X) Y^l \in F[Y].$$

At a point p where $|S(p)| = h$, the pointwise matrix is invertible. Cramer's rule expresses the rational coefficients $c_l(X)$ using the determinant of the global matrix as a common denominator. At p , that denominator becomes the nonzero determinant of the pointwise matrix, so the coefficients can be evaluated. Their evaluations solve the original pointwise system, and uniqueness therefore gives

$$\text{ev}_p(G)(Y) := Y^h + \sum_{l=0}^{h-1} c_l(p) Y^l = \prod_{w \in S(p)} (Y - w).$$

Thus one global polynomial recovers the maximum-size pointwise sets.

Why the roots may require an extension. Although G has coefficients in $F = \mathbb{K}(X)$, it need not split over F . Its individual roots therefore need not be rational functions of X . Instead, pass to a splitting field E/F and write

$$\mathcal{R} = \{\alpha \in E : G(\alpha) = 0\}.$$

These roots are nevertheless algebraic functions of X . Indeed, if $D(X) \in \mathbb{K}[X]$ is a nonzero common denominator of the coefficients of G , then

$$\tilde{G}(X, Y) := D(X)G(Y) \in \mathbb{K}[X, Y], \quad \tilde{G}(X, \alpha) = 0 \quad (\alpha \in \mathcal{R}).$$

Thus each root satisfies a polynomial equation over $\mathbb{K}[X]$, even though it need not belong to $\mathbb{K}(X)$.

4.2 The algebraic reconstruction lemma

The following lemma gives the precise reconstruction statement.

Lemma 10 (Algebraic reconstruction of bounded polynomial moments). *Let $\mathbb{K}$ be a finite field, let $P \subseteq \mathbb{K}$, and let*

$$\mathcal{S} = (S(p))_{p \in P}, \quad S(p) \subseteq \mathbb{K},$$

be a (d, T, K) -bounded polynomial-moment family with moment polynomials $\mu_0, \dots, \mu_T \in \mathbb{K}[X]$. In other words, assume that

$$\deg \mu_j \leq dj, \quad \mu_j(p) = \sum_{w \in S(p)} w^j \quad (p \in P, 0 \leq j \leq T),$$

and

$$|S(p)| \leq K \quad (p \in P).$$

Put $F = \mathbb{K}(X)$, and suppose that

$$T \geq 2K - 1, \quad |P| - dK(K - 1) > 2dK^2T.$$

Define the maximum fiber size

$$h := \max_{p \in P} |S(p)| \leq K$$

and its Hankel determinant

$$\Delta_h(X) := \begin{cases} 1, & h = 0, \\ \det(\mu_{i+j}(X))_{0 \leq i, j < h}, & h > 0. \end{cases}$$

Then $\Delta_h \neq 0$, and there exist a monic separable polynomial

$$G(Y) = Y^h + \sum_{l=0}^{h-1} c_l Y^l \in F[Y],$$

and a finite separable splitting field E/F of G . Write

$$\mathcal{R} = \{\alpha_1, \dots, \alpha_h\} = \{\alpha \in E : G(\alpha) = 0\},$$

so that these objects have the following properties:

(i) Moment identities. For every $0 \leq j \leq T$,

$$\mu_j(X) = \sum_{\alpha \in \mathcal{R}} \alpha^j. \quad (6)$$

(ii) Maximum-size fibers. For every $p \in P$,

$$\Delta_h(p) \neq 0 \iff |S(p)| = h.$$

In particular,

$$\#\{p \in P : |S(p)| = h\} \geq |P| - dh(h - 1).$$

(iii) Coefficients and discriminant. The polynomial Δ_h is nonzero and satisfies

$$\deg \Delta_h \leq dh(h - 1).$$

If $h > 0$, the coefficients $c_0, \dots, c_{h-1} \in F$ are the unique solution of

$$\sum_{l=0}^{h-1} \mu_{i+l}(X) c_l = -\mu_{i+h}(X) \quad (0 \leq i < h). \quad (7)$$

For every $0 \leq l < h$, there is an $n_l \in \mathbb{K}[X]$ such that

$$c_l = \frac{n_l}{\Delta_h}, \quad \deg n_l \leq d(h^2 - l) \leq dh^2.$$

Furthermore,

$$\Delta_h = \prod_{1 \leq r < s \leq h} (\alpha_s - \alpha_r)^2. \quad (8)$$

(iv) Specialization. For every $p \in P$ satisfying $|S(p)| = h$, one has $\Delta_h(p) \neq 0$. Thus the coefficientwise specialization

$$\text{ev}_p(G)(Y) := Y^h + \sum_{l=0}^{h-1} c_l(p)Y^l \in \mathbb{K}[Y]$$

is defined and satisfies

$$\text{ev}_p(G)(Y) = \prod_{w \in S(p)} (Y - w). \quad (9)$$

When $h = 0$, the empty-sum and empty-product conventions give $G = 1$, $\mathcal{R} = \emptyset$, $\Delta_0 = 1$, and $\mu_j = 0$ for every $0 \leq j \leq T$. The conclusions hold in every characteristic, including characteristic two.

Proof. If $h = 0$, then $S(p) = \emptyset$ for every $p \in P$. Consequently μ_j vanishes at all $|P|$ evaluation points. Since

$$\deg \mu_j \leq dT < |P| \quad (0 \leq j \leq T),$$

the polynomial root bound gives $\mu_j = 0$ for every j . Taking $G = 1$, $E = F$, and $\mathcal{R} = \emptyset$ proves all the conclusions.

Suppose henceforth that $h > 0$. Since $2h - 2 \leq 2K - 2 \leq T$, all moments in Δ_h are defined. Expansion of its determinant and the inequalities $\deg \mu_{i+j} \leq d(i+j)$ give

$$\deg \Delta_h \leq dh(h-1). \quad (10)$$

At any point with

$$S(p) = \{w_1, \dots, w_h\},$$

the moment identities give

$$(\mu_{i+j}(p))_{0 \leq i, j < h} = V_p V_p^T, \quad (V_p)_{i,s} = w_s^i.$$

Since the w_s are distinct, V_p is an invertible Vandermonde matrix, and therefore

$$\Delta_h(p) = \det(V_p)^2 = \prod_{1 \leq s < t \leq h} (w_t - w_s)^2 \neq 0. \quad (11)$$

Because h is the maximum fiber size, at least one such point exists. Hence Δ_h is a nonzero polynomial.

At every point with $|S(p)| < h$, the size- h moment matrix factors as

$$(\mu_{i+j}(p))_{0 \leq i, j < h} = W_p W_p^T, \quad (W_p)_{i,w} = w^i \quad (0 \leq i < h, w \in S(p)).$$

The matrix W_p has fewer than h columns, so its product has rank less than h , and $\Delta_h(p) = 0$. By (10), the nonzero polynomial Δ_h has at most $dh(h-1)$ zeros. Since every fiber has size at most h , it follows that

$$\#\{p \in P : |S(p)| = h\} \geq |P| - dh(h-1) \geq |P| - dK(K-1) > 2dK^2T. \quad (12)$$

Since $\Delta_h \neq 0$, the Hankel system (7) has a unique solution over F . For that solution, set

$$G(Y) = Y^h + \sum_{l=0}^{h-1} c_l Y^l.$$

Cramer's rule gives a common denominator Δ_h . Replacing column l by $(\mu_{i+h})_{0 \leq i < h}$ gives a numerator of degree at most

$$d(h(h-1) + h - l) = d(h^2 - l) \leq dh^2.$$

Fix any point satisfying

$$S(p) = \{w_1, \dots, w_h\}.$$

Equation (11) gives $\Delta_h(p) \neq 0$, so the coefficients c_l have no poles at p . Specializing (7) therefore gives, for $0 \leq i < h$,

$$\begin{aligned} 0 &= \mu_{i+h}(p) + \sum_{l=0}^{h-1} \mu_{i+l}(p) c_l(p) \\ &= \sum_{s=1}^h w_s^i \left(w_s^h + \sum_{l=0}^{h-1} c_l(p) w_s^l \right) \\ &= \sum_{s=1}^h w_s^i \operatorname{ev}_p(G)(w_s). \end{aligned}$$

Equivalently,

$$V_p \begin{pmatrix} \operatorname{ev}_p(G)(w_1) \\ \vdots \\ \operatorname{ev}_p(G)(w_h) \end{pmatrix} = 0.$$

Since V_p is invertible, every w_s is a root of $\operatorname{ev}_p(G)(Y)$. Both sides of (9) are monic of degree h , so the claimed specialization follows.

Let $\alpha_1, \dots, \alpha_h$ be the roots of G in a splitting field, initially listed with multiplicity. Newton's identities express their j -th power sum as a polynomial

$$\mathcal{P}_j(c_0, \dots, c_{h-1})$$

of total degree at most j . For $j > h$, use the recurrence given by the monic equation $G(\alpha_s) = 0$. Crucially, this direction of Newton's identities computes power sums from coefficients without dividing by j ; it remains valid in characteristic two. For $1 \leq j \leq T$, the expression

$$\Delta_h^j(\mu_j - \mathcal{P}_j(c_0, \dots, c_{h-1}))$$

is a polynomial of degree at most

$$2dh^2j \leq 2dK^2T.$$

By (9), it vanishes at every point with $|S(p)| = h$. There are strictly more than $2dK^2T$ such points by (12), so the polynomial is identically zero. Consequently,

$$\mu_j(X) = \sum_{s=1}^h \alpha_s^j \quad (1 \leq j \leq T).$$

For $j = 0$, the polynomial μ_0 is constant. At any point with $|S(p)| = h$, it equals

$$\mu_0(p) = h \cdot 1 = \sum_{s=1}^h \alpha_s^0 \quad \text{in } \mathbb{K},$$

which proves the remaining moment identity without mistaking parity for integer cardinality.

Since $2h - 2 \leq T$, the root moment identities give

$$\Delta_h = \det(V_\alpha V_\alpha^\top) = \det(V_\alpha)^2, \quad (V_\alpha)_{i,s} = \alpha_s^i.$$

As $\Delta_h \neq 0$, the roots $\alpha_1, \dots, \alpha_h$ are distinct. Therefore G is separable and its root set $\mathcal{R}$ satisfies both (6) and (8). $\square$

Relation to Reed–Solomon list reconstruction and recovery. Ar, Lipton, Rubinfeld, and Sudan studied reconstruction of several algebraic functions from unordered, mixed evaluations [ALRS98]. Reed–Solomon list recovery has the same pointwise input—small sets $S(p)$ —but usually seeks all low-degree polynomials $f \in \mathbb{K}[X]$ satisfying $f(p) \in S(p)$ at sufficiently many points [Sud97, GS99, GR06]. Generalized key-equation and block-Hankel methods give a further technical connection [RR00, ZGA11]. Recent work sharpens the achievable agreement and output-list bounds [GST23, GLSTW24, LS25], and power sums have also appeared in Reed–Solomon decoding and lattice hardness [GGG18, BP23].

The additional hypothesis here is that the power sums of the entire pointwise sets are low-degree polynomials. For $h > 0$, the reconstruction produces the bivariate interpolation polynomial

$$Q(X, Y) := \Delta_h(X)G(Y) = \Delta_h(X)Y^h + \sum_{l=0}^{h-1} n_l(X)Y^l \in \mathbb{K}[X, Y].$$

At a maximum-size point, Q vanishes on $S(p)$ by specialization. At a smaller point, the Hankel matrix and each of its column replacements factor through a Vandermonde matrix with fewer than h columns, so $\Delta_h(p)$ and every $n_l(p)$ vanish. Hence

$$Q(p, w) = 0 \quad (p \in P, w \in S(p)).$$

Moreover, for every $f \in \mathbb{K}[X]$ of degree at most d , the coefficient bounds give

$$\deg_X Q(X, f(X)) \leq dh^2.$$

Therefore, if $f(p) \in S(p)$ at more than dh^2 points, the polynomial root bound gives $G(f) = 0$. Thus every such list-recovery candidate is one of at most h rational roots of G . The reconstruction also retains its possibly nonrational algebraic roots, together with the exact whole-fiber moment and specialization identities that remain valid in characteristic two.

5 Soundness of the nearest-codeword reduction

We now return to the binary affine system (3). Recall that its coordinate $x_{\tau,p,w}$ indicates whether the field element w is selected at evaluation point p in the table indexed by τ . The index $\tau = 0$ denotes the global table, while $\tau = (C, \beta)$ denotes the table for a satisfying local assignment $\beta \in \mathcal{B}_C$ to clause C .

We prove that a sufficiently low-weight binary solution determines a satisfying Boolean assignment.

5.1 Pointwise support sets of a low-weight solution

A support set consists of the field values with nonzero indicators at one evaluation point.

Suppose that a binary solution $x \in \mathbb{F}_2^M$ satisfies

$$Hx = b, \quad \text{wt}(x) \leq 4M^{1/200}R, \quad R = (\ell + 1)|P|.$$

For each $\tau \in \Theta$ and $p \in P$, define

$$S_\tau(p) = \{w \in \mathbb{K} : x_{\tau,p,w} = 1\}.$$

Every nonzero coordinate of x contributes to exactly one set $S_\tau(p)$. Therefore

$$\sum_{\tau \in \Theta} \sum_{p \in P} |S_\tau(p)| = \text{wt}(x) \leq 4M^{1/200}R.$$

For each table index, retain the evaluation points with bounded support sets:

$$P_\tau := \{p \in P : |S_\tau(p)| \leq K\}, \quad K = N^4.$$

This restriction affects only the analysis: the affine system and all ordinary-moment, shifted-moment, and clause constraints remain indexed by the original evaluation set P . Markov's inequality bounds the number of discarded evaluation points:

$$\begin{aligned} |P \setminus P_\tau| &= \#\{p \in P : |S_\tau(p)| > K\} \\ &\leq \frac{1}{K} \sum_{p \in P} |S_\tau(p)| \\ &\leq \frac{4M^{1/200}R}{N^4} \\ &\leq 4 \cdot 40^{1/200} q N^{-199/200} < \frac{q}{20} \quad (N \geq 100). \end{aligned}$$

Here we used $M \leq 40N^{401}$ and $R \leq Nq$. In particular, every type retains a large evaluation set:

$$|P_\tau| > |P| - \frac{q}{20} \quad (\tau \in \Theta). \quad (13)$$

The parameter choices also give

$$|P| - \frac{q}{20} - N^9 > q/2, \quad q/2 > 2N^{39}. \quad (14)$$

These estimates use $|P| = q - m$, $m \leq N$, and $q \geq N^{200}$. The sets P_τ may differ; no common retained evaluation set is required.

5.2 Applying the algebraic reconstruction lemma

Fix $\tau \in \Theta$. The ordinary-moment constraints apply to its support sets on the full evaluation set. Condition (C3) says that, for every $0 \leq j \leq T$, the vector

$$\left(\sum_{w \in S_\tau(p)} w^j \right)_{p \in P}$$

belongs to $\text{RS}(P, dj)$. Consequently there is a unique polynomial

$$\mu_{\tau,j}(X) \in \mathbb{K}[X], \quad \deg \mu_{\tau,j} \leq dj,$$

satisfying

$$\mu_{\tau,j}(p) = \sum_{w \in S_\tau(p)} w^j \quad (p \in P).$$

These polynomials come from the original constraints on P . Restricting the evaluation set does not change them. For each table, define the family of support sets

$$\mathcal{S}_\tau := (S_\tau(p))_{p \in P_\tau}.$$

Since $P_\tau \subseteq P$, the same polynomials satisfy

$$\mu_{\tau,j}(p) = \sum_{w \in S_\tau(p)} w^j \quad (p \in P_\tau),$$

which verifies condition (M2). By the definition of P_τ ,

$$|S_\tau(p)| \leq K \quad (p \in P_\tau),$$

which verifies condition (M1).

The parameters of Section 3 give

$$T = N^{30} \geq 2N^4 - 1 = 2K - 1,$$

and (13) and (14) give

$$\begin{aligned} |P_\tau| - dK(K-1) &> |P| - \frac{q}{20} - N^9 \\ &> q/2 \\ &> 2N^{39} \\ &\geq 2dK^2T. \end{aligned}$$

In particular,

$$dT \leq N^{31} < q/2 < |P_\tau|.$$

Thus $\mathcal{S}_\tau$ is a (d, T, K) -bounded polynomial-moment family on P_τ , and it satisfies every hypothesis of Lemma 10.

Apply that lemma separately to each $\mathcal{S}_\tau$. It gives a separable monic polynomial

$$G_\tau(Y) \in F[Y], \quad h_\tau = \max_{p \in P_\tau} |S_\tau(p)| = \deg_Y G_\tau \leq K, \quad F = \mathbb{K}(X).$$

Lemma 3 supplies one finite separable extension E/F in which all of these polynomials split. Write

$$\mathcal{R}_\tau = \{\alpha \in E : G_\tau(\alpha) = 0\}.$$

In particular, $\mathcal{R}_0$ corresponds to the table indexed by 0, and $\mathcal{R}_{(C,\beta)}$ corresponds to the table indexed by (C, β) . Each root set consists of h_τ distinct elements. The ordinary moment polynomials of the original pointwise sets are exactly the power sums of these roots:

$$\mu_{\tau,j}(X) = \sum_{\alpha \in \mathcal{R}_\tau} \alpha^j \quad (\tau \in \Theta, 0 \leq j \leq T). \quad (15)$$

The polynomials $\mu_{\tau,j}$ are defined on P ; only the application of Lemma 10 uses P_τ .

The global parity constraint and the zeroth root-moment identity imply the following.

Corollary 11. *The root set $\mathcal{R}_0$ of the global type is nonempty.*

Proof. Condition (C1) gives $\mu_{0,0} = 1$. Equation (15) with $j = 0$ therefore gives

$$1 = \sum_{\alpha \in \mathcal{R}_0} 1 = |\mathcal{R}_0| \cdot 1 \quad \text{in } \mathbb{K}.$$

Thus $|\mathcal{R}_0|$ is odd, and in particular $\mathcal{R}_0 \neq \emptyset$. □

5.3 Valuative recovery of local assignments

For the table indexed by $\tau = (C, \beta)$, the bit β_i is specified for each $i \in I_C$. The anchor a_i lies outside P . For roots algebraic over $\mathbb{K}(X)$, fix a valuation above $X - a_i$ for each variable, using the same valuation across all clause tables containing that variable.

