Orthogonal basis for functions over a slice of the Boolean hypercube

Yuval Filmus

Technion — Israel Institute of Technology, Haifa, Israel

January 12, 2016

Abstract

We present a simple, explicit orthogonal basis of eigenvectors for the Johnson and Kneser graphs, based on Young's orthogonal representation of the symmetric group. Our basis can also be viewed as an orthogonal basis for the vector space of all functions over a slice of the Boolean hypercube (a set of the form $\{(x_1, \ldots, x_n) \in \{0,1\}^n : \sum_i x_i = k\}$), which refines the eigenspaces of the Johnson association scheme; our basis is orthogonal with repsect to *any* exchangeable measure. More concretely, our basis is an orthogonal basis for all multilinear polynomials $\mathbb{R}^n \to \mathbb{R}$ which are annihilated by the differential operator $\sum_i \partial/\partial x_i$. As an application of the last point of view, we show how to lift low-degree functions from a slice to the entire Boolean hypercube while maintaining properties such as expectation, variance and L^2-norm.

As an application of our basis, we streamline Wimmer's proof of Friedgut's theorem for the slice. Friedgut's theorem, a fundamental result in the analysis of Boolean functions, states that a Boolean function on the Boolean hypercube with low total influence can be approximated by a Boolean junta (a function depending on a small number of coordinates). Wimmer generalized this result to slices of the Boolean hypercube, working mostly over the symmetric group, and utilizing properties of Young's orthogonal representation. Using our basis, we show how the entire argument can be carried out directly on the slice.

1 Introduction

Functions over the Boolean hypercube $\{0,1\}^n$ are often studied using the tools of *Fourier analysis* (see O'Donnell's excellent recent monograph [19]). The crucial idea is to study functions from the point of view of the *Fourier basis*, an orthonormal basis of functions over the Boolean hypercube. In this work, we consider functions on a different domain, a *slice of the Boolean hypercube* $\binom{[n]}{k} = \{(x_1, \ldots, x_n) \in \{0,1\}^n : \sum_i x_i = k\}$; we always assume that $k \leq n/2$. Such functions arise naturally in coding theory, in the context of constant-weight codes, and have recently started appearing in theoretical computer science as well. In this work we provide an explicit orthogonal basis for the vector space of functions on a slice.

The slice has been studied in algebraic combinatorics under the name *Johnson association scheme*, and in spectral graph theory in relation to the Johnson and Kneser graphs. Our basis is the analog of the Fourier basis for the scheme, and it refines the decomposition induced by the primitive idempotents. Our basis is also an orthogonal basis for the eigenvectors of the Johnson and Kneser graphs, and any other graph belonging to the Bose–Mesner algebra of the Johnson association scheme. Such (weighted) graphs arise in Lovász's proof of the Erdős–Ko–Rado theorem [16], and in Wilson's proof [25] of a t-intersecting version of the theorem.

Despite the name, it is perhaps best to view the slice $\binom{[n]}{k}$ as the set of cosets of $S_k \times S_{n-k}$ inside S_n. This point of view suggests "lifting" an orthogonal basis from the symmetric group to the slice. Following Bannai and Ito [1], the relevant representations of the symmetric group are those corresponding to partitions $(n-d) + d$ for $d \leq k$. Our basis arises from Young's orthogonal representation of the symmetric group. However, we present the basis and prove its properties without reference to the symmetric group at all. One

feature that is inherited from the symmetric group is the lack of a canonical basis: our basis relies on the ordering of the coordinates.

Dunkl [2] showed that the space of functions over the slice $\binom{[n]}{k}$ can be identified with the space of multilinear polynomials in n variables $x_1, \ldots, x_n$ which are annihilated by the differential operator $\sum_{i=1}^{n} \partial/\partial x_i$; the input variable x_i is an indicator variable for the event that i belongs to the input set. Functions annihilated by this operator were termed *harmonic* by Dunkl [4]. Our basis forms an orthogonal basis for the space of harmonic multilinear polynomials for *every* exchangeable measure (a measure invariant under the action of S_n). As a consequence, we show how to lift a low-degree function from the slice $\binom{[n]}{k}$ to the Boolean cube (under an appropriate measure) while maintaining some of its properties such as expectation, variance and L^2-norm.

Wimmer [26] recently generalized a fundamental theorem of Friedgut [9] from the Boolean hypercube to the slice. Friedgut's theorem, sometimes known as Friedgut's junta theorem, states that a Boolean function on the Boolean hypercube with low total influence is close to a Boolean junta (a function depending on a small number of variables). Although Wimmer's main theorem is a statement about functions on the slice, Wimmer lifts the given function to the symmetric group, where most of his argument takes place, exploiting essential properties of Young's orthogonal representation. Eventually, a hypercontractive property of the slice (due to Lee and Yau [15]) is invoked to complete the proof. As an application of our basis, we give a streamlined version of Wimmer's proof in which our basis replaces the appeal to the symmetric group and to Young's orthogonal representation.

Note added in proof   Since writing this paper, we have learned that the same basis has been constructed by Srinivasan [23] in a beautiful paper. Srinivasan in fact constructs an extended basis for the entire Boolean cube $\{0, 1\}^n$, which he identifies with the canonical Gelfand–Tsetlin basis [24], and shows that it is orthogonal with respect to all exchangeable measures. However, he provides neither an explicit description of the basis elements, nor even a canonical indexing scheme for the basis elements. Instead, he gives a recursive algorithm that constructs the basis. We believe that both approaches have merit.

Related work   Apart from Friedgut's theorem, several other classical results in Fourier analysis of Boolean functions have recently been generalized to the slice. O'Donnell and Wimmer [21, 22] generalized the Kahn–Kalai–Linial theorem [12] to the slice, and deduced a robust version of the Kruskal–Katona theorem. Filmus [5] generalized the Friedgut–Kalai–Naor theorem [10] to the slice.

Filmus, Kindler, Mossel and Wimmer [6] and Filmus and Mossel [7] generalized the invariance principle [17] to the slice. The invariance principle on the slice compares the behavior of low-degree harmonic multilinear polynomials on a slice $\binom{[n]}{k}$ and on the Boolean hypercube $\{0, 1\}^n$ with respect to the corresponding product measure $\mu_{k/n}$. If the harmonic multilinear polynomial f has degree d and unit variance, then the invariance principle states that for any Lipschitz functional φ,

$$|\mathop{\mathbb{E}}_{\sigma}[\varphi(f)] - \mathop{\mathbb{E}}_{\mu_p}[\varphi(f)]| = \tilde{O}\left(\sqrt{\frac{d}{\sqrt{p(1-p)n}}}\right),$$

where σ is the uniform distribution on the slice. The invariance principle can be used to lift results such as the Kindler–Safra theorem [13, 14] and Majority is Stablest from the Boolean hypercube to the slice.

Filmus and Mossel also give basis-free proofs for some of the results appearing in this paper. For example, they give a basis-free proof for the fact that the L^2-norm of a low-degree harmonic multilinear polynomial on the slice is similar to its L^2-norm on the Boolean cube under the corresponding product measure.

Synopsis   We describe the space of harmonic multilinear polynomials in Section 2. Our basis is defined in Section 3, in which we also compute the norms of the basis elements. We show that our basis forms a basis for functions on the slice in Section 4, in which we also show how to lift low-degree functions from the slice to the entire hypercube, and explain why our basis is an orthogonal basis of eigenvectors for the Johnson and Kneser graphs. Section 5 and Section 6 are devoted to the proof of the Wimmer–Friedgut theorem.

2

Notation We use the notation $[n] = \{1, \ldots, n\}$. The cardinality of a set S is denoted $|S|$. If $S \subseteq [n]$ and $\pi \in S_n$ (the symmetric group on $[n]$) then $S^\pi = \{\pi(x) : x \in S\}$. We use the same notation in other similar circumstances. We compose permutations from right to left, so $\beta\alpha$ means apply α then β. We use the falling power notation: $n^{\underline{k}} = n(n-1)\cdots(n-k+1)$ (the number of terms is k). For example, $\binom{n}{k} = n^{\underline{k}}/k!$. A function is *Boolean* if its values are in $\{0, 1\}$.

2 Harmonic multilinear polynomials

We construct our basis as a basis for the vector space of *harmonic multilinear polynomials* over $x_1, \ldots, x_n$, a notion defined below. For simplicity, we only consider polynomials over $\mathbb{R}$, but the framework works just as well over any field of characteristic zero.

Definition 2.1. A polynomial $P \in \mathbb{R}[x_1, \ldots, x_n]$ is *multilinear* if all monomials of P are squarefree (not divisible by any x_i^2).

A multilinear polynomial $P \in \mathbb{R}[x_1, \ldots, x_n]$ is *harmonic* if

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

We denote the vector space of harmonic multilinear polynomials over $x_1, \ldots, x_n$ by $\mathcal{H}_n$.

The *degree* of a non-zero multilinear polynomial is the maximal number of variables in any monomial. We denote the subspace of $\mathcal{H}_n$ consisting of polynomials of degree at most d by $\mathcal{H}_{n,d}$.

A polynomial has *pure degree* d if all its monomials have degree d. We denote the subspace of $\mathcal{H}_n$ consisting of polynomials of degree exactly d by $\mathcal{H}'_{n,d}$.

The following lemma calculates the dimension of the vector space of harmonic multilinear polynomials of given degree.

Lemma 2.1. *All polynomials in $\mathcal{H}_n$ have degree at most $n/2$. For $d \leq n/2$,*

$$\dim \mathcal{H}_{n,d} = \binom{n}{d}, \quad \dim \mathcal{H}'_{n,d} = \binom{n}{d} - \binom{n}{d-1},$$

where $\binom{n}{-1} = 0$.

Proof. We start by proving the upper bound on the degree of polynomials in $\mathcal{H}_n$. Let $P \in \mathcal{H}_n$ have degree $\deg P = d$. The pure degree d part of P is also in H_n, and so we can assume without loss of generality that P has pure degree d. For any $\mathbf{y} = y_1, \ldots, y_n$, the univariate polynomial $P(t\mathbf{1} + \mathbf{y})$ (where $\mathbf{1}$ is the constant vector) doesn't depend on t, since

$$\frac{dP(\mathbf{y} + t\mathbf{1})}{dt} = \sum_{i=1}^{n} \frac{\partial P}{\partial x_i}(\mathbf{y} + t\mathbf{1}) = 0.$$

In particular, if M is any monomial in P with coefficient $\alpha \neq 0$ and $\mathbf{y}$ is the vector with $y_i = -1$ whenever x_i appears in M and $y_i = 0$ otherwise, then $P(\mathbf{v} + \mathbf{1}) = P(\mathbf{v}) = (-1)^d \alpha \neq 0$, showing that P must contain

some monomial supported on variables not appearing in M, since $y_i = 1$ only for x_i not appearing in M. In particular, $2d \le n$.

