Sparse juntas on the biased hypercube

Irit Dinur, Yuval Filmus, Prahladh Harsha

May 26, 2023

Abstract

We give a structure theorem for Boolean functions on the biased hypercube which are ϵ-close to degree d in L_2, showing that they are close to *sparse juntas*. Our structure theorem implies that such functions are $O(\epsilon^{C_d} + p)$-close to constant functions. We pinpoint the exact value of the constant C_d.

1 Introduction

Let $f\colon \{0,1\}^n \to \{0,1\}$ be a Boolean function on the hypercube, where $\{0,1\}^n$ is endowed with the uniform measure. If f has degree 1 then it is a dictator, and if it has degree d then it is a junta [NS94]. Friedgut, Kalai and Naor [FKN02] showed that if f is ϵ-close to degree 1 (that is, $\|f^{>1}\|^2 \leq \epsilon$) then f is $O(\epsilon)$-close to a dictator (that is, $\Pr[f \neq g] = O(\epsilon)$ for some dictator g). Kindler and Safra [KS02, Kin03] (see also [KK20]) showed that if f is ϵ-close to degree d then f is $O(\epsilon)$-close to a $2^{O(d)}$-junta.

Kindler and Safra proved their result in the more general setting of the biased hypercube, which is $\{0,1\}^n$ endowed with the biased measure μ_p. They showed that if f is a Boolean function on the p-biased hypercube which is ϵ-close to degree d then f is $O(\epsilon + \epsilon^{5/4}/p^{4d})$-close to an $O(1/p^{4d})$-junta, assuming $p \leq 1/2$. When p is constant, this gives the same statement as before, but when p is small, the result is no longer tight.

The biased hypercube in the small p regime is interesting for at least two reasons. First, it is the regime which is most relevant for understanding threshold phenomena in the Erdős–Rényi random graph model $G(n,p)$. Second, it is a good model for other domains with sparse inputs, such as the symmetric group and the Grassmann scheme, the latter of which was recently used to prove the 2-to-2 games conjecture [KMS23].

In this paper, we extend the FKN and Kindler–Safra theorems to functions on the biased hypercube, proving a statement which does not deteriorate as p gets small. For every d, ϵ and $p \leq 1/2$ we find an explicit set of Boolean functions $\mathcal{F}_{d,\epsilon,p}$ such that

(a) If $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is ϵ-close to degree d then f is $O(\epsilon)$-close to a function in $\mathcal{F}_{d,\epsilon,p}$.

(b) Every function in $\mathcal{F}_{d,\epsilon,p}$ is $O(\epsilon)$-close to degree d with respect to μ_p.

Let us demonstrate what we mean by *explicit* using the case $d = 1$, which appears in earlier work of Filmus [Fil16]. Filmus showed that we can take

$$\mathcal{F}_{1,\epsilon,p} = \left\{ \bigvee_{i \in S} y_i, \neg \bigvee_{i \in S} y_i : |S| \leq \max(1, C\sqrt{\epsilon}/p) \right\},$$

for some constant C, where $(y_1, \ldots, y_n) \in \{0,1\}^n$ is the input point.

The functions in $\mathcal{F}_{1,\epsilon,p}$ are $O(\sqrt{\epsilon})$-close to dictators and so $O(\sqrt{\epsilon} + p)$-close to constants, and the same holds for every Boolean function which is ϵ-close to degree 1. This shows that if we consider Boolean functions which are ϵ-close to degree 1 up to an error of magnitude $\sqrt{\epsilon}$, then all we see is dictators. If we look closer, narrowing the magnitude of the error to ϵ, then a more nuanced picture emerges. We will be interested in both points of view. We state our results formally in Section 1.1.

The proof of Kindler and Safra deteriorates with p due to the use of hypercontractivity. In contrast, our proof works by reduction to the unbiased setting, and so does not suffer from this deterioration. In more detail, we express the biased hypercube as a weighted average of unbiased hypercubes of smaller dimension:

$$(\{0,1\}^n, \mu_p) \approx (\{0,1\}^{\mathbf{S}}, \mu_{1/2}), \text{ where } \mathbf{S} \sim \mu_{2p}.$$

We apply the *unbiased* Kindler–Safra theorem on each unbiased hypercube $\{0,1\}^S$, and combine the results using an *agreement theorem*. We describe our methods in more detail in Section 1.2.

1.1 Main results

Here is the generalization of the result of Filmus [Fil16] to arbitrary d. The statement uses various terms which we explain after it. The statement is heavily inspired by the FKN theorem for functions on the symmetric group due to Filmus [Fil21].

Theorem 1.1 (Main). *Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is ϵ-close to degree d, where $p \leq 1/2$. Then f is $O(\epsilon)$-close to* round$(g, \{0,1\})$*, where g is a degree d polynomial satisfying the following properties, for some constant C depending only on d:*

(i) *If $y \in \{0,1\}^n$ has weight at most d then $g(y) \in \{0,1\}$.*

(ii) *For every T and e, the monomial support of g contains at most C/p^e sets $U \supseteq T$ of size $|T| + e$.*

(iii) *For all e, the number of inputs $y \in \{0,1\}^n$ of weight e such that $g(y) \notin \{0,1\}$ but $g(z) \in \{0,1\}$ for all $z < y$ is at most $C\epsilon/p^e$.*

Conversely, if g is a degree d polynomial satisfying these properties then round$(g, \{0,1\})$ *is $O(\epsilon)$-close to degree d.*

Rather than specifying the list $\mathcal{F}_{d,\epsilon,p}$ directly, we instead specify a list $\mathcal{G}_{d,\epsilon,p}$ of degree d functions. The functions in $\mathcal{F}_{d,\epsilon,p}$ are obtained by rounding the functions in $\mathcal{G}_{d,\epsilon,p}$ to Boolean. The corresponding list for $d = 1$ is

$$\mathcal{G}_{1,\epsilon,p} = \left\{ \sum_{i \in S} y_i, 1 - \sum_{i \in S} y_i : |S| \leq \max(1, C\sqrt{\epsilon}/p) \right\}.$$

A degree d polynomial g is in $\mathcal{G}_{d,\epsilon,p}$ if it satisfies the three properties listed in the theorem. To explain these properties, let us introduce the *monomial expansion* of g, which is its unique expansion as a linear combination of monomials $y_S = \prod_{i \in S} y_i$. We denote the coefficient of y_S by $\tilde{g}(S)$; these are the *monomial coefficients* of g.

The monomial expansion differs from the more usual Fourier expansion, in which instead of the monomials y_S we have the Fourier characters $\omega_S = \prod_{i \in S} \frac{y_i - p}{\sqrt{p(1-p)}}$. Using the monomial expansion guarantees that the functions in $\mathcal{G}_{d,\epsilon,p}$ have the same form regardless of p (cf. $\sum_{i \in S} y_i$ to $\sum_{i \in S} \sqrt{p(1-p)} \frac{y_i - p}{\sqrt{p(1-p)}} + p|S|$). Concretely, the monomial coefficients appearing in the functions in $\mathcal{G}_{d,\epsilon,p}$ depend only on d, a property which fails for the Fourier expansion.

Item i states that $g(y) \in \{0,1\}$ whenever y has (Hamming) weight at most d. Equivalently, if $|S| \leq d$ then

$$\sum_{T \subseteq S} \tilde{g}(T) \in \{0,1\}.$$

This implies that the monomial coefficients of g are quantized. As an example, when $d = 1$, the constant coefficient is either 0 or 1; in the former case, all other coefficients are 0 or 1; and in the latter case, all other coefficients are 0 or -1.

Item ii concerns the *monomial support* of g, which is the set

$$\text{supp}(g) = \{S \subseteq [n] : \tilde{g}(S) \neq 0\}.$$

Choosing $T = \emptyset$, the property states that the expected number of monomials in $\mathrm{supp}(g)$ which evaluate to 1 under a random input is $O(1)$. Considering other T, the same holds even under the condition $y_T = 1$. For these reasons, we call a polynomial satisfying this property a *sparse junta*.

This item fails for the Fourier expansion. For example, the Fourier support of the function $\sum_{i=1}^{C/p^2} y_{2i-1} y_{2i}$ contains $2C/p^2$ singletons rather than just $O(1/p)$.

Item iii concerns *minimal non-Boolean inputs*. An input $y \in \{0,1\}^n$ is a minimal non-Boolean input of g if $g(y) \notin \{0,1\}$ but $g(z) \in \{0,1\}$ for all inputs z strictly below y, that is, obtained from y by switching one or more 1s to 0s. Equivalently, if $y = 1_S$ then the restriction $g|_S$ obtained by zeroing out all coordinates outside of S is Boolean except for the input 1_S. Due to the second property, every minimal non-Boolean input occurring in g has weight at least $d + 1$.

We can think of a minimal non-Boolean input as a forbidden configuration, which we can describe by specifying the function $g|_S$. When $d = 1$, there are two minimal non-Boolean inputs, corresponding to the functions

$$y_1 + y_2 \text{ and } 1 - y_1 - y_2.$$

Accordingly, if $\tilde{g}(\emptyset) = 0$ and $|\mathrm{supp}(g)| = m$ then $\binom{m}{2} = O(\epsilon/p^2)$, which implies that either $m \leq 1$ or $m = O(\sqrt{\epsilon}/p)$.

We list all minimal non-Boolean inputs for $d = 2$ in Section 2.2, where we also show that there is a finite number of them for every d.

Monotone version

When the function f is monotone, we can guarantee that the function g is monotone as well.

Theorem 1.2. *Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is a monotone function which is ϵ-close to degree d, where $p \leq 1/2$. Then f is $O(\epsilon)$-close to a monotone width d DNF g, where g satisfies the following properties, for some constant C depending only on g:*

(i) *For every T and e, the DNF g contains at most C/p^e minterms of $U \supseteq T$ of size $|U| + e$.*

(ii) *For all e, the number of sets S of size e such that $\deg(g|_S) > d$ but $\deg(g|_T) \leq d$ for all $T \subsetneq S$ is at most $C\epsilon/p^e$.*

Conversely, if g is a monotone width d DNF satisfying these properties then it is $O(\epsilon)$-close to degree d.

When $d = 1$, the unique forbidden configuration (corresponding to **Item ii**) is $y_1 \vee y_2$. We list all forbidden configurations for $d = 2$ in Section 2.2, where we also show that there is a finite number of them for every d.

Junta approximation

The FKN theorem of Filmus [Fil16] implies that if $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is ϵ-close to degree 1 then it is $O(\sqrt{\epsilon})$-close to a dictator and $O(\sqrt{\epsilon} + p)$-close to a constant. Theorem 1.1 implies a similar result for arbitrary d, where $\sqrt{\epsilon}$ is replaced by an appropriate power of ϵ.

Theorem 1.3 (Junta approximation). *Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is ϵ-close to degree d, where $p \leq 1/2$. There are parameters m, μ, depending only on d, such that:*

(a) *f is $O(\epsilon^{1/\mu})$-close to a Boolean degree d function.*

(b) *f is $O(\epsilon^{1/m} + p)$-close to a Boolean constant function.*

When $d = 1$, the optimal value of both m and μ is 2, which is tight when f is obtained by rounding the following function to Boolean:

$$g = y_1 + \cdots + y_{\sqrt{\epsilon}/p}.$$

3

We can generalize this construction to arbitrary d. Let P be a degree d polynomial such that $P(0), \ldots, P(m-1) \in \{0, 1\}$ and $P(0) \neq P(1)$ (von zur Gathen and Roche [GR97] studied such non-constant polynomials without the restriction $P(0) \neq P(1)$), and take the degree d polynomial

$$g = P(y_1 + \cdots + y_{\epsilon^{1/m}/p}).$$

The input to P behaves roughly like a Poisson random variable with expectation $\lambda = \epsilon^{1/m}$, and so $\Pr[g \notin \{0, 1\}] = O(\lambda^m) = O(\epsilon)$. This implies that $\mathrm{round}(g, \{0, 1\})$ is $O(\epsilon)$-close to degree d.

Since $P(1) \neq P(0)$, a similar calculation shows that $\Pr[g = P(0)] \approx 1 - \Theta(\lambda)$ and $\Pr[g = P(1)] \approx \Theta(\lambda)$, and so g is $\Theta(\lambda)$-far from all constant functions.

Using a Ramsey-theoretic argument, we show that this kind of construction is best possible for m.

Theorem 1.4. *Given d, the optimal value of m in Theorem 1.1 is the maximum m such that there exists a degree d polynomial P satisfying $P(0), \ldots, P(m-1) \in \{0, 1\}$ and $P(0) \neq P(1)$.*

As an example, when $d = 2$ we get $m = 4$, corresponding to the polynomial $P(z) = \frac{z(3-z)}{2}$.

It is an intriguing open question to find the optimal value of μ. Another interesting exponent is the best ν such that f is $O(\epsilon^{1/\nu})$-close to some Boolean junta or to some degree d junta (both exponents are the same).

1.2 Proof sketch

We prove Theorem 1.1 by reduction to the Kindler–Safra theorem, in the following form.

Theorem 1.5 (Kindler–Safra). *If $f \colon (\{0, 1\}^n, \mu_{1/2}) \to \{0, 1\}$ is ϵ-close to degree d then there is a Boolean degree d function g which is $O(\epsilon)$-close to f.*

(While Kindler and Safra did not state their theorem in this way, this formulation easily follows from their result.)

Suppose we are given a function $f \colon (\{0, 1\}^n, \mu_p) \to \{0, 1\}$ which is ϵ-close to degree d. We look at restrictions $f|_S$ obtained by zeroing out the coordinates outside of S. If $S \sim \mu_{2p}$ then on average, $f|_S$ is close to degree d with respect to the uniform measure on $\{0, 1\}^S$:

$$\mathop{\mathbb{E}}_{S \sim \mu_{2p}} \left[\mathop{\mathbb{E}}_{\mu_{1/2}} [(f|_S - f^{\leq d}|_S)^2] \right] = \mathop{\mathbb{E}}_{\mu_p} [(f - f^{\leq d})^2] \leq \epsilon.$$

Indeed, if $S \sim \mu_{2p}$ and $y \sim \mu_{1/2}(\{0, 1\}^S)$ then the marginal distribution of y is μ_p.

Applying the Kindler–Safra theorem for every $f|_S$, we obtain a Boolean degree d function g_S such that

$$\mathop{\mathbb{E}}_{S \sim \mu_{2p}} \left[\mathop{\mathbb{E}}_{\mu_{1/2}} [(f|_S - g_S)^2] \right] = O(\epsilon). \tag{1}$$

We would like to construct the function g by pasting together the various functions g_S, and for that we need to know that the g_S agree with each other, on average. This sort of pasting is achieved using *agreement theorems*, and in this case we use the following theorem, with K being the maximal size of the monomial support of a Boolean degree d function, $q = 2p$, and $c = \sqrt{1/2}$.

Theorem 1.6 (Junta agreement theorem). *Fix the following parameters: integer K (junta size), $q \in (0, 1)$ (bias), $c \in (0, 1)$ (fractional intersection size).*

For each $S \subseteq [n]$, let g_S be a degree d polynomial whose monomial support contains at most K monomials. Suppose that the following agreement condition holds:

$$\mathop{\Pr}_{\substack{T \sim \mu_{cq} \\ T_1, T_2 \sim \mu_r}} [g_{T \cup T_1}|_T \neq g_{T \cup T_2}|_T] \leq \epsilon, \ \textit{where } r = \frac{(1-c)q}{1-cq}.$$

(The value of r is chosen so that the marginal distribution of $S_1 := T \cup T_1$ and $S_2 := T \cup T_2$ is μ_q.)

4

Then there exists a degree d polynomial g such that

$$\Pr_{S \sim \mu_q} [g|_S \neq g_S] = O_{K,c}(\epsilon).$$

(The error bound depends on K and c but is independent of q.)

Moreover, for each A, the monomial coefficient $\tilde{g}(A)$ is chosen by majority decoding: it maximizes

$$\Pr_{\substack{S \sim \mu_{2p} \\ S \supseteq A}} [\tilde{g}_S(A) = \tilde{g}(A)].$$

The junta agreement theorem follows from the more general agreement theorem in our earlier work [DFH17], which does not require a bound on the size of the monomial support. However, since the proof of the junta agreement theorem is much easier than the proof of the general agreement theorem, we include it in this paper.

In order to apply Theorem 1.6, we need to show that the condition in the theorem holds. Since $f_{S_1}|_T = f_{S_2}|_T$, Equation (1) implies that

$$\mathbb{E}_{S_1,S_2} [(g_{S_1}|_T - g_{S_2}|_T)^2] = O(\epsilon).$$

Since g_{S_1}, g_{S_2} are both juntas, either $g_{S_1}|_T = g_{S_2}|_T$ or $\Pr[g_{S_1}|_T \neq g_{S_2}|_T] = \Omega(1)$, and this implies the agreement condition of Theorem 1.6.