Lemma 12 (Local valuation extraction). *Let $\tau = (C, \beta)$, $i \in I_C$, and let E/F be the common finite separable splitting field fixed in the preceding subsection. Fix once and for all an extension v_i of ord_{X-a_i} to E . For every $\alpha \in \mathcal{R}_\tau$,*

$$v_i(\alpha - \beta_i) \geq 1. \quad (16)$$

Proof. Write $\ell_i = X - a_i$, set $h = h_\tau$, and enumerate the distinct roots as $\alpha_1, \dots, \alpha_h$. There is nothing to prove if $h = 0$. For $h > 0$, abbreviate the Hankel determinant of this particular type by

$$\Delta_h(X) = \det(\mu_{\tau, r+s}(X))_{0 \leq r, s < h}.$$

Recall from (2) that $\eta_{\tau, i, j}(X)$ is the unique polynomial of degree at most $(d-1)j$ satisfying

$$\eta_{\tau, i, j}(p) = \sum_{w \in S_\tau(p)} \left(\frac{w - \beta_i}{p - a_i} \right)^j \quad (p \in P).$$

The pointwise binomial identity and uniqueness of Reed–Solomon interpolation imply, for $0 \leq j \leq T$,

$$\ell_i^j \eta_{\tau, i, j}(X) = \sum_{l=0}^j \binom{j}{l} (-\beta_i)^{j-l} \mu_{\tau, l}(X). \quad (17)$$

Both sides have degree at most $dj < |P|$ and agree on all $p \in P$, so this is an identity in $\mathbb{K}[X]$, not merely an agreement on the sampled grid.

Put

$$y_s = \frac{\alpha_s - \beta_i}{\ell_i}.$$

Combining (17) with (15) gives

$$\sum_{s=1}^h y_s^j = \eta_{\tau, i, j}(X) \quad (0 \leq j \leq T). \quad (18)$$

In particular, each displayed power sum has nonnegative v_i -valuation.

Write

$$G_\tau(Y) = Y^h + \sum_{l < h} c_l Y^l, \quad D_0 = dh^2 + h.$$

Lemma 10 gives $v_i(c_l) \geq -dh^2$. The y_s are the roots of

$$\ell_i^{-h} G_\tau(\ell_i Y + \beta_i) = Y^h + \sum_{l < h} b_l Y^l,$$

where $b_h = 1$ and, for $0 \leq l < h$,

$$b_l = \ell_i^{l-h} \left[\binom{h}{l} \beta_i^{h-l} + \sum_{k'=l}^{h-1} c_{k'} \binom{k'}{l} \beta_i^{k'-l} \right], \quad v_i(b_l) \geq -D_0.$$

If $y_s = 0$, then $v_i(y_s) = +\infty$, and the desired conclusion already holds for that root. Otherwise write $t_s = v_i(y_s) \in \mathbb{Q}$. Since $v_i(b_l) \geq -D_0$ for $l < h$, the monic equation implies

$$t_s \geq -D_0. \quad (19)$$

Indeed, if $t_s < -D_0$, then y_s^h would be the unique minimum-valuation term of the equation. If $t_s < 0$, the minimum of

$$v_i(b_l) + lt_s \quad (0 \leq l \leq h, b_l \neq 0)$$

must be attained by at least two indices. Thus there exist $0 \leq l < k' \leq h$ with $b_l, b_{k'} \in F^\times$ and

$$(k' - l)t_s = v_i(b_l) - v_i(b_{k'}) \in \mathbb{Z}.$$

Consequently

$$t_s < 0 \implies t_s \leq -1/h. \quad (20)$$

Let $V = (y_s^j)_{0 \leq j < h, 1 \leq s \leq h}$. By (8),

$$(\det V)^2 = \ell_i^{-h(h-1)} \Delta_h.$$

Since $\Delta_h \in \mathbb{K}[X] \setminus \{0\}$,

$$0 \leq v_i(\Delta_h) = \text{ord}_{X-a_i}(\Delta_h) \leq \deg \Delta_h \leq dh(h-1).$$

Thus V is invertible and

$$2v_i(\det V) = v_i(\Delta_h) - h(h-1), \quad v_i(\det V) \leq \frac{1}{2}(d-1)h(h-1) \leq dh^2.$$

By (19), each entry of V has valuation at least $-hD_0$, and every cofactor has valuation at least $-h^2D_0$. Therefore

$$v_i((V^{-1})_{s,j}) \geq -U_h, \quad U_h = h^2D_0 + dh^2 + 1 \leq 4N^{17}. \quad (21)$$

Choose

$$z = hU_h + 1.$$

The required moments exist since

$$z + h - 1 \leq 5N^{21} < T = N^{30}.$$

Equation (18) gives

$$V \begin{pmatrix} y_1^z \\ \vdots \\ y_h^z \end{pmatrix} = \begin{pmatrix} \eta_{\tau,i,z} \\ \eta_{\tau,i,z+1} \\ \vdots \\ \eta_{\tau,i,z+h-1} \end{pmatrix}.$$

The right-hand vector has nonnegative valuation. By (21),

$$v_i(y_s^z) \geq -U_h \quad (1 \leq s \leq h).$$

If $t_s < 0$, however, (20) gives

$$v_i(y_s^z) = zt_s \leq -\frac{hU_h + 1}{h} < -U_h,$$

a contradiction. Therefore $v_i(y_s) \geq 0$ for every root, which is precisely (16). $\square$

5.4 Clause matching

For each clause C , the sets $\mathcal{R}_{(C,\beta)}$ indexed by $\beta \in \mathcal{B}_C$ are called its *clause subtype root sets*. The clause constraint expresses every indicator of the table indexed by 0 as the sum modulo two of the corresponding indicators of the tables indexed by (C, β) . Together with the root-moment identities, this gives the following matching property.

Lemma 13 (Parity matching of clause roots). *For each clause C , every $\alpha \in \mathcal{R}_0$ belongs to $\mathcal{R}_{(C,\beta)}$ for at least one $\beta \in \mathcal{B}_C$.*

Proof. The pointwise clause constraint and Reed–Solomon interpolation give, for $0 \leq j \leq T$,

$$\mu_{0,j}(X) = \sum_{\beta \in \mathcal{B}_C} \mu_{(C,\beta),j}(X).$$

Substitute the reconstructed moments from (15) and move the right-hand side to the left. Since $\text{char } \mathbb{K} = 2$,

$$\sum_{\alpha \in \mathcal{R}_0} \alpha^j + \sum_{\beta \in \mathcal{B}_C} \sum_{\alpha \in \mathcal{R}_{(C,\beta)}} \alpha^j = 0.$$

Let W_C be the set of distinct roots occurring an odd number of times in these sets. Then

$$\sum_{\alpha \in W_C} \alpha^j = 0 \quad (0 \leq j \leq T), \quad |W_C| \leq (1 + |\mathcal{B}_C|)K \leq 9K < T + 1.$$

If $W_C \neq \emptyset$, the first $|W_C|$ equations form an invertible Vandermonde matrix applied to the all-ones vector, an impossibility. Thus $W_C = \emptyset$. A root appearing once in $\mathcal{R}_0$ must therefore occur at least once among the clause subtype root sets. $\square$

5.5 Completing the soundness proof

Fix $\alpha \in \mathcal{R}_0$. For each clause C , this element belongs to some $\mathcal{R}_{(C,\beta)}$; the resulting satisfying assignments agree on shared variables.

Proposition 14 (Binary nearest-codeword soundness). *If the affine system (3) has a binary solution x satisfying*

$$\text{wt}(x) \leq 4M^{1/200}R, \quad R = (\ell + 1)|P|,$$

then φ is satisfiable.

Proof. Fix such a binary solution. As verified above, its restricted support families satisfy Definition 9 and the numerical hypotheses of Lemma 10. Apply the reconstruction lemma separately to each family and take the common splitting field E . By Corollary 11, choose one $\alpha \in \mathcal{R}_0$. For every clause C , Lemma 13 supplies a satisfying local tuple $\beta(C)$ with

$$\alpha \in \mathcal{R}_{(C,\beta(C))}.$$

Fix an index i appearing in two clauses C, C' . Apply the same valuation v_i on the common splitting field in both applications of Lemma 12. This gives

$$v_i(\alpha - \beta(C)_i) \geq 1, \quad v_i(\alpha - \beta(C')_i) \geq 1.$$

If the two bits differed, the ultrametric inequality would imply

$$v_i(1) = v_i(\beta(C)_i - \beta(C')_i) \geq 1,$$

contradicting $v_i(1) = 0$. All locally selected tuples therefore agree on shared variables. Since unused variables were discarded, they assemble into one global Boolean assignment; each clause is satisfied by construction. $\square$

Together with Lemma 8, this soundness implication completes the binary nearest-codeword gap:

$$\varphi \text{ satisfiable} \implies W(H, b) \leq R,$$

and, whenever $Hx = b$ is consistent,

$$\varphi \text{ unsatisfiable} \implies W(H, b) > 4M^{1/200}R.$$

It remains to transfer this binary gap to CVP using Lemma 7.

6 Transfer to CVP and complexity

First handle the formula-preprocessing cases from the start of the construction. If no clauses remain, return $(B, t, r) = ((1), (0), 1)$, a one-dimensional YES instance. If an empty clause remains, return

$$B = (2), \quad t = (1), \quad r = \frac{1}{2}.$$

This one-dimensional instance has distance 1, hence is a strict NO instance because $1^c r = 1/2$.

For every other formula, construct the explicit affine system (3). Gaussian elimination over $\mathbb{F}_2$ tests whether the affine system is consistent; it does not decide whether the input formula is satisfiable.

If the affine system is inconsistent, return the same one-dimensional NO instance. By Lemma 8, inconsistency cannot occur for a satisfiable formula.

On the consistent branch, find any $u \in \{0, 1\}^M$ solving $Hu = b$. The parity-lift lattice has the square basis (1), with the inverse coordinate permutation restoring the original order. Its squared distance to this target equals the minimum binary weight. Output

$$(B, t, r) = (B_H, u, \lceil \sqrt{(\ell + 1)|P|} \rceil).$$

On this branch its dimension is $n = M$, its basis is nonsingular, and its target and radius are binary-encoded integers. Furthermore,

$$\text{dist}_2(u, L(B_H)) = \min_{v \in u - \Lambda_H} \|v\|_2.$$

For a satisfiable formula, Lemmas 8 and 7 give

$$\text{dist}_2(u, L(B_H)) \leq \sqrt{R} \leq r.$$

For an unsatisfiable formula, suppose for contradiction that

$$\text{dist}_2(u, L(B_H)) \leq M^c r.$$

Since $c = 1/400$, Lemma 7 gives

$$W(H, b) = \text{dist}_2(u, L(B_H))^2 \leq M^{1/200} r^2 \leq 4M^{1/200} R,$$

where $r = \lceil \sqrt{R} \rceil$ implies $r^2 \leq 4R$. Proposition 14 would then make φ satisfiable, a contradiction. Therefore

$$\text{dist}_2(u, L(B_H)) > M^c r.$$

The inequality is strict, as required by the promise definition.

The finite field has extension degree $e = O(\log N)$, and its deterministic construction uses $N^{O(1)}$ bit operations. Equations (4) and (5) bound both dimensions of H by $O(N^{401})$. Consequently, the dense binary system and the square output basis each have $O(N^{802})$ bits. Field arithmetic, Reed–Solomon parity checks, consistency testing, and basis construction are deterministic polynomial-time operations.

7 Further consequences

The same binary affine gap gives coding-theoretic and fixed-finite-norm consequences. In each statement, n denotes the output block length or lattice rank.

Recall the binary nearest-codeword and syndrome-decoding promise problems defined in Section 1.1.

Corollary 15 (Hardness of binary nearest codeword and syndrome decoding). *Binary nearest codeword and binary syndrome decoding are NP-hard to approximate within $n^{1/200}$ under deterministic polynomial-time many-one reductions.*

Proof. First suppose that (3) is consistent, compute a binary solution u , and put

$$\mathcal{C} = \ker_{\mathbb{F}_2} H, \quad d_H(u, \mathcal{C}) = \min\{\text{wt}(x) : x \in \mathbb{F}_2^n, Hx = b\}.$$

Use the integer radius $R = (\ell + 1)|P|$. Completeness follows directly from Lemma 8. Conversely, if a binary solution satisfies

$$\text{wt}(x) \leq n^{1/200} R,$$

then, since $n = M$, it also satisfies

$$\text{wt}(x) \leq 4M^{1/200} R.$$

Proposition 14 therefore implies that φ is satisfiable. Consequently an unsatisfiable formula gives the strict NO promise

$$d_H(u, \mathcal{C}) > n^{1/200} R.$$

Gaussian elimination computes either a generator for $\mathcal{C}$ or an independent parity-check matrix, with at most $O(n^2)$ output bits. For an inconsistent affine system, output the fixed binary code $\{00\}$, target 11, and radius 1: its distance 2 is strictly greater than $2^{1/200}$. By Lemma 8, this branch is never reached by a satisfiable formula. The same system and radius give the equivalent syndrome-decoding formulation. $\square$

For a fixed rational $p \geq 1$, define

$$\text{dist}_p(t, L) = \min_{z \in L} \|t - z\|_p.$$

The norm parameter p is fixed independently of the input; the lattice basis, target, and threshold constitute the promise instance.

For a parity-lift lattice, raising this distance to the norm exponent gives the minimum binary Hamming weight.

Corollary 16 (Hardness for fixed finite rational norms). *Fix a rational $p \geq 1$. There is a deterministic polynomial-time many-one reduction from 3SAT to the full-rank ℓ_p closest-vector promise problem with approximation factor*

$$n^{1/(200p)}.$$

The lattice basis is square and integral, and the target and radius have polynomial-size exact rational encodings.

Proof. Write $p = a/b$ in lowest terms and set the fixed integer $A = \lceil 4p \rceil$. For $R = (\ell + 1)|P|$, compute

$$j_p = \min\{j \in \mathbb{Z}_{\geq 0} : j^a \geq A^a R^b\}, \quad r_p = \frac{j_p}{A}.$$

Fixed-degree integer exponentiation and binary search compute this radius in deterministic polynomial time. Moreover,

$$R^{1/p} \leq r_p \leq (1 + 1/A)R^{1/p}, \quad r_p^p \leq e^{p/A} R \leq e^{1/4} R < 2R.$$

For a consistent system, apply Lemma 7 to obtain the same integral square basis B_H and binary target u , now equipped with the ℓ_p norm and radius r_p . If φ is satisfiable, completeness and the exact ℓ_p distance identity give

$$\text{dist}_p(u, L(B_H))^p = W(H, b) \leq R \leq r_p^p.$$

For soundness, suppose that

$$\text{dist}_p(u, L(B_H)) \leq n^{1/(200p)} r_p.$$

Lemma 7 then yields

$$W(H, b) = \text{dist}_p(u, L(B_H))^p \leq n^{1/200} r_p^p < 2n^{1/200} R \leq 4n^{1/200} R.$$

Since $n = M$, Proposition 14 would imply that φ is satisfiable. Thus an unsatisfiable input satisfies the strict promise

$$\text{dist}_p(u, L(B_H)) > n^{1/(200p)} r_p.$$

On an inconsistent system, the one-dimensional fixed output $(B, t, r) = ((2), (1), 1/2)$ has ℓ_p distance $1 > 1/2$ for every finite p , and is a strict NO instance. $\square$

References

- [AGMZ26] D. Aggarwal, R. Gupta, A. Morolia, and C. Zhang. Mind the gap? Not for SVP hardness under ETH! In *53rd International Colloquium on Automata, Languages, and Programming*, volume 374 of *Leibniz International Proceedings in Informatics*, pages 8:1–8:24, 2026. doi:10.4230/LIPIcs.ICALP.2026.8.
- [AR05] D. Aharonov and O. Regev. Lattice problems in $\text{NP} \cap \text{coNP}$. *Journal of the ACM*, 52(5):749–765, 2005. doi:10.1145/1089023.1089025.
- [Ajt96] M. Ajtai. Generating hard instances of lattice problems. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pages 99–108, 1996. doi:10.1145/237814.237838.
- [AKKV11] M. Alekhnovich, S. Khot, G. Kindler, and N. K. Vishnoi. Hardness of approximating the closest vector problem with pre-processing. *Computational Complexity*, 20(4):741–753, 2011. doi:10.1007/s00037-011-0031-3.
- [ALRS98] S. Ar, R. J. Lipton, R. Rubinfeld, and M. Sudan. Reconstructing algebraic functions from mixed data. *SIAM Journal on Computing*, 28(2):487–510, 1998. doi:10.1137/S0097539796297577.
- [ABSS97] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, 1997. doi:10.1006/jcss.1997.1472.
- [Bab86] L. Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. doi:10.1007/BF02579403.
- [BP23] H. Bennett and C. Peikert. Hardness of the (approximate) shortest vector problem: A simple proof via Reed–Solomon codes. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2023)*, volume 275 of *Leibniz International Proceedings in Informatics*, pages 37:1–37:20, 2023. doi:10.4230/LIPIcs.APPROX/RANDOM.2023.37.

- [BMVT78] E. R. Berlekamp, R. J. McEliece, and H. C. A. van Tilborg. On the inherent intractability of certain coding problems. *IEEE Transactions on Information Theory*, 24(3):384–386, 1978. doi:10.1109/TIT.1978.1055873.
- [BGLR25] V. Bhattiprolu, V. Guruswami, E. Lee, and X. Ren. Inapproximability of finding sparse vectors in codes, subspaces, and lattices. In *Proceedings of the 66th IEEE Annual Symposium on Foundations of Computer Science*, pages 1295–1303, 2025. doi:10.1109/FOCS63196.2025.00068.
- [BGR25] V. Bhattiprolu, V. Guruswami, and X. Ren. PCP-free APX-hardness of nearest codeword and minimum distance. *Electronic Colloquium on Computational Complexity*, Report TR25-029, 2025. <https://eccc.weizmann.ac.il/report/2025/029/>. Subsequently incorporated into [BGLR25].
- [DKRS03] I. Dinur, G. Kindler, R. Raz, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. *Combinatorica*, 23(2):205–243, 2003. doi:10.1007/s00493-003-0019-y.
- [DKS98] I. Dinur, G. Kindler, and S. Safra. Approximating-CVP to within almost-polynomial factors is NP-hard. In *Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, pages 99–111, 1998. doi:10.1109/SFCS.1998.743433.
- [FM04] U. Feige and D. Micciancio. The inapproximability of lattice and coding problems with preprocessing. *J. Comput. Syst. Sci.*, 69(1):45–67, 2004. doi:10.1016/j.jcss.2004.01.002.
- [GGG18] V. Gandikota, B. Ghazi, and E. Grigorescu. NP-hardness of Reed–Solomon decoding, and the Prouhet–Tarry–Escott problem. *SIAM Journal on Computing*, 47(4):1547–1584, 2018. doi:10.1137/16M110349X.
- [GST23] E. Goldberg, C. Shangguan, and I. Tamo. List-decoding and list-recovery of Reed–Solomon codes beyond the Johnson radius for every rate. *IEEE Transactions on Information Theory*, 69(4):2261–2268, 2023. doi:10.1109/TIT.2022.3222877.
- [GG00] O. Goldreich and S. Goldwasser. On the limits of nonapproximability of lattice problems. *Journal of Computer and System Sciences*, 60(3):540–563, 2000. doi:10.1006/jcss.1999.1686.
- [GGH97] O. Goldreich, S. Goldwasser, and S. Halevi. Public-key cryptosystems from lattice reduction problems. In *Advances in Cryptology—CRYPTO ’97*, volume 1294 of *Lecture Notes in Computer Science*, pages 112–131. Springer, 1997. doi:10.1007/BFb0052231.
- [GLSTW24] Z. Guo, R. Li, C. Shangguan, I. Tamo, and M. Wootters. Improved list-decodability and list-recoverability of Reed–Solomon codes via tree packings. *SIAM Journal on Computing*, 53(2):389–430, 2024. doi:10.1137/21M1463707.
- [GR06] V. Guruswami and A. Rudra. Limits to list decoding Reed–Solomon codes. *IEEE Transactions on Information Theory*, 52(8):3642–3649, 2006. doi:10.1109/TIT.2006.878164.

