CAMBRIDGE
UNIVERSITY PRESS

ARTICLE

FKN theorem for the multislice, with applications

Yuval Filmus[†]

Computer Science Department, Technion – Israel Institute of Technology
Email: yuvalfi@cs.technion.ac.il

Abstract

The Friedgut–Kalai–Naor (FKN) theorem states that if f is a Boolean function on the Boolean cube which is close to degree one, then f is close to a *dictator*, a function depending on a single coordinate. The author has extended the theorem to the *slice*, the subset of the Boolean cube consisting of all vectors with fixed Hamming weight. We extend the theorem further, to the *multislice*, a multicoloured version of the slice.

As an application, we prove a stability version of the edge-isoperimetric inequality for settings of parameters in which the optimal set is a dictator.

1. Introduction

The classical Friedgut–Kalai–Naor (FKN) theorem [10] is a basic structural result in Boolean function analysis. It is a stability version of the following trivial result: the only Boolean functions on the Boolean cube $\{0, 1\}^n$ which have degree one are dictators, that is, functions depending on a single coordinate.[1] The FKN theorem can be stated in two equivalent ways.

(1) If $f\colon \{0, 1\}^n \to \{0, 1\}$ is ε-close to degree one, that is, $\|f^{>1}\|^2 = \varepsilon$, then f is $O(\varepsilon)$-close to a Boolean dictator, that is, $\mathbb{P}\,[f \neq g] = O(\varepsilon)$ for some Boolean dictator $g\colon \{0, 1\}^n \to \{0, 1\}$.

(2) If $f\colon \{0, 1\}^n \to \mathbb{R}$ is a degree one function which is ε-close to Boolean, that is,

$$\mathbb{E}\,[\mathrm{dist}\,(f, \{0, 1\})^2] = \varepsilon,$$

then f is $O(\varepsilon)$-close to a Boolean dictator, that is, $\mathbb{E}\,[(f - g)^2] = O(\varepsilon)$ for some Boolean dictator $g\colon \{0, 1\}^n \to \{0, 1\}$.

In fact, the error bound can be improved from $O(\varepsilon)$ to $\varepsilon + \tilde{O}(\varepsilon^2)$; see [12, 14].

The FKN theorem has been extended to many other domains: to graph products [1], to the biased Boolean cube [12, 13], to sums of functions on disjoint variables [12, 15], and to non-product domains: the symmetric group [4, 5], the slice [7], and high-dimensional expanders [3].

In this paper we extend it to the multislice, a generalization of the slice recently considered by O'Donnell, Wu and the author [9].

[†]Taub Fellow, supported by the Taub Foundations. The research was funded by ISF grant 1337/16.
[1]This result is much less trivial on other domains, such as the symmetric group [6] and the Grassmann scheme [8]. See [8] for a survey, which includes a simple proof of this result for the multislice.

Given positive integers $\kappa_1, \ldots, \kappa_\ell$ summing to n, the *multislice* $\mathcal{U}_\kappa$ consists of all vectors in $[\ell]^n$ in which the number of coordinates equal to i is κ_i. When $\ell = 2$, this domain is known as the *slice*, and when $\ell = n$, we obtain the symmetric group. In this paper, we focus on the case in which ℓ is constant, and furthermore the multislice is unbiased: $\kappa_1, \ldots, \kappa_n \geqslant \rho n$ for some constant $\rho > 0$. The biased case, in which the weights $\kappa_1, \ldots, \kappa_\ell$ are allowed to become arbitrarily small, is more difficult, since in this case the approximating function need not be a dictator; see [7] for more details.

In order to formulate the FKN theorem for the multislice, we need to generalize the concept of degree one function. There are several different routes to this generalization, all yielding the same class of functions.

(1) *Representation theory of the symmetric group.* The multislice can be viewed as a permutation module of S_n. The representation theory of S_n decomposes the space of functions on S_n to isotypical components indexed by partitions of n, which are partially ordered according to majorization. In the case of the slice, the degree d functions are those supported on the isotypical components corresponding to partitions in which the first part contains at least $n - d$ boxes. We can use the same definition on the multislice.

(2) *Polynomial degree.* We can view the input to a function on the multislice as consisting of ℓn Boolean variables x_{ji} encoding the input vector u via $x_{ji} = 1_{u_j = i}$. A function on the Boolean cube or on the slice has degree d if it can be represented as a polynomial of degree d over these variables. This definition carries over to the multislice.

(3) *Junta degree.* A function on the Boolean cube or on the slice has degree d if it is a linear combination of d-juntas, that is, functions depending on d coordinates. The same definition works on the multislice.

Armed with the concept of degree one function, we can state our main theorem.

Theorem 1.1. *Fix an integer $\ell \geqslant 2$ and a parameter $\rho > 0$. There exists a constant $N = N(\ell, \rho)$ for which the following hold. Let $\kappa_1, \ldots, \kappa_\ell \geqslant \rho n$ be integer weights summing to $n \geqslant N$.*

If $f \colon \mathcal{U}_\kappa \to \mathbb{R}$ is a degree one function which satisfies $\mathbb{E}\left[\operatorname{dist}(f, \{0, 1\})^2\right] = \varepsilon$, then there exists a Boolean function $g \colon \mathcal{U}_\kappa \to \{0, 1\}$, depending on a single coordinate, such that $\mathbb{E}\left[(f - g)^2\right] \leqslant \varepsilon + O_{\ell, \rho}(\varepsilon^2)$.

If $F \colon \mathcal{U}_\kappa \to \{0, 1\}$ satisfies $\|F^{>1}\|^2 = \varepsilon$, then there exists a Boolean function $g \colon \mathcal{U}_\kappa \to \{0, 1\}$, depending on a single coordinate, such that $\mathbb{P}\left[F \neq g\right] \leqslant 4\varepsilon + O_{\ell, \rho}(\varepsilon^2)$.[2]

(The definition of $F^{>1}$ appears at the end of Section 2.1.)

1.1 Application to edge isoperimetry

Let A be an arbitrary subset of the multislice $\mathcal{U}_\kappa$. The (edge) *expansion* of A is

$$\Phi(A) = \mathbb{P}_{\substack{u \sim A \\ \tau \sim \operatorname{Trans}(n)}} [u^\tau \notin A],$$

where u is a random point chosen from A, $\tau = (j_1 \, j_2)$ is a random transposition in S_n, and u^τ is obtained from u by switching the values of u_{j_1} and u_{j_2}. In words, the expansion of A is the probability that if we choose a random point of A and switch two of its coordinates at random, we reach a point not in A.

