

Harmonicity and Invariance on Slices of the Boolean Cube

Yuval Filmus¹ and Elchanan Mossel²

¹Department of Computer Science, Technion — Israel Institute of Technology ,
yuvalfi@cs.technion.ac.il

²Mathematics and IDSS, Massachusetts Institute of Technology , elmos@mit.edu

December 21, 2018

Abstract

In a recent work with Kindler and Wimmer we proved an invariance principle for the slice for low-influence, low-degree harmonic multilinear polynomials (a polynomial in x_1, \dots, x_n is *harmonic* if it is annihilated by $\sum_{i=1}^n \frac{\partial}{\partial x_i}$). Here we provide an alternative proof for *general* low-degree harmonic multilinear polynomials, with no constraints on the influences. We show that any real-valued harmonic multilinear polynomial on the slice whose degree is $o(\sqrt{n})$ has approximately the same distribution under the slice and cube measures.

Our proof is based on ideas and results from the representation theory of S_n , along with a novel decomposition of random increasing paths in the cube in terms of martingales and reverse martingales. While such decompositions have been used in the past for stationary reversible Markov chains, our decomposition is applied in a non-stationary non-reversible setup. We also provide simple proofs for some known and some new properties of harmonic functions which are crucial for the proof.

Finally, we provide independent simple proofs for the known facts that 1) one cannot distinguish between the slice and the cube based on functions of $o(n)$ coordinates and 2) Boolean symmetric functions on the cube cannot be approximated under the uniform measure by functions whose sum of influences is $o(\sqrt{n})$.

Contents

1	Introduction	3
1.1	Low-degree polynomials	3
1.2	Influences, symmetric functions and circuits	4
1.3	Other results comparing the cube to the slice	5
2	Definitions	6
3	Harmonic functions	7
3.1	Canonical representation	7
3.2	Orthogonality of homogeneous parts	10
3.3	Poincaré inequalities	11
3.4	Cometric property	13
4	Invariance principle	14
4.1	The proof	16
4.2	High-degree functions	22
5	Approximately Boolean functions	23
6	Multilinear functions	26
6.1	Blekherman’s theorem	27
6.2	Vandermonde interpolation	28
6.3	Norm invariance	29
6.4	Interpolation invariance	33
7	Functions depending on few coordinates	36
8	Functions with low total influence	38
9	Decomposing the slice	40
9.1	Coarse decomposition	40
9.2	Gelfand–Tsetlin basis	41
10	Why harmonic functions?	42
10.1	Commutative algebra	43
10.2	Representation theory	45

1 Introduction

The basic question motivating our work is the following:

Question 1.1. *Assume n is even. How distinguishable are the uniform measure μ on $\{0, 1\}^n$ and the measure ν given by the uniform measure on $\{0, 1\}^n$ conditioned on $\sum_i x_i = n/2$?*

More generally: how distinguishable are the product measure μ_p on $\{0, 1\}^n$ in which each coordinate takes the value 1 independently with probability p and ν_{pn} given by the uniform measure on $\{0, 1\}^n$ conditioned on $\sum_i x_i = pn$ (assuming pn is an integer)?

Note that the two measures are easily distinguished using the simple sum-of-coordinates test. How does the answer change if we “do not allow” the sum-of-coordinate test? From a computational perspective we might be interested in restricted families of tests, such as *low-depth circuits* or *low-degree polynomials*. Furthermore, it turns out that the canonical representation of functions with respect to ν_{pn} , called harmonic representation, does not include the function $\sum x_i$ (indeed, under ν_{pn} it can be represented as a constant).

We call $\{0, 1\}^n$ the *cube*, the support of the distribution ν_{pn} the *slice*, and the support of ν the *middle slice*. For exposition purposes, the introduction will only discuss the middle slice, though all results (previous and ours) extend for the case of μ_p and ν_{pn} for every fixed p .

1.1 Low-degree polynomials

In a recent joint work with Kindler and Wimmer [18] we provided a partial answer to Question 1.1 by extending the non-linear invariance principle of [29]. As mentioned earlier, any function on the slice has a canonical representation as a *harmonic*¹ multilinear polynomial. The harmonic property means that $\sum_{i=1}^n \frac{\partial f}{\partial x_i} = 0$. Suppose that f is a harmonic low-degree low-influence multilinear polynomial. The invariance principle of [18] establishes that the distribution of f under the measure ν is close to its distribution under the product measure μ on $\{0, 1\}^n$, as well as to its distribution under the product space \mathbb{R}^n equipped with the product Gaussian measure $\mathcal{G} = N(1/2, 1/4)^{\otimes n}$.

The restriction to multilinear harmonic polynomials is quite natural in the slice — as every function on the slice has a unique representation as a harmonic multilinear polynomial (this fact, due to Dunkl [14], is proved in Section 3). It is the analog of the implicit restriction to multilinear polynomials in the original non-linear invariance principle. Further motivations, from commutative algebra and representation theory, are described in Section 10.

Both the invariance principle proven in [29] and the one proven in [18] require that the functions have low influences. Indeed, a function like x_1 has a rather different distribution under μ compared to \mathcal{G} . Similarly, the function $x_1 - x_2$ has a rather different distribution under ν compared to \mathcal{G} . However, note that the distribution of $x_1 - x_2$ under ν is quite similar to its distribution under μ . It is natural to speculate that low-degree harmonic functions have similar distributions under ν and μ . Unfortunately, the proof of the invariance principle in [18] goes through Gaussian space, rendering the low-influence condition necessary even when comparing ν and μ .

Our main result in this paper is a direct proof of the invariance principle on the slice, showing that the distribution of a low-degree harmonic function on the slice is close to its distribution on the corresponding cube. Our results do not require the condition of low influences.

Theorem 1.2. *Let $f: \{-1, 1\}^n \rightarrow \mathbb{R}$ be a harmonic multilinear polynomial of degree $o(\sqrt{n})$ and variance 1. For any 1-Lipschitz function φ (i.e., one satisfying $|\varphi(x) - \varphi(y)| \leq |x - y|$),*

$$|\mathbb{E}_{\nu}[\varphi(f)] - \mathbb{E}_{\mu}[\varphi(f)]| = o(1),$$

and the Lévy distance² between the distribution of f under μ and its distribution under ν is $o(1)$.

¹This somewhat unfortunate terminology is borrowed from Bergeron [4, Section 8.4], in which an S_n -harmonic polynomial is one which is annihilated by $\sum_{i=1}^n \frac{\partial^k}{\partial x_i^k}$ for all k . For multilinear polynomials, both definitions coincide.

²The Lévy distance between two real random variables X, Y is the infimum value of ϵ such that for all $t \in \mathbb{R}$ it holds that $\Pr[X \leq t - \epsilon] - \epsilon \leq \Pr[Y \leq t] \leq \Pr[X \leq t + \epsilon] + \epsilon$.

See Theorem 4.6 as well as Corollary 4.7 for more quantitative bounds and more general statements (which apply in particular to any i.i.d. measure on the cube and the corresponding slice). In Subsection 4.2 we show that the results cannot be extended to polynomials whose degree is much bigger than \sqrt{n} .

While Theorem 1.2 cannot hold for arbitrary multilinear polynomials, we are able to recover a similar theorem for such polynomials by considering the distribution of f on several coupled slices at once; see Theorem 6.17 and its corollaries.

Between representation theory and probability An interesting aspect of our work is the interplay between the L_2 theory of the slice and probabilistic arguments based on coupling.³ The L_2 theory of the slice, which was developed by Dunkl [14, 15], is intimately related to representations of S_n . In Section 3 we provide an elementary and self-contained approach to this theory. The L_2 theory is used to control the norm of low-degree harmonic functions with respect to several different measures, in particular by using a two-sided Poincaré inequality. This, in turn, allows us to introduce the main probabilistic argument.

Novel proof ingredients The basic idea of the proof is to use the coupling method by showing that the distribution of f on different slices of the cube is almost identical, as long as the slices are of distance at most roughly \sqrt{n} (here the distance between the k th slice and the ℓ th slice is $|k - \ell|$). In fact, for our proof to work we crucially need to allow distances which are somewhat larger than \sqrt{n} .

To construct the coupling, we use a uniform random increasing path to couple level ℓ to level k above or below it. The main novel technique is representing the difference between the two levels as a difference of two martingales. Such representations have been used before in the analysis of *stationary reversible* Markov chains in Banach spaces [30], and even earlier in the analysis of stochastic integrals [28]. However, all previous decompositions were for stationary reversible chains, while ours is neither. Our novel representation of the differences might be of interest in other applications. We outline this argument in the introduction to Section 4, and carry it out in the rest of the section.

Applications Except for the natural interpretation of Theorem 1.2 in terms of distinguishing between distributions, it can be used to prove results in extremal combinatorics in the same way the main result of [18] is used. For example, in Proposition 5.5 we give a proof of the Kindler–Safta theorem on the slice, first proved in [18].

1.2 Influences, symmetric functions and circuits

We prove a few other results that give partial answers to Question 1.1:

- Using direct computation of the total variation distance we prove the following theorem:

Theorem 7.2. Let f be a function on $\{0, 1\}^n$ depending on $o(n)$ coordinates and satisfying $\|f\|_\infty \leq 1$. Then

$$|\mathbb{E}_\nu[f] - \mathbb{E}_\mu[f]| = o(1).$$

- We prove that symmetric functions cannot be approximated by functions whose total influence is $o(\sqrt{n})$ (see Proposition 8.5 for a more general formulation):

Proposition 1.3. *There exists a constant $\delta > 0$ such that if f is a symmetric Boolean function such that $\frac{1}{3} \leq \mathbb{E}_{\mu_p}[f] \leq \frac{2}{3}$ then $\Pr_{\mu_p}[f \neq g] > \delta$ for every Boolean function g satisfying $\text{Inf}[g] = o(\sqrt{n})$.*

Since it is well-known [7], based on arguments from [27, 19], that a depth d size m circuit has total influence at most $O((\log m)^{d-1})$, our result immediately implies circuit lower bounds for such functions.

³The informal term “ L_2 theory” refers to studying functions via L_2 norms and inner products. For example, Fourier analysis is an L_2 theory, since its central concept is the Fourier basis, which is an orthonormal basis with respect to a given inner product. In contrast, total variation distance, which is inherently related to coupling, is an L_1 concept.

However, much better bounds are known, see e.g. [39] for general symmetric functions and [32] for the case of the majority function. Nevertheless, Proposition 1.3 is more general as it holds for functions f that are not necessarily the majority function and for functions g that are not necessarily in AC^0 . Moreover, the proof of Proposition 1.3 is based on a new and simple probabilistic argument.

1.3 Other results comparing the cube to the slice

Question 1.1 is not a new question. So we conclude the introduction with a few classical results related to this question:

- The limiting behaviors of the partial sum $W(s) = \frac{1}{\sqrt{s}} \sum_{i=1}^s (x_i - \frac{1}{2})$ as $s \rightarrow \infty$ under the cube and the slice measures are well-studied. It is well-known that under the cube measure $W(s)$ converges to Brownian motion, while under the slice measure it converges to a Brownian bridge.
- It is well-known that the partial sums $W(s)$ are at least as concentrated in the slice as they are in the cube [20].
- It is well-known that Lipschitz functions of the random variables x_1, \dots, x_n are concentrated both in the cube and in the slice. The results for the slice follow from the hypercontractive estimates by Lee and Yau [26]. These are also needed in our proofs.

Paper organization Following Section 2, which contains several useful definitions, the paper is composed of four major parts:

1. **Harmonic analysis on the slice (Section 3).** In this part we provide a self-contained introduction to harmonic analysis on the slice. Most of the results proved in this section are not new, but our proofs are novel and elementary.
2. **Invariance principles (Sections 4–6).** In this part we prove several invariance principles. Section 4 contains the major result of this paper, an invariance principle for Lipschitz functions. This theorem applies to polynomials of degree $o(\sqrt{n})$, a condition discussed in Section 4.2. Section 5 extends the invariance principle to the non-Lipschitz function $\varphi(x) = (|x| - 1)^2$, appearing in many applications. As an application of this invariance principle, we give an alternative proof of a non-optimal Kindler–Safra theorem for the slice, first proved in [18] (an optimal version of the theorem has since been proved by Keller and Klein [22]). Both of these results are for harmonic polynomials. The case of non-harmonic polynomials is tackled in Section 6, in which an analog of the invariance principle is proved for arbitrary multilinear polynomials.
3. **Minor invariance principles (Sections 7–8).** In this part we prove two simpler invariance principles for restricted classes of functions: bounded functions depending on $o(n)$ coordinates (Section 7) and functions whose total influence is $o(\sqrt{n})$ (Section 8).
4. **Connections (Sections 9–10).** In this final part, we provide some connections to other mathematical fields. Section 9 provides a representation-theoretic angle on the slice. This point of view serves to demystify some of the important properties of the slice proved in Section 3. Section 10 justifies our use of harmonic polynomials to “lift” functions from slice to cube. Two justifications are given: a Sperner-theoretic justification (Section 10.1) and a representation-theoretic justification (Section 10.2). Section 10.1 also contains an alternative proof of the important property that every function on the slice has a canonical representation as a harmonic multilinear polynomial. In fact, we prove a stronger result, Blekherman’s theorem [5], which is crucial in Section 6. A corollary of Blekherman’s theorem is also used in the proof of the Kindler–Safra theorem in Section 5.

Acknowledgements

Both authors would like to thank the referees for their extensive and helpful comments.

Y.F. would like to mention that this material is based upon work supported by the National Science Foundation under agreement No. DMS-1128155. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors, and do not necessarily reflect the views of the National Science Foundation. Part of the work was done while at the Institute for Advanced Study, Princeton, NJ. The research was also funded by ISF grant 1337/16. The author is a Taub Fellow, and supported by the Taub Foundations.

E.M. would like to acknowledge the support of the following grants: NSF grants DMS 1106999 and CCF 1320105, DOD ONR grant N00014-14-1-0823, and grant 328025 from the Simons Foundation.

2 Definitions

Notation We employ the falling power notation $n^k = n(n-1)\cdots(n-k+1)$. The notation $\mathbf{1}_E$ equals 1 if the condition E holds, and 0 otherwise. The sign function is denoted sgn . The L_2 triangle inequality is $(a+b)^2 \leq 2(a^2+b^2)$ or its generalization $(\sum_{i=1}^n a_i)^2 \leq n \sum_{i=1}^n a_i^2$.

A monomial is *squarefree* if it is not divisible by a square of a variable. (Thus there are 2^n squarefree monomials on n variables.) A polynomial is *multilinear* if all monomials are squarefree. A polynomial is *homogeneous* if all monomials have the same total degree. The d th homogeneous part of a polynomial $f = \sum c_m m$, denote $f^{=d}$, is the sum of $c_m m$ over all monomials m of total degree d . A polynomial f over x_1, \dots, x_n is *harmonic* if $\sum_{i=1}^n \frac{\partial f}{\partial x_i} = 0$.

A univariate function f is *C-Lipschitz* if $|f(x) - f(y)| \leq C|x - y|$. A function is *Lipschitz* if it is 1-Lipschitz.

The expectation and variance of a random variable are denoted \mathbb{E}, \mathbb{V} , and $\|\cdot\|$ denotes its L_2 norm $\|X\| = \sqrt{\mathbb{E}[X^2]}$. To signify that expectation is taken with respect to a distribution α , we write $\mathbb{E}_\alpha[X]$, $\mathbb{V}_\alpha[x]$, and $\|\cdot\|_\alpha$. A normal distribution with mean μ and variance σ^2 is denoted $N(\mu, \sigma^2)$. A binomial distribution with n trials and success probability p is denoted $B(n, p)$.

The symmetric group on $[n] = \{1, \dots, n\}$ is denoted S_n . A distribution on \mathbb{R}^n is *exchangeable* if it is invariant under the action of S_n (that is, under permutation of the coordinates); a discrete distribution is exchangeable if the probability of (x_1, \dots, x_n) depends only on $x_1 + \dots + x_n$. For a function f on \mathbb{R}^n and a permutation π , we define $f^\pi(x_1, \dots, x_n) = f(x_{\pi(1)}, \dots, x_{\pi(n)})$. We compose permutations left-to-right, so that $(f^\alpha)^\beta = f^{\alpha\beta}$.

Asymptotic notation The notation $O(f)$ means a function g such that $g \leq Cf$ for some positive constant $C > 0$. In particular, $g = O(f)$ means that $g \leq Cf$ for some positive constant $C > 0$. The notation $\Omega(f)$ means a function g such that $g \geq cf$ for some positive constant $c > 0$. The notation $\Theta(f)$ means a function g such that $cf \leq g \leq Cf$ for some positive constants $C \geq c > 0$.

We write $g(n) = o(f(n))$ if $\lim_{n \rightarrow \infty} g(n)/f(n) = 0$. The quantity n should be clear from context (in some cases, for example, it is $p(1-p)n$). We write $g(n) = \omega(f(n))$ if $\lim_{n \rightarrow \infty} g(n)/f(n) = \infty$.

The notation $O_a(f)$ means a function g such that $g \leq C(a)f$ for some everywhere positive function $C(a)$. We similarly define $O_{a,b}(f)$, $\Omega_a(f)$, and so on. For example, $n/p(1-p) = O_p(n)$.

The slice The n -dimensional Boolean cube is the set $\{0, 1\}^n$. For an integer $0 \leq k \leq n$, the k th slice of the n -dimensional Boolean cube is the set

$$\binom{[n]}{k} = \left\{ (x_1, \dots, x_n) \in \{0, 1\}^n : \sum_{i=1}^n x_i = k \right\}.$$

Probability measures Our work involves two main probability measures, where n is always understood:

- The uniform measure on the slice $\binom{[n]}{k}$ is ν_k .

- The product measure μ_p on the Boolean cube is given by $\mu_p(x) = p^{\sum_i x_i} (1-p)^{\sum_i (1-x_i)}$.

Note that $\nu_k, \mu_{k/n}$ have the same marginal distributions.

3 Harmonic functions

A basic and easy result states that every function on $\{-1, 1\}^n$ has a unique representation as a multilinear polynomial, known as the *Fourier expansion*. It is easy to see that a multilinear polynomial has the same mean and variance with respect to the uniform measure on $\{-1, 1\}^n$ and with respect to the standard n -dimensional Gaussian measure. In this section we describe the corresponding canonical representation on the slice, due to Dunkl [14, 15] and elaborated by Srinivasan [36], Filmus [17] and Ambainis et al. [1]. Most of the results in this section are already known, though the proofs presented in this section are novel. A possible exception is the two-sided Poincaré inequality for derivatives, Lemma 3.16.

The canonical representation of functions on the slice is described in Subsection 3.1. We decompose this representation into orthogonal parts in Subsection 3.2, where we also deduce that the mean and variance of a low-degree function is similar on the slice and on the Boolean cube. The analog of the Poincaré inequality is proved in Subsection 3.3 alongside results of a similar flavor. Finally, we prove that degree is subadditive with respect to multiplication, and monotone with respect to substitution, in Subsection 3.4.

3.1 Canonical representation

Every function on the slice $\binom{[n]}{k}$ can be represented as a multilinear polynomial, but this representation is not unique. However, as found by Dunkl [14, 15], we can make it unique by demanding that it be *harmonic* in the sense of the following definition.

Definition 3.1. A polynomial P over x_1, \dots, x_n is *harmonic* if

$$\sum_{i=1}^n \frac{\partial P}{\partial x_i} = 0.$$

In other words, P is harmonic if $\Delta P = 0$, where Δ is the differential operator $\sum_{i=1}^n \frac{\partial}{\partial x_i}$.

Definition 3.2. A *basic function* is a (possibly empty) product of factors $x_i - x_j$ on disjoint indices. A function is *elementary* if it is a linear combination of basic functions.

Most, but not all, of the harmonic polynomials we consider will be multilinear. In particular, notice that all elementary functions are multilinear. Here are some basic properties of harmonic polynomials.

Lemma 3.3. *The set of harmonic polynomials is an algebra of polynomials, and is closed under partial derivatives, under permutations of the coordinates, and under taking homogeneous parts. In particular, all elementary functions are harmonic.*

Proof. Suppose f, g are harmonic. Then $\Delta(\alpha f + \beta g) = \alpha \Delta f + \beta \Delta g = 0$; $\Delta(fg) = f \Delta g + g \Delta f = 0$; $\Delta \frac{\partial f}{\partial x_i} = \frac{\partial \Delta f}{\partial x_i} = 0$; and $\Delta(f^\pi) = (\Delta f)^\pi = 0$. Finally, since $\Delta(\sum_{d=0}^n f^{=d}) = \sum_{d=0}^n \Delta f^{=d}$ and $\Delta f^{=d}$ is homogeneous of degree $d-1$, we see that $\Delta f^{=d} = 0^{=d-1} = 0$. \square

Lemma 3.4. *A polynomial f is harmonic if and only if for all x_1, \dots, x_n, c we have*

$$f(x_1 + c, \dots, x_n + c) = f(x_1, \dots, x_n).$$

In particular, if $f^\pi = f$ for all $\pi \in S_n$ then $\deg f = 0$.

Proof. Given x_1, \dots, x_n , define a function

$$\phi(x_1, \dots, x_n, c) = f(x_1 + c, \dots, x_n + c).$$

The chain rule implies that $\frac{\partial \phi}{\partial c} = \Delta f$. Hence $\Delta f = 0$ iff ϕ is independent of c .

If f moreover satisfies $f^\pi = f$ for all $\pi \in S_n$, then f is a symmetric polynomial, and so $f = \phi(x_1 + \dots + x_n)$ for some univariate polynomial ϕ . The polynomial ϕ satisfies $\phi(x) = \phi(x + nc)$ for all c , and so both ϕ and f must be constant. \square

Our proofs of Theorem 3.6 and Lemma 3.7 proceed by analyzing “derivatives” of the form $f - f^{(i j)}$, and using the following simple lemma to conclude the same property for f itself.

Lemma 3.5. *Let X be a ring of characteristic zero with an S_n action, and let $V \subseteq X$ be a vector space satisfying the following properties:*

(a) *If $f \in X$ then $f - f^{(i j)} \in V$ for all $i \neq j$.*

(b) *If $f \in X$ then $S(f) := \mathbb{E}_\pi[f^\pi] \in V$.*

Then $V = X$.

Proof. Every permutation $\pi \in S_n$ can be written as $\pi = (i_1 j_1) \cdots (i_\ell j_\ell)$. Since V is a vector space, the first property shows that

$$f - f^\pi = \sum_{t=1}^{\ell} f^{(i_1 j_1) \cdots (i_{t-1} j_{t-1})} - f^{(i_1 j_1) \cdots (i_t j_t)} \in V.$$

Taking average over all $\pi \in S_n$, we get that $f - S(f) \in V$. The second property now implies that $f \in V$. \square

Our first theorem states that every function on the slice has a unique representation as a harmonic multilinear polynomial of degree at most $\min(k, n - k)$.

Theorem 3.6. *Let $0 \leq k \leq n$. Every function on the slice $\binom{[n]}{k}$ has a unique representation as a harmonic multilinear polynomial of degree at most $\min(k, n - k)$.*

Proof. We use the notation $f \equiv g$ (read f agrees with g) to denote that f agrees with g on the slice $\binom{[n]}{k}$.

We start by proving, by induction on $\min(k, n - k)$, that every function on a slice $\binom{[n]}{k}$ has some representation as a harmonic multilinear polynomial. The base cases, $k = 0$ and $k = n$, are trivial, since in these cases all functions are constant.