- [GS99] V. Guruswami and M. Sudan. Improved decoding of Reed–Solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999. doi:10.1109/18.782097.
- [HKW26] J. A. Huang, Y. K. Ko, and C. Wang. On the (classical and quantum) fine-grained complexity of approximate CVP and max-cut. In *53rd International Colloquium on Automata, Languages, and Programming*, volume 374 of *Leibniz International Proceedings in Informatics*, pages 111:1–111:17, 2026. doi:10.4230/LIPIcs.ICALP.2026.111.
- [LS71] J. Leech and N. J. A. Sloane. Sphere packings and error-correcting codes. *Canadian Journal of Mathematics*, 23(4):718–745, 1971. doi:10.4153/CJM-1971-081-3.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. doi:10.1007/BF01457454.
- [LS25] R. Li and N. Shagrithaya. Near-optimal list-recovery of linear code families. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025)*, volume 353 of *Leibniz International Proceedings in Informatics*, pages 53:1–53:14, 2025. doi:10.4230/LIPIcs.APPROX/RANDOM.2025.53.
- [Mas69] J. L. Massey. Shift-register synthesis and BCH decoding. *IEEE Transactions on Information Theory*, 15(1):122–127, 1969. doi:10.1109/TIT.1969.1054260.
- [Mos15] D. Moshkovitz. The projection games conjecture and the NP-hardness of $\ln n$ -approximating set-cover. *Theory of Computing*, 11(7):221–235, 2015. doi:10.4086/toc.2015.v011a007. See Section 5.1.
- [Muk22] P. Mukhopadhyay. The projection games conjecture and the hardness of approximation of super-SAT and related problems. *Journal of Computer and System Sciences*, 123:186–201, 2022. doi:10.1016/j.jcss.2021.09.002.
- [NIST24a] National Institute of Standards and Technology. Module-lattice-based key-encapsulation mechanism standard. *Federal Information Processing Standards Publication 203*, 2024. doi:10.6028/NIST.FIPS.203.
- [NIST24b] National Institute of Standards and Technology. Module-lattice-based digital signature standard. *Federal Information Processing Standards Publication 204*, 2024. doi:10.6028/NIST.FIPS.204.
- [RS60] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, 8(2):300–304, 1960. doi:10.1137/0108018.
- [Reg04] O. Regev. Improved inapproximability of lattice and coding problems with preprocessing. *IEEE Transactions on Information Theory*, 50(9):2031–2037, 2004. doi:10.1109/TIT.2004.833350.
- [Reg09] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):34:1–34:40, 2009. doi:10.1145/1568318.1568324.

- [RR00] R. M. Roth and G. Ruckenstein. Efficient decoding of Reed–Solomon codes beyond half the minimum distance. *IEEE Transactions on Information Theory*, 46(1):246–257, 2000. doi:[10.1109/18.817522](https://doi.org/10.1109/18.817522).
- [Sud97] M. Sudan. Decoding of Reed Solomon codes beyond the error-correction bound. *Journal of Complexity*, 13(1):180–193, 1997. doi:[10.1006/jcom.1997.0439](https://doi.org/10.1006/jcom.1997.0439).
- [vEB81] P. van Emde Boas. Another NP-complete partition problem and the complexity of computing short vectors in a lattice. Technical Report MI-UvA-81-04, Mathematical Institute, University of Amsterdam, 1981. <https://staff.fnwi.uva.nl/p.vanemdeboas/vectors/mi8104c.html>.
- [Var97] A. Vardy. The intractability of computing the minimum distance of a code. *IEEE Transactions on Information Theory*, 43(6):1757–1766, 1997. doi:[10.1109/18.641542](https://doi.org/10.1109/18.641542).
- [ZGA11] A. Zeh, C. Gentner, and D. Augot. An interpolation procedure for list decoding Reed–Solomon codes based on generalized key equations. *IEEE Transactions on Information Theory*, 57(9):5946–5959, 2011. doi:[10.1109/TIT.2011.2162160](https://doi.org/10.1109/TIT.2011.2162160).

CHAPTER 8

The Sharp Inequality in Ehrhart's Volume Conjecture

ABSTRACT. We prove the sharp bound that an n -dimensional convex body whose barycenter is its unique interior lattice point has volume at most $(n+1)^n/n!$.

CONTENTS

1. Introduction
 2. The real potential and lattice Bergman spaces
 3. Vanishing orders and the lower slope
 4. Bergman convexity and the upper slope
- References

1. INTRODUCTION

Ehrhart asked whether, among convex bodies whose barycenter is their only interior lattice point, the centered simplex maximizes volume [Ehr64]. Let $\Delta_n = \text{conv}\{0, e_1, \dots, e_n\}$ and write vol for ordinary Euclidean volume. The precise inequality is the following.

Theorem 1.1. *Let $K \subset \mathbb{R}^n$ be a full-dimensional compact convex body with barycenter 0. If $\text{int}(K) \cap \mathbb{Z}^n = \{0\}$, then*

$$\text{vol}(K) \leq \frac{(n+1)^n}{n!}. \quad (1)$$

The bound is sharp: the simplex $(n+1)\Delta_n - (1, \dots, 1)$ has barycenter 0, contains no other interior lattice point, and has volume $(n+1)^n/n!$. The equality refinement predicts that every extremizer is a unimodular image of this simplex under an integer linear map of determinant ± 1 [NP14, Conj. 1.1]. We do not determine whether these are the only equality cases.

Previous work. Ehrhart proved the conjecture for planar convex bodies and for simplices in every dimension [Ehr55, Ehr79]. For arbitrary centered bodies, Milman–Pajor symmetrization and Minkowski’s theorem give $\text{vol}(K) \leq 4^n$, as observed by Henk, Henze, and Hernández Cifre [MP00, HHH16]. Using thin-shell estimates, Huang, Slomka, Tkocz, and Vritsiou improved this to $\text{vol}(K) \leq 4^n e^{-c\sqrt{n}}$ for a universal $c > 0$ [HSTV22, Prop. 6.2]. Campos, van Hintum, Morris, and Tiba subsequently obtained $\text{vol}(K) \leq 4^n \exp(-cn/(\log n)^8)$ using bounds for the isotropic constant [CHMT24]. Combining their argument with Klartag and Lehec’s solution of Bourgain’s slicing problem gives $\text{vol}(K) \leq 4^n e^{-cn}$ for a universal $c > 0$ [CHMT24, KL25].

The conjecture is also known for certain classes of convex bodies. Berman and Berndtsson connect the sharp bound with the anticanonical volume of Kähler–Einstein toric varieties [BB17]. Their analytic argument combines Bergman convexity and a Moser–Trudinger inequality with a Green function at a torus-fixed point. In particular, it proves the conjecture for centered rational polytopes with facet presentation

$$P = \{y : \langle \ell_F, y \rangle \geq -a_F \text{ for every facet } F\}, \quad \ell_F \in \mathbb{Z}^n, \quad 0 < a_F \leq 1,$$

where each ℓ_F is primitive, meaning that its coordinates have no common factor [BB17, Cor. 1.4]. They also establish the sharp bound for bodies in the positive orthant with barycenter $(1, \dots, 1)$ and record Klartag’s direct derivation from Grünbaum’s barycentric half-space inequality [BB17, Thm. 1.5 and Rem. 3.2]; see [Gru60]. The displayed facet conditions also give an affine reduction to the orthant case. Write

$$Q^* = \{u : \langle u, y \rangle \leq 1 \text{ for every } y \in Q\}$$

for the polar of a convex body Q containing 0. Nill and Paffenholz extend this argument to convex bodies in the polar of a lattice polytope and classify the equality cases [NP14, Thm. 1.4]. Their hypothesis does not cover every centered rational polytope with unique interior lattice point 0: they give the example

$$P_0 = \text{conv}\{\pm(3/2, 1/4), \pm(3/2, 5/4)\}, \quad P_0^* \cap \mathbb{Z}^2 = \{0, \pm(1, -2)\},$$

so P_0^* contains no full-dimensional lattice polygon.

Idea of the proof. By a theorem of Berman and Berndtsson [BB13, Thm. 1.1], there is a smooth convex potential ϕ whose gradient transports $e^{-\phi(x)} dx$ to Lebesgue measure on K ; Cordero-Erausquin and Klartag also studied this transport problem [CK15]. On $X = (\mathbb{C}^*)^n$, regard ϕ as a function of the logarithmic radii:

$$\phi(z) = \phi(\log |z_1|^2, \dots, \log |z_n|^2).$$

Let $d\nu = dx d\theta$, where $x_i = \log |z_i|^2$ and $d\theta$ is normalized angular measure. The space $\mathcal{H}_k$ of holomorphic functions on X that are square-integrable with weight $e^{-k\phi} d\nu$ has Laurent monomial basis

$$\mathcal{H}_k = \text{span}\{z^m : m \in \mathbb{Z}^n \cap \text{int}(kK)\}, \quad \dim \mathcal{H}_k = \#(\mathbb{Z}^n \cap \text{int}(kK)).$$

The unique-interior-lattice-point hypothesis is therefore exactly the assertion $\mathcal{H}_1 = \mathbb{C}$.

At higher levels, fix $p = (1, \dots, 1)$. Every monomial equals 1 at p , so choose an orthonormal basis (s_a) adapted to vanishing order at p . Vanishing to order j forces all Taylor coefficients of degree less than j to vanish, imposing at most $\binom{n+j-1}{n}$ linear conditions. Write q_a for the order of s_a , truncated at $c_K k$, where $c_K = (n! \operatorname{vol}(K))^{1/n}$, and form the finite-level potentials $u_t^k = k^{-1} \log \sum_a |s_a|^2 e^{t q_a}$. Their regularized limit is a *ray of potentials*, a one-parameter family $(\psi_t)_{t \geq 0}$ with $\psi_0 = \phi$ and bounded $\psi_t - \phi$ for each t .

Introduce the normalized partition function and its logarithm:

$$Z(t) = \frac{1}{\operatorname{vol}(K)} \int_X e^{-\psi_t} d\nu, \quad L(t) = -\log Z(t).$$

Since $\mathcal{H}_1 = \mathbb{C}$, the weighted holomorphic space associated with ψ_t still consists only of constants, so its Bergman kernel equals $1/(\operatorname{vol}(K)Z(t))$. Berndtsson's positivity theorem makes L convex [Ber06], as in the one-dimensional Bergman mechanism of Berman and Berndtsson [BB17, Thm. 2.3]. Counting the vanishing conditions gives the lower bound on its initial slope. A separate local Schwarz estimate uses the decay of functions vanishing at p to bound ψ_t from above on a ball of radius proportional to $e^{-t/2}$; the volume of that ball gives the upper bound. Together, these arguments yield

$$\frac{n}{n+1} c_K \leq L'_+(0) \leq n. \quad (2)$$

Their comparison gives $c_K \leq n+1$, proving (1).

Fujita previously used the same point filtration and vanishing-order count to obtain a sharp volume bound in Fano geometry [Fuj18, Thms. 2.3 and 5.1]. Related constructions include filtrations of Laurent polynomials in graded Ehrhart theory [RR24, Cav26], toric and filtered Bergman kernels [Zel09, PS14], and rays from filtered linear series [BC11, RW14].

2. THE REAL POTENTIAL AND LATTICE BERGMAN SPACES

We first construct the transport potential, identify lattice points with an orthogonal basis of holomorphic monomials, and establish the convergence needed to pass from those finite-dimensional spaces to a limiting potential.

Let $h_K(x) = \sup_{y \in K} \langle y, x \rangle$ be the support function of K . Since its barycenter is 0, the existence theorem of Berman and Berndtsson [BB13, Thm. 1.1] gives a smooth strictly convex potential $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ whose gradient transports its associated log-concave measure to Lebesgue measure on K :

$$(\nabla \phi)_\#(e^{-\phi(x)} dx) = \mathbf{1}_K(y) dy, \quad |\phi(x) - h_K(x)| \leq C. \quad (3)$$

Moreover, $\nabla \phi$ is a diffeomorphism onto $\operatorname{int}(K)$, so the transport identity is equivalent to

$$\nabla \phi(\mathbb{R}^n) = \operatorname{int}(K), \quad \det D^2 \phi = e^{-\phi}.$$

In particular, the transported measure has exactly the volume of K :

$$\int_{\mathbb{R}^n} e^{-\phi(x)} dx = \int_{\mathbb{R}^n} \det D^2 \phi(x) dx = \operatorname{vol}(K). \quad (4)$$

To expose the lattice, attach independent angular coordinates to the original real variables. Write $d\theta$ for the Haar probability measure on $(S^1)^n$, so that

$$X = (\mathbb{C}^*)^n, \quad z_i = e^{x_i/2 + i\theta_i}, \quad d\nu = dx d\theta.$$

Let $d\lambda$ denote Euclidean Lebesgue measure on $\mathbb{C}^n$. The change to logarithmic coordinates gives

$$d\nu(z) = \pi^{-n} \prod_{i=1}^n |z_i|^{-2} d\lambda(z). \quad (5)$$

Here $x \in \mathbb{R}^n$ is unrestricted; in particular, the $|z_i|$ need not be bounded or bounded away from 0. Whenever ϕ is evaluated on X , it denotes the pullback

$$\phi(z) = \phi(\log |z_1|^2, \dots, \log |z_n|^2).$$

Since the angular variables have total mass one, (4) becomes

$$\int_X e^{-\phi} d\nu = \text{vol}(K), \quad d\mu = \frac{e^{-\phi}}{\text{vol}(K)} d\nu. \quad (6)$$

Write $\mathcal{O}(X)$ for the space of holomorphic functions on X . For any real weight b on X , its weighted holomorphic space and Bergman kernel are

$$\begin{aligned} \mathcal{H}(b) &= \left\{ f \in \mathcal{O}(X) : \int_X |f|^2 e^{-b} d\nu < \infty \right\}, \\ B_b(z) &= \sup_{f \in \mathcal{H}(b) \setminus \{0\}} \frac{|f(z)|^2}{\int_X |f|^2 e^{-b} d\nu}. \end{aligned} \quad (7)$$

For $k \geq 1$, specialize to $b = k\phi$:

$$\mathcal{H}_k = \mathcal{H}(k\phi), \quad \|s\|_k^2 = \int_X |s|^2 e^{-k\phi} d\nu.$$

The point of the complex torus is that its holomorphic Laurent monomials are indexed by $\mathbb{Z}^n$. By (3), z^m has finite weighted norm exactly when $m/k \in \text{int}(K)$: only then does its integrand decay exponentially in every direction. Since K is bounded, there are only finitely many such exponents.

Lemma 2.1. *The monomials z^m , with $m \in \mathbb{Z}^n \cap \text{int}(kK)$, form an orthogonal basis of $\mathcal{H}_k$. In particular, $\mathcal{H}_k$ is finite-dimensional; writing $d_k = \dim \mathcal{H}_k$, one has*

$$d_k = \#(\mathbb{Z}^n \cap \text{int}(kK)), \quad \frac{d_k}{k^n} \longrightarrow \text{vol}(K), \quad \mathcal{H}_1 = \mathbb{C}. \quad (8)$$

More generally, $\mathcal{H}(b) = \mathbb{C}$ whenever $b - \phi$ is bounded on X .

Proof. For $u = m/k$, integration over $\theta \in (S^1)^n$ gives

$$\|z^m\|_k^2 = I_k(u) := \int_{\mathbb{R}^n} e^{k(\langle u, x \rangle - \phi(x))} dx. \quad (9)$$

If $u \in \text{int}(K)$, there is $\delta > 0$ for which $h_K(x) - \langle u, x \rangle \geq \delta|x|$. Thus (3) implies $I_k(u) < \infty$. If $u \notin \text{int}(K)$, take a supporting direction $v \neq 0$ with $\langle u, v \rangle \geq h_K(v)$. On a tube of fixed radius about the ray $\mathbb{R}_{\geq 0}v$, $\langle u, x \rangle - \phi(x)$ is bounded below, so $I_k(u) = \infty$. For the Laurent expansion $s = \sum_{m \in \mathbb{Z}^n} c_m z^m$, Parseval's identity for the integral over θ and Tonelli's theorem give

$$\|s\|_k^2 = \sum_{m \in \mathbb{Z}^n} |c_m|^2 I_k(m/k),$$

with zero terms omitted. Hence an integrable s has nonzero coefficients only for $m \in \mathbb{Z}^n \cap \text{int}(kK)$. This set is finite because K is bounded, proving both the basis statement and finite-dimensionality. The lattice-counting limit follows from the fact that ∂K has measure zero. At $k = 1$, the only admissible exponent is 0. Finally, a bounded change of weight does not affect integrability, so $\mathcal{H}(b) = \mathcal{H}(\phi) = \mathbb{C}$ whenever $b - \phi$ is bounded. $\square$

We next approximate the probability measure $d\mu = e^{-\phi} d\nu / \text{vol}(K)$ using the lattice monomials. By Lemma 2.1 and (9), an orthonormal basis of $\mathcal{H}_k$ is given explicitly by

$$s_m(z) = \frac{z^m}{\sqrt{I_k(m/k)}}, \quad m \in \mathbb{Z}^n \cap \text{int}(kK).$$

Define the Bergman kernel B_k and the associated probability measure μ_k by

$$\begin{aligned} B_k(z) &= B_{k\phi}(z) = \sum_{m \in \mathbb{Z}^n \cap \text{int}(kK)} |s_m(z)|^2 = \sum_{m \in \mathbb{Z}^n \cap \text{int}(kK)} \frac{|z^m|^2}{I_k(m/k)}, \\ d\mu_k &= \frac{B_k e^{-k\phi}}{d_k} d\nu. \end{aligned} \quad (10)$$

The kernel B_k is independent of the orthonormal basis, so we may later replace the monomials by a basis adapted to vanishing at a point without changing B_k or μ_k . The next lemma supplies

the two convergence statements needed for the limiting construction. The logarithmic kernels recover ϕ , so the limiting deformation starts at the original transport potential; convergence of μ_k in total variation transfers averaged vanishing orders to the limiting measure μ . All limits in the next lemma are as $k \rightarrow \infty$, and $o_k(1)$ denotes an error tending to zero in that limit. Whenever uniformity is asserted, it is on the specified compact set.