Applying Theorem 1.6, we obtain a function g such that

$$\Pr_{S \sim \mu_p} [g|_S \neq g_S] = O(\epsilon).$$

A short argument now shows that $\Pr[f \neq g] = O(\epsilon)$.

Since g is obtained by majority decoding and each g_S is a junta, we can show that g is a sparse junta, proving **Item ii**. This allows us to show that $\mathbb{E}_{\mu_p}[(f - g)^2] = O(\epsilon)$, a step which we highlight below.

Next, let us show why **Item iii** holds. Let 1_S be a minimal non-Boolean input of g. Since g is a sparse junta, the probability that $y_S = 1$ and $y_T = 0$ for all other minimal non-Boolean inputs 1_T is $\Omega(p^{|S|})$. These events are disjoint, and so

$$\Pr[f \neq g] \geq \Pr[g \notin \{0,1\}] \geq \sum_S \Omega(p^{|S|}),$$

where the sum goes over all minimal non-Boolean inputs of g. This implies Item i.

Finally, let us address **Item i**. The original function g doesn't necessarily satisfy it. However, an argument along the lines of the preceding paragraph shows that the number of offending inputs is small, allowing us to slightly perturb g so that it satisfies Item i.

From L_0 to L_2 using the reverse union bound In the proof of Theorem 1.1, we deduce the L_2 guarantee $\mathbb{E}_{\mu_p}[(f - g)^2] = O(\epsilon)$ from the L_0 guarantee $\Pr[f \neq g] = O(\epsilon)$ using the fact that g is a sparse junta. The argument uses a technical tool, the *reverse union bound*, which we would like to highlight here.

Since $\Pr[f \neq g] = O(\epsilon)$ and f is Boolean, it follows that $\Pr[g \notin \{0,1\}] = O(\epsilon)$, and so $\Pr[g(g-1) \neq 0] = O(\epsilon)$. We will deduce that $\mathbb{E}_{\mu_p}[g^2(g-1)^2] = O(\epsilon)$, which implies $\mathbb{E}_{\mu_p}[(f - g)^2] = O(\epsilon)$ via a short argument.

Since g is a sparse junta whose monomial coefficients are quantized, the same holds for $h = g(g - 1)$, and also for h^2. We will show that $\Pr[h \neq 0] = O(\epsilon)$ implies that $\mathbb{E}_{\mu_p}[h^2] = O(\epsilon)$.

Since the monomial coefficients of h^2 are quantized, it suffices to show that the monomial support of h^2 is sparse, in the sense that it contains $O(\epsilon/p^e)$ sets of size e for all $e \leq 2d$. Since h is a sparse junta, in order to show that the monomial support of h^2 is sparse, it suffices to show that the monomial support of h itself is sparse.

The union bound shows that

$$\Pr[h \neq 0] \leq \sum_{S \in \text{supp}(h)} p^{|S|}.$$

5

If the inequality were in the other direction then it would follow that the monomial support of h is sparse, since $\Pr[h \neq 0] = O(\epsilon)$. In general, we cannot reverse the inequality (consider for example $h = \sum_{i=1}^{n} y_i$, where the left-hand side tends to 1 while the right-hand side tends to infinity for constant p). However, when h is a sparse junta, the *reverse union bound* states that we can indeed reverse the direction of the inequality:

$$\Pr[h \neq 0] = \Omega\left(\sum_{S \in \text{supp}(h)} p^{|S|}\right).$$

The proof of the reverse union bound is not difficult. It uses the Harris–FKG inequality.

Differences from conference version

This paper originates in an extended abstract [DFH19], which covers both the material eventually published in [DFH17] and part of the material in the present work. The full version of [DFH19] is available on arXiv (1711.09428) and ECCC (TR17-180).

The full version of [DFH19] contains weaker versions of Theorems 1.1 and 1.2, which describe only some properties of the approximating function g, and in particular do not constitute a full characterization of all Boolean functions which are close to degree d. The current versions of these theorems are inspired by [Fil21]. In addition, Theorem 1.4 is completely new. It is inspired by [Fil23].

In contrast, the main theorems in the full version of [DFH19] apply in the more general *A-valued* setting, which we describe briefly in Section 7, as well as in the setting of the slice. We omit these generalizations in this version to make it more concise, and since these generalizations only require a few new ideas, the most interesting of which is that whereas the monomial expansion of functions on the slice is not unique, the sparse monomial expansion of a junta is unique.

Paper organization

We start with various preliminaries in Section 2. We formally define sparse juntas and describe some of their properties in Section 3. We prove Theorem 1.6, the junta agreement theorem, in Section 4. We prove the main structure theorem, Theorem 1.1, in Section 5, and its monotone version, Theorem 1.2, in Section 6. We prove the junta approximation theorem, Theorems 1.3 and 1.4, in Section 7.

Acknowledgements

2   Preliminaries

We make constant use of the inequality $(a + b)^2 \leq 2a^2 + 2b^2$.

We sometimes use the notation $[n] = \{1, \ldots, n\}$.

We use 1_S for the characteristic vector of the set S or for the indicator of the event S. We also use the notation $[S]$ for the latter to improve legibility.

The *weight* of $y \in \{0, 1\}^n$ is the number of coordinates equal to 1.

A function on $\{0, 1\}^n$ is an *M-junta* if it depends on at most M coordinates.

Monotone functions   If $x, y \in \{0, 1\}^n$ then $x \leq y$ if $x_i \leq y_i$ for all $i \in [n]$. We write $x < y$ if $x \leq y$ and $x \neq y$.

A function $f \colon \{0, 1\}^n \to \mathbb{R}$ is *monotone* if $f(x) \leq f(y)$ whenever $x \leq y$.

If $f \colon \{0, 1\}^n \to \{0, 1\}$ is monotone then a *minterm* of f is an input x such that $f(x) = 1$ and $f(y) = 0$ for all $y < x$. The set of minterms of f constitutes its *minterm support*.

A monotone function $f\colon \{0,1\}^n \to \{0,1\}$ is a *width d monotone DNF* if all its minterms have weight at most d.

Biased measure   For $n \in \mathbb{N}$ and $p \in \{0,1\}$, the measure μ_p on $\{0,1\}^n$ is given by

$$\mu_p(y_1,\ldots,y_n) = p^{\sum_i y_i}(1-p)^{\sum_i (1-y_i)}.$$

Equivalently, if we identify an element in $\{0,1\}^n$ with the corresponding subset, then a subset $S \sim \mu_p$ is sampled by including each element independently with probability p.

When the domain $\{0,1\}^n$ is not clear from context, we write it explicitly, for example $\mu_p(\{0,1\}^S)$ or $\mu_p(S)$.

Given a function $f\colon \{0,1\}^n \to \mathbb{R}$, we often think of the domain as endowed with the measure μ_p for some p. In this case, we write $f\colon (\{0,1\}^n, \mu_p) \to \mathbb{R}$. If we then write $\mathbb{E}[f]$ or $\Pr[f \in \{0,1\}]$, then the underlying measure is μ_p.

Closeness   Our basic notion of closeness is L_2. Two functions f,g are ϵ-*close with respect to* μ_p if $\mathbb{E}_{\mu_p}[(f-g)^2] \leq \epsilon$. If μ_p is clear from context then we omit its mention.

Monomial expansion   Every function $f\colon \{0,1\}^n \to \mathbb{R}$ has a unique representation as a multilinear polynomial:

$$f(y_1,\ldots,y_n) = \sum_{S \subseteq [n]} \tilde{f}(S)y_S, \text{ where } y_S = \prod_{i \in S} y_i,$$

where $y_1,\ldots,y_n \in \{0,1\}$. We call this representation the *monomial expansion*. We stress that it differs from the Fourier expansion.

The functions y_S are called *monomials*, and the coefficients $\tilde{f}(S)$ are called *monomial coefficients*. The *monomial support* of f is

$$\mathrm{supp}(f) = \{S \subseteq [n] : \tilde{g}(S) \neq 0\}.$$

The *degree* of f, denoted $\deg f$, is the maximal size of a set in $\mathrm{supp}(f)$ (if $\mathrm{supp}(f) = \emptyset$ then $\deg f = 0$). This notion of degree coincides with the Fourier-theoretic notion of degree.

Low degree functions   A function $f\colon \{0,1\}^n \to \mathbb{R}$ has *degree d* if $\deg f \leq d$. A function is ϵ-*close to degree d (with respect to μ_p)* if $\mathbb{E}_{\mu_p}[(f-g)^2] \leq \epsilon$ for some degree d function g. Equivalently, $\mathbb{E}_{\mu_p}[(f-f^{\leq d})^2] = \mathbb{E}_{\mu_p}[(f^{>d})^2] \leq \epsilon$, where $f^{\leq d}, f^{>d}$ are obtained from the p-biased Fourier expansion of f by retaining all levels up to d and all levels beyond d, respectively.

Function restriction   We denote the restriction of $y \in \{0,1\}^n$ to the coordinates in a subset $S \subseteq [n]$ by $y|_S \in \{0,1\}^S$.

Given a function $f\colon \{0,1\}^n \to \mathbb{R}$ and a subset $S \subseteq [n]$, the restriction $f|_S \colon \{0,1\}^S \to \mathbb{R}$ is obtained by substituting zero for all coordinates outside of S:

$$f|_S(y) = \sum_{T \subseteq S} \tilde{f}(T)y_T.$$

Function substitution   Given a function $f\colon \{0,1\}^n \to \mathbb{R}$ and a subset $S \subseteq [n]$, the substitution $f|_{y_S \leftarrow 1} \colon \{0,1\}^{\overline{S}} \to \mathbb{R}$ is obtained by substituting $y_i = 1$ for all $i \in S$:

$$f|_{y_S \leftarrow 1} = \sum_T \tilde{f}(T)y_{T \setminus S}.$$

If $S = \{i\}$ then we use the notation $f|_{y_i \leftarrow 1}$.

Boolean functions A function is *Boolean* if it is $\{0,1\}$-valued. If f, g are Boolean then they are ϵ-close iff $\Pr[f \neq g] \leq \epsilon$. In other words, for Boolean functions the L_2 and L_0 notions of distance coincide.

If $f \colon \{0,1\}^n \to \mathbb{R}$ then $\mathrm{round}(f, \{0,1\})$ is the Boolean function obtained by rounding each $f(y)$ to the nearest value among $\{0,1\}$. The original function f is *ϵ-close to Boolean* if

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

where $\mathrm{dist}(f, \{0,1\}) = \min(|f - 0|, |f - 1|)$. Equivalently, $\mathbb{E}[(f - \mathrm{round}(f, \{0,1\}))^2] \leq \epsilon$.

2.1 Unbiased structure theorems

Nisan–Szegedy Nisan and Szegedy [NS94] proved the following theorem.

Theorem 2.1 (Nisan–Szegedy)**.** *Every Boolean degree d function on $\{0,1\}^n$ is an M_d-junta, where $M_d = d2^{d-1}$.*

The number of coordinates in the junta was improved to $O(2^d)$ in [CHS20, Wel22].

Kindler–Safra In unpublished work, Kindler and Safra [KS02, Kin03] proved the following result, which we state in its formulation due to Keller and Klein [KK20].

Theorem 2.2 (Kindler–Safra)**.** *There exists a universal constant $c > 0$ such that the following holds for all $d \in \mathbb{N}$.*
If $f \colon (\{0,1\}^n, \mu_{1/2}) \to \{0,1\}$ is ϵ-close to degree d, where $\epsilon < c^d$, then f is 2ϵ-close to a Boolean degree d function.

This formulation follows by combining Theorem 1.4 and Proposition 5.6 in [KK20]. (Keller and Klein also prove a stronger version in which the closeness is improved to the optimal expression, $\epsilon + \tilde{O}(\epsilon^2)$.)

For our purposes, we need a version of the Kindler–Safra theorem which holds for all ϵ.

Theorem 2.3. *There exists a universal constant $C > 0$ such that the following holds for all $d \in \mathbb{N}$.*
If $f \colon (\{0,1\}^n, \mu_{1/2}) \to \{0,1\}$ is ϵ-close to degree d then f is $C^d\epsilon$-close to a Boolean degree d function.

Proof. We take $C = \max(1/c, 2)$, where c is the constant in Theorem 2.2.

Suppose that $f \colon (\{0,1\}^n, \mu_{1/2}) \to \{0,1\}$ is ϵ-close to degree d. If $\epsilon < c^d$ then the result follows directly from Theorem 2.2. Otherwise, $C^d\epsilon \geq 1$, and so f is $C^d\epsilon$-close to the constant zero function. □

Monotone Kindler–Safra When the function f is monotone, we can guarantee that the approximating function in Theorem 2.3 is monotone as well.

We start with an auxiliary lemma.

Lemma 2.4. *Suppose that $f \colon (\{0,1\}^n, \mu_{1/2}) \to \mathbb{R}$ is a monotone function and $g \colon (\{0,1\}^n, \mu_{1/2}) \to \mathbb{R}$ is an M-junta.*
Either g is monotone or $\Pr[f \neq g] \geq 2^{-M}$.

Proof. Suppose without loss of generality that g depends on the first M coordinates. We separate accordingly the inputs to f and g into two parts: a point in $\{0,1\}^M$, and a point in $\{0,1\}^{n-M}$.

If g is not monotone then $g(x, 0) > g(y, 0)$ for some $x < y$. For each $z \in \{0,1\}^{n-M}$ we have $g(x, z) > g(y, z)$ while $f(x, z) \leq f(y, z)$, and so either $f(x, z) \neq g(x, z)$ or $f(y, z) \neq g(y, z)$. Since there are 2^{n-M} choices for z,

$$\Pr[f \neq g] \geq \frac{2^{n-M}}{2^n} = 2^{-M}. \qquad \square$$

We can now prove a monotone version of the Kindler–Safra theorem.

Theorem 2.5. *For every $d \in \mathbb{N}$ there is a constant $K_d > 0$ such that the following holds.*

If $f : (\{0,1\}^n, \mu_{1/2}) \to \{0,1\}$ is monotone and ϵ-close to degree d then f is $K_d\epsilon$-close to a monotone Boolean degree d function.

Proof. Suppose that $f : (\{0,1\}^n, \mu_{1/2}) \to \{0,1\}$ is monotone and ϵ-close to degree d. Applying Theorem 2.3, f is $C^d\epsilon$-close to a Boolean degree d function g.

If g is monotone then we are done. Otherwise, Theorem 2.1 shows that g is an M_d-junta, and so Lemma 2.4 shows that $C^d\epsilon \geq 2^{-M_d}$. Consequently $\epsilon \geq 2^{-M_d}/C^d$, and so f is $C^d 2^{M_d}\epsilon$-close to the zero function. $\qquad \square$

Curiously, the proof of the biased version of Theorem 2.5 only uses Lemma 2.4.

2.2 Forbidden configurations

Our work involves three different types of *forbidden configurations*. These are configurations of coefficients which cannot occur in Boolean functions, but do arise in our context — but only sparingly. For example, the biased FKN theorem [Fil16] involves functions of the form $f = \sum_i c_i y_i$. If f is Boolean then we cannot have $c_i = c_j = 1$, but if f is ϵ-close to Boolean then there could be up to $O(\epsilon/p^2)$ copies of this configuration.

Minimal non-Boolean inputs   Let $f : \{0,1\}^n \to \mathbb{R}$. An input x is a *minimal non-Boolean input* of f if $f(x) \notin \{0,1\}$ but $f(y) \in \{0,1\}$ for all $y < x$.

Suppose that f has degree 1, say

$$f = c_0 + \sum_{i=1}^n c_i y_i.$$

If $c_0 \notin \{0,1\}$ then 0 is the only minimal non-Boolean input. Otherwise, suppose without loss of generality that $c_0 = 0$. If $c_i \notin \{0,1\}$ then $1_{\{i\}}$ is a minimal non-Boolean input. If $c_i = c_j = 1$ then $1_{\{i,j\}}$ is a minimal non-Boolean input. There are no other types of minimal non-Boolean inputs.

In addition, if x is a minimal non-Boolean input then this is due to some non-zero monomial coefficients: if $1_{\{i\}}$ is a minimal non-Boolean input then $c_i \neq 0$, and if $1_{\{i,j\}}$ is a minimal non-Boolean input of f then $c_i, c_j \neq 0$. In both cases, the minimal non-Boolean input $x = 1_S$ is such that S is the union of sets in $\operatorname{supp}(f)$: $\{i\}$ in the former case, and $\{i\}, \{j\}$ in the latter case.

The following lemma shows that a similar picture holds for all d.

Lemma 2.6. *For every d there is a constant L_d such that every minimal non-Boolean input of a degree d function $f : \{0,1\}^n \to \mathbb{R}$ has weight at most L_d.*

Furthermore, if 1_S is a minimal non-Boolean input of f then S can be written as the union of at most L_d sets in $\operatorname{supp}(f)$.