Consider now a function f on the slice $\binom{[n]}{k}$, where $0 < k < n$. Define a function g on the slice $\binom{[n-2]}{k-1}$ by $g(x_1, \dots, x_{n-2}) = f(x_1, \dots, x_{n-2}, 0, 1) - f(x_1, \dots, x_{n-2}, 1, 0)$. By the induction hypothesis, g can be represented as a harmonic multilinear polynomial $G \equiv g$ of degree at most $\min(k - 1, (n - 2) - (k - 1)) = \min(k, n - k) - 1$. Let $H = (x_n - x_{n-1})G$. Note that H is a multilinear polynomial of degree at most $\min(k, n - k)$, and by Lemma 3.3 it is harmonic. We claim that $f - f^{(n-1 n)} \equiv H$ (recall that $f^{(n-1 n)}$ is obtained from f by permuting coordinates $n - 1$ and n). Indeed, when $x_{n-1} = x_n$, both sides vanish, and when $(x_{n-1}, x_n) = (0, 1)$ or $(x_{n-1}, x_n) = (1, 0)$, this is true by definition (checking both cases separately). Thus $f - f^{(n-1 n)}$ can be represented as a harmonic multilinear polynomial of degree at most $\min(k, n - k)$. For short, we say that $f - f^{(n-1 n)}$ is *representable*.

The same argument implies that $f - f^{(i j)}$ is representable for all i, j , satisfying the first property in Lemma 3.5, with X consisting of all functions on the slice $\binom{[n]}{k}$, and V consisting of all representable functions. Since $S(f) = \mathbb{E}[f]$ is a constant and so representable, the lemma shows that all functions are representable.

It remains to prove that the representation is unique. To that end, it is enough to show that if P is a harmonic multilinear polynomial of degree at most $\min(k, n - k)$ such that $P \equiv 0$ then $P = 0$. We prove this by induction on $\min(k, n - k)$. The base cases, $k = 0$ and $k = n$, are trivial, since if P is a constant polynomial agreeing with the zero function, then $P = 0$.

Suppose now that $P \equiv 0$ on the slice $\binom{[n]}{k}$ for some harmonic multilinear polynomial P of degree at most $\min(k, n - k)$, where $0 < k < n$. Write $P = A + x_{n-1}B + x_nC + x_{n-1}x_nD$, where A, B, C, D are polynomials over x_1, \dots, x_{n-2} , and notice that $P - P^{(n-1\ n)} = (x_{n-1} - x_n)(B - C)$. Since $\Delta[(x_{n-1} - x_n)(B - C)] = (x_{n-1} - x_n)\Delta(B - C)$, we see that $B - C$ is harmonic. Considering the substitution $(x_{n-1}, x_n) = (0, 1)$, we see that $B - C$ is a harmonic multilinear polynomial of degree at most $\min(k, n - k) - 1$ which agrees with 0 on the slice $\binom{[n-2]}{k-1}$. The induction hypothesis implies that $B - C = 0$, and so $P = P^{(n-1\ n)}$.

The same argument implies that $P = P^{(i\ j)}$ for all i, j , and so $P = P^\pi$ for all π . Lemma 3.4 implies that P is constant. Since P agrees with 0, we must have $P = 0$, completing the proof. \square

Similarly, we can prove that every harmonic multilinear polynomial is elementary.

Lemma 3.7. *A multilinear polynomial is harmonic iff it is elementary. In particular, a harmonic multilinear polynomial over x_1, \dots, x_n has degree at most $n/2$.*

Proof. Lemma 3.3 implies that every elementary polynomial is harmonic. We now prove that every multilinear harmonic polynomial is elementary by induction on n . If $n = 0$ then any polynomial is constant and so elementary. If $n = 1$ then any harmonic polynomial is constant and so elementary. Consider now any $n \geq 2$ and any harmonic multilinear polynomial f over x_1, \dots, x_n . For any $i \neq j$ we can write $f = A + x_iB + x_jC + x_ix_jD$, where A, B, C, D don't involve x_i, x_j , so that $f - f^{(i\ j)} = (x_i - x_j)(B - C)$. Notice that $f - f^{(i\ j)}$ is harmonic, and since $\Delta[(x_i - x_j)(B - C)] = (x_i - x_j)\Delta(B - C)$, we see that $B - C$ is also harmonic. By induction, we get that $B - C$ is elementary, and so $f - f^{(i\ j)} = (x_i - x_j)(B - C)$ is elementary.

This shows that the first property in Lemma 3.5 is satisfied, where X consists of all multilinear harmonic polynomials, and V consists of all elementary polynomials. Lemma 3.4 shows that $S(f)$ is constant and so elementary, and so the second property is satisfied as well. Lemma 3.5 therefore implies that every multilinear harmonic polynomial is elementary. \square

The algebraically inclined reader may enjoy the following reformulation of Theorem 3.6.

Corollary 3.8. *Let $0 \leq k \leq n$, and fix a field \mathbb{F} of characteristic zero. Consider the polynomial ideal*

$$I = \left\langle x_1^2 - x_1, \dots, x_n^2 - x_n, \sum_{i=1}^n x_i - k \right\rangle = I\left(\binom{[n]}{k}\right),$$

where we think of the slice $\binom{[n]}{k}$ as an affine variety. Then $\mathbb{F}[x_1, \dots, x_n]/I$ is isomorphic to the ring of harmonic multilinear polynomials of degree at most $\min(k, n - k)$ over x_1, \dots, x_n .

We discuss this reformulation in Section 10.1, giving an alternative proof of the theorem.

We conclude this section by computing the dimension of the space of harmonic multilinear polynomials of given degree.

Corollary 3.9. *Let $d \leq n/2$. Then*

- (a) *The linear space $H_{\leq d}$ of harmonic multilinear polynomials of degree at most d has dimension $\binom{n}{d}$.*
- (b) *The linear space H_d of harmonic multilinear polynomials which are homogeneous of degree d has dimension $\binom{n}{d} - \binom{n}{d-1}$, where $\binom{n}{-1} = 0$.*

Proof. The first item follows directly from Theorem 3.6, taking $k = d$. Since $\dim H_{\leq d} = \sum_{e=0}^d \dim H_e$, the second item follows from the first. \square

For another proof, see the proof of Theorem 10.1.

3.2 Orthogonality of homogeneous parts

As stated in the introduction to this section, multilinear polynomials enjoy the useful property of having the same mean and variance with respect to all product measures with fixed marginal mean and variance. The corresponding property for harmonic multilinear polynomials is stated in the following theorem, which also follows from the work of the first author [17]. A representation-theoretic proof of the theorem is outlined in Section 9.1.

Theorem 3.10. *Let f, g be homogeneous harmonic multilinear polynomials of degree d_f, d_g , respectively, and let α be an exchangeable measure. If $d_f \neq d_g$ then $\mathbb{E}_\alpha[fg] = 0$. If $d_f = d_g = d$ then there exists a constant $C_{f,g}$ independent of α such that*

$$\mathbb{E}_\alpha[fg] = C_{f,g} \mathbb{E}_\alpha[(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2].$$

Proof. Let $h = fg$, and note that Lemma 3.3 implies that h is harmonic. Let $H = \mathbb{E}_{\pi \in S_n}[h^\pi]$, and note that H is also harmonic. Since α is exchangeable, $\mathbb{E}_\alpha[fg] = \mathbb{E}_\alpha[H]$.

We first note that H is a linear combination of the functions

$$b_t = \mathbb{E}_{\pi \in S_n} x_{\pi(1)}^2 \cdots x_{\pi(t)}^2 x_{\pi(t+1)} \cdots x_{\pi(d_f+d_g-t)},$$

say $H = \sum_t \beta_t(f, g) b_t$.

Suppose first that $d_f \neq d_g$. It is easy to check that $\mathbb{E}_{N(0,1)}[b_t] = 0$ for all t , and so $\mathbb{E}_{N(0,1)}[H] = 0$. Since H is harmonic, Lemma 3.4 implies that $\mathbb{E}_{N(\mu,1)}[H]$ doesn't depend on μ , and so

$$0 = \mathbb{E}_{N(0,1)}[H] = \mathbb{E}_{N(\mu,1)}[H] = \sum_{t=0}^{\min(d_f, d_g)} \beta_t(f, g) (1 + \mu^2)^t \mu^{d_f+d_g-2t}.$$

The polynomial $P_t(\mu) = (1 + \mu^2)^t \mu^{d_f+d_g-2t}$ has minimal degree term $\mu^{d_f+d_g-2t}$, and so the polynomials $P_0, \dots, P_{\min(d_f, d_g)}$ are linearly independent. This shows that $\beta_t(f, g) = 0$ for all t , and so $H = 0$. In particular, $\mathbb{E}_\alpha[fg] = \mathbb{E}_\alpha[H] = 0$.

When $d_f = d_g = d$, it is still true that $\mathbb{E}_{N(0,1)}[b_t] = 0$ for all $t < d$, but now $\mathbb{E}_{N(0,1)}[b_d] = 1$. Therefore the same argument as before shows that

$$\beta_d(f, g) = \sum_{t=0}^d \beta_t(f, g) P_t.$$

The linear independence of P_0, \dots, P_d implies that $\beta_t(f, g)/\beta_d(f, g)$ depends only on d , and so

$$H = \beta_d(f, g) \sum_{t=0}^d \kappa_t b_t$$

for constants $\kappa_0, \dots, \kappa_t$ depending only on d . In particular,

$$\mathbb{E}_\alpha[fg] = \mathbb{E}_\alpha[H] = \beta_d(f, g) \mathbb{E}_\alpha \left[\sum_{t=0}^d \kappa_t b_t \right].$$

Applying the same argument to $f' = g' = (x_1 - x_2) \cdots (x_{2d-1} - x_{2d})$, we obtain

$$\mathbb{E}_\alpha[(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2] = \beta'_d \mathbb{E}_\alpha \left[\sum_{t=0}^d \kappa_t b_t \right],$$

where β'_d depends only on d . The theorem follows with $C_{f,g} = \beta_d(f, g)/\beta'_d$. \square

Corollary 3.11. *Let f be a harmonic multilinear polynomial of degree at most d with constant coefficient $f^{=0}$. Suppose that α, β are exchangeable measures and $C > 0$ is a constant that for $t \leq d$ satisfies*

$$\mathbb{E}_\alpha[(x_1 - x_2)^2 \cdots (x_{2t-1} - x_{2t})^2] \leq C \mathbb{E}_\beta[(x_1 - x_2)^2 \cdots (x_{2t-1} - x_{2t})^2].$$

Then $\mathbb{E}_\alpha[f] = f^{=0}$, $\|f\|_\alpha^2 \leq C\|f\|_\beta^2$, and $\mathbb{V}_\alpha[f] \leq C\mathbb{V}_\beta[f]$.

Proof. Write $f = \sum_{t=0}^d f^{=t}$, where $f^{=t}$ is homogeneous of degree t . Theorem 3.10 implies that $f^{=t_1}, f^{=t_2}$ are orthogonal with respect to all exchangeable measures. This implies that $\mathbb{E}_\alpha[f] = f^{=0}$ and

$$\|f\|_\alpha^2 = \sum_{t=0}^d \|f^{=t}\|_\alpha^2.$$

The theorem also implies that for some K_0, \dots, K_d we moreover have

$$\|f\|_\alpha^2 = \sum_{t=0}^d K_t \mathbb{E}_\alpha[(x_1 - x_2)^2 \cdots (x_{2t-1} - x_{2t})^2] \leq C \sum_{t=0}^d K_t \mathbb{E}_\beta[(x_1 - x_2)^2 \cdots (x_{2t-1} - x_{2t})^2] = C\|f\|_\beta^2.$$

Finally, since $f - f^{=0}$ is also harmonic, we deduce that $\mathbb{V}[f]_\alpha \leq C\mathbb{V}[f]_\beta$. \square

The following lemma computes $\mathbb{E}[(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2]$ for the measures ν_k, μ_p .

Lemma 3.12. *Let $p = k/n$. We have*

$$\begin{aligned} \mathbb{E}_{\nu_k}[(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2] &= 2^d \frac{k^d(n-k)^d}{n^{2d}} = (2p(1-p))^d \left(1 \pm O\left(\frac{d^2}{p(1-p)n}\right)\right), \\ \mathbb{E}_{\mu_p}[(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2] &= (2p(1-p))^d. \end{aligned}$$

This straightforward computation appears in [17, Theorem 4.1] and [18, Lemma 2.9]. Qualitatively, the lemma states that the norm of a low-degree basic function is similar in both ν_k and μ_p . This is not surprising: the coordinates in the slice are almost independent, and a low-degree basic function depends only on a small number of them.

3.3 Poincaré inequalities

We proceed by proving the so-called two-sided Poincaré inequality, starting with the following fact.

Lemma 3.13. *Let f be a harmonic multilinear polynomial. Then*

$$\sum_{i < j} f^{(i j)} = \sum_{d=0}^{n/2} \left[\binom{n}{2} - d(n-d+1) \right] f^{=d},$$

where $f^{=d}$ is the d th homogeneous part of f .

Proof. In view of Lemma 3.7, f is elementary. Therefore it is enough to consider the case that f is a basic function of some degree d , say $f = \prod_{t=1}^d (x_{a_t} - x_{b_t})$. We split the transpositions into four kinds. The first kind is transpositions which do not involve any a_t, b_t . There are $\binom{n-2d}{2}$ such transpositions $(i j)$, and they all satisfy $f^{(i j)} = f$. The second kind is transpositions of the form $(a_t b_t)$. There are d of these, and they satisfy $f^{(a_t b_t)} = -f$. The third kind is transpositions of the form $(a_t j)$ or $(b_t j)$, where $j \neq a_s, b_s$. There are $d(n-2d)$ pairs of these. Since

$$(x_{a_t} - x_{b_t})^{(a_t j)} + (x_{a_t} - x_{b_t})^{(b_t j)} = x_{a_t} - x_{b_t},$$

each such pair contributes one multiple of f . The fourth kind is transpositions involving two pairs $(a_t b_t), (a_s b_s)$, which we group in the obvious way into $\binom{d}{2}$ quadruples. Direct computation shows that

$$\sum_{\pi \in \{(a_t a_s), (a_t b_s), (b_t a_s), (b_t b_s)\}} [(x_{a_t} - x_{b_t})(x_{a_s} - x_{b_s})]^\pi = 2(x_{a_t} - x_{b_t})(x_{a_s} - x_{b_s}),$$

and so the contribution of each such quadruple is two multiples of f . In total, we obtain

$$\sum_{i < j} f^{(i j)} = \left[\binom{n-2d}{2} - d + d(n-2d) + 2 \binom{d}{2} \right] f = \left[\binom{n}{2} - d(n-d+1) \right] f. \quad \square$$

Lemma 3.14. *Let f be a harmonic multilinear polynomial of degree at most d . Then with respect to any exchangeable measure,*

$$n \mathbb{V}[f] \leq \frac{1}{2} \sum_{i < j} \|f - f^{(i j)}\|^2 \leq d(n-d+1) \mathbb{V}[f].$$

Proof. Write $f = \sum_{t=0}^d f^{=t}$. Theorem 3.10 implies that the homogeneous parts are orthogonal. We have

$$\frac{1}{2} \sum_{i < j} \|f - f^{(i j)}\|^2 = \binom{n}{2} \|f\|^2 - \left\langle f, \sum_{i < j} f^{(i j)} \right\rangle.$$

Lemma 3.13 implies that

$$\frac{1}{2} \sum_{i < j} \|f - f^{(i j)}\|^2 = \sum_{t=0}^d t(n-t+1) \|f^{=t}\|^2.$$

The lemma now follows from the observation that for $1 \leq t \leq d$ we have $n \leq t(n-t+1) \leq d(n-d+1)$, since $t(n-t+1)$ is increasing for $t \leq (n+1)/2$. \square

Finally, we prove another two-sided Poincaré inequality, this time for derivatives. We start with the following surprising corollary of Theorem 3.10.

Lemma 3.15. *Let f, g be homogeneous harmonic multilinear polynomials of degree d . Then for any exchangeable measure α ,*

$$\frac{\sum_{i=1}^n \mathbb{E}_\alpha \left[\frac{\partial f}{\partial x_i} \frac{\partial g}{\partial x_i} \right]}{\mathbb{E}_\alpha [fg]} = 2d \frac{\mathbb{E}_\alpha [(x_1 - x_2)^2 \cdots (x_{2d-3} - x_{2d-2})^2]}{\mathbb{E}_\alpha [(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2]}.$$

Proof. Since the functions $\frac{\partial f}{\partial x_i}, \frac{\partial g}{\partial x_i}$ are harmonic, Theorem 3.10 implies that there are constants $C_{f,g}, D_{f,g}$ such that

$$\frac{\sum_{i=1}^n \mathbb{E}_\alpha \left[\frac{\partial f}{\partial x_i} \frac{\partial g}{\partial x_i} \right]}{\mathbb{E}_\alpha [fg]} = \frac{D_{f,g}}{C_{f,g}} \frac{\mathbb{E}_\alpha [(x_1 - x_2)^2 \cdots (x_{2d-3} - x_{2d-2})^2]}{\mathbb{E}_\alpha [(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2]}.$$

Here $D_{f,g} = \sum_{i=1}^n C_{\frac{\partial f}{\partial x_i}, \frac{\partial g}{\partial x_i}}$. We can evaluate the ratio $D_{f,g}/C_{f,g}$ by considering the distribution $\alpha = N(0, I_n)$:

$$\frac{\sum_{i=1}^n \mathbb{E}_{N(0, I_n)} \left[\frac{\partial f}{\partial x_i} \frac{\partial g}{\partial x_i} \right]}{\mathbb{E}_{N(0, I_n)} [fg]} = \frac{D_{f,g}}{C_{f,g}} \frac{2^{d-1}}{2^d} = \frac{D_{f,g}}{2C_{f,g}}.$$

On the other hand, with respect to $N(0, I_n)$ we have

$$\sum_{i=1}^n \mathbb{E} \left[\frac{\partial f}{\partial x_i} \frac{\partial g}{\partial x_i} \right] = \sum_{i=1}^n \sum_{S \ni i} \hat{f}(S) \hat{g}(S) = \sum_S |S| \hat{f}(S) \hat{g}(S) = d \mathbb{E}[fg].$$

We conclude that $D_{f,g}/(2C_{f,g}) = d$, and so $D_{f,g}/C_{f,g} = 2d$. \square

We deduce the following two-sided Poincaré inequality.

Lemma 3.16. *Let f be a harmonic multilinear polynomial of degree d , and let α be an exchangeable measure. Suppose that for $1 \leq t \leq d$ we have*

$$m \leq 2t \frac{\mathbb{E}_\alpha[(x_1 - x_2)^2 \cdots (x_{2t-3} - x_{2t-2})^2]}{\mathbb{E}_\alpha[(x_1 - x_2)^2 \cdots (x_{2t-1} - x_{2t})^2]} \leq M.$$

Then also

$$m \mathbb{V}[f] \leq \sum_{i=1}^n \left\| \frac{\partial f}{\partial x_i} \right\|^2 \leq M \mathbb{V}[f].$$

Proof. Write $f = \sum_{t=0}^d f^{=t}$, and notice that $\frac{\partial f}{\partial x_i} = \sum_{t=1}^d \frac{\partial f^{=t}}{\partial x_i}$, the latter sum being a decomposition into homogeneous parts. Theorem 3.10 implies that the homogeneous parts are orthogonal with respect to α , and so

$$\sum_{i=1}^n \left\| \frac{\partial f}{\partial x_i} \right\|^2 = \sum_{t=1}^d \sum_{i=1}^n \left\| \frac{\partial f^{=t}}{\partial x_i} \right\|^2.$$

Using Lemma 3.15 we can upper bound

$$\sum_{t=1}^d \sum_{i=1}^n \left\| \frac{\partial f^{=t}}{\partial x_i} \right\|^2 \leq M \sum_{t=1}^d \|f^{=t}\|^2 = M \mathbb{V}[f].$$

The lower bound is obtained in the same way. □

The following lemma computes m, M for the measures ν_k, μ_p .

Lemma 3.17. *Let $p = k/n$. We have*

$$\begin{aligned} 2d \frac{\mathbb{E}_{\nu_k}[(x_1 - x_2)^2 \cdots (x_{2d-3} - x_{2d-2})^2]}{\mathbb{E}_{\nu_k}[(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2]} &= d \frac{(n - 2d + 2)(n - 2d + 1)}{(k - d + 1)(n - k - d + 1)} \\ &= \frac{d}{p(1-p)} \left(1 \pm O\left(\frac{d}{p(1-p)n}\right) \right), \\ 2d \frac{\mathbb{E}_{\mu_p}[(x_1 - x_2)^2 \cdots (x_{2d-3} - x_{2d-2})^2]}{\mathbb{E}_{\mu_p}[(x_1 - x_2)^2 \cdots (x_{2d-1} - x_{2d})^2]} &= \frac{d}{p(1-p)}. \end{aligned}$$

The proof is a straightforward corollary of Lemma 3.12.

3.4 Cometric property

Theorem 3.6 states that every function on the slice $\binom{[n]}{k}$ can be represented as a harmonic multilinear polynomial. The following result shows that if the original function can be represented as a polynomial of degree d , then its harmonic representation has degree at most d .

Lemma 3.18. *Let $0 \leq k \leq n$. Let f be a polynomial of degree d on the variables x_1, \dots, x_n . The unique harmonic multilinear polynomial agreeing with f on the slice $\binom{[n]}{k}$ has degree at most d .*

Proof. We can assume without loss of generality that f is multilinear, since on the slice $x_i^2 = x_i$. Hence it suffices to prove the theorem for the monomial $f = x_1 \cdots x_d$. Theorem 3.6 states that there is a unique harmonic polynomial g of degree at most $\min(k, n - k)$ agreeing with f on the slice $\binom{[n]}{k}$. If $d \geq \min(k, n - k)$ then the lemma is trivial, so assume that $d < \min(k, n - k)$.

Decompose g into its homogeneous parts: $g = g^{=0} + \cdots + g^{=\min(k, n - k)}$. We will show that if e satisfies $d < e \leq \min(k, n - k)$ and h is a basic function of degree e then $\langle f, h \rangle = 0$, where the inner product is with respect to the uniform measure on the slice. Since f agrees with g on the slice, it follows that $\langle g, h \rangle = 0$.

Lemma 3.7 shows that $g^{=e}$ is a linear combination of basic functions of degree e (since basic functions are homogeneous), and so $\langle g, g^{=e} \rangle = 0$. Theorem 3.10 implies that $g^{=e} = 0$. Since this is the case for all $e > d$, we conclude that $\deg g \leq d$.

Consider therefore a basic function $h = (x_{a_1} - x_{b_1}) \cdots (x_{a_e} - x_{b_e})$. We have

$$\binom{n}{k} \langle f, h \rangle = \sum_{\substack{x \in \binom{[n]}{k}: \\ x_1 = \cdots = x_d = 1}} (x_{a_1} - x_{b_1}) \cdots (x_{a_e} - x_{b_e}).$$

Since $e > d$, for some i it holds that $a_i, b_i > d$. The only non-zero terms in the sum (if any) are those for which $x_{a_i} \neq x_{b_i}$. We can match each term in which $x_{a_i} - x_{b_i} = 1$ with a term in which $x_{a_i} - x_{b_i} = -1$, obtained by switching x_{a_i} and x_{b_i} . The two terms have opposite signs and so cancel. This shows that the entire sum vanishes, completing the proof. \square

As a corollary, we deduce the so-called *Q-polynomial* or *cometric* property of the slice. For a function f on a slice, denote by $\deg f$ the degree of its unique harmonic multilinear representation on the slice.

Corollary 3.19. *Let f, g be functions on a slice. Then $\deg(fg) \leq \deg f + \deg g$.*

Similarly, we have the following substitution property.