Lemma 2.2. *There exists $C > 0$ such that, for every $k \geq 2$ and every $z \in X$,*

$$\frac{1}{k} \log B_k(z) \leq \phi(z) + C \frac{\log k}{k}. \quad (11)$$

Moreover, as $k \rightarrow \infty$,

$$\frac{1}{k} \log B_k \rightarrow \phi \quad \text{locally uniformly on } X, \quad \|\mu_k - \mu\|_{\text{TV}} \rightarrow 0.$$

Proof. Put $A = \max_{y \in K} |y|$. Since $\nabla \phi(\mathbb{R}^n) \subset K$, the function ϕ is A -Lipschitz on all of $\mathbb{R}^n$. For $x \in \mathbb{R}^n$, $u \in K$, and $|v - x| < 1/k$, it follows that

$$\langle u, v \rangle - \phi(v) \geq \langle u, x \rangle - \phi(x) - \frac{2A}{k}.$$

Integrating (9) over this real ball gives

$$I_k(u) \geq \omega_n k^{-n} e^{k(\langle u, x \rangle - \phi(x)) - 2A},$$

where ω_n is the volume of the unit ball in $\mathbb{R}^n$. For $z \in X$ with $x_i = \log |z_i|^2$ and $u = m/k$, we have $|z^m|^2 = e^{k\langle u, x \rangle}$. Therefore (10) gives

$$B_k(z) \leq \frac{e^{2A}}{\omega_n} k^n d_k e^{k\phi(z)} \leq C k^{2n} e^{k\phi(z)},$$

where the last inequality follows from $d_k/k^n \rightarrow \text{vol}(K)$. Taking logarithms proves (11).

To obtain both remaining conclusions from a single estimate, introduce the unnormalized densities

$$\rho_k = k^{-n} B_k e^{-k\phi}, \quad \rho = e^{-\phi}, \quad \rho_k d\nu = \frac{d_k}{k^n} d\mu_k, \quad \rho d\nu = \text{vol}(K) d\mu.$$

If $x = (\log |z_1|^2, \dots, \log |z_n|^2)$, the monomial formula (10) gives

$$\rho_k(x) = \frac{1}{k^n} \sum_{m \in \mathbb{Z}^n \cap \text{int}(kK)} \frac{e^{k(\langle m/k, x \rangle - \phi(x))}}{I_k(m/k)}.$$

For fixed x , the summands with m/k within $O(k^{-1/2})$ of $\nabla \phi(x)$ already suffice for both convergence statements. We approximate those summands by a Gaussian.

Write $H = D^2 \phi$ and define the Legendre transform

$$\phi^*(u) = \sup_{x \in \mathbb{R}^n} (\langle u, x \rangle - \phi(x)).$$

For $u \in \text{int}(K)$, the exponent in (9) has its unique maximum at $x_u = (\nabla \phi)^{-1}(u)$, with Hessian $-H(x_u)$. The multivariate Laplace estimate therefore gives

$$I_k(u) = e^{k\phi^*(u)} \frac{(2\pi/k)^{n/2}}{\sqrt{\det H(x_u)}} (1 + o_k(1)). \quad (12)$$

The error is uniform on every compact $S \subset \text{int}(K)$. Indeed, with $\delta_S = \text{dist}(S, \partial K) > 0$,

$$h_K(x) - \langle u, x \rangle \geq \delta_S |x| \quad (u \in S).$$

Together with (3), this controls the tails uniformly; the x_u remain in a compact set where $D^2 \phi$ is uniformly positive definite.

Fix $R > 0$, put $y = \nabla \phi(x)$, and retain only the monomials with $|m/k - y| \leq R/\sqrt{k}$. For $u = m/k$ and $q = \sqrt{k}(u - y)$, Taylor's formula gives

$$k(\phi^*(u) + \phi(x) - \langle u, x \rangle) = \frac{1}{2} \langle H(x)^{-1} q, q \rangle + o_k(1),$$

uniformly for $|q| \leq R$ and x in any fixed compact subset of $\mathbb{R}^n$. Substituting (12) into the monomial sum for ρ_k now gives a Gaussian Riemann sum:

$$\begin{aligned}\rho_k(x) &\geq F_R(x) + o_k(1), \\ F_R(x) &= \frac{\sqrt{\det H(x)}}{(2\pi)^{n/2}} \int_{|q| \leq R} e^{-\langle H(x)^{-1}q, q \rangle / 2} dq,\end{aligned}\tag{13}$$

again locally uniformly in x .

This one estimate gives both desired limits. On any compact $Q \subset \mathbb{R}^n$, F_R has a strictly positive minimum. Together with the global upper bound already proved, this gives constants $b_Q, C > 0$ such that

$$b_Q \leq \rho_k(x) \leq Ck^n \quad (x \in Q, k \text{ sufficiently large}).$$

Consequently,

$$\frac{1}{k} \log B_k(z) - \phi(z) = \frac{n \log k + \log \rho_k(x)}{k} \xrightarrow[k \rightarrow \infty]{} 0$$

uniformly for $x \in Q$, proving local uniform convergence.

For the measures, let $R \rightarrow \infty$ in (13). The full Gaussian integral equals $(2\pi)^{n/2} \sqrt{\det H(x)}$, so $\det H = e^{-\phi}$ gives

$$\liminf_{k \rightarrow \infty} \rho_k(x) \geq \det H(x) = \rho(x).$$

On the other hand, (8) and (6) give the exact mass limit

$$\int_X \rho_k d\nu = \frac{d_k}{k^n} \xrightarrow[k \rightarrow \infty]{} \text{vol}(K) = \int_X \rho d\nu.\tag{14}$$

Since $0 \leq \min(\rho_k, \rho) \leq \rho$ and $\min(\rho_k, \rho) \rightarrow \rho$ pointwise, dominated convergence and (14) give

$$\|\rho_k - \rho\|_{L^1(d\nu)} = \int_X \rho_k d\nu + \int_X \rho d\nu - 2 \int_X \min(\rho_k, \rho) d\nu \xrightarrow[k \rightarrow \infty]{} 0.$$

Since $d_k/k^n \rightarrow \text{vol}(K)$ as $k \rightarrow \infty$, normalizing these densities proves $\|\mu_k - \mu\|_{\text{TV}} \rightarrow 0$. $\square$

3. VANISHING ORDERS AND THE LOWER SLOPE

We first count the vanishing orders forced by the dimensions of $\mathcal{H}_k$. We then realize that count as the initial slope of a limiting family of potentials.

Fix $p = (1, \dots, 1) \in X$. For a nonzero holomorphic function s , expand around p as

$$s(p + \zeta) = \sum_{\alpha \in \mathbb{Z}_{\geq 0}^n} c_\alpha \zeta^\alpha, \quad \text{ord}_p(s) = \min\{|\alpha| : c_\alpha \neq 0\}, \quad |\alpha| = \alpha_1 + \dots + \alpha_n.$$

Thus $\text{ord}_p(s) \geq j$ means that all Taylor coefficients of total degree less than j vanish; set $\text{ord}_p(0) = \infty$. Filter $\mathcal{H}_k$ by these conditions:

$$F_k^j = \{s \in \mathcal{H}_k : \text{ord}_p(s) \geq j\}, \quad \mathcal{H}_k = F_k^0 \supseteq F_k^1 \supseteq F_k^2 \supseteq \dots$$

The number of Taylor coefficients of degree less than j controls the codimension of F_k^j .

Lemma 3.1. *For every $j, k \geq 1$,*

$$d_k - \dim F_k^j \leq \binom{n+j-1}{n}.\tag{15}$$

Proof. Taking the Taylor polynomial of total degree less than j defines a linear map on $\mathcal{H}_k$ with kernel F_k^j and target of dimension $\binom{n+j-1}{n}$. $\square$

Since $\mathcal{H}_k$ is finite-dimensional and a holomorphic function vanishing to every order is zero, $F_k^j = 0$ for all sufficiently large j . For each j , choose an orthonormal basis of $F_k^j \cap (F_k^{j+1})^\perp$, where orthogonality is taken in $\mathcal{H}_k$. Combining these bases gives an orthonormal basis $s_1, \dots, s_{d_k}$ of $\mathcal{H}_k$. Writing $j_a = \text{ord}_p(s_a)$, it satisfies

$$F_k^j = \text{span}\{s_a : j_a \geq j\}, \quad \dim F_k^j = \#\{a : j_a \geq j\}.\tag{16}$$

Replacing the monomial basis by this adapted basis does not change B_k or μ_k .

For fixed $s > 0$ and $j = \lfloor sk \rfloor$, the lattice count (8) and the vanishing estimate (15) give

$$\frac{\dim F_k^{\lfloor sk \rfloor}}{k^n} \geq \text{vol}(K) - \frac{s^n}{n!} + o_k(1).$$

The leading-order lower bound is positive exactly when $s < c_K = (n! \text{vol}(K))^{1/n}$, so we truncate the vanishing orders at this scale:

$$N_k = \lfloor c_K k \rfloor, \quad q_a = \min\{j_a, N_k\}.$$

For a real parameter t , multiply the squared contribution of a basis element with truncated order q_a by e^{tq_a} . The resulting potential and its initial slope are

$$u_t^k(z) = \frac{1}{k} \log \sum_{a=1}^{d_k} |s_a(z)|^2 e^{tq_a}, \quad g_k(z) = \frac{\sum_a q_a |s_a(z)|^2}{kB_k(z)}.$$

Thus $u_0^k = k^{-1} \log B_k$ and $g_k = \partial_t u_t^k|_{t=0}$; at each z , $g_k(z)$ is the average of q_a/k with probabilities $|s_a(z)|^2/B_k(z)$. In particular,

$$0 \leq g_k \leq c_K, \quad |u_t^k - u_0^k| \leq c_K |t|. \quad (17)$$

The average of g_k under μ_k records exactly the truncated vanishing orders. Since the adapted basis is orthonormal, (16) gives

$$\int_X g_k d\mu_k = \frac{1}{kd_k} \sum_{a=1}^{d_k} q_a = \frac{1}{kd_k} \sum_{j=1}^{N_k} \dim F_k^j \geq \frac{N_k d_k - \binom{n+N_k}{n+1}}{kd_k}.$$

The inequality follows from (15) and $\sum_{j=1}^{N_k} \binom{n+j-1}{n} = \binom{n+N_k}{n+1}$. By (8), its large- k limit is the same sharp integral as in Fujita's point-filtration argument [Fuj18, proof of Thm. 5.1]:

$$\liminf_{k \rightarrow \infty} \int_X g_k d\mu_k \geq \frac{1}{\text{vol}(K)} \int_0^{c_K} \left(\text{vol}(K) - \frac{s^n}{n!} \right) ds = \frac{n}{n+1} c_K. \quad (18)$$

The bound (18) involves a different potential u_t^k and initial slope g_k at each level k . We now construct a single, k -independent family of potentials ψ_t with $\psi_0 = \phi$, then transfer this bound to its initial slope $g = \dot{\psi}_{0+}$. The construction uses the complex analogue of convexity.

Recall that a function $a : D \rightarrow [-\infty, \infty)$ on a domain $D \subset \mathbb{C}^N$ is *plurisubharmonic* if it is upper semicontinuous and its restriction to every complex affine line satisfies the submean inequality

$$a(w) \leq \frac{1}{2\pi} \int_0^{2\pi} a(w + re^{i\theta}v) d\theta$$

for every $w \in D$, $v \in \mathbb{C}^N$, and $r > 0$ for which the closed complex disk $\{w + \zeta v : |\zeta| \leq r\}$ lies in D . When a is twice differentiable, this is equivalent to positive semidefiniteness of its complex Hessian $(\partial^2 a / \partial w_i \partial \bar{w}_j)_{i,j=1}^N$.

For $t = \log |\tau|^2$, put

$$U_k(z, \tau) = u_t^k(z) = \frac{1}{k} \log \sum_{a=1}^{d_k} |s_a(z) \tau^{q_a}|^2.$$

This function is plurisubharmonic on $X \times \mathbb{C}^*$: it is k^{-1} times the logarithm of the squared norm of a holomorphic vector. By Lemma 2.2 and (17), the functions U_k are locally uniformly bounded above. Their regularized tail envelopes therefore define a plurisubharmonic limit:

$$\Psi_m = \text{usc}_{(z,\tau)} \left(\sup_{k \geq m} U_k(z, \tau) \right), \quad \Psi = \lim_{m \rightarrow \infty} \Psi_m, \quad \psi_t(z) = \Psi(z, e^{t/2}). \quad (19)$$

Here $\text{usc} f(z, \tau) = \limsup_{(z', \tau') \rightarrow (z, \tau)} f(z', \tau')$ denotes upper-semicontinuous regularization in both variables. For a compact neighborhood Q of z , set $\epsilon_m(Q) = \sup_{k \geq m} \sup_Q |u_0^k - \phi| \rightarrow 0$. By (17), for $z' \in Q$ and $k \geq m$, $|U_k(z', \tau) - \phi(z')| \leq \epsilon_m(Q) + c_K |\log |\tau|^2|$. Taking the regularized

supremum at $(z, 1)$ gives $|\Psi_m(z, 1) - \phi(z)| \leq \epsilon_m(Q)$ and therefore $\psi_0 = \phi$. For $t \geq 0$, taking the same regularized limit in $u_0^k \leq u_t^k \leq u_0^k + c_K t$ gives

$$\psi_0 = \phi, \quad \phi \leq \psi_t \leq \phi + c_K t \quad (t \geq 0). \quad (20)$$

Since a radial plurisubharmonic function is convex in its logarithmic radius, $t \mapsto \psi_t(z)$ is convex. Define its initial right-hand velocity by

$$g(z) := \dot{\psi}_{0+}(z) := \lim_{t \downarrow 0} \frac{\psi_t(z) - \phi(z)}{t}. \quad (21)$$

In particular, $0 \leq g \leq c_K$.

The envelope construction immediately gives $\limsup_{k \rightarrow \infty} u_t^k \leq \psi_t$. Convexity now relates the finite-level velocities g_k to the limiting velocity g : for every fixed $t > 0$,

$$g_k(z) \leq \frac{u_t^k(z) - u_0^k(z)}{t}.$$

Since $u_0^k \rightarrow \phi$ locally uniformly, taking upper limits in this inequality and then letting $t \downarrow 0$ gives $\limsup_{k \rightarrow \infty} g_k \leq g$. Moreover, $0 \leq g_k \leq c_K$. Total-variation convergence of μ_k to μ and reverse Fatou therefore give

$$\limsup_{k \rightarrow \infty} g_k \leq g, \quad \limsup_{k \rightarrow \infty} \int_X g_k d\mu_k \leq \int_X g d\mu. \quad (22)$$

Proposition 3.2. *For $c_K = (n! \operatorname{vol}(K))^{1/n}$, the initial velocity (21) satisfies*

$$\int_X g d\mu \geq \frac{n}{n+1} c_K.$$

Proof. Combine the finite-level estimate (18) with the limiting comparison (22). $\square$

4. BERGMAN CONVEXITY AND THE UPPER SLOPE

It remains to turn the limiting potential into a convex logarithmic partition function and control its growth near the point p . A function is *pluriharmonic* if both it and its negative are plurisubharmonic; adding such a function preserves plurisubharmonicity. In particular, $\log |z_i|^2$ is pluriharmonic where $z_i \neq 0$.

On $D = (\mathbb{C}^*)^n \times \{\tau \in \mathbb{C} : |\tau| > 1\}$, Berndtsson's positivity theorem says that the logarithm of the Bergman kernel associated with any plurisubharmonic weight and Euclidean measure is itself plurisubharmonic [Ber06, Thm. 1.1]. The same conclusion holds for the kernels (7) defined using $d\nu$: if b is plurisubharmonic on D , then

$$(z, \tau) \mapsto \log B_{b_\tau}(z), \quad b_\tau(z) = b(z, \tau), \quad (23)$$

is plurisubharmonic. Indeed, by (5), the Euclidean weight $a(z, \tau) = b(z, \tau) + \sum_{i=1}^n \log |z_i|^2 + n \log \pi$ satisfies $e^{-a} d\lambda = e^{-b} d\nu$. The added logarithms are pluriharmonic, so a is plurisubharmonic whenever b is. More generally, Berndtsson's theorem holds on any *pseudoconvex* domain.

We now turn the ray into the one-dimensional convex function required by (2). Set

$$Z(t) = \frac{1}{\operatorname{vol}(K)} \int_X e^{-\psi_t} d\nu, \quad L(t) = -\log Z(t), \quad t \geq 0.$$

In particular, $Z(0) = 1$ and $L(0) = 0$. In general, if $1 = f_0, f_1, \dots, f_r$ is a basis of a weighted holomorphic space, its Gram matrix and Bergman kernel are

$$G_{ij}(t) = \int_X f_i \bar{f}_j e^{-\psi_t} d\nu, \quad B_{\psi_t}(z) = v(z)^* G(t)^{-1} v(z), \quad v(z) = (f_0(z), \dots, f_r(z))^T.$$

The unnormalized partition function is $G_{00}(t)$, so positivity of $\log B_{\psi_t}$ does not generally imply convexity of $-\log G_{00}(t)$. The unique-interior-lattice-point hypothesis is what reduces $G(t)$ to a 1×1 matrix and identifies these two quantities.

Lemma 4.1. *The function L is finite and convex, and*

$$L'_+(0) = \int_X g d\mu. \quad (24)$$

Proof. By (20), $\psi_t - \phi$ is bounded for each $t \geq 0$, so Lemma 2.1 gives $\mathcal{H}(\psi_t) = \mathbb{C}$. Taking the supremum over constant functions in (7) therefore identifies the Bergman kernel with the partition function:

$$B_{\psi_t}(z) = \left(\int_X e^{-\psi_t} d\nu \right)^{-1} = \frac{e^{L(t)}}{\text{vol}(K)}.$$

Apply (23) to the plurisubharmonic weight $b(z, \tau) = \Psi(z, \tau)$. It follows that $\log B_{\psi_t}(z) = L(\log |\tau|^2) - \log \text{vol}(K)$ is plurisubharmonic. A radial subharmonic function is convex in its logarithmic radius, so L is convex on $(0, \infty)$; (20) extends the convexity continuously to 0.

Finally, (20) and the definition of g give $(1 - e^{-(\psi_t - \phi)})/t \rightarrow g$ pointwise as $t \downarrow 0$, with the quotient bounded between 0 and c_K . Dominated convergence and (6) therefore give

$$\frac{1 - Z(t)}{t} = \int_X \frac{1 - e^{-(\psi_t - \phi)}}{t} d\mu \longrightarrow \int_X g d\mu.$$

Since $Z(t) \rightarrow 1$ and $L = -\log Z$, this proves (24). $\square$

To bound the other side of the initial slope, it suffices to look near the point p where the filtration was defined. High vanishing order cancels the exponential filtration weight on a ball whose radius decreases like $e^{-t/2}$.