Proof. According to Theorem 2.1, every Boolean degree d function is an M_d-junta. We will show that we can take $L_d = M_d + d$.

Suppose that 1_S is a minimal non-Boolean input of f. By minimality, $f|_S$ depends on all coordinates, and so S is a union of all sets in $\operatorname{supp}(f|_S)$. We can clearly write it as a union of at most $|S|$ such sets, say $B_1, \ldots, B_{|S|}$.

If $|S| \leq M_d$ then we are done. Otherwise, let m be the minimal index such that $|B_1 \cup \cdots \cup B_m| > M_d$. By minimality, $f|_{B_1 \cup \cdots \cup B_{m-1}}$ is Boolean, and so $|B_1 \cup \cdots \cup B_{m-1}| \leq M_d$ according to Theorem 2.1. Since $f|_{B_1 \cup \cdots \cup B_m}$ depends on all inputs and $|B_1 \cup \cdots \cup B_m| > M_d$, Theorem 2.1 shows that it is not Boolean, and so $B_1 \cup \cdots \cup B_m = S$. Therefore

$$|S| \leq |B_1 \cup \cdots \cup B_{m-1}| + |B_m| \leq M_d + d. \qquad \square$$

Looking ahead, we will be interested in minimal non-Boolean inputs of degree d functions whose weight is larger than d. If 1_S is a minimal non-Boolean input of f then $f|_S$ is a function which is Boolean except

9

for the input 1_S. Listing such *minimal non-Boolean functions* is a convenient way to describe all minimal non-Boolean inputs. For example, when $d = 1$, the minimal non-Boolean functions are

$$y_1 + y_2, 1 - y_1 - y_2.$$

In general, the minimal non-Boolean inputs come in pairs $f, 1 - f$, and it suffices to list those which satisfy $f(0) = 0$.

Here are the minimal non-Boolean functions satisfying $f(0) = 0$ for $d = 2$:

$$y_1(y_2 + y_3)$$
$$y_1 y_2 + y_1 y_3 + y_2 y_3$$
$$y_1(1 - y_2 - y_3)$$
$$y_1 + y_2 y_3$$
$$(y_1 + y_2)(1 - y_3) - 2 y_1 y_2$$
$$y_1(1 - y_3) + y_2 - 2 y_1 y_2$$
$$(y_1 + y_2)(1 - y_3) - y_1 y_2$$
$$y_1 + y_2 + y_3 - 2 y_1 y_2 - 2 y_1 y_3 - 2 y_2 y_3$$
$$y_1 + y_2 + y_3 - 2 y_1 y_2 - 2 y_1 y_3 - y_2 y_3$$
$$y_1 + y_2 + y_3 - 2 y_1 y_2 - y_1 y_3 - y_2 y_3$$
$$y_1 y_2 + y_3 y_4$$
$$y_1(1 - y_3) + y_2(1 - y_4) - y_1 y_2$$
$$y_1 + y_2 + y_3 - y_1 y_2 - y_1 y_3 - y_2 y_3 - y_1 y_4$$
$$y_1 + y_2 + y_3 + y_4 - y_1 y_2 - y_1 y_3 - y_1 y_4 - y_2 y_3 - y_2 y_4 - y_3 y_4$$

Minimal non-monotone inputs Let $f \colon \{0,1\}^n \to \{0,1\}$. An input x is *minimal non-monotone* if $f(x) < f(y)$ for some $y < x$, and $f(y) \le f(z)$ whenever $y \le z < x$.

Theorem 2.1 implies that each such input has weight at most M_d.

Lemma 2.7. *Every minimal non-monotone input of a degree d function $f \colon \{0,1\}^n \to \{0,1\}$ has weight at most M_d, where M_d is the constant from Theorem 2.1.*

Furthermore, if 1_S is a minimal non-monotone input of f then S can be written as the union of at most M_d sets in $\mathrm{supp}(f)$.

Proof. By minimality, if 1_S is a minimal non-monotone input of f then $f|_S$ depends on all inputs, and so $|S| \le M_d$ by Theorem 2.1. Moreover, S is the union of all sets in $\mathrm{supp}(f|_S)$, and so it is the union of at most $|S|$ sets in $\mathrm{supp}(f)$. $\qquad\square$

We extend the definition on minimal non-monotone inputs to non-Boolean functions as follows. If $f \colon \{0,1\}^n \to \mathbb{R}$ then an input $x = 1_S$ is *minimal non-monotone Boolean* if $f|_S$ is Boolean and x is a minimal non-monotone input of $f|_S$.

Lemma 2.7 holds also for minimal non-monotone Boolean inputs of a degree d function $g \colon \{0,1\}^n \to \mathbb{R}$, since if 1_S is a minimal non-monotone Boolean input of g then it is also a minimal non-monotone input of the Boolean function $g|_S$.

As in the case of minimal non-Boolean inputs, we will be interested in minimal non-monotone inputs of degree d functions whose weight is larger than d, which we can describe using *minimal non-monotone functions*. When $d = 1$, there are no minimal non-monotone functions. When $d = 2$, the unique minimal non-monotone function is

$$y_1 + y_2 + y_3 - y_1 y_2 - y_1 y_3 - y_2 y_3.$$

Minimal high-degree inputs Let $f\colon \{0,1\}^n \to \{0,1\}$, and let d be a parameter. An input 1_S is *minimal high-degree* if $\deg(f|_S) > d$ but $\deg(f|_T) \leq d$ for all $T \subsetneq S$.

Lemma 2.8. *Every minimal high-degree input of a monotone width d DNF $f\colon \{0,1\}^n \to \{0,1\}$ has weight at most L_d, where L_d is the constant from Lemma 2.6.*

Furthermore, if 1_S is a minimal high-degree input of f then S can be written as the union of at most L_d minterms of f.

Proof. By minimality, if 1_S is a minimal high-degree input of f then $f|_S$ depends on all coordinates, and so S is the union of at most $|S|$ minterms in $\operatorname{supp}(f|_S)$, say $B_1, \ldots, B_m$.

If $|S| \leq M_d$ then we are done. Otherwise, let m be the minimal index such that $|B_1 \cup \cdots \cup B_m| > M_d$. By minimality, $\deg(f|_{B_1 \cup \cdots \cup B_{m-1}}) \leq d$, and so $|B_1 \cup \cdots \cup B_{m-1}| \leq M_d$ by Theorem 2.1. Since $f|_{B_1 \cup \cdots \cup B_m}$ depends on all inputs and $|B_1 \cup \cdots \cup B_m| > M_d$, Theorem 2.1 shows that it cannot have degree d, and so $B_1 \cup \cdots \cup B_m = S$. Therefore $|S| \leq M_d + d$. $\qquad\square$

As in the case of minimal non-Boolean inputs and minimal non-monotone inputs, we can describe minimal high-degree inputs as minimal high-degree monotone DNFs. When $d = 1$, the unique minimal high-degree monotone DNF is $y_1 \vee y_2$, and when $d = 2$, the minimal high-degree monotone DNFs are

$$y_1 \wedge (y_2 \vee y_3)$$
$$(y_1 \wedge y_2) \vee (y_1 \wedge y_3) \vee (y_2 \wedge y_3)$$
$$y_1 \vee (y_2 \wedge y_3)$$
$$y_1 \vee y_2 \vee y_3$$
$$(y_1 \wedge y_2) \vee (y_3 \wedge y_4)$$

3 Sparsity

The approximating function g in Theorem 1.1 has sparse monomial support, and its monotone counterpart in Theorem 1.2 has sparse minterm support. This crucial property drives much of the proofs of these theorems.

Here is the notion of sparsity that comes up in both theorems.

Definition 3.1 (Sparse). A set system over a set V is a collection of subsets of V.

A set system $\mathcal{F}$ is (d, C, ϵ)-*sparse*, where $d \geq 0$ is an integer, $\epsilon \geq 0$, and $C \geq 1$, if the following three properties hold:

(i) All sets in $\mathcal{F}$ have size at most d.

(ii) For all $T \subseteq [n]$ and integer $e \geq 0$, the set system $\mathcal{F}$ contains at most C^e sets $U \supseteq T$ of size $|U| + e$.

(iii) For all integer $e \geq 0$, the set system $\mathcal{F}$ contains at most $C^e \epsilon$ sets U of size e.

A set system is (d, C)-sparse if it is $(d, C, 1)$-sparse; this corresponds to omitting the final property.

Definition 3.2 (Sparse junta). A function $f\colon \{0,1\}^n \to \mathbb{R}$ is a (d, C)-*sparse junta* if $\operatorname{supp}(f)$ is (d, C)-sparse.

A monotone function $f\colon \{0,1\}^n \to \{0,1\}$ is a (d, C)-*sparse DNF* if its minterm support is (d, C)-sparse.

We describe some elementary properties of sparse set systems in Section 3.1. Sparse set systems satisfy a crucial property we call the *reverse union bound*, which we explain in Section 3.2. Finally, we discuss *quantized* sparse juntas, which are sparse juntas in which the monomial coefficients belong to a fixed finite set. We show in Section 3.3 that if a quantized sparse junta f is $L - 0$-close to Boolean (that is, $\Pr[f \notin \{0,1\}] \leq \epsilon$) then it is also L_2-close to Boolean (that is, $\mathbb{E}[\operatorname{dist}(f, \{0,1\})^2] = O(\epsilon)$).

3.1 Elementary properties

We discuss the effect of three natural operations on the sparsity of set systems: union, sum and substitution.

The union of two set systems is their set-theoretic union, and it corresponds to the sum of functions:

$$\text{supp}(f + g) \subseteq \text{supp}(f) \cup \text{supp}(g),$$

where the inclusion is tight generically.

Definition 3.3 (Sum of set systems). If $\mathcal{F}, \mathcal{G}$ are set systems then their *sum* is

$$\mathcal{F} \sqcup \mathcal{G} = \{A \cup B : A \in \mathcal{F}, B \in \mathcal{G}\}.$$

The $k'th$ *iterated sum* of a set system $\mathcal{F}$ is the sum of k copies of $\mathcal{F}$:

$$\mathcal{F}^{\sqcup k} = \{A_1 \cup \cdots \cup A_k : A_1, \ldots, A_k \in \mathcal{F}\}.$$

Since $A_1, \ldots, A_k$ are not necessarily distinct, $\mathcal{F}^{\sqcup k}$ includes $\mathcal{F}^{\sqcup \ell}$ for all non-zero $\ell < k$.

The sum of two set systems corresponds to the product of functions:

$$\text{supp}(fg) \subseteq \text{supp}(f) \sqcup \text{supp}(g),$$

where the inclusion is tight generically.

Definition 3.4 (Contraction). If $\mathcal{F}$ is a set system over V and $S \subseteq V$ then

$$\mathcal{F}|_{S \leftarrow 1} = \{T \setminus S : T \in \mathcal{F}\}.$$

Contraction corresponds to substituting the value of variables:

$$\text{supp}(f|_{y_S \leftarrow 1}) \subseteq \text{supp}(f)|_{S \leftarrow 1},$$

and the inclusion is tight generically.

Union, sum and contraction all preserve sparsity, after adapting the parameters.

Lemma 3.5. *Let $\mathcal{F}, \mathcal{G}$ be (d, C)-sparse set systems.*

(a) $\mathcal{F} \cup \mathcal{G}$ is $(d, 2C)$-sparse.

(b) $\mathcal{F} \sqcup \mathcal{G}$ is $(2d, O_d(C))$-sparse.

(c) $\mathcal{F}^{\sqcup k}$ are $(kd, O_{k,d}(C))$-sparse.

(d) $\mathcal{F}|_{J \leftarrow 1}$ is $(d, O_{d,|J|}(C))$-sparse.

Proof. Every set system satisfies the definition of (d, C)-sparse for $e = 0$, and so it suffices to prove it for $e \geq 1$.

Item a. Let $\mathcal{F}, \mathcal{G}$ be (d, C)-sparse set systems. Given R and $e \geq 1$, there are at most C^e sets in each of $\mathcal{F}, \mathcal{G}$ of size $|R| + e$ which contain R. Therefore there are at most $2C^e \leq (2C)^e$ sets in $\mathcal{F} \cup \mathcal{G}$ of size $|R| + e$ which contain R.

Item b. Let $\mathcal{F}, \mathcal{G}$ be (d, C)-sparse set systems. Given R and e, we need to bound the number of sets $S = A \cup B$ such that $A \in \mathcal{F}$, $B \in \mathcal{G}$, $A \cup B \supseteq R$, and $|A \cup B| = |R| + e$. We can assume that $|R| \leq d$.

We will show that for every $a, a_r, a_b \leq |R|$ there are at most $4^d C^e$ pairs (A, B) such that:

(i) $A \in \mathcal{F}$, $B \in \mathcal{G}$, $A \cup B \supseteq R$, $|A \cup B| = |R| + e$.

(ii) $|A| = a$, $|A \cap R| = a_r$, and $|A \cap B| = a_b$.

Since there are at most $(d+1)^3$ many options for a, a_r, a_b, this will complete the proof.

There are $\binom{|R|}{a_r} \leq 2^d$ choices for $A \cap R$ such that $|A \cap R| = a_r$. Since $\mathcal{F}$ is (d, C)-sparse, there are at most C^{a-a_r} sets $A \in \mathcal{F}$ extending $A \cap R$ of size a. There are $\binom{|A|}{a_b} \leq 2^a \leq 2^d$ choices for $A \cap B$.

Notice that B needs to contain both $A \cap B$ and $R \setminus A$, two sets which are disjoint. Also, since $|A \cup B| = |R| + e$, we have $|B| = |A \cup B| + |A \cap B| - |A| = |R| + e + a_b - a$. Since $\mathcal{G}$ is (d, C)-sparse, the number of sets $B \in \mathcal{G}$ of size b extending $(A \cap B) \cup (R \setminus A)$ is at most $C^{(|R|+e+a_b-a)-a_b-(|R|-a_r)} = C^{e-a+a_r}$.

Altogether, given a, a_r, a_b, the number of pairs A, B satisfying the properties listed above is bounded by

$$\underbrace{2^d}_{A \cap R} \cdot \underbrace{C^{a-a_r}}_{A \text{ given } A \cap R} \cdot \underbrace{2^d}_{A \cap B \text{ given } A} \cdot \underbrace{C^{e-a+a_r}}_{B \text{ given } A \cap R, A \cap B} = 4^d C^e,$$

as required.

Item c. This item follows from Items a and b by induction on k.

Item d. Let $\mathcal{F}$ be a (d, C)-sparse set system, and let J be a set. Given R and e, we need to bound the number of sets $T \in \mathcal{F}$ such that $T \setminus J \supseteq R$ and $|T \setminus J| = |R| + e$. We can assume that R is disjoint from J.

We will bound the number of such sets T given $I = T \cap J$. Any such set T contains both I and R, and $|T| = |T \setminus J| + |T \cap J| = |R| + e + |I|$. Since $\mathcal{F}$ is (d, C)-sparse, there are at most $C^{(|R|+e+|I|)-(|I|+|R|)} = C^e$ sets $T \in \mathcal{F}$ containing $I \cup R$ and of size $|R| + e + |I|$. We conclude that the number of sets $T \in \mathcal{F}$ such that $T \setminus J \supseteq R$ and $|T \setminus J| = |R| + e$ is at most

$$\sum_{I \subseteq J} C^e \leq 2^{|J|} C^e \leq 2^d C^e \leq (2^d C)^e. \qquad \square$$

Lemma 3.6. *Let $\mathcal{F}, \mathcal{G}$ be (d, C, ϵ)-sparse.*

(a) $\mathcal{F} \cup \mathcal{G}$ is $(d, 2C, \epsilon)$-sparse.

(b) $\mathcal{F} \sqcup \mathcal{G}$ is $(2d, O_d(C), \epsilon)$-sparse. For this it suffices that $\mathcal{G}$ be (d, C)-sparse.

(c) $\mathcal{F}^{\sqcup k}$ are $(kd, O_{k,d}(C), \epsilon)$-sparse.

Proof.

Item a. Let $\mathcal{F}, \mathcal{G}$ be (d, C, ϵ)-sparse set systems. Lemma 3.5 shows that $\mathcal{F} \cup \mathcal{G}$ is $(d, 2C)$-sparse. Given $e \geq 1$, there are at most $C^e \epsilon$ sets of size e in each of $\mathcal{F}, \mathcal{G}$. Therefore there are at most $2C^e \epsilon \leq (2C)^e \epsilon$ sets of size e in $\mathcal{F} \cup \mathcal{G}$.