We proceed with the formula for $\mathcal{H}'_{n,d}$; the formula for $\mathcal{H}_{n,d}$ easily follows. When $d = 0$, the formula clearly holds, so assume $d \ge 1$. The vector space of all multilinear polynomial of pure degree d over $x_1, \ldots, x_n$ has dimension $\binom{n}{d}$. Denote by $\mathrm{cf}(P, M)$ the coefficient of the monomial M in P. Harmonicity is the set of conditions

$$\sum_{\substack{i \in [n]: \\ x_i \notin M}} \mathrm{cf}(P, x_i M) = 0 \text{ for all monomials } M \text{ of degree } d-1.$$

There are $\binom{n}{d-1}$ conditions, showing that $\dim \mathcal{H}'_{n,d} \ge \binom{n}{d} - \binom{n}{d-1}$. In order to prove equality, we need to show that the conditions are linearly independent. We do this by showing that there is a polynomial P having pure degree d satisfying all but one of them, that is

$$\sum_{i=1}^{n} \frac{\partial P}{\partial x_i} = x_1 \cdots x_{d-1}.$$

Such a polynomial is given by

$$P = \frac{1}{d} \sum_{t=1}^{d} (-1)^{t+1} \binom{d}{t} \mathop{\mathbb{E}}_{\substack{A \in [d-1]: \, |A| = d-t \\ B \in [n] \setminus [d-1]: \, |B| = t}} x_A x_B,$$

using the notation $x_S = \prod_{i \in S} x_i$. Indeed,

$$\sum_{i=1}^{n} \frac{\partial P}{\partial x_i} = x_1 \cdots x_{d-1} + \frac{1}{d} \sum_{t=2}^{d} \left[t \cdot (-1)^{t+1} \binom{d}{t} + (d-t+1) \cdot (-1)^t \binom{d}{t-1} \right] \mathop{\mathbb{E}}_{\substack{A \in [d-1]: \, |A| = d-t \\ B \in [n] \setminus [d-1]: \, |B| = t-1}} x_A x_B$$

$$= x_1 \cdots x_{d-1}. \qquad \square$$

Frankl and Graham [8] gave a basis for $\mathcal{H}_n$.

Definition 2.2. For $d \le n/2$, a *sequence of length d* is a sequence $S = s_1, \ldots, s_d$ of distinct numbers in $[n]$. The set of all sequence of length d is denoted by $\mathcal{S}_{n,d}$, and the set of all sequences is denoted by $\mathcal{S}_n$.

For any two disjoint sequences $A, B \in \mathcal{S}_{n,d}$ we define

$$\chi_{A,B} = \prod_{i=1}^{d} (x_{a_i} - x_{b_i}).$$

The basis functions will be $\chi_{A,B}$ for appropriate A, B.

Definition 2.3. For $d \le n/2$, let $A, B \in \mathcal{S}_{n,d}$ be disjoint. We say that *A is smaller than B*, written $A < B$, if $a_i < b_i$ for all $i \in [d]$. Similarly, we say that *A is at most B*, written $A \le B$, if $a_i \le b_i$ for all $i \in [d]$.

A sequence $B \in \mathcal{S}_n$ is a *top set* if B is increasing and for some disjoint sequence A of the same length, $A < B$. The set of top sets of length d is denoted by $\mathcal{B}_{n,d}$, and the set of all top sets is denoted by $\mathcal{B}_n$.

The following lemma is mentioned without proof in [8].

Lemma 2.2. For $0 \le d \le n/2$, $|\mathcal{B}_{n,d}| = \binom{n}{d} - \binom{n}{d-1}$, where $\binom{n}{-1} = 0$.

Proof. We encode each sequence $B \in \mathcal{S}_{n,d}$ as a ± 1 sequence $\beta_0, \ldots, \beta_d$ as follows. We put $\beta_0 = 1$, and for $i \in [d]$, $\beta_i = 1$ if $i \notin B$ and $\beta_i = -1$ if $i \in B$. It is not hard to check that B is a top set iff all running sums of β are positive. Each sequence β is composed of d entries -1 and $n - d + 1$ entries 1. The probability that such a sequence has all running sums positive is given by the solution to Bertrand's ballot problem: it is $\frac{(n-d+1)-d}{n+1}$. Therefore the total number of top sets is

$$\frac{(n-d+1)-d}{n+1} \binom{n+1}{d} = \frac{n-d+1}{n+1} \binom{n+1}{n-d+1} - \frac{d}{n+1} \binom{n+1}{d} = \binom{n}{n-d} - \binom{n}{d-1}. \qquad \square$$

We can now give Frankl and Graham's basis, which is given in [8] without proof.

Lemma 2.3. *For each $B \in \mathcal{B}_n$, let $\phi(B) \in \mathcal{S}_{n,|B|}$ be any sequence satisfying $\phi(B) < B$. The set $\{\chi_{\phi(B),B} : B \in \mathcal{B}_n\}$ is a basis for $\mathcal{H}_n$. Moreover, for $d \leq n/2$, the set $\{\chi_{\phi(B),B} : B \in \mathcal{B}_{n,d}\}$ is a basis for $\mathcal{H}'_{n,d}$.*

Proof. It is clearly enough to prove that $X_d = \{\chi_{\phi(B),B} : B \in \mathcal{B}_{n,d}\}$ is a basis for $\mathcal{H}'_{n,d}$. In view of Lemma 2.1 and Lemma 2.2, it is enough to prove that all functions in X_d belong to $\mathcal{H}'_{n,d}$, and that X_d is linearly independent. Clearly all functions in X_d are multilinear polynomial of pure degree d. To show that they are harmonic, notice that if $A < B$, where $|A| = |B| = d$, then

$$\sum_{i=1}^{n} \frac{\partial \chi_{A,B}}{\partial x_i} = \sum_{j=1}^{d} \left(\frac{\partial \chi_{A,B}}{\partial x_{a_j}} + \frac{\partial \chi_{A,B}}{\partial x_{b_j}} \right) = \sum_{j=1}^{d}(1-1)\frac{\chi_{A,B}}{x_{a_j} - x_{b_j}} = 0.$$

It remains to prove that X_d is linearly independent. For an increasing sequence $S \in \mathcal{S}_{n,d}$, let $\Pi(S)$ be the monomial $\Pi(S) = \prod_{i=1}^{d} x_{s_i}$. If $\Pi(S)$ appears in $\chi_{\phi(B),B}$ then $S \leq B$. Consider now the matrix representing X_d in the basis $\{\Pi(S) : S \in \mathcal{S}_{n,d}\}$ arranged in an order compatible with the partial order of $\mathcal{S}_{n,d}$. The resulting matrix is in echelon form and so has full rank, showing that X_d is a linearly independent set. $\square$

We comment that $\dim \mathcal{H}'_{n,d} = \binom{n}{d} - \binom{n}{d-1}$ is the dimension of the irreducible representation of S_n corresponding to the partition $(n-d) + d$, as easily calculated using the hook formula. This is not a coincidence: indeed, the set of standard Young tableaux of shape $(n-d), d$ is in bijection with $\mathcal{B}_{n,d}$ by identifying the top sets with the contents of the second row of each such tableau (the tableau can be completed uniquely).

3 Young's orthogonal basis

In this section we will construct an orthogonal basis for $\mathcal{H}_n$, and calculate the norms of the basis elements. Our basis will be orthogonal with respect to a wide class of measures.

Definition 3.1. A probability distribution over the variables $x_1, \ldots, x_n$ is *exchangeable* if it is invariant under permutations of the indices. Given an exchangeable distribution μ, we define an inner product on $\mathcal{H}_n$ by

$$\langle f, g \rangle = \mathbb{E}_{\mu}[fg].$$

The norm of $f \in \mathcal{H}_n$ is $\|f\| = \sqrt{\langle f, f \rangle}$.

We are now ready to define the basis.

Definition 3.2. For $B \in \mathcal{B}_{n,d}$, define

$$\chi_B = \sum_{\substack{A \in \mathcal{S}_{n,d}: \\ A < B}} \chi_{A,B}.$$

For $d \leq n/2$, we define

$$\chi_d = \chi_{2,4,\ldots,2d} = \chi_{1,3,\ldots,2d-1;2,4,\ldots,2d} = \prod_{i=1}^{d}(x_{2i-1} - x_{2i}).$$

Young's orthogonal basis for $\mathcal{H}_n$ is

$$\mathcal{Y}_n = \{\chi_B : B \in \mathcal{B}_n\}.$$

Young's orthogonal basis for $\mathcal{H}'_{n,d}$ is

$$\mathcal{Y}_{n,d} = \{\chi_B : B \in \mathcal{B}_{n,d}\}.$$

We stress that the sequences A in the definition of χ_B need not be increasing.

The following theorem justifies the name "orthogonal basis".

Theorem 3.1. *The set $\mathcal{Y}_n$ is an orthogonal basis for $\mathcal{H}_n$ with respect to any exchangeable measure. The set $\mathcal{Y}_{n,d}$ is an orthogonal basis for $\mathcal{H}'_{n,d}$ with respect to any exchangeable mesure. In particular, the subspaces $\mathcal{H}'_{n,d}$ for different d are mutually orthogonal.*

Proof. Lemma 2.3 shows that each $\chi_B \in \mathcal{Y}_{n,d}$ lies in $\mathcal{H}'_{n,d}$. The technique used in proving the lemma shows that $\mathcal{Y}_{n,d}$ is a basis for $\mathcal{H}'_{n,d}$, but this will also follow from Lemma 2.1 once we prove that the functions in $\mathcal{Y}_{n,d}$ are pairwise orthogonal. In fact, we will prove the following more general claim: if $B_1, B_2 \in \mathcal{B}_n$ and $B_1 \neq B_2$ then χ_{B_1}, χ_{B_2} are orthogonal. This will complete the proof of the theorem.

Consider any $B_1 \in \mathcal{B}_{n,d_1}$ and $B_2 \in \mathcal{B}_{n,d_2}$, where $B_1 \neq B_2$. Our goal is to prove that $\langle \chi_{B_1}, \chi_{B_2} \rangle = 0$. We have

$$\langle \chi_{B_1}, \chi_{B_2} \rangle = \sum_{\substack{A_1 < B_1 \\ A_2 < B_2}} \mathbb{E}[\chi_{A_1, B_1} \chi_{A_2, B_2}].$$

We call each of the terms $\chi_{A_1, B_1} \chi_{A_2, B_2}$ appearing in this expression a *quadratic product*. We will construct a sign-flipping involution among the quadratic products, completing the proof. The involution is also allowed to have fixed points; in this case, the expectation of the corresponding quadratic product vanishes.

Consider a quadratic product

$$\chi_{A_1, B_1} \chi_{A_2, B_2} = \prod_{i=1}^{d_1} (x_{a_{1,i}} - x_{b_{1,i}}) \prod_{j=1}^{d_2} (x_{a_{2,j}} - x_{b_{2,j}}).$$

We can represent this quadratic product as a directed graph G on the vertex set $A_1 \cup A_2 \cup B_1 \cup B_2$. For each factor $x_i - x_j$ in the quadratic product, we draw an edge from i to j; all edges point in the direction of the larger vertex (the vertex having a larger index). We further annotate each edge with either 1 or 2, according to which of $\chi_{A_1, B_1}, \chi_{A_2, B_2}$ it corresponds to. Every variable x_i appears in at most two factors, and so the total degree of each vertex is at most 2. Therefore the graph decomposes as an undirected graph into a disjoint union of paths and cycles. The annotations on the edges alternate on each connected component. Every directed graph in which edges point in the direction of the larger vertex, the total degree of each vertex is at most 2, and the annotations on the edges alternate in each connected component, is the graph corresponding to some quadratic product $\chi_{A'_1, B'_1} \chi_{A'_2, B'_2}$. The value of A'_1, B'_1, A'_2, B'_2 can be read using the annotations on the edges.