---

[2]This bound can probably be improved to $\varepsilon + O_{\ell, \rho}(\varepsilon^2 \log(1/\varepsilon))$, along the lines of [12, Theorem 5.3] and [14, Theorem 5.33], but we have not attempted to do so.

The edge-isoperimetry question is the following:

Given $0 < \alpha < 1$, which sets of size $\alpha |\mathcal{U}_\kappa|$ minimize the expansion?

When $\alpha n = \sum_{i \in S} \kappa_i$ for some $S \subseteq [\ell]$, it is natural to conjecture that the sets of the form $A = \{u : u_j \in S\}$ minimize the expansion, and this is indeed the case. Using our FKN theorem, we are able to show a stability version of this result: if a set of size αn has almost minimal expansion, then it is close to a set with minimal expansion.

Preliminaries

We use $\mathbb{E}$ to denote expectation, and $\mathbb{P}$ to denote probability. The distance of an element x to a set S is $\text{dist}(x, S) = \min_{y \in S} |x - y|$. For a set S, the notation $S \pm \varepsilon$ stands for $\{x : \text{dist}(x, S) \leqslant \varepsilon\}$. A function is *Boolean* if it is $\{0, 1\}$-valued. The L^2 triangle inequality is the inequality $(a + b)^2 \leqslant 2(a^2 + b^2)$.

Let $\kappa_1, \ldots, \kappa_\ell$ be positive integers summing to n. The *multislice* $\mathcal{U}_\kappa$ consists of all vectors $u \in [\ell]^n$ in which the number of coordinates equal to i is κ_i, for all $i \in [\ell]$. The multislice is ρ-*balanced* if $\kappa_1, \ldots, \kappa_\ell \geqslant \rho n$.

We endow the multislice with the uniform measure. If f is a function on the multislice, then its L^2 norm is $\|f\| = \sqrt{\mathbb{E}[f^2]}$. We say that two functions f, g are ε-close if $\|f - g\|^2 \leqslant \varepsilon$.

We can think of a function on the multislice as being defined over the set of Boolean variables $(x_{ji})_{j \in [n], i \in [\ell]}$, which encode an element $u \in \mathcal{U}_\kappa$ in the following way: $x_{ji} = 1$ if $u_j = i$. Thus $\sum_{i=1}^{\ell} x_{ji} = 1$ for all $j \in [n]$, and $\sum_{j=1}^{n} x_{ji} = \kappa_i$ for all $i \in [\ell]$. (When $\ell = n$, the multislice is the symmetric group S_n, and the x_{ji} are the entries of the permutation matrix representing the input permutation.)

Since $x_{j\ell} = 1 - \sum_{i=1}^{\ell-1} x_{ji}$, we do not need to include $x_{1\ell}, \ldots, x_{n\ell}$ explicitly as inputs. This is the usual convention in the case of the slice ($\ell = 2$), in which the input consists of just n Boolean variables $x_1, \ldots, x_n$.

2. Degree one functions

In this section we propose several different definitions of degree one functions, and show that they are all equivalent. While similar results hold for degree d functions for arbitrary d, we concentrate here on the case $d = 1$.

Throughout the section, we fix a multislice $\mathcal{U}_\kappa$ on n points and $\ell \geqslant 2$ colours.

2.1 Spectral definition

A partition of n is a non-increasing sequence of positive integers summing to n. We represent a partition as a finite sequence, or as an infinite sequence $(\lambda_i)_{i=1}^{\infty}$ where all but finitely many entries are zero. We can think of κ as a partition of n by sorting it accordingly. We say that a partition λ *majorizes* a partition μ, in symbols $\lambda \succeq \mu$, if $\lambda_1 + \cdots + \lambda_i \geqslant \mu_1 + \cdots + \mu_i$ holds for all $i \geqslant 1$.

The multislice $\mathcal{U}_\kappa$ can be viewed as a permutation module of the symmetric group. The representation theory of the symmetric group gives an orthogonal decomposition of the vector space of real-valued functions on the multislice:

$$\mathbb{R}^{\mathcal{U}_\kappa} = \bigoplus_{\lambda \succeq \kappa} V^\lambda,$$

where λ goes over all partitions of n majorizing κ. Furthermore, it is known that $V^{(n)}$ consists of all constant functions, and $V^{(n-1,1)}$ is spanned by functions of the form $x_{ji_1} - x_{ji_2}$ [16, Chapter 2].

Definition. A function on the multislice has *spectral degree one* if it lies in $V^{(n)} \oplus V^{(n-1,1)}$.

The orthogonal decomposition corresponds to the level decomposition of functions on the Boolean cube. In particular, we will use the following notation, for a function f on the multislice:

(1) $f^{=0}$ is the projection of f to $V^{(n)}$,
(2) $f^{=1}$ is the projection of f to $V^{(n-1,n)}$,
(3) $f^{\leqslant 1} = f^{=0} + f^{=1}$, and $f^{>1} = f - f^{\leqslant 1}$.

Since $V^{(n)}$ consists of all constant functions, $f^{=0}$ is the constant function $\mathbb{E}[f]$.

2.2 Polynomial definition

We can view the multislice as a function in the Boolean variables x_{ji}, where j ranges over $[n]$ and i ranges over $[\ell]$, given by $x_{ji} = 1_{u_j=i}$.

Definition. A function on the multislice has *polynomial degree one* if it can be represented as a polynomial of degree at most 1 in the variables x_{ji}.

Note that since $x_{j\ell} = 1 - \sum_{i=1}^{\ell-1} x_{ji}$, we can assume that the variables $x_{1\ell}, \ldots, x_{n\ell}$ do not appear in the polynomial representation.

Lemma 2.1. *A function on the multislice has spectral degree one if and only if it has polynomial degree one.*

Proof. If a function has spectral degree one, then it is a linear combination of the constant function 1 and functions of the form $x_{ji_1} - x_{ji_2}$, so it has polynomial degree one.

