Oligarchy testing

Yuval Filmus, Noam Lifshitz, Dor Minzer, Elchanan Mossel

May 24, 2020

1 Introduction

1.1 Linearity testing

Let $f\colon \mathbb{Z}_2^n \to \mathbb{Z}_2$ satisfy the identity

$$f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y})$$

for all $\mathbf{x},\mathbf{y} \in \mathbb{Z}_2^n$, where $\mathbf{x}+\mathbf{y}$ is elementwise addition. This identity states that f is linear, that is, of the form

$$f(\mathbf{x}) = \sum_{i \in S} x_i$$

for some subset $S \subseteq [n]$. This is easy to see directly:

$$f(\mathbf{x}) = \sum_{i=1}^n x_i f(\mathbf{e_i}),$$

where e_i is the i'th basis vector.

What can we say if we only know that the identity holds most of the time? That is, suppose that

$$\Pr_{\mathbf{x},\mathbf{y}\sim\mathbb{Z}_2^n}[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y})] = 1-\epsilon,$$

where $\mathbf{x},\mathbf{y}$ are chosen according to the uniform distribution. What can we say about f? A famous result in theoretical computer science states that in this case, there is a linear function ℓ such that

$$\Pr[f = \ell] = 1 - O(\epsilon).$$

1.2 Arrow's theorem

Consider a ranked ballot among three candidates A, B, C. That is, each voter has to submit a ranking of the three candidates. We encode this ranking in the following way. Does the voter prefer A to B? B to C? C to A? We thus associate with each voter a triplet of Boolean values. For these values to be consistent with some ranking, they need to be not all equal (NAE).

Suppose that we want to aggregate the votes by aggregating each of the issues separately. That is, we apply some function f only on the A vs B votes, to come up with a consensus value for this issue. Similarly, we apply f to decide whether B is preferred over C, and whether C is preferred over A. We ask for f to be symmetric: $f(\mathbf{1}-\mathbf{x}) = \mathbf{1} - f(\mathbf{x})$. In order for the results to make sense, we want them to correspond to some ranking, that is, to satisfy the NAE property.

We are thus led to the following question. Which symmetric functions $f\colon \{0,1\}^n \to \{0,1\}$ satisfy the following: whenever $(x_1,y_1,z_1),\ldots,(x_n,y_n,z_n) \in \mathsf{NAE}$, then $(f(\mathbf{x}),f(\mathbf{y}),f(\mathbf{z})) \in \mathsf{NAE}$? Arrow [Arr50] showed that the only admissible functions are dictators, that is, $f(\mathbf{x}) = x_i$ and $f(\mathbf{x}) = 1 - x_i$!

Kalai [Kal02] proved that even if f satisfies this criterion only with probability $1-\epsilon$ over random votes, then f is $O(\epsilon)$-close to a dictator.

1.3 Doctrinal paradox

Arrow's paradox is only one of many paradoxes in social choice theory, more specifically, in judgment aggregation theory. Another one is the doctrinal paradox. Let $f\colon \mathbb{Z}_2^n \to \mathbb{Z}_2$ satisfy the identity

$$f(\mathbf{x} \cdot \mathbf{y}) = f(\mathbf{x}) \cdot f(\mathbf{y})$$

for all $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$, where $\mathbf{x} \cdot \mathbf{y}$ denotes elementwise multiplication. Note first that $f(\mathbf{x} \cdot \mathbf{y}) \leq f(\mathbf{x})$, and this easily shows that f is monotone. If $f \neq 0$, let $\mathbf{y}$ be a minterm of f. If $\mathbf{x} \not\geq \mathbf{y}$ then

$$f(\mathbf{x}) = f(\mathbf{x}) \cdot f(\mathbf{y}) = f(\mathbf{x} \cdot \mathbf{y}) = 0,$$

since $\mathbf{x} \cdot \mathbf{y} < \mathbf{y}$. Conversely, if $\mathbf{x} \geq \mathbf{y}$ then

$$f(\mathbf{x}) = f(\mathbf{x}) \cdot f(\mathbf{y}) = f(\mathbf{x} \cdot \mathbf{y}) = f(\mathbf{y}) = 1.$$

We conclude that either $f = 0$, or

$$f(\mathbf{x}) = \prod_{i \in S} x_i$$

for some subset $S \subseteq [n]$. This is called the *doctrinal paradox* due to a judgement aggregation interpretation. The solutions are known as *oligarchies*, but we will often call them ANDs instead.