Since $B_1 \neq B_2$, some connected component must have a vertex with in-degree 1. Choose the connected component C satisfying this property having the largest vertex. We construct a sequence of intervals inside C, with the property that the endpoints of each interval are either endpoints of C or are connected to the rest of C with outgoing edges. Furthermore, each interval, other than possibly the last one, contains some vertex with in-degree 1. The sequence terminates with an interval containing an odd number of edges.

When C is a path, the first interval I_0 is the entire path. When C is a cycle with maximal vertex M, the first interval I_0 is the path obtained by removing M from C. Given an interval I_t with an even number of edges, we can break it into two (possibly empty) subintervals terminating at the maximal point M_t of I_t: $I_t = J_t \to M_t \leftarrow K_t$. Note that not both J_t, K_t can be empty since I_t contains some vertex with in-degree 1. If J_t is empty then we define $I_{t+1} = K_t$, which terminates the sequence. Similarly, if K_t is empty then we define $I_{t+1} = J_t$, which terminates the sequence. If both J_t, K_t are non-empty then at least one of them has a vertex with in-degree 1. We let I_{t+1} be the sub-interval among J_t, K_t with the larger maximal point.

Since the intervals decrease in size, the sequence eventually terminates at some interval I_t having an odd number of edges. Let C' be the component obtained by reversing I_t, and let G' be the graph obtained by replacing C with C'. We claim that G' is the graph of some quadratic product $\chi_{A'_1, B_1} \chi_{A'_2, B_2}$ (which could be the same as the original quadratic product). Indeed, the property of the endpoints of I_t guarantees that in G', edges always point in the direction of the larger vertex. Since I_t has an odd number of edges, the annotations on the edges still alternate in C'. Clearly every vertex has total degree at most 2, and so G'

is the graph of some quadratic product. Furthermore, the annotations of the edges pointing at each vertex are the same, again using the property of the endpoints of I_t. Therefore G' is indeed the graph of some quadratic product $\chi_{A'_1,B_1}\chi_{A'_2,B_2}$. Moreover, if we apply the same algorithm on the graph G' then we will obtain the graph G, so the algorithm is an involution on the quadratic products $\chi_{\cdot,B_1}\chi_{\cdot,B_2}$, possibly having fixed points.

Let $v_1,\ldots,v_\ell$ be the vertices of I_t, and consider the permutation π which maps v_i to $v_{\ell+1-i}$ and fixes all other points in $[n]$. If we apply π on G and flip all edges in I_t, then we obtain the graph G'. Since the measure is exchangeable and the number of edges in I_t is odd, this shows that $\mathbb{E}[\chi_{A_1,B_1}\chi_{A_2,B_2}] = -\mathbb{E}[\chi_{A'_1,B_1}\chi_{A'_2,B_2}]$. We conclude that the mapping $\chi_{A_1,B_1}\chi_{A_2,B_2} \Rightarrow \chi_{A'_1,B_1}\chi_{A'_2,B_2}$ is a sign-flipping involution, completing the proof. $\qquad\square$

In order to complete the picture, we need to evaluate the norms of the basis elements χ_B, which necessarily depend on the measure.

Theorem 3.2. *Let $B \in \mathcal{B}_{n,d}$. The squared norm of χ_B is $\|\chi_B\|^2 = c_B\|\chi_d\|^2$ with respect to any exchangeable measure, where*

$$c_B = \prod_{i=1}^{d} \frac{(b_i - 2(i-1))(b_i - 2(i-1) - 1)}{2}.$$

(Recall that $\chi_d = \chi_{2,4,\ldots,2d}$.)

Proof. We consider first the case in which the exchangeable measure is the measure ν_p for some $p \in [0,1]$. Under this measure, the variables $x_1,\ldots,x_n$ are independent, with $\Pr[x_i = -p] = 1 - p$ and $\Pr[x_i = 1 - p] = p$. The expectation of each x_i is $\mathbb{E}[x_i] = (1-p)(-p) + p(1-p) = 0$, while the variance is $\mathbb{E}[x_i^2] = (1-p)p^2 + p(1-p)^2 = p(1-p)$. The squared norm of χ_B is

$$\|\chi_B\|^2 = \langle \chi_B, \chi_B \rangle = \sum_{\substack{A_1,A_2\in\mathcal{S}_d:\\ A_1,A_2<B}} \mathbb{E}[\chi_{A_1,B}\chi_{A_2,B}].$$

In the proof of Theorem 3.1 we associated a directed graph with each *quadratic product* $\chi_{A_1,B}\chi_{A_2,B}$: the vertices are $A_1 \cup A_2 \cup B$, and the edges point from $a_{1,i}$ and $a_{2,i}$ to b_i for each $i \in [d]$, annotated by 1 or 2 according to whether they came from $\chi_{A_1,B}$ or from $\chi_{A_2,B}$. Since each vertex appears at most twice, the graph decomposes as a sum of paths and cycles. The edges point from A_1, A_2 to B, and so each vertex either has in-degree 0 (if it is in $A_1 \cup A_2$) or in-degree 2 (if it is in B). Therefore the paths and cycles have the following forms, respectively:

$$\alpha_1 \to \beta_1 \leftarrow \alpha_2 \to \beta_2 \leftarrow \cdots \leftarrow \alpha_\ell \to \beta_\ell \leftarrow \alpha_{\ell+1},$$
$$\alpha_1 \to \beta_1 \leftarrow \alpha_2 \to \beta_2 \leftarrow \cdots \leftarrow \alpha_\ell \to \beta_\ell \leftarrow \alpha_1.$$

Here the α_i belong to $A_1 \cup A_2$, and the β_i belong to B. The corresponding factors of $\chi_{A_1,B}\chi_{A_2,B}$ are, respectively:

$$(x_{\alpha_1} - x_{\beta_1})(x_{\alpha_2} - x_{\beta_1})(x_{\alpha_2} - x_{\beta_2})\cdots(x_{\alpha_\ell} - x_{\beta_\ell})(x_{\alpha_{\ell+1}} - x_{\beta_\ell}),$$
$$(x_{\alpha_1} - x_{\beta_1})(x_{\alpha_2} - x_{\beta_1})(x_{\alpha_2} - x_{\beta_2})\cdots(x_{\alpha_\ell} - x_{\beta_\ell})(x_{\alpha_1} - x_{\beta_\ell}).$$

We proceed to calculate the expectation of each of these factors under ν_p. The expectation of a monomial is zero unless each variable appears exactly twice, in which case the expectation is $(p(1-p))^\ell$ (since each monomial has total degree 2ℓ). In the case of a path, there is exactly one such monomial, namely $x_{\beta_1}^2 x_{\beta_2}^2 \cdots x_{\beta_\ell}^2$. In the case of a cycle, there are two such monomials: $x_{\beta_1}^2 x_{\beta_2}^2 \cdots x_{\beta_\ell}^2$ and $x_{\alpha_1}^2 x_{\alpha_2}^2 \cdots x_{\alpha_\ell}^2$. Both monomials appear with unit coefficient. Notice that ℓ is the size of the subset of B appearing in the path or cycle. Hence the expectation of the entire quadratic product is $2^C(p(1-p))^d$, where $C = C(A_1, A_2)$ is the number of cycles. In total, we get

$$\|\chi_B\|_{\nu_p}^2 = \sum_{\substack{A_1,A_2\in\mathcal{S}_d:\\ A_1,A_2<B}} 2^{C(A_1,A_2)}(p(1-p))^d.$$

We proceed to show that

$$\sum_{\substack{A_1,A_2 \in \mathcal{S}_d: \\ A_1,A_2 < B}} 2^{C(A_1,A_2)} = 2^d c_B = \prod_{i=1}^{d}(b_i - 2(i-1))(b_i - 2(i-1) - 1).$$

The quantity on the right enumerates the sequences $\alpha_{1,1}, \alpha_{2,1}, \alpha_{1,2}, \alpha_{2,2}, \ldots, \alpha_{1,d}, \alpha_{2,d} \in \mathcal{S}_{n,2d}$ in which $\alpha_{1,i}, \alpha_{2,i} \le b_i$ for all $i \in [d]$, which we call *legal sequences*. We show how to map legal sequences into quadratic products in such a way that $\chi_{A_1,B}\chi_{A_2,B}$ has exactly $2^{C(A_1,A_2)}$ preimages.

Let $\alpha_{1,1}, \alpha_{2,1}, \alpha_{1,2}, \alpha_{2,2}, \ldots, \alpha_{1,d}, \alpha_{2,d} \in \mathcal{S}_{n,2d}$ be a legal sequence. We construct a quadratic product $\chi_{A_1,B}\chi_{A_2,B}$ alongside its associated directed graph. We maintain the following invariant: after having processed $\alpha_{1,i}, \alpha_{2,i}$ (and so adding b_i to the graph), all vertices belonging to cycles have appeared in the sequence, and out of each path, exactly one vertex (from B) has not appeared previously in the sequence. Furthermore, the endpoints of each path do not belong to B (and so have incoming edges) and have different annotations.

We start with the empty product and graph. At step i we process $\alpha_{1,i}, \alpha_{2,i}$. Suppose first that $\alpha_{1,i}, \alpha_{2,i} \ne b_i$. If $\alpha_{1,i} \notin B$, then we set $a_{1,i} = \alpha_i$ and add the edge $\alpha_i \xrightarrow{1} b_i$. If $\alpha_{1,i} \in B$ then necessarily $\alpha_{1,i} < b_i$, and so it appears in some component C. We locate the endpoint x whose incoming edge is labelled 2, set $a_{1,i} = x$ and add the edge $x \xrightarrow{1} b_i$. Note that the other endpoint of C is labelled 1. We do the same for $\alpha_{2,i}$. It is routine to check that we have maintained the invariant.

Suppose next that one of $\alpha_{1,i}, \alpha_{2,i}$ equals b_i, say $\alpha_{2,i} = b_i$. We process $\alpha_{1,i}$ as in the preceding step. Let x be the other endpoint of the path containing b_i. We set $a_{2,i} = x$ and connect x to b_i, completing the cycle. This completes the description of the mapping.

We proceed to describe the multivalued inverse mapping, from a quadratic product to a legal sequence. We process the quadratic product in d steps, updating the graph by removing each vertex mentioned in the legal sequence. We maintain the invariant that each original path remains a path, and each original cycle either remains a cycle or disappears after processing the largest vertex. Furthermore, each edge still points at the larger vertex, the annotations alternate, a vertex b_i not yet processed has two incoming edges, and other vertices have no incoming edges.

