Electronic Colloquium on Computational Complexity, Report No. 180 (2026)

A Resolution of Friedgut's Conjecture on Influential Coalitions

Eshan Chattopadhyay*
Cornell University
eshan@cs.cornell.edu

Mohit Gurumukhani*
Cornell University
mgurumuk@cs.cornell.edu

Abstract

We prove that, for every constant $\varepsilon > 0$ and every function $f : \Sigma^n \to \{0,1\}$, there is a coalition of $O(n/\sqrt{\log n})$ coordinates and a target output $b \in \{0,1\}$ such that, after the remaining coordinates are sampled uniformly and independently, the coalition can choose its values to make the output equal to b with probability at least $1 - \varepsilon$. The bound is independent of the alphabet size and also holds for monotone Boolean functions on $[0,1]^n$, resolving a conjecture of Friedgut (Combinatorics, Probability and Computing, 2004). Unlike the Boolean cube setting, where Kahn, Kalai, and Linial (FOCS, 1988) give a coalition bound of $O(n/\log n)$, no sublinear bound independent of the alphabet size was previously known.

In collective coin flipping, our result gives the first sublinear bound on the number of bad players needed to force a fixed output with probability at least $1-\varepsilon$ in any one-round protocol with independent uniform messages, regardless of the message length.

A key ingredient in our proof is an encoding that lets us relate the influence of a function on a product space to the p-biased influence of the encoded function. We then rely on a structure theorem of Hatami (Annals of Mathematics, 2012) for functions with small p-biased influence to bias the encoded function.

1 Introduction

The Kahn–Kalai–Linial (KKL) theorem [KKL88] is a fundamental result in the analysis of Boolean functions. It states that every balanced Boolean function $f : \{0,1\}^n \to \{0,1\}$ has a coordinate with influence at least $\Omega(\log n/n)$, where the influence of a coordinate is the probability that changing its value changes the output. An important consequence is that a small coalition of coordinates can almost determine the output. More precisely, for every constant $\varepsilon > 0$ and any target output $b \in \{0,1\}$, there is a set of $O(n/\log n)$ coordinates such that, after the remaining coordinates are sampled uniformly and independently, the coalition can choose its values to make the output of f equal to b with probability at least $1 - \varepsilon$.

The same question can be asked for functions $f : \Sigma^n \to \{0,1\}$ over an arbitrary finite alphabet Σ. Does a sublinear coalition still suffice, with a bound independent of $|\Sigma|$? In this setting, the influence of a coordinate is the probability that, after the remaining coordinates are sampled uniformly and independently, its value can be chosen to make the output either 0 or 1. Bourgain, Kahn, Kalai, Katznelson, and Linial (BKKKL) [BKKKL92] extended the KKL theorem to these product spaces. Their theorem states that every balanced function has a coordinate with influence at least $\Omega(\log n/n)$, with no dependence on the alphabet size. More generally, their theorem also applies when each coordinate is uniform over $[0,1]$ and $f : [0,1]^n \to \{0,1\}$.[1]

*This work was supported in part by NSF Award CCF-2514586.

[1] In the continuous setting, we consider Boolean functions on $[0,1]^n$ that are either monotone or Borel measurable, with probabilities taken under Lebesgue measure.

ISSN 1433-8092

However, the consequence for coalitions no longer holds: a small coalition need not be able to force an arbitrary target output with probability close to one. To see this, consider a finite-alphabet version of the example discussed by Friedgut [Fri04]. Set $\Sigma = \{0, 1, \dots, n-1\}$ and define $f : \Sigma^n \to \{0,1\}$ by

$$f(x) = 1 \iff x_i \neq 0 \text{ for every } i \in [n].$$

The probability that $f(x) = 1$ is $(1 - 1/n)^n$, which tends to $1/e$. Every coordinate has influence $(1 - 1/n)^{n-1}$: it can determine the output exactly when all the other coordinates are nonzero. Nevertheless, a coalition of $s = o(n)$ coordinates can make the output equal to 1 with probability at most $(1 - 1/n)^{n-s}$, which still tends to $1/e$. This is because every coordinate outside the coalition must be nonzero, regardless of how the coalition chooses its values. In contrast, a single coordinate can force the output to be 0 by setting its value to 0.

This example illustrates why repeatedly applying BKKKL does not give the same coalition guarantee. For this discussion, it suffices to consider monotone functions (see Claim 2.2). For a monotone Boolean function on $\{0,1\}^n$, fixing a coordinate to 1 gives a function on the remaining coordinates whose expectation is larger by half the influence of that coordinate. Fixing it to 0 instead gives a function whose expectation is smaller by the same amount. We can then apply KKL to the restricted function and repeat, making progress towards the same target output at each step. For a monotone Boolean function over a larger ordered alphabet, fixing a coordinate to its largest value increases the expectation of the restricted function, while fixing it to its smallest value decreases it. The sum of the increase and the decrease still equals the influence, but one may be much smaller than the other. In the example above, fixing a coordinate to a nonzero value gives a function whose expectation exceeds the original expectation by only $(1 - 1/n)^{n-1}/n$, despite the coordinate having constant influence. Thus, BKKKL may give an influential coordinate for which no restriction makes substantial progress towards our chosen target. Choosing whichever target gives the larger change in expectation does not resolve the issue: after restricting a coordinate, the next step may move the expectation in the opposite direction.