Ilan Nehama [Neh13] asked what happens if we only know the following:

$$\Pr_{\mathbf{x},\mathbf{y}}[f(\mathbf{x} \cdot \mathbf{y}) = f(\mathbf{x}) \cdot f(\mathbf{y})] \geq 1 - \epsilon.$$

He showed that f is $16(n\epsilon)^{1/3}$-close to an oligarchy. Our goal is to prove a similar result, but without any dependence of n.

2 Approaches that fail

Perhaps one of the approaches used for linearity testing would generalize? How about Kalai's proof of the robust Arrow's theorem?

Let us briefly go over these proofs, one by one.

Linearity testing via self-correction The first proof of linearity testing, due to Blum et al. [BLR93], has as its starting point the following observation:

$$f(\mathbf{x} + \mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y}) \iff f(\mathbf{x}) = f(\mathbf{y}) + f(\mathbf{x} + \mathbf{y}).$$

In our case, this equation only holds with probability $1 - \epsilon$. However, a good "guess" for $f(\mathbf{x})$ would be the majority value of $f(\mathbf{y}) + f(\mathbf{x} + \mathbf{y})$. Amazingly, if ϵ is small enough (smaller than some constant), then the majority value function is itself linear, and $O(\epsilon)$-close to f!

In our case, the corresponding observation would be the nonsensical

$$f(\mathbf{x} \cdot \mathbf{y}) = f(\mathbf{x}) \cdot f(\mathbf{y}) \iff f(\mathbf{x}) = f(\mathbf{x} \cdot \mathbf{y}) / f(\mathbf{y}).$$

The problem is that we cannot divide by zero!

Linearity testing via Fourier analysis Bellare et al. [BCH+96] rephrased the defining equation in a different way. They replaced $\mathbb{Z}_2$ with $\{-1, 1\}$, transforming the premise to the following form:

$$\Pr_{\mathbf{x},\mathbf{y}}[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) \cdot f(\mathbf{y})] \geq 1-\epsilon,$$

which is the same as

$$\mathop{\mathbb{E}}_{\mathbf{x},\mathbf{y}}[f(\mathbf{x})f(\mathbf{y})f(\mathbf{x}+\mathbf{y})] \geq 1-\epsilon.$$

The left-hand side has an expression in terms of the *Fourier coefficients* of f:

$$1-\epsilon \leq \sum_S \hat{f}(S)^3.$$

Since $\sum_S \hat{f}(S)^2 = 1$, we see that some Fourier coefficient must be close to 1, and so f has correlation close to 1 with the corresponding character, implying that f is close to that character.

We can try to do the same in our case, but unfortunately the formula for $\mathbb{E}[f(\mathbf{x})f(\mathbf{y})f(\mathbf{x}\cdot\mathbf{y})]$ is not as nice. Moreover, whereas linear functions happen to coincide with Fourier characters, the same is unfortunately not true for oligarchies.