At step i, we process b_i. Let x_1, x_2 be the neighbors of b_i labelled $1, 2$, respectively. Suppose first that $x_1 \ne x_2$. We put $\alpha_{1,i} = x_1$ and $\alpha_{2,i} = x_2$, and remove the vertices x_1, x_2. The other neighbors of x_1, x_2, if any, are connected to b_i with edges pointing away from b_i with annotations $2, 1$, respectively. These neighbors had incoming edges and so are b_j, b_k for $j, k > i$. It follows that the invariant is maintained. The case $x_1 = x_2$ corresponds to a cycle whose largest vertex is b_i. We either put $\alpha_{1,i} = x_1$ and $\alpha_{2,i} = b_i$ or $\alpha_{1,i} = b_i$ and $\alpha_{2,i} = x_2$, deleting the entire cycle in both cases.

It is routine to check that the two mappings we have described are inverses. Furthermore, the multivalued mapping from quadratic products to legal sequences has valency $2^{C(A_1,A_2)}$ when processing $\chi_{A,B_1}\chi_{A,B_2}$. This completes the proof of the formula for $2^d c_B$.

Having considered the measure ν_p, we consider a related measure μ_p. Under this measure the x_i are independent, $\Pr[x_i = 0] = 1 - p$, and $\Pr[x_1 = 1] = p$. Note that $(x_1 - p, \ldots, x_n - p) \sim \nu_p$. Since $(x_i - p) - (x_j - p) = x_i - x_j$, we conclude that

$$\|\chi_B\|_{\mu_p}^2 = \|\chi_B\|_{\nu_p}^2 = (2p(1-p))^d c_B.$$

Consider now a general exchangeable measure m. Exchangeability implies that for some integers $\gamma_0, \ldots, \gamma_d$,

$$\|\chi_B\|_m^2 = \sum_{k=0}^{d} \gamma_k \mathop{\mathbb{E}}_m\left[\prod_{i=1}^{d-k} x_i^2 \prod_{i=1}^{2k} x_{d-k+i}\right],$$

since all monomials in $\mathbb{E}[\chi_B^2]$ have total degree $2d$. Substituting the measure μ_p, which satisfies $\mathbb{E}[x_i] = \mathbb{E}[x_i^2] = p$, we obtain

$$\sum_{k=0}^{d} \gamma_k p^{d+k} = \|\chi_B\|_{\mu_p}^2 = (2p(1-p))^d c_B.$$

8

Reading off the coefficient of γ_k, we deduce

$$\gamma_k = (-1)^k \binom{d}{k} 2^d c_B,$$

and so

$$\|\chi_B\|_m^2 = 2^d c_B N_d, \quad N_d = \sum_{k=0}^{d} (-1)^k \binom{d}{k} \underset{m}{\mathbb{E}} \left[\prod_{i=1}^{d-k} x_i^2 \prod_{i=1}^{2k} x_{d-k+i} \right].$$

In order to evaluate N_d, we consider $B = 2, 4, \dots, 2d$. In this case, $b_i = 2i$ and so

$$c_B = \prod_{i=1}^{d} \frac{(b_i - 2(i-1))(b_i - 2(i-1) - 1)}{2} = 1.$$

We conclude that $2^d N_d = \|\chi_d\|^2$, and the theorem follows. $\qquad \square$

Having verified that $\mathcal{Y}_n$ is a basis for $\mathcal{H}_n$, we can define the corresponding expansion.

Definition 3.3. Let $f \in \mathcal{H}_n$. The *Young–Fourier expansion* of f is the unique representation

$$f = \sum_{B \in \mathcal{B}_n} \hat{f}(B) \chi_B.$$

The following simple lemma gives standard properties of this expansion.

Lemma 3.3. *Let $f \in \mathcal{H}_n$. The following hold with respect to any exchangeable measure. For each $B \in \mathcal{B}_n$, we have $\hat{f}(B) = \langle f, \chi_B \rangle / \|\chi_B\|^2$. The mean, variance and L2 norm of f are given by*

$$\mathbb{E}[f] = \hat{f}(\emptyset), \quad \mathbb{V}[f] = \sum_{\substack{B \in \mathcal{B}_n: \\ B \neq \emptyset}} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2, \quad \mathbb{E}[f^2] = \sum_{B \in \mathcal{B}_n} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2,$$

where $\emptyset$ is the empty sequence.

In particular, if $f \in \mathcal{H}_n$ then for every exchangeable measure π we can write $\mathbb{E}[f^2] = \sum_d \|\chi_d\|^2 \mathbf{W}^d[f]$, where $\mathbf{W}^d[f] = \sum_{|B|=d} c_B \hat{f}(B)^2$ depends only on f. Filmus and Mossel [7] give a basis-free proof of this important fact.

The familiar Fourier basis for the Boolean hypercube gives a simple criterion for a function to depend on a variable. The matching criterion in our case is also simple but not as powerful.

Lemma 3.4. *Let $f \in \mathcal{H}_n$, and for $m \in [n]$ let $g = \mathbb{E}_{\pi \in S_m}[f^\pi]$ be the result of averaging f over all permutations of the first m coordinates. Then*

$$g = \sum_{\substack{B \in \mathcal{B}_n: \\ B \cap [m] = \emptyset}} \hat{f}(B) \chi_B.$$

In particular, f is invariant under permutation of the first m coordinates iff $\hat{f}(B) = 0$ whenever B intersects $[m]$.

Proof. It is enough to prove the formula for basis functions $f \in \mathcal{Y}_n$. If $f = \chi_B$ where $B \cap [m] = \emptyset$ then $f^\pi = f$ for all $\pi \in S_m$ since $A < B$ iff $A^\pi < B$, and so $g = f$. Suppose next that $f = \chi_B$ where B intersects $[m]$, say $b_j \in [m]$. For $a < b_j$ not belonging to B, define

$$\chi_{a,B} = \sum_{\substack{A \in \mathcal{S}_{n,|B|}: \\ A < B, a_j = a}} \chi_{A,B}.$$

Since $\chi_B = \sum_{a \in [b_j] \setminus B} \chi_{a,B}$, it is enough to show that $\mathbb{E}_{\pi \in S_m}[\chi_{a,B}^\pi]$ vanishes. Since the mapping $\pi \mapsto (a \, b_j)\pi$ is an involution on S_m, it is enough to notice that $\mathbb{E}[\chi_{a,B}^{(a \, b_j)\pi}] = -\mathbb{E}[\chi_{a,B}^\pi]$. $\qquad \square$

4   Slices of the Boolean hypercube

Harmonic multilinear polynomials appear naturally in the context of *slices of the Boolean hypercube.*

Definition 4.1. Let n be an integer and let $k \leq n/2$ be an integer. The (n, k) *slice of the Boolean hypercube* is

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

We also identify $\binom{[n]}{k}$ with subsets of $[n]$ of cardinality k. We endow the slice $\binom{[n]}{k}$ with the uniform measure, which is clearly exchangeable.

A *function over the slice* is a function $f \colon \binom{[n]}{k} \to \mathbb{R}$. Every function $f \colon \mathbb{R}^n \to \mathbb{R}$ can be interpreted as a function on the slice in the natural way.

We proceed to show that $\mathcal{Y}_{n,0} \cup \cdots \cup \mathcal{Y}_{n,k}$ is an orthogonal basis for the slice $\binom{[n]}{k}$.

Theorem 4.1. *Let n and $k \leq n/2$ be integers, and put $p = k/n$. The set $\{\chi_B : B \in \mathcal{B}_{n,d} \text{ for some } d \leq k\}$ is an orthogonal basis for the vector space of functions on the slice, and for $B \in \mathcal{B}_{n,d}$,*

$$\|\chi_B\|^2 = c_B 2^d \frac{k^{\underline{d}}(n-k)^{\underline{d}}}{n^{\underline{2d}}} = c_B (2p(1-p))^d \left(1 \pm O_p\left(\frac{d^2}{n}\right)\right).$$

Proof. We prove the formula and estimate for $\|\chi_B\|^2$ below. When $d \leq k$, the formula shows that $\|\chi_B\|^2 \neq 0$ and so $\chi_B \neq 0$ as a function on the slice. Hence the set $\{\chi_B : B \in \mathcal{B}_{n,d} \text{ for some } d \leq k\}$ consists of non-zero mutually orthogonal vectors. Lemma 2.2 shows that the number of vectors in this set is $\binom{n}{k}$, matching the dimension of the vector space of functions on the slice. Hence this set constitutes a basis for the vector space of functions on the slice.

In order to prove the formula for $\|\chi_B\|^2$, we need to compute

$$\|\chi_d\|^2 = \mathbb{E}\left[\prod_{i=1}^d (x_{2i-1} - x_{2i})^2\right].$$

The quantity $\prod_{i=1}^d (x_{2i-1} - x_{2i})^2$ is non-zero for a subset $S \in \binom{[n]}{k}$ if S contains exactly one out of each pair $\{x_{2i-1}, x_{2i}\}$, in which case the quantity has value 1. Therefore

$$\|\chi_d\|^2 = \Pr_{S \in \binom{[n]}{k}}[|S \cap \{x_{2i-1}, x_{2i}\}| = 1 \text{ for all } i \in [d]]$$

$$= 2^d \Pr_{S \in \binom{[n]}{k}}[x_1, x_3, \ldots, x_{2d-1} \in S \text{ and } x_2, x_4, \ldots, x_{2d} \notin S]$$

$$= 2^d \frac{\binom{n-2d}{k-d}}{\binom{n}{k}} = 2^d \frac{(n-2d)!k!(n-k)!}{(k-d)!(n-k-d)!n!} = 2^d \frac{k^{\underline{d}}(n-k)^{\underline{d}}}{n^{\underline{2d}}}.$$

This yields the formula for $\|\chi_B\|^2$. We can estimate this expression as follows:

$$2^d \frac{k^{\underline{d}}(n-k)^{\underline{d}}}{n^{\underline{d}}} = (2p(1-p))^d \frac{\left(1 - \frac{1}{k}\right)\left(1 - \frac{2}{k}\right) \cdots \left(1 - \frac{d-1}{k}\right) \cdot \left(1 - \frac{1}{n-k}\right)\left(1 - \frac{2}{n-k}\right) \cdots \left(1 - \frac{d-1}{n-k}\right)}{\left(1 - \frac{1}{n}\right)\left(1 - \frac{2}{n}\right) \cdots \left(1 - \frac{2d-1}{n}\right)}$$

$$= (2p(1-p))^d \frac{\left(1 - \frac{O(d^2)}{k}\right)\left(1 - \frac{O(d^2)}{n-k}\right)}{\left(1 - \frac{O(d^2)}{n}\right)},$$

implying the stated estimate. $\qquad\square$

The theorem shows that every function on $\binom{[n]}{k}$ can be represented uniquely as a harmonic multilinear polynomial of degree at most k (see Filmus and Mossel [7] for a basis-free proof). The quantity $(2p(1-p))^d$ is the squared norm of χ_d under the μ_p measure (defined below) over the entire Boolean hypercube, and this allows us to lift low-degree functions from the slice to the entire Boolean hypercube while preserving properties such as the expectation and variance.