The example leaves open whether a small coalition can always force at least one of the two outputs. Friedgut [Fri04] conjectured that this is indeed the case, formulating the question for Boolean functions on $[0,1]^n$. In the finite-alphabet setting, the question asks whether, for every constant $\varepsilon > 0$ and every function $f : \Sigma^n \to \{0,1\}$, there is a set of $o(n)$ coordinates, with a bound independent of $|\Sigma|$, and a target output $b \in \{0,1\}$ such that, after the remaining coordinates are sampled uniformly and independently, the coalition can choose its values to make the output equal to b with probability at least $1 - \varepsilon$. Towards this, Filmus, Hambardzumyan, Hatami, Hatami, and Zuckerman [FHHHZ19] showed that, under any product distribution on $\{0,1\}^n$, a coalition of $O(n \log\log n / \log n)$ coordinates can force some fixed output with probability at least $1 - \varepsilon$, for any constant $\varepsilon > 0$. This allows arbitrary coordinate biases, but encoding larger alphabets as bits introduces a dependence on the message length.

We resolve Friedgut's conjecture, obtaining a coalition of size $O(n/\sqrt{\log n})$ in both the finite-alphabet and continuous settings.

Theorem 1 (Informal version of Theorem 3.1 and Corollary 3.2)**.** *For every constant $\varepsilon > 0$, every finite alphabet Σ, and every function $f : \Sigma^n \to \{0,1\}$, there is a coalition of $O(n/\sqrt{\log n})$ coordinates and a target output $b \in \{0,1\}$ such that, after the remaining coordinates are sampled uniformly and independently, the coalition can choose its values to make the output equal to b with probability at least $1 - \varepsilon$. The implicit constant depends only on ε. The same guarantee holds for monotone functions as well as Borel measurable Boolean functions on $[0,1]^n$.*

Recall that, on the Boolean cube, KKL gives a coalition of size $O(n/\log n)$. In the other direction, Ajtai and Linial [AL93] showed the existence of almost-balanced functions $f : \{0,1\}^n \to \{0,1\}$ that are resilient to coalitions of size $\Omega(n/(\log n)^2)$. More precisely, for every constant $\delta > 0$, there is a constant $c > 0$ such that every coalition of at most $cn/(\log n)^2$ coordinates can change the probability of either output by at most δ. This leaves a gap between the lower bound of $\Omega(n/(\log n)^2)$, which already holds on the Boolean cube, and our upper bound of $O(n/\sqrt{\log n})$.

Connection to collective coin flipping. The study of influential coalitions originated in collective coin flipping, introduced by Ben-Or and Linial [BL85], where players try to agree on a random bit despite adversarial faults. This is a fundamental problem in fault-tolerant distributed computing and has been extensively studied (see the survey of Dodis [Dod06]). In the full information model, players broadcast messages using independent private randomness. An unbounded adversary controls a fixed set of bad players and chooses their messages after observing the good players' messages in the current round. For one-round protocols with uniform messages from Σ, this is exactly the setting discussed above: each coordinate is a player's message, the function $f : \Sigma^n \to \{0, 1\}$ determines the output, and controlling a player allows the adversary to choose that coordinate after observing all coordinates outside the coalition.

Thus, for single-bit messages, KKL shows that $O(n/\log n)$ bad players suffice to force a fixed output with probability at least $1-\varepsilon$, for any constant $\varepsilon > 0$. Our result gives a bound of $O(n/\sqrt{\log n})$ bad players, regardless of the message length.

1.1 Proof Overview

We explain the main ideas in the proof of Theorem 1 for monotone functions $f : [r]^n \to \{0, 1\}$, where $[r] = \{1, \ldots, r\}$. Throughout this overview, we fix a constant error $\varepsilon > 0$. As discussed above, repeatedly choosing an influential coordinate need not make progress towards the same output. Our approach is to decrease the expectation towards 0 for as long as we can make sufficient progress, and then show that a small coalition can force 1 if this process stops.

For a monotone function $h : [r]^n \to \{0, 1\}$, let $\mathbf{J}^0_j(h)$ denote the increase in the probability of output 0 when coordinate j can be chosen after observing the remaining coordinates, which are sampled uniformly and independently. By monotonicity, this is exactly the decrease in expectation obtained by fixing coordinate j to its smallest value. Set $\mathbf{J}^0(h) = \sum_j \mathbf{J}^0_j(h)$. (See Section 2.1 for formal definitions.) Our key result is that, if $\mathbb{E}[h] \geq \varepsilon$ and $\mathbf{J}^0(h) \leq \tau$ for $\tau > 0$, then a coalition of $\exp(O(\tau^2))$ coordinates can force 1 with probability at least $1 - \varepsilon$. The implicit constant depends only on ε. This is formally stated as Lemma 3.3. Thus, small total influence towards 0 gives a small coalition that can force the opposite output.