The case $e = 0$ requires special treatment. If $\epsilon < 1$ then $\mathcal{F}, \mathcal{G}$ do not contain $\emptyset$, and so $\mathcal{F} \cup \mathcal{G}$ also doesn't contain $\emptyset$. If $\epsilon \geq 1$ then clearly $\mathcal{F} \cup \mathcal{G}$ contains at most ϵ sets of size 0.

Item b. Let $\mathcal{F}$ be a (d, C, ϵ)-sparse set system and let $\mathcal{G}$ be a (d, C)-sparse set system. Lemma 3.5 shows that $\mathcal{F} \sqcup \mathcal{G}$ is $(d, O_d(C))$-sparse.

Given e, we need to bound the number of sets $S = A \cup B$ of size e, where $A \in \mathcal{F}$ and $B \in \mathcal{G}$. For this, we show that for each $a, a_b \leq d$, there are at most $2^d C^e \epsilon$ sets $A \cup B$ of size e such that $A \in \mathcal{F}$, $B \in \mathcal{G}$, $|A| = a$, and $|A \cap B| = a_b$.

Indeed, there are at most $C^a \epsilon$ sets $A \in \mathcal{F}$ of size a. Given A, there are at most $\binom{|A|}{a_b} \leq 2^d$ choices for $A \cap B$. Since $|A \cup B| = e$, we have $|B| = |A \cup B| + |A \cap B| - |A| = e + a_b - a$. Therefore given $A \cap B$, there are at most C^{e-a} sets $B \in \mathcal{G}$ including $A \cap B$ such that $|A \cup B| = e$. Altogether, given a, a_b, the number of pairs A, B satisfying the properties listed above is bounded by

$$\underbrace{C^a \epsilon}_{A} \cdot \underbrace{2^d}_{A \cap B \text{ given } A} \cdot \underbrace{C^{e-a}}_{B \text{ given } A \cap B} = 2^d C^e \epsilon,$$

as claimed.

Since there are at most $(d+1)^2$ many options for a, a_b, this shows that the number of sets $S \in \mathcal{F} \sqcup \mathcal{G}$ of size $e \geq 1$ is at most $(d+1)^2 2^d C^e \epsilon \leq ((d+1)^2 2^d C)^e \epsilon$, completing the proof except for the case $e = 0$.

In order to handle the case $e = 0$, we again consider whether $\epsilon < 1$ or not. If $\epsilon < 1$ then $\emptyset \notin \mathcal{F}$ and so $\emptyset \notin \mathcal{F} \sqcup \mathcal{G}$. If $\epsilon \geq 1$ then $\mathcal{F} \sqcup \mathcal{G}$ trivially contains at most ϵ sets of size 0.

Item c. This item follows from Items a and b by induction on k. $\qquad\square$

3.2 Reverse union bound

Let $\mathcal{F}$ be a set system. The union bound shows that

$$\Pr_{Y \sim \mu_p}[Y \supseteq S \text{ for some } S \in \mathcal{F}] \leq \sum_{S \in \mathcal{F}} p^{|S|}.$$

In general, the union bound is not tight. For example, suppose that $\mathcal{F} = \{\{1\}, \ldots, \{n\}\}$. As $n \to \infty$, the left-hand side tends to 1 while the right-hand side tends to infinity. More generally, if $\mathcal{F} = \{[m] \cup \{m+1\}, \ldots, [m] \cup \{n\}\}$ then the left-hand side tends to p^m while the right-hand side tends to infinity.

In this section, we show that the union bound is tight up to a constant factor provided that $\mathcal{F}$ is sparse.

Lemma 3.7 (Reverse union bound). *Let $p \leq 1/2$, let $\mathcal{G}$ be a $(d, C/p)$-sparse set system, and let $\mathcal{F} \subseteq \mathcal{G}$.*

For $S \in \mathcal{F}$, let $\mathcal{E}_S$ denote the event that $y_S = 1$ and $y_R = 0$ for all $R \in \mathcal{G}$ such that $R \not\subseteq S$.

Let $\mathcal{E}$ denote the union of the events $\mathcal{E}_S$ for all $S \in \mathcal{F}$. Then

$$\Pr[\mathcal{E}] = \Theta_{d,C}\left(\sum_{S \in \mathcal{F}} p^{|S|}\right).$$

Proof. Let $\mathcal{M}$ consist of the inclusion-minimal sets in $\mathcal{F}$. Since $\mathcal{M}$ is an antichain, the events $\mathcal{E}_S$ for $S \in \mathcal{M}$ are disjoint, and so

$$\Pr[\mathcal{E}] = \Pr\left[\bigcup_{S \in \mathcal{M}} \mathcal{E}_S\right] = \sum_{S \in \mathcal{M}} \Pr[\mathcal{E}_S] \leq \sum_{S \in \mathcal{M}} p^{|S|}.$$

For every $S \in \mathcal{M}$,

$$\Pr[\mathcal{E}_S] = p^{|S|} \Pr[y_T = 0 \text{ for all } \emptyset \neq T \in \mathcal{F}|_{S \leftarrow 1}] \overset{(*)}{\geq} \prod_{\emptyset \neq T \in \mathcal{F}|_{S \leftarrow 1}} \Pr[y_T = 0],$$

where $(*)$ follows from the Harris–FKG inequality since the events $y_T = 0$ are anti-monotone. Lemma 3.5 shows that $\mathcal{F}|_{S \leftarrow 1}$ is $(d, K/p)$-sparse, where K depends only on d, C. Therefore

$$\Pr[\mathcal{E}_S] \geq p^{|S|} \prod_{e=1}^{d} (1 - p^e)^{K^e/p^e} \geq p^{|S|} \prod_{e=1}^{d} (1 - 2^{-e})^{(2K)^e} = \Omega_{d,C}(p^{|S|}),$$

14

since $(1-x)^{1/x}$ is decreasing in $[0,1]$. Therefore

$$\Pr[\mathcal{E}] \geq \Omega_{d,C}\left(\sum_{S\in\mathcal{M}} p^{|S|}\right).$$

So far, we have shown that

$$\Pr[\mathcal{E}] = \Theta_{d,C}\left(\sum_{S\in\mathcal{M}} p^{|S|}\right).$$

In order to complete the proof, we will show that

$$\sum_{S\in\mathcal{F}} p^{|S|} = \Theta_{d,C}\left(\sum_{S\in\mathcal{M}} p^{|S|}\right).$$

Since $\mathcal{M} \subseteq \mathcal{F}$, clearly

$$\sum_{S\in\mathcal{F}} p^{|S|} \geq \sum_{S\in\mathcal{M}} p^{|S|}.$$

In the other direction, if $T \in \mathcal{F}$ then $T \supseteq S$ for some $S \in \mathcal{M}$, and so

$$\sum_{T\in\mathcal{F}} p^{|T|} \leq \sum_{S\in\mathcal{M}} \sum_{e=|S|}^{d} \sum_{\substack{T\in\mathcal{F}\\ T\supseteq S, |T|=e}} p^e \leq \sum_{S\in\mathcal{M}} \sum_{e=|S|}^{d} (C/p)^{e-|S|} p^e \leq \sum_{S\in\mathcal{M}} (d+1)C^d p^{|S|} = \Theta_{d,C}\left(\sum_{S\in\mathcal{M}} p^{|S|}\right). \quad \square$$

When $\mathcal{F} = \mathcal{G}$, we recover the statement that the union bound is tight up to a constant factor for sparse set systems. Indeed, if $\mathcal{E}$ happens then in particular $Y \supseteq S$ for some $S \in \mathcal{F}$, and so the Lemma shows that

$$\Pr_{Y\sim\mu_p}[Y \supseteq S \text{ for some } S \in \mathcal{F}] \geq \Pr[\mathcal{E}] = \Omega_{d,C}\left(\sum_{S\in\mathcal{F}} p^{|S|}\right).$$

3.3   Quantized sparse juntas

The sparse juntas that we consider later on will arise from applying the junta agreement theorem, Theorem 1.6. Such functions are *quantized* in the following sense.

Definition 3.8 (Quantized functions)**.** A function $f\colon \{0,1\}^n \to \mathbb{R}$ is B-*quantized*, where B is a finite set, if all monomial coefficients of f belong to B.

Quantization is preserved under basic operations.

Lemma 3.9. *Let $f,g\colon \{0,1\}^n \to \mathbb{R}$ be B-quantized functions of degree d, where B is a finite set.*

(a) For any $c \in \mathbb{R}$, the function cf is cB-quantized, where $cB = \{cb : b \in B\}$.

(b) The function $f + g$ is $(B+B)$-quantized, where $B + B = \{b_1 + b_2 : b_1, b_2 \in B\}$.

(c) The function fg is C-quantized, where C is a finite set depending only on d, B.

Proof. Items a and b are clear. For Item c, we use the formula

$$\widetilde{fg}(U) = \sum_{S\cup T=U} \tilde{f}(S)\tilde{g}(T),$$

which shows that $\widetilde{fg}(U)$ is the sum of $3^{|U|} \leq 3^{2d}$ products of two elements from B. $\quad \square$

15

The junta agreement theorem will result in a quantized sparse junta g which is close to Boolean in an L_0 sense: $\Pr[g \notin \{0,1\}] = O(\epsilon)$. In this section, we show how to conclude closeness in an L_2 sense: $\Pr[\text{dist}(g, \{0,1\})^2] = O(\epsilon)$.

We start with a similar result which compares the L_0 and L_2 norms.

Lemma 3.10. *If $f \colon (\{0,1\}^n, \mu_p) \to \mathbb{R}$ is a B-quantized $(d, C/p)$-sparse junta, where $p \leq 1/2$, then*

$$\Pr[f \neq 0] = \Theta_{d,C,B}(\mathbb{E}[f^2]).$$

Proof. If $f \neq 0$ then clearly $y_S = 1$ for some $S \in \text{supp}(f)$, and so the union bound shows that

$$\Pr[f \neq 0] \leq \sum_{S \in \text{supp}(f)} p^{|S|}.$$

We can get a matching lower bound using Lemma 3.7. Let $\mathcal{F}$ be all inclusion-minimal sets in $\text{supp}(f)$, and let $\mathcal{G} = \text{supp}(f)$. If the event $\mathcal{E}_S$ in the Lemma happens for some $S \in \mathcal{F}$ then $f = \tilde{f}(S) \neq 0$, and so

$$\Pr[f \neq 0] \geq \Pr[\mathcal{E}] = \Omega_{d,C}\left(\sum_{S \in \mathcal{F}} p^{|S|}\right).$$

The calculation at the very end of the proof of Lemma 3.7 shows that

$$\sum_{S \in \text{supp}(f)} p^{|S|} = O_{d,C}\left(\sum_{S \in \mathcal{F}} p^{|S|}\right),$$

and we conclude that

$$\Pr[f \neq 0] \geq \Pr[\mathcal{E}] = \Omega_{d,C}\left(\sum_{S \in \text{supp}(f)} p^{|S|}\right). \tag{2}$$

Similarly,

$$\mathbb{E}[f^2] \geq \min_{\substack{b \in B \\ b \neq 0}} b^2 \Pr[\mathcal{E}] = \Omega_{d,C,B}\left(\sum_{S \in \text{supp}(f)} p^{|S|}\right) = \Omega_{d,C,B}(\Pr[f \neq 0]).$$

Equation (2) implies that $\text{supp}(f)$ is $(d, C/p, \epsilon)$-sparse, where $\epsilon = O_{d,C}(\Pr[f \neq 0])$. Lemma 3.6 shows that $\text{supp}(f^2) \subseteq \text{supp}(f)^{\sqcup 2}$ is $(2d, K/p, \epsilon)$-sparse, where $K = O_{d,C}(1)$, and Lemma 3.9 shows that f^2 is D-quantized for some finite B' depending only on d, B. Therefore

$$\mathbb{E}[f^2] \leq \max_{b \in B'} b^2 \sum_{e=0}^{2d} (K/p)^e \epsilon = O_{d,C,B}(\Pr[f \neq 0]). \qquad \square$$

Using this, we derive the main result of this section.

Lemma 3.11. *If $f \colon (\{0,1\}^n, \mu_p) \to \mathbb{R}$ is a B-quantized $(d, C/p)$-sparse junta, where $p \leq 1/2$, then*

$$\Pr[f \notin \{0,1\}] = \Theta_{d,C,B}(\mathbb{E}[\text{dist}(f, \{0,1\})^2]).$$

Proof. Let $g = f(f-1)$. Since $\text{supp}(g) \subseteq \text{supp}(f)^{\sqcup 2}$, Lemmas 3.5 and 3.9 show that g is a B'-quantized $(d, K/p)$-sparse junta, where B', K depend on d, C, B. Therefore Lemma 3.10 implies that

$$\Pr[f \notin \{0,1\}] = \Pr[g \neq 0] = \Theta_{d,C,B}(\mathbb{E}[g^2]).$$

We show that $\mathbb{E}[\text{dist}(f, \{0,1\})^2] = O_{d,C,B}(\Pr[f \notin \{0,1\}])$ by relating $\mathbb{E}[\text{dist}(f, \{0,1\})^2]$ to $\mathbb{E}[g^2]$. Consider any input y. If $\text{round}(f, \{0,1\}) = a$ then $|y - (1-a)| \geq 1/2$ and so $(f(y) - (1-a))^2 \geq 1/4$. Consequently,

$$\text{dist}(f(y), \{0,1\})^2 = (f(y) - a)^2 \leq 4(f(y) - a)^2 (f(y) - (1-a))^2 = 4g(y)^2.$$

It follows that
$$\mathbb{E}[\mathrm{dist}(f,\{0,1\})^2] \le 4\,\mathbb{E}[g^2] = O_{d,C,B}(\Pr[f \notin \{0,1\}]).$$

The proof of the other direction, $\Pr[f \notin \{0,1\}] = O_{d,C,B}(\mathbb{E}[\mathrm{dist}(f,\{0,1\})^2])$, uses the notion of minimal non-Boolean input, defined in Section 2.2. If $f(y) \notin \{0,1\}$ then $y \ge 1_S$ for some minimal non-Boolean input 1_S of f, implying that $y_S = 1$. Denoting the set of minimal non-Boolean inputs of f by $\mathcal{F}$, this shows that
$$\Pr[f \notin \{0,1\}] \le \sum_{1_S \in \mathcal{F}} \Pr[y_S = 1] = \sum_{1_S \in \mathcal{F}} p^{|S|}.$$

We would like to bound the right-hand side using Lemma 3.7, but it is not necessarily the case that $\mathcal{F} \subseteq \mathrm{supp}(f)$. However, Lemma 2.6 shows that $\mathcal{F} \subseteq \mathrm{supp}(f)^{\sqcup L_d}$. Lemma 3.9 shows that $\mathrm{supp}(f)^{\sqcup L_d}$ is $(d, K'/p)$-sparse, where K' depends on d, C. Therefore Lemma 3.7, applied with $\mathcal{F}$ and $\mathcal{G} = \mathrm{supp}(f)^{\sqcup L_d}$, shows that
$$\Pr[f \notin \{0,1\}] \le \sum_{1_S \in \mathcal{F}} p^{|S|} = \Theta_{d,C}(\Pr[\mathcal{E}]),$$

where $\mathcal{E}$ is the event that $y_S = 1$ for some $S \in \mathcal{F}$, and $y_R = 0$ for all $R \in \mathrm{supp}(f)^{\sqcup L_d}$ such that $R \subseteq S$. If the event $\mathcal{E}$ happens then $f(y) = f(1_S) \notin \{0,1\}$. Moreover, since $|S| \le L_d$ by Lemma 2.6 and f is B-quantized, $\mathrm{dist}(f(y), \{0,1\})^2 = \Omega_{d,B}(1)$. Therefore
$$\Pr[f \notin \{0,1\}] = O_{d,C,B}(\mathbb{E}[\mathrm{dist}(f,\{0,1\})^2]). \qquad \square$$

4 Junta agreement theorem

We prove our main structure theorems, Theorems 1.1 and 1.2, by reduction to the Kindler–Safra theorem on $(\{0,1\}^n, \mu_{1/2})$. As we explain in Section 1.2, this involves applying the Kindler–Safra theorem to $f|_S$ for $S \sim \mu_{2p}$, obtaining approximating juntas g_S, and pasting the juntas together to a global function g.

In order to show that the juntas g_S can be pasted together, we need to assume that they agree with each other, in the sense that
$$\Pr_{(S_1,S_2,T)\sim\nu}[g_{S_1}|_T \ne g_{S_2}|_T] \le \epsilon$$

for an appropriate distribution ν supported on triplets (S_1, S_2, T) such that $S_1, S_2 \supseteq T$.

Since we are interested in the behavior of g_S for $S \sim \mu_{2p}$, we need the marginals of S_1 and S_2 to be μ_{2p}. Moreover, we want S_1 and S_2 to be independent given T. This naturally leads to the product distribution $\mu_{q,r}$ (where $q = 2p$ and $r \le q$), which is defined as follows:

1. T includes each $i \in [n]$ with probability r.

2. For $\ell \in \{1,2\}$, S_ℓ includes each $i \in T$ with probability 1, and each $i \notin T$ with probability $\frac{q-r}{1-r}$.