Conversely, suppose that f has polynomial degree one, so that

$$f = c + \sum_{j=1}^{n} \sum_{i=1}^{\ell} c_{ji} x_{ji}.$$

Let $c_j = \sum_{i=1}^{\ell} c_{ji}/\ell$ for all $j \in [n]$. Since $\sum_{i=1}^{\ell} x_{ji} = 1$, we have

$$f = c + \sum_{j=1}^{n} \ell c_j + \sum_{j=1}^{n} \sum_{i=1}^{\ell} (c_{ji} - c_j) x_{ji}.$$

By construction,

$$\sum_{i=1}^{\ell} (c_{ji} - c_j) = 0,$$

hence

$$c_{j\ell} - c_j = -\sum_{i=1}^{\ell-1} (c_{ji} - c_j).$$

Therefore

$$f = c + \sum_{j=1}^{n} \ell c_j + \sum_{j=1}^{n} \sum_{i=1}^{\ell-1} (c_{ji} - c_j)(x_{ji} - x_{j\ell}).$$

This shows that f has spectral degree one. $\qquad\square$

2.3 Junta definition

A *dictator* is a function depending on a single coordinate. This also includes constant functions.

Definition. A function on the multislice has *junta degree one* if it can be represented as a linear combination of dictators.

Lemma 2.2. *A function on the multislice has polynomial degree one if and only if it has junta degree one.*

Proof. The functions $1, x_{ji}$ are dictators, so if a function has polynomial degree one then it has junta degree one. Conversely, if f depends only on the jth coordinate, then $f = \sum_{i=1}^{\ell} c_i x_{ji}$ for some constants $c_1, \ldots, c_\ell$, so f has polynomial degree one. Therefore a function having junta degree one also has polynomial degree one. ☐

In view of Lemmas 2.1 and 2.2, we define a function on the multislice to have degree one if it satisfies any of the definitions given above.

2.4 Normal form

We close this section by describing a normal form for degree one functions.

Lemma 2.3. *Every degree one function on the multislice has a unique representation of the form*

$$f = c + \sum_{j=1}^{n} \sum_{i=1}^{\ell} c_{ji} x_{ji},$$

where $\sum_{i=1}^{\ell} c_{ji} = 0$ for all $j \in [n]$ and $\sum_{j=1}^{n} c_{ji} = 0$ for all $i \in [\ell]$.

Proof. We start by showing that if f has degree one then it has a representation as required by the lemma. By linearity, it suffices to show this for the function x_{11} (the lemma clearly holds for constant functions). Since $\sum_{i=1}^{\ell} x_{1i} = 1$, we have

$$x_{11} = \frac{1}{\ell} \sum_{i=1}^{\ell} (x_{11} - x_{1i}) + \frac{1}{\ell}.$$

Similarly, since $\sum_{j=1}^{n} x_{ji} = \kappa_i$, we have

$$x_{1i} = \frac{1}{n} \sum_{j=1}^{n} (x_{1i} - x_{ji}) + \frac{\kappa_i}{n}.$$

Combining both expressions, we obtain

$$x_{11} = \frac{1}{\ell n} \sum_{i=1}^{\ell} \sum_{j=1}^{n} (x_{11} - x_{j1} - x_{1i} + x_{ji}) + \frac{1}{\ell n} \sum_{i=1}^{\ell} (\kappa_1 - \kappa_i) + \frac{1}{\ell}$$

$$= \frac{1}{\ell n} \sum_{i=1}^{\ell} \sum_{j=1}^{n} (x_{11} - x_{j1} - x_{1i} + x_{ji}) + \frac{\kappa_1}{n}.$$

It is not hard to check that $x_{11} - x_{j1} - x_{1i} + x_{ji}$ satisfies the requisite properties for all $j' \in [n]$ and $i' \in [\ell]$, hence so does the expression given for x_{11}.

Next, we show that the representation is unique. It suffices to show that the only representation of the zero function is the zero polynomial. In other words, we have to show that if

$$0 = c + \sum_{j=1}^{n} \sum_{i=1}^{\ell} c_{ji} x_{ji},$$

where the c_{ji} satisfy all the constraints in the lemma, then $c = 0$ and $c_{ji} = 0$ for all $j \in [n]$ and $i \in [\ell]$.

Choose any two indices $j_1 \neq j_2$ and any two colours $i_1 \neq i_2$. Consider an arbitrary point u in the multislice satisfying $u_{j_1} = i_1$ and $u_{j_2} = i_2$, and the point obtained by switching i_1 and i_2. Subtracting the corresponding right-hand sides, we deduce

$$0 = c_{j_1 i_1} - c_{j_1 i_2} - c_{j_2 i_1} + c_{j_2 i_2}.$$

This identity also holds when $j_1 = j_2$. Averaging over all values of j_2 and using $\sum_{j=1}^{n} c_{ji_1} = \sum_{j=1}^{n} c_{ji_2} = 0$, we deduce that $c_{j_1 i_1} = c_{j_1 i_2}$ for all $i_1 \neq i_2$. Since $\sum_{i=1}^{\ell} c_{j_1 i} = 0$, this implies that $c_{j_1 i} = 0$ for all $i \in [\ell]$ and all $j_1 \in [n]$. It follows that also $c = 0$, so the only representation of zero is the zero polynomial, completing the proof of uniqueness. □

3. FKN theorem

In this section we prove Theorem 1.1, by induction on the number of colours. The actual statement that we will prove by induction is the following.

Theorem 3.1. *Fix an integer $\ell \geqslant 2$ and a parameter $\rho > 0$. There exists a constant $N = N(\ell, \rho)$ such that, for every ρ-balanced multislice on $n \geqslant N$ points and ℓ colours, the following holds.*

If $f \colon \mathcal{U}_\kappa \to \mathbb{R}$ is a degree one function which satisfies $\mathbb{E}\left[\mathrm{dist}\,(f, \{0, 1\})^2\right] = \varepsilon$, then there exists a Boolean dictator g such that $\mathbb{E}\left[(f - g)^2\right] = O_{\ell,\rho}(\varepsilon)$.

Theorem 1.1 follows from this formulation using the following argument.