Before proving the above key result, we show how it directly implies our theorem. Starting with f, we repeatedly fix coordinates to their smallest value to decrease the expectation. Set $\delta = c\sqrt{\log n}/n$ for a sufficiently small constant $c > 0$. As long as the expectation exceeds ε and some coordinate has $\mathbf{J}^0_j \geq \delta$, we fix that coordinate to its smallest value, namely 1. Each step decreases the expectation by at least δ, so there are at most $1/\delta = O(n/\sqrt{\log n})$ such steps. If the expectation falls to at most ε, then the coordinates we fixed form a coalition forcing the output 0 with high probability, as desired. Otherwise, the remaining function has expectation greater than ε and total influence towards 0 at most $n\delta = c\sqrt{\log n}$. Applying Lemma 3.3, we obtain a coalition of size $\exp(O(c^2 \log n))$ that can force the output to be 1 with probability at least $1 - \varepsilon$. For sufficiently small c, this coalition has size at most $\sqrt{n}$. By monotonicity, this coalition biases the output towards 1 even without controlling the previously fixed coordinates. Thus, in either case, we obtain a coalition of size $O(n/\sqrt{\log n})$.

We next explain how to prove Lemma 3.3, the central ingredient in our argument. Our starting point is a theorem of Hatami [Hat12] that gives a small coalition whenever the total p-biased influence is small.

For a Boolean function on independent bits that each equal 1 with probability p, the p-biased influence of a bit is the probability that independently resampling the bit from the same distribution changes the output. The total p-biased influence is the sum of these probabilities over all bits. (See Definition 2.1 for a formal definition.) The consequence of Hatami's theorem that we use states that for a monotone function on independent p-biased bits (where $0 < p \leq 1/2$), with expectation at least ε and total p-biased influence at most τ, a coalition of $\exp(O(\tau^2))$ bits can force the output to be 1 with probability at least $1 - \varepsilon$. The bound depends only on τ and ε, even if there are arbitrarily many bits and p is arbitrarily small. (We refer the reader to Theorem 3.4 for the precise statement.)

This gives a route to our lemma. Recall that we are given a monotone function $h : [r]^n \to \{0, 1\}$ with $\mathbb{E}[h] \geq \varepsilon$ and $\mathbf{J}^0(h) \leq \tau$. We encode each coordinate of h by independent p-biased bits so that the resulting function has total p-biased influence at most $O(\tau)$. We must also approximately preserve the expectation and ensure that a coalition of bits biasing the encoded function gives a coalition of at most as many original coordinates biasing h. Applying Hatami's theorem to the encoded function would then give the desired bound.

To obtain such an encoding, we represent each coordinate by independent bits that equal 1 with probability p. To do this, each bit is assigned a value in $[r]$, and the encoding outputs the largest value assigned to a bit that equals 1 (or 1 if all bits are 0). With suitable repetitions of the assigned values and sufficiently small p, the output approximates the uniform distribution on $[r]$. The important property is that every (non-constant) threshold function over $[r]$ corresponds to an OR function over a subset of the p-biased bits. Since fixing all other coordinates of a monotone h leaves a threshold or a constant function, this lets us bound the total p-biased influence of the encoded function by $2\,\mathbf{J}^0(h)$, up to an arbitrarily small approximation error. This is proved in Lemma 3.5. Now applying Theorem 3.4 gives the desired coalition of bits. These bits belong to at most as many original coordinates, so controlling those coordinates gives a coalition for h with no dependence on r, as desired.

Organization We introduce necessary background in Section 2. In Section 3, we prove our main result that bounds the size of an influential coalition, using a reduction to the biased Boolean cube. We present and analyze this reduction in Section 4.

2 Preliminaries

Notation. All logarithms, unless specified otherwise, have base 2. For $b \in \{0, 1\}$, we write $\overline{b}$ to mean $1 - b$. For a string $x \in \Sigma^n$ and a set $S \subset [n]$, let x_S denote its restriction to the coordinates in S, and let x_{-S} denote its restriction to the remaining coordinates. We use $\mathbf{U}_n^r$ to denote the uniform distribution over $[r]^n$. When r is clear from context, we write $\mathbf{U}_n$. For x, y in either $\{0, 1\}^n$ or $[r]^n$ or $[0, 1]^n$, we write $x \leq y$ if $x_i \leq y_i$ for every $i \in [n]$. A Boolean function f on either domain is monotone if $x \leq y$ implies $f(x) \leq f(y)$.

2.1 Influence of Boolean Functions

We use two notions of influence: Influence on the biased Boolean cube and influence towards a fixed output in product spaces.

p-biased influence. Let μ_p^n denote the distribution on $\{0, 1\}^n$ obtained by independently setting each bit to 1 with probability p. When n is clear from context, we will just write μ_p.

We use the following notion of influence with respect to the p-biased distribution, as in Hatami [Hat12].