Theorem 4.2. *Let n and $k \leq n/2$ be integers, and put $p = k/n$. Let $f \colon \binom{[n]}{k} \to \mathbb{R}$ be a function whose representation as a harmonic multilinear polynomial has degree d, and define $\tilde{f} \colon \{0,1\}^{[n]} \to \mathbb{R}$ by interpreting this polynomial over $\{0,1\}^{[n]}$.*

If we endow $\{0,1\}^{[n]}$ with the measure μ_p given by $\mu_p(x_1, \ldots, x_n) = p^{\sum_{i=1}^n x_i}(1-p)^{\sum_{i=1}^n (1-x_i)}$ then

$$\mathbb{E}[\tilde{f}] = \mathbb{E}[f], \qquad \|\tilde{f}\|^2 = \|f\|^2 \left(1 \pm O_p\left(\frac{d^2}{n}\right)\right), \qquad \mathbb{V}[\tilde{f}] = \mathbb{V}[f]\left(1 \pm O_p\left(\frac{d^2}{n}\right)\right).$$

Proof. The definition of $\tilde{f}$ directly implies that for all $B \in \mathcal{B}_n$, $\hat{\tilde{f}}(B) = \hat{f}(B)$, and furthermore $\hat{f}(B) = 0$ for $|B| > d$. According to Lemma 3.3, $\mathbb{E}[\tilde{f}] = \hat{\tilde{f}}(\emptyset) = \hat{f}(\emptyset) = \mathbb{E}[f]$. In order to prove the estimate on the norms, we compute the squared norm of χ_ℓ with respect to μ_p:

$$\|\chi_\ell\|_{\mu_p}^2 = \mathop{\mathbb{E}}_{\mu_p}\left[\prod_{i=1}^\ell (x_{2i-1} - x_{2i})^2\right] = (2p(1-p))^\ell.$$

Denoting by σ the uniform measure on $\binom{[n]}{k}$, Lemma 3.3 and Theorem 4.1 imply that

$$\|\tilde{f}\|^2 = \sum_{B \in \mathcal{B}_n} \hat{\tilde{f}}(B)\|\chi_{|B|}\|_{\mu_p}^2 = \sum_{B \in \mathcal{B}_n} \hat{f}(B)(2p(1-p))^{|B|},$$

$$\|f\|^2 = \sum_{B \in \mathcal{B}_n} \hat{f}(B)\|\chi_{|B|}\|_\sigma^2 = \sum_{B \in \mathcal{B}_n} \hat{f}(B)(2p(1-p))^{|B|}\left(1 \pm O_p\left(\frac{d^2}{n}\right)\right),$$

implying the estimate on the norms. The estimate on the variance follows analogously. $\square$

Johnson association scheme The Johnson association scheme is an association scheme whose underlying set is $\binom{[n]}{k}$. Instead of describing the scheme itself, we equivalently describe its Bose–Mesner algebra.

Definition 4.2. Let n, k be integers such that $k \leq n/2$. A square matrix M indexed by $\binom{[n]}{k}$ belongs to the *Bose–Mesner algebra* of the (n, k) *Johnson association scheme* if $M_{S,T}$ depends only on $|S \cap T|$.

While it is not immediately obvious, the Bose–Mesner algebra is indeed an algebra of matrices, that is, it is closed under multiplication. Furthermore, it is a commutative algebra, and so all matrices have common eigenspaces. In particular, the algebra is spanned by a basis of primitive idempotents $J_0, \ldots, J_k$. As the following lemma shows, these idempotents correspond to the bases $\mathcal{Y}_{n,0}, \ldots, \mathcal{Y}_{n,k}$.

Lemma 4.3. *Let n, k be integers such that $k \leq n/2$, and let M belong to the Bose–Mesner algebra of the (n, k) Johnson association scheme. The bases $\mathcal{Y}_{n,0}, \ldots, \mathcal{Y}_{n,k}$ span the eigenspaces of M.*

Proof. Since all matrices in the Bose–Mesner algebra have the same eigenspaces, it is enough to consider a particular matrix in the algebra which has $k + 1$ distinct eigenvalues. Let M be the matrix corresponding to the linear operator

$$M \colon f \mapsto \sum_{1 \leq i < j \leq n} f^{(i\ j)}.$$

More explicitly, it is not hard to calculate that $M(S, S) = \binom{k}{2} + \binom{n-k}{2}$, $M(S, T) = 1$ if $|S \cap T| = k - 1$, and $M(S, T) = 0$ otherwise. In particular, M belongs to the Bose–Mesner algebra. Lemma 5.5, which we prove in Section 5, shows that for $d \leq k$, the subspace spanned by $\mathcal{Y}_{n,d}$ is an eigenspace of M corresponding to the eigenvalue $\binom{n}{2} - d(n + 1 - d)$. All eigenvalues are distinct (since the maximum of the parabola $d(n + 1 - d)$ is at $d = (n + 1)/2$), completing the proof. $\square$

As an immediate corollary, we deduce that $\mathcal{Y}_d$ is an orthogonal basis for the dth eigenspace in the (n,k) Johnson graph and (n,k) Kneser graph, as well as any other graph on $\binom{[n]}{k}$ in which the weight on an edge (S,T) depends only on $|S \cap T|$. In the Johnson graph, two sets are connected if their intersection has size $k-1$, and in the Kneser graph, two sets are connected if they are disjoint.

The slice $\binom{[n]}{k}$ can be identified as the set of cosets of $S_k \times S_{n-k}$ inside S_n. Bannai and Ito [1], following Dunkl [3, 4], use this approach to determine the idempotents of the Johnson association scheme from the representations of the symmetric group S_n. They obtain the idempotent J_d from the representation corresponding to the partition $(n-d)+d$. The basis $\mathcal{Y}_{n,d}$ can be derived from Young's orthogonal representation corresponding to the partition $(n-d)+d$, but we do not develop this connection here; such a construction appears in Srinivasan [23].

5 Influences

Throughout this section, we fix some arbitrary exchangeable measure. All inner products and norms are with respect to this measure.

One of the most important quantities arising in the analysis of functions on the hypercube is the influence. In this classical case, influence is defined with respect to a single coordinate. In our case, the basic quantity is the influence of a *pair* of coordinates.

Definition 5.1. Let $f \in \mathcal{H}_n$. For $i, j \in [n]$, define a function $f^{(i\,j)} \in \mathcal{H}_n$ by $f^{(i\,j)}(x) = f(x^{(i\,j)})$, where $x^{(i\,j)}$ is obtained by switching x_i and x_j. The *influence* of the pair (i,j) is

$$\mathrm{Inf}_{ij}[f] = \tfrac{1}{2}\|f^{(i\,j)} - f\|^2.$$

The mth *total influence* of the function is

$$\mathrm{Inf}^m[f] = \frac{1}{m} \sum_{1 \le i < j \le m} \mathrm{Inf}_{ij}[f].$$

When $m = n$, we call the resulting quantity the *total influence*, denoted $\mathrm{Inf}[f]$.

We start with a triangle inequality for influences (cf. [26, Lemma 5.4] for the Boolean case, in which the constant $\frac{9}{2}$ can be improved to $\frac{3}{2}$).

Lemma 5.1. *Let $f \in \mathcal{H}_n$. For distinct $i, j, k \in [n]$ we have*

$$\mathrm{Inf}_{ij}[f] \le \tfrac{9}{2}(\mathrm{Inf}_{ik}[f] + \mathrm{Inf}_{jk}[f]).$$

Proof. The Cauchy–Schwartz inequality implies that $(a+b+c)^2 \le 3(a^2+b^2+c^2)$. Since $(i\,j) = (i\,k)(j\,k)(i\,k)$,

$$
\begin{aligned}
\mathrm{Inf}_{ij}[f] &= \tfrac{1}{2}\|f - f^{(i\,k)(j\,k)(i\,k)}\|^2 \\
&\le 3(\tfrac{1}{2}\|f - f^{(i\,k)}\|^2 + \tfrac{1}{2}\|f^{(i\,k)} - f^{(j\,k)(i\,k)}\|^2 + \tfrac{1}{2}\|f^{(j\,k)(i\,k)} - f^{(i\,k)(j\,k)(i\,k)}\|^2 \\
&= 6\,\mathrm{Inf}_{ik}[f] + 3\,\mathrm{Inf}_{jk}[f].
\end{aligned}
$$

The lemma is obtained by averaging with the similar inequality $\mathrm{Inf}_{ij}[f] \le 6\,\mathrm{Inf}_{jk}[f] + 3\,\mathrm{Inf}_{ik}[f]$ obtained using $(i\,j) = (j\,k)(i\,k)(j\,k)$. $\square$

As a consequence of the triangle inequality, we can identify a set of "important" coordinates for functions with low total influence (cf. [26, Proposition 5.3]).

Lemma 5.2. *Let $f \in \mathcal{H}_n$. For every $\tau > 0$ there exists a set $S \subseteq [n]$ of size $O(\mathrm{Inf}[f]/\tau)$ such that $\mathrm{Inf}_{ij}[f] < \tau$ whenever $i, j \notin S$.*

Proof. Construct a graph G on the vertex set $[n]$ in which two vertices i, j are connected if $\mathrm{Inf}_{ij}[f] \geq \tau$. Let M be a maximal matching in G. For each $(i, j) \in M$, Lemma 5.1 shows that for all $k \neq i, j$, $\mathrm{Inf}_{ik}[f] + \mathrm{Inf}_{jk}[f] = \Omega(\tau)$. Summing over all edges in M we obtain $\mathrm{Inf}[f] = \Omega(\tau|M|)$ and so $|M| = O(\mathrm{Inf}[f]/\tau)$. The endpoints of M form a vertex cover S in G of size $2|M| = O(\mathrm{Inf}[f]/\tau)$. Since S is a vertex cover, whenever $\mathrm{Inf}_{ij}[f] \geq \tau$ then (i, j) is an edge and so either $i \in S$ or $j \in S$ (or both). It follows that $\mathrm{Inf}_{ij}[f] < \tau$ whenever $i, j \notin S$. $\square$

Our goal in the rest of this section is to give a formula for $\mathrm{Inf}^m[f]$. Our treatment closely follows Wimmer [26]. We start with a formula for $f^{(m\ m+1)}$.

Lemma 5.3. *Let $f \in \mathcal{H}_n$. For $m \in [n-1]$,*

$$f^{(m\ m+1)} = \sum_{\substack{B \in \mathcal{B}_n: \\ m,m+1 \notin B \ or \\ m,m+1 \in B}} \hat{f}(B)\chi_B + \sum_{\substack{B \in \mathcal{B}_n: \\ b_{i+1}=m,m+1 \notin B}} \left(\frac{1}{m-2i}\hat{f}(B) + \frac{m-2i+1}{m-2i}\hat{f}(B^{(m\ m+1)}) \right) \chi_B$$

$$+ \sum_{\substack{B \in \mathcal{B}_n: \\ b_{i+1}=m+1,m \notin B}} \left(-\frac{1}{m-2i}\hat{f}(B) + \frac{m-2i-1}{m-2i}\hat{f}(B^{(m\ m+1)}) \right) \chi_B.$$

It might happen that $B \in \mathcal{B}_n$ but $B^{(m\ m+1)} \notin \mathcal{B}_n$, but in that case, the coefficient in front of $\hat{f}(B^{(m\ m+1)})$ vanishes.