Corollary 3.20. *Let f be a function on a slice, and let g be the function on a smaller slice obtained by fixing one of the coordinates. Then $\deg g \leq \deg f$.*

Both results are non-trivial since the product of two harmonic multilinear polynomials need not be multilinear, and substitution doesn't preserve harmonicity.

4 Invariance principle

In this section we prove an invariance principle showing that the distribution of a low-degree harmonic function on a slice $\binom{[n]}{k}$ is similar to its distribution on the Boolean cube with respect to the measure $\mu_{k/n}$. For convenience, we analyze the similarity in distribution via Lipschitz test functions, and derive similarity in more conventional terms as a corollary. The basic idea is to show that the distribution of a low degree function on a given slice $\binom{[n]}{k}$ is similar to its distribution on nearby slices $\binom{[n]}{\ell}$. If we can show this for all slices satisfying $|k - \ell| \leq B$ for some $B = \omega(\sqrt{n})$, then the invariance follows by decomposing the Boolean cube into a union of slices.

The argument Before giving the formal proof, we provide the intuition underlying the argument. Our argument concerns the following objects:

- A harmonic multilinear polynomial f of degree d and unit norm. We think of d as “small”.
- A Lipschitz function φ .
- A slice $\binom{[n]}{pn}$. We think of p as constant, though the argument even works for sub-constant p .

Our goal is to show that $\mathbb{E}_{\mu_p}[\varphi(f)] \approx \mathbb{E}_{\nu_{pn}}[\varphi(f)]$. The first step is to express μ_p as a mixture of ν_ℓ for various ℓ :

$$\mathbb{E}_{\mu_p}[\varphi(f)] = \sum_{\ell=0}^n \binom{n}{\ell} p^\ell (1-p)^{n-\ell} \mathbb{E}_{\nu_\ell}[\varphi(f)].$$

Applying the triangle inequality, this shows that

$$|\mathbb{E}_{\mu_p}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]| \leq \sum_{\ell=0}^n \binom{n}{\ell} p^\ell (1-p)^{n-\ell} |\mathbb{E}_{\nu_\ell}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]|.$$

In general we expect $|\mathbb{E}_{\nu_\ell}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]|$ to grow with $|\ell - pn|$, and our strategy is to consider separately slices close to pn , say $|pn - \ell| \leq \delta$, and slices far away from pn , say $|pn - \ell| > \delta$. We will bound the contribution of slices close to pn directly. If δ is large enough then we expect the contribution of slices far away from pn to be small, essentially since μ_p is concentrated on slices close to pn . For this argument to work, we need to choose δ so that $\delta = \omega(\sqrt{n})$.

It remains to bound $|\mathbb{E}_{\nu_\ell}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]|$ for ℓ close to pn . One strategy to obtain such a bound is to bound instead $|\mathbb{E}_{\nu_s}[\varphi(f)] - \mathbb{E}_{\nu_{s+1}}[\varphi(f)]|$ for various s , and use the triangle inequality. To this end, it is natural to consider the following coupling: let $(\mathbf{X}(s), \mathbf{X}(s+1)) \in \binom{[n]}{s} \times \binom{[n]}{s+1}$ be chosen uniformly at random under the constraint $\mathbf{X}(s) \subset \mathbf{X}(s+1)$. We can then bound

$$\begin{aligned} |\mathbb{E}_{\nu_s}[\varphi(f)] - \mathbb{E}_{\nu_{s+1}}[\varphi(f)]| &= |\mathbb{E}[\varphi(f(\mathbf{X}(s))) - \varphi(f(\mathbf{X}(s+1)))]| \leq \\ &\mathbb{E}[|\varphi(f(\mathbf{X}(s))) - \varphi(f(\mathbf{X}(s+1)))|] \leq \mathbb{E}[|f(\mathbf{X}(s)) - f(\mathbf{X}(s+1))|]. \end{aligned}$$

Denoting $\pi(s+1) = \mathbf{X}(s+1) \setminus \mathbf{X}(s)$ and using the multilinearity of f , this shows that

$$|\mathbb{E}_{\nu_s}[\varphi(f)] - \mathbb{E}_{\nu_{s+1}}[\varphi(f)]| \leq \mathbb{E} \left[\left| \frac{\partial f}{\partial x_{\pi(s+1)}}(\mathbf{X}(s)) \right| \right] = \mathbb{E} \left[\frac{1}{n-s} \sum_{i \notin \mathbf{X}(s)} \left| \frac{\partial f}{\partial x_i}(\mathbf{X}(s)) \right| \right].$$

While we cannot bound $\sum_i \left| \frac{\partial f}{\partial x_i} \right|$ directly, Lemma 3.16 implies that $\sum_i \left(\frac{\partial f}{\partial x_i} \right)^2 = O(d)$. Applying Cauchy-Schwarz, we get that for s close to pn ,

$$\begin{aligned} |\mathbb{E}_{\nu_s}[\varphi(f)] - \mathbb{E}_{\nu_{s+1}}[\varphi(f)]| &\leq \frac{1}{\Theta(n)} \mathbb{E} \left[\sum_{i=1}^n \left| \frac{\partial f}{\partial x_i}(\mathbf{X}(s)) \right| \right] \\ &\leq \frac{1}{\Theta(n)} \mathbb{E} \left[\sqrt{n} \sqrt{\sum_{i=1}^n \left(\frac{\partial f}{\partial x_i}(\mathbf{X}(s)) \right)^2} \right] = O \left(\sqrt{\frac{d}{n}} \right). \end{aligned}$$

Recall now that our original goal was to bound $|\mathbb{E}_{\nu_\ell}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]|$ for $|\ell - pn| \leq \delta$, and our intended δ satisfied $\delta = \omega(\sqrt{n})$. Unfortunately, the idea just described only gives a bound of the form $|\mathbb{E}_{\nu_\ell}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]| = O(\delta\sqrt{d/n})$, which is useless for our intended δ .

One way out is to take $\delta = C\sqrt{n}$. This allows us to obtain meaningful bounds both on the contribution of slices close to pn and on the contribution of slices far away from pn . Although this only gives a constant upper bound on $|\mathbb{E}_{\mu_p}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]|$ if applied directly, this idea can be used in conjunction with the invariance principle for the Boolean cube [29] to give an invariance principle for the slice, and this is the route chosen in the prequel [18]. One drawback of this approach is that the invariance principle for the Boolean cube requires all influences to be small.

Our approach, in contrast, considers a coupling $(\mathbf{X}(0), \dots, \mathbf{X}(n))$ of *all* slices. Analogous to $f(\mathbf{X}(s+1)) - f(\mathbf{X}(s))$, we consider the quantity

$$\mathbf{C}(s) = (n-s)(f(\mathbf{X}(s+1)) - f(\mathbf{X}(s))) - s(f(\mathbf{X}(s-1)) - f(\mathbf{X}(s))).$$

As before, we can bound $\mathbb{E}[|\mathbf{C}(s)|] = O(\sqrt{dn})$. Moreover,

$$\sum_{u=s}^t \mathbf{C}(u) = (n-t)f(\mathbf{X}(t+1)) + (t+1)f(\mathbf{X}(t)) - (n-s+1)f(\mathbf{X}(s)) - sf(\mathbf{X}(s-1)),$$

and so we can bound $|\mathbb{E}_{\nu_\ell}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)]|$ by bounding the expectation of $\sum_{u=pn}^\ell \mathbf{C}(u)$ or of $\sum_{u=\ell}^{pn} \mathbf{C}(u)$. The triangle inequality gives $|\sum_{u=s}^t \mathbf{C}(u)| = O(|s-t|\sqrt{dn})$, which suffers from the same problem that we

encountered above. However, by expressing $\mathbf{C}(s)$ as a difference of two martingales, we are able to improve on the triangle inequality, showing that

$$\left| \sum_{u=s}^t \mathbf{C}(u) \right| = O(\sqrt{|s-t|dn}),$$

a bound which is useful for $|s-t| = o(n/d)$ rather than for $|s-t| = o(\sqrt{n/d})$ as before.

In more detail, we define

$$\begin{aligned} \mathbf{U}(u) &= f(\mathbf{X}(u+1)) - f(\mathbf{X}(u)) - \mathbb{E}[f(\mathbf{X}(u+1)) - f(\mathbf{X}(u)) | \mathbf{X}(u)], \\ \mathbf{D}(u) &= f(\mathbf{X}(u-1)) - f(\mathbf{X}(u)) - \mathbb{E}[f(\mathbf{X}(u-1)) - f(\mathbf{X}(u)) | \mathbf{X}(u)], \end{aligned}$$

both martingales by construction, $\mathbf{U}(u)$ for increasing u , and $\mathbf{D}(u)$ for decreasing u . We claim that $\mathbf{C}(u) = (n-u)\mathbf{U}(u) - u\mathbf{D}(u)$. If this holds, then using the fact that $\mathbb{E}[\mathbf{U}(u)\mathbf{U}(v)] = \mathbb{E}[\mathbf{D}(u)\mathbf{D}(v)] = 0$ for $u \neq v$ and the L_2 triangle inequality $(a+b)^2 \leq 2a^2 + 2b^2$, we get

$$\begin{aligned} \mathbb{E} \left[\left(\sum_{u=s}^t \mathbf{C}(u) \right)^2 \right] &\leq 2 \mathbb{E} \left[\left(\sum_{u=s}^t (n-u)\mathbf{U}(u) \right)^2 \right] + 2 \mathbb{E} \left[\left(\sum_{u=s}^t u\mathbf{D}(u) \right)^2 \right] \\ &= 2 \sum_{u=s}^t (n-u)^2 \mathbb{E}[\mathbf{U}(u)^2] + 2 \sum_{u=s}^t u^2 \mathbb{E}[\mathbf{D}(u)^2]. \end{aligned}$$

This shows that $\mathbb{E}[(\sum_{u=s}^t \mathbf{C}(u))^2]$ scales linearly in $t-s$ rather than quadratically in $t-s$, which is what we would get if we just applied the triangle inequality. Since the L_1 norm is bounded by the L_2 norm, we conclude that $\mathbb{E}[|\sum_{u=s}^t \mathbf{C}(u)|] = O(\sqrt{|s-t|dn})$.

Finally, let us explain why $\mathbf{C}(u) = (n-u)\mathbf{U}(u) - u\mathbf{D}(u)$. In view of our previous expression for $\mathbf{C}(u)$, this boils down to proving that

$$(n-u) \mathbb{E}[f(\mathbf{X}(u+1)) - f(\mathbf{X}(u)) | \mathbf{X}(u)] - u \mathbb{E}[f(\mathbf{X}(u-1)) - f(\mathbf{X}(u)) | \mathbf{X}(u)] = 0.$$

We can rewrite the left-hand side as

$$\mathbb{E} \left[\sum_{i \notin \mathbf{X}(u)} [f(\mathbf{X}(u) \cup \{i\}) - f(\mathbf{X}(u))] - \sum_{i \in \mathbf{X}(u)} [f(\mathbf{X}(u) \setminus \{i\}) - f(\mathbf{X}(u))] \right].$$

Since f is multilinear, we can replace the differences with derivatives:

$$\mathbb{E} \left[\sum_{i \notin \mathbf{X}(u)} \frac{\partial f}{\partial x_i}(\mathbf{X}(u)) - \sum_{i \in \mathbf{X}(u)} -\frac{\partial f}{\partial x_i}(\mathbf{X}(u)) \right] = \mathbb{E} \left[\sum_{i=1}^n \frac{\partial f}{\partial x_i}(\mathbf{X}(u)) \right].$$

However, the last expression clearly vanishes, since f is harmonic. This completes the outline of the proof.

4.1 The proof

The basic setup of our argument is described in the following definition.

Definition 4.1. We are given a harmonic multilinear polynomial f of degree $d \geq 1$.

Let $\boldsymbol{\pi} \in S_n$ be chosen uniformly at random, and define random variables $\mathbf{X}(s) \in \binom{[n]}{s}$ for $0 \leq s \leq n$ as follows:

$$\mathbf{X}(s)_i = \begin{cases} 1 & \text{if } i \in \boldsymbol{\pi}(\{1, \dots, s\}), \\ 0 & \text{if } i \in \boldsymbol{\pi}(\{s+1, \dots, n\}). \end{cases}$$

For $0 \leq s \leq n$, define random variables $\mathbf{U}(s)$ (for $s \neq n$), $\mathbf{D}(s)$ (for $s \neq 0$) and $\mathbf{C}(s)$ by

$$\begin{aligned}\mathbf{U}(s) &= f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)) - \mathbb{E}[f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)) | \mathbf{X}(s)], \\ \mathbf{D}(s) &= f(\mathbf{X}(s-1)) - f(\mathbf{X}(s)) - \mathbb{E}[f(\mathbf{X}(s-1)) - f(\mathbf{X}(s)) | \mathbf{X}(s)], \\ \mathbf{C}(s) &= (n-s)\mathbf{U}(s) - s\mathbf{D}(s).\end{aligned}$$

In words, $\mathbf{X}(0), \dots, \mathbf{X}(n)$ is a random maximal chain in the Boolean cube. These random variables form a coupling of the uniform distributions on all slices. The random variable $\mathbf{U}(s)$ measures the effect of moving up from $\mathbf{X}(s)$ to $\mathbf{X}(s+1)$ on f , normalized so that it has zero mean given $\mathbf{X}(s)$. The random variable $\mathbf{D}(s)$ similarly measures the effect of moving down from $\mathbf{X}(s)$ to $\mathbf{X}(s-1)$. Finally, $\mathbf{C}(s)$ measures the effect of moving away from $\mathbf{X}(s)$. The usefulness of this representation stems from the fact that $\mathbf{U}(s)$ and $\mathbf{D}(s)$ are martingale differences and are therefore orthogonal while $\mathbf{C}(s)$ is on one hand easily expressed in terms of $\mathbf{D}(s)$ and $\mathbf{U}(s)$ and on the other hand is useful for bounding differences of f . This representation of a random directed path in terms of martingales and reverse martingales is inspired by previous work using such representations for stationary reversible Markov chains and for stochastic integrals [28, 30]. Part of the novelty of our proof is using such a representation in a non-stationary setup. In particular, the formula for the representation is significantly different from previous applications of the method. The basic properties of the representation are stated formally in the following lemma.

Lemma 4.2. *The following properties hold:*

(a) *If $s \neq t$ then $\mathbb{E}[\mathbf{U}(s)\mathbf{U}(t)] = \mathbb{E}[\mathbf{D}(s)\mathbf{D}(t)] = 0$.*

(b) *For all s we can bound*

$$\begin{aligned}\mathbb{E}[\mathbf{U}(s)^2] &\leq 4\mathbb{E}[(f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)))^2] \leq \frac{4}{n-s} \sum_{i=1}^n \left\| \frac{\partial f}{\partial x_i} \right\|_{\nu_s}^2, \\ \mathbb{E}[\mathbf{D}(s)^2] &\leq 4\mathbb{E}[(f(\mathbf{X}(s-1)) - f(\mathbf{X}(s)))^2] \leq \frac{4}{s} \sum_{i=1}^n \left\| \frac{\partial f}{\partial x_i} \right\|_{\nu_s}^2.\end{aligned}$$

(c) *For all s we have $\mathbf{C}(s) = (n-s)f(\mathbf{X}(s+1)) - (n-2s)f(\mathbf{X}(s)) - sf(\mathbf{X}(s-1))$.*

(d) *For all $s \leq t$ we have*

$$\sum_{u=s}^t \mathbf{C}(u) = (n-t)f(\mathbf{X}(t+1)) + (t+1)f(\mathbf{X}(t)) - (n-s+1)f(\mathbf{X}(s)) - sf(\mathbf{X}(s-1)).$$

Proof. For (a) note that if $s < t$ then $\mathbb{E}[\mathbf{U}(t) | \mathbf{X}(s+1), \mathbf{X}(s)] = 0$, and so $\mathbb{E}[\mathbf{U}(t)\mathbf{U}(s)] = 0$. For similar reasons we have $\mathbb{E}[\mathbf{D}(t)\mathbf{D}(s)] = 0$.

For (b), note first that the L_2 triangle inequality implies that $\mathbb{E}[\mathbf{U}(s)^2] \leq 4\mathbb{E}[(f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)))^2]$, using the bound

$$\mathbb{E}_{\mathbf{X}(s)}[\mathbb{E}[f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)) | \mathbf{X}(s)]^2] \leq \mathbb{E}_{\mathbf{X}(s)}[\mathbb{E}[(f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)))^2 | \mathbf{X}(s)]] = \mathbb{E}[(f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)))^2].$$

Since f is multilinear, we have $f(\mathbf{X}(s+1)) = f(\mathbf{X}(s)) + \frac{\partial f}{\partial x_{\pi(s+1)}}(\mathbf{X}(s))$, and so

$$\begin{aligned}\mathbb{E}[\mathbf{U}(s)^2] &\leq 4\mathbb{E}\left[\left(\frac{\partial f}{\partial x_{\pi(s+1)}}(\mathbf{X}(s))\right)^2\right] \\ &= \frac{4}{n-s}\mathbb{E}\left[\sum_{i: \mathbf{X}(s)_i=0} \left(\frac{\partial f}{\partial x_i}(\mathbf{X}(s))\right)^2\right] \\ &\leq \frac{4}{n-s}\sum_{i=1}^n \mathbb{E}\left[\left(\frac{\partial f}{\partial x_i}(\mathbf{X}(s))\right)^2\right].\end{aligned}$$

This implies our estimate for $\mathbb{E}[\mathbf{U}(s)^2]$. The estimate for $\mathbb{E}[\mathbf{D}(s)^2]$ is obtained in the same way.

For (c), we again use the formula $f(\mathbf{X}(s+1)) = f(\mathbf{X}(s)) + \frac{\partial f}{\partial x_{\pi(s+1)}}(\mathbf{X}(s))$ to obtain

$$\mathbb{E}[f(\mathbf{X}(s+1)) - f(\mathbf{X}(s)) | \mathbf{X}(s)] = \frac{1}{n-s} \mathbb{E} \left[\sum_{i: \mathbf{X}(s)_i=0} \frac{\partial f}{\partial x_i}(\mathbf{X}(s)) \right].$$

Similarly, using the formula $f(\mathbf{X}(s-1)) = f(\mathbf{X}(s)) - \frac{\partial f}{\partial x_{\pi(s-1)}}(\mathbf{X}(s))$ instead, we have

$$\mathbb{E}[f(\mathbf{X}(s-1)) - f(\mathbf{X}(s)) | \mathbf{X}(s)] = -\frac{1}{s} \mathbb{E} \left[\sum_{i: \mathbf{X}(s)_i=1} \frac{\partial f}{\partial x_i}(\mathbf{X}(s)) \right].$$

Combining these formulas together, we obtain

$$\begin{aligned} \mathbf{C}(s) &= (n-s)f(\mathbf{X}(s+1)) - (n-2s)f(\mathbf{X}(s)) - sf(\mathbf{X}(s-1)) + \mathbb{E} \left[\sum_{i=1}^n \frac{\partial f}{\partial x_i}(\mathbf{X}(s)) \right] \\ &= (n-s)f(\mathbf{X}(s+1)) - (n-2s)f(\mathbf{X}(s)) - sf(\mathbf{X}(s-1)), \end{aligned}$$

since f is harmonic.

Part (d) follows by simple computation, whose highlight is noticing that the coefficient of $f(\mathbf{X}(u))$ for $s < u < t$ in the sum is $(n-u+1) - (n-2u) - (u+1) = 0$. \square

The exact definition of $\mathbf{C}(s)$ is aimed at the cancellation of the derivative terms in the proof of Lemma 4.2(c), in which we use the harmonicity of f . We also use the harmonicity of f to bound the variance of f with respect to various slices.

The next step is bounding the quantity appearing in Lemma 4.2(b) in terms of d , the degree of f . This is a simple application of Lemma 3.16 on page 13.

Lemma 4.3. *For every integer $0 \leq qn \leq n$ such that $d \leq q(1-q)n$ we have*

$$\sum_{i=1}^n \left\| \frac{\partial f}{\partial x_i} \right\|_{\nu_{qn}}^2 = O \left(\frac{d}{q(1-q)} \right) \mathbb{V}[f]_{\nu_{qn}}.$$

Proof. Combine Lemma 3.16 with Lemma 3.17. \square

Our work so far has not focused on any specific slice. Now we turn to the following question. Suppose that f has unit variance on the slice $\binom{[n]}{pn}$. What can we say about its behavior on a nearby slice $\binom{[n]}{qn}$?

Lemma 4.4. *There exists a constant $K > 0$ such that the following holds. Let $0 \leq pn, qn \leq n$ be integers such that $|p-q| \leq p(1-p)/(2d)$, and suppose that $\mathbb{V}[f]_{\nu_{pn}} = 1$. If $d^2 \leq Kp(1-p)n$ then*

$$\mathbb{E}[(f(\mathbf{X}(pn)) - f(\mathbf{X}(qn)))^2] = O \left(\frac{|q-p|d}{p(1-p)} \right).$$

Proof. Let $k = pn$ and $\ell = qn$, and assume that $k < \ell$ (the other case is very similar). Note that $|(d/d\rho)\rho(1-\rho)| = |1-2\rho| \leq 1$ for $\rho \in [0, 1]$, and so $|p(1-p) - q(1-q)| \leq |p-q| \leq p(1-p)/(2d)$. Therefore for $K \leq 1/2$ the assumption $d^2 \leq Kp(1-p)n$ implies that $d^2 \leq r(1-r)n$ for all $r \in [p, q]$, since $r(1-r) \geq p(1-p) - p(1-p)/(2d) \geq p(1-p)/2$.

Lemma 4.2(d) implies that

$$\Delta \triangleq \sum_{u=k}^{\ell} \mathbf{C}(u) = (n-\ell)f(\mathbf{X}(\ell+1)) + (\ell+1)f(\mathbf{X}(\ell)) - (n-k+1)f(\mathbf{X}(k)) - kf(\mathbf{X}(k-1)).$$

As a first step, we bound $\mathbb{E}[\Delta^2]$. The L_2 triangle inequality implies that

$$\mathbb{E}[\Delta^2] \leq 2 \underbrace{\mathbb{E} \left[\left(\sum_{u=k}^{\ell} (n-u) \mathbf{U}(u) \right)^2 \right]}_{\Delta_U^2} + 2 \underbrace{\mathbb{E} \left[\left(\sum_{u=k}^{\ell} u \mathbf{D}(u) \right)^2 \right]}_{\Delta_D^2}.$$

Lemma 4.2(a) and Lemma 4.2(b) show that

$$\Delta_U^2 = \sum_{u=k}^{\ell} (n-u)^2 \mathbb{E}[\mathbf{U}(u)^2] \leq \sum_{u=k}^{\ell} 4(n-u) \sum_{i=1}^n \left\| \frac{\partial f}{\partial x_i} \right\|_{\nu_u}^2.$$

Applying Lemma 4.3 (using $d^2 \leq r(1-r)n$ for all $r \in [p, q]$), we obtain

$$\Delta_U^2 \leq \sum_{u=k}^{\ell} O \left(\frac{d(n-u)}{(u/n)(1-u/n)} \right) \mathbb{V}[f]_{\nu_u} \leq \sum_{u=k}^{\ell} O \left(\frac{dn}{p(1-p)} \right) \mathbb{V}[f]_{\nu_u}.$$

Putting $u = rn$, Corollary 3.11 together with Lemma 3.12 shows that

$$\mathbb{V}[f]_{\nu_u} \leq \max_{e \leq d} \frac{(2r(1-r))^e (1 \pm O(\frac{e^2}{p(1-p)n}))}{(2p(1-p))^e (1 \pm O(\frac{e^2}{p(1-p)n}))},$$

using the observation $r(1-r) \geq p(1-p)/2$ as well as $\mathbb{V}[f]_{\nu_{pn}} = 1$. For small enough $K > 0$ we have $1 \pm O(\frac{e^2}{p(1-p)n}) = \Theta(1)$. As observed above, $|p(1-p) - r(1-r)| \leq |p-r| \leq p(1-p)/2d$, and so

$$\mathbb{V}[f]_{\nu_u} \leq O(1) \cdot \max_{e \leq d} \left(1 + \frac{1}{2d} \right)^e = O(1).$$

Therefore

$$\Delta_U^2 \leq \sum_{u=k}^{\ell} O \left(\frac{dn}{p(1-p)} \right) = O \left(\frac{|\ell - k|dn}{p(1-p)} \right).$$

We can bound Δ_D^2 similarly, and conclude that

$$\mathbb{E}[\Delta^2] = O \left(\frac{|\ell - k|dn}{p(1-p)} \right).$$

The triangle inequality implies that

$$(n+1) \|f(\mathbf{X}(\ell)) - f(\mathbf{X}(k))\| \leq \|\Delta\| + (n-\ell) \|f(\mathbf{X}(\ell+1)) - f(\mathbf{X}(\ell))\| + k \|f(\mathbf{X}(k-1)) - f(\mathbf{X}(k))\|.$$

The latter two terms can be bounded in the same way that we bounded $\mathbb{E}[\mathbf{U}(u)^2]$ above, and we conclude that

$$(n+1) \|f(\mathbf{X}(\ell)) - f(\mathbf{X}(k))\| = O \left(\sqrt{\frac{|\ell - k|dn}{p(1-p)}} \right).$$

This implies that

$$\mathbb{E}[(f(\mathbf{X}(pn)) - f(\mathbf{X}(qn)))^2] = O \left(\frac{|\ell - k|d}{p(1-p)n} \right) = O \left(\frac{|q-p|d}{p(1-p)} \right). \quad \square$$

Combining Lemma 4.4 with Chernoff's bound, we obtain our main theorem.