Definition 2.1 (p-biased influence)**.** *For any $n \in \mathbb{N}$, $p \in [0,1]$, and function $f : \{0,1\}^n \to \{0,1\}$, we define the p-biased influence of variable i on f as:*

$$\mathbf{I}_i^p(f) = \Pr_{x,y \sim \mu_p}\big[f(x_1,\dots,x_{i-1},x_i,x_{i+1},\dots,x_n) \neq f(x_1,\dots,x_{i-1},y_i,x_{i+1},\dots,x_n)\big].$$

With this, we define the total p-biased influence of f as:

$$\mathbf{I}^p(f) = \sum_{i=1}^{n} \mathbf{I}_i^p(f).$$

Influence towards a fixed output. We next define the influence towards a direction. Before that, we need notation for the probability that a coalition can force a given output.

Throughout, let Σ be a finite set, let $n \in \mathbb{N}$, let $f : \Sigma^n \to \{0,1\}$ be a function, and let σ be an arbitrary distribution on Σ. Let σ^n be the product distribution obtained by sampling each coordinate independently from σ.

For $i \in [n]$ and $b \in \{0,1\}$, we define the following quantity:

$$\mathbf{P}_i^b(f,\sigma) = \Pr_{x \sim \sigma^n}\big[\exists j \in \Sigma : f(x_1,\dots,x_{i-1},j,x_{i+1},\dots,x_n) = b\big].$$

We also extend this notion to sets. For $S \subseteq [n]$ and $b \in \{0,1\}$, define

$$\mathbf{P}_S^b(f,\sigma) = \Pr_{x \sim \sigma^n}\big[\exists j \in \Sigma^S : f(j, x_{-S}) = b\big].$$

Often we do not mention the distribution σ, in which case the underlying distribution is uniform over the domain. We use the same definitions for monotone or Borel measurable functions $f : [0,1]^n \to \{0,1\}$, with each coordinate sampled uniformly from $[0,1]$ and coalition assignments chosen from $[0,1]^S$. These probabilities are well defined when f is Borel measurable: coalition-success events are projections of Borel sets, and hence are analytic and Lebesgue measurable.[2] When f is monotone, the coalition can force output b precisely when setting every coordinate in S to b gives output b. The resulting restricted function is monotone and hence Lebesgue measurable.

Using these, define the influence towards direction $b \in \{0,1\}$ of coordinate $i \in [n]$ as:

$$\mathbf{J}_i^b(f,\sigma) = \mathbf{P}_i^b(f,\sigma) - \Pr_{x \sim \sigma^n}[f(x) = b].$$

We finally define $\mathbf{J}^1(f,\sigma) = \sum_i \mathbf{J}_i^1(f,\sigma)$ and $\mathbf{J}^0(f,\sigma) = \sum_i \mathbf{J}_i^0(f,\sigma)$. As above, when we omit σ, the underlying distribution will be the uniform distribution.

2.2 Monotonizing and Discretizing Boolean functions

We present results that let us without loss of generality assume that our given function is monotone and is over finite alphabet.

[2] This is not true for functions f that are merely Lebesgue measurable. Hence, all our statements are stated only for monotone f or Borel measurable f.

Monotonizing Boolean Functions When proving lower bounds on the influence of coalitions, it suffices to consider monotone Boolean functions. This follows from the following monotonization lemma, proved using a sifting argument similar to those in [BL85, KKL88, BKKKL92].

Claim 2.2. *For every Borel measurable function $f : [0,1]^n \to \{0,1\}$, there exists a monotone function $g : [0,1]^n \to \{0,1\}$ such that $\mathbb{E}[g] = \mathbb{E}[f]$ and, for all $S \subseteq [n]$ and $b \in \{0,1\}$, it holds that $\mathbf{P}_S^b(g) \le \mathbf{P}_S^b(f)$. Furthermore, if $f : [r]^n \to \{0,1\}$, then such a monotone function $g : [r]^n \to \{0,1\}$ exists with these properties.*

Proof. This follows by a standard sifting argument. Write $D = [0,1]$ or $D = [r]$, as appropriate, and let μ denote the uniform probability measure on D. For each $i \in [n]$ and assignment $x_{-i} \in D^{n-1}$, let $\ell_{i,x}(t) = f(x_{-i}, t)$. We replace this function by a monotone function $m_{i,x} : D \to \{0,1\}$ with the same expectation, as follows.

If $\mathbb{E}[\ell_{i,x}] = b$ for $b \in \{0,1\}$, then we define $m_{i,x} \equiv b$. Otherwise, when $D = [0,1]$, define $m_{i,x}(t) = 1$ if and only if $t \ge 1 - \mathbb{E}[\ell_{i,x}]$. When $D = [r]$, set $k = |\{t \in [r] : \ell_{i,x}(t) = 1\}|$ and define $m_{i,x}(t) = 1$ if and only if $t > r - k$. Thus, in the discrete case, the ones occupy precisely the largest k elements of $[r]$.

We monotonize f coordinate by coordinate, with the operation in coordinate i replacing every $\ell_{i,x}$ by $m_{i,x}$. Each operation preserves the expectation. It also preserves monotonicity in every previously treated coordinate: if two assignments to the other coordinates are ordered in such a coordinate, then their original functions are pointwise ordered. Their expectations are therefore ordered, and the replacements remain pointwise ordered, including when one of the original functions has expectation 0 or 1. Consequently, the final function g is monotone and satisfies $\mathbb{E}[g] = \mathbb{E}[f]$. We argue each step preserves Borel measurability, so that the operations can be iterated, below.