Proof. For brevity, define $\pi = (m\ m+1)$. We start by showing that if $B \in \mathcal{B}_n$ but $B^\pi \notin \mathcal{B}_n$ then the coefficient in front of $\hat{f}(B^\pi)$ vanishes. Clearly, this case can happen only if $m+1 \in B$, say $b_{i+1} = m+1$, and $m \notin B$. Since $B \in \mathcal{B}_n$, $m - i = |[m+1] \setminus B| \geq i+1$. Since $B^\pi \notin \mathcal{B}_n$, $m - 1 - i = |[m] \setminus B^\pi| \leq i$. We conclude that $m = 2i+1$, and indeed the corresponding coefficient is $\frac{m-2i-1}{m-2i} = 0$.

Similarly, if $b_{i+1} \in \{m, m+1\}$ for some $B \in \mathcal{B}_n$ then $b_{i+1} - i - 1 = |[b_{i+1}] \setminus B| \geq i+1$, and in particular $m - 2i \geq b_{i+1} - 1 - 2i \geq 1$. This shows that the expressions appearing in the expansion of f^π are well-defined.

Consider some $B \in \mathcal{B}_n$. If $m, m+1 \notin B$ then $A < B$ iff $A^\pi < B$, and so

$$\chi_B^\pi = \sum_{\substack{A \in \mathcal{S}_n: \\ A < B}} \chi_{A^\pi, B} = \sum_{\substack{A \in \mathcal{S}_n: \\ A^\pi < B}} \chi_{A, B} = \chi_B.$$

If $m, m+1 \in B$, say $b_i = m$ and $b_{i+1} = m+1$, then $\chi_{A,B}^\pi = \chi_{A',B}$, where A' is obtained from A by switching a_i and a_{i+1}. Again $A < B$ iff $A' < B$, and so $\chi_B^\pi = \chi_B$ as in the preceding case.

Finally, consider some set $B \in \mathcal{B}_{n,d}$ such that $B^\pi \in \mathcal{B}_n$, $b_{i+1} = m$ and $m+1 \notin B$. The set B doesn't necessarily belong to $\mathcal{B}_n$, and in that case we define $\chi_B = 0$; under this convention, the formula $\chi_B = \sum_{A<B} \chi_{A,B}$ still holds (vacuously). We define a function ϕ which maps a sequence $A < B$ to a sequence $\phi(A) < B^\pi$ so that the following equation holds:

$$\chi_{A,B}^\pi - \chi_{A,B} = \chi_{\phi(A),B^\pi}. \tag{1}$$

The function ϕ is given by

$$\phi(A) = \begin{cases} a_1, \ldots, a_i, m, a_{i+2}, \ldots, a_d & \text{if } m+1 \notin A, \\ a_1, \ldots, a_i, m, a_{i+2}, \ldots, a_j, a_{i+1}, a_{j+2}, \ldots, a_d & \text{if } a_{j+1} = m+1. \end{cases}$$

Since $b_{i+1} = m$ and $A < B$, in the second case necessarily $j > i$. It is not hard to verify that indeed $\phi(A) < B^\pi$.

We proceed to verify equation (1). Suppose first that $m+1 \notin A$. Then

$$\chi_{A,B}^\pi - \chi_{A,B} = \prod_{k=1}^{i}(x_{a_k} - x_{b_k}) \times [(x_{a_{i+1}} - (m+1)) - (x_{a_{i+1}} - m)] \times \prod_{k=i+2}^{d}(x_{a_k} - x_{b_k})$$

$$= \prod_{k=1}^{i}(x_{a_k} - x_{b_k}) \times (m - (m+1)) \times \prod_{k=i+2}^{d}(x_{a_k} - x_{b_k}) = \chi_{\phi(A),B^\pi}.$$

Suppose next that $a_{j+1} = m + 1$. Then

$$\chi_{A,B}^{\pi} - \chi_{A,B} = \prod_{\substack{k=1 \\ k \neq i+1, j+1}}^{n} (x_{a_k} - x_{b_k}) \times [(x_{a_{i+1}} - (m+1))(m - x_{b_{j+1}}) - (x_{a_{i+1}} - m)(m+1 - x_{b_{j+1}})].$$

Using $(\alpha - 1)\beta - \alpha(\beta + 1) = -\alpha - \beta$,

$$(x_{a_{i+1}} - (m+1))(m - x_{b_{j+1}}) - (x_{a_{i+1}} - m)(m+1 - x_{b_{j+1}}) = -(x_{a_{i+1}} - m) - (m - x_{b_{j+1}}) = -(x_{a_{i+1}} - x_{b_{j+1}}).$$

Therefore

$$\chi_{A,B}^{\pi} - \chi_{A,B} = \prod_{\substack{k=1 \\ k \neq i+1, j+1}}^{n} (x_{a_k} - x_{b_k}) \times (m - (m+1)) \times (x_{a_{i+1}} - x_{b_{j+1}}) = \chi_{\phi(A),B^{\pi}}.$$

This completes the proof of equation (1).

Every sequence in $A \in \operatorname{im} \phi$ satisfies $a_{i+1} = m$. Let $A < B^{\pi}$ be any sequence satisfying $a_{i+1} = m$. We proceed to determine $|\phi^{-1}(A)|$. For every $t \in [m-1] \setminus (A \cup B)$, the sequence $a_1, \ldots, a_i, t, a_{i+2}, \ldots, a_d$ is in $\phi^{-1}(A)$. For each $j > i$ such that $a_{j+1} < m$, another sequence in $\phi^{-1}(A)$ is $a_1, \ldots, a_i, a_{j+1}, a_{i+2}, \ldots, a_j, m + 1, a_{j+2}, \ldots, a_d$. If the number of such latter indices is J, then $|[m-1] \setminus (A \cup B)| = m - 1 - 2i - J$. In total, $|\phi^{-1}(A)| = m - 1 - 2i$. We conclude that

$$\chi_B^{\pi} - \chi_B = (m - 2i - 1) \sum_{A \in \operatorname{im} \phi} \chi_{A,B^{\pi}}. \tag{2}$$

We proceed to calculate $\chi_{B^{\pi}} - \chi_B$. If $A < B$ then clearly $A^{\pi} < B^{\pi}$. Conversely, if $A^{\pi} < B^{\pi}$ then $A < B$ unless $a_{i+1}^{\pi} = m$, in which case $A^{\pi} \in \operatorname{im} \phi$. The preceding paragraph shows that every $A^{\pi} < B^{\pi}$ satisfying $a_{i+1}^{\pi} = m$ belongs to $\operatorname{im} \phi$, and so

$$\chi_{B^{\pi}} - \chi_B = \sum_{\substack{A \in \mathcal{S}_{n,d}: \\ A < B}} (\chi_{A^{\pi}, B^{\pi}} - \chi_{A,B}) + \sum_{A \in \operatorname{im} \phi} \chi_{A,B^{\pi}}$$

$$= \sum_{\substack{A \in \mathcal{S}_{n,d}: \\ A < B}} (\chi_{A,B}^{\pi} - \chi_{A,B}) + \sum_{A \in \operatorname{im} \phi} \chi_{A,B^{\pi}} = (m - 2i) \sum_{A \in \operatorname{im} \phi} \chi_{A,B^{\pi}}, \tag{3}$$

using equation (2). Combining equations (2),(3) together, we deduce

$$\chi_B^{\pi} = \chi_B + \frac{m - 2i - 1}{m - 2i}(\chi_{B^{\pi}} - \chi_B) = \frac{1}{m - 2i}\chi_B + \frac{m - 2i - 1}{m - 2i}\chi_{B^{\pi}}.$$

Applying π to both sides of equations (2),(3), we get

$$\chi_B - \chi_B^{\pi} = (m - 2i - 1) \sum_{A \in \operatorname{im} \phi} \chi_{A,B^{\pi}}^{\pi}, \quad \chi_{B^{\pi}} - \chi_B^{\pi} = (m - 2i) \sum_{A \in \operatorname{im} \phi} \chi_{A,B^{\pi}}^{\pi}.$$

Therefore

$$\chi_{B^{\pi}}^{\pi} = \chi_B^{\pi} + \frac{m - 2i}{m - 2i - 1}(\chi_B - \chi_B^{\pi}) = \frac{m - 2i}{m - 2i - 1}\chi_B - \frac{1}{m - 2i - 1}\chi_B^{\pi} = -\frac{1}{m - 2i}\chi_{B^{\pi}} + \frac{m - 2i + 1}{m - 2i}\chi_B. \quad \square$$

Lemma 5.3 allows us to come up with a remarkable formula for $\sum_{1 \leq i < m} f^{(i\ m)}$, showing that the basis vectors χ_B are eigenvectors of this operator (cf. [26, Proposition 4.2], in which $\lambda_m(B)$ is the content of m in the Young tableau of shape $n - |B|, |B|$ having bottom row B).

Lemma 5.4. *Let $f \in \mathcal{H}_n$. For $m \in [n]$,*

$$\sum_{1 \le i < m} f^{(i\ m)} = \sum_{B \in \mathcal{B}_n} \lambda_m(B)\hat{f}(B)\chi_B, \quad \text{where } \lambda_m(B) = \begin{cases} i - 2 & \text{if } b_i = m, \\ m - i & \text{if } b_{i-1} < m < b_i, \end{cases}$$

using the conventions $b_0 = -\infty$ and $b_{d+1} = \infty$ for $B \in \mathcal{B}_{n,d}$.