Proposition 4.5. *Let $0 < p, \epsilon < 1$ satisfy $\epsilon \leq p(1-p)$. Then*

$$\Pr[|B(n, p) - np| > \epsilon n] \leq 2e^{-\epsilon^2 n / (6p(1-p))}.$$

Proof. Suppose first that $p \leq 1/2$, and let $X \sim B(n, p)$ and $\mu = np$. One common version of Chernoff's bound states that for $0 < \delta < 1$ we have

$$\Pr[|X - \mu| > \delta \mu] \leq 2e^{-\delta^2 \mu / 3}.$$

Choose $\delta = \epsilon/p$, and note that $\delta \leq 1-p < 1$. Since $\delta^2 \mu = (\epsilon^2/p^2)(np) = \epsilon^2 n/p \geq \epsilon^2 n / (2p(1-p))$, the bound follows in this case.

When $p \geq 1/2$, we look at $X \sim B(n, 1-p)$ instead, using the fact that $n - X \sim B(n, p)$. \square

Theorem 4.6. *There exists a constant $K > 0$ such that the following holds, whenever $p(1-p)n \geq 3 \log n$.*

Let f be a harmonic multilinear polynomial of degree d satisfying $d \leq K \sqrt{p(1-p)n}$ such that $\mathbb{V}[f]_{\nu_{pn}} = 1$. For any Lipschitz function φ ,

$$\left| \mathbb{E}_{\nu_{pn}}[\varphi(f)] - \mathbb{E}_{\mu_p}[\varphi(f)] \right| = O\left(\sqrt{\frac{d}{p(1-p)n}}\right).$$

Proof. We can assume, without loss of generality, that $\mathbb{E}[f] = 0$ (Corollary 3.11 implies that the expectation is the same with respect to both ν_{pn} and μ_p), so that $\mathbb{E}[f^2] = 1$, and that $\varphi(0) = 0$. Since φ is Lipschitz, it follows that $|\varphi(x)| \leq |x|$ for all x .

Let $\delta := p(1-p)/(2d)$. Lemma 4.4 implies that whenever $|q-p| \leq \delta$,

$$\begin{aligned} \left| \mathbb{E}_{\nu_{pn}}[\varphi(f)] - \mathbb{E}_{\nu_{qn}}[\varphi(f)] \right| &= \left| \mathbb{E}[\varphi(f(\mathbf{X}(pn))) - \varphi(f(\mathbf{X}(qn)))] \right| \leq \left| \mathbb{E}[f(\mathbf{X}(pn)) - f(\mathbf{X}(qn))] \right| \\ &\leq \|f(\mathbf{X}(pn)) - f(\mathbf{X}(qn))\| = O\left(\sqrt{\frac{|q-p|d}{p(1-p)}}\right). \end{aligned}$$

Expressing μ_p as a mixture of distributions of the form ν_{qn} , we obtain

$$\begin{aligned} \left| \mathbb{E}_{\nu_{pn}}[\varphi(f)] - \mathbb{E}_{\mu_p}[\varphi(f)] \right| &= \left| \mathbb{E}_{\nu_{pn}}[\varphi(f)] - \sum_{\ell=0}^n \binom{n}{\ell} p^\ell (1-p)^{n-\ell} \mathbb{E}_{\nu_\ell}[\varphi(f)] \right| \\ &\leq \sum_{\ell=0}^n \binom{n}{\ell} p^\ell (1-p)^{n-\ell} \left| \mathbb{E}_{\nu_{pn}}[\varphi(f)] - \mathbb{E}_{\nu_\ell}[\varphi(f)] \right| \\ &\leq O\left(\sqrt{\frac{d}{p(1-p)}}\right) \underbrace{\mathbb{E}_{\ell \sim B(n,p)} \left[\sqrt{\left| \frac{\ell}{n} - p \right|} \right]}_C + \\ &\quad \underbrace{\Pr[|B(n, p) - pn| > \delta n]}_{\epsilon_1} \cdot \underbrace{\mathbb{E}_{\nu_{pn}}[|\varphi(f)|] + \mathbb{E}_{x \sim \mu_p}[|\varphi(f(x))| \cdot \mathbf{1}_{|\sum_i x_i - pn| > \delta n}]}_{\epsilon_2}. \end{aligned}$$

The coefficient C can be bounded using Proposition 4.5 and the formula $\mathbb{E}[X] = \int_0^\infty \Pr[X \geq t] dt$ (for $X \geq 0$):

$$\begin{aligned} C &= \int_0^\infty \Pr[\sqrt{|B(n, p)/n - p|} \geq t] dt = \int_0^\infty \Pr[|B(n, p) - pn| \geq t^2 n] dt \\ &\leq \int_0^{\sqrt{p(1-p)}} 2e^{-t^4 n / (6p(1-p))} dt + \int_{\sqrt{p(1-p)}}^1 2e^{-p(1-p)n/6} dt. \end{aligned}$$

Substitute $s = \sqrt[4]{n/6p(1-p)} \cdot t$ to get

$$C \leq 2 \sqrt[4]{\frac{6p(1-p)}{n}} \int_0^\infty e^{-s^4} ds + O(e^{-p(1-p)n/6}) = O\left(\sqrt[4]{\frac{p(1-p)}{n}}\right),$$

since $e^{-p(1-p)n/6} \leq \frac{1}{\sqrt{n}}$ whereas $\sqrt[4]{\frac{p(1-p)}{n}} = \sqrt[4]{\frac{p(1-p)n}{n^2}} \geq \frac{\sqrt[4]{3 \log n}}{\sqrt{n}} \geq \frac{1}{\sqrt{n}}$.

We proceed to bound the error terms ϵ_1, ϵ_2 . With respect to ν_{pn} , $\mathbb{E}[|\varphi(f)|] \leq \mathbb{E}[|f|] \leq \|f\| \leq 1$, and so Proposition 4.5 implies that

$$\epsilon_1 \leq 2e^{-\delta^2 n/(6p(1-p))}.$$

For the second error term, let $M > 0$ be a parameter to be determined. We have

$$\epsilon_2 \leq M \Pr[|B(n, p) - pn| > \delta n] + \mathbb{E}_{x \sim \mu_p} [|\varphi(f(x))| \cdot \mathbf{1}_{|\varphi(f(x))| > M}].$$

If $|\varphi(f(x))| > M$ then certainly $|f(x)| \geq |\varphi(f(x))| > M$, and so

$$\mathbb{E}_{x \sim \mu_p} [|\varphi(f(x))| \cdot \mathbf{1}_{|\varphi(f(x))| > M}] \leq \mathbb{E}_{x \sim \mu_p} [|f(x)| \cdot \mathbf{1}_{|f(x)| > M}] \leq \frac{1}{M} \mathbb{E}_{\mu_p} [f^2] = O\left(\frac{1}{M}\right),$$

since Lemma 3.12 and Corollary 3.11 imply that $\|f\|_{\mu_p} = O(\|f\|_{\nu_{pn}}) = O(1)$ for small enough K . Hence

$$\epsilon_2 \leq 2Me^{-\delta^2 n/(6p(1-p))} + O\left(\frac{1}{M}\right).$$

Choosing $M = e^{-\delta^2 n/(12p(1-p))}$, we conclude that

$$\epsilon_2 = O(e^{-\delta^2 n/(12p(1-p))}).$$

Putting everything together, we obtain

$$|\mathbb{E}_{\nu_{pn}} [\varphi(f)] - \mathbb{E}_{\mu_p} [\varphi(f)]| = O\left(\sqrt{\frac{d}{\sqrt{p(1-p)n}}} + e^{-\frac{\delta^2 n}{12p(1-p)}}\right).$$

Substituting $\delta = p(1-p)/(2d)$, we deduce

$$|\mathbb{E}_{\nu_{pn}} [\varphi(f)] - \mathbb{E}_{\mu_p} [\varphi(f)]| = O\left(\underbrace{\sqrt{\frac{d}{\sqrt{p(1-p)n}}}}_A + \underbrace{e^{-\frac{p(1-p)n}{48d^2}}}_B\right).$$

Note that $B = e^{-1/(48A^4)}$. When $A < 0.38$, calculation shows that $B < A$. Since $A \leq \sqrt{K}$, choosing $K \leq 0.38^2$ ensures that $B < A$. \square

As a corollary, we can estimate the Lévy distance between $f(\nu_{pn})$ and $f(\mu_p)$, along the lines of [29, Theorem 3.19(28)].

Corollary 4.7. *Suppose that $p(1-p)n \geq 3 \log n$, and let f be a harmonic multilinear polynomial of degree d satisfying $d \leq K \sqrt{p(1-p)n}$ such that $\mathbb{V}[f]_{\nu_{pn}} = 1$, where $K > 0$ is the constant from Theorem 4.6. The Lévy distance between $f(\nu_{pn})$ and $f(\mu_p)$ is at most*

$$\epsilon = O\left(\sqrt[4]{\frac{d}{\sqrt{p(1-p)n}}}\right).$$

That is, for all y it holds that

$$\Pr_{\nu_{pn}} [f \leq y - \epsilon] - \epsilon \leq \Pr_{\mu_p} [f \leq y] \leq \Pr_{\nu_{pn}} [f \leq y + \epsilon] + \epsilon.$$

Proof. Given y , define a $(1/\epsilon)$ -Lipschitz function φ by

$$\varphi(t) = \begin{cases} 1 & \text{if } t \leq y, \\ \frac{y+\epsilon-t}{\epsilon} & \text{if } y \leq t \leq y + \epsilon, \\ 0 & \text{if } t \geq y + \epsilon. \end{cases}$$

It is easy to check that $\mathbf{1}_{t \leq y} \leq \varphi(t) \leq \mathbf{1}_{t \leq y + \epsilon}$, and so

$$\Pr_{\mu_p}[f \leq y] - \Pr_{\nu_{pn}}[f \leq y + \epsilon] \leq \mathbb{E}_{\mu_p}[\varphi(f)] - \mathbb{E}_{\nu_{pn}}[\varphi(f)] = \frac{1}{\epsilon} O\left(\sqrt{\frac{d}{\sqrt{p(1-p)}n}}\right) = \epsilon,$$

using Theorem 4.6 and the correct choice of ϵ . We get the bound in the other direction in the same way. \square

We conjecture that $f(\nu_{pn})$ and $f(\mu_p)$ are also close in CDF distance. Unfortunately, the method of proof of [29, Theorem 3.19(30)] relies on the anticoncentration of multivariate Gaussian distributions [8], whereas both $f(\nu_{pn})$ and $f(\mu_p)$ are discrete distributions. We consider it an interesting open problem to extend Corollary 4.7 to CDF distance.

Question 4.8. *Suppose that $p(1-p)n$ is “large” and d is “small”, compared to n . Is it true that for every harmonic multilinear polynomial f of degree d satisfying $\mathbb{V}[f]_{\nu_{pn}} = 1$, the CDF distance between $f(\nu_{pn})$ and $f(\mu_p)$ is $o(1)$?*

4.2 High-degree functions

Theorem 4.6 requires that $d = O(\sqrt{p(1-p)n})$. Indeed, Lemma 3.12, which implies that the norm of a low-degree function is approximately the same under both μ_p and ν_{pn} , already requires the degree to be $O(\sqrt{p(1-p)n})$. For $d = \omega(\sqrt{p(1-p)n})$ and constant $p \neq \frac{1}{2}$ we exhibit below a $0/\pm C$ -valued function f (for some C depending on p and d) which satisfies $\|f\|_{\mu_p} = 1$ while $\|f\|_{\nu_{pn}} = o(1)$. This shows that for constant $p \neq \frac{1}{2}$ the dependence on the degree is essential in Theorem 4.6, since $|\mathbb{E}_{\nu_{pn}}[|f|] - \mathbb{E}_{\mu_p}[|f|]| = \|f\|_{\mu_p}^2 - \|f\|_{\nu_{pn}}^2 = 1 - o(1)$. We do not know whether this dependence is necessary for $p = \frac{1}{2}$. Indeed, Lemma 3.12 can be extended above \sqrt{n} in this case, as the calculation below shows.

Let $d = \omega(\sqrt{p(1-p)n})$, and assume further that $d = o((p(1-p)n)^{2/3})$. We consider the function $f = (2p(1-p))^{-d/2}(x_1 - x_2) \cdots (x_{2d-1} - x_{2d})$, whose μ_p -norm is 1 according to Lemma 3.12. The lemma also gives its ν_k -norm (where $k = pn$) as

$$\|f\|_{\nu_k}^2 = (p(1-p))^{-d} \frac{k^d (n-k)^d}{n^{2d}}.$$

We estimate this expression using Stirling’s approximation, starting with k^d :

$$k^d = \frac{k!}{(k-d)!} = \left(\frac{k}{k-d}\right)^{k-d+1/2} \frac{k^d}{e^d} e^{O(1/k) - O(1/(k-d))} = \left(1 + \frac{d}{k-d}\right)^{k-d} \frac{k^d}{e^d} (1 \pm o(1)).$$

The Taylor series $\log(1+x) = x - x^2/2 + O(x^3)$ shows that

$$\left(1 + \frac{d}{k-d}\right)^{k-d} = \exp\left[d - \frac{d^2}{2(k-d)} + o(1)\right] = \exp\left[d - \frac{d^2}{2k} + o(1)\right],$$

and so $k^d = k^d e^{-d^2/2k} (1 \pm o(1))$. We can similarly estimate $(n-k)^d = (n-k)^d e^{-d^2/2(n-k)} (1 \pm o(1))$ and $n^{2d} = n^{2d} e^{-2d^2/n} (1 \pm o(1))$, concluding that

$$\begin{aligned} \|f\|_{\nu_k}^2 &= (p(1-p))^{-d} \frac{k^d (n-k)^d}{n^{2d}} e^{-d^2/2k - d^2/2(n-k) + 2d^2/n} (1 \pm o(1)) \\ &= \exp\left[\frac{d^2}{2p(1-p)n} (-1 + 4p(1-p)) \pm o(1)\right]. \end{aligned}$$

If $p \neq \frac{1}{2}$ is fixed, we immediately conclude that $\|f\|_{\nu_k} = o(1)$.

The constant $C = (2p(1-p))^{-d/2}$ is unbounded as a function of d . We do not know whether Theorem 4.6 can be extended to higher degrees for bounded functions.

5 Approximately Boolean functions

Theorem 4.6 only applies to Lipschitz test functions, but in many applications we are interested in functions which grow faster, for example the squared-distance-from- $\{-1, 1\}$ function $\varphi(x) = (|x| - 1)^2$. In this section we show how to handle such functions using hypercontractivity.

Proposition 5.1. *Denote by $\|\cdot\|_r$ the L_r norm. For a multilinear polynomial f , let T_ρ denote the operator*

$$T_\rho f = \sum_{i=0}^n \rho^i f^{=i}.$$

For $r \geq 2$ and with respect to μ_p ,

$$\|T_\rho f\|_r \leq \|f\|_2, \text{ where } \rho = \sqrt{\frac{p(1-p)}{r-1}}.$$

For a harmonic multilinear polynomial f , let H_ρ denote the operator

$$H_\rho f = \sum_{i=0}^{n/2} \rho^{i(1-(i-1)/n)} f^{=i}.$$

For $r \geq 2$ and with respect to ν_{pn} ,

$$\|H_\rho f\|_r \leq \|f\|_2, \text{ where } \rho = (r-1)^{\Theta(\log(p(1-p)))} = (p(1-p))^{\Theta(\log(r-1))}.$$

Proof. The first result is classical, appearing in [6, 3] for $p = 1/2$ and in [40, 33] for general p . The second result is due to Lee and Yau [26]. They proved the corresponding log-Sobolev inequality, which implies hypercontractivity as shown in [13]. \square

These results imply that for low-degree functions, the L_r norm and the L_2 norm are comparable.

Lemma 5.2. *Let f be a multilinear polynomial of degree d , and let $r \geq 2$ be a constant. With respect to μ_p ,*

$$\|f\|_r \leq O(p(1-p))^{-O(d)} \|f\|_2.$$

If f is also harmonic, then with respect to ν_{pn} ,

$$\|f\|_r \leq O(p(1-p))^{-O(d)} \|f\|_2.$$

Proof. In both cases, we apply Proposition 5.1 to $T_{1/\rho}f$ or to $H_{1/\rho}f$, where ρ is given by the proposition. In the first case, we get

$$\|f\|_r^2 \leq \|T_{1/\rho}f\|_2^2 = \sum_{i=0}^d \rho^{-2i} \|f^{=i}\|_2^2 \leq \rho^{-2d} \|f\|_2^2.$$

In the second case, we get

$$\|f\|_r^2 \leq \|H_{1/\rho}f\|_2^2 = \sum_{i=0}^d \rho^{-2i(1-(i-1)/n)} \|f^{=i}\|_2^2 \leq \rho^{-2d} \|f\|_2^2.$$

We obtain the stated bounds by substituting the values of ρ . \square

We can now obtain our invariance principle for $\varphi(x) = (|x| - 1)^2$.

Theorem 5.3. *Suppose that $p(1-p)n \geq 3 \log n$, and let f be a harmonic multilinear polynomial of degree d satisfying $d \leq K \sqrt{p(1-p)n}$ such that $\|f\|_{\nu_{pn}} = 1$, where K is the constant in Theorem 4.6. We have*

$$\left| \mathbb{E}_{\nu_{pn}} [(|f| - 1)^2] - \mathbb{E}_{\mu_p} [(|f| - 1)^2] \right| = O \left(\sqrt[4]{\frac{d}{\sqrt{p(1-p)n}}} (p(1-p))^{-O(d)} \right).$$

Proof. Note that $\|f\|_{\nu_{pn}} = 1$ implies that $\mathbb{V}[f]_{\nu_{pn}} \leq 1$ and $\|f\|_{\mu_p} = O(1)$, due to Corollary 3.11 and Lemma 3.12.

Let $M \geq 1$ be a parameter to be decided, and define the $2(M-1)$ -Lipschitz function ψ by

$$\psi(x) = \begin{cases} (|x| - 1)^2 & \text{if } |x| \leq M, \\ (M - 1)^2 & \text{if } |x| \geq M. \end{cases}$$

When $|x| \geq M$, we have $(|x| - 1)^2 \leq x^2$, and so with respect to any measure we have

$$\mathbb{E}[(|f| - 1)^2] \leq \mathbb{E}[\psi(f)] + \mathbb{E}[f^2 \mathbf{1}_{|f| \geq M}].$$

With respect to either μ_p or ν_{pn} , using Lemma 5.2 we can bound

$$\mathbb{E}[f^2 \mathbf{1}_{|f| \geq M}] \leq \frac{1}{M} \mathbb{E}[|f|^3] = \frac{1}{M} \|f\|_3^3 \leq \frac{1}{M} O(p(1-p))^{-O(d)}.$$

Theorem 4.6 therefore implies that

$$\left| \mathbb{E}_{\nu_{pn}} [(|f| - 1)^2] - \mathbb{E}_{\mu_p} [(|f| - 1)^2] \right| \leq MO \left(\sqrt{\frac{d}{\sqrt{p(1-p)n}}} \right) + \frac{1}{M} O(p(1-p))^{-O(d)}.$$

Choosing M to be the geometric mean of both terms appearing above results in the statement of the theorem. \square

As an illustration of this theorem, we give an alternative proof of [18, Theorem 7.5], a Kindler–Safra theorem for the slice.

Definition 5.4. A function f on a given domain is *Boolean* if on the domain it satisfies $f \in \{\pm 1\}$. If the domain is a cube, we use the term *cube-Boolean*. If it is a slice, we use the term *slice-Boolean*.

Proposition 5.5. *Let f be a multilinear polynomial of degree d such that $\mathbb{E}_{\mu_p} [(|f| - 1)^2] = \epsilon$. There exists a cube-Boolean function g on $(p(1-p))^{-O(d)}$ coordinates such that $\|f - g\|_{\mu_p}^2 = O((p(1-p))^{-O(d)} \epsilon)$.*

Proof. This is essentially proved in [24, 23]. Explicitly, they prove the same result without the guarantee that g is cube-Boolean. In order to get our version, let $F = \text{sgn } f$ and $G = \text{sgn } g$. By definition $\|F - f\|^2 = \epsilon$, and so $\|F - g\|^2 = O((p(1-p))^{-O(d)} \epsilon)$. Since F is cube-Boolean, this implies that $\|F - G\|^2 = O((p(1-p))^{-O(d)} \epsilon)$. We conclude that $\|f - G\|^2 = O((p(1-p))^{-O(d)} \epsilon)$. \square

Theorem 5.6. *Let f be a slice-Boolean harmonic multilinear polynomial such that $\|f^{>d}\|_{\nu_{pn}}^2 = \epsilon$, where $f^{>d} = \sum_{i>d} f^{=i}$. There exists a slice-Boolean harmonic multilinear polynomial h depending on $(p(1-p))^{-O(d)}$ coordinates (that is, invariant to permutations of the other coordinates) satisfying*

$$\|f - h\|_{\nu_{pn}}^2 \leq O((p(1-p))^{-O(d)} \epsilon) + O_{p,d} \left(\frac{1}{n^{1/8}} \right).$$

Before we can prove this theorem, we need an auxiliary result [18, Theorem 3.3], which we prove here (simplifying the original proof) to make the proof self-contained. The proof uses Corollary 10.4, whose self-contained proof appears in Subsection 10.1.