Linearity testing via induction David et al. [DDG+17] came up with a new proof of linearity testing, using induction. Here is their idea in a nutshell. Suppose that

$$\Pr_{\mathbf{x},\mathbf{y}}[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y})] \geq 1-\epsilon.$$

We can find a setting of x_n, y_n for which

$$\Pr_{\mathbf{x}_{-n},\mathbf{y}_{-n}}[f(\mathbf{x}_{-n}+\mathbf{y}_{-n}, x_n+y_n) = f(\mathbf{x}_{-n}, x_n) + f(\mathbf{y}_{-n}, y_n)] \geq 1-\epsilon.$$

Now, let us define a new function $g\colon \mathbb{Z}_2^{n-1} \to \mathbb{Z}_2$:

$$g(\mathbf{z}) = f(\mathbf{z}, x_n) + f(\mathbf{z}, y_n) + f(\mathbf{z}, x_n+y_n) = f(\mathbf{z}, x_n \cdot y_n).$$

It turns out that

$$\Pr_{\mathbf{z},\mathbf{w}}[g(\mathbf{z}+\mathbf{w}) = g(\mathbf{z}) + g(\mathbf{w})] \geq 1-O(\epsilon),$$

and so g is $O(\epsilon)$-close to a linear function (by induction). From here, one can deduce that f is $O(\epsilon)$-close to a linear function.

Once again, the construction of g from f relies on the fact that $+$ is a group operation.

Kalai's proof Let $f\colon \{-1, 1\}^n \to \{-1, 1\}$ be the function used to aggregate votes on the relative ranking of two candidates. Kalai [Kal02] calculated the probability that random votes produce a valid ranking: it is

$$\frac{3}{4} - \frac{3}{4}\sum_S (-1/3)^{|S|} \hat{f}(S)^2.$$

Since $\sum_S \hat{f}(S)^2 = 1$ and $\hat{f}(\emptyset) = 0$, this can only be 1 if all of the Fourier mass of f lies on Fourier coefficients of size 1, which easily implies that f is a dictator, thus recovering Arrow's theorem. Similarly, if this probability is $1-\epsilon$ then almost all of the Fourier mass lies on the first level, and so f is $O(\epsilon)$-close to a dictator by the FKN theorem [FKN02] (proved simultaneously).

As already commented above, the expression we get in our case is not as nice.

3 Reduction to an eigenfunction problem

Suppose that we know that

$$\Pr_{\mathbf{x},\mathbf{y}}[f(\mathbf{x}\cdot\mathbf{y}) = f(\mathbf{x})\cdot f(\mathbf{y})] \geq 1-\epsilon.$$

It will be slightly more convenient to write this as follows:

$$\mathop{\mathbb{E}}_{\mathbf{x},\mathbf{y}}[|f(\mathbf{x}\cdot\mathbf{y}) - f(\mathbf{x})\cdot f(\mathbf{y})|] \leq \epsilon.$$

Applying the triangle inequality to move the absolute value outside of $\mathbb{E}_{\mathbf{y}}$, we get

$$\mathop{\mathbb{E}}_{x}\Big|\mathop{\mathbb{E}}_{y}[f(\mathbf{x}\cdot\mathbf{y})] - \mathop{\mathbb{E}}_{y}[f(\mathbf{x})\cdot f(\mathbf{y})]\Big| \leq \epsilon.$$

Let us try to make sense of the two inner expectations. On the right-hand side we have

$$\mathop{\mathbb{E}}_{\mathbf{y}}[f(\mathbf{x})f(\mathbf{y})] = f(\mathbf{x})\cdot\mathop{\mathbb{E}}_{\mathbf{y}}[f(\mathbf{y})].$$

On the left-hand side, we have a kind of averaging of f over all vectors below $\mathbf{x}$:

$$\mathop{\mathbb{E}}_{\mathbf{y}}[f(\mathbf{x}\cdot\mathbf{y})] = \mathop{\mathbb{E}}_{\mathbf{z}\leq\mathbf{x}}[f(\mathbf{z})].$$

Let us denote the right-hand side of the last equation by $T_\downarrow f$; we call $T_\downarrow$ a *one-sided noise operator*, to differentiate it from the more usual operator T_ρ, which is two-sided (can change zeros to ones and vice versa). Then altogether, denoting $\mu = \mathbb{E}[f]$, we get

$$\|T_\downarrow f - \mu f\|_1 \leq \epsilon.$$

In other words, f is an approximate eigenfunction of $T_\downarrow$!