It is easy to check that the marginal distributions of S_1, S_2, T are μ_q, μ_q, μ_r, respectively.

We can now state the junta agreement theorem, which is a formal statement of Theorem 1.6.

Theorem 4.1 (Junta agreement theorem). *Fix an integer K and $c \in (0,1)$. The following holds for all $q \in (0,1)$.*

Suppose that for each $S \subseteq [n]$ we are given a degree d function $g_S \colon \{0,1\}^n \to \{0,1\}$ whose monomial support contains at most K monomials (the junta assumption*). Assume that the functions g_S satisfy the following* agreement condition*:*
$$\Pr_{(S_1,S_2,T)\sim\mu_{q,cq}}[g_{S_1}|_T \ne g_{S_2}|_T] \le \epsilon.$$

Define a degree d function $g \colon \{0,1\}^n \to \{0,1\}$ as follows: for each $|A| \le d$, let $\tilde{g}(A)$ be any value a maximizing
$$\Pr_{\substack{S \sim \mu_q \\ S \supseteq A}}[\tilde{g}_S(A) = a].$$

17

The function g is a $(d, O_{K,c}(1/q))$-sparse junta satisfying

$$\Pr_{S \sim \mu_q}[g|_S \neq g_S] = O_{K,c}(\epsilon).$$

In the statement of the Theorem and below, we use the convention that in $\Pr$ and $\mathbb{E}$, we take the distribution on the first line of the subscript conditioned on the constraints in the following lines. For example, in the second display, S is chosen by including all elements in A with probability 1 and all elements outside of A with probability q.

Our previous work [DFH17] proves a stronger version of Theorem 4.1, which does not require the g_S to be juntas but has the same conclusion (without the guarantee that g is a sparse junta, which follows from the g_S being juntas). The proof of this stronger version is somewhat non-intuitive, using the notion of good and excellent sets which is borrowed from [IKW12]. It turns out that the proof simplifies dramatically when the g_S are juntas, and this is the proof that we present here.

The proof of Theorem 4.1 relies on the following expansion lemma.

Lemma 4.2. *Assume the setup of Theorem 4.1. There exists a constant T_c, depending only on c, such that for every $|A| \leq d$,*

$$\Pr_{\substack{S_1, S_2 \sim \mu_q \\ S_1, S_2 \supseteq A}}[\tilde{g}_{S_1}(A) \neq \tilde{g}_{S_2}(A)] \leq T_c \Pr_{\substack{(S_1, S_2, T) \sim \mu_{q,cq} \\ T \supseteq A}}[\tilde{g}_{S_1}(A) \neq \tilde{g}_{S_2}(A)].$$

We first derive Theorem 4.1 assuming Lemma 4.2, and then prove Lemma 4.2 using elementary Fourier analysis.

Proof of Theorem 4.1.

Agreement   Our first goal is to show that $\Pr[g|_S \neq g_S] = O_{K,c}(\epsilon)$. We start by applying the union bound to relate the probability to disagreement on a particular monomial coefficient:

$$\Pr_{S \sim \mu_q}[g|_S \neq g_S] \leq \sum_{|A| \leq d} \Pr_{S \sim \mu_q}[S \supseteq A \text{ and } \tilde{g}(A) \neq \tilde{g}_S(A)] = \sum_{|A| \leq d} q^{|A|} \Pr_{\substack{S \sim \mu_q \\ S \supseteq A}}[\tilde{g}(A) \neq \tilde{g}_S(A)].$$

We chose $\tilde{g}(A)$ as a value a minimizing $\Pr[\tilde{g}_S(A) \neq a \mid S \supseteq A]$, and so

$$\Pr_{S \sim \mu_q}[g|_S \neq g_S] \leq \sum_{|A| \leq d} q^{|A|} \Pr_{\substack{S_1, S_2 \sim \mu_q \\ S \supseteq A}}[\tilde{g}_{S_1}(A) \neq \tilde{g}_{S_2}(A)].$$

We can relate this to the agreement condition using Lemma 4.2:

$$\Pr_{S \sim \mu_q}[g|_S \neq g_S] \leq T_c \sum_{|A| \leq d} q^{|A|} \Pr_{\substack{(S_1, S_2, T) \sim \mu_{q,cq} \\ T \supseteq A}}[\tilde{g}_{S_1}(A) \neq \tilde{g}_{S_2}(A)]$$

$$= T_c \sum_{|A| \leq d} c^{-|A|} \Pr_{(S_1, S_2, T) \sim \mu_{q,cq}}[T \supseteq A \text{ and } \tilde{g}_{S_1}(A) \neq \tilde{g}_{S_2}(A)].$$

In order to use the junta assumption, we move the sum over A inside:

$$\Pr_{S \sim \mu_q}[g|_S \neq g_S] \leq T_c c^{-d} \mathbb{E}_{(S_1, S_2, T) \sim \mu_{q,cq}}[\text{disagreements}(g_{S_1}|_T, g_{S_2}|_T)],$$

where disagreements counts the number of different monomial coefficients. The junta assumption implies that there are at most $2K$ disagreements, and so

$$\Pr_{S \sim \mu_q}[g|_S \neq g_S] \leq T_c c^{-d} 2K \Pr_{(S_1, S_2, T) \sim \mu_{q,cq}}[g_{S_1}|_T \neq g_{S_2}|_T] = O_{K,c}(\epsilon).$$

Sparse junta It remains to show that g is a $(d, O_{K,c}(1/q))$-sparse junta. The main observation is that if $\tilde{g}(A) \neq 0$ then

$$\Pr_{\substack{S \sim \mu_q \\ S \supseteq A}} [\tilde{g}_S(A) = 0] \leq \Pr_{\substack{S \sim \mu_q \\ S \supseteq A}} [\tilde{g}_S(A) = \tilde{g}(A)] \leq \Pr_{\substack{S \sim \mu_q \\ S \supseteq A}} [\tilde{g}_S(A) \neq 0],$$

implying that

$$\Pr_{\substack{S \sim \mu_q \\ S \supseteq A}} [\tilde{g}_S(A) \neq 0] \geq \frac{1}{2},$$

and so

$$\Pr_{S \sim \mu_q} [S \supseteq A \text{ and } \tilde{g}_S(A) \neq 0] \geq \frac{1}{2} q^{|A|}.$$

Given $|R| \leq d$, we deduce that

$$\sum_{\substack{A \in \mathrm{supp}(g) \\ A \supseteq R}} \frac{1}{2} q^{|A|} \leq \sum_{\substack{A \in \mathrm{supp}(g) \\ A \supseteq R}} \Pr_{S \sim \mu_q} [S \supseteq A \text{ and } \tilde{g}_S(A) \neq 0] \leq \mathbb{E}_{S \sim \mu_q} [|\mathrm{supp}(g_S)| \cdot 1_{S \supseteq R}] \leq q^{|R|} K.$$

We conclude that the number of sets $A \in \mathrm{supp}(g)$ containing R of size $|R| + e$ is at most $2Kq^{-e} \leq (2K/q)^e$, as needed. $\square$

4.1 Proof of Lemma 4.2

It remains to prove Lemma 4.2. We first observe that we can sample S_1, S_2 on the left by sampling $S_1', S_2' \sim \mu_q(\{0,1\}^{\overline{A}})$ and then taking $S_b = S_b' \cup A$. We can similarly sample S_1, S_2, T on the right via $(S_1', S_2', T') \sim \mu_{q,cq}(\{0,1\}^{\overline{A}})$. Henceforth we will omit the domain $\{0,1\}^{\overline{A}}$ in such expressions to avoid clutter.

For every $\sigma \in \mathbb{R}$ we define the indicator function $h_\sigma \colon \{0,1\}^{\overline{A}} \to \{0,1\}$ by

$$h_\sigma(S') = [\tilde{g}_{S' \cup A}(A) = \sigma].$$

We can rephrase Lemma 4.2 in these terms as follows:

$$\sum_\sigma \mathbb{E}_{S_1', S_2' \sim \mu_q} [h_\sigma(S_1)(1 - h_\sigma(S_2))] \leq T_c \sum_\sigma \mathbb{E}_{(S_1', S_2', T) \sim \mu_{q,cq}} [h_\sigma(S_1)(1 - h_\sigma(S_2))].$$

Lemma 4.2 thus reduces to the following statement.

Lemma 4.3. *Let $h \colon \{0,1\}^{n'} \to \{0,1\}$ be an indicator function. There exists a constant T_c, depending only on c, such that*

$$\mathbb{E}_{S_1', S_2' \sim \mu_q} [h(S_1')(1 - h(S_2'))] \leq T_c \mathbb{E}_{(S_1', S_2', T) \sim \mu_{q,cq}} [h(S_1')(1 - h(S_2'))].$$

At this point, we invoke Fourier analysis on $(\{0,1\}^{n'}, \mu_q)$. The Fourier basis for this domain is given by the functions

$$\omega_S(y) = \prod_{i \in S} \omega(y), \text{ where } \omega(y) = \frac{y - q}{\sqrt{q(1-q)}}.$$

This basis is orthonormal with respect to μ_q, and the corresponding Fourier expansion of h is

$$h(y) = \sum_U \hat{h}(U) \omega_S(y),$$

where $\hat{h}(\emptyset) = \mathbb{E}[h]$. Since h is Boolean, we furthermore have

$$\mathbb{E}[h] = \mathbb{E}[h^2] = \sum_U \hat{h}(U)^2.$$

19

This allows us to express the left-hand side of Lemma 4.3 as

$$\underset{S_1',S_2'\sim\mu_q}{\mathbb{E}}[h(S_1')(1-h(S_2'))] = \mathbb{E}[h] - \mathbb{E}[h]^2 = \sum_{U\neq\emptyset}\hat{h}(U)^2.$$

In order to express the right-hand side in a similar form, we need to compute

$$\rho_c = \underset{(s_1,s_2,t)\sim\mu_{q,cq}}{\mathbb{E}}[\omega(s_1)\omega(s_2)].$$

Given the value of ρ_c, the expectation on the right-hand side of Lemma 4.3 is

$$\underset{(S_1',S_2',T)\sim\mu_{q,cq}}{\mathbb{E}}[h(S_1')(1-h(S_2'))] = \sum_U\hat{h}(U)^2 - \sum_U\rho_c^{|U|}\hat{h}(U)^2 \geq (1-\rho_c)\sum_{U\neq\emptyset}\hat{h}(U)^2.$$

Therefore Lemma 4.3 holds as long as $T_c \geq 1/(1-\rho_c)$.

In order to complete the proof of Lemma 4.3, we compute ρ_c. If $(s_1,s_2,t) \sim \mu_{q,cq}$ then $\Pr[s_1 = 1] = \Pr[s_2 = 1] = q$ and

$$\Pr[s_1 = s_2 = 1] = cq + (1-cq)\frac{(q-cq)^2}{(1-cq)^2} = \frac{q^2 - 2cq^2 + cq}{1-cq}.$$

Therefore

$$\rho_c = \underset{(s_1,s_2,t)\sim\mu_{q,cq}}{\mathbb{E}}\left[\frac{(s_1-q)(s_2-q)}{q(1-q)}\right] = \frac{1}{q(1-q)}\left[\frac{q^2-2cq^2+cq}{1-cq}-q^2\right] = \frac{cq(1-q)}{(1-cq)q} \leq c,$$

implying that we can take $T_c = 1/(1-c)$.

5 Structure theorem

In this section we prove our main structure theorem, Theorem 1.1, which we reformulate using terminology defined in Section 2.

Theorem 5.1. *Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is ϵ-close to degree d, where $p \leq 1/2$. Then $\Pr[f \neq g] = O(\epsilon)$ and $\mathbb{E}[(f-g)^2] = O(\epsilon)$ for some function $g\colon (\{0,1\}^n, \mu_p) \to \mathbb{R}$ satisfying the following properties:*

(i) If $y \in \{0,1\}^n$ has weight at most d then $g(y) \in \{0,1\}$.

(ii) g is a $(d, O_d(1/p))$-sparse junta.

(iii) For every e, the number of minimal non-Boolean inputs of g of weight e is $O_d(\epsilon/p^e)$.

Conversely, if g satisfies these properties, then it is $O_d(\epsilon)$-close to Boolean.

Corollary 5.2. *Let f be as in the theorem. Then f is $O_d(\epsilon)$-close to $\mathrm{round}(g, \{0,1\})$, where g is the function in the theorem.*

Conversely, if g satisfies the conditions in the theorem, then $\mathrm{round}(g, \{0,1\})$ is $O_d(\epsilon)$-close to degree d.

Proof of Corollary 5.2. Let f be as in the theorem. Then

$$\Pr[f \neq \mathrm{round}(g, \{0,1\})] \leq \Pr[f \neq g] = O(\epsilon).$$

Conversely, suppose that g satisfies the conditions in the theorem. Then g has degree d and

$$\mathbb{E}[(\mathrm{round}(g, \{0,1\}) - g)^2] = O_d(\epsilon)$$

since g is $O_d(\epsilon)$-close to Boolean. $\qquad\square$

Henceforth, all big O constants will depend on d, and we do not mention this explicitly.

Most of the effort will be concentrated on proving the direct part of the theorem, which we will do in two steps:

1. We first construct a sparse junta h which is $O(\epsilon)$-close to f.

2. We then modify h to a sparse junta g which satisfies the first two properties in the theorem, and show that it satisfies the third property as well.

We then prove the easier converse part of the theorem.

5.1 Step 1: Agreement

Since f is ϵ-close to degree d, there is a degree d function $f^{\leq d}$ such that $\mathbb{E}_{\mu_p}[(f - f^{\leq d})^2] \leq \epsilon$.

For $S \subseteq [n]$, define
$$\epsilon_S = \mathop{\mathbb{E}}_{\mu_{1/2}} [(f|_S - f^{\leq d}|_S)^2],$$

where $\mu_{1/2}$ refers to the distribution $\mu_{1/2}(\{0,1\}^S)$; below we usually omit the domain. Since sampling $S \sim \mu_{2p}$ and $y \sim \mu_{1/2}(\{0,1\}^S)$ is the same as sampling $y \sim \mu_p$,
$$\mathop{\mathbb{E}}_{S \sim \mu_{2p}} [\epsilon_S] = \mathop{\mathbb{E}}_{\mu_p}[(f - f^{\leq d})^2] \leq \epsilon.$$

Applying the Kindler–Safra theorem (Theorem 2.3), there exist Boolean degree d functions g_S such that $\mathbb{E}_{\mu_{1/2}}[(f|_S - g_S)^2] = O(\epsilon_S)$, and so
$$\mathop{\mathbb{E}}_{S \sim \mu_{2p}} \left[\mathop{\mathbb{E}}_{\mu_{1/2}} [(f|_S - g_S)^2] \right] = O\left(\mathop{\mathbb{E}}_{S \sim \mu_{2p}} [\epsilon_S] \right) = O(\epsilon).$$

We would like to paste the functions g_S to a global function h using the junta agreement theorem (Theorem 4.1), which we will apply with $K = M_d$ (the value from Theorem 2.1), $c = \sqrt{1/2}$, and $q = 2p$. Theorem 2.1 shows that each g_S is a K-junta. In order to apply the junta agreement theorem, we need to verify the agreement condition:
$$\mathop{\Pr}_{(S_1, S_2, T) \sim \mu_{2p, \sqrt{2}p}} [g_{S_1}|_T \neq g_{S_2}|_T] = O(\epsilon).$$

To see this, first notice that since g_{S_1} and g_{S_2} are both K-juntas, if $g_{S_1}|_T \neq g_{S_2}|_T$ then $\Pr_{\mu_{\sqrt{1/2}}}[g_{S_1}|_T \neq g_{S_2}|_T] \geq \sqrt{1/2}^{2K} = 2^{-K}$. Since g_{S_1} and g_{S_2} are Boolean, this shows that
$$\mathop{\Pr}_{(S_1, S_2, T) \sim \mu_{2p, \sqrt{2}p}} [g_{S_1}|_T \neq g_{S_2}|_T] \leq 2^K \mathop{\mathbb{E}}_{(S_1, S_2, T) \sim \mu_{2p, \sqrt{2}p}} \left[\mathop{\mathbb{E}}_{\mu_{\sqrt{1/2}}} [(g_{S_1}|_T - g_{S_2}|_T)^2] \right].$$

Applying the inequality $(g_{S_1}|_T - g_{S_2}|_T)^2 \leq 2(g_{S_1}|_T - f|_T)^2 + 2(g_{S_2}|_T - f|_T)^2$, we obtain
$$\mathop{\Pr}_{(S_1, S_2, T) \sim \mu_{2p, \sqrt{2}p}} [g_{S_1}|_T \neq g_{S_2}|_T] \leq 2^{K+1} \mathop{\mathbb{E}}_{(S_1, S_2, T) \sim \mu_{2p, \sqrt{2}p}} \left[\mathop{\mathbb{E}}_{\mu_{\sqrt{1/2}}} [(g_{S_1}|_T - f|_T)^2] \right].$$

We can choose (S_1, T) in the following way: choose $S_1 \sim \mu_{2p}$, and choose $T \sim \mu_{\sqrt{1/2}}(S_1)$. If we then choose $y \sim \mu_{\sqrt{1/2}}(T)$ then $y \sim \mu_{1/2}(S_1)$. This shows that
$$\mathop{\Pr}_{(S_1, S_2, T) \sim \mu_{2p, \sqrt{2}p}} [g_{S_1}|_T \neq g_{S_2}|_T] \leq 2^{K+1} \mathop{\mathbb{E}}_{S \sim \mu_{2p}} \left[\mathop{\mathbb{E}}_{\mu_{1/2}} [(g_S - f|_S)^2] \right] = O(\epsilon),$$

and so the agreement condition is satisfied.