Theorem 5.7. *Let f be a multilinear polynomial depending on $M \leq \min(p, 1-p)n$ variables, and let \tilde{f} be the unique harmonic multilinear polynomial of degree at most $\min(p, 1-p)n$ agreeing with f on the slice $\binom{[n]}{pn}$ (the harmonic projection of f on the slice $\binom{[n]}{pn}$). Then*

$$\|f - \tilde{f}\|_{\mu_p}^2 = O\left(\frac{M^2 2^M}{p(1-p)n}\right) \|f\|_{\mu_p}^2.$$

Proof. The Fourier expansion of f with respect to μ_p is

$$f = \sum_{S \subseteq [M]} \hat{f}(S) \omega_S, \text{ where } \omega_S = \prod_{i \in S} \frac{x_i - p}{\sqrt{p(1-p)}}.$$

The characters ω_S are orthogonal and have unit norm with respect to μ_p . This shows that $\|f\|_{\mu_p}^2 = \sum_{S \subseteq [M]} \hat{f}(S)^2$. Harmonic projection is a linear operator, and so denoting the harmonic projection of ω_S by $\tilde{\omega}_S$, we have

$$f - \tilde{f} = \sum_{S \subseteq [M]} \hat{f}(S) (\omega_S - \tilde{\omega}_S).$$

We show below that

$$\|\omega_S - \tilde{\omega}_S\|_{\mu_p}^2 = O\left(\frac{|S|^2}{p(1-p)n}\right). \quad (1)$$

Assuming without loss of generality that f depends on the first M variables, the L_2 triangle inequality then implies that

$$\|f - \tilde{f}\|_{\mu_p}^2 \leq 2^M \sum_{S \subseteq [M]} \hat{f}(S)^2 \|\omega_S - \tilde{\omega}_S\|_{\mu_p}^2 \leq O\left(\frac{M^2 2^M}{p(1-p)n}\right) \sum_{S \subseteq [M]} \hat{f}(S)^2 = O\left(\frac{M^2 2^M}{p(1-p)n}\right) \|f\|_{\mu_p}^2.$$

It remains to prove (1). For definiteness, consider $S = \{1, \dots, d\}$, where $d \leq \min(p, 1-p)n$. The first step is to consider the related function $\chi_S = \prod_{i=1}^d x_i$ and its harmonic projection $\tilde{\chi}_S$. We will be particularly interested in the coefficient of the monomial $\prod_{i=1}^d x_i$ in $\tilde{\chi}_S$. According to Corollary 10.4, the coefficient of $\prod_{i=1}^d x_i$ in $\tilde{\chi}_S$ is independent of p , as long as $d \leq pn \leq n - d$. This suggests considering the harmonic projection of $\prod_{i=1}^d x_i$ on the slice $\binom{[n]}{d}$. On that slice, $\prod_{i=1}^d x_i$ is the indicator function of the point $\mathbf{p} = \{1, \dots, d\}$, and we denote its harmonic projection on $\binom{[n]}{d}$ by $\mathbf{1}_{\mathbf{p}}$. The coefficient of $\prod_{i=1}^d x_i$ in $\mathbf{1}_{\mathbf{p}}$ clearly equals its coefficient in $\mathbf{1}_{\mathbf{p}}^{\perp d}$. Since all monomials in $\mathbf{1}_{\mathbf{p}}^{\perp d}$ have degree d , this coefficient is simply $\mathbf{1}_{\mathbf{p}}^{\perp d}(\mathbf{p})$.

Let E_d be the operator projecting a function ϕ on $\binom{[n]}{d}$ to $\phi^{\perp d}$. The trace of E_d is clearly the dimension of the space of harmonic multilinear polynomials which are homogeneous of degree d , which is $\binom{n}{d} - \binom{n}{d-1}$ by Corollary 3.9. Symmetry shows that the diagonal elements of E_d are all equal to

$$\frac{\binom{n}{d} - \binom{n}{d-1}}{\binom{n}{d}} = 1 - \frac{d}{n-d+1}.$$

This is exactly the value of $\mathbf{1}_{\mathbf{p}}^{\perp d}(\mathbf{p})$. We conclude that the coefficient of $\prod_{i=1}^d x_i$ in $\tilde{\chi}_S$ is $1 - \frac{d}{n-d+1}$, and so

$$\chi_S - \tilde{\chi}_S = \frac{d}{n-d+1} \chi_S + \dots,$$

where the dots hide a linear combination of other monomials.

If we substitute $x_i := \frac{x_i - p}{\sqrt{p(1-p)}}$ in $\tilde{\chi}_S$ then we get a harmonic multilinear polynomial of degree d which agrees with ω_S on the slice $\binom{[n]}{pn}$, and so equals $\tilde{\omega}_S$. Substituting this in the preceding formula, we deduce that

$$\omega_S - \tilde{\omega}_S = \frac{d}{n-d+1} \omega_S + \psi,$$

where ψ is orthogonal to ω_S with respect to μ_p .

Since $\tilde{\omega}_S$ agrees with ω_S on the slice, we can compute $\|\tilde{\omega}_S\|_{\nu_{pn}}^2$ explicitly. Simple estimates show that $\|\tilde{\omega}_S\|_{\nu_{pn}}^2 \leq 1 + O(\frac{d^2}{p(1-p)n})$ (see [18, Lemma 4.2] for the details). Corollary 3.11 and Lemma 3.12 imply the same bound on $\|\tilde{\omega}_S\|_{\mu_p}^2$. It follows that

$$\begin{aligned} \|\omega_S - \tilde{\omega}_S\|_{\mu_p}^2 &= \frac{d}{n-d+1} + \|\psi\|_{\mu_p}^2 = \frac{d}{n-d+1} + \|\tilde{\omega}_S\|_{\mu_p}^2 - \left(1 - \frac{d}{n-d+1}\right) \leq \\ &= \frac{2d}{n-d+1} + O\left(\frac{d^2}{p(1-p)n}\right) = O\left(\frac{d^2}{p(1-p)n}\right). \end{aligned}$$

This completes the proof. \square

Armed with this result, we can prove Theorem 5.6.

Proof. Let $\check{f} = f^{\leq d}$ (that is, $\check{f} = \sum_{i \leq d} f^{=i}$), so that $\mathbb{E}_{\nu_{pn}}[(|\check{f}| - 1)^2] \leq \mathbb{E}_{\nu_{pn}}[(\check{f} - f)^2] = \epsilon$. Notice that \check{f} is a harmonic multilinear polynomial of degree at most d . We would like to apply Theorem 5.3 to \check{f} , which is possible if $n \geq d^2/(p(1-p)K^2)$. If this is not the case then $n \leq (p(1-p))^{-O(d)}$ (using $1/p(1-p) \geq 2$), and so the theorem is trivial (we can take $h = f$).

Assume therefore that n is large enough. Theorem 5.3 implies that

$$\mathbb{E}_{\mu_p}[(|\check{f}| - 1)^2] \leq \epsilon + \underbrace{O\left(\sqrt[4]{\frac{d}{\sqrt{p(1-p)n}}}(p(1-p))^{-O(d)}}\right)}_{\epsilon_1}.$$

Proposition 5.5 implies that there exists a cube-Boolean function g on a set J of $M = O((p(1-p))^{-O(d)})$ coordinates such that $\epsilon_2 \triangleq \mathbb{E}_{\mu_p}[(\check{f} - g)^2] = O((p(1-p))^{-O(d)}\epsilon_1)$. The function g is also slice-Boolean, but it is not necessarily harmonic. Let \tilde{g} be its harmonic projection on $\binom{[n]}{pn}$; this will be our choice for h . Note that \tilde{g} also depends only on the coordinates in J , and in particular it has degree at most M (in fact, Lemma 3.18 implies that $\deg \tilde{g} \leq \deg g$). Invoking Theorem 5.7, we see that $\|g - \tilde{g}\|_{\mu_p}^2 = O(\frac{M^2 2^M}{p(1-p)n})$, and so

$$\epsilon_3 \triangleq \|\check{f} - \tilde{g}\|_{\mu_p}^2 = O\left(\frac{M^2 2^M}{p(1-p)n} + \epsilon_2\right).$$

Corollary 3.11 and Lemma 3.12 imply that $\|\check{f} - \tilde{g}\|_{\nu_{pn}}^2 = O(\epsilon_3)$, using the fact that $\deg(\check{f} - \tilde{g}) \leq M$. The L_2 triangle inequality shows that

$$\|f - \tilde{g}\|_{\nu_{pn}}^2 \leq O(\|f - \check{f}\|_{\nu_{pn}}^2 + \epsilon_3) = O((p(1-p))^{-O(d)}\epsilon) + O\left(\sqrt[4]{\frac{d}{\sqrt{p(1-p)n}}}(p(1-p))^{-O(d)}\right) + O\left(\frac{M^2 2^M}{p(1-p)n}\right). \quad \square$$

The proof of [18, Theorem 7.5] contains an additional argument guaranteeing that $\deg h \leq d$. The same argument can be applied here. The idea is that there are finitely many Boolean functions on $(p(1-p))^{-O(d)}$ coordinates, and each of them of degree larger than d has (as $n \rightarrow \infty$) constant distance from all Boolean functions of degree at most d . Hence if ϵ is small enough, g must have degree at most d . We refer the reader to [18] for the complete details.

6 Multilinear functions

Theorem 4.6 only applies to harmonic multilinear polynomials. The harmonicity condition is crucial here. Indeed, a polynomial such as $\sum_{i=1}^n x_i$ behaves very differently on the Boolean cube (where it has a non-trivial

distribution) and on a slice (where it is constant). Nevertheless, we are able to recover similar theorems by looking at several slices at once. Our first invariance result, which we call *norm invariance*, states that if two low-degree multilinear polynomials f, g are close in L_2 norm for enough “well-separated” slices, then they are close in L_2 norm over the entire Boolean cube. Our second invariance result, which we call *interpolation invariance*, gives a recipe for constructing the distribution of a low-degree multilinear polynomial f from its profile on several “well-separated” slices, where by *profile* we mean the distribution of f on several *coupled* slices, just as in Section 4.

Our main technical tool is a theorem of Blekherman [5], which states that any degree d multilinear polynomial P corresponds uniquely to a degree d polynomial Q in the variables x_1, \dots, x_n, S such that

- (a) $P(x_1, \dots, x_n) = Q(x_1, \dots, x_n, x_1 + \dots + x_n)$ for any point in the Boolean cube $\{0, 1\}^n$.
- (b) For each e , the coefficient of S^e is a *harmonic* multilinear polynomial (of degree at most $d - e$).

This theorem allows us to reduce the analysis of arbitrary multilinear polynomials to that of harmonic ones.

We state Blekherman’s theorem in Section 6.1. After preparatory work in Section 6.2, we prove our norm invariance theorems in Section 6.3, and our interpolation invariance theorems in Section 6.4.

There are several principal results in this section, which we now highlight.

Norm invariance (§6.3) The main results are Theorem 6.8, which bounds the L_2 norm of a low-degree multilinear polynomial on the Boolean cube with respect to its L_2 norm on several well-separated slices, and Theorem 6.9, which goes in the other direction, bounding the L_2 norm of a low-degree multilinear polynomial on a slice in terms of its L_2 norm on the Boolean cube and the centrality of the slice.

Both results are combined in Corollary 6.10, which states (roughly) that a low-degree multilinear polynomial has small L_2 norm on the Boolean cube if and only if it has small L_2 norm on several well-separated slices. Another conclusion, Corollary 6.11, states that two low-degree multilinear polynomials are close in L_2 norm on the Boolean cube if and only if they are close in L_2 norm on several well-separated slices.

Interpolation invariance (§6.4) The main results are Theorem 6.13 and Corollary 6.14, which show how to estimate the distribution of a low-degree multilinear polynomial on a given slice given its distribution on several coupled well-separated slices, and Theorem 6.17 and Corollary 6.18, which similarly show how to estimate the distribution of a low-degree multilinear polynomial on the entire Boolean cube given its distribution on several coupled well-separated slices.

These results imply that if two low-degree multilinear polynomials have a similar distribution on several coupled well-separated slices then they have a similar distribution on other slices and on the entire Boolean cube, as we show in Corollary 6.16 and in Corollary 6.19, respectively.

6.1 Blekherman’s theorem

Our starting point is a theorem of Blekherman [5] quoted in Lee et al. [25]. For completeness, we prove this theorem in Section 10.1.

Theorem 6.1. *Let f be a multilinear polynomial over x_1, \dots, x_n of degree $d \leq n/2$, and define $S := \sum_{i=1}^n x_i$. There exist harmonic multilinear polynomials f_0, \dots, f_d over x_1, \dots, x_n , where $\deg f_i \leq d - i$, such that*

$$f(x_1, \dots, x_n) \equiv \sum_{i=0}^d f_i(x_1, \dots, x_n) S^i \pmod{I}, \text{ where } I = \langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle,$$

or equivalently, both sides agree on every point of the Boolean cube $\{0, 1\}^n$. Moreover, this representation is unique.

For our purposes, it will be better to consider f as a polynomial in $(S - np)/\sqrt{np(1 - p)}$ rather than in S .

Corollary 6.2. Let f be a multilinear polynomial over x_1, \dots, x_n of degree $d \leq n/2$, let $p \in (0, 1)$, and define $\tilde{S} := (\sum_{i=1}^n x_i - np) / \sqrt{np(1-p)}$. There exist harmonic multilinear polynomials f_0, \dots, f_d over x_1, \dots, x_n , where $\deg f_i \leq d - i$, such that

$$f(x_1, \dots, x_n) \equiv \sum_{i=0}^d f_i(x_1, \dots, x_n) \tilde{S}^i \pmod{I}, \text{ where } I = \langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle,$$

or equivalently, both sides agree on every point of the Boolean cube $\{0, 1\}^n$. Moreover, this representation is unique.

Proof. Follows from the fact that S and \tilde{S} are affine shifts of one another. \square

We call the representation of f in Corollary 6.2 its *Blekherman expansion* with respect to p , and f_0, \dots, f_d its *Blekherman coefficients* with respect to p . If we substitute $\tilde{S} = \sigma$ in the Blekherman expansion then we get a harmonic multilinear polynomial of degree at most d which agrees with f on the slice $\binom{[n]}{k}$, where $k = np + \sigma\sqrt{np(1-p)}$. We denote this function by $f[\sigma]$. Note that this notation depends on p .

The Blekherman expansion is *linear* in the sense that if $h = \alpha f + \beta g$ then $h_e = \alpha f_e + \beta g_e$ and $h[\sigma] = \alpha f[\sigma] + \beta g[\sigma]$. This immediately follows from its uniqueness and the fact that harmonic functions are closed under taking linear combinations.

6.2 Vandermonde interpolation

Our arguments will involve extracting the Blekherman coefficients f_i given $f[\sigma]$ for various values of σ . We will consider the simple setting in which we are given $d+1$ values of σ , and in that case the problem translates to solving a system of linear equations whose coefficient matrix is a Vandermonde matrix. The quality of the reconstruction will depend on the magnitude of the entries in the inverse matrix, which we estimate using a result of Turner [41].

Proposition 6.3 (Turner [41]). Let ξ_1, \dots, ξ_n be arbitrary real numbers, and consider the $n \times n$ Vandermonde matrix V given by $V_{ij} = \xi_i^{j-1}$, where $1 \leq i, j \leq n$. The inverse of V is given by $V^{-1} = UL$, where

$$L_{ij} = \begin{cases} 0 & \text{if } i < j, \\ \prod_{\substack{k=1 \\ k \neq j}}^i \frac{1}{\xi_j - \xi_k} & \text{otherwise,} \end{cases} \quad U_{ij} = \begin{cases} 0 & \text{if } i > j, \\ 1 & \text{if } i = j, \\ U_{(i-1)(j-1)} - \xi_{j-1} U_{i(j-1)} & \text{if } i < j, \text{ where } U_{0(j-1)} = 0. \end{cases}$$

This proposition implies the following interpolation result.

Theorem 6.4. Suppose that f is a multilinear polynomial over x_1, \dots, x_n of degree $d \leq n/2$ with Blekherman coefficients f_0, \dots, f_d with respect to some $p \in (0, 1)$. Let $\sigma_1, \dots, \sigma_{d+1}$ be $d+1$ distinct values, and define

$$\eta = \min(1, \min_{i \neq j} |\sigma_i - \sigma_j|), \quad M = \max(1, \max_i |\sigma_i|).$$

For $0 \leq e \leq d$ and $1 \leq i \leq d+1$ there exist coefficients c_{ei} of magnitude $|c_{ei}| \leq (4M/\eta)^d$ such that for all $0 \leq e \leq d$,

$$f_e = \sum_{i=1}^{d+1} c_{ei} f[\sigma_i].$$

Proof. Let V be the Vandermonde matrix for $\sigma_1, \dots, \sigma_{d+1}$, so that

$$\begin{bmatrix} f[\sigma_1] \\ \vdots \\ f[\sigma_{d+1}] \end{bmatrix} = V \begin{bmatrix} f_0 \\ \vdots \\ f_d \end{bmatrix}.$$

Inverting V , this shows that

$$V^{-1} \begin{bmatrix} f[\sigma_1] \\ \vdots \\ f[\sigma_{d+1}] \end{bmatrix} = \begin{bmatrix} f_0 \\ \vdots \\ f_d \end{bmatrix}.$$

We can thus take $c_{ei} = V_{(e+1)i}^{-1}$. It remains to bound the magnitude of the entries of V^{-1} .

Let L, U be the matrices in Proposition 6.3. The formula for L implies that all of its entries have magnitude at most $(1/\eta)^d$. As for U , we will prove by induction that when $i \leq j$, $|U_{ij}| \leq \binom{j-1}{i-1} M^{j-i}$. This is clearly true when $i = j$. When $i = 1$, $|U_{1j}| = \prod_{k=1}^{j-1} |\sigma_k| \leq M^{j-1}$. The inductive step follows from

$$|U_{ij}| \leq |U_{(i-1)(j-1)}| + M|U_{i(j-1)}| \leq \binom{j-2}{i-2} M^{j-i} + M \cdot \binom{j-2}{i-1} M^{j-1-i} = \binom{j-1}{i-1} M^{j-i}.$$

It follows that all entries of U are bounded in magnitude by $2^d M^d \leq \frac{4^d}{d+1} \cdot M^d$. The theorem follows. \square

6.3 Norm invariance

We are now ready to prove our norm invariance principle. Our argument will require a few auxiliary results. We start with an estimate on the central moments of binomial distributions.

Lemma 6.5. *Let $S \sim \text{B}(n, p)$ and $\tilde{S} = \frac{S - np}{\sqrt{np(1-p)}}$. For all $d \geq 0$,*

$$\mathbb{E}[\tilde{S}^{2d}] \leq \frac{2d!}{(2p(1-p))^d}.$$

Proof. Hoeffding's inequality states that

$$\Pr[|S - np| \geq t] \leq 2e^{-2t^2/n} \implies \Pr[|\tilde{S}| \geq t] \leq 2e^{-2p(1-p)t^2}.$$

Plugging this in the general formula $\mathbb{E}[|X|] = \int_0^\infty \Pr[|X| \geq t] dt$, we get

$$\mathbb{E}[\tilde{S}^{2d}] = \int_0^\infty \Pr[\tilde{S}^{2d} \geq t] dt \leq 2 \int_0^\infty e^{-2p(1-p)t^{1/d}} dt.$$

Substituting $s = 2p(1-p)t^{1/d}$, we have $t = (s/2p(1-p))^d$ and so $dt/ds = d(s/2p(1-p))^{d-1}/2p(1-p)$, implying

$$\mathbb{E}[\tilde{S}^{2d}] \leq 2 \cdot (2p(1-p))^{-d} \int_0^\infty e^{-s} ds^{d-1} ds = 2 \cdot (2p(1-p))^{-d} d!,$$

using the classical integral formula for the Γ function:

$$\int_0^\infty e^{-s} s^{d-1} ds = \Gamma(d) = (d-1)!. \quad \square$$

We comment that for every fixed p and d , as $n \rightarrow \infty$, $\mathbb{E}[\tilde{S}^{2d}]$ converges to $\mathbb{E}[\text{N}(0, 1)^{2d}] = \frac{(2d)!}{2^d d!}$.

We also need an anti-concentration result for binomial distributions, which follows from the Berry–Esseen theorem.

Lemma 6.6. *Let $S \sim \text{B}(n, p)$ and $\tilde{S} = \frac{S - np}{\sqrt{np(1-p)}}$. For all $a < b$,*

$$\Pr[\tilde{S} \in (a, b)] \geq \frac{b-a}{\sqrt{2\pi}} e^{-\max(a^2, b^2)/2} - \frac{1}{\sqrt{np(1-p)}}.$$

Proof. The Berry–Esseen theorem states that the cumulative distribution functions of \tilde{S} and of the standard normal distribution $N(0, 1)$ differ by at most $C\rho/\sigma^3\sqrt{n}$, where $C < 1$ is an absolute constant, $\rho = \mathbb{E}[|\text{Ber}(p) - p|^3] = p(1-p)(p^2 + (1-p)^2)$, and $\sigma^2 = \mathbb{V}[\text{Ber}(p)] = p(1-p)$. The result follows from $\rho \leq p(1-p)/2$. \square

We are now ready to tackle norm invariance. We start by bounding the norm of f on the Boolean cube given its norm on several well-separated slices.

Definition 6.7. Fix $p \in (0, 1)$ and n . A set $0 \leq k_1, \dots, k_r \leq n$ is said to be an (η, M) -system, for $\eta \leq 1$ and $M \geq 1$, if the following two conditions hold for $\sigma_i = \frac{k_i - np}{\sqrt{np(1-p)}}$:

(a) For every $i \neq j$, $|\sigma_i - \sigma_j| \geq \eta$.

(b) For every i , $|\sigma_i| \leq M$.

Theorem 6.8. Let $p \in (0, 1)$, and let f be a multilinear polynomial over x_1, \dots, x_n of degree $d \leq \sqrt{np(1-p)}$. Let k_1, \dots, k_{d+1} be an (η, M) -system, where $M \leq \sqrt{np(1-p)}/2$, and suppose that $\|f\|_{\nu_{k_i}} \leq 1$ for all $1 \leq i \leq d+1$. Then

$$\|f\|_{\mu_p} = O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)}.$$

Moreover, the Blekherman coefficients f_0, \dots, f_d of f with respect to p satisfy, for $0 \leq e \leq d$,

$$\|f_e\|_{\mu_p} \leq O(M/\eta)^d.$$

Proof. If $d = 0$ then f is constant, so we can assume that $d \geq 1$.