Two-sided noise Let's imagine for a moment replacing $T_\downarrow$ with T_ρ, the standard two-sided noise operator given by

$$T_\rho f = \sum_S \rho^{|S|}\hat{f}(S)\chi_S.$$

Suppose that we are given a function $f\colon\{0,1\}^n \to \{0,1\}$ such that

$$\|T_\rho f - \mu f\|_1 \leq \epsilon.$$

Since $\|T_\rho f - f\|_\infty \leq 1$, this implies

$$\|T_\rho f - \mu f\|_2^2 \leq \epsilon.$$

Applying Parseval's identity,

$$\sum_S (\rho^{|S|} - \mu)^2 \hat{f}(S)^2 \leq \epsilon.$$

This shows that $\mu \approx \rho^d$ for some d, and f is concentrated on the d'th level.

Unfortunately, in the case of the one-sided noise operator we don't have a "Parseval's identity", since $T_\downarrow$ is not normal, and so its eigenfunctions are not orthogonal! Still, it does make sense to compute the exact eigenfunctions of $T_\downarrow$.

Exact eigenfunctions Let $\mathsf{AND}_S = \prod_{i\in S} x_i$. It is not hard to check that

$$T_\downarrow \mathsf{AND}_S = 2^{-|S|}\mathsf{AND}_S.$$

This shows that the functions AND_S are eigenfunctions of $T_\downarrow$. Since they form a basis, these are in fact all the eigenfunctions. However, they are not orthogonal (as mentioned above). For example, $\langle x_1, x_2\rangle = 1/4$.

4 Towards a junta

One step which does apply to $T_\downarrow$ is the one where we replace L_1 norm with L_2 norm:

$$\|T_\downarrow f - \mu f\|_2^2 \leq \epsilon.$$

Does this tell us anything about f?

Here we make a crucial observation: when computing the L_2 norm, we are feeding inputs of weight roughly $n/2$ to μf and to $T_\downarrow f$. In the latter case, this translates to feeding inputs of weight roughly $n/4$ to f. This hints that we should look at the Fourier expansion of f with respect to $\mu_{1/4}$.

The 1/4-biased Fourier expansion of f is the representation of f in a tensorial basis ω_S which is orthonormal with respect to the $\mu_{1/4}$ measure. It turns out that

$$\omega_S = \prod_{i \in S} \frac{x_i - 1/4}{\sqrt{3}/4}.$$

The $T_\downarrow$ operator has the following effect on this basis:

$$T_\downarrow \frac{x_i - 1/4}{\sqrt{3}/4} = \frac{1}{2}\left(\frac{x_i - 1/4}{\sqrt{3}/4} + \frac{-1/4}{\sqrt{3}/4}\right) = \frac{1}{\sqrt{3}} \frac{x_i - 1/2}{1/2}.$$

In other words,

$$T_\downarrow \omega_i = \frac{1}{\sqrt{3}} \chi_i,$$

where χ_S is the usual Fourier basis. Consequently, if we denote by $\tilde{f}$ the 1/4-biased Fourier expansion of f, then

$$T_\downarrow f = \sum_S \sqrt{3}^{-|S|} \tilde{f}(S) \chi_S.$$

Since f is Boolean, $\sum \tilde{f}(S)^2 \leq 1$. This means that $T_\downarrow f$ is concentrated on the low levels. More precisely,

$$\|f^{>d}\|^2 \leq \frac{2(\|(T_\downarrow f)^{>d}\|^2 + \epsilon)}{\mu} = O\left(\frac{3^{-d/2} + \epsilon}{\mu}\right).$$

We can assume that μ is not too small, since otherwise we can approximate f by zero. Let us therefore ignore μ. So $\|f^{>d}\|^2$ is "small". Since f is Boolean, this means that f is close to a *junta*, by a theorem of Bourgain [Bou02], perfected by Kindler, Kirshner and O'Donnell [KKO18]. In other words, there exists a "small" set J, whose size depends only on ϵ, and a Boolean function $g \colon \{0,1\}^J \to \{0,1\}$, such that

$$\Pr_{\substack{\mathbf{x} \in \{0,1\}^J \\ \mathbf{y} \in \{0,1\}^{\overline{J}}}} [f(\mathbf{x}, \mathbf{y}) \neq g(\mathbf{x})] \text{ is small.}$$