Proof. The proof is by induction on m. When $m = 1$, the sum in question vanishes, and so we need to prove that $\lambda_1(B) = 0$ for all $B \in \mathcal{B}_n$. Indeed, $b_1 \ge 2$ for all $B \in \mathcal{B}_n$ and so $\lambda_1(B) = 0$. Suppose now that the formula holds for some m. Let $\pi = (m\ m+1)$. Since $(i\ m+1) = \pi(i\ m)\pi$,

$$\sum_{1 \le i < m+1} f^{(i\ m+1)} = \left(\sum_{1 \le i < m} (f^\pi)^{(i\ m)} \right)^\pi + f^\pi.$$

Lemma 5.3 shows that

$$f^\pi = \sum_{\substack{B \in \mathcal{B}_n: \\ m,m+1 \notin B \text{ or} \\ m,m+1 \in B}} \hat{f}(B)\chi_B$$

$$+ \sum_{\substack{B^\pi \in \mathcal{B}_n: \\ b_{i+1}=m,m+1 \notin B}} \left(\frac{1}{m-2i}\hat{f}(B) + \frac{m-2i+1}{m-2i}\hat{f}(B^\pi) \right) \chi_B$$

$$+ \left(-\frac{1}{m-2i}\hat{f}(B^\pi) + \frac{m-2i-1}{m-2i}\hat{f}(B) \right) \chi_{B^\pi}.$$

The induction hypothesis implies that

$$\sum_{1 \le i < m} (f^\pi)^{(i\ m)} = \sum_{\substack{B \in \mathcal{B}_n: \\ m,m+1 \notin B \text{ or} \\ m,m+1 \in B}} \lambda_m(B)\hat{f}(B)\chi_B$$

$$+ \sum_{\substack{B^\pi \in \mathcal{B}_n: \\ b_{i+1}=m,m+1 \notin B}} (i-1)\left(\frac{1}{m-2i}\hat{f}(B) + \frac{m-2i+1}{m-2i}\hat{f}(B^\pi) \right) \chi_B$$

$$+ (m-i-1)\left(-\frac{1}{m-2i}\hat{f}(B^\pi) + \frac{m-2i-1}{m-2i}\hat{f}(B) \right) \chi_{B^\pi}.$$

Another application of Lemma 5.3 gives

$$\left(\sum_{1 \le i < m} (f^\pi)^{(i\ m)} \right)^\pi = \sum_{\substack{B \in \mathcal{B}_n: \\ m,m+1 \notin B \text{ or} \\ m,m+1 \in B}} \lambda_m(B)\hat{f}(B)\chi_B$$

$$+ \sum_{\substack{B^\pi \in \mathcal{B}_n: \\ b_{i+1}=m,m+1 \notin B}} \left[\frac{(m-2i)(m-i-1)-1}{m-2i}\hat{f}(B) - \frac{m-2i+1}{m-2i}\hat{f}(B^\pi) \right] \chi_B$$

$$+ \left[-\frac{m-2i-1}{m-2i}\hat{f}(B) + \frac{(m-2i)(i-1)+1}{m-2i}\hat{f}(B^\pi) \right] \chi_{B^\pi}.$$

15

We conclude that

$$\sum_{1 \le i < m+1} f^{(i\ m+1)} = \left(\sum_{1 \le i < m} (f^\pi)^{(i\ m)} \right)^\pi + f^\pi$$

$$= \sum_{\substack{B \in \mathcal{B}_n: \\ m,m+1 \notin B \text{ or} \\ m,m+1 \in B}} (\lambda_m(B) + 1) \hat{f}(B) \chi_B + \sum_{\substack{B^\pi \in \mathcal{B}_n: \\ b_{i+1} = m, m+1 \notin B}} (m - i - 1) \hat{f}(B) \chi_B + (i-1) \hat{f}(B^\pi) \chi_{B^\pi}.$$

It remains to verify that these coefficients match $\lambda_{m+1}(B)$. If $m, m+1 \notin B$ then $\lambda_m(B) = m - i$, where $b_{i-1} < m < b_i$. Since $m + 1 \notin B$, also $b_{i-1} < m + 1 < b_i$, and so $\lambda_{m+1}(B) = m + 1 - i = \lambda_m(B) + 1$. If $m, m+1 \in B$, say $b_i = m$ and $b_{i+1} = m + 1$, then $\lambda_{m+1}(B) = (i + 1) - 2 = (i - 2) + 1 = \lambda_m(B) + 1$. Finally, suppose that $b_{i+1} = m$ and $m + 1 \notin B$. In this case, $b_{i+1} < m + 1 < b_{i+2}$ and so $\lambda_{m+1}(B) = (m + 1) - (i + 2) = m - i - 1$, while $b_{i+1}^\pi = m + 1$ and so $\lambda_{m+1}(B^\pi) = (i + 1) - 2 = i - 1$. All cases match the coefficients in the displayed formula, completing the proof. $\qquad\square$

Lemma 5.4 shows that the elements of $\mathcal{Y}_n$ are eigenvectors of the operators $f \mapsto \sum_{1 \le i < m} f^{(i\ m)}$ for all m. In the terminology of Vershik and Okounkov [24], this makes $\mathcal{Y}_n$ a Gelfand–Tsetlin basis. Srinivasan [23] shows that the Gelfand–Tsetlin basis is unique in our setting, and so the basis he constructs, which is also a Gelfand–Tsetlin basis, is the same as $\mathcal{Y}_n$.

Lemma 5.4 allows us to give a formula for the mth total influence of a function (cf. [26, Definition 4.3]).

Lemma 5.5. *Let $f \in \mathcal{H}_n$. For $m \in [n]$,*

$$\sum_{1 \le i < j \le m} f^{(i\ j)} = \sum_{B \in \mathcal{B}_n} \tau_m(B) \chi_B, \quad \text{where } \tau_m(B) = \frac{m(m-1)}{2} - |B \cap [m]|(m + 1 - |B \cap [m]|).$$

Moreover,

$$\text{Inf}^m[f] = \sum_{B \in \mathcal{B}_n} \frac{|B \cap [m]|(m + 1 - |B \cap [m]|)}{m} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2.$$

In particular,

$$\text{Inf}[f] = \sum_{B \in \mathcal{B}_n} \frac{|B|(n + 1 - |B|)}{n} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2.$$

Proof. Lemma 5.4 proves the first formula with $\tau_m(B) = \sum_{k=2}^m \lambda_k(B)$. We can write $\lambda_k(B) = \lambda_k'(B) + \lambda_k''(B)$, where $\lambda_k'(B)$ is non-zero when $k \in B$, and $\lambda_k''(B)$ is non-zero when $k \notin B$. It follows from the definition that

$$\lambda_{b_1}'(B), \lambda_{b_2}'(B), \ldots, \lambda_{b_{|B|}}'(B) = -1, 0, \ldots, |B| - 2.$$

In particular,

$$\sum_{i=1}^m \lambda_i'(B) = -1 + \cdots + (|B \cap [m]| - 2) = \frac{|B \cap [m]|(|B \cap [m]| - 3)}{2}.$$

Consider the sequence $\lambda_1''(B), \ldots, \lambda_n''(B)$, in which we omit $\lambda_k''(B)$ for all $k \in B$. The sequence starts at $\lambda_1''(B) = 1 - 1 = 0$. Let $\lambda_k''(B) = k - i$, where $b_{i-1} < k < b_i$. If $b_i \ne k + 1$ then $\lambda_{k+1}''(B) = (k + 1) - i = \lambda_k''(B) + 1$. If $b_i = k + 1, \ldots, b_{i+\ell-1} = k + \ell$ and $k + \ell + 1 \le n$ then $\lambda_{k+1}''(B) = \cdots = \lambda_{k+\ell}'' = 0$ and $\lambda_{k+\ell+1}''(B) = (k + \ell + 1) - (i + \ell) = \lambda_k''(B) + 1$. In other words,

$$\lambda_1''(B), \ldots, \lambda_n''(B) = 0, 1, \ldots, n - |B| - 1.$$

In particular,

$$\sum_{i=1}^m \lambda_i''(B) = 0 + \cdots + (m - |B \cap [m]| - 1) = \frac{(m - |B \cap [m]|)(m - |B \cap [m]| - 1)}{2}.$$

16

In total,

$$\tau_m(B) = \frac{|B \cap [m]|(|B \cap [m]| - 3) + (m - |B \cap [m]|)(m - |B \cap [m]| - 1)}{2}$$

$$= \frac{m(m-1)}{2} - |B \cap [m]|(m + 1 - |B \cap [m]|),$$

completing the proof of the first formula.

In order to compute the mth total influence, notice that

$$\mathrm{Inf}_{ij}[f] = \tfrac{1}{2}\|f^{(i\,j)} - f\|^2 = \tfrac{1}{2}\|f\|^2 + \tfrac{1}{2}\|f^{(i\,j)}\|^2 - \langle f, f^{(i\,j)}\rangle = \|f\|^2 - \langle f, f^{(i\,j)}\rangle,$$

using the exchangeability of the measure. Therefore

$$\mathrm{Inf}^m[f] = \frac{1}{m}\left(\frac{m(m-1)}{2}\|f\|^2 - \left\langle f, \sum_{1 \le i < j \le m} f^{(i\,j)}\right\rangle\right).$$

The formula for $\mathrm{Inf}^m[f]$ now immediately follows from the orthogonality of the basis (Theorem 3.1) and the norms stated in Theorem 3.2. $\qquad\square$

As a simple corollary, we obtain a version of Poincaré's inequality.

Lemma 5.6. *For any $f \in \mathcal{H}_{n,d}$ we have*

$$\mathbb{V}[f] \le \mathrm{Inf}[f] \le d\,\mathbb{V}[f].$$

Proof. Lemma 3.3 and Lemma 5.5 give the formulas

$$\mathbb{V}[f] = \sum_{\substack{B \in \mathcal{B}_n: \\ B \ne \emptyset}} \hat{f}(B)^2 c_B \|\chi_{2,\dots,2|B|}\|^2,$$

$$\mathrm{Inf}[f] = \sum_{\substack{B \in \mathcal{B}_n: \\ B \ne \emptyset}} \frac{|B|(n + 1 - |B|)}{n} \hat{f}(B)^2 c_B \|\chi_{2,\dots,2|B|}\|^2.$$

The left inequality follows from $|B|(n + 1 - |B|) \ge n$, and the right inequality from $n + 1 - |B| \le n$. $\qquad\square$

6 Wimmer–Friedgut theorem

For a Boolean function f on the Boolean hypercube $\{0,1\}^n$, the influence of variable i is

$$\mathrm{Inf}_i[f] = \Pr_{x_1,\dots,x_n \in \{0,1\}^n}[f(x_1,\dots,x_n) \ne f(x_1,\dots,1 - x_i,\dots,x_n)],$$

under the uniform distribution over the hypercube. The total influence of the function is $\mathrm{Inf}[f] = \sum_{i=1}^n \mathrm{Inf}_i[f]$. Friedgut [9] proved that for every $\epsilon > 0$, every Boolean function f is ϵ-close to a function g depending on $2^{O(\mathrm{Inf}[f]/\epsilon)}$ coordinates, that is, $\Pr_{x \in \{0,1\}^n}[f(x) \ne g(x)] \le \epsilon$.

Friedgut's theorem can be seen as an "inverse theorem" corresponding to the easy fact that a Boolean function depending on d coordinates has total influence at most d. Since the total influence is bounded by the degree, Friedgut's theorem can also be seen as a strengthening of the result of Nisan and Szegedy [18] that a Boolean function of degree d depends on at most $d2^{d-1}$ coordinates.

Wimmer [26] proved an analog of Friedgut's theorem for functions on a slice of the Boolean cube. His proof takes place mostly on the symmetric group, and uses properties of Young's orthogonal representation. We rephrase his proof in terms of Young's orthogonal basis for the slice.

The proof relies crucially on a hypercontractivity property due to Lee and Yau [15]. Before stating the property, we need to define the noise operator.

Definition 6.1. The *Laplacian* operator on functions $f \in \mathcal{H}_n$ is given by

$$Lf = f - \frac{1}{\binom{n}{2}} \sum_{1 \leq i < j \leq n} f^{(i\, j)}.$$

The *noise operator* H_t is given by $H_t = e^{-tL}$.

The Laplacian corresponds to the Markov chain applying a random transposition $(i\ j)$. Moreover, $L = I - K$ where K is the transition matrix of the Markov chain. We can expand the noise operator as

$$H_t = e^{t(K-I)} = e^{-t} \sum_{\ell=0}^{\infty} \frac{t^\ell}{\ell!} K^\ell.$$

In words, H_t corresponds to applying $P(t)$ many random transpositions $(i\ j)$, where $P(t)$ is the Poisson distribution with mean t.