Applying the junta agreement theorem, we obtain a $(d, O(1/p))$-sparse junta h such that

$$\Pr_{S \sim \mu_{2p}}[h|_S \neq g_S] = O(\epsilon).$$

Furthermore, each monomial coefficient of h is a monomial coefficient of some Boolean K-junta. Since there are finitely many Boolean K-juntas, h is B-quantized for some finite set B depending only on d.[1]

To conclude this step of the proof, we need to show that $\mathbb{E}_{\mu_p}[(f-h)^2] = O(\epsilon)$. We first show that $\Pr_{\mu_p}[f \neq h] = O(\epsilon)$. Indeed,

$$\Pr_{\mu_p}[f \neq h] = \mathop{\mathbb{E}}_{S \sim \mu_{2p}}\left[\Pr_{\mu_{1/2}}[f|_S \neq h|_S]\right] \leq \mathop{\mathbb{E}}_{S \sim \mu_{2p}}\left[\Pr_{\mu_{1/2}}[f|_S \neq g_S] + [g_S \neq h|_S]\right] = O(\epsilon).$$

(The term $[g_S \neq g|_S]$ is an indicator.)

In order to convert this L_0 guarantee to an L_2 guarantee, we use Lemma 3.11, which implies that

$$\mathop{\mathbb{E}}_{\mu_p}[\mathrm{dist}(h, \{0,1\})^2] = O(\Pr_{\mu_p}[h \notin \{0,1\}]) = O(\Pr_{\mu_p}[h \neq f]) = O(\epsilon).$$

This implies that

$$\mathop{\mathbb{E}}_{\mu_p}[(f-h)^2] \leq 2\mathop{\mathbb{E}}_{\mu_p}[(f - \mathrm{round}(h, \{0,1\}))^2] + 2\mathop{\mathbb{E}}_{\mu_p}[(\mathrm{round}(h, \{0,1\}) - h)^2] = 2\Pr_{\mu_p}[f \neq h] + O(\epsilon) = O(\epsilon).$$

Concluding, in this step we have constructed a function h satisfying the following properties:

(i) h is a $(d, C/p)$-sparse junta.

(ii) h is B-quantized.

(iii) $\Pr[h \neq f] = O(\epsilon)$.

(iv) $\mathbb{E}[(h-f)^2] = O(\epsilon)$.

In the final two properties, the underlying distribution is μ_p, which we assume henceforth.

5.2   Step 2: Fixing coefficients

In this step we modify h to a function g which satisfies all properties listed in Theorem 5.1. We do this by constructing a sequence of functions $h = h_{-1}, h_0, \ldots, h_d = g$, where h_e satisfies the following properties:

(i) h_e is a $(d, C_e/p)$-sparse junta, where C_e depends only on d, e.

(ii) h_e is B_e-quantized for some finite B_e depending only on d, e.

(iii) $\Pr[h_e \neq f] = O_{d,e}(\epsilon)$.

(iv) $\mathbb{E}[(h_e - f)^2] = O_{d,e}(\epsilon)$.

(v) $h_e(y) \in \{0,1\}$ if y has weight at most e.

The function h clearly satisfies the properties for $e = -1$. We now show how to construct h_e from h_{e-1} for $e \in \{0, \ldots, d\}$.

Given h_{e-1}, let $\mathcal{F}$ consist of all minimal non-Boolean inputs of h_{e-1} of weight e. We define h_e by modifying $\tilde{h}_e(S)$ for each $1_S \in \mathcal{F}$ so that $h_e(1_S) = 0$:

$$\tilde{h}_e(S) = -\sum_{T \subsetneq S} \tilde{h}_{e-1}(T).$$

All other monomial coefficients stay the same.

---

[1]In fact, we can take $B = \{a/2^{d-1} : a \in \{-2^{d-1}, \ldots, 2^{d-1}\}\}$.

Item i. Lemma 2.6 shows that $\mathcal{F} \subseteq \mathrm{supp}(h_{e-1})^{\sqcup L_d}$. Lemma 3.5 shows that $\mathcal{G} = \mathrm{supp}(h_{e-1})^{\sqcup L_d}$ is $(L_d d, C'_e/p)$-sparse, for some constant C'_e. Since $\mathrm{supp}(h_e) \subseteq \mathrm{supp}(h_{e-1}) \cup \mathcal{F}$, Lemma 3.5 shows that h_e is a $(d, C_e/p)$-sparse junta, for some constant C_e.

Item ii. Each modified monomial coefficient is the negated sum of $2^e - 1$ monomial coefficients of h_{e-1}, and so h_e is B_e-quantized for some finite set B_e.

Item iii. If $h_e(y) \neq h_{e-1}(y)$ then $y_S = 1$ for some $S \in \mathcal{F}$, and so

$$\Pr[h_e \neq h_{e-1}] \leq \sum_{S \in \mathcal{F}} p^{|S|}.$$

In order to bound this, we apply the reverse union bound (Lemma 3.7), concluding that

$$\Pr[h_e \neq h_{e-1}] = O(\Pr[\mathcal{E}]),$$

where $\mathcal{E}$ is the event that for some $S \in \mathcal{F}$, it holds that $y_S = 1$ and $y_R = 0$ for all $R \in \mathcal{G}$ such that $R \not\subseteq S$. If $\mathcal{E}$ happens then $h_{e-1}(y) = h_{e-1}(1_S) \notin \{0, 1\}$, and so

$$\Pr[h_e \neq h_{e-1}] = O(\Pr[h_{e-1} \notin \{0,1\}]) = O(\Pr[h_{e-1} \neq f]) = O(\epsilon).$$

Therefore

$$\Pr[h_e \neq f] \leq \Pr[h_e \neq h_{e-1}] + \Pr[h_{e-1} \neq f] = O(\epsilon).$$

Item iv. Lemmas 3.5 and 3.9 show that $h_e - h_{e-1}$ is a B_e^+-quantized $(d, O(1/p))$-sparse junta, for some finite set B_e^+. Therefore Lemma 3.10, applied with $f = h_e - h_{e-1}$, shows that

$$\mathbb{E}[(h_{e-1} - h_e)^2] = O(\Pr[h_{e-1} \neq h_e]) = O(\epsilon).$$

Consequently

$$\mathbb{E}[(h_e - f)^2] \leq 2\,\mathbb{E}[(h_e - h_{e-1})^2] + 2\,\mathbb{E}[(h_{e-1} - f)^2] = O(\epsilon).$$

Item v. Since we only modified monomial coefficients of size e, clearly $h_e(y) \in \{0, 1\}$ for y of weight less than e. By construction, $h_e(y) \in \{0, 1\}$ for y of weight exactly e.

Concluding the proof

Taking $g = h_d$, the function g satisfies all properties stated in Theorem 5.1, except for Item iii, which states that the number of minimal non-Boolean inputs of g of weight e is $O(\epsilon/p^e)$.

Let $\mathcal{F}$ consist of all minimal non-Boolean inputs of g. Lemma 2.6 shows that $\mathcal{F} \subseteq \mathcal{G} = \mathrm{supp}(g)^{\sqcup L_d}$. Applying Lemma 3.5 and the reverse union bound just as above, we conclude that

$$\sum_{S \in \mathcal{F}} p^{|S|} = O(\Pr[g \notin \{0,1\}]) = O(\Pr[g \neq f]) = O(\epsilon).$$

Therefore $\mathcal{F}$ contains at most $O(\epsilon/p^e)$ inputs of weight e.

5.3 Converse part

We conclude the proof of Theorem 5.1 by proving the converse direction. Suppose that g satisfies the three properties in the statement of the theorem:

(i) $g(y) \in \{0, 1\}$ whenever $y \in \{0, 1\}^n$ has weight at most d.

23

(ii) g is a $(d, O(1/p))$-sparse junta.

(iii) g has $O(\epsilon/p^e)$ minimal non-Boolean inputs of weight e, for every e.

We need to show that g is $O(\epsilon)$-close to Boolean.

We start by showing that the monomial coefficients of g are quantized. Theorem 2.1 implies that Boolean degree d functions are B-quantized for some finite B depending only on d. If $|A| \leq d$ then $g|_A$ is a Boolean degree d function by Item i, and so $\tilde{g}(A) \in B$. In other words, g is also B-quantized.

If $g(y) \notin \{0,1\}$ then $y \geq z$ for some minimal non-Boolean input z. If z has weight e then the probability that $y \geq z$ is p^e. Lemma 2.6 shows that z has weight at most L_d, and so

$$\Pr[g \notin \{0,1\}] \leq \sum_{e \leq L_d} p^e \cdot O(\epsilon/p^e) = O(\epsilon).$$

Lemma 3.11 concludes that g is $O(\epsilon)$-close to Boolean.

6 Monotone structure theorems

In this section we prove Theorem 1.2. We start by proving a monotone version of Theorem 5.1, which is interesting in its own right, and then deduce Theorem 1.2 from it.

6.1 Monotone version of Theorem 5.1

We start by proving a monotone version of Theorem 5.1.

Theorem 6.1. *Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is monotone and ϵ-close to degree d, where $p \leq 1/2$. Then f is $O(\epsilon)$-close to a function $g\colon (\{0,1\}^n, \mu_p) \to \mathbb{R}$ satisfying the following properties:*

(i) If $|S| \leq d$ then $g|_S$ is monotone and Boolean.

(ii) g is a $(d, O_d(1/p))$-sparse junta.

(iii) For every e, the number of minimal non-Boolean inputs of g of weight e is $O_d(\epsilon/p^e)$.

(iv) For every e, the number of minimal non-monotone Boolean inputs of g of weight e is $O_d(\epsilon/p^e)$.

Conversely, if g satisfies these properties then it is $O_d(\epsilon)$-close to some monotone Boolean function.

Corollary 6.2. *Let f be as in theorem. Then f is $O_d(\epsilon)$-close to $\mathrm{round}(g, \{0,1\})$, where g is the function in the theorem.*

Conversely, if g satisfies the conditions in the theorem, then $\mathrm{round}(g, \{0,1\})$ is $O_d(\epsilon)$-close to a monotone Boolean function which is $O_d(\epsilon)$-close to degree d.

Proof of Corollary 6.2. The first statement follows as in the proof of Corollary 5.2. Now suppose that g satisfies the conditions in the theorem. According to the theorem, $\mathbb{E}[(g - f)^2] = O_d(\epsilon)$ for some monotone Boolean function f. Since g has degree d, f is $O_d(\epsilon)$-close to degree d. $\square$

The proof of the direct part of Theorem 6.1 closely resembles the proof of the direct part of Theorem 5.1, while the proof of the converse part is a little bit more tricky. For both proofs, we suppress the dependence of big O constants on d.

Direct part

Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is monotone and ϵ-close to degree d. The argument in Section 5.1 applies, constructing a function h which satisfies the properties stated there (we repeat them below as part of the proof). We now run an argument in the style of Section 5.2, constructing a sequence of functions $h = h_{-1}, h_0, \ldots, h_d = g$, where h_e satisfies the following properties:

  (i) h_e is a $(d, C_e/p)$-sparse junta, where C_e depends only on d, e.

  (ii) h_e is B_e-quantized for some finite B_e depending only on d, e.

  (iii) $\Pr[h_e \neq f] = O_{d,e}(\epsilon)$.

  (iv) $\mathbb{E}[(h_e - f)^2] = O_{d,e}(\epsilon)$.

  (v) $h_e|_S$ is a monotone Boolean function if $|S| \leq e$.

The function h constructed in Section 5.1 satisfies these properties for $e = -1$. We now show how to construct h_e from h_{e-1} for $e \in \{0, \ldots, d\}$, using a modification of the construction in Section 5.2.

Given h_{e-1}, let $\mathcal{F}_b$ consist of all minimal non-Boolean inputs of h_{e-1} of weight e, let $\mathcal{F}_m$ consist of all minimal non-monotone Boolean inputs of h_{e-1} of weight e, and let $\mathcal{F} = \mathcal{F}_b \cup \mathcal{F}_m$. We define h_e by modifying $\tilde{h}_e(S)$ for each $1_S \in \mathcal{F}$ so that $h_e(1_S) = 1$. All other monomial coefficients stay the same.

Item i follows just as in Section 5.2, using Lemma 2.7 in addition to Lemma 2.6. We need to take $\mathcal{G} = \mathrm{supp}(f)^{\sqcup \max(L_d, M_d)}$. Item ii follows as in Section 5.2.

Item iii requires a new argument. If $h_e(y) \neq h_{e-1}(y)$ then $y_S = 1$ for some $S \in \mathcal{F}$, and so

$$\Pr[h_e \neq h_{e-1}] \leq \sum_{S \in \mathcal{F}_b} p^{|S|} + \sum_{S \in \mathcal{F}_m} p^{|S|}.$$

The argument in Section 5.2 shows that the first term on the right is $O(\epsilon)$. In order to bound the second term, we apply the reverse union bound (Lemma 3.7) with $p := 2p$, concluding that

$$\sum_{S \in \mathcal{F}_m} p^{|S|} \leq \sum_{S \in \mathcal{F}_m} (2p)^{|S|} = O(\Pr_{\mu_{2p}}[\mathcal{E}]),$$

where $\mathcal{E}$ is the event that for some $S \in \mathcal{F}$, it holds that $y_S = 1$ and $y_R = 0$ for all $R \in \mathrm{supp}(\mathcal{G})$ such that $R \nsubseteq S$.

If $\mathcal{E}$ happens for some S, an event we denote by $\mathcal{E}_S$, then $h_{e-1}|_S$ is a non-monotone Boolean function of degree d, which is an M_d-junta by Theorem 2.1. Therefore Lemma 2.4 implies that $\Pr_{\mu_{1/2}}[h_{e-1}|_S \neq f|_S] \geq 2^{-M_d}$. Thus

$$\Pr_{\mu_p}[h_{e-1} \neq f] = \Pr_{\substack{S \sim \mu_{2p} \\ y \sim \mu_{1/2}(S)}}[h(y) \neq f(y)] \geq 2^{-M_d} \Pr_{\mu_{2p}}[\mathcal{E}].$$

Since the left-hand side is $O(\epsilon)$, we conclude that

$$\sum_{S \in \mathcal{F}_m} p^{|S|} = O(\Pr_{\mu_{2p}}[\mathcal{E}]) = O(\epsilon).$$

This shows that $\Pr[h_e \neq h_{e-1}] = O(\epsilon)$, and so $\Pr[h_e \neq f] \leq \Pr[h_e \neq h_{e-1}] + \Pr[h_{e-1} \neq f] = O(\epsilon)$.

Item iv follows from Item iii as in Section 5.2. Item v follows by construction.

Concluding the proof   Taking $g = h_d$, the function g satisfies all properties stated in Theorem 6.1 except for Items iii and iv.

Item iii follows as in Section 5.2. In order to prove Item iv, let $\mathcal{F}_m$ denote all minimal non-monotone Boolean inputs of g. The proof of Item iii above shows that

$$\sum_{S \in \mathcal{F}_m} p^{|S|} = O(\epsilon),$$

from which Item iv of the theorem immediately follows.

25

Converse part

We conclude the proof of Theorem 6.1 by proving the converse direction. Suppose that g satisfies the properties in the statement of the theorem:

(i) $g(y) \in \{0,1\}$ whenever $y \in \{0,1\}^n$ has weight at most d.

(ii) g is a $(d, O(1/p))$-sparse junta.

(iii) g has $O(\epsilon/p^e)$ minimal non-Boolean inputs of weight e, for all e.

(iv) g has $O(\epsilon/p^e)$ minimal non-monotone Boolean inputs of weight e, for all e.

We need to show that g is $O(\epsilon)$-close to some monotone Boolean function f.