It remains to show that a single operation, taking f to f' by monotonizing coordinate i, satisfies $\mathbf{P}_S^b(f') \le \mathbf{P}_S^b(f)$ for every $S \subseteq [n]$ and $b \in \{0,1\}$ and that f' is Borel measurable. First, suppose that $i \in S$. The replacement does not introduce an output value that was absent from the original one-variable function. Indeed, if the expectation of this restricted function were $e \in \{0,1\}$, then it must attain the value e over a measure 1 set, and hence changing this to constant e does not introduce any new value; otherwise, the output values available to the functions are $\{0,1\}$ and are left unchanged. Thus, every assignment outside S that allows the coalition to force b under f' also allows it to force b under f.

Now suppose that $i \notin S$. Fix an assignment α to the coordinates in $[n] \setminus (S \cup \{i\})$. For each $\beta \in D^S$, define

$$C_\beta = \{t \in D : f(t,\alpha,\beta) = b\}, \qquad C'_\beta = \{t \in D : f'(t,\alpha,\beta) = b\}.$$

Conditioned on α, the probability that the coalition can force b under f is $\mu(\bigcup_\beta C_\beta)$. The replacement preserves $\mu(C'_\beta) = \mu(C_\beta)$. Moreover, the sets C'_β are all upper intervals when $b = 1$ and all lower intervals when $b = 0$, and hence are nested. For either choice of D, this gives

$$\mu\left(\bigcup_\beta C'_\beta\right) = \sup_\beta \mu(C'_\beta) = \sup_\beta \mu(C_\beta) \le \mu\left(\bigcup_\beta C_\beta\right).$$

Averaging over α, we obtain $\mathbf{P}_S^b(f') \le \mathbf{P}_S^b(f)$. Applying this at every step proves the claim in both cases.

Lastly, we show that f' thus obtained is Borel measurable. First, the map $x_{-i} \mapsto \mathbb{E}[\ell_{i,x}]$ is Borel measurable by the Fubini–Tonelli theorem. Second, $m_{i,x}$ depends on $\ell_{i,x}$ only through $\theta = \mathbb{E}[\ell_{i,x}]$. Letting m_θ be the corresponding function, we see that the map $(\theta, t) \mapsto m_\theta(t)$ is the indicator of a Borel subset of $[0,1] \times D$. As $f'(x) = m_{\mathbb{E}[\ell_{i,x}]}(x_i)$, a composition of these Borel maps, it is Borel measurable as claimed. □

Discretizing Boolean Functions We next show how to discretize a monotone function while approximately preserving its expectation and the probability that any coalition can force either output.

Claim 2.3. *For all $n, r \in \mathbb{N}$, with $r \geq 2$, the following holds. For any monotone function $f : [0,1]^n \to \{0,1\}$, there exists a monotone function $g : [r]^n \to \{0,1\}$ such that $|\mathbb{E}[f] - \mathbb{E}[g]| \leq \frac{n}{r}$, and for any $S \subseteq [n]$ and $b \in \{0,1\}$, we have that $\left|\mathbf{P}^b_S(g) - \mathbf{P}^b_S(f)\right| \leq \frac{n}{r}$.*

Proof. We define g as follows:

$$g(y) = f\left(\frac{y_1 - 1}{r-1}, \ldots, \frac{y_n - 1}{r-1}\right).$$

As f is monotone, we immediately have that g too is monotone. We then proceed by an appropriate hybrid argument, discretizing one coordinate at a time. When discretizing a single coordinate, by monotonicity of f, the change in expectation is at most $\frac{1}{r}$. This can be seen as follows: on fixing all other coordinates, the restricted function is either constant or a threshold. Thus, its probability under the uniform distribution on $\{0, 1/(r-1), \ldots, 1\}$ differs from its probability under the uniform distribution on $[0,1]$ by at most $1/r$. By triangle inequality, we infer that $|\mathbb{E}[f] - \mathbb{E}[g]| \leq \frac{n}{r}$ as desired.

We similarly argue regarding $\mathbf{P}^b_S(g)$ and $\mathbf{P}^b_S(f)$ for any $S \subseteq [n]$ and $b \in \{0,1\}$. First, since f and g are monotone and discretization maps values 1 and r to 0 and 1 respectively, we see that discretizing coordinates of S preserves these quantities, and we assume that those coordinates are fixed to 1 or r for $b = 0$ and $b = 1$ respectively. Then, for other coordinates, we argue as earlier and discretize coordinate by coordinate, observing that the change in probability is at most $\frac{1}{r}$. Applying appropriate triangle inequalities then, we obtain that $\left|\mathbf{P}^b_S(g) - \mathbf{P}^b_S(f)\right| \leq \frac{n-|S|}{r} \leq \frac{n}{r}$, as desired. □

2.3 Basic Probability notions

For distributions $\mathbf{X}, \mathbf{Y}$ on the same finite set Ω, let $|\mathbf{X} - \mathbf{Y}| = \frac{1}{2}\sum_{x \in \Omega} |\Pr[\mathbf{X} = x] - \Pr[\mathbf{Y} = x]|$ denote their statistical distance. Next, let $\|\mathbf{X} - \mathbf{Y}\|_\infty = \max_{w \in \Omega} |\Pr[\mathbf{X} = w] - \Pr[\mathbf{Y} = w]|$ denote the ℓ_∞ distance between them.