We are given that for each i , $\|f[\sigma_i]\|_{\nu_{k_i}} \leq 1$. Since $f[\sigma_i]$ is a harmonic multilinear polynomial of degree at most d , Corollary 3.11 and Lemma 3.12 show that for $q = k_i/n$,

$$\|f[\sigma_i]\|_{\mu_p}^2 \leq \frac{(2p(1-p))^d}{(2q(1-q))^d} \left(1 - O\left(\frac{d^2}{q(1-q)n}\right)\right)^{-1} = O(1) \cdot \left(\frac{p(1-p)}{q(1-q)}\right)^d.$$

Since $M \leq \sqrt{np(1-p)}/2$, we have

$$|p - q| \leq \frac{M\sqrt{np(1-p)}}{n} \leq \frac{p(1-p)}{2},$$

which implies that

$$\frac{p}{q} \leq \frac{p}{p - p(1-p)/2} = \frac{2}{1+p} \leq 2.$$

Similarly, $(1-p)/(1-q) \leq 2$. This shows that

$$\|f[\sigma_i]\|_{\mu_p}^2 \leq O(1)^d.$$

Theorem 6.4 shows that there exist coefficients c_{ei} of magnitude at most $(4M/\eta)^d$ such that for each e , $f_e = \sum_i c_{ei} f[\sigma_i]$. It follows that for each e ,

$$\|f_e\|_{\mu_p} \leq \sum_{i=1}^{d+1} |c_{ei}| \|f[\sigma_i]\|_{\mu_p} \leq (d+1) \cdot O(M/\eta)^d \cdot O(1)^d = O(M/\eta)^d.$$

Since $f = \sum_{e=0}^d \tilde{S}^e f_e$, the Cauchy–Schwarz inequality implies that

$$\|f\|_{\mu_p} \leq \sum_{e=0}^d \|\tilde{S}^e f_e\|_{\mu_p} \leq \sum_{e=0}^d \sqrt{\|\tilde{S}^{2e}\|_{\mu_p}} \sqrt{\|f_e^2\|_{\mu_p}}.$$

Lemma 6.5 shows that

$$\sqrt{\|\tilde{S}^{2e}\|_{\mu_p}} = \sqrt[4]{\mathbb{E}_{\mu_p}[\tilde{S}^{4e}]} = O\left(\frac{d}{p(1-p)}\right)^{d/2}.$$

Hypercontractivity, in the form of Lemma 5.2, implies that with respect to μ_p ,

$$\sqrt{\|f_e^2\|} = \|f_e\|_4 \leq O(p(1-p))^{-O(\epsilon)} \|f_e\|_2 = \left(\frac{M}{p(1-p)\eta}\right)^{O(d)}.$$

In total, we deduce that

$$\|f\|_{\mu_p} \leq (d+1) \cdot O\left(\frac{d}{p(1-p)}\right)^d \cdot \left(\frac{M}{p(1-p)\eta}\right)^{O(d)} = \left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)}. \quad \square$$

We can also go in the other direction. In the statement of Theorem 6.9 and similar results below, we allow big O constants to depend on the fixed value of p (and below, on the fixed value of other parameters).

Theorem 6.9. *Fix $p \in (0, 1)$, and let f be a multilinear polynomial over x_1, \dots, x_n of degree $d \leq \sqrt{\log[np(1-p)/30]}$ satisfying $\|f\|_{\mu_p} \leq 1$. If $k = np + \sqrt{np(1-p)} \cdot \sigma$ then*

$$\|f\|_{\nu_k} \leq e^{O(d^2)} (1 + |\sigma|)^d e^{d|\sigma|/\sqrt{np(1-p)}}.$$

Proof. If $d = 0$ then f is constant, so we can assume that $d \geq 1$.

For $0 \leq i \leq d$, let $a_i = i - 1/2$ and $b_i = i$. Lemma 6.6 shows that

$$\Pr[\tilde{S} \in (a_i, b_i)] \geq \frac{1}{2\sqrt{2\pi}} e^{-d^2/2} - \frac{1}{\sqrt{np(1-p)}}.$$

The first summand is at least

$$\frac{1}{2\sqrt{2\pi}} e^{-\log[np(1-p)/30]/2} = \frac{1}{2\sqrt{2\pi}} \cdot \frac{1}{\sqrt{np(1-p)/30}} > \frac{1.09}{\sqrt{np(1-p)}},$$

and so

$$\Pr[\tilde{S} \in (a_i, b_i)] = \Omega(e^{-d^2/2}).$$

In particular, the norm of f restricted to $\tilde{S} \in (a_i, b_i)$ is $O(e^{d^2/2})$, and so there must exist $k_i = np + \sqrt{np(1-p)} \cdot \sigma_i$ such that $\sigma_i \in (a_i, b_i)$ and $\|f\|_{\nu_{k_i}} = O(e^{d^2/2})$.

The resulting system k_0, \dots, k_d is a $(1/2, d)$ -system, and so Theorem 6.8 shows that the Blekherman coefficients f_0, \dots, f_d satisfy

$$\|f_e\|_{\mu_p} \leq O(e^{d^2/2}) \cdot O(d)^d = e^{O(d^2)}.$$

Substituting a given value of σ , we deduce that

$$\|f[\sigma]\|_{\mu_p} \leq \sum_{e=0}^d |\sigma|^e \|f_e\|_{\mu_p} \leq (1 + |\sigma|)^d e^{O(d^2)}.$$

Corollary 3.11 and Lemma 3.12 show that for $q = k/n$,

$$\|f[\sigma]\|_{\nu_k} \leq (1 + |\sigma|)^d e^{O(d^2)} \frac{(2q(1-q))^{d/2}}{(2p(1-p))^{d/2}} \left(1 + O\left(\frac{d^2}{p(1-p)n}\right)\right)^{1/2} = (1 + |\sigma|)^d e^{O(d^2)} \cdot \left(\frac{q(1-q)}{p(1-p)}\right)^{d/2}.$$

In order to estimate the final factor, note that $q = p + \sqrt{\frac{p(1-p)}{n}} \sigma$ and so

$$\frac{q}{p} = 1 + \sqrt{\frac{1-p}{p}} \frac{\sigma}{\sqrt{n}} \leq e^{\sqrt{\frac{1-p}{p}} \frac{\sigma}{\sqrt{n}}} \leq e^{\sigma/\sqrt{np(1-p)}}.$$

Similarly $(1 - q)/(1 - p) \leq e^{-\sigma/\sqrt{np(1-p)}}$. Therefore

$$\left(\frac{q(1-q)}{p(1-p)}\right)^{d/2} \leq e^{d|\sigma|/\sqrt{np(1-p)}},$$

and the theorem follows. \square

We can combine both results to obtain the following clean corollary.

Corollary 6.10. *Fix $p \in (0, 1)$ and $d \geq 1$, and suppose that \mathcal{F} is a collection of multilinear polynomials of degree at most d on at least $\frac{30}{p(1-p)}e^{d^2}$ variables (different functions could depend on a different number of variables). The following three conditions are equivalent:*

- (a) *There exists a constant C_1 such that $\|f\|_{\mu_p} \leq C_1$ for all $f \in \mathcal{F}$.*
- (b) *There exists a constant C_2 such that $\|f\|_{\nu_k} \leq (C_2(1 + \frac{|k-np|}{\sqrt{n}}))^d$ for all $f \in \mathcal{F}$ over x_1, \dots, x_n and for all $0 \leq k \leq n$.*
- (c) *There exists a constant C_3 such that $\|f\|_{\nu_k} \leq C_3$ for all $f \in \mathcal{F}$ over x_1, \dots, x_n and for $k = \lfloor np + \sqrt{np(1-p)} \cdot \sigma \rfloor$ for $\sigma \in \{0, \dots, d\}$.*

Proof. Suppose first that condition (a) holds, and let $\sigma = \frac{k-np}{\sqrt{np(1-p)}}$, so that $|\sigma| \leq \sqrt{n/p(1-p)}$. Then Theorem 6.9 shows that

$$\|f\|_{\nu_k} \leq C_1 \cdot e^{O(d^2)} \cdot (1 + |\sigma|)^d \cdot e^{d/p(1-p)} \leq (C_1 e^{O(d^2)} (1 + |\sigma|))^d,$$

which implies condition (b).

Suppose next that condition (b) holds. In particular, for $\sigma_1, \dots, \sigma_{d+1} = 0, \dots, d$ and $k_i = np + \sqrt{np(1-p)} \cdot \sigma_i$ it is the case that $\|f\|_{\nu_{k_i}} \leq (C_2(1 + d\sqrt{p(1-p)}))^d$, which implies condition (c) with $C_3 = (C_2(1 + d\sqrt{p(1-p)}))^d$.

Finally, suppose that condition (c) holds. Since $0, 1, \dots, d$ is a $(1, d)$ -system, Theorem 6.8 shows that

$$\|f\|_{\mu_p} = O\left(\frac{d^2}{p(1-p)}\right)^{O(d)} C_3,$$

which depends only on p and d , implying condition (a). \square

Here is a different interpretation of these results.

Corollary 6.11. *Fix $p \in (0, 1)$ and $d \geq 1$. There is a constant $C > 0$ such that the following implications hold for any two multilinear polynomials of degree at most d on $n \geq \frac{30}{p(1-p)}e^{d^2}$ variables:*

- (a) *If $\|f - g\|_{\mu_p} = \epsilon$ then for all $0 \leq k \leq n$, $\|f - g\|_{\nu_k} \leq (C(1 + \frac{|k-np|}{\sqrt{n}}))^d \epsilon$.*
- (b) *If k_1, \dots, k_{d+1} is an (η, M) -system and $\|f - g\|_{\nu_{k_i}} \leq \epsilon$ for $1 \leq i \leq d+1$ then $\|f - g\|_{\mu_p} \leq (CM/\eta)^{O(d)} \epsilon$.*

Proof. The first statement follows directly from Theorem 6.9 (using $|\sigma| \leq \sqrt{n/p(1-p)}$), and the second statement follows directly from Theorem 6.8. In both cases, we apply the theorems to $f - g$. \square

6.4 Interpolation invariance

We move on to the interpolation invariance principle. Definition 4.1 describes a coupling $\mathbf{X}(0), \dots, \mathbf{X}(n)$ of the distributions ν_0, \dots, ν_n , which we will use to define the *profile* of a function.

Definition 6.12. Fix parameters $p \in (0, 1)$, $d, n \geq 0$, and let k_1, \dots, k_{d+1} be an (η, M) -system. The *profile* of a multilinear function f of degree at most d with respect to the system k_1, \dots, k_{d+1} is the joint distribution of $f(\mathbf{X}(k_1)), \dots, f(\mathbf{X}(k_{d+1}))$, which we denote by $\mathbf{f}_1, \dots, \mathbf{f}_{d+1}$.

The profile of a function f allows us to recover its distribution on arbitrary slices, as reflected by Lipschitz functions.

Theorem 6.13. *There exists a constant $K > 0$ such that the following holds. Let $p \in (0, 1)$, let f be a multilinear polynomial on x_1, \dots, x_n of degree $1 \leq d \leq \sqrt{Knp(1-p)}$, let k_1, \dots, k_{d+1} be an (η, M) -system for $M \leq \sqrt{np(1-p)}/2$, and let $\mathbf{f}_1, \dots, \mathbf{f}_{d+1}$ be the profile of f with respect to this system. Suppose that $\|f\|_{\nu_{k_i}} \leq 1$ for $1 \leq i \leq d+1$. For every slice $k = np + \sqrt{np(1-p)} \cdot \sigma$ such that $L := \max_i |\sigma - \sigma_i| \leq \frac{\sqrt{np(1-p)}}{8d}$, and for any Lipschitz function φ ,*

$$\left| \mathbb{E}_{\nu_k}[\varphi(f)] - \mathbb{E}[\varphi(\mathbf{f}[k])] \right| = O\left((1 + |\sigma|) \frac{M}{\eta} \right)^d \sqrt{\frac{L}{\sqrt{np(1-p)}}}, \quad \text{where } \mathbf{f}[k] = \sum_{i=1}^{d+1} \gamma_{ki} \mathbf{f}_i,$$

for some constants γ_{ki} depending on p, n, d satisfying

$$|\gamma_{ki}| = O\left((1 + |\sigma|) \frac{M}{\eta} \right)^d.$$

Proof. Let $p_i = k_i/n = p + \sqrt{\frac{p(1-p)}{n}} \sigma_i$ (where $\sigma_i = \frac{k_i - np}{\sqrt{np(1-p)}}$), and let $q = k/n = p + \sqrt{\frac{p(1-p)}{n}} \sigma$. The condition on L guarantees that $|p_i - q| = \sqrt{\frac{p(1-p)}{n}} |\sigma_i - \sigma| \leq \frac{p(1-p)}{8d}$. The condition on M guarantees that $|p_i - p| \leq \sqrt{\frac{p(1-p)}{n}} M \leq p(1-p)/2$, and so $\frac{1}{4}p(1-p) \leq p_i(1-p_i) \leq \frac{9}{4}p(1-p)$. In particular, $|p_i - q| \leq \frac{p(1-p)}{8d} \leq \frac{p_i(1-p_i)}{2d}$. Lemma 4.4 (applied with $p := p_i$ and $q := q$) shows that

$$\mathbb{E}[|f(\mathbf{X}(k_i)) - f(\mathbf{X}(k))|]^2 \leq \mathbb{E}[(f(\mathbf{X}(k_i)) - f(\mathbf{X}(k)))^2] = O\left(\frac{dL}{\sqrt{np(1-p)}} \right).$$

Let f_0, \dots, f_d be the Blekherman coefficients of f , and let c_{ei} be the coefficients given by Theorem 6.4, so that $|c_{ei}| \leq (4M/\eta)^d$. Theorem 6.4 shows that for $0 \leq e \leq d$,

$$\mathbb{E}\left[\left| f_e(\mathbf{X}(k)) - \sum_{i=1}^{d+1} c_{ei} f(\mathbf{X}(k_i)) \right| \right] \leq \sum_{i=1}^{d+1} |c_{ei}| \mathbb{E}[|f(\mathbf{X}(k_i)) - f(\mathbf{X}(k))|] \leq (d+1) \cdot \left(\frac{4M}{\eta} \right)^d \cdot O\left(\sqrt{\frac{dL}{\sqrt{np(1-p)}}} \right).$$

Since $f[\sigma] = \sum_{e=0}^d \sigma^e f_e$, we conclude that

$$\mathbb{E}\left[\left| f(\mathbf{X}(k)) - \sum_{e=0}^d \sigma^e \sum_{i=1}^{d+1} c_{ei} f(\mathbf{X}(k_i)) \right| \right] \leq (d+1)^2 \cdot (1 + |\sigma|^d) \cdot \left(\frac{4M}{\eta} \right)^d \cdot O\left(\sqrt{\frac{dL}{\sqrt{np(1-p)}}} \right).$$

Rearrangement yields the statement of the theorem, with $\gamma_{ki} = \sum_{e=0}^d \sigma^e c_{ei}$. \square

We immediately obtain a corollary for the Lévy distance.

Corollary 6.14. *Under the setting of Theorem 6.13, the Lévy distance between $f(\nu_k)$ and $\mathbf{f}[k]$ is at most*

$$O\left((1+|\sigma|)\frac{M}{\eta}\right)^{d/2} \sqrt[4]{\frac{L}{\sqrt{np(1-p)}}}.$$

Proof. The proof is identical to the proof of Corollary 4.7. \square

A striking form of this corollary compares two different profiles.

Definition 6.15. Let $X = (X_1, \dots, X_r), Y = (Y_1, \dots, Y_r)$ be two r -dimensional real distributions. The Lévy–Prokhorov distance between these distributions is the infimum value of ϵ such that for all Borel sets A ,

1. $\Pr[X \in A] \leq \Pr[Y \in A^\epsilon] + \epsilon$, and
2. $\Pr[Y \in A] \leq \Pr[X \in A^\epsilon] + \epsilon$,

where A^ϵ consists of all points at distance at most ϵ from A in the L^∞ metric.

In fact, for all our results it suffices to consider *halfspaces* for A , that is, sets of the form $\{(x_1, \dots, x_r) : \sum_{i=1}^r a_i x_i \leq b\}$. If we restrict the definition of Lévy–Prokhorov distance in this way, then it coincides with the usual Lévy distance in the one-dimensional case.

Corollary 6.16. *Under the setting of Theorem 6.13 for two functions f, g , if the Lévy–Prokhorov distance between the profiles of f and g is ϵ , then the Lévy distance between $f(\nu_k)$ and $g(\nu_k)$ is at most*

$$O\left((1+|\sigma|)\frac{M}{\eta}\right)^{d/2} \sqrt[4]{\frac{L}{\sqrt{np(1-p)}}} + O\left((1+|\sigma|)\frac{M}{\eta}\right)^d \epsilon.$$

Proof. We start by bounding the Lévy distance between $\mathbf{f}[k]$ and $\mathbf{g}[k]$. Given t , define the halfspace A_t by

$$A_t = \{(x_1, \dots, x_{d+1}) : \sum_{i=1}^{d+1} \gamma_{ki} x_i \leq t\},$$

and notice that $A_t^\epsilon \subseteq A_{t+B\epsilon}$, where

$$B = \sum_{i=1}^{d+1} |\gamma_{ki}| = O\left((1+|\sigma|)\frac{M}{\eta}\right)^d.$$

Since the Lévy–Prokhorov distance between the profiles of f and g is at most ϵ , it follows that

$$\begin{aligned} \Pr[\mathbf{f}[k] \leq t] &= \Pr[(\mathbf{f}_1, \dots, \mathbf{f}_{d+1}) \in A_t] \leq \\ &\Pr[(\mathbf{g}_1, \dots, \mathbf{g}_{d+1}) \in A_t^\epsilon] + \epsilon \leq \Pr[(\mathbf{g}_1, \dots, \mathbf{g}_{d+1}) \in A_{t+B\epsilon}] + \epsilon = \Pr[\mathbf{g}[k] \leq t + B\epsilon] + \epsilon. \end{aligned}$$

This shows that the Lévy distance between $\mathbf{f}[k]$ and $\mathbf{g}[k]$ is at most $B\epsilon$.

The result now follows from Corollary 6.14 and the triangle inequality for the Lévy distance. \square

By combining different slices, we can obtain a similar result for μ_p .

Theorem 6.17. *There exists a constant $K > 0$ such that the following holds. Let $p \in (0, 1)$, let f be a multilinear polynomial on x_1, \dots, x_n of degree $1 \leq d \leq K \sqrt{np(1-p)/\log[np(1-p)]}$, let k_1, \dots, k_{d+1} be an (η, M) -system for $M \leq \sqrt{np(1-p)}/(9d)$, and let $\mathbf{f}_1, \dots, \mathbf{f}_{d+1}$ be the profile of f with respect to this system. Suppose that $\|f\|_{\nu_{k_i}} \leq 1$ for $1 \leq i \leq d+1$.*

For $s := \sqrt{3 \log[np(1-p)]}$, define a distribution \mathbf{f} as follows: choose $\sigma \sim \frac{B(n,p)-np}{\sqrt{np(1-p)}}$ conditioned on $|\sigma| \leq s$, and let $\mathbf{f} = \mathbf{f}[\sigma]$, as in Theorem 6.13. Then for any Lipschitz function φ ,

$$|\mathbb{E}_{\mu_p}[\varphi(f)] - \mathbb{E}[\varphi(\mathbf{f})]| = O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)} \frac{1}{\sqrt[4]{np(1-p)}}.$$

Proof. We can assume, without loss of generality, that $\varphi(0) = 0$.

Since $M + s \leq \frac{\sqrt{np(1-p)}}{8d}$ due to the bound on M and our choice of s , Theorem 6.13 implies that

$$\sum_{\substack{k=np+\sqrt{np(1-p)}\cdot\sigma: \\ |\sigma|\leq s}} \Pr[\mathbf{B}(n, p) = k] \mathbb{E}_{\nu_k}[\varphi(f)] - \mathbb{E}[\varphi(\mathbf{f})] = O\left(\frac{M}{\eta}\right)^d \frac{1}{\sqrt[4]{np(1-p)}} \underbrace{\mathbb{E}[(1+|\sigma|)^d \sqrt{M+|\sigma|}]}_C,$$

where $\sigma \sim \frac{\mathbf{B}(n, p) - np}{\sqrt{np(1-p)}}$. When $|\sigma| \leq 1$, $C \leq 2^{d+1}\sqrt{M}$, and when $|\sigma| \geq 1$, $C \leq 2^d|\sigma|^d(\sqrt{M} + \sqrt{|\sigma|}) \leq 2^{d+1}|\sigma|^{2d}\sqrt{M}$. Therefore Lemma 6.5 implies that

$$\mathbb{E}[C] \leq 2^{d+1}\sqrt{M}(1 + \mathbb{E}[|\sigma|^{2d}]) = O\left(\frac{d}{p(1-p)}\right)^d \sqrt{M}.$$

We conclude that

$$\left| \sum_{\substack{k=np+\sqrt{np(1-p)}\cdot\sigma: \\ |\sigma|\leq s}} \Pr[\mathbf{B}(n, p) = k] \mathbb{E}_{\nu_k}[\varphi(f)] - (1 - \epsilon_1) \mathbb{E}[\varphi(\mathbf{f})] \right| = O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^d \frac{1}{\sqrt[4]{np(1-p)}},$$

where $\epsilon_1 := \Pr[|\mathbf{B}(n, p) - np| > \sqrt{np(1-p)} \cdot s]$. Dividing this bound by $1 - \epsilon_1$ and using $1 - \frac{1}{1-\epsilon_1} = O(\epsilon_1)$, we deduce that

$$|\mathbb{E}_{\mu_p}[\varphi(f)] - \mathbb{E}[\varphi(\mathbf{f})]| \leq O(\epsilon_1) \mathbb{E}[\varphi(f)] + \underbrace{\mathbb{E}_{x \sim \mu_p} [\varphi(f) \mathbf{1}_{|\sum_i x_i - np| > s\sqrt{np(1-p)}}]}_{\epsilon_2} + O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^d \frac{1}{\sqrt[4]{np(1-p)}}.$$

As in the proof of Theorem 4.6, Proposition 4.5 implies that $\epsilon_1 \leq 2e^{-s^2/6}$. Theorem 6.8 shows that

$$\|f\|_{\mu_p} = O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)}.$$

This allows us to bound the first error term above, since $\mathbb{E}[\varphi(f)] \leq \mathbb{E}[|f|] \leq \|f\|$. The other error term is bounded in the proof of Theorem 4.6 by $O(e^{-s^2/12})$. Altogether, we obtain

$$|\mathbb{E}_{\nu_k}[\varphi(f)] - \mathbb{E}[\varphi(\mathbf{f})]| \leq O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)} \left[\frac{1}{\sqrt[4]{np(1-p)}} + e^{-s^2/12} \right].$$

Substituting the value for s , we deduce that

$$|\mathbb{E}_{\nu_k}[\varphi(f)] - \mathbb{E}[\varphi(\mathbf{f})]| \leq O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)} \frac{1}{\sqrt[4]{np(1-p)}}. \quad \square$$

Just as before, we can obtain corollaries for the Lévy distance.

Corollary 6.18. *Under the setting of Theorem 6.17, the Lévy distance between $f(\mu_p)$ and \mathbf{f} is*

$$O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)} \frac{1}{\sqrt[8]{np(1-p)}}.$$

Proof. The proof is identical to the proof of Corollary 6.14. □

Corollary 6.19. *Under the setting of Theorem 6.17 for two functions f, g , if the Lévy distance between the profiles of f and g is $\epsilon \leq 1/2$, then the Lévy distance between $f(\mu_p)$ and $g(\mu_p)$ is at most*

$$O\left(\frac{d}{p(1-p)} \cdot \frac{M}{\eta}\right)^{O(d)} \frac{1}{\sqrt[8]{np(1-p)}} + O\left(\sqrt{\frac{\log(1/\epsilon) M}{p(1-p) \eta}}\right)^d \epsilon.$$

Proof. As in the proof of Corollary 6.16, the idea is to bound the Lévy distance between \mathbf{f} and \mathbf{g} . Let σ have the distribution in Theorem 6.17, and let τ be a threshold to be determined. The argument of Corollary 6.16 shows that

$$\Pr[\mathbf{f}[\sigma] \leq t] \leq \Pr[\mathbf{g}[\sigma] \leq t + B_\sigma \epsilon] + \epsilon, \text{ where } B_\sigma = O\left((1 + |\sigma|) \frac{M}{\eta}\right)^d.$$

Therefore

$$\begin{aligned} \Pr[\mathbf{f} \leq t] &= \mathbb{E}_\sigma[\Pr[\mathbf{f}[\sigma] \leq t]] \leq \mathbb{E}_\sigma[\Pr[\mathbf{g}[\sigma] \leq t + B_\sigma \epsilon] \mathbf{1}_{|\sigma| \leq \tau}] + \Pr[|\sigma| > \tau] + \epsilon \\ &\leq \Pr[\mathbf{g} \leq t + B_\tau \epsilon] + \Pr[|\sigma| > \tau] + \epsilon. \end{aligned}$$

As in the proof of Lemma 6.5, Hoeffding's inequality shows that

$$\Pr[|\sigma| > \tau] = O(\Pr[|\tilde{S}| > \tau]) = O(e^{-2p(1-p)\tau^2}),$$

showing that the Lévy distance between \mathbf{f} and \mathbf{g} is at most $\max(B_\tau \epsilon, O(e^{-2p(1-p)\tau^2}) + \epsilon)$. Choosing $\tau = \sqrt{\frac{\log(1/\epsilon)}{2p(1-p)}}$, we conclude that the Lévy distance between \mathbf{f} and \mathbf{g} is at most

$$O\left(\sqrt{\frac{\log(1/\epsilon) M}{p(1-p) \eta}}\right)^d \epsilon.$$

The result now follows from Corollary 6.18 and the triangle inequality for the Lévy distance. \square

Theorem 6.17 states that we can recover the distribution of a low-degree multilinear polynomial on the Boolean cube from its distribution on a few coupled slices. Interpolation in the other direction is not possible: the distribution of a low-degree polynomial function on the Boolean cube doesn't determine its distribution on the various slices. For example, consider the following two functions, for even n :

$$f_1 = \frac{\sum_{i=1}^n x_i - np}{\sqrt{np(1-p)}}, \quad f_2 = \frac{\sum_{i=1}^{n/2} (x_{2i-1} - x_{2i})}{\sqrt{np(1-p)}}.$$

The central limit theorem shows that with respect to μ_p , the distribution of both functions is close to $N(0, 1)$. However, on the slice ν_{pn} , the first function vanishes, while the second function also has a distribution close to $N(0, 1)$, due to Theorem 4.6.