Lemma 4.2.

$$L(t) \leq nt \quad (t \geq 0). \quad (25)$$

Proof. Choose a branch of the coordinates $\zeta_i = \log z_i$ near p , so that p corresponds to $\zeta = 0$ and $d\nu = \pi^{-n} d\lambda(\zeta)$. Let B_{2r} be the Euclidean ball of radius $2r$ in these coordinates and put $M_r = \sup_{B_{2r}} \phi$. For $s \in \mathcal{H}_k$ with $\|s\|_k = 1$, the submean inequality on B_{2r} bounds s on B_r by $C_r e^{kM_r/2}$. If s vanishes to order j at p , applying the one-variable Schwarz estimate on each complex line through 0 improves this to

$$|s(\zeta)| \leq C_r e^{kM_r/2} \left(\frac{|\zeta|}{r} \right)^j \quad (|\zeta| < r).$$

If $|\zeta| < re^{-t/2}$, then $q_a \leq j_a$ gives $(|\zeta|/r)^{2j_a} e^{tq_a} \leq 1$. Thus $u_t^k(\zeta) \leq M_r + k^{-1} \log(C_r^2 d_k)$ on the open joint region $|\zeta| |\tau| < r$, $|\tau| > 1$. The bound therefore survives the upper-semicontinuous regularization in (19). Since $d_k = O(k^n)$, taking the regularized tail limit gives $\psi_t \leq M_r$ on $B_{re^{-t/2}}$, whose $d\nu$ -volume is $\pi^{-n} \lambda(B_r) e^{-nt} = c_r e^{-nt}$, where $c_r > 0$. Consequently, for some constant D independent of t ,

$$Z(t) \geq \frac{c_r}{\text{vol}(K)} e^{-M_r - nt}, \quad L(t) \leq nt + D.$$

For $0 < t < T$, convexity and $L(0) = 0$ now give

$$\frac{L(t)}{t} \leq \frac{L(T)}{T} \leq n + \frac{D}{T}.$$

Letting $T \rightarrow \infty$ proves (25). $\square$

Together, Proposition 3.2 and Lemmas 4.1 and 4.2 give (2), proving Theorem 1.1.

REFERENCES

- [BB13] R. J. Berman and B. Berndtsson, *Real Monge–Ampère equations and Kähler–Ricci solitons on toric log Fano varieties*, Ann. Fac. Sci. Toulouse Math. (6) **22** (2013), 649–711.
- [BB17] R. J. Berman and B. Berndtsson, *The volume of Kähler–Einstein Fano varieties and convex bodies*, J. Reine Angew. Math. **723** (2017), 127–152.
- [Ber06] B. Berndtsson, *Subharmonicity properties of the Bergman kernel and some other functions associated to pseudoconvex domains*, Ann. Inst. Fourier (Grenoble) **56** (2006), 1633–1662.
- [BC11] S. Boucksom and H. Chen, *Okounkov bodies of filtered linear series*, Compos. Math. **147** (2011), 1205–1229.
- [CHMT24] M. Campos, P. van Hintum, R. Morris, and M. Tiba, *Towards Hadwiger’s conjecture via Bourgain slicing*, Int. Math. Res. Not. IMRN **2024** (2024), no. 10, 8282–8295, doi:10.1093/imrn/rnad198.
- [Cav26] I. Cavey, *Graded Ehrhart theory and toric geometry*, Proc. Amer. Math. Soc. **154** (2026), 1859–1866.
- [CK15] D. Cordero-Erausquin and B. Klartag, *Moment measures*, J. Funct. Anal. **268** (2015), 3834–3866.

- [Ehr55] E. Ehrhart, *Une généralisation du théorème de Minkowski*, C. R. Acad. Sci. Paris **240** (1955), 483–485.
- [Ehr64] E. Ehrhart, *Une généralisation probable du théorème fondamental de Minkowski*, C. R. Acad. Sci. Paris **258** (1964), 4885–4887.
- [Ehr79] E. Ehrhart, *Volume réticulaire critique d'un simplexe*, J. Reine Angew. Math. **305** (1979), 218–220.
- [Fuj18] K. Fujita, *Optimal bounds for the volumes of Kähler–Einstein Fano manifolds*, Amer. J. Math. **140** (2018), 391–414.
- [Gru60] B. Grünbaum, *Partitions of mass-distributions and of convex bodies by hyperplanes*, Pacific J. Math. **10** (1960), 1257–1261.
- [HHH16] M. Henk, M. Henze, and M. A. Hernández Cifre, *Variations of Minkowski’s theorem on successive minima*, Forum Math. **28** (2016), 311–325, doi:10.1515/forum-2014-0093.
- [HSTV22] H. Huang, B. A. Slomka, T. Tkocz, and B.-H. Vritsiou, *Improved bounds for Hadwiger’s covering problem via thin-shell estimates*, J. Eur. Math. Soc. **24** (2022), 1431–1448.
- [KL25] B. Klartag and J. Lehec, *Affirmative resolution of Bourgain’s slicing problem using Guan’s bound*, Geom. Funct. Anal. **35** (2025), 1147–1168, doi:10.1007/s00039-025-00718-w.
- [MP00] V. D. Milman and A. Pajor, *Entropy and asymptotic geometry of non-symmetric convex bodies*, Adv. Math. **152** (2000), 314–335.
- [NP14] B. Nill and A. Paffenholz, *On the equality case in Ehrhart’s volume conjecture*, Adv. Geom. **14** (2014), 579–586.
- [PS14] F. T. Pokorný and M. Singer, *Toric partial density functions and stability of toric varieties*, Math. Ann. **358** (2014), 879–923.
- [RR24] V. Reiner and B. Rhoades, *Harmonics and graded Ehrhart theory*, J. Combin. Algebra, to appear; arXiv:2407.06511.
- [RW14] J. Ross and D. Witt Nyström, *Analytic test configurations and geodesic rays*, J. Symplectic Geom. **12** (2014), 125–169.
- [Zel09] S. Zelditch, *Bernstein polynomials, Bergman kernels and toric Kähler varieties*, J. Symplectic Geom. **7** (2009), 51–76.

CHAPTER 9

Super-exponential lower bounds for $R(3, \dots, 3)$

ABSTRACT. Let $R_k(3) = R(\underbrace{3, \dots, 3}_k)$ denote the multicolor Ramsey number for k colors, that is, the least N for which every k -coloring of the edges of K_N contains a monochromatic triangle. We prove that there exists an absolute constant $c > 0$ such that, for every integer $k \geq 2$,

$$R_k(3) \geq \left(\frac{ck^{1/3}}{\log k} \right)^k.$$

Together with the classical factorial upper bound, this establishes $R_k(3) = k^{\Theta(k)}$. In particular, the Shannon capacity of graphs with independence number 2 is unbounded.

CONTENTS

1. Introduction
 2. Saturated matrices, coordinate covers, and palettes
 3. Recursive triangle-free colorings
- References

1. INTRODUCTION

Write

$$R_k(3) = R(\underbrace{3, \dots, 3}_k)$$

for the least integer N such that every coloring of $E(K_N)$ with k colors contains a monochromatic triangle. Via the standard product construction

$$R_{k+\ell}(3) - 1 \geq (R_k(3) - 1)(R_\ell(3) - 1)$$

and Fekete's lemma, the limit

$$L = \lim_{k \rightarrow \infty} R_k(3)^{1/k} \quad (1)$$

exists in $[1, \infty]$.

Previous lower bounds were obtained by tensoring small triangle-free colorings and by constructions from sum-free partitions [GG55, Chu73, CG83, Exo94, FS00]. After a series of improvements, the best known lower bound before this work was

$$R_k(3) \geq 380^{k/5} - O(1)$$

[ACPPRT21, Blo183]. On the upper-bound side, refinements [Whi73, Wan97, XXC02, Blo183, Rad] of the standard monochromatic-neighborhood recurrence led to constant-factor improvements of the upper bound. The current record is

$$R_k(3) \leq \left(e - \frac{1}{6}\right) k! + 1 \quad (k \geq 4). \quad (2)$$

The question whether $R_k(3)$ grows superexponentially was recorded by Graham, Rothschild, and Spencer [GRS90, p. 146]. The connection between the multicolor Ramsey problem and Shannon capacity appears implicitly in work of Erdős, McEliece, and Taylor [EMT71] and was made explicit by Alon and Orlitsky [AO95, §2.1], who also raised the analogous question for $R(c, \dots, c)$ with fixed c . The gap between these exponential lower bounds and factorial upper bounds was later highlighted explicitly by Conlon, Fox, and Sudakov [CFS15, §2.1]; see also [CG83, Rad]. For the broader interaction between structure and randomness, see Gowers [Gow00]. Erdős offered \$250 for determining the value of (1) and \$100 for deciding whether it is finite [CG83, Blo183, CG].

Theorem 1.1. *There exists an absolute constant $c > 0$ such that, for every integer $k \geq 2$,*

$$R(\underbrace{3, \dots, 3}_k) = R_k(3) \geq \left(\frac{ck^{1/3}}{\log k}\right)^k. \quad (3)$$

Together with (2), this gives

$$k^{(1/3 - o(1))k} \leq R_k(3) \leq k^{(1 + o(1))k}, \quad R_k(3) = k^{\Theta(k)}.$$

In particular,

$$\lim_{k \rightarrow \infty} R_k(3)^{1/k} = +\infty. \quad (4)$$

The Ramsey–Shannon correspondence [EMT71, AO95] also gives an equivalent formulation of (4). For a graph G , write

$$\Theta(G) = \sup_{m \geq 1} \alpha(G^{\boxtimes m})^{1/m}$$

for its Shannon capacity, where $\alpha(G)$ is the independence number and $\boxtimes$ denotes the strong graph product. Given a triangle-free k -coloring of K_N , let H_i be the graph of color i and set

$$G = \overline{H_1} \vee \dots \vee \overline{H_k},$$

where $\vee$ denotes the complete join. Then $\alpha(G) = 2$, and the N diagonal words form an independent set in $G^{\boxtimes k}$. Thus $\Theta(G) \geq N^{1/k}$. Taking $N = R_k(3) - 1$ and applying (4),

we obtain graphs with independence number 2 and arbitrarily large Shannon capacity. In particular, Shannon capacity cannot be bounded above by any function of the independence number.

1.1. Proof outline. The argument adapts the random-matrix and coordinate-covering ingredients of Alon, Ben-Eliezer, Shangquan, and Tamo [ABST20, Lemmas 3.4 and 4.1]. Their saturated-matrix argument in turn builds on work of Chakraborty, Radhakrishnan, Raghunathan, and Sasatte on zero-error list decoding [CRRS06]. We give direct proofs of the matrix and two-sided coordinate-covering statements without the hat-guessing terminology. The saturated-matrix construction is not new; its application to $R(3, \dots, 3)$ is.

Fix a stage parameter H . A union bound produces an H -colored $s \times H^m$ matrix, where $m \asymp H \log H$ and $s \asymp H^2 \log^2 H$, such that among any $m+1$ columns some row contains every color. This property yields maps $f, g : [H]^s \rightarrow [H]^s$, fixed once for the entire construction, such that for every $x, y \in [H]^s$ there is a $d \in [s]$ with

$$x_d = (f(y))_d \quad \text{or} \quad y_d = (g(x))_d.$$

We construct an edge-coloring recursively while maintaining the stronger invariant that, at stage j , each color graph is properly $(j+1)$ -colorable. Divide the vertices into blocks indexed by palettes $P \subseteq [jt]$ of size $t = s \lceil \log H \rceil$. The palette P records the colors missing from its block; all other colors are active. Inside that block, place a copy of the preceding coloring with colors relabeled by $[jt] \setminus P$. Each active color then has a proper internal vertex labeling in $[j]$. A maximal packing provides many palettes while ensuring that distinct palettes differ in at least s colors in each direction.

For blocks $P < Q$, choose colors $a_1, \dots, a_s \in Q \setminus P$ and $b_1, \dots, b_s \in P \setminus Q$. The internal labels of $u \in V_P$ in the a_d and of $v \in V_Q$ in the b_d form words $x(u), y(v) \in [H]^s$. Apply the fixed coordinate cover to these words: if $x_d(u) = (f(y(v)))_d$, color uv with a_d ; otherwise color it with a b_d for which $y_d(v) = (g(x(u)))_d$. Every cross-edge color is therefore active at exactly one endpoint block and missing at the other. Crucially, the proper label at the active endpoint is determined by the opposite endpoint.

This last property excludes triangles on two blocks: two edges of the same color from one outside vertex into an active block force the same internal label, so the edge between their endpoints cannot have that color. If the color is missing from the block, there is no internal edge of that color in the first place. A triangle on three blocks would require a color to belong to each of $P \triangle Q$, $Q \triangle R$, and $P \triangle R$: by symmetry, if $c \in P$, the first two memberships give $c \notin Q$ and $c \in R$, contradicting the third. Finally, label each active block using its proper internal labels and each missing block with the new label $j+1$. Since every cross edge of a fixed color joins an active block to a missing block, these labels properly color its entire color graph and propagate the inductive invariant.

Multiplying the sizes of the packed palette families over the H stages gives a triangle-free coloring with $k_H \asymp H^3 \log^3 H$ colors and at least $(c_0 H)^{k_H}$ vertices. Rounding down to the nearest such color count changes only the absolute constant. Since $H \gg k^{1/3} / \log k$, this proves (3) for all sufficiently large k ; decreasing the constant covers the remaining $k \geq 2$.

All logarithms are natural, $[r] = \{1, \dots, r\}$, and $[0]$ is empty. We do not optimize the absolute constants.

2. SATURATED MATRICES, COORDINATE COVERS, AND PALETTES

We first show that, in a thin, wide matrix randomly colored with H colors, every sufficiently large set of columns contains a row in which all H colors appear. This is the saturated-matrix construction of [ABST20, Lemma 3.4], whose underlying family was studied earlier in [CRRS06].

Lemma 2.1. *Let $H \geq 2$, and define*

$$m = \lceil 2H \log H \rceil, \quad s = m(m+1) + 1. \quad (5)$$

There is a matrix

$$A = (A_{r,z})_{r \in [s], z \in [H]^m} \quad \text{with } A_{r,z} \in [H]$$

such that, for every $T \subseteq [H]^m$ with $|T| = m + 1$, some row $r \in [s]$ satisfies

$$\{A_{r,z} : z \in T\} = [H]. \quad (6)$$

Proof. Choose all entries independently and uniformly from $[H]$. For a fixed set T of $m + 1$ columns, a given row fails (6) only if one of the H symbols is absent. Hence

$$\mathbb{P}(\{A_{r,z} : z \in T\} \neq [H]) \leq H \left(1 - \frac{1}{H}\right)^{m+1} < H \exp\left(-\frac{m}{H}\right) \leq \frac{1}{H}.$$

Independence bounds the probability that all s rows fail for this T by H^{-s} . Since there are at most $\binom{H^m}{m+1} \leq H^{m(m+1)}$ choices of T , a second union bound gives a total failure probability at most

$$H^{m(m+1)-s} = H^{-1} < 1.$$

Thus one matrix works simultaneously for all sets of $m + 1$ columns. $\square$

The next lemma converts this matrix property into a two-sided coordinate cover. One function handles all but a small exceptional set of words; the other handles the remaining words by assigning them distinct coordinates. The construction adapts [ABST20, Lemma 4.1] to the present symmetric formulation.

Lemma 2.2 (Two-sided coordinate cover). *For H, m, s as in (5), there are fixed maps*

$$f, g : [H]^s \longrightarrow [H]^s$$

such that, for every $x, y \in [H]^s$,

$$\exists d \in [s] : x_d = (f(y))_d \quad \text{or} \quad y_d = (g(x))_d. \quad (7)$$

Proof. Fix a matrix A from Lemma 2.1. Write $\bar{x} = (x_1, \dots, x_m)$, and set

$$(g(x))_r = A_{r,\bar{x}} \quad (r \in [s]). \quad (8)$$

For $y \in [H]^s$, let

$$E_y = \{z \in [H]^m : A_{r,z} \neq y_r \text{ for every } r \in [s]\}. \quad (9)$$

Then $|E_y| \leq m$. Otherwise take $m + 1$ members of E_y . The matrix property gives a row containing every symbol on those columns and, in particular, a column z for which $A_{r,z} = y_r$ for some r , contrary to (9).

Enumerate E_y in a fixed order as $z^{(1)}, \dots, z^{(\rho)}$, where $\rho \leq m$, and define

$$(f(y))_d = \begin{cases} z_d^{(d)}, & 1 \leq d \leq \rho, \\ 1, & \rho < d \leq s. \end{cases}$$

The first case makes sense because $d \leq \rho \leq m$. If $\bar{x} \notin E_y$, then (9) and (8) give some $d \in [s]$ such that

$$y_d = A_{d,\bar{x}} = (g(x))_d.$$

If instead $\bar{x} \in E_y$, then $\bar{x} = z^{(d)}$ for some $d \leq \rho$, and

$$x_d = z_d^{(d)} = (f(y))_d.$$

In either case, (7) follows. $\square$

We next pack the palettes that will specify the colors missing from each block. Fix $H \geq 3$, use the parameters of (5), and put

$$t = s \lceil \log H \rceil, \quad M_j = jt \quad (0 \leq j \leq H). \quad (10)$$

In particular, $t \geq 2s$. At stage j , a palette is a member of $\binom{[M_j]}{t}$. A block with palette P will omit exactly the colors in P .

Lemma 2.3 (Separated palettes). *For every $1 \leq j \leq H$, there is a family*

$$\mathcal{P}_j \subseteq \binom{[jt]}{t}$$

such that distinct $P, Q \in \mathcal{P}_j$ satisfy

$$|P \setminus Q| = |Q \setminus P| \geq s. \quad (11)$$

Writing $B_j = |\mathcal{P}_j|$, one can arrange that $B_j \geq 1$ and

$$B_j \geq \frac{\binom{jt}{t}}{\sum_{d=0}^{s-1} \binom{t}{d} \binom{(j-1)t}{d}} \geq \frac{j^t}{s(e^2 j \lceil \log H \rceil^2)^s}. \quad (12)$$

Proof. Take a maximal family satisfying (11). Every t -subset of $[jt]$ then lies at difference less than s from at least one selected palette. Counting these covering balls gives the first inequality in (12). Since $t = s \lceil \log H \rceil \geq 2s$, the estimates

$$\binom{jt}{t} \geq j^t, \quad \sum_{d=0}^{s-1} \binom{t}{d} \binom{(j-1)t}{d} \leq s \binom{t}{s} \binom{jt}{s}, \quad \binom{N}{r} \leq \left(\frac{eN}{r}\right)^r$$

immediately give the second inequality in (12). A maximal family is nonempty, so $B_j \geq 1$. $\square$

3. RECURSIVE TRIANGLE-FREE COLORINGS

For an edge-coloring κ , write $\Gamma_\kappa(c)$ for the spanning graph whose edges are those of color c . We construct colorings for which every $\Gamma_\kappa(c)$ has a short proper vertex coloring. This stronger invariant is what makes a triangle-free recursion possible.