We record the following useful inequality.

Fact 2.4 (Data Processing Inequality)**.** *Let Ω, Ω' be arbitrary finite sets. Let $\mathbf{X}, \mathbf{Y} \sim \Omega$ be such that $|\mathbf{X} - \mathbf{Y}| \leq \varepsilon$. Then, for any $f : \Omega \to \Omega'$, we have that $|f(\mathbf{X}) - f(\mathbf{Y})| \leq \varepsilon$.*

From this, we easily obtain the following inequality.

Claim 2.5. *Let $\mathbf{X}, \mathbf{Y} \sim \Omega$ and $\varepsilon > 0$ be such that $|\mathbf{X} - \mathbf{Y}| \leq \varepsilon$. Then, for any $f : \Omega \to \{0,1\}$, we have that*

$$\left|\mathop{\mathbb{E}}_{x \sim \mathbf{X}}[f(x)] - \mathop{\mathbb{E}}_{y \sim \mathbf{Y}}[f(y)]\right| \leq \varepsilon.$$

Proof. Fix any such $f, \mathbf{X}, \mathbf{Y}$. Then,

$$\begin{aligned}\left|\mathop{\mathbb{E}}_{x \sim \mathbf{X}}[f(x)] - \mathop{\mathbb{E}}_{y \sim \mathbf{Y}}[f(y)]\right| &= |f(\mathbf{X}) - f(\mathbf{Y})| \\ &\leq \varepsilon.\end{aligned}$$

where last step follows by Fact 2.4. □

3 A bound on influential coalition size

In this section, we will prove the following main result regarding biasing a function over $[r]^n$.

Theorem 3.1. *For all constant $\varepsilon > 0$, there exists a constant $C = C(\varepsilon) > 0$ such that the following holds. For all $r, n \in \mathbb{N}$ and function $f : [r]^n \to \{0,1\}$, there exists $b \in \{0,1\}$, and $B \subseteq [n]$ with $|B| \leq C \cdot \frac{n}{\sqrt{\log(n)}}$ such that*

$$\mathbf{P}_B^b(f) \geq 1 - \varepsilon.$$

We obtain the same coalition bound for monotone or Borel measurable Boolean functions on $[0,1]^n$.

Corollary 3.2. *For all constant $\varepsilon > 0$, there exists a constant $C = C(\varepsilon) > 0$ such that the following holds. For any $n \in \mathbb{N}$, consider $[0,1]^n$ with the uniform measure. Let $f : [0,1]^n \to \{0,1\}$ be an arbitrary monotone or Borel measurable function. Then, there exists $b \in \{0,1\}$ and $B \subseteq [n]$ with $|B| \leq C \cdot \frac{n}{\sqrt{\log(n)}}$ such that*

$$\mathbf{P}_B^b(f) \geq 1 - \varepsilon.$$

Proof. If f is non-monotone (and so by assumption Borel measurable), then we apply Claim 2.2 and without loss of generality assume that f is a monotone function. Then, we apply Claim 2.3 with $r = \max\{2, \lceil 2n/\varepsilon \rceil\}$ to obtain a function $g : [r]^n \to \{0,1\}$ such that if there exists $B \subseteq [n]$ and $b \in \{0,1\}$ such that $\mathbf{P}_B^b(g) \geq 1 - \frac{\varepsilon}{2}$, then $\mathbf{P}_B^b(f) \geq 1 - \varepsilon$. We finally apply Theorem 3.1 to g with error parameter $\frac{\varepsilon}{2}$ and find such b, B to obtain the desired result. $\square$

To prove Theorem 3.1, we will need the following result that shows that if total influence towards 0 of a function is small, then there exists a small coalition, biasing a function towards 1.

Lemma 3.3. *For all constant $\varepsilon \in (0,1)$, there exists a constant $C = C(\varepsilon) \geq 1$ such that the following holds. For all $n, r \in \mathbb{N}$, $\tau > 0$ and a monotone function $f : [r]^n \to \{0,1\}$ such that $\mathbb{E}_{x \sim \mathbf{U}_n}[f(x)] \geq \varepsilon$ and $\mathbf{J}^0(f) \leq \tau$, there exists $T \subseteq [n]$ with $|T| \leq \exp(C \cdot \tau^2)$ so that*

$$\mathbf{P}_T^1(f) \geq 1 - \varepsilon.$$

We prove this key lemma in Section 3.1. We first show how it implies Theorem 3.1.

Proof of Theorem 3.1. We apply Claim 2.2 to assume without loss of generality that f is monotone. We will first show how to find such b and B, and then analyze and show that it has the claimed properties.