7 Functions depending on few coordinates

In this section we prove an invariance principle for bounded functions depending on $o(n)$ coordinates. In contrast to our work so far, the functions in question need not be harmonic multilinear polynomials. The invariance principle immediately follows from the following bound on total variation distance.

Lemma 7.1. *Let $p(1-p)n \rightarrow \infty$ and $m = o(p(1-p)n)$. Denote the projection of μ_p and ν_{pn} into the first m coordinates by μ'_p and ν'_{pn} . The total variation distance between μ'_p and ν'_{pn} is $o(1)$.*

Proof. Let $k = pn$, and consider the ratio $\rho(\ell)$ between the probability of a set of size ℓ under ν'_k and under μ'_p :

$$\rho(\ell) = \frac{\nu'_k([\ell])}{\mu'_p([\ell])} = \frac{\binom{n-m}{k-\ell} / \binom{n}{k}}{p^\ell (1-p)^{m-\ell}}.$$

To understand the behavior of $\rho(\ell)$, we compute the ratio $\rho(\ell+1)/\rho(\ell)$:

$$\frac{\rho(\ell+1)}{\rho(\ell)} = \frac{1-p}{p} \frac{k-\ell}{n-m-k+\ell+1}.$$

Thus $\rho(\ell+1) > \rho(\ell)$ iff $(1-p)(k-\ell) > p(n-m-k+\ell+1)$ iff $\ell < p(m-1)$. We deduce that the largest value of $\rho(\ell)$ is obtained for $\ell_0 = pm$ (assuming for simplicity that this is indeed an integer).

At the point ℓ_0 we can estimate, using Stirling's approximation,

$$\begin{aligned} \rho(\ell_0) &= \frac{\binom{n-m}{p(n-m)}}{\binom{n}{pn}} p^{-pm} (1-p)^{-(1-p)m} \\ &= \frac{p^{-p(n-m)} (1-p)^{-(1-p)(n-m)} e^{O(1/p(1-p)(n-m))}}{\sqrt{2\pi p(1-p)(n-m)}} \\ &= \frac{p^{-pn} (1-p)^{-(1-p)n} e^{O(1/p(1-p)n)}}{\sqrt{2\pi p(1-p)n}} p^{-pm} (1-p)^{-(1-p)m} \\ &= \sqrt{\frac{n}{n-m}} e^{O(1/p(1-p)(n-m)) - O(1/p(1-p)n)} = 1 + O\left(\frac{m}{p(1-p)n}\right). \end{aligned}$$

Altogether, we deduce that for all ℓ , $\rho(\ell) \leq 1 + O(m/p(1-p)n)$. Therefore the total variation distance is

$$\sum_{\ell: \rho(\ell) > 1} \binom{m}{\ell} (\nu'_k([\ell]) - \mu'_p([\ell])) = \sum_{\ell: \rho(\ell) > 1} \binom{m}{\ell} \mu'_p([\ell]) (\rho(\ell) - 1) = O\left(\frac{m}{p(1-p)n}\right). \quad \square$$

As an immediate corollary, we obtain an invariance principle for bounded functions depending on $o(n)$ coordinates.

Theorem 7.2. *Let f be a function on $\{0, 1\}^n$ depending on $o(p(1-p)n)$ coordinates and satisfying $\|f\|_\infty \leq 1$. As $p(1-p)n \rightarrow \infty$,*

$$|\mathbb{E}_{\nu_{pn}}[f] - \mathbb{E}_{\mu_p}[f]| = o(1).$$

Proof. Suppose that f depends on $m = o(p(1-p)n)$ coordinates. Applying the triangle inequality shows that

$$|\mathbb{E}_{\nu_{pn}}[f] - \mathbb{E}_{\mu_p}[f]| \leq \sum_{x \in \{0,1\}^m} |\nu'_{pn}(x) - \mu'_p(x)| |f(x)| \leq \sum_{x \in \{0,1\}^m} |\nu'_{pn}(x) - \mu'_p(x)|,$$

the last expression being exactly the total variation distance between ν'_{pn} and μ'_p . □

Corollary 7.3. *Let f be a function on $\{0, 1\}^n$ depending on $o(p(1-p)n)$ coordinates. As $p(1-p)n \rightarrow \infty$, the CDF distance between $\nu_{pn}(f)$ and $\mu_p(f)$ tends to zero, that is,*

$$\sup_{t \in \mathbb{R}} |\Pr_{\nu_{pn}}[f < t] - \Pr_{\mu_p}[f < t]| = o(1).$$

Proof. Consider the functions $f_t = \mathbf{1}_{f < t}$, which satisfy $\|f_t\|_\infty = 1$ for all $t \in \mathbb{R}$. □

8 Functions with low total influence

In this section we prove an invariance principle for Boolean functions whose total influence is $o(\sqrt{p(1-p)n})$. In other words, we show that Boolean functions with total influence $o(\sqrt{p(1-p)n})$ cannot distinguish the slice from the cube.

For us a Boolean function is a function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, and total influence is defined as follows: $\text{Inf}[f] = p(1-p)n \Pr[f(x) \neq f(y)]$, where $x \sim \mu_p$ and y is a random neighbor of x , chosen uniformly from all n neighbors (this differs from the normalization in [18] by a factor of $p(1-p)$). As in Section 7, the functions we consider in this section need not be harmonic multilinear polynomials.

The upper bound $\sqrt{p(1-p)n}$ is necessary. Indeed, consider a threshold function with threshold pn . This function has total influence $\Theta(\sqrt{p(1-p)n})$, is constant on the slice, and is roughly balanced on the cube. The condition that the function is Boolean is also necessary (total influence can be extended to arbitrary functions, see for example [31, Section 8.4]). The function $(x_1 + \dots + x_n)/\sqrt{p(1-p)n}$ has unit variance and unit total influence on the cube, is constant on the slice, and has a non-trivial distribution on the cube. Invariance does hold, however, for non-Boolean harmonic multilinear polynomials of degree $o(\sqrt{p(1-p)n})$, as Theorem 4.6 shows.

We will use part of the setup of Definition 4.1, namely the random variables $\mathbf{X}(s) \in \binom{[n]}{s}$ for $0 \leq s \leq n$. We also define two new random variables: $\mathbf{s} \sim B(n, p)$ and $\mathbf{t} \sim B(n-1, p)$, where $B(n, p)$ is the binomial distribution with n trials and success probability p . The basic observation is that $\mathbf{X}(\mathbf{s})$ is distributed according to the measure μ_p on the cube $\{0, 1\}^n$ whereas $\mathbf{X}(np)$ is distributed uniformly on the slice $\binom{[n]}{pn}$, and so our goal would be to bound the probability $\Pr[f(\mathbf{X}(\mathbf{s})) \neq f(\mathbf{X}(np))]$.

The following lemma bounds $\Pr[f(\mathbf{X}(\mathbf{s})) \neq f(\mathbf{X}(np))]$ using a hybrid argument.

Lemma 8.1. *For every Boolean function f ,*

$$\Pr[f(\mathbf{X}(\mathbf{s})) \neq f(\mathbf{X}(np))] \leq \sum_{s=np}^{n-1} \Pr[\mathbf{s} \geq s+1] \Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1))] + \sum_{s=0}^{np-1} \Pr[\mathbf{s} \leq s] \Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1))].$$

Proof. For an event E , let $\llbracket E \rrbracket$ denote the corresponding indicator. Let \mathbf{I} denote the interval whose endpoints are \mathbf{s} and np (inclusive). Then

$$\begin{aligned} \llbracket f(\mathbf{X}(\mathbf{s})) \neq f(\mathbf{X}(np)) \rrbracket &\leq \sum_{\substack{s: \\ \{s, s+1\} \subseteq \mathbf{I}}} \llbracket f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1)) \rrbracket = \\ &\sum_{s=0}^{n-1} \llbracket \{s, s+1\} \subseteq \mathbf{I} \rrbracket \cdot \llbracket f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1)) \rrbracket. \end{aligned}$$

Since \mathbf{I} is independent of \mathbf{X} , taking expectations we get

$$\begin{aligned} \Pr[f(\mathbf{X}(\mathbf{s})) \neq f(\mathbf{X}(np))] &\leq \sum_{s=0}^{n-1} \Pr[\{s, s+1\} \subseteq \mathbf{I}] \Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1))] \\ &= \sum_{s=np}^{n-1} \Pr[\mathbf{s} \geq s+1] \Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1))] \\ &\quad + \sum_{s=0}^{np-1} \Pr[\mathbf{s} \leq s] \Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1))]. \quad \square \end{aligned}$$

We can write the total influence in a very similar form.

Lemma 8.2. *For every Boolean function f ,*

$$\text{Inf}[f] = \sum_{s=0}^{n-1} p(1-p)n \Pr[\mathbf{t} = s] \Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(s+1))].$$

Proof. Let (x, y) be a random edge of the cube oriented so that $|x| < |y|$, where $|x|$ is the Hamming weight of x . We can pick (x, y) in the following way. Pick $z \sim \mu_p$ and a random coordinate i , and let $x = z|_{i=0}$ and $y = z|_{i=1}$ (that is, x is obtained from z by setting the i th coordinate to zero, and y by setting it to one). The Hamming weight of x thus has the same distribution of \mathbf{t} . Therefore $(x, y) \sim (\mathbf{X}(\mathbf{t}), \mathbf{X}(\mathbf{t} + 1))$, which directly implies the formula given above. \square

In order to bound $\Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(np))]$, it remains to analyze the ratio between the coefficients in the two lemmas.

Lemma 8.3. *If $p(1-p)n \rightarrow \infty$ then for every Boolean function f ,*

$$\Pr[f(\mathbf{X}(s)) \neq f(\mathbf{X}(np))] = O\left(\frac{1}{\sqrt{p(1-p)n}} \text{Inf}[f]\right).$$

Proof. In view of Lemma 8.1 and Lemma 8.2, it remains to bound from above the following two ratios:

$$\begin{aligned} \rho_1(s) &= \frac{\Pr[\mathbf{s} \geq s+1]}{p(1-p)n \Pr[\mathbf{t} = s]} \quad (pn \leq s \leq n-1), \\ \rho_2(s) &= \frac{\Pr[\mathbf{s} \leq s]}{p(1-p)n \Pr[\mathbf{t} = s]} \quad (0 \leq s \leq pn-1). \end{aligned}$$

It is not hard to check that

$$\rho_1(s) \leq \frac{\Pr[\mathbf{t} \geq s]}{p(1-p)n \Pr[\mathbf{t} = s]}, \quad \rho_2(s) \leq \frac{\Pr[\mathbf{t} \leq s]}{p(1-p)n \Pr[\mathbf{t} = s]}.$$

Log-concavity of the binomial coefficients shows that ρ_1 is decreasing while ρ_2 is increasing, and so in the given ranges, $\rho_1(s) \leq \rho_1(np)$ and $\rho_2(s) \leq \rho_2(np)$. It is known that the median of the binomial distribution $B(n-1, p)$ is one of $np-1, np$, and so the numerator of both $\rho_1(np)$ and $\rho_2(np)$ is $1/2 \pm o(1)$. The local limit theorem shows that the common denominator is $(1 \pm o(1))\sqrt{p(1-p)n}/2\pi$. Therefore

$$\rho_1(s), \rho_2(s) \leq \sqrt{\frac{\pi}{2} \frac{1}{p(1-p)n}}. \quad \square$$

Our main result in this section immediately follows.

Theorem 8.4. *If $p(1-p)n \rightarrow \infty$ and f is a Boolean function satisfying $\text{Inf}[f] = o(\sqrt{p(1-p)n})$ then $|\mathbb{E}_{\mu_p}[f] - \mathbb{E}_{\nu_{pn}}[f]| = o(1)$.*

As a corollary, we prove that balanced symmetric Boolean functions cannot be approximated by Boolean functions whose influence is $o(\sqrt{p(1-p)n})$.

Proposition 8.5. *There exists a constant $\delta > 0$ such that for all $p \in (0, 1)$ and large enough n , if f is a symmetric Boolean function such that $\frac{1}{3} \leq \mathbb{E}_{\mu_p}[f] \leq \frac{2}{3}$ then $\Pr_{\mu_p}[f \neq g] > \delta$ for every Boolean function g satisfying $\text{Inf}[g] = o(\sqrt{p(1-p)n})$.*

Our result can be considered a generalization of two recent results (with much weaker bounds):

- O'Donnell and Wimmer [32] proved much tighter results (in terms of δ) in the special case of the majority function.

- Tal [39] proved tighter bounds (in terms of δ) for general symmetric functions assuming g is an AC^0 circuit whose influence sum is $o(\sqrt{p(1-p)n})$.

Proof. It is known that the total variation distance between μ_p and μ_q in the regime $|p - q| = o(1)$ is $O(|p - q|\sqrt{\frac{n}{p(1-p)}})$. Choose $C > 0$ so that $|p - q| \leq C\sqrt{p(1-p)/n}$ implies that the variation distance between μ_p and μ_q is at most $1/4$. Call a q satisfying this bound *good*.

Let g be a Boolean function satisfying $\text{Inf}[g] = o(\sqrt{p(1-p)n})$. Suppose that $\Pr_{\mu_p}[f \neq g] \leq \delta$. If $x \sim \mu_p$ then $|x|/n$ is good with constant probability, and so there exists a slice $\binom{[n]}{qn}$ with q good such that $\Pr_{\nu_{qn}}[f \neq g] = O(\delta)$.

Theorem 8.4 implies that $|\mathbb{E}_{\mu_q}[g] - \mathbb{E}_{\nu_{qn}}[g]| = o(1)$. On the one hand, $|\mathbb{E}_{\nu_{qn}}[g] - \mathbb{E}_{\nu_{qn}}[f]| = O(\delta)$, and so the symmetry of f implies that either $\mathbb{E}_{\nu_{qn}}[g] = O(\delta)$ or $\mathbb{E}_{\nu_{qn}}[g] = 1 - O(\delta)$. On the other hand, $|\mathbb{E}_{\mu_q}[g] - \mathbb{E}_{\mu_p}[f]| \leq |\mathbb{E}_{\mu_q}[g] - \mathbb{E}_{\mu_p}[g]| + |\mathbb{E}_{\mu_p}[g] - \mathbb{E}_{\mu_p}[f]| \leq 1/4 + \delta$. We conclude that either $\mathbb{E}_{\mu_p}[f] \leq 1/4 + O(\delta) + o(1)$ or $\mathbb{E}_{\mu_p}[f] \geq 3/4 - O(\delta) - o(1)$. For an appropriate choice of δ , this contradicts the assumption $1/3 \leq \mathbb{E}_{\mu_p}[f] \leq 2/3$ for large enough n . \square

9 Decomposing the slice

In this section we describe two ways of decomposing the space of functions over the slice. One decomposition arises naturally from the harmonic representation: $f = \sum_{d \leq n/2} f^{=d}$. The corresponding decomposition of the space of functions is into the subspaces of homogeneous harmonic multilinear polynomials of degree d , for $0 \leq d \leq n/2$. We call this decomposition the *coarse decomposition*, and discuss it further in Section 9.1. This decomposition is the underlying reason for results such as Theorem 3.10, Lemma 3.13, and Lemma 3.15

The coarse decomposition can be refined into a distinguished basis of the space of functions over the slice, known as the *Gelfand–Tsetlin basis*, which is described in Section 9.2. This decomposition has been used by the first author [17] to simplify the proof of Friedgut’s junta theorem for the slice, due to Wimmer [42].

This section assumes basic knowledge of the representation theory of the symmetric group, such as the one provided by Chapter 2 and Chapter 8 of Diaconis [12].

9.1 Coarse decomposition

Every harmonic function f can be decomposed as a sum of homogeneous parts:

$$f = \sum_{d \leq n/2} f^{=d}.$$

This decomposition naturally arises from the representation theory of the slice. In this subsection we discuss this connection, and obtain a representation-theoretic proof of Theorem 3.10 as a result.

So far we have viewed the slice as a subset of the Boolean cube. However, it is also possible to view the slice $\binom{[n]}{k}$ as the set of cosets of $S_k \times S_{n-k}$ inside S_n . In other words, we identify a subset $A \in \binom{[n]}{k}$ with the set of permutations π such that $\pi(\{1, \dots, k\}) = A$. From this point of view it is natural to consider the action of S_n on $\binom{[n]}{k}$ by permutation of the coordinates. The resulting module $\mathbb{R}[\binom{[n]}{k}]$ is isomorphic to the permutation module $M^{(n-k,k)}$ (assuming $k \leq n/2$) whose decomposition into irreducible modules (Specht modules) can easily be computed to be

$$\mathbb{R} \left[\binom{[n]}{k} \right] \approx M^{(n-k,k)} \approx S^{(n)} \oplus S^{(n-1,1)} \oplus \dots \oplus S^{(n-k,k)}.$$

Since each irreducible representation appears once, $(S_n, S_k \times S_{n-k})$ forms a *Gelfand pair*. Classical representation theory of the symmetric group (for example, [35, §2.9–2.10]) allows us to identify $S^{(n-d,d)}$ with the space of all homogeneous harmonic multilinear polynomials of degree d over the slice.

Consider now the middle slice $\binom{[n]}{[n/2]}$. We can uniquely identify every harmonic multilinear polynomial on x_1, \dots, x_n with a function over the middle slice. If f, g are harmonic multilinear polynomials and α

is an exchangeable measure then $\mathbb{E}_\alpha[fg] = f^T Ag$, where A commutes with the action of S_n (since α is exchangeable). Since A commutes with the action, Schur's lemma implies that it acts on each irreducible $S^{(n-d,d)}$ by scalar multiplication (since each irreducible representation appears only once). This implies that for some constants $\lambda_0, \dots, \lambda_{\lfloor n/2 \rfloor}$ depending only on α ,

$$f^T Ag = \sum_{d \leq n/2} \lambda_d \langle f^{=d}, g^{=d} \rangle,$$

where $f^{=d}$ is the component of f in $S^{(n-d,d)}$ (defined uniquely since there is only one copy of this irreducible representation) and the inner product is given by $\langle x, y \rangle = x^T y$; there are no mixed summands $\langle f^{=d}, g^{=e} \rangle$ since the decomposition into irreducible representations is orthogonal. Theorem 3.10 follows by taking, for each $d \leq n/2$,

$$C_{f,g} = \frac{\langle f^{=d}, g^{=d} \rangle}{\|(x_1 - x_2) \dots (x_{2d-1} - x_{2d})\|^2}.$$

A similar argument explains why the decomposition into homogeneous parts appears in some of our other results such as Lemma 3.13 and Lemma 3.15. More generally, if A is any operator on $\mathbb{R}[\binom{[n]}{k}]$ which commutes with the action of S_n then

$$Af = \sum_{d \leq \min(k, n-k)} \lambda_d f^{=d},$$

where $\lambda_0, \dots, \lambda_d$ are the eigenvalues of A .

The decomposition into homogeneous parts appears in several other guises:

Association schemes The operators on $\mathbb{R}[\binom{[n]}{k}]$ which commute with the action of S_n form a commutative algebra known as the *Bose–Mesner algebra* of the *Johnson association scheme*. This algebra has two important bases: the spatial basis corresponds to the characterization of the algebra as the set of matrices A in which $A(S, T)$ depends only on $|S \cap T|$, and the spectral basis corresponds to the decomposition into homogeneous parts. For more on association schemes, see Bannai and Ito [2].

Regular semilattices and spherical posets For each k , we can consider the subset of the n -dimensional Boolean cube consisting of all sets of cardinality at most k as a lattice under the set inclusion order. The truncated Boolean lattice is an example of a *regular semilattice* and, when $k \leq n/2$, of a *spherical lattice*. Delsarte [11] and Stanton [38] gave general constructions which recover the decomposition of $\mathbb{R}[\binom{[n]}{k}]$ into its irreducible components. For more, see Ceccherini-Silberstein et al. [9, §8.2–8.4].

Differential posets We can also consider the entire Boolean cube as a lattice. The Boolean lattice is an example of a μ -*differential poset*. Stanley [37] gave a general construction, which in the special case of the Boolean lattice decomposes it into the various slices, and each slice into the irreducible components. This decomposition corresponds to the *Terwilliger algebra* of the *binary Hamming association scheme*. For more on this area, see Engel [16, §6.2].

9.2 Gelfand–Tsetlin basis

So far we have discussed the decomposition of each function on the slice into its homogeneous parts. This decomposition has a similar counterpart for functions on the Boolean cube, the so-called *levels* of the Fourier expansion. For functions on the Boolean cube, the decomposition can be refined to the *Fourier basis*. The Fourier basis is uniquely defined (up to scalar multiplication) as the set of characters of the group \mathbb{Z}_2^n . In the case of the slice, we can obtain such a canonical basis by specifying an order on the coordinates, say the standard order x_1, \dots, x_n .

Recall the decomposition into irreducible S_n -modules of $\mathbb{R}[\binom{[n]}{k}]$ discussed in the preceding subsection:

$$\mathbb{R} \left[\binom{[n]}{k} \right] \approx M^{(n-k,k)} \approx S^{(n)} \oplus \dots \oplus S^{(n-k,k)}.$$

If we consider $S^{(n-d,d)}$ (the space of homogeneous harmonic multilinear polynomials of degree d) as a module over S_{n-1} rather than over S_n , then it is no longer irreducible. Instead, it decomposes into two irreducibles: $S^{(n-1-d,d)} \oplus S^{(n-d,d-1)}$. If we then consider each of these irreducibles as modules over S_{n-2} they decompose even further, and continuing in this way, eventually we get a decomposition into irreducibles of dimension 1. The corresponding basis (defined up to scalar multiplication) is known as the *Gelfand–Tsetlin basis*.

Srinivasan [36] described an inductive construction of this basis using ideas from Sperner theory, which are related to the theory of differential posets mentioned above. Ambainis et al. [1] gave an inductive construction which closely follows the definition, in the context of quantum computation. Filmus [17] gave an explicit formula for the basis elements, which we now describe.