Theorem 5.1 shows that g is $O(\epsilon)$-close to the Boolean function $G = \mathrm{round}(g, \{0,1\})$. However, G is not necessarily monotone. Accordingly, we define

$$f(y) = \max_{z \leq y} G(z),$$

which is monotone by construction. It remains to show that g is $O(\epsilon)$-close to f. We will show that G is $O(\epsilon)$-close to f. This would imply that

$$\mathbb{E}[(g-f)^2] \leq 2\,\mathbb{E}[(g-G)^2] + 2\,\mathbb{E}[(G-f)^2] = O(\epsilon) + 2\Pr[G \neq f] = O(\epsilon).$$

If $y = 1_S \in \{0,1\}^n$ is such that no $z \leq y$ is a minimal non-Boolean input or a minimal non-monotone Boolean input of g then $G|_S = g|_S$ and $g|_S$ is monotone, implying that $f(y) = G(y)$. Therefore if $f(y) \neq G(y)$ then $y \geq z$ for some input $z = 1_T$ which is either a minimal non-Boolean input of g or a minimal non-monotone Boolean input of g. Since $y \geq z$ with probability $p^{|T|}$, if we let $\mathcal{F}$ consist of all minimal non-Boolean and minimal non-monotone Boolean inputs of g then

$$\Pr[G \neq f] \leq \sum_{1_T \in \mathcal{F}} p^{|T|} \leq \sum_{e \leq \max(L_d, M_d)} p^e \cdot O(\epsilon/p^e) = O(\epsilon),$$

using Lemmas 2.6 and 2.7 to bound the size of an input in $\mathcal{F}$. This concludes the proof.

6.2 Approximation by monotone DNF

We now deduce Theorem 1.2, which we rephrase as follows.

Theorem 6.3. *Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is monotone and ϵ-close to degree d, where $p \leq 1/2$. Then f is $O(\epsilon)$-close to a monotone DNF F which satisfies the following properties:*

(i) *F is a $(d, O_d(1/p))$-sparse DNF.*

(ii) *For every e, the number of minimal high-degree inputs of F of weight e is $O_d(\epsilon/p^e)$.*

Conversely, if F is a monotone DNF satisfying these properties then it is $O_d(\epsilon)$-close to degree d.

We derive Theorem 6.3 from Theorem 6.1. In the proof, we suppress the dependence of the big O constants on d.

Direct part

Suppose that $f\colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ is monotone and ϵ-close to degree d. Applying the direct part of Theorem 6.1, we obtain a function g satisfying the stated properties which is $O(\epsilon)$-close to f. Applying the converse part of Theorem 6.1, we obtain a monotone Boolean function F which is $O(\epsilon)$-close to g, and so $O(\epsilon)$-close to f.

We first show that all minterms of F have weight at most d, using the way that F is constructed in the proof of the converse part of Theorem 6.1. Recall that $G(y) = \text{round}(g(y), \{0,1\})$. Suppose that 1_S is an input of weight larger than d such that $F(z) = 0$ for all $z < 1_S$, and so $G(z) = 0$ for all $z < 1_S$. If $z = 1_T$ has weight at most d then $g|_T = G|_T = 0$, and so $\tilde{g}(A) = 0$ for all $A \subseteq T$. It follows that $\tilde{g}(A) = 0$ for all $A \subseteq S$, and so $G(z) = 0$ for all $z \leq 1_S$, implying that $F(1_S) = 0$. Thus 1_S cannot be a minterm of F.

Thus F is a monotone width d DNF, whose minterms are all y of weight at most d such that $g(y) = 1$ and $g(z) = 0$ for all $z < y$. If $y = 1_S$ is such a minterm then $g|_S = y_S$, and so $\tilde{g}(S) = 1$. Therefore the minterm support of F is contained in the monomial support of g. Since g is $(d, O(1/p))$-sparse, it follows that F is also $(d, O(1/p))$-sparse, proving Item i.

It remains to prove Item ii. We first show that $F|_S = g|_S$ if and only if $F|_S$ has degree d. If $F|_S = g|_S$ then $F|_S$ has degree d. Conversely, suppose that $F|_S$ has degree d. Since $g|_A$ is Boolean and monotone whenever $|A| \leq d$, we see that $F|_A = g|_A$ and so $\tilde{f}(A) = \tilde{g}_A$ for all $A \subseteq S$ of size at most d. Since $F|_S$ has degree d, this means that $F|_S = g|_S$.

We conclude that if y is a minimal high-degree input of F then it is a minimal input such that $F(y) \neq g(y)$. In particular, $g|_T = F|_T$ is Boolean and monotone for all $1_T < y$. Since $F(y) \neq g(y)$, either $g(y) \neq G(y)$ or $g(y) = G(y) \neq F(y)$. If $g(y) \neq G(y)$ then $g(y) \notin \{0,1\}$, and so y is a minimal non-Boolean input of g. If $g(y) = G(y) \neq F(y)$ then y is a minimal non-monotone Boolean input of g. Item ii now follows from Items iii and iv of Theorem 6.1.

Converse part

Let F be a monotone DNF satisfying the properties listed in Theorem 6.3. Define

$$g = \sum_{|A| \leq d} \tilde{F}(A) y_A.$$

Observe that $g|_S = F|_S$ if and only if $F|_S$ has degree d. In particular, if $g(y) \neq F(y)$ then $y \geq z$ for some minimal high-degree input z of g. Denoting by $\mathcal{H}$ the set of all minimal high-degree inputs of g, this implies that

$$\Pr[g \neq F] \leq \sum_{1_T \in \mathcal{H}} p^{|T|} \leq \sum_{e \leq L_d} p^e \cdot O(\epsilon/p^e) = O(\epsilon),$$

using Lemma 2.8 to bound the weight of a minimal high-degree input.

In order to upgrade this to an L_2 bound, we appeal to Lemma 3.11. First, we need to show that g is a quantized sparse junta.

Since every monomial coefficient of g is a monomial coefficient of some (monotone) Boolean function on d coordinates, g is B-quantized for some finite B.

In order to show that g is a sparse junta, let $\mathcal{M}$ be the set of minterms of F, and observe that

$$F = \sum_e (-1)^e \sum_{S_1,\ldots,S_e \in \mathcal{M}} y_{S_1 \cup \cdots \cup S_e}.$$

In particular, $\text{supp}(g) \subseteq \text{supp}(F)^{\sqcup d}$, where $\text{supp}(F)$ is the monomial support of F. Since F is a $(d, O(1/p))$-sparse DNF, Lemma 3.5 implies that g is a $(d, O(1/p))$-sparse junta.

Applying Lemma 3.11, we see that

$$\mathbb{E}[\text{dist}(g, \{0,1\})^2] = O(\Pr[g \notin \{0,1\}]) = O(\Pr[g \neq F]) = O(\epsilon).$$

Therefore

$$\mathbb{E}[(g - F)^2] \leq 2\,\mathbb{E}[(g - \text{round}(g, \{0,1\})^2] + 2\,\mathbb{E}[(\text{round}(g, \{0,1\}) - F)^2] \leq O(\epsilon) + 2\Pr[g \neq F] = O(\epsilon).$$

27

7 Junta approximation

In this section we prove Theorems 1.3 and 1.4. Before restating them, we formally define the parameter appearing in them.

Definition 7.1. Given $d \in \mathbb{N}$, the parameter m_d is the maximum m such that there exists a degree d polynomial P satisfying $P(0), \ldots, P(m-1) \in \{0,1\}$ and $P(0) \neq P(1)$.

It might not be clear from the definition that $m_d < \infty$. In fact, $m_d \leq 2d$. To see this, suppose that P is a degree d polynomial P such that $P(0), \ldots, P(m-1) \in \{0,1\}$ and $P(0) \neq P(1)$. Since $P(0) \neq P(1)$, the polynomial P is not constant, and so there can be at most d values x such that $P(x) = 0$ and at most d values y such that $P(y) = 1$. Consequently $m \leq 2d$.

We can now restate the two theorems from the introduction, as a single theorem.

Theorem 7.2. *Suppose that $f : (\{0,1\}^n, \mu_p) \to \{0,1\}$ is ϵ-close to degree d, where $p \leq 1/2$. Then for some constant μ_d depending only on d,*

(a) f is $O_d(\epsilon^{1/m_d} + p)$-close to a constant Boolean function.

(b) f is $O_d(\epsilon^{1/\mu_d})$-close to a Boolean degree d function.

Furthermore, the first item becomes false if we replace m_d by any smaller constant (for any value of the constant hidden in the big O).

In the rest of this section, we suppress the dependence of big O constants on d.

The constant μ_d. We do not know the optimal value of μ_d, and in particular, whether it equals m_d. If we consider A-*valued* functions instead of Boolean functions, a setting we define below, then it is easy to construct examples in which $m_{d,A} \neq \mu_{d,A}$, where $m_{d,A}, \mu_{d,A}$ are the corresponding parameters in the A-valued setting.

Let $A \subseteq \mathbb{R}$ be a finite set. If $f(x) \in A$ for all x in the domain of A then we say that f is A-valued. A $\{0,1\}$-valued function is thus the same as a Boolean function. Everything in this paper generalizes to the A-valued setting, once we generalize the Kindler–Safra theorem to the A-valued setting (this can be done in a black-box way, as shown in the full version of [DFH19], available on arXiv (1711.09428) and ECCC (TR17-180)).

When $d = 1$, it is not hard to check that $m_{1,A}$ is the size of the largest arithmetic progression in A. Therefore if a, b, c are linearly independent over $\mathbb{Q}$ then the alphabet

$$A = \{0, a, b, c, a+b, a+c, b+c\}$$

satisfies $m_{1,A} = 2$. Now consider the function

$$f = ay_1 + by_2 + cy_3.$$

By construction, f is $\Theta(p^3)$-close to A, and so $F = \text{round}(f, A)$ is $\Theta(p^3)$-close to degree 1. All A-valued degree 1 functions depend on at most two coordinates, and so are $\Omega(p)$-far from f. This shows that $\mu_{d,A} \geq 3$.

7.1 Converse part

The converse part of Theorem 7.2 is easy to prove, and serves to motivate the statement.

Proof of converse part. Let P be a degree d polynomial such that $P(0), \ldots, P(m_d - 1) \in \{0,1\}$ and $P(0) \neq P(1)$. For any $p \leq 1/2$ and $\epsilon \leq 1$, define

$$f_{p,\epsilon} = P(y_1 + \cdots + y_N), \text{ where } N = \epsilon^{1/m_d}/p.$$

We assume that N is an integer, and in particular, $\epsilon^{1/m_d} \geq p$.

If fewer than m_d of $y_1, \ldots, y_N$ evaluate to 1 then $f_{p,\epsilon} \in \{0,1\}$, and so

$$\Pr[f_{p,\epsilon} \notin \{0,1\}] \leq N^{m_d} p^{m_d} = \epsilon.$$

In order to upgrade this to an L_2 bound, we use Lemma 3.11. Since $y_1 + \cdots + y_N$ is a $(1, 1/p)$-sparse junta, Lemma 3.5 shows that $f_{p,\epsilon}$ is a $(d, O(1/p))$-sparse junta. In order to show that $f_{p,\epsilon}$ is quantized, notice that we can find coefficients c_e such that

$$P(n) = \sum_{e \leq d} c_e \binom{n}{e},$$

and so

$$f_{p,\epsilon} = \sum_{e \leq d} c_e \sum_{\substack{S \subseteq [N] \\ |S|=e}} y_S,$$

showing that $f_{p,\epsilon}$ is $\{c_0, \ldots, c_d\}$-quantized. Therefore Lemma 3.11 shows that $f_{p,\epsilon}$ is $O(\epsilon)$-close to Boolean.

We can now show that Item a cannot be improved. Notice that

$$\Pr[f_{p,\epsilon} = P(0)] \geq (1-p)^N \geq 4^{-\epsilon^{1/m_d}} = \Omega(1),$$

using the fact that $(1-p)^{1/p}$ is monotone decreasing. Similarly,

$$\Pr[f_{p,\epsilon} = P(1)] \geq Np(1-p)^{N-1} = \Omega(\epsilon^{1/m_d}).$$

It follows that for any constant c,

$$\mathbb{E}[(f_{p,\epsilon} - c)^2] \geq \frac{1}{4} \Pr[f_{p,\epsilon} \neq \mathrm{round}(c, \{0,1\})] = \Omega(\epsilon^{1/m_d}).$$

Choosing a sequence of p, ϵ with $\epsilon \to 0$ and $p \leq \epsilon^{1/m_d}$, this shows that the constant m_d cannot be improved in Item a. $\qquad\square$

7.2 Direct part

Let $f \colon (\{0,1\}^n, \mu_p) \to \{0,1\}$ be ϵ-close to degree d, where $p \leq 1/2$. Applying Theorem 5.1, we obtain a function $g \colon (\{0,1\}^n, \mu_p) \to \mathbb{R}$ which is $O(\epsilon)$-close to f and satisfies the following properties:

(i) g is a $(d, O(1/p))$-sparse junta.

(ii) $g(y) \in \{0,1\}$ whenever y as weight at most d, and consequently g is B-quantized for some finite set B depending only on d.

(iii) The number of minimal non-Boolean inputs of g of weight e is $O(\epsilon/p^e)$, for all e.

Item ii implies that $\tilde{g}(\emptyset) = g(0) \in \{0,1\}$. We assume without loss of generality that $g(0) = 0$: otherwise, replace f, g, and the approximating function h by their "negations" $1 - f, 1 - g, 1 - h$.

Let $\mathcal{M}$ consist of all inclusion-minimal sets in $\mathrm{supp}(g)$, and let $\mathcal{M}_e = \{S \in \mathcal{M} : |S| = e\}$. The main part of the proof is the following Ramsey-theoretic lemma.

Lemma 7.3. *For every $e \leq d$, either (i) $|\mathcal{M}_e| = O(1)$ or (ii) there exists a degree d polynomial P such that $P(0) = g(0)$, $P(1) \neq P(0)$, and for every $m \leq m_d$ such that $m \geq 1$, there are $s \leq em$ and a collection $\mathcal{F}$ of $\Omega(|\mathcal{M}_e|^m)$ sets such that each $S \in \mathcal{F}$ satisfies the following:*

(a) S has size s.

(b) S is the union of m sets in $\mathcal{M}_e$.

(c) $g(1_S) = P(m)$.

Let us see how this lemma implies the theorem. We start with the following corollary.

Corollary 7.4. *For every $e \leq d$,*
$$|\mathcal{M}_e| = O(1 + \epsilon^{1/m_d}/p^e).$$

Proof. Let $e \leq d$, and apply Lemma 7.3. If Case (i) holds then we are done, so suppose that Case (ii) holds, for some P.

By the definition of m_d, there exists $m \leq m_d$ such that $P(m) \notin \{0,1\}$. Let $s, \mathcal{F}$ be as promised by the Lemma for m. We would like to deduce a lower bound on $\Pr[g \notin \{0,1\}]$ from the lower bound on $|\mathcal{F}|$, and for this we apply the reverse union bound, Lemma 3.7.

Lemma 7.3 guarantees that $\mathcal{F} \subseteq \operatorname{supp}(g)^{\sqcup m}$. Lemma 3.5 shows that $\mathcal{G} = \operatorname{supp}(g)^{\sqcup m}$ is $(md, O(1/p))$-sparse, and so the reverse union bound shows that
$$\Pr[\mathcal{E}] = \Omega(|\mathcal{F}|p^s) = \Omega(|\mathcal{F}|p^{em}),$$
where $\mathcal{E}$ is the event that $y_S = 1$ for some $S \in \mathcal{F}$ and $y_R = 0$ for all $R \in \mathcal{G}$ such that $R \not\subseteq S$. If $\mathcal{E}$ happens then $g(y) = g(1_S) = P(m) \notin \{0,1\}$, and so
$$\Pr[\mathcal{E}] \leq \Pr[g \notin \{0,1\}] \leq \Pr[g \neq f] = O(\epsilon).$$

Consequently $|\mathcal{F}| = O(\epsilon/p^{em})$. Since $|\mathcal{F}| = \Omega(|\mathcal{M}_e|^m)$, we conclude that $|\mathcal{M}_e| = O(\epsilon^{1/m}/p^e) = O(\epsilon^{1/m_d}/p^e)$. $\square$

We can now finish the proof of Item a of Theorem 7.2.

Proof of Item a. If $g(y) \neq 0$ then $y_S = 1$ for some $S \in \mathcal{M}$. Therefore Corollary 7.4 shows that
$$\Pr[g \neq 0] \leq \sum_{e=1}^{d} p^e |\mathcal{M}_e| = \sum_{e=1}^{d} O(p^e + \epsilon^{1/m_d}) = O(p + \epsilon^{1/m_d}),$$
and so $\Pr[f \neq 0] \leq \Pr[f \neq g] + \Pr[g \neq 0] = O(\epsilon^{1/m_d} + p)$.