Proposition 3.1 (Recursive coloring). *Fix $H \geq 3$, let t and M_j be as in (10), and, for $1 \leq r \leq H$, let $\mathcal{P}_r$ be a palette family provided by Lemma 2.3. Write $B_r = |\mathcal{P}_r|$. For each $0 \leq j \leq H$, there is a coloring*

$$\kappa_j : E(K_{n_j}) \longrightarrow [M_j], \quad n_j = \prod_{r=1}^j B_r,$$

with the following properties:

- (1) κ_j contains no monochromatic triangle;
- (2) for each $c \in [M_j]$,

$$\chi(\Gamma_{\kappa_j}(c)) \leq j + 1.$$

Proof. Fix once and for all the maps $f, g : [H]^s \rightarrow [H]^s$ provided by Lemma 2.2; the same pair will be used at every stage and for every pair of blocks. For $j = 0$, take the edgeless graph on one vertex. Suppose the statement holds at stage $j - 1$, and order $\mathcal{P}_j$ once and for all.

Internal edges and labels. For every $P \in \mathcal{P}_j$, take a block V_P of n_{j-1} vertices. Since

$$|[M_j] \setminus P| = (j - 1)t = M_{j-1},$$

place a copy of κ_{j-1} on V_P , relabeling its color set bijectively by $[M_j] \setminus P$. A color c is *active* on V_P if $c \notin P$ and *missing* if $c \in P$. For each active c , the inductive invariant supplies a proper coloring

$$\ell_c^P : V_P \longrightarrow [j]$$

of the internal graph of color c . In particular,

$$\kappa_j(uu') = c \implies \ell_c^P(u) \neq \ell_c^P(u') \quad (u, u' \in V_P). \quad (13)$$

Cross edges. For each ordered pair $P < Q$, fix distinct colors

$$a_1, \dots, a_s \in Q \setminus P, \quad b_1, \dots, b_s \in P \setminus Q,$$

which is possible by [Lemma 2.3](#). The a_d are active on V_P and missing on V_Q ; the b_d have the opposite status. For $u \in V_P$ and $v \in V_Q$, form

$$x(u) = (\ell_{a_d}^P(u))_{d=1}^s, \quad y(v) = (\ell_{b_d}^Q(v))_{d=1}^s.$$

These belong to $[H]^s$ because $j \leq H$.

If $x_d(u) = (f(y(v)))_d$ for some d , assign uv the color a_d for the least such d . Otherwise [Lemma 2.2](#) supplies a d with $y_d(v) = (g(x(u)))_d$; assign uv the color b_d for the least such d . In particular,

$$\kappa_j(uv) \in P \triangle Q. \quad (14)$$

More importantly, the label at the active endpoint is determined by the opposite endpoint:

$$\kappa_j(uv) = a_d \implies \ell_{a_d}^P(u) = (f(y(v)))_d, \quad (15)$$

$$\kappa_j(uv) = b_d \implies \ell_{b_d}^Q(v) = (g(x(u)))_d. \quad (16)$$

Triangles in one or two blocks. A triangle inside one block is not monochromatic by induction. Suppose next that $u, u' \in V_P$ and $v \in V_Q$, with $\kappa_j(uv) = \kappa_j(u'v) = c$. If $c \in P$, the color is missing internally on V_P , so $\kappa_j(uu') \neq c$.

If instead $c \notin P$, it is active on V_P . When $P < Q$, the distinctness of the a_d gives a unique d with $c = a_d$, and (15) yields

$$\ell_c^P(u) = (f(y(v)))_d = \ell_c^P(u').$$

When $Q < P$, the distinctness of the relevant b_d and (16) give exactly the same conclusion. Now (13) shows that the internal edge uu' cannot have color c .

Triangles in three blocks. Suppose a triangle of color c meets three distinct blocks V_P, V_Q, V_R . By (14),

$$c \in P \triangle Q, \quad c \in Q \triangle R, \quad c \in P \triangle R.$$

By symmetry, suppose $c \in P$. The first inclusion gives $c \notin Q$, and the second then gives $c \in R$. But this contradicts $c \in P \triangle R$. The coloring is therefore triangle-free.

The next proper-coloring invariant. For a fixed $c \in [M_j]$, define

$$L_c(v) = \begin{cases} \ell_c^P(v), & v \in V_P, c \notin P, \\ j+1, & v \in V_P, c \in P. \end{cases}$$

Inside an active block, L_c properly colors the c -edges; inside a missing block, there are no c -edges. By (14), every cross edge of color c joins an active block to a missing block. Its endpoints therefore receive labels in $[j]$ and $\{j+1\}$, respectively. Thus L_c is a proper $(j+1)$ -coloring of $\Gamma_{\kappa_j}(c)$. Finally, the B_j blocks each have n_{j-1} vertices, proving $n_j = B_j n_{j-1}$ and completing the induction. $\square$

Proof of Theorem 1.1. By decreasing c , it suffices to prove (3) for sufficiently large k . For $H \geq 3$, write

$$m(H) = \lceil 2H \log H \rceil, \quad s(H) = m(H)(m(H) + 1) + 1, \quad k_H = Hs(H) \lceil \log H \rceil.$$

These special color counts satisfy $k_H \asymp H^3 \log^3 H$ and $k_{H+1}/k_H = 1 + O(1/\log H)$. If $k_H \leq k < k_{H+1}$, monotonicity therefore shows that a bound $R_{k_H}(3) \geq (c_0 H)^{k_H}$ implies $R_k(3) \geq (c_1 H)^k$ for another absolute constant $c_1 > 0$: the factor $k_H/k = 1 - O(1/\log H)$ changes $\log(c_0 H)$ by only an additive constant. Moreover, $k < k_{H+1} = O(H^3 \log^3 H)$ implies $H \gg k^{1/3}/\log k$. Hence, after decreasing c again, it suffices to consider

$$k = k_H = Hs \lceil \log H \rceil$$

for sufficiently large H , where $s = s(H)$.

At stage H , [Proposition 3.1](#) produces a triangle-free k -coloring on $\prod_{j=1}^H B_j$ vertices. Recall that $t = s\lceil \log H \rceil$, $s = O(H^2 \log^2 H)$, $k \asymp H^3 \log^3 H$, and $\log(H!) \geq H \log H - H$. Multiplying [\(12\)](#) therefore gives, for sufficiently large k and a sufficiently small absolute constant $c > 0$,

$$R_k(3) \geq \frac{(H!)^{t-s}}{s^H (e^2 \lceil \log H \rceil^2)^{sH}} \geq \left(\frac{ck^{1/3}}{\log k} \right)^k.$$

This proves [\(3\)](#). Since $k^{1/3}/\log k \rightarrow \infty$, [\(4\)](#) follows. $\square$

REFERENCES

- [ACPPRT21] R. Ageron, P. Casteras, T. Pellerin, Y. Portella, A. Rimmel, and J. Tomasik, *New lower bounds for Schur and weak Schur numbers*, 2021, [arXiv:2112.03175](#).
- [ABST20] N. Alon, O. Ben-Eliezer, C. Shangguan, and I. Tamo, *The hat guessing number of graphs*, J. Combin. Theory Ser. B **144** (2020), 119–149, [doi:10.1016/j.jctb.2020.01.003](#); see also [arXiv:1812.09752](#).
- [AO95] N. Alon and A. Orlitsky, *Repeated communication and Ramsey graphs*, IEEE Trans. Inform. Theory **41** (1995), 1276–1289.
- [Blo183] T. F. Bloom, *Erdős problem #183*, <https://www.erdosproblems.com/183>; see also <https://www.erdosproblems.com/latex/183>.
- [CRRS06] S. Chakraborty, J. Radhakrishnan, N. Raghunathan, and P. Sasatte, *Zero error list-decoding capacity of the $q/(q-1)$ channel*, in *Foundations of Software Technology and Theoretical Computer Science*, Lecture Notes in Comput. Sci. **4337**, Springer, 2006, 129–138.
- [Chu73] F. R. K. Chung, *On the Ramsey numbers $N(3, 3, \dots, 3; 2)$* , Discrete Math. **5** (1973), 317–321, [doi:10.1016/0012-365X\(73\)90125-8](#).
- [CG] F. Chung and R. Graham, *Multi-color Ramsey number for triangles*, in *Erdős Problems*, <https://mathweb.ucsd.edu/~erdosproblems/erdos/newproblems/MulticolorR3.html>.
- [CG83] F. R. K. Chung and C. M. Grinstead, *A survey of bounds for classical Ramsey numbers*, J. Graph Theory **7** (1983), 25–37, [doi:10.1002/jgt.3190070105](#).
- [CFS15] D. Conlon, J. Fox, and B. Sudakov, *Recent developments in graph Ramsey theory*, in *Surveys in Combinatorics 2015*, London Math. Soc. Lecture Note Ser. **424**, Cambridge University Press, 2015, [arXiv:1501.02474](#).
- [EMT71] P. Erdős, R. J. McEliece, and H. Taylor, *Ramsey bounds for graph products*, Pacific J. Math. **37** (1971), 45–46.
- [Exo94] G. Exoo, *A lower bound for Schur numbers and multicolor Ramsey numbers*, Electron. J. Combin. **1** (1994), R8.
- [FS00] H. Fredricksen and M. M. Sweet, *Symmetric sum-free partitions and lower bounds for Schur numbers*, Electron. J. Combin. **7** (2000), R32.
- [Gow00] W. T. Gowers, *Rough structure and classification*, Geom. Funct. Anal., Special Volume, Part I (2000), 79–117; reprinted in *Visions in Mathematics*, Birkhäuser, 2010, [doi:10.1007/978-3-0346-0422-2_4](#).
- [GRS90] R. L. Graham, B. L. Rothschild, and J. H. Spencer, *Ramsey Theory*, 2nd ed., Wiley-Interscience, 1990.
- [GG55] R. E. Greenwood and A. M. Gleason, *Combinatorial relations and chromatic graphs*, Canad. J. Math. **7** (1955), 1–7, [doi:10.4153/CJM-1955-001-4](#).
- [Rad] S. P. Radziszowski, *Small Ramsey numbers*, Electron. J. Combin., Dynamic Survey DS1, <https://www.combinatorics.org/ojs/index.php/eljc/article/view/DS1>.
- [Wan97] H. Wan, *Upper bounds for Ramsey numbers $R(3, 3, \dots, 3)$ and Schur numbers*, J. Graph Theory **26** (1997), 119–122.
- [Whi73] E. G. Whitehead, Jr., *The Ramsey number $N(3, 3, 3, 3; 2)$* , Discrete Math. **4** (1973), 389–396.
- [XXC02] X.-D. Xu, Z. Xie, and Z. Chen, *Upper bounds for Ramsey numbers $R_n(3)$ and Schur numbers*, Math. Econ. **19** (2002), 81–84.

CHAPTER 10

Counterexamples to the Compactness and Degeneracy Conjectures for Extremal Numbers

ABSTRACT. We give two complementary counterexamples in extremal graph theory. First, we construct a finite family $\mathcal{F}$ of connected bipartite graphs, each containing a cycle, for which

$$\text{ex}(n, \mathcal{F}) = O(n^{4/3-1/48}) \quad \text{but} \quad \text{ex}(n, F) = \Omega(n^{4/3}) \quad (F \in \mathcal{F}).$$

This disproves the Erdős–Simonovits compactness conjecture. Second, we construct a fixed connected bipartite 2-degenerate graph H and fixed constants $c, \varepsilon > 0$ such that

$$\text{ex}(n, H) \geq cn^{3/2+\varepsilon} \quad \text{for all sufficiently large } n.$$

This disproves a conjecture of Erdős on extremal numbers of r -degenerate graphs.

CONTENTS

1. Introduction
 2. The forbidden graphs for compactness
 3. Proof of the upper bound in Theorem 1.1
 4. The generalized-quadrangle witnesses
 5. A binary entropy inequality
 6. The layered graph and the parameter thresholds
 7. A sampled Hamming-ball graph
 8. Excluding the forbidden graph
- References

1. INTRODUCTION

We disprove the compactness conjecture of Erdős and Simonovits and a degeneracy conjecture of Erdős in extremal graph theory. The compactness conjecture [ES82, Wig] asks whether forbidding a finite family of graphs, each containing a cycle, can reduce the extremal number by more than a constant factor compared with forbidding any individual member. The degeneracy conjecture [Erd67] predicts that every bipartite r -degenerate graph has extremal number $O(n^{2-1/r})$, which we disprove even when $r = 2$.

Throughout, all graphs are finite, simple, and undirected, and implicit constants may depend on fixed forbidden graphs or families. For a graph G , let $V(G)$ and $E(G)$ denote its vertex and edge sets, respectively. Given a family $\mathcal{F}$ of graphs, we say G is $\mathcal{F}$ -free if it contains no member of $\mathcal{F}$ as a subgraph. The extremal number of $\mathcal{F}$ is

$$\text{ex}(n, \mathcal{F}) := \max\{|E(G)| : |V(G)| = n, G \text{ is } \mathcal{F}\text{-free}\}.$$

For a single graph H , we abbreviate $\text{ex}(n, \{H\})$ to $\text{ex}(n, H)$.

1.1. The compactness conjecture. Erdős and Simonovits conjectured that the extremal number of a finite family of graphs is, up to a constant factor, the extremal number of one of its members [ES82, FS13, Con26, Blo575]. The original formulation admits simple counterexamples [Wig]. For example, the folklore family $\{K_{1,2}, 2K_2\}$ [Blo180] satisfies, for $n \geq 4$,

$$\text{ex}(n, \{K_{1,2}, 2K_2\}) = 1, \quad \text{ex}(n, K_{1,2}) = \lfloor n/2 \rfloor, \quad \text{ex}(n, 2K_2) = n - 1.$$

The corrected form of the conjecture (see [Wig]) asks: for every finite nonempty family of graphs $\mathcal{F}$, all of whose members contain cycles, do there exist $F \in \mathcal{F}$ and $C > 0$ such that

$$\text{ex}(n, \mathcal{F}) \leq C \text{ex}(n, F) \quad \text{for all sufficiently large } n? \quad (1)$$

Conlon, Mulrenin, and Pohoata [CMP26] recently disproved a conjecture of Verstraëte [Ver16, Conj. VIII], a stronger host-graph analog of compactness for even cycles, but their examples do not resolve the compactness conjecture itself. We disprove this conjecture even when every member of $\mathcal{F}$ is connected and bipartite.

Theorem 1.1 (Failure of compactness). *There exists a finite nonempty family $\mathcal{F}$ of connected bipartite graphs, every member of which contains a cycle, such that, for $\varepsilon = 1/48$,*

$$\text{ex}(n, \mathcal{F}) = O(n^{4/3-\varepsilon}) \quad \text{and} \quad \text{ex}(n, F) = \Omega(n^{4/3}) \quad (F \in \mathcal{F}). \quad (2)$$

In particular, no member of $\mathcal{F}$ satisfies (1).

1.2. The degeneracy conjecture. A graph H is r -degenerate if every nonempty subgraph of H has a vertex of degree at most r . Erdős [Erd67, Erd97, Blo146] conjectured that every fixed bipartite r -degenerate graph H satisfies

$$\text{ex}(n, H) = O(n^{2-1/r}). \quad (3)$$

Alon, Krivelevich, and Sudakov proved the weaker general estimate $\text{ex}(n, H) = O(n^{2-1/(4r)})$ for bipartite r -degenerate graphs [AKS03, Thm. 3.5].

The conjecture is known in several cases: when one bipartition class has maximum degree at most r [Fur91, AKS03]; for r -degenerate blow-ups of trees [GJN22]; and, when $r = 2$, for grids [BJST23] and certain critical 2-degenerate graphs [DGL25].

A related conjecture of Erdős [Erd81, Blo113] asserts that a bipartite graph H is 2-degenerate if and only if $\text{ex}(n, H) = O(n^{3/2})$. Janzer [Jan23] disproved the reverse implication by constructing, for every $\eta > 0$, a 3-regular bipartite graph H with $\text{ex}(n, H) = O(n^{4/3+\eta})$. Janzer's construction does not address the forward implication, which is the $r = 2$ case of (3).

We disprove Erdős's degeneracy conjecture by constructing a counterexample for $r = 2$, thereby also refuting the forward implication of the related conjecture.

Theorem 1.2 (Failure of the 2-degenerate bound). *There exist a fixed connected bipartite 2-degenerate graph H and constants $c, \varepsilon > 0$ such that*

$$\text{ex}(n, H) \geq cn^{3/2+\varepsilon}$$

for all sufficiently large n .

1.3. Proof strategy. For the compactness conjecture, we select $\mathcal{F}$ to consist of C_4 , C_6 , and certain admissible quotients of two fixed templates built from subdivided complete bipartite graphs, in analogy with the rooted-power construction of Bukh and Conlon [BC18]. We count short paths in a graph avoiding this family to obtain the upper bound $\text{ex}(n, \mathcal{F}) = O(n^{21/16})$. For each individual forbidden graph $F \in \mathcal{F}$, we obtain the lower bound $\text{ex}(n, F) = \Omega(n^{4/3})$ from the incidence graph of a generalized quadrangle, allowing the characteristic of the underlying field to depend on the forbidden graph. We develop this counterexample in Sections 2 to 4.

For the degeneracy conjecture, we construct H in layers, adjoining a vertex for every pair of vertices in the preceding layer. This construction is related to the complete degenerate graphs of Grzesik, Janzer, and Nagy [GJN22]. To obtain a lower bound on $\text{ex}(n, H)$, we independently retain vertices of a bipartite graph defined by Hamming distance. An embedding of H would increase a bounded entropy potential by a fixed amount at each layer, which is impossible after sufficiently many layers. A second-moment argument shows that the sampled graph has $\Omega(n^{3/2+\varepsilon})$ edges, and padding extends the construction to every sufficiently large order. We develop this counterexample in Sections 5 to 8.

2. THE FORBIDDEN GRAPHS FOR COMPACTNESS

The compactness family arises from two properly 2-colored bipartite templates.

Definition 2.1 (Subdivisions). For $k \in \{2, 3\}$, form S_k from $K_{3,k}$ by replacing each edge with a two-edge path; therefore $|V(S_2)| = 11$ and $|V(S_3)| = 15$. Call the three original vertices *bases*, the other k original vertices *centers*, and the inserted vertices *subdivision vertices*.

Definition 2.2 (Admissible identifications). Let T be a properly two-colored graph with distinguished subgraphs T_1, T_2 . An equivalence relation $\approx$ on $V(T)$ is *admissible* if

- (1) $v \approx w$ implies that v and w have the same color; and
- (2) for $i \in \{1, 2\}$ and $v, w \in V(T_i)$, $v \approx w$ implies $v = w$.

Write $T/\approx$ for the simple graph whose vertices are the equivalence classes and in which $[v][w]$ is an edge whenever some representatives are adjacent in T . Repeated edges are suppressed.

The first condition keeps $T/\approx$ bipartite and prevents loops. The second ensures that the maps $T_i \rightarrow T/\approx$ are injective.

Definition 2.3 (The first template). Form J_0 from two copies of S_2 with base triples $\{x, y, z\}$ and $\{x', y, z\}$ by identifying y, z , keeping all other vertices distinct, and adjoining a color-one vertex λ adjacent to x, x' . Thus, $|V(J_0)| = 21$. Let $\mathcal{J}$ consist of the uncolored admissible quotients $J_0/\approx$ satisfying $x \not\approx x'$. This condition preserves all four distinguished bases; λ may be identified whenever admissibility permits.