Proof of Theorem 1.1. We start with the first part of the theorem. Let $f \colon \mathcal{U}_\kappa \to \mathbb{R}$ be a degree one function which satisfies $\mathbb{E}\left[\mathrm{dist}\,(f, \{0, 1\})^2\right] = \varepsilon$. Theorem 3.1 shows that there exists a Boolean dictator g such that $\mathbb{E}\left[(f - g)^2\right] = O_{\ell,\rho}(\varepsilon)$. Let $h = f - g$. Since g is Boolean, $\mathbb{E}\left[\mathrm{dist}\,(h, \{0, \pm 1\})^2\right] \leqslant \varepsilon$. When $|h| \leqslant 1/2$, we have $\mathrm{dist}\,(h, \{0, \pm 1\})^2 = h^2$, and hence

$$\varepsilon \geqslant \mathbb{E}\left[\mathrm{dist}\,(h, \{0, \pm 1\})^2\right] \geqslant \mathbb{E}\left[h^2 1_{|h| \leqslant 1/2}\right] = \mathbb{E}\left[h^2\right] - \mathbb{E}\left[h^2 1_{|h| > 1/2}\right].$$

When $|h| > 1/2$, we have $h^4 > h^2/4$, so

$$\mathbb{E}\left[h^2 1_{|h| > 1/2}\right] < 4\,\mathbb{E}\left[h^4\right].$$

In [9] it is shown that a ρ-biased multislice is hypercontractive for any constant ρ and constant number of colours, hence $\mathbb{E}\left[h^4\right] = O_{\ell,\rho}(\mathbb{E}\left[h^2\right]^2) = O_{\ell,\rho}(\varepsilon^2)$, since h has degree one. This shows that $\mathbb{E}\left[h^2\right] < \varepsilon + 4\,\mathbb{E}\left[h^4\right] = \varepsilon + O_{\ell,\rho}(\varepsilon^2)$, completing the proof of the first part of the theorem.

We continue with the second part of the theorem, which is very similar. Let $F \colon \mathcal{U}_\kappa \to \mathbb{R}$ be a Boolean function which satisfies $\|F^{>1}\|^2 = \varepsilon$, and let $f = F^{\leqslant 1}$. Since $\mathbb{E}\left[\mathrm{dist}\,(f, \{0, 1\})^2\right] \leqslant \mathbb{E}\left[(f - F)^2\right] = \varepsilon$, the first part gives a Boolean dictator g satisfying $\mathbb{E}\left[(f - g)^2\right] = \varepsilon + O_{\ell,\rho}(\varepsilon^2)$. The L^2 triangle inequality implies that $\mathbb{E}\left[(F - g)^2\right] \leqslant 4\varepsilon + O_{\ell,\rho}(\varepsilon^2)$. Since both F and g are Boolean, $\mathbb{P}\left[F \neq g\right] = \mathbb{E}\left[(F - g)^2\right]$, completing the proof. □

For brevity, in the rest of the section we use $O(\cdot)$ for $O_{\ell,\rho}(\cdot)$.

3.1 Base case

The base case of our inductive proof is when $\ell = 2$, and it follows from the main result of [7], whose statement reads as follows.

Theorem 3.2. *Suppose that* $f : \mathcal{U}_{k,n-k} \to \{0,1\}$ *satisfies* $\| f^{>1} \|^2 = \varepsilon$, *where* $2 \leqslant k \leqslant n/2$. *Then either* f *or* $1 - f$ *is* $O(\varepsilon)$-*close to a function of the form* $\max_{i \in S} x_{i1}$, *where* $S \subseteq [n]$ *has cardinality at most* $\max(1, O(\sqrt{\varepsilon}/(k/n)))$.

From this theorem, we deduce the base case of Theorem 3.1.

Proof of Theorem 3.1 in the case $\ell = 2$. Let $f : \mathcal{U}_{k,n-k} \to \mathbb{R}$ be a degree one function satisfying $\mathbb{E}\left[\text{dist}\,(f,\{0,1\})^2\right] = \varepsilon$, and assume without loss of generality that $k \leqslant n/2$. Let F be the function obtained by rounding f to $\{0,1\}$. By definition, $\mathbb{E}\left[(F - f)^2\right] = \varepsilon$, so $\|F^{>1}\|^2 \leqslant \varepsilon$ (since $F^{\leqslant 1}$ is the degree one function which is closest to F).

By choosing N appropriately, we can ensure that $k \geqslant 2$, and hence Theorem 3.2 applies, showing that either f or $1 - f$ is $O(\varepsilon)$-close to a function depending on at most $\max(1, m)$ coordinates, where $m = O(\sqrt{\varepsilon}/(k/n)) = O_\rho(\sqrt{\varepsilon})$.

We now consider two cases. The first case is when $m \leqslant 1$. In this case, F is $O(\varepsilon)$-close to a dictator. Since $\mathbb{E}\left[(F - f)^2\right] = \varepsilon$, it follows via the L^2 triangle inequality that f is also $O(\varepsilon)$-close to the same dictator.

When $m > 1$, we can lower-bound $\varepsilon \geqslant e_\rho$ for some constant $e_\rho > 0$ depending on ρ. The L^2 triangle inequality implies that

$$\mathbb{E}\left[f^2 1_{F=1}\right] = \mathbb{E}\left[((f-1)+1)^2 1_{F=1}\right] \leqslant 2\,\mathbb{E}\left[(f-1)^2 1_{F=1}\right] + 2.$$

Therefore

$$
\begin{aligned}
\varepsilon &= \mathbb{E}\left[\text{dist}\,(f,\{0,1\})^2\right] \\
&= \mathbb{E}\left[f^2 1_{F=0}\right] + \mathbb{E}\left[(f-1)^2 1_{F=1}\right] \\
&\geqslant \mathbb{E}\left[f^2 1_{F=0}\right] + \frac{1}{2}\,\mathbb{E}\left[f^2 1_{F=1}\right] - 1 \\
&\geqslant \frac{1}{2}\,\mathbb{E}\left[f^2\right] - 1.
\end{aligned}
$$

In other words, $\mathbb{E}\left[f^2\right] \leqslant 2(1 + \varepsilon)$. This implies that

$$\|f - 0\|^2 \leqslant \frac{2(1+\varepsilon)}{\varepsilon}\varepsilon \leqslant \frac{2(1 + e_\rho)}{e_\rho}\varepsilon,$$

completing the proof in this case. $\square$

3.2 Inductive step

We now assume that Theorem 3.1 holds for a certain value of $\ell \geqslant 2$, and will prove it for $\ell + 1$.

We start with a simple comment. Theorem 3.1 is trivial for large ε, using the same argument used to derive the second part of Theorem 1.1. Indeed, suppose that $\varepsilon \geqslant \varepsilon_0$. The aforementioned argument shows that

$$\mathbb{E}\left[(f - 0)^2\right] \leqslant \frac{2(1+\varepsilon)}{\varepsilon}\varepsilon \leqslant \frac{2(1 + \varepsilon_0)}{\varepsilon_0}\varepsilon.$$

Since 0 is a dictator, we see that when $\varepsilon \geqslant \varepsilon_0$, Theorem 3.1 trivially holds. Therefore, from now on we may assume that ε is small enough (as a function of ℓ and ρ).