Finding such b and B Let $\delta = \frac{1}{100 \cdot C_{Lemma\ 3.3}} \cdot \frac{\sqrt{\log(n)}}{n}$ where $C_{Lemma\ 3.3} \geq 1$ is the constant in Lemma 3.3 for the parameter ε.

We consider the following algorithm to find such a B.

- Let $Z_0 = \emptyset$ and let $f_0 = f$.
- For $i \geq 0$, proceed as follows.
  - If $\Pr[f_i(x) = 0] \geq 1 - \varepsilon$, then let $b = 0$, $B = Z_i$ and terminate.
  - Otherwise, if there exists $j \in [n] \setminus Z_i$ such that $\mathbf{J}_j^0(f_i) \geq \delta$, then define $Z_{i+1} = Z_i \cup \{j\}$, and $f_{i+1}(x) : [r]^{n-i-1} \to \{0,1\}$ as the function obtained from fixing variable j in f_i to 1, and proceed to the next iteration.
  - Otherwise, if for all $j \in [n] \setminus Z_i$, $\mathbf{J}_j^0(f_i) < \delta$, then set $b = 1$, and apply Lemma 3.3 (with $\tau = \delta n$) to f_i to obtain the desired coalition B.

Analysis We now analyze the above algorithm.

As $\mathbf{J}_j^0(f_i) \geq \delta$, we easily see that $\Pr[f_{i+1} = 0] \geq \Pr[f_i = 0] + \delta$. This shows that our algorithm must terminate in finite steps. If our algorithm terminates by setting $b = 0$, then we will have that $|B| \leq \frac{1}{\delta} = O\left(\frac{n}{\sqrt{\log(n)}}\right)$ as desired.

Otherwise, our algorithm terminates by setting $b = 1$ in iteration i for some $i \geq 0$. In that case, we have that $|B| \leq \exp(C_{\textit{Lemma 3.3}} \cdot (\delta \cdot n)^2) = o(n/\sqrt{\log(n)})$, by our choice of constants in δ. To argue about correctness, by monotonicity of f, we see that

$$\mathbf{P}_B^1(f) = \Pr_x[f(x_{-B}, r^B) = 1] \geq \Pr_x[f(x_{-B-Z_i}, r^B, 1^i) = 1] \geq 1 - \varepsilon.$$

□

3.1 Proving the direction switching lemma

In this subsection, we will prove Lemma 3.3, our result that if total influence of a function towards 0 is small, then small coalition can bias it towards 1.

To prove this, we will need the following two ingredients. First, we will need the following result of Hatami regarding functions with small p-biased influence.

Theorem 3.4 (Simplified version of Corollary 2.10 from [Hat12]). *For all constants $\varepsilon \in (0,1)$, there exists a constant $C = C(\varepsilon) \geq 1$ such that the following holds. For all $n \in \mathbb{N}, 0 < p \leq \frac{1}{2}$, and a monotone function $f : \{0,1\}^n \to \{0,1\}$ with $\mathbb{E}_{x \sim \mu_p}[f(x)] \geq \varepsilon$, there exists $S \subseteq [n]$ with $|S| \leq \exp\left(C \cdot (\mathbf{I}^p(f))^2\right)$ such that*

$$\mathbf{P}_S^1(f, \mu_p) \geq 1 - \varepsilon.$$

Second, we need the following result that relates p-biased influence of a function to bias towards 0 of the original function.

Lemma 3.5. *For $r, n \in \mathbb{N}$, let $f : [r]^n \to \{0,1\}$ be a monotone function. Then, for all $\varepsilon > 0$, $p \in (0, 1/2)$ there exists $m \in \mathbb{N}$ and a function $g : \{0,1\}^m \to \{0,1\}$ with the following properties.*

1. *g is monotone.*
2. *$\left|\mathbb{E}_{y \sim \mu_p}[g(y)] - \mathbb{E}_{x \sim \mathbf{U}_n}[f(x)]\right| \leq prn$.*
3. *$\mathbf{I}^p(g) \leq 2 \cdot \mathbf{J}^0(f) + 4pr \cdot n^2$.*
4. *For all $S \subseteq [m]$, if*
$$\mathbf{P}_S^1(g, \mu_p) \geq 1 - \varepsilon,$$
then there exists $T \subseteq [n]$ with $|T| \leq |S|$ such that
$$\mathbf{P}_T^1(f) \geq 1 - \varepsilon - prn.$$

We will prove Lemma 3.5 in Section 4. We now show how we can combine these two ingredients and prove our direction switching lemma.