Since g is a $(d, O(1/p))$-sparse junta, Lemma 3.10 implies that $\mathbb{E}[(g-0)^2] = O(\epsilon^{1/m_d} + p)$, and so $\mathbb{E}[(f-0)^2] \leq 2\mathbb{E}[(f-g)^2] + 2\mathbb{E}[(g-0)^2] = O(\epsilon^{1/m_d} + p)$. $\square$

The proof of Item b will proceed by bounding $|\operatorname{supp}_e(g)|$, where $\operatorname{supp}_e(g)$ consists of all sets in $\operatorname{supp}(g)$ of size e.

Lemma 7.5. *There is a constant N_d, depending only on d, such that for every $e \leq d$,*
$$|\operatorname{supp}_e(g)| = O(1 + \epsilon^{1/N_d}/p^e).$$

Proof. Theorem 2.1 shows that any degree d function depending on more than M_d coordinates cannot be Boolean. The union of $N_d = \binom{M_d}{\leq d} + 1$ distinct sets from $\operatorname{supp}_e(g)$ cannot have size at most M_d, and so if $A_1, \ldots, A_N \in \operatorname{supp}(g)$ are distinct then $g|_{A_1 \cup \cdots \cup A_N}$ cannot be Boolean.

If $|\operatorname{supp}_e(g)| < N_d$ then we are done, so suppose that $|\operatorname{supp}_e(g)| \geq N_d$. Then there are $\Omega(|\operatorname{supp}_e(g)|^{N_d})$ subsets of $\operatorname{supp}_e(g)$ of size N_d. Let $\mathcal{F}$ consist of all unions of N_d distinct sets from $\operatorname{supp}_e(g)$. Any such has size at most eN_d and so can be written as a union of N_d sets in at most $2^{eN_d^2}$ ways. Consequently, $|\mathcal{F}| = \Omega(|\operatorname{supp}_e(g)|^{N_d})$.

By construction, $\mathcal{F} \subseteq \mathcal{G} = \operatorname{supp}_e(g)^{\sqcup N_d}$. Since g is a $(d, O(1/p))$-sparse junta, Lemma 3.5 shows that $\mathcal{G}$ is $(eN_d, O(1/p))$-sparse. Applying the reverse union bound (Lemma 3.7) with $p := 2p$, this shows that
$$\Pr_{\mu_{2p}}[\mathcal{E}] \geq |\mathcal{F}|(2p)^{eN_d},$$

where $\mathcal{E}$ is the event that for some $S \in \mathcal{F}$ we have $y_S = 1$ and $y_R = 0$ for all $R \in \mathcal{G}$ such that $R \not\subseteq S$. Suppose that this event happens for S. Denoting $y = 1_T$, we have $g|_T = g|_S$, and so $g|_T$ is not Boolean. Furthermore, since $g|_T$ depends on $|S| \le eN_d$ coordinates, we have $\Pr_{\mu_{1/2}}[g \notin \{0,1\}] \ge 2^{-eN_d}$. This shows that

$$\Pr_{\mu_p}[g \notin \{0,1\}] \ge \Omega(|\mathcal{F}|p^{eN}) = \Omega(|\operatorname{supp}_e(g)|^N p^{eN}).$$

Since $\Pr[g \notin \{0,1\}] \le \Pr[g \ne f] = O(\epsilon)$, we conclude that

$$|\operatorname{supp}_e(g)| = O(\epsilon^{1/N_d}/p). \qquad \square$$

We can now prove Item b of Theorem 7.2.

Proof of Item b. Lemma 7.5 shows that for some constant K_d depending only on d, we have $|\operatorname{supp}_e(g)| \le K_d(1 + \epsilon^{1/N_d}/p^e)$. Write

$$g = \underbrace{\sum_{e:\, \epsilon^{1/N_d} \le p_e} g^{=e}}_{g_J} + \underbrace{\sum_{e:\, \epsilon^{1/N_d} > p_e} g^{=e}}_{h},$$

where $g^{=e}$ is obtained from the monomial expansion of g by retaining only monomials of degree e.

By construction, $|\operatorname{supp}_e(g)| \le 2K_d$ for all e in the first sum, and so g_J is an R_d-junta, where $R_d = 2(d+1)K_d$. Similarly, $|\operatorname{supp}_e(g)| \le 2K_d\epsilon^{1/N_d}/p^e$ for all e in the first sum, and so

$$\Pr[g \ne g_J] = \Pr[h \ne 0] \le \sum_{e:\, \epsilon^{1/N_d} > p_e} p^e \cdot O(\epsilon^{1/N_d}/p^e) = O(\epsilon^{1/N_d}).$$

It follows that $\Pr[f \ne g_J] \le \Pr[f \ne g] + \Pr[g \ne g_J] = O(\epsilon^{1/N_d})$.

If g_J is Boolean then we are done, assuming $\mu_d \ge N_d$. Otherwise, since g_J is an R_d-junta, $\Pr[g_J \ne f] \ge \Pr[g_J \notin \{0,1\}] \ge p^{R_d}$. Since $\Pr[g_J \ne f] = O(\epsilon^{1/N_d})$, we conclude that $p = O(\epsilon^{1/R_dN_d})$. Applying Item a, this shows that f is $O(\epsilon^{1/\mu_d})$-close to the constant zero function, where $\mu_d = \max(m_d, R_dN_d)$. $\qquad \square$

7.3 Proof of Lemma 7.3

In order to complete the proof of Theorem 7.2, we prove Lemma 7.3.

If $A_1, \ldots, A_m \in \mathcal{M}_e$ then we can write $g(A_1 \cup \cdots \cup A_m)$ as the sum of all monomial coefficients supported by $A_1 \cup \cdots \cup A_m$. Taking as an example the case $m = 2$, this corresponds to taking e_1 elements B_1 from A_1 and e_2 elements B_2 from A_2, where $e_1 + e_2 \le d$, and summing over all $\tilde{g}(B_1 \cup B_2)$ obtained in this way. If A_1 and A_2 are not disjoint then we might be summing the same monomial coefficient several times, and so we would like A_1, A_2, and more generally $A_1, \ldots, A_m$, to be disjoint for "many" tuples $A_1, \ldots, A_m$. However, this need not be the case: for example, it might be that all $A \in \mathcal{M}_e$ contain some element a. In the first step of the proof, we extract such a "core" inductively.

Lemma 7.6. *Let e, m be parameters. Let $\mathcal{S}$ be a collection of sets of size e. There exist constants $C_{e,m}, p_{e,m} > 0$ such that the following holds.*

Either (i) $|\mathcal{S}| \le C_{e,m}$, or (ii) there is a set I such that with probability at least $p_{e,m}$, if we choose $A_1, \ldots, A_m \in \mathcal{S}$ at random then $A_1, \ldots, A_m \supseteq I$ and $A_1 \setminus I, \ldots, A_m \setminus I$ are disjoint.

Proof. The proof is by induction on e. If $e = 0$ then case (i) clearly holds for $C_{e,m} = 1$, so assume that $e \ge 1$.

Suppose first that if we choose $A_1, \ldots, A_m \in \mathcal{S}$ at random, the probability that $A_1, \ldots, A_m$ are disjoint is at least $1/2$. If this happens then case (ii) holds for $I = \emptyset$ assuming $p_{e,m} \le 1/2$.

Suppose next that if we choose $A_1, \ldots, A_m \in \mathcal{S}$ at random, the probability that $A_1, \ldots, A_m$ are not all disjoint is at least $1/2$. If we now choose $i, j \in [m]$ such that $i \ne j$ then $A_i \cap A_j \notin \emptyset$ with probability at least $1/\binom{m}{2}$.

31

We can find a fixed choice $i = I$ for which the above holds. Thus if we choose $A \in \mathcal{S}$ at random, $A \cap A_I \neq \emptyset$ with probability at least $1/2\binom{m}{2}$. Consequently, one of the elements $a \in A$ is contained in a $q_{e,m} = 1/2e\binom{m}{2}$ fraction of the sets in $\mathcal{S}$.

Let $\mathcal{T}$ consist of $A \setminus \{a\}$ for all $A \in \mathcal{S}$ containing a, and apply the induction hypothesis. Thus either (i') $|\mathcal{T}| \leq C_{e-1,m}$, or (ii') there exists J such that if we choose $A_1 \setminus \{a\}, \ldots, A_m \setminus \{a\}$ at random from $\mathcal{T}$ then with probability at least $p_{e-1,m}$, all of these sets contain J, and are disjoint otherwise.

If case (i') holds then $|\mathcal{S}| \leq |\mathcal{T}|/q_{e,m} \leq C_{e-1,m}/q_{e,m}$, and so case (i) holds for $C_{e,m} = C_{e-1,m}/q_{e,m}$. If case (ii') holds then case (ii) holds for $I = J \cup \{a\}$ and $p_{e,m} = \min(1/2, q_{e,m}^m p_{e-1,m})$. $\qquad\square$

Let R be a constant to be chosen later, and apply Lemma 7.6 with e, $m = R$, and $\mathcal{S} = \mathcal{M}_e$. If $|\mathcal{M}_e| \leq C_{e,R}$ then case (i) holds, so suppose that there exists I such that if we choose $A_1, \ldots, A_R \in \mathcal{M}_e$ at random then with probability at least $p_{e,R}$, all of $A_1, \ldots, A_R$ contain I, and $A_1 \setminus I, \ldots, A_R \setminus I$ are disjoint. If $|I| = e$ then the probability that all of $A_1, \ldots, A_R$ contain I is $1/|\mathcal{M}_e|^R$, and so $|\mathcal{M}_e| \leq (1/p_{e,R})^{1/R}$, and again case (i) holds. We will show that if $|I| < e$ then case (ii) holds. In this case, since I is strictly contained in some element of $\mathcal{M}_e$, the definition of $\mathcal{M}_e$ shows that $g(1_I) = 0$.

Suppose that $A_1, \ldots, A_R$ satisfy the constraints above, and let $A_1', \ldots, A_m'$ be any $m \geq 1$ distinct elements chosen from $A_1, \ldots, A_R$. Then

$$g(A_1' \cup \cdots \cup A_m') = \sum_{\ell=0}^{d} \sum_{1 \leq i_1 < \cdots < i_\ell \leq m} \sum_{\substack{J \subseteq I \\ B_{i_1} \subseteq A_{i_1}' \setminus I, \ldots, B_{i_\ell} \subseteq A_{i_\ell}' \setminus I \\ B_{i_1}, \ldots, B_{i_\ell} \neq \emptyset \\ |B_{i_1}| + \cdots + |B_{i_\ell}| + |J| \leq d}} \tilde{g}(B_{i_1} \cup \cdots \cup B_{i_\ell} \cup J).$$

This suggests coloring all tuples of up to d elements with the value of the inner sum. Each such sum is the sum of $O(1)$ monomial coefficients of g, which is B-quantized, and therefore there are C_d colors, for some constant C_d depending only for d. The hypergraph Ramsey theorem now implies that for an appropriate value of R depending only on d, there is a subset $A' \subseteq \{A_1, \ldots, A_R\}$ of size m_d such that all tuples of $\ell \leq d$ elements from A' have the same color $\gamma_\ell(A_1, \ldots, A_R)$.

Let $\gamma_0, \ldots, \gamma_\ell$ be a random choice of colors. If we choose $A_1, \ldots, A_R \in \mathcal{M}_e$ at random, then with probability at least $p_{e,R}$, all of them contain I, and $A_1 \setminus I, \ldots, A_R \setminus I$ are disjoint. If this happens, then choose $A' \subseteq \{A_1, \ldots, A_R\}$ of size m_d at random. With probability at least $1/\binom{R}{m_d}$, there are $\gamma_1', \ldots, \gamma_\ell'$ such that all tuples of $\ell \leq d$ elements from A' have color γ_ℓ'. Furthermore, $\gamma_\ell' = \gamma_\ell$ for all ℓ with probability $1/C_d^{d+1}$.

Therefore there is a choice of $\gamma_0, \ldots, \gamma_\ell$ such that if we choose $A_1, \ldots, A_{m_d} \in \mathcal{M}_e$ at random then with probability at least $p_{e,R}/\binom{R}{m_d}C_d^{d+1}$, all tuples of $\ell \leq d$ elements from $A' = \{A_1, \ldots, A_{m_d}\}$ have color γ_ℓ. In particular, for all $m \leq m_d$ such that $m \geq 1$,

$$g(A_1 \cup \cdots \cup A_m) = \sum_{\ell \leq d} \binom{m}{\ell} \gamma_\ell.$$

We call such $A_1, \ldots, A_m$ *good*. The number of good tuples is thus $\Omega(|\mathcal{M}_e|^m)$.

Define the polynomial P by

$$P(\mu) = \sum_{\ell=0}^{d} \binom{\mu}{\ell} \gamma_\ell.$$

This is a polynomial of degree d such that $P(0) = g(1_I) = 0$ and $P(1) = g(1_{A_1}) \neq 0$, since $A_1 \in \mathcal{M}_e$. The union $A_1 \cup \cdots \cup A_m$ has size $s = |I| + m(e - |I|) \leq em$.

Let $\mathcal{F}$ be the collection of all sets of the form $A_1 \cup \cdots \cup A_m$ for good $A_1, \ldots, A_m$. For any $S \in \mathcal{F}$, if $A_1 \cup \cdots \cup A_m = S$ then each A_i is a subset of S, and so at most $2^{sm} \leq 2^{em^2}$ good tuples correspond to the same S. Since there are $\Omega(|\mathcal{M}_e|^m)$ good tuples, it follows that $|\mathcal{F}| = \Omega(|\mathcal{M}_e|^m)$.

References

[CHS20] John Chiarelli, Pooya Hatami, and Michael E. Saks. *An asymptotically tight bound on the number of relevant variables in a bounded degree Boolean function*. Combinatorica, 40(2):237–244, 2020. arXiv:1801.08564. 8

[DFH17] Irit Dinur, Yuval Filmus, and Prahladh Harsha. *Agreement tests on graphs and hypergraphs*, 2017. (manuscript). arXiv:1711.09426, eccc:2017/TR17-181. 5, 6, 18

[DFH19] ———. *Analyzing Boolean functions on the biased hypercube via higher-dimensional agreement tests*. In Timothy M. Chan, ed., *Proc. 30th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA)*, pages 2124–2133. 2019. arXiv:1711.09428,1711.09426, eccc:2017/TR17-180,2017/TR17-181. 6, 28

[Fil16] Yuval Filmus. *Friedgut-Kalai-Naor theorem for slices of the Boolean cube*. Chic. J. Theoret. Comput. Sci., 2016(14), 2016. arXiv:1410.7834. 1, 2, 3, 9

[Fil21] ———. *Boolean functions on S_n which are nearly linear*. Discrete Analysis, 25:1–27, 2021. arXiv:2107.07833. 2, 6

[Fil23] ———. *Junta threshold for low degree Boolean functions on the slice*. Electron. J. Comb., 30(1):P1.55, 2023. arXiv:2203.04760. 6

[FKN02] Ehud Friedgut, Gil Kalai, and Assaf Naor. *Boolean functions whose Fourier transform is concentrated on the first two levels and neutral social choice*. Adv. Appl. Math., 29(3):427–437, 2002. 1

[GR97] Joachim von zur Gathen and James R. Roche. *Polynomials with two values*. Combinatorica, 17(3):345–362, 1997. 4

[IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. *New direct-product testers and 2-query PCPs*. SIAM J. Comput., 41(6):1722–1768, 2012. (Preliminary version in *41st STOC*, 2009). eccc:2009/TR09-090. 18

[Kin03] Guy Kindler. *Property Testing, PCP and Juntas*. Ph.D. thesis, Tel Aviv University, 2003. eccc:2016/TR16-007. 1, 8

[KK20] Nathan Keller and Ohad Klein. *A structure theorem for almost low-degree functions on the slice*. Israel J. Math., 240(1):179–220, 2020. arXiv:1901.08839. 1, 8

[KMS23] Subhash Khot, Dor Minzer, and Muli Safra. *Pseudorandom sets in Grassmann graph have near-perfect expansion*. To appear in Ann. of Math., 2023. (Preliminary version in *59th FOCS*, 2018). eccc:2018/TR18-006. 1

[KS02] Guy Kindler and Shmuel Safra. *Noise-resistant Boolean functions are juntas*, 2002. (manuscript). 1, 8

[NS94] Noam Nisan and Mario Szegedy. *On the degree of Boolean functions as real polynomials*. Comput. Complexity, 4:301–313, 1994. (Preliminary version in *24th STOC*, 1992). 1, 8

[Wel22] Jake Wellens. *A tighter bound on the number of relevant variables in a bounded degree Boolean function*, 2022. arXiv:1903.08214. 8