Next, we need a criterion that guarantees that the approximating function in Theorem 3.1 is constant. We will use the concept of *influence*: given two coordinates $j_1, j_2 \in [n]$ and a function $f : \mathcal{U}_\kappa \to \mathbb{R}$,

$$\mathrm{Inf}_{j_1 j_2}[f] = \mathop{\mathbb{E}}_{u \sim \mathcal{U}_\kappa} \left[\left(f(u) - f(u^{(j_1 j_2)}) \right)^2 \right].$$

Lemma 3.3. *Let $\mathcal{U}_\kappa$ be a ρ-balanced multislice with ℓ colours. There exists a constant $\eta = \eta(\rho)$ such that the following holds for all $\varepsilon \leqslant \eta$. If $f : \mathcal{U}_\kappa \to \mathbb{R}$ is a degree one function which satisfies $\mathbb{E}\left[\mathrm{dist}\,(f, \{0,1\})^2\right] = \varepsilon$ and $\mathrm{Inf}_{j_1 j_2}[f] \leqslant \eta$ for all $j_1, j_2 \in [n]$ then there exists a constant $C \in \{0,1\}$ such that $\mathbb{E}\left[(f - C)^2\right] = O(\varepsilon)$ and $|\mathbb{E}[f] - C| = O(\sqrt{\varepsilon})$.*

Proof. Theorem 3.1 shows the existence of a Boolean dictator g satisfying $\mathbb{E}\left[(f - g)^2\right] = O(\varepsilon)$. The L^2 triangle inequality shows that $\mathrm{Inf}_{j_1 j_2}[g] = O(\mathrm{Inf}_{j_1 j_2}[f] + \varepsilon) = O(\eta)$. Suppose, for the sake of contradiction, that g is not constant. Then there exists a coordinate j_1 and colours i_1, i_2 such that $g(u) = 0$ if $u_{j_1} = i_1$ and $g(u) = 1$ if $u_{j_1} = i_2$. Let j_2 be any other coordinate. A random u chosen from the multislice satisfies $u_{j_1} = i_1$ and $u_{j_2} = i_2$ with probability larger than ρ^2. When that happens, $(g(u) - g(u^{(j_1 j_2)}))^2 = 1$. Therefore $\mathrm{Inf}_{j_1 j_2}[g] > \rho^2$. By choosing $\eta = \rho^2$, we reach a contradiction. We conclude that $g = C$ for some constant $C \in \{0,1\}$.

The L^1–L^2 norm inequality implies that $\mathbb{E}\left[|f - C|\right]^2 \leqslant \mathbb{E}\left[(f - C)^2\right] = O(\varepsilon)$, and hence $|\mathbb{E}[f] - C| \leqslant \mathbb{E}\left[|f - C|\right] = O(\sqrt{\varepsilon})$, completing the proof. $\qquad\square$

3.2.1 Isolating the dictatorial coordinate

The first step in the argument is to identify the dictatorial coordinate, if any. We do this by looking at the degree one expansion of f:

$$f = c + \sum_{j=1}^{n} \sum_{i=1}^{\ell} c_{ji} x_{ji}.$$

Note that although there are $\ell + 1$ colours, using the identity $x_{j(\ell+1)} = 1 - \sum_{i=1}^{\ell} x_{ji}$ we can eliminate all variables involving the last colour.

Let $j_1 \neq j_2$ be two arbitrary coordinates, and let $i \neq \ell + 1$ be an arbitrary colour. Suppose that u is an element of the multislice satisfying $u_{j_1} = i$ and $u_{j_2} = \ell + 1$. A short calculation shows that

$$f(u) - f(u^{(j_1 j_2)}) = c_{j_1 i} - c_{j_2 i}.$$

When choosing u at random from the multislice, the event $u_{j_1} = i$ and $u_{j_2} = \ell + 1$ occurs with probability larger than ρ^2. Therefore the L^2 triangle inequality implies that

$$4\varepsilon \geqslant \mathbb{E}\left[\mathrm{dist}\,(f(u) - f(u^{(j_1 j_2)}), \{0, \pm 1\})^2\right] > \rho^2 \, \mathrm{dist}\,(c_{j_1 i} - c_{j_2 i}, \{0, \pm 1\})^2,$$

implying that $\mathrm{dist}\,(c_{j_1 i} - c_{j_2 i}, \{0, \pm 1\}) = O(\sqrt{\varepsilon})$. Choosing $c_i := \min_j c_{ji}$, we deduce that $c_{ji} \in \{c_i, c_i + 1\} \pm O(\sqrt{\varepsilon})$ for all $j \in [n]$.

We associate with each coordinate $j \in [n]$ a vector $\gamma_j \in \{0,1\}^\ell$ such that $|c_{ji} - c_i - \gamma_{ji}| = O(\sqrt{\varepsilon})$. Assuming $n > 2^\ell$, there exists a vector $v \in \{0,1\}^\ell$ which is realized by at least two coordinates j_1, j_2. Our goal now is to show that v is realized by all but at most one coordinate. To this end, let us assume that $\gamma_{k_1}, \gamma_{k_2} \neq v$ for some coordinates $k_1 \neq k_2$. Let $i_1, i_2 \neq \ell + 1$ be colours such that $\gamma_{k_1 i_1} \neq v_{i_1}$ and $\gamma_{k_2 i_2} \neq v_{i_2}$.

If an element u of the multislice satisfies $\{u_{j_1}, u_{k_1}\} = \{i_1, \ell + 1\}$, then $f(u) - f(u^{(j_1 k_1)}) = \pm 1 \pm O(\sqrt{\varepsilon})$, and similarly for j_2, k_2. Hence we can find a constraint on $u_{j_1}, u_{k_1}, u_{j_2}, u_{k_2}$ which implies $f(u) - f(u^{(j_1 k_1)(j_2 k_2)}) = 2 \pm O(\sqrt{\varepsilon})$. For small enough ε, this guarantees that $\mathrm{dist}\,(f(u) -$

$f(u^{(j_1\,k_1)(j_2\,k_2)}), \{0,\pm1\})^2 \geqslant 1/2$. A random $u \sim \mathcal{U}_\kappa$ satisfies the constraint with probability $\Omega(\rho^4)$, so

$$4\varepsilon \geqslant \mathbb{E}\left[\text{dist}\left(f(u) - f(u^{(j_1\,k_1)(j_2\,k_2)}), \{0,\pm1\}\right)^2\right] = \Omega(\rho^4),$$

which is impossible if ε is small enough.