Definition 2.4 (The second template). Let K_0 be obtained from disjoint copies $S_3^{(1)}, S_3^{(2)}$ of S_3 by reversing the coloring on $S_3^{(2)}$ and adding an edge $d_1 d_2$ between specified centers of the two copies. Thus, $|V(K_0)| = 30$. Let $\mathcal{K}$ consist of the uncolored graphs $K_0/\approx$, where $\approx$ is admissible for $S_3^{(1)}, S_3^{(2)}$.

The two templates are shown in Figure 1.

Definition 2.5 (The forbidden family). Let

$$\mathcal{F} := \{C_4, C_6\} \cup \mathcal{J} \cup \mathcal{K}. \quad (4)$$

As each member of $\mathcal{J}$ contains S_2 and each member of $\mathcal{K}$ contains S_3 , every graph in $\mathcal{F}$ contains a cycle.




FIGURE 1. The graphs J_0 and K_0 .

3. PROOF OF THE UPPER BOUND IN THEOREM 1.1

Throughout this section, B is a bipartite graph containing no C_4 or C_6 , and S is either side of the bipartition of B . Let R_S be the auxiliary *common-neighbor graph* on S , defined by

$$uv \in E(R_S) \iff u \neq v \text{ and } N_B(u) \cap N_B(v) \neq \emptyset.$$

The graph R_S records pairs joined by length-two paths in B , which are central to the forbidden configurations.

For distinct $u, v \in S$, write $u \sim v$ when $uv \in E(R_S)$, and call u, v *related* or *unrelated* according as they are adjacent or nonadjacent in R_S . Call a set of vertices R_S -*independent* if no pair is adjacent in R_S . For an R_S -independent triple $T \subseteq S$, put

$$L(T) := \{c \in S : c \sim t \text{ for all } t \in T\} \text{ and } r(T) := |L(T)|.$$

Lemma 3.1 (Common-neighbor geometry). *The following assertions hold:*

- (1) *Every related pair has a unique common neighbor. If three distinct members of S are pairwise related, the three pairs have the same common neighbor.*
- (2) *If $T \subseteq S$ is an R_S -independent triple, then $L(T)$ is R_S -independent and disjoint from T . If $r(T) \geq k$, for $k \in \{2, 3\}$, then B contains a copy of S_k with base set T . Conversely, the bases of any such copy form an R_S -independent triple with at least k common centers.*
- (3) *If $u, v \in S$ are distinct and unrelated, then any two distinct vertices related to both u and v are unrelated. Equivalently, the set*

$$Z_{uv} = \{w \in S : w \sim u, w \sim v\}$$

is R_S -independent.

Proof. A related pair has a unique common neighbor, since two would form a C_4 . For three pairwise related vertices, their common neighbors either all coincide or form a C_6 . This proves (1).

If $u \not\sim v$ and $w, w' \in Z_{uv}$ were related, (1) would make their common neighbor adjacent to both u and v , a contradiction. This proves (3).

For distinct $u, v \in T$, we have $L(T) \subseteq Z_{uv}$. Thus, (3) makes $L(T)$ independent, while irreflexivity gives $L(T) \cap T = \emptyset$. If $r(T) \geq k$, the unique common neighbors of the $3k$ pairs between T and any k centers in $L(T)$ are distinct: a repetition would relate two bases or two

centers. This forms a copy of S_k . Conversely, (1) forces related bases in such a copy to share a subdivision vertex at every center, a contradiction. Thus, its bases are independent and its k centers belong to $L(T)$, proving (2). $\square$

Lemma 3.2 (Non-backtracking four-edge-walk count). *If $\delta(B) \geq d \geq 2$, then, for every $u \in S$,*

$$\sum_{\substack{v \in S \setminus \{u\} \\ v \not\sim u}} |Z_{uv}| \geq d(d-1)^3. \quad (5)$$

Proof. There are at least $d(d-1)^3$ non-backtracking walks u, a, w, b, v . Since B has girth at least eight, each satisfies $v \neq u$ and $v \not\sim u$; otherwise it closes to a cycle of length at most six. By Lemma 3.1(1), these walks correspond bijectively to pairs (v, w) with $v \neq u$, $v \not\sim u$, and $w \in Z_{uv}$. Counting these pairs proves (5). $\square$

Lemma 3.3 (Minimum-degree reduction). *Let G be an n -vertex graph without C_4 or C_6 , and put $m = |E(G)| > 0$. There is a nonempty bipartite subgraph B of order $N \leq n$ and minimum degree d such that*

$$m \leq 2nd \quad \text{and} \quad \Delta(B)(d-1)^2 \leq N.$$

Proof. A maximum cut yields a bipartite subgraph with at least $m/2$ edges. Deleting vertices of degree $< m/(2n)$ cannot exhaust it, since fewer than $m/2$ edges would be removed. Thus, the remainder B satisfies $d = \delta(B) \geq m/(2n)$, and hence $m \leq 2nd$.

Since B has girth at least eight, the radius-three neighborhood of a maximum-degree vertex is a tree. Its third level contains at least $\Delta(B)(d-1)^2$ vertices, so $\Delta(B)(d-1)^2 \leq N$. $\square$

The main estimate uses the two forbidden families in complementary ways. Excluding $\mathcal{J}$ bounds the number of vertices that fail to be centers of a subdivided $K_{3,3}$, while excluding $\mathcal{K}$ forces every edge to meet one of these vertices.

Proposition 3.4 (Key counting estimate). *For the family $\mathcal{F}$ in (4),*

$$\text{ex}(n, \mathcal{F}) = O(n^{21/16}) = O(n^{4/3-1/48}).$$

Proof. Let G be an n -vertex $\mathcal{F}$ -free graph. We may assume $m := |E(G)| \geq Cn^{21/16}$, where $C \geq 1$ is sufficiently large. Let B be the bipartite subgraph supplied by Lemma 3.3, with order $N \leq n$ and minimum degree d . Put

$$d \geq \frac{m}{2n} \gg CN^{5/16}, \quad p = d(d-1)^3, \quad R = \frac{p}{2N}.$$

Fix a bipartition class S and $u \in S$, and write $t_{uv} = |Z_{uv}|$. Lemma 3.1 shows that $C_u := \sum_{\substack{v \neq u \\ v \not\sim u}} \binom{t_{uv}}{3}$ counts the copies of S_2 with center u . By discrete convexity of binomial coefficients, we have

$$C_u \geq N \binom{N^{-1} \sum_{\substack{v \neq u \\ v \not\sim u}} t_{uv}}{3} \gg \frac{d^{12}}{N^2},$$

where Lemma 3.2 gives $\sum_{\substack{v \neq u \\ v \not\sim u}} t_{uv} \geq d(d-1)^3$.

We use the exclusion of $\mathcal{J}$ to bound the number of possible base triples. For an unrelated pair $y, z \in S$, let

$$A_{yz} := \{x \in S \setminus \{y, z\} : \{x, y, z\} \text{ is independent in } R_S, r(\{x, y, z\}) \geq 2\}.$$

This is the set of possible third bases for copies of S_2 whose other bases are y, z . If two distinct vertices of A_{yz} were related, their witnessing copies of S_2 , together with their common neighbor, would produce a member of $\mathcal{J}$. Thus, A_{yz} is R_S -independent; its vertices therefore have pairwise disjoint neighborhoods, and $|A_{yz}| \leq N/d$.

Writing $\mathcal{T}_S$ for the set of independent triples in S with at least two common centers, we obtain

$$|\mathcal{T}_S| = \frac{1}{3} \sum_{\substack{\{y, z\} \subseteq S \\ y \not\sim z}} |A_{yz}| \leq \frac{N}{3d} \binom{N}{2} \leq \frac{N^3}{6d}. \quad (6)$$

Let $U \subseteq V(B)$ be the set of vertices that are not centers of any copy of S_3 . For every bipartition class S and every $u \in S$, we have

$$C_u = \sum_{\substack{T \in \mathcal{T}_S \\ u \in L(T)}} (r(T) - 1).$$

Indeed, for each triple T with $u \in L(T)$, choosing another center $v \in L(T) \setminus \{u\}$ determines a copy of S_2 with base triple T and centers u, v ; conversely, every copy counted by C_u arises uniquely in this way.

If $u \in U \cap S$ and $u \in L(T)$, then necessarily $r(T) = 2$: otherwise, u together with two other members of $L(T)$ would be the centers of a copy of S_3 , by [Lemma 3.1](#). Therefore, additionally using (6), we have

$$|U \cap S| \frac{d^{12}}{N^2} \ll \sum_{u \in U \cap S} C_u \leq 2|\mathcal{T}_S| \ll \frac{N^3}{d}.$$

Summing over the two bipartition classes gives

$$|U| \ll \frac{N^5}{d^{13}}. \quad (7)$$

We claim that U is a vertex cover of B . Otherwise, there would be an edge $uv \in E(B)$ with $u, v \notin U$. Choose a copy of S_3 having u as a center and a copy of S_3 having v as a center. Together with the edge uv , these copies form an admissible quotient of K_0 , and hence give a member of $\mathcal{K}$, contrary to the choice of B . Thus, every edge of B meets U .

By [Lemma 3.3](#) we have $\Delta(B) \ll \frac{N}{d^2}$. Since U is a vertex cover, (7) implies

$$Nd \leq 2e(B) \leq 2|U|\Delta(B) \ll \frac{N^6}{d^{15}}.$$

Therefore $d^{16} \ll N^5$, which is a contradiction provided that $C \geq 1$ is a sufficiently large constant. $\square$

4. THE GENERALIZED-QUADRANGLE WITNESSES

We now proceed to the lower bound side of [Theorem 1.1](#). Since the required constructions already appear in the literature (e.g. [\[PT09\]](#) and [\[BHKT18\]](#)), we keep the proof brief.

Let q be a prime power, let $V = \mathbb{F}_q^4$, and fix a nondegenerate alternating bilinear form

$$\langle \cdot, \cdot \rangle : V \times V \longrightarrow \mathbb{F}_q.$$

The symplectic generalized quadrangle $W(q)$ has point and line classes

$$\mathcal{P} = \{P \leq V : \dim P = 1\} \quad \text{and} \quad \mathcal{M} = \{M \leq V : \dim M = 2, \langle M, M \rangle = 0\},$$

with incidence given by containment. Let

$$I_q := I(W(q))$$

denote its bipartite incidence graph: its vertex set is $\mathcal{P} \sqcup \mathcal{M}$, and $P \in \mathcal{P}$ is adjacent to $M \in \mathcal{M}$ precisely when $P \subseteq M$. This is the standard symplectic construction of $W(q)$; see [\[PT09, §3.2.1\]](#) and [\[BHKT18, §4\]](#).

The quadrangle has order (q, q) : every point lies on $q + 1$ lines, every line contains $q + 1$ points, and

$$|\mathcal{P}| = |\mathcal{M}| = (q + 1)(q^2 + 1).$$

Its incidence graph has girth eight [\[PT09, §1.1, §1.2.1, and §3.1.1\]](#). Consequently,

$$n_q = 2(q + 1)(q^2 + 1), \quad e_q = (q + 1)^2(q^2 + 1) \geq 2^{-4/3} n_q^{4/3}. \quad (8)$$

For $S \in \{\mathcal{P}, \mathcal{M}\}$, retain the common-neighbor graph R_S , the sets $L(T)$, and the quantities $r(T)$ from the preceding section, with $B = I_q$. Explicitly, for distinct vertices in the same bipartition class,

$$P \sim P' \iff \langle P, P' \rangle = 0 \quad (P, P' \in \mathcal{P}),$$

and

$$M \sim M' \iff M \cap M' \neq \{0\} \quad (M, M' \in \mathcal{M}).$$

We will also use two classical facts. The dual generalized quadrangle $W(q)^D$ is isomorphic to the parabolic quadrangle $Q(4, q)$, and $W(q)$ is self-dual when q is even [PT09, §3.2.1]; see also [BHKT18, §4].

Definition 4.1 (Local patterns). A J -pattern in S consists of independent triples $\{x, y, z\}$ and $\{x', y, z\}$, each with at least two common centers, such that $x \sim x'$. An S_3 -pattern in S is an R_S -independent triple with at least three common centers.

A member of $\mathcal{J}$ forces a J -pattern in one bipartition class, while a member of $\mathcal{K}$ forces an S_3 -pattern in both bipartition classes.

Proposition 4.2 (Quadrangle witnesses). *For even q , I_q is $\mathcal{J}$ -free; for odd q , it is $\mathcal{K}$ -free. In either case, it contains no C_4 or C_6 .*

Proof. Let $y, z \in \mathcal{P}$ be noncollinear, and put $U = y + z$. Since y and z are not orthogonal, U is a nondegenerate two-dimensional subspace of V . Their common centers are precisely the projective points of $U^\perp$. If a third point x has at least two common centers with y, z , then x is orthogonal to two distinct projective points of $U^\perp$, and therefore to all of $U^\perp$. Hence,

$$x \subseteq (U^\perp)^\perp = U.$$

Any two distinct projective points of the nondegenerate symplectic plane U are nonorthogonal. Therefore, any two possible third bases are unrelated, and $\mathcal{P}$ contains no J -pattern.

If q is even, the self-duality of $W(q)$ [PT09, §3.2.1] gives the same conclusion for $\mathcal{M}$. Thus, neither bipartition class contains a J -pattern, and I_q is $\mathcal{J}$ -free.

If q is odd, identify $\mathcal{M}$ with the point class of the dual quadrangle $Q(4, q)$. Every triad in $Q(4, q)$ has either zero or two centers [BHKT18, Proposition 4.5]; see also [PT09, §1.3.6(iii) and §3.3.1]. Consequently, $\mathcal{M}$ contains no S_3 -pattern, and therefore I_q is $\mathcal{K}$ -free.

Finally, the girth-eight property excludes C_4 and C_6 . □

Proposition 4.3 (Individual lower bounds). *For every $F \in \mathcal{F}$,*

$$\text{ex}(n, F) = \Omega(n^{4/3}).$$

Proof. If $F \in \{C_4, C_6\} \cup \mathcal{J}$, take $q = 2^j$; if $F \in \mathcal{K}$, take $q = 3^j$. In either case, Proposition 4.2 makes I_q F -free. For $t \in \{2, 3\}$,

$$n_{tq} = 2(tq + 1)(t^2q^2 + 1) \leq t^3n_q.$$

Choose the largest admissible q for which $n_q \leq n$. The displayed inequality gives $n_q \gg n$. Padding I_q with isolated vertices preserves F -freeness, since F is connected and contains an edge. Hence, (8) gives

$$\text{ex}(n, F) \geq e_q \gg n^{4/3}. \quad \square$$

Proof of Theorem 1.1. The family in (4) is finite and nonempty, and each of its members is connected, bipartite, and contains a cycle. Propositions 3.4 and 4.3 give (2); in particular, (1) cannot hold for any $F \in \mathcal{F}$. □

5. A BINARY ENTROPY INEQUALITY

We now proceed to the proof of Theorem 1.2. The graph H consists of layers $L_1, L_2, \dots$, with every vertex in L_i having two parents in L_{i-1} . The host graph is a variant of the Hamming cube in which points at most a fixed linear distance apart are adjacent. This construction naturally leads to the entropy inequalities considered in this section.

All logarithms and entropies are to base two. Write

$$h(x) := -x \log_2 x - (1 - x) \log_2 (1 - x), \quad h(0) = h(1) = 0,$$

and put

$$\kappa := \frac{3}{2} - \frac{3}{4} \log_2 3.$$

The following inequality converts a bound on the average disagreement between a child and its parents into a bound on its conditional entropy.

Lemma 5.1 (Independent-parent entropy). *Let X, Y be independent Bernoulli variables with common parameter q , and let Z be any jointly distributed binary variable. Write*

$$v = \mathbb{P}(Z = 1), \quad d = \frac{\mathbb{P}(X \neq Z) + \mathbb{P}(Y \neq Z)}{2}.$$

Then

$$H(Z | X, Y) \leq \kappa + (\log_2 3)d + \frac{h(v) - h(q)}{2}. \quad (9)$$

Proof. For a probability vector (r_0, r_1) and nonnegative numbers t_0, t_1 , the Gibbs inequality gives

$$-\sum_{z \in \{0,1\}} r_z \log_2 r_z + \sum_{z \in \{0,1\}} r_z \log_2 t_z \leq \log_2(t_0 + t_1). \quad (10)$$

Indeed, after writing $T = t_0 + t_1$, the difference between the right-hand and left-hand sides is

$$\sum_{z \in \{0,1\}} r_z \log_2 \left(\frac{r_z}{t_z/T} \right) \geq 0.$$

Put

$$c(x, y, z) = \frac{\mathbf{1}_{x \neq z} + \mathbf{1}_{y \neq z}}{2}, \quad \pi_z = \mathbb{P}(Z = z).$$

For fixed x, y , apply (10) with

$$r_z = \mathbb{P}(Z = z | X = x, Y = y), \quad t_z = \sqrt{\pi_z} 3^{-c(x,y,z)}.$$

Averaging over X, Y gives

$$H(Z | X, Y) - \frac{h(v)}{2} - (\log_2 3)d \leq \mathbb{E} \log_2 \left(\sum_{z=0}^1 \sqrt{\pi_z} 3^{-c(X,Y,z)} \right). \quad (11)$$

It remains to bound the right-hand side. Set

$$a = \sqrt{1-v}, \quad b = \sqrt{v}, \quad u = 2q - 1,$$

so that $a^2 + b^2 = 1$. The three possible types for $\{X, Y\}$, namely 00, 11, and 01, have probabilities

$$\alpha_0 = \frac{(1-u)^2}{4}, \quad \alpha_1 = \frac{(1+u)^2}{4}, \quad \alpha_* = \frac{1-u^2}{2},$$

and give respective expressions

$$t_0 = a + \frac{b}{3}, \quad t_1 = \frac{a}{3} + b, \quad t_* = \frac{a+b}{\sqrt{3}}.$$

Thus, after adding $h(q)/2$, the right-hand side of (11) becomes

$$F = \frac{1}{2}h(q) + \alpha_0 \log_2 t_0 + \alpha_1 \log_2 t_1 + \alpha_* \log_2 t_*.$$

Since

$$\log_2 t \leq \log_2 \lambda + (\log_2 e) \left(\frac{t}{\lambda} - 1 \right) \quad (t, \lambda > 0),$$

we have

$$\begin{aligned} \log_2 \left(a + \frac{b}{3} \right) &\leq \log_2 \left(\frac{4}{3\sqrt{2}} \right) + (\log_2 e) \left(\frac{3\sqrt{2}}{4} \left(a + \frac{b}{3} \right) - 1 \right), \\ \log_2 \left(\frac{a}{3} + b \right) &\leq \log_2 \left(\frac{4}{3\sqrt{2}} \right) + (\log_2 e) \left(\frac{3\sqrt{2}}{4} \left(\frac{a}{3} + b \right) - 1 \right), \\ \log_2 \left(\frac{a+b}{\sqrt{3}} \right) &\leq \log_2 \left(\sqrt{\frac{2}{3}} \right) + (\log_2 e) \left(\frac{a+b}{\sqrt{2}} - 1 \right). \end{aligned}$$

Multiplying these inequalities by

$$\alpha_0 = \frac{(1-u)^2}{4}, \quad \alpha_1 = \frac{(1+u)^2}{4}, \quad \alpha_* = \frac{1-u^2}{2},$$

respectively, and summing, gives

$$\begin{aligned} & \alpha_0 \log_2 \left(a + \frac{b}{3} \right) + \alpha_1 \log_2 \left(\frac{a}{3} + b \right) + \alpha_* \log_2 \left(\frac{a+b}{\sqrt{3}} \right) \\ & \leq 1 - \frac{3}{4} \log_2 3 + \frac{u^2}{4} \log_2 \frac{4}{3} \\ & \quad + (\log_2 e) \left(\frac{(2-u)a + (2+u)b}{2\sqrt{2}} - 1 \right) \\ & \leq 1 - \frac{3}{4} \log_2 3 + \frac{u^2}{4} \log_2 \frac{4}{3} + (\log_2 e) \left(\sqrt{1 + \frac{u^2}{4}} - 1 \right), \end{aligned}$$

where the final inequality follows from Cauchy–Schwarz and $a^2 + b^2 = 1$.