Lemma 5.5 gives a formula for Lf and $H_t f$.

Lemma 6.1. *Let $f \in \mathcal{H}_n$. For every t,*

$$Lf = \sum_{B \in \mathcal{B}_n} \frac{2|B|(n+1-|B|)}{n(n-1)} \hat{f}(B)\chi_B, \quad H_t f = \sum_{B \in \mathcal{B}_n} \exp\left(-t\frac{2|B|(n+1-|B|)}{n(n-1)}\right) \hat{f}(B)\chi_B.$$

The hypercontractivity result of Lee and Yau [15] gives for all $p < q$ a value of t such that $\|H_t f\|_q \leq \|f\|_p$.

Proposition 6.2. *Let n, k be integers such that $1 \leq k \leq n-1$. The log-Sobolev constant ρ of the Markov chain corresponding to the Laplacian L is given by*

$$\rho^{-1} = \Theta\left(n \log \frac{n^2}{k(n-k)}\right).$$

Consequently, for every $t \geq 0$ and $1 \leq p \leq q \leq \infty$ satisfying $\frac{q-1}{p-1} \leq \exp(2\rho t)$ and all $f \in \mathcal{H}_n$, $\|H_t f\|_q \leq \|f\|_p$.

Proof. The first result is [15, Theorem 5]. Their parameter t is scaled by a fraction of n. Furthermore, their log-Sobolev constant is the reciprocal of ours. The second result is due to Gross [11], and is quoted from [20, Theorem 2.4]. □

We can now state and prove the Wimmer–Friedgut theorem [26, Theorem 1.3].

Definition 6.2. A function $f \in \mathcal{H}_n$ *depends (only) on a set* $S \subseteq [n]$ if $f = f^{(i\, j)}$ whenever $i, j \notin S$.

Theorem 6.3. *Let n, k be integers such that $1 \leq k \leq n/2$, and define $p = k/n$. For every Boolean function f on the slice and any $\epsilon > 0$ there exists a Boolean function g depending on $O(p^{-O(\mathrm{Inf}[f]/\epsilon)})$ coordinates such that $\Pr[f \neq g] \leq \epsilon$.*

Proof. Let $\tau > 0$ be a parameter to be determined. Lemma 5.2 shows that there exists a set $S \subseteq [n]$ of size $m = O(\mathrm{Inf}[f]/\tau)$ such that $\mathrm{Inf}_{ij}[f] < \tau$ whenever $i, j \notin S$. Without loss of generality, we can assume that $S = \{n-m+1, \ldots, n\}$. Let h be the function obtained from f by averaging over all permutations of $[n-m]$, and let g be the Boolean function obtained from rounding h to $\{0, 1\}$. Note that h and g both depend only on the last m coordinates. Lemma 3.4 shows that the Young–Fourier expansion of h is obtained from that of f by dropping all terms $\hat{f}(B)\chi_B$ for which $B \cap [n-m] \neq \emptyset$. Therefore Lemma 3.3 shows that

$$\Pr[f \neq g] = \|f - g\|^2 \leq 2\|f - h\|^2 = 2 \sum_{\substack{B \in \mathcal{B}_n : \\ B \cap [n-m] \neq \emptyset}} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2.$$

We bound the sum on the right-hand side by considering separately large sets and small sets. Let $d \leq n/2$ be a parameter to be determined. For the large sets, we have

$$\sum_{\substack{B \in \mathcal{B}_n: \\ |B| \geq d}} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2 \leq \frac{n}{d(n+1-d)} \sum_{\substack{B \in \mathcal{B}_n: \\ |B| \geq d}} \frac{|B|(n+1-|B|)}{n} c_B \|\chi_{|B|}\|^2 = \frac{n}{d(n+1-d)} \operatorname{Inf}[f],$$

using Lemma 5.5.

In order to bound the part corresponding to small sets, we need to apply hypercontractivity. Choose $(p,q) = (\frac{4}{3}, 2)$ in Proposition 6.2 to deduce that for $t = \Theta(n \log \frac{n^2}{k(n-k)}) = \Theta(n \log p^{-1})$ we have $\|H_t F\|_2 \leq \|F\|_{4/3} = \|F\|_2^{3/2}$ for every function F on the slice which takes the values $0, \pm 1$. In particular, choosing $F = f - f^{(i\,j)}$ we obtain $(2 \operatorname{Inf}_{ij}[H_t f])^{1/2} \leq (2 \operatorname{Inf}_{ij}[f])^{3/4}$ and so $\operatorname{Inf}_{ij}[H_t f] \leq \sqrt{2} \operatorname{Inf}_{ij}[f]^{3/2}$. Therefore

$$\operatorname{Inf}^{n-m}[H_t f] \leq \frac{\sqrt{2}}{n-m} \sum_{1 \leq i < j \leq n-m} \operatorname{Inf}_{ij}[f]^{3/2} \leq \frac{\sqrt{2}n}{n-m} \sqrt{\tau} \operatorname{Inf}[f].$$

On the other hand, Lemma 5.5 and Lemma 6.1 show that

$$
\begin{aligned}
\operatorname{Inf}^{n-m}[H_t f] &= \sum_{B \in \mathcal{B}_n} \exp\left(-t \frac{2|B|(n+1-|B|)}{n(n-1)}\right) \frac{|B \cap [n-m]|(n-m+1-|B \cap [n-m]|)}{n-m} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2 \\
&\geq \sum_{\substack{B \in \mathcal{B}_n: \\ |B| < d \\ B \cap [n-m] \neq \emptyset}} \exp\left(-t \frac{2d(n+1-d)}{n(n-1)}\right) \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2 \\
&= \sum_{\substack{B \in \mathcal{B}_n: \\ |B| < d \\ B \cap [n-m] \neq \emptyset}} p^{\Theta(d)} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2.
\end{aligned}
$$

Altogether,

$$\sum_{\substack{B \in \mathcal{B}_n: \\ |B| < d \\ B \cap [n-m] \neq \emptyset}} \hat{f}(B)^2 c_B \|\chi_{|B|}\|^2 \leq \sqrt{2} p^{-\Theta(d)} \frac{n}{n-m} \sqrt{\tau} \operatorname{Inf}[f].$$

Putting both bounds together, we deduce

$$\Pr[f \neq g] \leq 2 \frac{n}{d(n+1-d)} \operatorname{Inf}[f] + 2\sqrt{2} p^{-\Theta(d)} \frac{n}{n-m} \sqrt{\tau} \operatorname{Inf}[f].$$

We now choose $d = \operatorname{Inf}[f]/(8\epsilon)$ and $\tau = p^{Cd}$ for an appropriate constant $C > 0$, so that $m = O(\operatorname{Inf}[f] p^{-O(\operatorname{Inf}[f]/\epsilon)}) = O(p^{-O(\operatorname{Inf}[f]/\epsilon)})$. We can assume that $d, m \leq n/2$, since otherwise the theorem is trivial. We can choose C so that

$$\Pr[f \neq g] \leq \frac{\epsilon}{2} + O(p^{Cd/3} \operatorname{Inf}[f]) = \frac{\epsilon}{2} + O(2^{-Cd/3} \operatorname{Inf}[f]) = \frac{\epsilon}{2} + O(2^{-Cd/3} d) \frac{\epsilon}{2}.$$

For an appropriate choice of C, the second term is at most $\epsilon/2$, completing the proof. □

References

[1] Eiichi Bannai and Tatsuro Ito. *Algebraic Combinatorics I: Association schemes.* Benjamin/Cummings Pub. Co., 1984.

[2] Charles F. Dunkl. A Krawtchouk polynomial addition theorem and wreath products of symmetric groups. *Indiana Univ. Math. J.*, 25:335–358, 1976.

[3] Charles F. Dunkl. An addition theorem for Hahn polynomials: the spherical functions. *SIAM J. MATH. Anal.*, 9:627–637, 1978.

[4] Charles F. Dunkl. Orthogonal functions on some permutation groups. In *Relations between combinatorics and other parts of mathematics*, volume 34 of *Proc. Sump. Pure Math.*, pages 129–147, Providence, RI, 1979. Amer. Math. Soc.

[5] Yuval Filmus. Friedgut–Kalai–Naor theorem for slices of the Boolean cube, 2014. Submitted.

[6] Yuval Filmus, Guy Kindler, Elchanan Mossel, and Karl Wimmer. Invariance principle on the slice, 2015. Submitted.

[7] Yuval Filmus and Elchanan Mossel. Harmonicity and invariance on slices of the Boolean cube, 2015. Submitted.

[8] Péter Frankl and Ron L. Graham. Old and new proofs of the Erdős-Ko-Rado theorem. *J. of Sichuan Univ. Natural Science Edition*, 26:112–122, 1989.

[9] Ehud Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):27–36, 1998.

[10] Ehud Friedgut, Gil Kalai, and Assaf Naor. Boolean functions whose Fourier transform is concentrated on the first two levels. *Adv. Appl. Math.*, 29(3):427–437, 2002.

[11] Leonard Gross. Logarithmic Sobolev inequalities. *Amer. J. Math.*, 97:1061–1083, 1975.

[12] Jeff Kahn, Gil Kalai, and Nati Linial. The influence of variables on Boolean functions. In *Proceedings of the 29th Symposium on the Foundations of Computer Science*, pages 68–80, White Plains, 1988.

[13] Guy Kindler. *Property testing, PCP and Juntas*. PhD thesis, Tel-Aviv University, 2002.

[14] Guy Kindler and Shmuel Safra. Noise-resistant Boolean functions are juntas, 2004. Manuscript.

[15] Tzong-Yau Lee and Horng-Tzer Yau. Logarithmic Sobolev inequality for some models of random walks. *Ann. Prob.*, 26(4):1855–1873, 1998.

[16] László Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25:1–7, 1979.

[17] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Ann. Math.*, 171:295–341, 2010.

[18] Noam Nisan and Mario Szegedy. On the degree of Boolean functions as real polynomials. *Comp. Comp.*, 4(4):301–313, 1994.

[19] Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

[20] Ryan O'Donnell and Karl Wimmer. KKL, Kruskal–Katona, and monotone nets. In *50th Annual Symposium on Foundations of Computer Science (FOCS 2009)*, pages 725–734, 2009.

[21] Ryan O'Donnell and Karl Wimmer. KKL, Kruskal–Katona, and monotone nets. *SIAM J. Comput.*, 42(6):2375–2399, 2013.

[22] Ryan O'Donnell and Karl Wimmer. Sharpness of KKL on Schreier graphs. *Elec. Comm. Prob.*, 18:12:1–12, 2013.

[23] 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.

[24] Anatoly Vershik and Andrei Okounkov. A new approach to the representation theory of the symmetric groups — II. *J. Math. Sci.*, 131:5471–5494, 2005.

[25] Richard M. Wilson. The exact bound in the Erdős-Ko-Rado theorem. *Combinatorica*, 4:247–257, 1984.

[26] 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.