We conclude that $\gamma_j = v$ for all but at most a single coordinate. Without loss of generality, let the exceptional coordinate (if any) be the last coordinate.

3.2.2 Constant pieces

Our strategy now is to consider restrictions of f obtained by fixing the value of the last coordinate. For $n \geqslant 2/\rho$, fixing the last coordinate to colour $\chi \in [\ell + 1]$ will result in a function f_χ on a $(\rho/2)$-balanced multislice $\kappa^{(\chi)}$ on $n - 1$ points and $\ell + 1$ colours. Let $\varepsilon_\chi = \mathbb{E}\left[\text{dist}\left(f_\chi, \{0, 1\}\right)^2\right]$. We will show that each f_χ is nearly constant by applying Lemma 3.3, and later on put all the pieces together. Just as above, we can assume that ε_χ is small enough, since $\varepsilon \geqslant \rho\varepsilon_\chi$.

Let us start by noting that

$$f_\chi = c^{(\chi)} + \sum_{j=1}^{n-1} \sum_{i=1}^{\ell} c_{ji} x_{ji},$$

where the coefficients c_{ji} are the same as before. Suppose now that $S \subseteq [n-1]$ is a set of $\kappa_{\ell+1}^{(\chi)}$ coordinates, and let $S' = S \cup \{n\}$. Let $f_{\chi,S}$ be the function obtained by setting all coordinates in S to the value $\ell + 1$:

$$f_{\chi,S} = c^{(\chi)} + \sum_{j \notin S'} \sum_{i=1}^{\ell} c_{ji} x_{ji}.$$

This is a function on a $(\rho/2)$-balanced multislice $\kappa^{(\chi,S)}$ on ℓ colours, so we can apply Lemma 3.3. In preparation for such an application, let us define $\varepsilon_{\chi,S} = \mathbb{E}\left[\text{dist}\left(f_{\chi,S}, \{0, 1\}\right)^2\right]$.

By construction, for each $i \in [\ell]$ there exists a value $d_i \in \{c_i, c_i + 1\}$ such that $|c_{ji} - d_i| = O(\sqrt{\varepsilon})$ for all $j \in [n-1]$. This allows us to upper-bound $\text{Inf}_{j_1 j_2}[f_{i,S}]$ for all coordinates j_1, j_2. Indeed, if $u_{j_1} = i_1$ and $u_{j_2} = i_2$, then

$$|f_{\chi,S}(u) - f_{\chi,S}(u^{(j_1\,j_2)})| = |c_{j_1 i_1} + c_{j_2 i_2} - c_{j_1 i_2} - c_{j_2 i_1}| = O(\sqrt{\varepsilon}).$$

This shows that $\text{Inf}_{j_1 j_2}[f_\chi] = O(\varepsilon)$. For small enough ε, this allows us to apply Lemma 3.3 in order to conclude that there is a constant $C_{\chi,S} \in \{0, 1\}$ such that $\mathbb{E}\left[(f_{\chi,S} - C_{\chi,S})^2\right] = O(\varepsilon_{\chi,S})$ and $|\mathbb{E}[f_{\chi,S}] - C_{\chi,S}| = O(\sqrt{\varepsilon_{\chi,S}})$.

We apply the foregoing to a random choice S. The next step is to show that $C_{\chi,S}$ is concentrated. To this end, we calculate

$$\mathbb{E}[f_{\chi,S}] = c^{(\chi)} + \sum_{j \notin S'} \sum_{i=1}^{\ell} c_{ji} \frac{\kappa_i^{(\chi)}}{m},$$

where $m = \sum_{i=1}^{\ell} \kappa_i^{(\chi)}$. We can view $\mathbb{E}[f_{\chi,S}]$ as a function on the multislice $\mathcal{U}_{m,\kappa_{\ell+1}^{(\chi)}}$. Denoting it by μ and using a different parametrization of the slice, we have

$$\mu = c^{(\chi)} + \sum_{j=1}^{n-1} x_j \sum_{i=1}^{\ell} c_{ji} \frac{\kappa_i^{(\chi)}}{m}.$$

This is a degree one function, and it satisfies

$$\mathbb{E}\left[\text{dist}\left(\mu, \{0, 1\}\right)^2\right] \leqslant \mathbb{E}\left[(\mu - C_{\chi,S})^2\right] = O(\mathbb{E}[\varepsilon_{\chi,S}]) = O(\varepsilon_\chi).$$

Furthermore, for each $j_1 \neq j_2$ we have

$$\text{Inf}_{j_1 j_2}[\mu] \leqslant \left(\sum_{i=1}^{\ell} (c_{j_1 i} - c_{j_2 i}) \frac{\kappa_i^{(\chi)}}{m} \right)^2 = O(\varepsilon).$$

For small enough ε, we can thus apply Lemma 3.3 (for two colours) to deduce that $\mathbb{E}\left[(\mu - C_\chi)^2\right] = O(\varepsilon_\chi)$ for some constant $C_\chi \in \{0, 1\}$.

Without loss of generality, let us suppose that $C_\chi = 0$. Then $\mathbb{E}[\mu^2] = O(\varepsilon_\chi)$, and hence $\mathbb{P}[\mu \geqslant 1/2] = O(\varepsilon_\chi)$. Since $\mathbb{E}[\varepsilon_{\chi,S}] = \varepsilon_\chi$, also $\mathbb{P}[\varepsilon_{\chi,S} > \delta] = O(\varepsilon_\chi)$, for any constant $\delta > 0$. Recalling that $|\mu - C_{\chi,S}| = O(\sqrt{\varepsilon_{\chi,S}})$, this implies that with probability $1 - O(\varepsilon_\chi)$,

$$C_{\chi,S} \leqslant \mu + |\mu - C_{\chi,S}| \leqslant \frac{1}{2} + O(\delta) < 1,$$

for an appropriate choice of δ. This shows that $C_{\chi,S} = 1$ with probability $O(\varepsilon_\chi)$. Therefore

$$
\begin{aligned}
\mathbb{E}[f_\chi^2] &= \mathbb{E}_S\left[\mathbb{E}[f_{\chi,S}^2]\right] \\
&= \mathbb{E}_S\left[\mathbb{E}[(f_{\chi,S} - C_{\chi,S})^2 1_{C_{\chi,S}=0}]\right] + \mathbb{E}_S\left[\mathbb{E}[(f_{\chi,S} - C_{\chi,S} + 1)^2 1_{C_{\chi,S}=1}]\right] \\
&\leqslant 2\,\mathbb{E}_S\left[\mathbb{E}[(f_{\chi,S} - C_{\chi,S})^2]\right] + 2\,\mathbb{P}[C_{\chi,S} = 1] \\
&= O\left(\mathbb{E}_S[\varepsilon_{\chi,S}]\right) + O(\varepsilon_\chi) = O(\varepsilon_\chi).
\end{aligned}
$$