Let us try to use this function g to simplify our situation. The idea is to average the inequality $\|T_\downarrow f - \mu f\|_2^2 \leq \epsilon$ over $\overline{J}$. If we denote by $\mathsf{A}_{1/2} f$ the function from $\{0,1\}^J \to \mathbb{R}$ given by

$$\mathsf{A}_{1/2} f(\mathbf{x}) = \mathop{\mathbb{E}}_{\mathbf{y} \in \{0,1\}^{\overline{J}}} [f(\mathbf{x}, \mathbf{y})],$$

then since $\mathsf{A}_{1/2}$ is contracting,

$$\|\mathsf{A}_{1/2} T_\downarrow f - \mu \mathsf{A}_{1/2} f\|_2^2 \leq \epsilon.$$

We can replace $\mathsf{A}_{1/2} f$ by g using the L_2^2 triangle inequality:

$$\|\mathsf{A}_{1/2} T_\downarrow f - \mu g\|_2^2 \text{ is small.}$$

What about $\mathsf{A}_{1/2}T_\downarrow f$? On an input $\mathbf{x} \in \{0,1\}^J$, it is equal to

$$\mathop{\mathbb{E}}_{\mathbf{y}\in\{0,1\}^{\overline{J}}}[T_\downarrow f(\mathbf{x},\mathbf{y})] = \mathop{\mathbb{E}}_{\substack{\mathbf{z}\leq\mathbf{x}\\ \mathbf{y}\in\{0,1\}^{\overline{J}}\\ \mathbf{w}\leq\mathbf{y}}}[f(\mathbf{z},\mathbf{w})] = T_\downarrow \mathsf{A}_{1/4} f(\mathbf{x}),$$

where $\mathsf{A}_{1/4} f$ is defined just as above, but with respect to the $\mu_{1/4}$ measure. In total, we get

$$\|T_\downarrow \mathsf{A}_{1/4} f - \mu g\|_2^2 \text{ is small.}$$

Note that whereas g is Boolean-valued, we are only guaranteed that $\mathsf{A}_{1/4} f$ is $[0,1]$-valued.

5 The generalized eigenfunction problem

The foregoing suggests studying the following generalized eigenvalues problem. Which functions $f\colon \{0,1\}^m \to [0,1]$ and $g\colon \{0,1\}^m \to \{0,1\}$ satisfy the following equation?

$$T_\downarrow f = \lambda g.$$

We are actually interested in approximate solutions of this equation, but it is prudent to first consider the exact version of the problem.

We already know of some solutions: $T_\downarrow \mathsf{AND}_S = 2^{-|S|}\mathsf{AND}_S$. Are there any others? Here are two examples.

Example 1 Suppose that $g(\mathbf{x}) = x_1 \vee x_2$ and $f(\mathbf{x}) = x_1 \oplus x_2$. I claim that $T_\downarrow f = \frac{1}{2} g$. Indeed:

- If $x_1 = 1$ (or $x_2 = 1$) then $g(\mathbf{x}) = 1$, and for $\mathbf{y} \leq \mathbf{x}$, $f(\mathbf{y}) = 1$ with probability exactly $1/2$ (since this is the XOR of one or two random bits).
- If $x_1 = x_2 = 0$ then clearly $g(\mathbf{x}) = f(\mathbf{x}) = 0$.

Example 2 Take $g \equiv 1$ and $f \equiv \lambda$. Clearly $T_\downarrow f = \lambda g$.