Therefore

$$F \leq \frac{1}{2} h \left(\frac{1+u}{2} \right) + 1 - \frac{3}{4} \log_2 3 + \frac{u^2}{4} \log_2 \frac{4}{3} + (\log_2 e) \left(\sqrt{1 + \frac{u^2}{4}} - 1 \right). \quad (12)$$

A direct optimization shows that the right-hand side of (12) is maximized at $u = 0$, where it equals κ . Together with (11), this proves (9). $\square$

In the layered construction, the two parents are sampled without replacement. The resulting error vanishes as the layer size tends to infinity.

Lemma 5.2 (Without-replacement correction). *Let $L \geq 4$, let $x_1, \dots, x_L \in \{0, 1\}$, and sample an ordered pair of distinct indices uniformly. Let X, Y be the corresponding bits, and let Z have any conditional binary law given X, Y . Put*

$$q = \frac{1}{L} \sum_{a=1}^L x_a, \quad v = \mathbb{P}(Z = 1), \quad d = \frac{\mathbb{P}(X \neq Z) + \mathbb{P}(Y \neq Z)}{2}.$$

Then

$$H(Z | X, Y) \leq \kappa + (\log_2 3)d + \frac{h(v) - h(q)}{2} + \eta(L),$$

where

$$\eta(L) = O \left(\frac{\log_2 L}{L} \right) = o(1)$$

uniformly over all choices of the bits and the conditional law of Z .

Proof. Couple sampling with and without replacement so that

$$\mathbb{P}((X, Y) \neq (X', Y')) \leq \frac{1}{L},$$

where X', Y' are independent Bernoulli variables with parameter q . Give Z' , conditionally on X', Y' , the same conditional law used to define Z . Then

$$|H(Z | X, Y) - H(Z' | X', Y')| \leq \frac{1}{L},$$

while

$$|d - d'| \leq \frac{1}{L}, \quad |v - v'| \leq \frac{1}{L}, \quad |h(v) - h(v')| \leq h(1/L).$$

Now apply Lemma 5.1 to (X', Y', Z') . $\square$

6. THE LAYERED GRAPH AND THE PARAMETER THRESHOLDS

The construction depends on a Hamming radius $\tau \in (0, 1/2)$ and a sampling exponent $\beta \in (0, 1)$. Define

$$A(\tau) := \kappa + \tau \log_2 3, \quad C(\tau) := 2h(\tau) - 1.$$

The threshold $A(\tau)$ controls exclusion of the layered graph, while $C(\tau)$ controls whether the sampled host has more than $n^{3/2}$ edges.

For now, assume that

$$A(\tau) < \beta < C(\tau). \quad (13)$$

We verify at the end that such parameters exist.

Choose

$$0 < \delta < \frac{\beta - A(\tau)}{4}.$$

Since $\eta(L) = o(1)$ and $\binom{L}{2} \asymp L^2$, fix $L_0 \geq 4$ sufficiently large that, for every $L \geq L_0$,

$$\eta(L) < \delta, \quad L + 3 \log_2 \left(\binom{L}{2} + 1 \right) - \delta \binom{L}{2} < -1. \quad (14)$$

Then choose an integer $s \geq 2$ such that

$$2s(\beta - A(\tau) - 2\delta) > 1. \quad (15)$$

Take disjoint layers $V_0, \dots, V_s$ with

$$|V_0| = L_0, \quad V_i = \binom{V_{i-1}}{2} \quad (1 \leq i \leq s).$$

Join each vertex $\{a, b\} \in V_i$ to its two parents $a, b \in V_{i-1}$, and let H be the resulting graph. Writing

$$L_i = |V_i|,$$

we have

$$L_i = \binom{L_{i-1}}{2} \geq L_{i-1} \geq L_0.$$

Fact 6.1. *The graph H is connected, bipartite, and 2-degenerate.*

7. A SAMPLED HAMMING-BALL GRAPH

For $m \geq 1$, put

$$U = \{0, 1\}^m, \quad Q = 2^m, \quad p = 2^{-\beta m}, \quad k = \lfloor \tau m \rfloor.$$

Take two disjoint copies U_L, U_R of U , joining vertices on opposite sides whenever their Hamming distance is at most k . Independently retain each vertex with probability p , and write G_m for the induced subgraph.

For a parent array

$$\mathbf{u} = (u_a)_{a=1}^L \in U^L$$

and a child array

$$\mathbf{z} = (z_{\{a,b\}})_{\{a,b\} \in \binom{[L]}{2}} \in U^M, \quad M = \binom{L}{2},$$

choose a uniformly random parent pair and orient it using an independent fair coin. In coordinate j , denote the resulting parent and child bits by X_j, Y_j, Z_j , and define

$$E(\mathbf{u}, \mathbf{z}) := \frac{1}{m} \sum_{j=1}^m H(Z_j \mid X_j, Y_j).$$

An array of conditional entropy E has at most $2^{mME + O(m \log_2 M)}$ realizations, while requiring its M children to survive sampling costs $2^{-\beta m M}$. Since $M \asymp L^2$, this dominates the 2^{mL} possible parent arrays whenever $E < \beta$.

Lemma 7.1 (Exclusion of low-entropy arrays). *With probability at least $1 - 2s2^{-m}$, simultaneously on both host sides and at every transition, there is no parent array of length L_{i-1} and pairwise distinct retained child array of length L_i satisfying*

$$E(\mathbf{u}, \mathbf{z}) \leq \beta - \delta. \quad (16)$$

Proof. Fix a host side, a transition, and a parent array. Write

$$L = L_{i-1}, \quad M = L_i = \binom{L}{2}.$$

For each coordinate j , partition the parent pairs according to their bit types

$$\mathcal{G} = \{00, 11, 01\},$$

where 01 includes both orientations of a mixed pair. Let $N_{g,j}$ be the size of group g , and let $b_{g,j}$ count its children having j th bit equal to one. Then

$$H(Z_j \mid X_j, Y_j) = \sum_{g \in \mathcal{G}} \frac{N_{g,j}}{M} h\left(\frac{b_{g,j}}{N_{g,j}}\right).$$

There are at most $(M+1)^{3m}$ possible profiles $(b_{g,j})_{g,j}$. For a fixed profile, the number of child arrays, allowing repeated words, is

$$\prod_{j=1}^m \prod_{g \in \mathcal{G}} \binom{N_{g,j}}{b_{g,j}} \leq 2^{mME(\mathbf{u}, \mathbf{z})},$$

by the estimate

$$\binom{N}{r} \leq 2^{Nh(r/N)}.$$

Therefore, at most

$$(M+1)^{3m} 2^{mM(\beta-\delta)}$$

child arrays satisfy (16).

For pairwise distinct children, the probability that they are all retained is

$$p^M = 2^{-\beta mM}.$$

There are $Q^L = 2^{mL}$ possible parent arrays, so (14) gives

$$\begin{aligned} \mathbb{P}(\text{a bad retained array}) &\leq Q^L (M+1)^{3m} 2^{mM(\beta-\delta)} p^M \\ &= 2^{m[L+3\log_2(M+1)-\delta M]} \\ &\leq 2^{-m}. \end{aligned}$$

Take a union bound over both host sides and all s transitions. $\square$

8. EXCLUDING THE FORBIDDEN GRAPH

Proposition 8.1 (Exclusion of the layered graph). *With probability $1 - o(1)$, the sampled graph G_m is H -free.*

Proof. Assume the event in Lemma 7.1, and suppose that

$$\iota : H \hookrightarrow G_m$$

is an embedding. Since H is connected and bipartite, consecutive layers occupy opposite sides of the host.

For the parent and child arrays at transition i , write

$$E_i = E(\mathbf{u}_i, \mathbf{z}_i).$$

The children are distinct and retained, so Lemma 7.1 gives

$$E_i > \beta - \delta. \quad (17)$$

Define the entropy potential of layer i by

$$\Phi_i = \frac{1}{m} \sum_{j=1}^m h \left(\frac{1}{L_i} \sum_{z \in \iota(V_i)} z(j) \right).$$

Since binary entropy belongs to $[0, 1]$,

$$0 \leq \Phi_i \leq 1.$$

At transition i , let

$$d_j = \frac{\mathbb{P}(X_j \neq Z_j) + \mathbb{P}(Y_j \neq Z_j)}{2}.$$

Every child is adjacent to both parents, and therefore

$$\begin{aligned} \frac{1}{m} \sum_{j=1}^m d_j &= \frac{1}{2mL_i} \sum_{\{a,b\} \in \binom{V_{i-1}}{2}} (\text{dist}(u_a, z_{\{a,b\}}) + \text{dist}(u_b, z_{\{a,b\}})) \\ &\leq \frac{k}{m} \leq \tau. \end{aligned}$$

Applying [Lemma 5.2](#) in each coordinate and averaging gives

$$E_i \leq A(\tau) + \frac{\Phi_i - \Phi_{i-1}}{2} + \eta(L_{i-1}). \quad (18)$$

Combining (17), (18), and (14), we obtain

$$\Phi_i - \Phi_{i-1} > 2(\beta - A(\tau) - 2\delta).$$

Hence, (15) gives

$$\Phi_s - \Phi_0 > 2s(\beta - A(\tau) - 2\delta) > 1,$$

contradicting $0 \leq \Phi_i \leq 1$. Therefore,

$$\mathbb{P}(H \subseteq G_m) \leq 2s 2^{-m} = o(1). \quad \square$$

Proof of Theorem 1.2. Continue assuming (13); the existence of suitable parameters will be checked at the end.

Put

$$D_m = \sum_{j=0}^{\lfloor \tau m \rfloor} \binom{m}{j}.$$

The standard Hamming-ball estimate gives

$$D_m = 2^{(h(\tau) + o(1))m}.$$

For

$$W_m = |V(G_m)|, \quad e_m = e(G_m),$$

we have

$$\mathbb{E}W_m = 2pQ, \quad \text{Var}(W_m) \leq 2pQ,$$

and

$$\mathbb{E}e_m = p^2QD_m, \quad \text{Var}(e_m) \leq p^2QD_m + 2p^3QD_m^2.$$

Since $\beta < C(\tau) < 1$,

$$pQ = 2^{(1-\beta)m} \longrightarrow \infty$$

and

$$p^2QD_m = 2^{(1+h(\tau)-2\beta+o(1))m} \longrightarrow \infty.$$

Consequently,

$$W_m \leq 3pQ, \quad e_m \geq \frac{1}{2}p^2QD_m$$

with probability $1 - o(1)$.

Combining these estimates with [Proposition 8.1](#), and padding with isolated vertices, gives

$$\text{ex}(N_m, H) \geq \frac{1}{2}p^2QD_m, \quad N_m = \left\lceil 3 \cdot 2^{(1-\beta)m} \right\rceil.$$

Moreover,

$$\begin{aligned}\frac{1 + h(\tau) - 2\beta}{1 - \beta} &= \frac{3}{2} + \frac{2h(\tau) - 1 - \beta}{2(1 - \beta)} \\ &= \frac{3}{2} + \frac{C(\tau) - \beta}{2(1 - \beta)} > \frac{3}{2}.\end{aligned}$$

Therefore, for any fixed

$$0 < \varepsilon < \frac{C(\tau) - \beta}{2(1 - \beta)}, \quad (19)$$

we have

$$\text{ex}(N_m, H) \geq N_m^{3/2+\varepsilon}$$

for all sufficiently large m .

Since $2^{1-\beta} < 2$,

$$N_{m+1} \leq 2N_m.$$

Thus, for every sufficiently large n , there is an m such that

$$N_m \leq n < N_{m+1} \leq 2N_m.$$

Padding once more gives

$$\text{ex}(n, H) \geq \text{ex}(N_m, H) \geq 2^{-3/2-\varepsilon} n^{3/2+\varepsilon}.$$

It remains to verify that the parameter window (13) is nonempty. Its width is

$$f(\tau) = C(\tau) - A(\tau) = 2h(\tau) - 1 - \kappa - \tau \log_2 3.$$

Since

$$f'(\tau) = 2 \log_2 \left(\frac{1 - \tau}{\tau} \right) - \log_2 3,$$

the strictly concave function f is maximized at

$$\tau = \frac{1}{1 + \sqrt{3}}.$$

Set $r = \sqrt{3}$. At this value of τ ,

$$2h(\tau) - \tau \log_2 3 = 2 \log_2 \left(1 + \frac{1}{r} \right),$$

while

$$\kappa = \frac{1}{2} + \frac{1}{2} \log_2(1 + r^2) - \frac{3}{2} \log_2 r.$$

Hence,

$$\begin{aligned}f(\tau) &= \frac{1}{2} \log_2 \left(\frac{(1 + r)^4}{8r(1 + r^2)} \right) \\ &= \frac{1}{2} \log_2 \left(1 + \frac{(r - 1)^4}{8r(1 + r^2)} \right) > 0.\end{aligned}$$

We may therefore choose any

$$A(\tau) < \beta < C(\tau),$$

and subsequently any ε satisfying (19). The remaining assertions follow from [Fact 6.1](#). $\square$

REFERENCES

- [AKS03] N. Alon, M. Krivelevich, and B. Sudakov, *Turán numbers of bipartite graphs and related Ramsey-type questions*, *Combin. Probab. Comput.* **12** (2003), 477–494, <https://doi.org/10.1017/S0963548303005741>.
- [BHKT18] D. Bartoli, T. Héger, G. Kiss, and M. Takáts, *On the metric dimension of affine planes, biaffine planes and generalized quadrangles*, *Australas. J. Combin.* **72** (2018), 226–248, https://ajc.maths.uq.edu.au/pdf/72/ajc_v72_p226.pdf.
- [Blo113] T. F. Bloom, *Erdős problem #113*, <https://www.erdosproblems.com/113>.
- [Blo146] T. F. Bloom, *Erdős problem #146*, <https://www.erdosproblems.com/146>.
- [Blo180] T. F. Bloom, *Erdős problem #180*, <https://www.erdosproblems.com/180>.
- [Blo575] T. F. Bloom, *Erdős problem #575*, <https://www.erdosproblems.com/575>.
- [BJST23] D. Bradač, O. Janzer, B. Sudakov, and I. Tomon, *The Turán number of the grid*, *Bull. Lond. Math. Soc.* **55** (2023), 194–204, <https://doi.org/10.1112/blms.12721>.
- [BC18] B. Bukh and D. Conlon, *Rational exponents in extremal graph theory*, *J. Eur. Math. Soc.* **20** (2018), 1747–1757, <https://doi.org/10.4171/JEMS/798>.
- [Con26] D. Conlon, *Combinatorial theorems relative to sparse sets*, manuscript, 2026; to appear in the Proceedings of the International Congress of Basic Science, [available online](#).
- [CMP26] D. Conlon, E. Mulrenin, and C. Pohoata, *Two counterexamples to a conjecture about even cycles*, 2026, [arXiv:2603.24515](https://arxiv.org/abs/2603.24515).
- [DGL25] Z. Dong, J. Gao, and H. Liu, *Bipartite Turán problems via graph gluing*, *Bull. Lond. Math. Soc.* **57** (2025), 3783–3796, <https://doi.org/10.1112/blms.70243>.
- [Erd67] P. Erdős, *Some recent results on extremal problems in graph theory*, in *Theory of Graphs* (International Symposium, Rome, 1966), Gordon and Breach, New York, 1967, 117–123.
- [Erd81] P. Erdős, *Problems and results in graph theory*, in *The Theory and Applications of Graphs* (Kalamazoo, 1980), Wiley, New York, 1981, 331–341, https://www.renyi.hu/~p_erdos/1981-20.pdf.
- [Erd97] P. Erdős, *Some of my favorite problems and results*, in *The Mathematics of Paul Erdős, I*, *Algorithms Combin.* **13**, Springer, Berlin, 1997, 47–67, https://doi.org/10.1007/978-3-642-60408-9_3.
- [ES82] P. Erdős and M. Simonovits, *Compactness results in extremal graph theory*, *Combinatorica* **2** (1982), 275–288, <https://doi.org/10.1007/BF02579234>.
- [Fur91] Z. Füredi, *On a Turán type problem of Erdős*, *Combinatorica* **11** (1991), 75–79, <https://doi.org/10.1007/BF01375476>.
- [FS13] Z. Füredi and M. Simonovits, *The history of degenerate (bipartite) extremal graph problems*, in *Erdős Centennial*, *Bolyai Soc. Math. Stud.* **25** (2013), 169–264, https://doi.org/10.1007/978-3-642-39286-3_7.
- [GJN22] A. Grzesik, O. Janzer, and Z. L. Nagy, *The Turán number of blow-ups of trees*, *J. Combin. Theory Ser. B* **156** (2022), 299–309, <https://doi.org/10.1016/j.jctb.2022.05.004>.
- [Jan23] O. Janzer, *Disproof of a conjecture of Erdős and Simonovits on the Turán number of graphs with minimum degree 3*, *Int. Math. Res. Not.* **2023** (2023), 8478–8494, <https://doi.org/10.1093/imrn/rnac076>.
- [PT09] S. E. Payne and J. A. Thas, *Finite Generalized Quadrangles*, 2nd ed., EMS Series of Lectures in Mathematics **9**, European Mathematical Society, 2009, <https://doi.org/10.4171/066>.
- [Ver16] J. Verstraëte, *Extremal problems for cycles in graphs*, in *Recent Trends in Combinatorics*, IMA Vol. Math. Appl. **159**, Springer, Cham, 2016, 83–116, https://doi.org/10.1007/978-3-319-24298-9_4.
- [Wig] Y. Wigderson, *The Erdős–Simonovits compactness conjecture needs more assumptions*, <https://ywigderson.math.ethz.ch/math/static/Compactness.pdf>.