Also taking the case $C_\chi = 1$ into account, we deduce

$$\mathbb{E}[(f_\chi - C_\chi)^2] = O(\varepsilon_\chi), \quad C_\chi \in \{0, 1\}.$$

3.2.3 Completing the proof

We can now complete the proof of Theorem 3.1. Let $g(u) = C_{u_n}$, a Boolean dictator. Let χ be the marginal distribution of u_n when $u \sim \mathcal{U}_\kappa$. Then

$$\mathbb{E}[(f - g)^2] = \mathbb{E}_\chi\left[\mathbb{E}[(f_\chi - C_\chi)^2]\right] = O\left(\mathbb{E}_\chi[\varepsilon_\chi]\right) = O(\varepsilon).$$

This completes the proof.

4. Edge isoperimetry

Consider a multislice $\mathcal{U}_\kappa$ on at least four points. Define the *volume* of a subset A of the multislice $\mathcal{U}_\kappa$ to be $\text{vol}(A) = |A|/|\mathcal{U}_\kappa|$. The goal of this section is to prove the following isoperimetric inequality: if $\text{vol}(A) = \alpha$ then

$$\Phi(A) \geqslant \frac{2(1 - \alpha)}{n - 1}.$$

We will also identify when this inequality is tight, and prove stability in these cases.

4.1 Spectral formula

For a partition $\lambda \succeq \kappa$, let $f^{=\lambda}$ denote the orthogonal projection of f to V^λ (see Section 2.1 for the appropriate definitions). Frobenius [11] proved the following formula:

$$\mathbb{E}_{\tau \sim \text{Trans}(n)}[f^\tau] = \sum_{\lambda \succeq \kappa} c_\lambda f^{=\lambda}, \quad \text{where } c_\lambda = \frac{1}{n(n-1)} \sum_{i=1}^{\ell} [\lambda_i^2 - (2i - 1)\lambda_i]. \tag{4.1}$$

A classical fact is that $c_\lambda > c_\mu$ if $\lambda \succ \mu$; see for example [2, Lemma 10]. This allows us to identify the minimal values of $1 - c_\lambda$.

Lemma 4.1. *We have $1 - c_{(n)} = 0$, $1 - c_{(n-1,1)} = 2/(n-1)$, and $1 - c_\lambda \geqslant 4/n$ for all $\lambda \neq (n)$, $(n-1,1)$.*

Proof. The largest three partitions in majorization order are (n), $(n-1,1)$, $(n-2,2)$. Calculation shows that $c_{(n-2,2)} = 4/n$, so the lemma follows from the observation that $c_\lambda > c_\mu$ if $\lambda \succ \mu$. □

The important formula of Frobenius allows us to deduce one for $\Phi(A)$.

Lemma 4.2. *For any $A \subseteq \mathcal{U}_\kappa$,*

$$\Phi(A) = \frac{1}{\mathrm{vol}\,(A)} \sum_{\lambda \unrhd \kappa} (1 - c_\lambda)\|1_A^{=\lambda}\|^2.$$

Proof. For a given element $u \in A$ and a given transposition $\tau \in \mathrm{Trans}(n)$, the element u^τ lies in A if $\mathbb{E}\,[1_{u^\tau} 1_A] = 1/|\mathcal{U}_\kappa|$, and otherwise $\mathbb{E}\,[1_{u^\tau} 1_A] = 0$. Hence

$$\mathop{\mathbb{P}}_{u \sim A}\,[u^\tau \in A] = \frac{|\mathcal{U}_\kappa|}{|A|}\,\mathbb{E}\,[1_A^\tau 1_A] = \frac{1}{\mathrm{vol}\,(A)}\langle 1_A^\tau, 1_A \rangle.$$

Averaging over τ, we get

$$\mathop{\mathbb{P}}_{\substack{u \sim A \\ \tau \sim \mathrm{Trans}(n)}}\,[u^\tau \in A] = \frac{1}{\mathrm{vol}\,(A)}\Big\langle \mathop{\mathbb{E}}_{\tau \sim \mathrm{Trans}(n)}\,[1_A^\tau], 1_A \Big\rangle.$$

Applying (4.1) and the orthogonality of the isotypical decomposition, we obtain

$$\mathop{\mathbb{P}}_{\substack{u \sim A \\ \tau \sim \mathrm{Trans}(n)}}\,[u^\tau \in A] = \frac{1}{\mathrm{vol}\,(A)} \sum_{\lambda \unrhd \kappa} \langle c_\lambda 1_A^{=\lambda}, 1_A^{=\lambda} \rangle = \frac{1}{\mathrm{vol}\,(A)} \sum_{\lambda \unrhd \kappa} c_\lambda \|1_A^{=\lambda}\|^2.$$

The lemma now follows from the identity $\mathrm{vol}\,(A) = \|1_A\|^2 = \sum_\lambda \|1_A^{=\lambda}\|^2$. □

4.2 Main argument

Lemma 4.2 implies an isoperimetric inequality, along the lines of Hoffman's bound.

Proposition 4.3. *If $A \subseteq \mathcal{U}_\kappa$ and $\mathrm{vol}\,(A) = \alpha$, then*

$$\Phi(A) \geqslant \frac{2(1-\alpha)}{n-1}.$$

Furthermore, if $\kappa_1, \ldots, \kappa_\ell \geqslant 2$ and the bound is tight, then A is a dictator (membership in A depends on the colour of a single coordinate).