A sequence $(a_1, b_1), \dots, (a_d, b_d)$ is *admissible* if:

1. All $2d$ numbers are distinct and belong to $[n]$.
2. $b_1 < \dots < b_d$.
3. $a_i < b_i$ for all i .

A set $B = \{b_1, \dots, b_d\}$ (where $b_1 < \dots < b_d$) is *admissible* if it can be completed to an admissible sequence (equivalently, it is the bottom row of a standard Young tableau of shape $(n-d, d)$). There are $\binom{n}{d} - \binom{n}{d-1}$ admissible sets of size d . For each admissible set B , we define a basis element

$$\chi_B = \sum_{\substack{a_1, \dots, a_d: \\ (a_1, b_1), \dots, (a_d, b_d) \text{ admissible}}} (x_{a_1} - x_{b_1}) \cdots (x_{a_d} - x_{b_d}).$$

In total there are $\binom{n}{k}$ admissible sets of size at most k , and the corresponding basis elements constitute the Gelfand–Tsetlin basis for $\binom{[n]}{k}$. Furthermore, if α is any exchangeable measure then

$$\|\chi_B\|_\alpha^2 = \binom{b_1}{2} \binom{b_2 - 2}{2} \cdots \binom{b_d - 2(d-1)}{2} \|(x_1 - x_2) \cdots (x_{2d-1} - x_{2d})\|_\alpha^2.$$

For proofs, see Filmus [17].

The Gelfand–Tsetlin basis can also be characterized (up to scalar multiplication) as the common eigenvectors of the Gelfand–Tsetlin algebra, which is generated by the Young–Jucys–Murphy elements

$$X_m = \sum_{i < m} (i \ m).$$

Filmus [17, Lemma 24] gives the following explicit formula:

$$\left[\sum_{\ell=1}^m X_\ell \right] \chi_B = \left[\binom{m}{2} - |B \cap [m]|(m+1 - |B \cap [m]|) \right] \chi_B.$$

Finally, this basis can be constructed in the same way as Young’s orthogonal basis. For each d , consider all standard Young tableaux of shape $(n-d, d)$, arranged in lexicographic order of the second row. With each tableau having first row a_1, \dots, a_{n-d} and second row b_1, \dots, b_d , associate the function $(x_{a_1} - x_{b_1}) \cdots (x_{a_d} - x_{b_d})$. Running the Gram–Schmidt orthogonalization process on these functions yields the degree d part of the Gelfand–Tsetlin basis described above.

10 Why harmonic functions?

Our invariance principle compares the distribution of harmonic multilinear polynomials on the slice and on the Boolean cube (with respect to an appropriate measure). This empirically validates the following

informal claim: *The correct way to lift a function from the slice to the Boolean cube is through its unique representation as a bounded degree harmonic multilinear polynomial.*

In the same way, the classical invariance principle of Mossel et al. [29] suggests that the correct way to lift a function from the Boolean cube to Gaussian space is through its unique representation as a multilinear polynomial. Alternatively, given a function on the Boolean cube, we lift it to a function on Gaussian space by linear interpolation in all axis-parallel directions.

In this section we discuss two other justifications for using harmonic multilinear polynomials, from commutative algebra and from representation theory. We briefly mention that James' intersecting kernel theorem [21] (see also [9, §10.9] and [10, §6.2]) draws together both points of view, and generalizes them from the slice to the *multislice*, in which the object of study is an arbitrary permutation module of the symmetric group. We hope to explore this connection in future work.

10.1 Commutative algebra

Let us consider first the case of the Boolean cube $\{-1, 1\}^n$. The classical invariance principle [29] compares the distribution of multilinear polynomials over the Boolean cube and over Gaussian space. Multilinear polynomials arise naturally when we consider $\{-1, 1\}^n$ as a variety defined by the ideal $I = \langle x_1^2 - 1, \dots, x_n^2 - 1 \rangle$. The coordinate ring over the variety consists of polynomials over x_1, \dots, x_n modulo the ideal I . It is immediately clear that every polynomial is equivalent to a multilinear polynomial modulo I , and a dimension computation shows that this representation is unique.

We can treat the case of the slice $\binom{[n]}{k}$ in a similar way. The slice is a variety defined by the ideal

$$I_k = \langle x_1^2 - x_1, \dots, x_n^2 - x_n, \sum_{i=1}^n x_i - k \rangle.$$

Theorem 3.6 shows that the coordinate ring is isomorphic to the space of harmonic multilinear polynomials of degree at most $\min(k, n - k)$. Intuitively speaking, the existence of $x_i^2 - x_i$ in the ideal allows us to reduce every polynomial to a multilinear polynomial; the existence of $\sum_{i=1}^n x_i - k$ corresponds to the harmonicity constraint; and the degree constraint follows from the fact that $x_{i_1} \dots x_{i_{k+1}} = 0$ (when $k \leq n/2$).

We can formalize this intuition to give another proof of Theorem 3.6, which uses arguments due to Blekherman [5] quoted in Lee et al. [25]. Similar arguments can be found in Engel [16, §6.2].

Theorem 10.1 (Reformulation of Theorem 3.6). *Let $0 \leq k \leq n$. For every polynomial P over x_1, \dots, x_n there exists a unique harmonic multilinear polynomial Q of degree at most $\min(k, n - k)$ such that $P \equiv Q \pmod{I_k}$.*

Proof. We start by showing that every polynomial is equivalent modulo I_k to a harmonic multilinear polynomial of degree at most k . We will use the following definitions:

- Given $S \subseteq [n]$, $x_S = \prod_{i \in S} x_i$.
- We denote by P_d the linear space of all homogeneous multilinear polynomials of degree d , and by $H_d \subset P_d$ its subspace consisting of harmonic polynomials.

The first step is replacing x_i^r (for $r > 1$) by x_i using the equation $x_i^2 = x_i$, which holds modulo I_k . We are left with a multilinear polynomial.

The second step is removing all monomials of degree larger than $\min(k, n - k)$. If $k \leq n/2$, then monomials of degree larger than k always evaluate to zero on the slice⁴. If $k \geq n/2$, then we get a similar simplification by replacing x_i with $1 - (1 - x_i)$, and using the fact that the product of more than $n - k$ different factors $1 - x_i$ always evaluates to zero on the slice. We are left with a multilinear polynomial of degree at most $\min(k, n - k)$.

⁴Formally speaking, this step requires us to work with $\sqrt{I_k}$. However, the ideal I_k is radical, see for example [34, Lemma 6.1].

Definition 3.1 defines a differential operator $\Delta = \sum_{i=1}^n \frac{\partial}{\partial x_i}$. When applied to multilinear polynomials, we can also think of it as a formal operator which maps x_S to $\sum_{i \in S} x_{S \setminus \{i\}}$, and extends linearly to multilinear polynomials (this is the Lefschetz lowering operator in the Boolean lattice). When restricted to P_d , the range of Δ is P_{d-1} . Linear algebra shows that P_d decomposes into $\ker \Delta$ and $\text{im } \Delta^T$, where Δ^T (the Lefschetz raising operator) is the operator on P_{d-1} which maps x_S to $\sum_{i \in [n] \setminus S} x_{S \cup \{i\}}$. Modulo I_k ,

$$\Delta^T x_S = \sum_{i \in [n] \setminus S} x_{S \cup \{i\}} = x_S \sum_{i \in [n] \setminus S} x_i = (k - |S|)x_S = (k - d + 1)x_S,$$

the third equality following logically since $x_S \neq 0$ only if $x_i = 1$ for all $i \in S$, and formally using $x_S x_i = x_S$ for all $i \in S$. Thus $\text{im } \Delta^T$ can be identified with P_{d-1} . We conclude that any polynomial in P_d can be written (modulo I_k) as the sum of a polynomial in H_d (belonging to $\ker \Delta$) and a polynomial in P_{d-1} (belonging to $\text{im } \Delta^T$).

Applying this construction recursively, we see that every homogeneous multilinear polynomial is equivalent modulo I_k to a harmonic multilinear polynomial, and moreover this operation doesn't increase the degree. This completes the proof that every polynomial is equivalent modulo I_k to a harmonic multilinear polynomial of degree at most $\min(k, n - k)$.

It remains to show that this representation is unique. We first show that Δ^T , considered as an operator from P_{d-1} to P_d for $d \leq n/2$, has full rank $\dim P_{d-1}$. Notice that for $S \subseteq [n]$,

$$\begin{aligned} \Delta \Delta^T x_S &= \Delta \sum_{i \in [n] \setminus S} x_{S \cup \{i\}} = (n - |S|)x_S + \sum_{i \in [n] \setminus S} \sum_{j \in S} x_{S \cup \{i\} \setminus \{j\}}, \\ \Delta^T \Delta x_S &= \Delta^T \sum_{j \in S} x_{S \setminus \{j\}} = |S|x_S + \sum_{i \in [n] \setminus S} \sum_{j \in S} x_{S \cup \{i\} \setminus \{j\}}. \end{aligned}$$

We conclude that on P_{d-1} , the following holds:

$$\Delta \Delta^T = \Delta^T \Delta + (2n - d + 1)I,$$

where I is the identity operator. Note that $\Delta^T \Delta$ is positive semidefinite. Since $d \leq n/2$, we see that $2n - d + 1 > 0$, and so $\Delta \Delta^T$ is positive definite, and in particular regular. It follows that Δ^T is also regular, and so has full rank.

The decomposition $\ker \Delta \oplus \text{im } \Delta^T$ of P_d shows that when $d > 0$, H_d has dimension $\dim \ker \Delta = \binom{n}{d} - \dim \text{im } \Delta^T = \binom{n}{d} - \binom{n}{d-1}$. When $d = 0$, the dimension is clearly 1. It follows that the space of all harmonic multilinear polynomials of degree at most $\min(k, n - k)$ has dimension $1 + \sum_{d=1}^{\min(k, n-k)} \left[\binom{n}{d} - \binom{n}{d-1} \right] = \binom{n}{k}$, matching the dimension of the space of functions on the slice. \square

Using very similar ideas, we can prove Blekherman's theorem stated in Section 6.1.

Theorem 10.2 (Reformulation of Theorem 6.1). *Every multilinear polynomial f over x_1, \dots, x_n of degree at most $d \leq n/2$ can be represented uniquely in the form*

$$f(x_1, \dots, x_n) \equiv \sum_{i=0}^d f_i(x_1, \dots, x_n) S^i \pmod{I},$$

where $I = \langle x_1^2 - x_1, \dots, x_n^2 - x_n \rangle$, $S = x_1 + \dots + x_n$, and f_i is a harmonic multilinear polynomial of degree at most $d - i$.

Proof. Throughout the proof we use the definitions of $x_S, P_d, H_d, \Delta, \Delta^T$ in the proof of Theorem 10.1. We also define $P_{\leq d}$ to be the space of all multilinear polynomials of degree at most d , and $H_{\leq d}$ to be the space of all harmonic multilinear polynomials of degree at most d .

We start by considering the case in which f is a homogeneous multilinear polynomial. We prove by induction on e that if $\deg f = e$ then for some f_0, \dots, f_e ,

$$f \equiv \sum_{i=0}^e f_i S^i \pmod{I}, \text{ where } f_i \in H_{\leq e-i}.$$

When $e = 0$ this is clear, since f is already harmonic, so suppose that $e > 0$. Recall that Δ maps P_e to P_{e-1} . The conjugate operator Δ^T maps $x_A \in P_{e-1}$ to

$$\sum_{i \notin A} x_{A \cup \{i\}} = x_A \cdot \sum_{i \notin A} x_i = x_A \cdot \left(S - \sum_{i \in A} x_i \right) \equiv (S - (e-1))x_A \pmod{I}.$$

Since $P_e = \ker \Delta \oplus \text{im } \Delta^T = H_d \oplus \text{im } \Delta^T$, it follows that $f \equiv h + (S - (e-1))g \pmod{I}$ for some $h \in H_e$ and $g \in P_{e-1}$. By the induction hypothesis, for some g_0, \dots, g_{e-1} ,

$$g \equiv \sum_{i=0}^{e-1} g_i S^i \pmod{I}, \text{ where } g_i \in H_{\leq e-1-i}.$$

Substituting this in the equation for f , we obtain

$$f \equiv h + \sum_{i=0}^{e-1} (S - (e-1))g_i S^i \equiv [h - (e-1)g_0] + \sum_{i=1}^e [g_{i-1} - (e-1)g_i] S^i \pmod{I},$$

where $g_e = 0$. Since $\deg[h - (e-1)g_0] \leq \max(e, e-1) = e$ and $\deg[g_{i-1} - (e-1)g_i] \leq \max((e-1) - (i-1), e-1-i) = e-i$, we have obtained the required representation.

When f is an arbitrary multilinear polynomial of degree d , we get the required representation by summing the representations of $f^{=0}, \dots, f^{=d}$.

Finally, we prove that the representation is unique by comparing dimensions:

$$\dim P_{\leq d} = \sum_{i=0}^d \binom{n}{i} = \sum_{i=0}^d \dim H_{\leq i} = \sum_{i=0}^d \dim H_{\leq d-i},$$

since $H_{\leq e}$ has dimension $\binom{n}{e}$ for $e \leq n/2$ by Theorem 3.6. \square

This implies a surprising corollary for *harmonic projections*.

Definition 10.3. Let f be a function over x_1, \dots, x_n . Its *harmonic projection* on the slice $\binom{[n]}{k}$ is the unique harmonic multilinear polynomial of degree at most $\min(k, n-k)$ which agrees with f on the slice.

Corollary 10.4. Let f be a multilinear polynomial over x_1, \dots, x_n of degree d , and let F_k be its harmonic projection on the slice $\binom{[n]}{k}$, where $d \leq k \leq n-d$. For k in this range, $F_k^{=d}$ does not depend on k .

Proof. Let f_0, \dots, f_d be the Blekherman decomposition of f , where f_i is a harmonic multilinear polynomial of degree at most $d-i$. The function $\sum_{i=0}^d f_i k^i$ is thus a harmonic multilinear polynomial which agrees with f on $\binom{[n]}{k}$. When $d \leq k \leq n-d$, this polynomial has degree at most $d \leq \min(k, n-k)$, and so it is the harmonic projection F_k of f on the slice $\binom{[n]}{k}$. Notice that $F_k^{=d} = f_0^{=d}$ does not depend on k . \square

10.2 Representation theory

As we have explained in Subsection 9.1, representation theory naturally leads to the decomposition of each function on the slice to its homogeneous parts. From the point of view of the representation theory of the symmetric group, it is natural to order the irreducible representations in increasing order of complexity:

$S^{(n)}, S^{(n-1,1)}, \dots, S^{(n-k,k)}$. This suggests defining the *degree* of a function f on the slice as the maximal d such that $f^{=d} \neq 0$; note that this definition doesn't make use of the harmonic multilinear representation of f .

Another definition which leads to the same parameter is the *junta degree*, which is the minimal d such that f can be written as a linear combination of functions depending on at most d coordinates. The cometric property (Corollary 3.19), which (apart from a few degenerate cases) only holds for this ordering of the irreducible representations, gives further credence to the notion of degree.

A similar notion of degree exists also for functions on the Boolean cube: the degree of a function on $\{0, 1\}^n$ is both its degree as a multilinear polynomial and its junta degree.

The following result shows that the lifting defined by the harmonic multilinear representation is the unique way to couple both notions of degree.

Theorem 10.5. *Let $0 \leq k \leq n/2$. Suppose L is a linear operator from $\mathbb{R}[\binom{[n]}{k}]$ to $\mathbb{R}[\{0, 1\}^n]$ such that*

(a) *For every function f over $\binom{[n]}{k}$, Lf is an extension of f , that is, $(Lf)(x) = f(x)$ for every $x \in \binom{[n]}{k}$.*

(b) *For every function f over $\binom{[n]}{k}$, $\deg Lf \leq \deg f$.*

(c) *For every function f over $\binom{[n]}{k}$ and for every permutation $\pi \in S_n$, $L(f^\pi) = (Lf)^\pi$. In other words, L commutes with the action of S_n .*

Then Lf is given by the unique harmonic multilinear polynomial of degree at most k which agrees with f on $\binom{[n]}{k}$.

Proof. In view of Lemma 3.7, it suffices to show that $Lf = f$ for $f = (x_1 - x_2) \cdots (x_{2d-1} - x_{2d})$. Suppose that $Lf = g$, where g is a multilinear polynomial of degree at most d . Since $f^{(1\ 2)} = -f$, also $g^{(1\ 2)} = -g$. If we write

$$g = x_1x_2r + x_1s + x_2t + u,$$

where r, s, t, u don't involve x_1, x_2 , then $g^{(1\ 2)} = x_1x_2r + x_1t + x_2s + u$, showing that $r = u = 0$ and $s = -t$. In other words, $g = (x_1 - x_2)s$ is a multiple of $x_1 - x_2$. Similarly g is a multiple of $x_3 - x_4, \dots, x_{2d-1} - x_{2d}$. Unique factorization forces $g = C(x_1 - x_2) \cdots (x_{2d-1} - x_{2d}) = Cf$ for some constant C . Since g extends f , the constant must be $C = 1$. \square

References

- [1] Andris Ambainis, Aleksandrs Belovs, Oded Regev, and Ronald de Wolf. Efficient quantum algorithms for (gapped) group testing and junta testing. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on discrete algorithms (SODA '16)*, pages 903–922, 2016.
- [2] Eiichi Bannai and Tatsuhiro Ito. *Algebraic Combinatorics I: Association Schemes*. Mathematics lecture notes series. Benjamin / Cummings, 1984.
- [3] William Beckner. Inequalities in Fourier analysis. *Ann. Math.*, 102:159–182, 1975.
- [4] François Bergeron. *Algebraic Combinatorics and Coinvariant Spaces*. CMS Treatises in Mathematics. A K Peters, 2009.
- [5] Greg Blekherman. Symmetric sums of squares on the hypercube, 2015. Manuscript in preparation.
- [6] Aline Bonami. Étude des coefficients Fourier des fonctions de $L^p(G)$. *Ann. Inst. Fourier*, 20(2):335–402, 1970.
- [7] Ravi B Boppana. The average sensitivity of bounded-depth circuits. *Information Processing Letters*, 63(5):257–261, 1997.

- [8] Anthony Carbery and James Wright. Distributional and L^q norm inequalities for polynomials over convex bodies in \mathbb{R}^n . *Math. Res. Lett.*, 3(8):233–248, 2001.
- [9] Tullio Ceccherini-Silberstein, Fabio Scarabotti, and Filippo Tolli. *Harmonic analysis on finite groups*, volume 108 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 2008.
- [10] Tullio Ceccherini-Silberstein, Fabio Scarabotti, and Filippo Tolli. *Representation theory of the symmetric groups*, volume 121 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 2010.
- [11] Phillipe Delsarte. Association schemes and t -designs in regular semilattices. *J. Comb. Theory Ser. A*, 20(2):230–243, 1976.
- [12] Persi Diaconis. *Group representations in probability and statistics*, volume 11 of *Institute of Mathematical Statistics Lecture Notes—Monograph Series*. Institute of Mathematical Statistics, Hayward, CA, 1988.
- [13] Persi Diaconis and Laurent Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Prob.*, 6(3):695–750, 1996.
- [14] Charles F. Dunkl. A Krawtchouk polynomial addition theorem and wreath products of symmetric groups. *Indiana Univ. Math. J.*, 25:335–358, 1976.
- [15] Charles F. Dunkl. Orthogonal functions on some permutation groups. In *Relations between combinatorics and other parts of mathematics*, volume 34 of *Proc. Symp. Pure Math.*, pages 129–147, Providence, RI, 1979. Amer. Math. Soc.
- [16] Konrad Engel. *Sperner Theory*, volume 65 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1997.
- [17] Yuval Filmus. An orthogonal basis for functions over a slice of the boolean hypercube. *Elec. J. Comb.*, 23(1):P1.23, 2016.
- [18] Yuval Filmus, Guy Kindler, Elchanan Mossel, and Karl Wimmer. Invariance principle on the slice. In *31st Conf. Comp. Comp.*, 2016.
- [19] Johan Håstad. Almost optimal lower bounds for small depth circuits. In Silvio Micali, editor, *Randomness and Computation*, volume 5 of *Advances in Computing Research*, pages 143–170. JAI Press, 1989.
- [20] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- [21] Gordon D. James. A characteristic free approach to the representation theory of S_n . *J. Algebra*, 46:430–450, 1977.
- [22] Nathan Keller and Ohad Klein. A structure theorem for almost low-degree functions on the slice. Manuscript.
- [23] Guy Kindler. *Property testing, PCP and Juntas*. PhD thesis, Tel-Aviv University, 2002.
- [24] Guy Kindler and Shmuel Safra. Noise-resistant Boolean functions are juntas, 2004. Unpublished manuscript.
- [25] Troy Lee, Anupam Prakash, Ronald de Wolf, and Henry Yuen. On the sum-of-squares degree of symmetric quadratic functions. In *Proceedings of the 31st Conference on Computational Complexity (CCC 2016)*, pages 17:1–17:31, 2016.
- [26] Tzong-Yau Lee and Horng-Tzer Yau. Logarithmic Sobolev inequality for some models of random walks. *Ann. Prob.*, 26(4):1855–1873, 1998.

- [27] N. Linial, Y. Mansour, and N. Nisan. Constant depth circuits, fourier transform and learnability. *Journal of the ACM*, 40(3):607–620, 1993.
- [28] Terry J. Lyons and T. S. Zhang. Decomposition of Dirichlet processes and its application. *Ann. Probab.*, 22(1):494–524, 1994.
- [29] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Ann. Math.*, 171:295–341, 2010.
- [30] Assaf Naor, Yuval Peres, Oded Schramm, and Scott Sheffield. Markov chains in smooth Banach spaces and Gromov-hyperbolic metric spaces. *Duke Math J.*, 134(1):165–197, 2006.
- [31] Ryan O’Donnell. *Analysis of Boolean functions*. Cambridge University Press, 2014.
- [32] Ryan O’Donnell and Karl Wimmer. Approximation by DNF: Examples and counterexamples. In *Automata, Languages and Programming*, volume 4596 of *Lecture Notes in Computer Science*, pages 195–206. Springer Berlin Heidelberg, 2007.
- [33] Krzysztof Oleszkiewicz. On a nonsymmetric version of the Khinchine–Kahane inequality. *Progr. Probab.*, 56:157–168, 2003.
- [34] Prasad Raghavendra and Ben Weitz. On the bit complexity of sum-of-squares proofs. *CoRR*, abs/1702.05139, 2017.
- [35] Bruce E. Sagan. *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*, volume 203 of *Graduate Texts in Mathematics*. Springer New York, 2001.
- [36] Murali K. Srinivasan. Symmetric chains, Gelfand–Tsetlin chains, and the Terwilliger algebra of the binary Hamming scheme. *J. Algebr. Comb.*, 34(2):301–322, 2011.
- [37] Richard P. Stanley. Variations on differential posets. *IMA Vol. Math. Appl.*, 19:145–165, 1990.
- [38] Dennis Stanton. Harmonics on posets. *J. Comb. Theory Ser. A*, 40(1):136–149, 1985.
- [39] Avishay Tal. Tight bounds on the Fourier spectrum of AC^0 . Manuscript, 2017.
- [40] Michel Talagrand. On Russo’s approximate zero-one law. *Ann. Probab.*, 22(3):1576–1587, 1994.
- [41] L. Richard Turner. Inverse of the Vandermonde matrix with applications. Technical Report NASA TN D-3547, Lewis Research Center, NASA, Cleveland, Ohio, August 1966.
- [42] Karl Wimmer. Low influence functions over slices of the Boolean hypercube depend on few coordinates. In *Conference on Computational Complexity (CCC 2014)*, pages 120–131, 2014.