General case It turns out that the two examples generate, in a sense, all solutions to the equation $T_\downarrow f = \lambda g$. The general solutions is

$$g = \bigwedge_{i=1}^{s} \bigvee_{j\in S_i} x_j,$$
$$f = 2^s\lambda \bigwedge_{i=1}^{s} \bigoplus_{j\in S_i} x_j,$$

where $2^s\lambda \leq 1$ and the sets $S_1,\ldots,S_s$ are disjoint. A mild generalization of the arguments above shows that indeed $T_\downarrow f = \lambda g$. The other direction is more challenging, and we will only outline it.

The first step is to show that g has to be monotone. Indeed, if $\mathbf{y} \leq \mathbf{x}$ and $g(\mathbf{x}) = 0$ then $f(\mathbf{z}) = 0$ for all $\mathbf{z} \leq \mathbf{x}$, and in particular $f(\mathbf{z}) = 0$ for all $\mathbf{z} \leq \mathbf{y}$, showing that $g(\mathbf{y}) = 0$.

Since g is monotone, it can be expressed as a disjunction of minterms. One shows that all minterms have the same width s, and that the hypergraph of the minterms is a complete s-partite hypergraph, giving the form of g. We deduce the form of f since $T_\downarrow$ is invertible.

Approximate solutions Suppose that we only know that $Tf \approx \lambda g$. If g is not of the form above ("AND-OR") then linear programming duality gives us a "proof" that this cannot happen for f which is $[0,1]$-valued, given that $Tf = \lambda g$. Using this proof, we can rule out such g whenever the approximation is good enough (as a function of m). (A direct argument gives better bounds.)

Having shown that g is an AND-OR, we can conclude that f is approximately equal to the corresponding "AND-XOR" by bounding the norm of $T_\downarrow^{-1}$.

6 Finishing the proof

Let us recap what has happened so far. We start with a function f satisfying

$$\Pr_{\mathbf{x},\mathbf{y}}[f(\mathbf{x}\cdot\mathbf{y}) = f(\mathbf{x})\cdot f(\mathbf{y})] = 1-\epsilon.$$

This implies that f is close to some junta g on some set of variables J, and furthermore

$$T_\downarrow \mathsf{A}_{1/4} f \approx \mu g,$$

where $\mu = \mathbb{E}[f]$. This, in turn, implies that g is close to some AND-OR, and $\mathsf{A}_{1/4}f$ to the corresponding AND-XOR.[1] It remains to show that g is actually an AND.

Suppose that g is an AND of s many ORs. Then $\mu \lesssim 2^{-s}$ (otherwise f won't be bounded in $[0,1]$). On the other hand, since $\mu = \mathbb{E}[f] \approx \mathbb{E}[g]$, we have $\mu \gtrsim 2^{-s}$, with (approximate) equality only if g is an AND. We conclude that g is indeed an AND, and so f is close to an AND.

Other noise rates So far we have considered only the equation $f(\mathbf{x}\cdot\mathbf{y}) = f(\mathbf{x})\cdot f(\mathbf{y})$. What about the following more general equation?

$$f(\mathbf{x}_1\cdots\mathbf{x}_\ell) = f(\mathbf{x}_1)\cdots f(\mathbf{x}_\ell).$$

It is easy to check that the exact solutions are still ANDs. If all we know is that the equation holds with probability $1-\epsilon$, then we can apply the preceding argument, reaching similar conclusions, with one difference: the operator $T_\downarrow$ is replaced by an operator T that zeroes a coordinate with probability $1-1/2^{\ell-1}$ (and $\mathsf{A}_{1/4}f$ is replaced by $\mathsf{A}_{1/2^\ell}f$). The only solutions of $Tf = \lambda g$ are ANDs, simplifying the proof.

Multiple functions Does the picture change if we replace f with three different functions? That is, suppose

$$\Pr_{\mathbf{x},\mathbf{y}}[f(\mathbf{x}\cdot\mathbf{y}) = f_x(\mathbf{x})\cdot f_y(\mathbf{y})] \ge 1-\epsilon.$$

What can we say about f, f_x, f_y? It turns out that we essentially get no new interesting solutions (other than $f = f_x = 0$ or $f = f_y = 0$ and their approximate counterparts).