Suppose now that the number of colours is bounded, and that the multislice is ρ-balanced for some constant ρ. If

$$\Phi(A) = (1 + \varepsilon)\frac{2(1-\alpha)}{n-1}$$

(where $\varepsilon > 0$), then there exists a Boolean dictator B such that

$$\frac{|A \triangle B|}{|\mathcal{U}_\kappa|} = O(\alpha(1-\alpha)\varepsilon).$$

Proof. Since $1_A^{=(n)} = \mathbb{E}\,[1_A]1 = \mathrm{vol}\,(A)1$, it follows that $\|1_A^{=(n)}\|^2 = \mathrm{vol}\,(A)^2$. Similarly,

$$\sum_{\lambda \succeq \kappa} \|1_A^{=\lambda}\|^2 = \|1_A\|^2 = \mathbb{E}\,[1_A] = \mathrm{vol}\,(A).$$

Hence, combining Lemmas 4.2 and 4.1, we have

$$\Phi(A) \geqslant \frac{1}{\mathrm{vol}\,(A)} \cdot \frac{2}{n-1} (\,\mathrm{vol}\,(A) - \mathrm{vol}\,(A)^2) = \frac{2(1 - \mathrm{vol}\,(A))}{n-1}.$$

This proves the upper bound. If the upper bound is tight, then 1_A is supported on $(n), (n-1, 1)$, so 1_A has degree one. This implies [8] that A is a dictator.

Suppose now that the number of colours is bounded, that the multislice is ρ-balanced, and that

$$\Phi(A) = (1 + \varepsilon)\frac{2(1 - \alpha)}{n-1}.$$

Let

$$\delta = \|1_A\|^2 - \|1_A^{=(n)}\|^2 - \|1_A^{=(n-1,1)}\|^2.$$

Then

$$\Phi(A) \geqslant \frac{1}{\mathrm{vol}\,(A)} \cdot \left[\frac{2}{n-1} (\,\mathrm{vol}\,(A) - \mathrm{vol}\,(A)^2) + \frac{2(n-2)}{n(n-1)}\delta \right]$$

$$= \frac{2(1 - \alpha)}{n-1} + \frac{2(n-2)}{n(n-1)}\frac{\delta}{\alpha}.$$

The assumption on $\Phi(A)$ thus implies an upper bound on δ:

$$\delta \leqslant \frac{n\alpha(1 - \alpha)}{(n-2)}\varepsilon = O(\alpha(1 - \alpha)\varepsilon).$$

Theorem 1.1 shows that if n is larger than some constant depending on ℓ and ρ, then 1_A is $O(\delta)$-close to a Boolean dictator 1_B, completing the proof. When n is small, compactness shows that 1_A is trivially $O(\delta)$-close to $\emptyset$, since there are only finitely many possible A, α, ε. $\square$

Corollary 4.4. *Suppose that $\alpha \in (0, 1)$ satisfies $\alpha n = \sum_{i \in S} \kappa_i$ for some $S \subseteq [\ell]$. Then the bound in Lemma 4.3 is tight for the families*

$$A_{j,S} = \{u : u_j \in S\}, \quad j \in [n].$$

Conversely, if the bound in Lemma 4.3 is tight for a family A, then there exists a set $S \subseteq [\ell]$ satisfying $\alpha n = \sum_{i \in S} \kappa_i$ and a coordinate $j \in [n]$ such that $A = A_{j,S}$.

Proof. The expansion of $A_{j,S}$ is the probability that a random transposition is of the form $(j\,k)$, where k is one of the $(1 - \alpha)n$ coordinates whose colour is not in S. Therefore

$$\Phi(A_{j,S}) = \frac{(1 - \alpha)n}{\binom{n}{2}} = \frac{2(1 - \alpha)}{n-1}.$$

This shows that the bound in Lemma 4.3 is tight for $A_{j,S}$.

Conversely, if the bound in Lemma 4.3 is tight for a family A, then A is a dictator and thus of the form $A_{j,S}$. Since $\mathrm{vol}\,(A_{j,S}) = \sum_{i \in S} \kappa_i/n$, we see that $\sum_{i \in S} \kappa_i = \alpha n$. $\square$

References

[1] Alon, N., Dinur, I., Friedgut, E. and Sudakov, B. (2004) Graph products, Fourier analysis and spectral techniques. *Geom. Funct. Anal.* **14** 913–940.

[2] Diaconis, P. and Shahshahani, M. (1981) Generating a random permutation with random transpositions. *Z. Wahrsch. Verw. Gebiete* **57** 159–179.

[3] Dikstein, Y., Dinur, I., Filmus, Y. and Harsha, P. (2018) Boolean function analysis on high-dimensional expanders. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques (APPROX/RANDOM 2018)*, Vol. 116 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 38:1–38:20.

[4] Ellis, D., Filmus, Y. and Friedgut, E. (2015) A quasi-stability result for dictatorships in S_n. *Combinatorica* **35** 573–618.

[5] Ellis, D., Filmus, Y. and Friedgut, E. (2015) A stability result for balanced dictatorships in S_n. *Random Struct. Alg.* **46** 494–530.

[6] Ellis, D., Friedgut, E. and Pilpel, H. (2011) Intersecting families of permutations. *J. Amer. Math. Soc.* **24** 649–682.

[7] Filmus, Y. (2016) Friedgut–Kalai–Naor theorem for slices of the Boolean cube. *Chic. J. Theoret. Comput. Sci.* **2016** 14:1–14:17.

[8] Filmus, Y. and Ihringer, F. (2019) Boolean degree 1 functions on some classical association schemes. *J. Combin. Theory Ser. A* **162** 241–270.

[9] Filmus, Y., O'Donnell, R. and Wu, X. (2018) A log-Sobolev inequality for the multislice, with applications. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, Vol. 124 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 34:1–34:12.

[10] Friedgut, E., Kalai, G. and Naor, A. (2002) Boolean functions whose Fourier transform is concentrated on the first two levels and neutral social choice. *Adv. Appl. Math.* **29** 427–437.

[11] Frobenius, G. (1900) Über die Charaktere der symmetrischen Gruppe. *Sitzungsberichte der Königlich preussischen Akademie der Wissenschaften zu Berlin*, pp. 516–534.

[12] Jendrej, J., Oleszkiewicz, K. and Wojtaszczyk, J. O. (2015) On some extensions of the FKN theorem. *Theory Comput.* **11** 445–469.

[13] Nayar, P. (2014) FKN theorem on the biased cube. *Colloq. Math.* **137** 253–261.

[14] O'Donnell, R. (2014) *Analysis of Boolean Functions*, Cambridge University Press.

[15] Rubinstein, A. (2012) Boolean functions whose Fourier transform is concentrated on pairwise disjoint subsets of the input. Master's thesis, Tel Aviv University, Israel.

[16] Sagan, B. E. (2001) *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*, second edition, Vol. 203 of Graduate Texts in Mathematics, Springer.