7 Future work

Our work leaves several interesting open questions for future research.

Quantitative aspects Our work shows that if $\Pr[f(\mathbf{x}\cdot\mathbf{y}) = f(\mathbf{x})\cdot f(\mathbf{y})] \ge 1-\epsilon$ then f is δ-close to an AND, where $\delta \to 0$ as $\epsilon \to 0$. Unfortunately, the dependence of δ on ϵ is super-polynomial. We conjecture that a polynomial dependence is possible (and this is what Ilan Nehama gets).

[1] We're actually cheating here: the parameters don't work out. The argument in the paper is more subtle.

Other functions Linearity testing and "AND testing" both generalize to ϕ-testing, which is the following problem. Given $\phi\colon \{0,1\}^\ell \to \{0,1\}$, which we think of as fixed, characterize all functions $f\colon \{0,1\}^n \to \{0,1\}$ satisfying

$$\Pr_{\mathbf{x}_1,\ldots,\mathbf{x}_\ell}[f(\phi(\mathbf{x}_1,\ldots,\mathbf{x}_\ell)) = \phi(f(\mathbf{x}_1),\ldots,f(\mathbf{x}_\ell))] \geq 1-\epsilon.$$

The case $\epsilon = 0$ was solved by Dokow and Holzman [DH09]: if ϕ is not equivalent to XOR or AND, then the only possible functions f are dictators and, possibly, constants. We conjecture that in the latter case, all approximate solutions are close to exact solutions. We have been able to prove this in some cases, like Majority and Anti-majority.

List-decoding version If $f\colon \mathbb{Z}_2^n \to \mathbb{Z}_2$ is a random function, then the probability that $f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y})$ is 1/2. Hence if

$$\Pr_{\mathbf{x},\mathbf{y}}[f(\mathbf{x}+\mathbf{y}) = f(\mathbf{x}) + f(\mathbf{y})] \geq \frac{1}{2} + \epsilon$$

then f is not random. How does this manifest itself? Bellare et al. [BCH$^+$96] shows that f must have nontrivial correlation with some linear function. Can we say something similar in our case?

References

[Arr50] K. Arrow. A difficulty in the theory of social welfare. *J. of Political Economy*, 58:328–346, 1950.

[BCH$^+$96] Mihir Bellare, Don Coppersmith, Johan Håstad, Marcos A. Kiwi, and Madhu Sudan. Linearity testing in characteristic two. *IEEE Transactions on Information Theory*, 42(6):1781–1795, 1996.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Computer and System Sciences*, 47(3):549–595, 1993.

[Bou02] J. Bourgain. On the distribution of the Fourier spectrum of boolean functions. *Israel Journal of Mathematics*, 131:269–276, 2002.

[DDG$^+$17] Roee David, Irit Dinur, Elazar Goldenberg, Guy Kindler, and Igor Shinkar. Direct sum testing. *SIAM J. Comput.*, 46(4):1336–1369, 2017.

[DH09] Elad Dokow and Ron Holzman. Aggregation of binary evaluations for truth-functional agendas. *Soc. Choice Welf.*, 32:221–241, 2009.

[FKN02] Ehud Friedgut, Gil Kalai, and Assaf Naor. Boolean functions whose fourier transform is concentrated on the first two levels. *Advances in Applied Mathematics*, 29(3):427–437, 2002.

[Kal02] G. Kalai. A Fourier-theoretic perspective on the condorcet paradox and arrow's theorem. *Adv. in Appl. Math.*, 29(3):412–426, 2002.

[KKO18] Guy Kindler, Naomi Kirshner, and Ryan O'Donnell. Gaussian noise sensitivity and fourier tails. *Israel Journal of Mathematics*, 225(1):71–109, Apr 2018.

[Neh13] Ilan Nehama. Approximately classic judgement aggregation. *Annals of Mathematics and Artificial Intelligence*, 68(1-3):91–134, 2013.