Proof of Lemma 3.3. We let

$$p = \min\left(\frac{\varepsilon}{2rn}, \frac{\tau}{2r \cdot n^2}\right)$$

so that $prn \le \frac{\varepsilon}{2}$ and $4prn \cdot^2 \le 2\tau$. With this, we apply Lemma 3.5 to f with parameters $\varepsilon_{\textit{Lemma 3.5}} = \frac{\varepsilon}{2}$ to obtain a function $g : \{0,1\}^m \to \{0,1\}$ for some $m \in \mathbb{N}$ with the following properties:

1. g is monotone.
2. $\mathbb{E}_{y \sim \mu_p}[g(y)] \ge \varepsilon - prn \ge \frac{\varepsilon}{2}$.
3. $\mathbf{I}^p(g) \le 2\tau + 4pr \cdot n^2 \le 4\tau$.
4. For all $S \subseteq [m]$, if $\mathbf{P}^1_S(g, \mu_p) \ge 1 - \frac{\varepsilon}{2}$, then there exists $T \subseteq [n]$ with $|T| \le |S|$ such that
$$\mathbf{P}^1_T(f) \ge 1 - \varepsilon.$$

With this, we apply Theorem 3.4 with parameter $\varepsilon_{\textit{Theorem 3.4}} = \frac{\varepsilon}{2}$ to g to find a $C = C(\varepsilon) \ge 1$ and a coalition $S \subseteq [m]$ with $|S| \le \exp(C \cdot \tau^2)$ such that $\mathbf{P}^1_S(g, \mu_p) \ge 1 - \frac{\varepsilon}{2}$. By property 4 of g above, the desired claim follows. □

4 From functions in product space to p-biased space

We use this section to prove Lemma 3.5. The main idea is to encode each coordinate so that every non-constant threshold of its encoded value is an OR of some of its bits. Since fixing all other coordinates of a monotone function leaves a threshold or a constant function, this will let us bound the p-biased influence of the encoded function by the influence of f towards 0. We first prove this influence bound, then choose the encoding to approximate the uniform distribution, and finally complete the proof of Lemma 3.5 at the end of Section 4.3. Throughout this section, fix a monotone function $f : [r]^n \to \{0,1\}$ and $p \in (0, 1/2)$.

4.1 The encoding and its influence bound

We assign each bit a value in $[r]$ and output the largest value assigned to a bit that equals 1, or 1 if all bits are 0. Formally, we define the following max selector function.

Definition 4.1. *For $\ell \in \mathbb{N}$ and a nondecreasing sequence $a = (a_1, \dots, a_\ell) \in [r]^\ell$, we define $\phi_a : \{0,1\}^\ell \to [r]$ as*

$$\phi_a(y) = \begin{cases} a_{\max\{i \in [\ell] : y_i = 1\}} & y \ne 0^\ell \\ 1 & y = 0^\ell \end{cases}$$

The useful property is that, for every $t > 1$,

$$\phi_a(y) \ge t \quad \Longleftrightarrow \quad \bigvee_{i \in [\ell] :\, a_i \ge t} y_i = 1.$$

Thus, every nonconstant threshold on $[r]$ becomes an OR of some of the encoding bits. This property holds for every choice of a. We will choose a to approximate the uniform distribution in Section 4.2.

For now, fix any nondecreasing sequence $a \in [r]^\ell$. Set $m = \ell \cdot n$ and define $g : \{0,1\}^m \to \{0,1\}$ as

$$g(y) = f(\phi_a(y_1, \dots, y_\ell), \dots, \phi_a(y_{m-\ell+1}, \dots, y_m)). \tag{1}$$

Let $\mathbf{A}$ be the distribution on $[r]$ obtained by sampling $y \sim \mu_p^\ell$ and outputting $\phi_a(y)$. The encoded blocks are independent, so their joint distribution is $\mathbf{A}^n$.

We now prove the influence bound.

Claim 4.2. *We have*

$$\mathbf{I}^p(g) \leq 2 \cdot \mathbf{J}^0(f, \mathbf{A}).$$

Proof. For $j \in [n]$, let $B_j = \{(j-1) \cdot \ell + 1, \ldots, j \cdot \ell\}$. It suffices to show that

$$\sum_{i \in B_j} \mathbf{I}_i^p(g) \leq 2 \cdot \mathbf{J}_j^0(f, \mathbf{A})$$

for every $j \in [n]$. Fix such a j and condition on all blocks except B_j. Let x_{-j} denote their encoded values, and let $h(y_{B_j}) = f(\phi_a(y_{B_j}), x_{-j})$ be the remaining function, where the encoded value is placed in coordinate j.

If fixing x_{-j} makes f constant, then h is constant as well, and both its total p-biased influence and the conditional influence of coordinate j towards 0 are zero.

Otherwise, by monotonicity, there is a value $t^* > 1$ such that $f(z, x_{-j}) = 1$ if and only if $z \geq t^*$. In particular, $f(1, x_{-j}) = 0$. The function h is therefore an OR of $k = |\{i \in [\ell] : a_i \geq t^*\}|$ bits. Each of these k bits changes the output exactly when the other $k-1$ bits are zero and the two independent samplings of the bit (from μ_p) give different values. Consequently,

$$\begin{aligned}
\mathbf{I}^p(h) &= k \cdot 2p(1-p) \cdot (1-p)^{k-1} \\
&= 2(1-p) \Pr[\text{exactly one of these } k \text{ bits is } 1] \\
&\leq 2 \Pr[h = 1] \\
&= 2 \Pr_{x_j \sim \mathbf{A}}[f(x_j, x_{-j}) = 1] \\
&= 2\big(\Pr_{x_j \sim \mathbf{A}}[f(x_j, x_{-j}) = 1] - f(1, x_{-j})\big),
\end{aligned}$$

where the final equality use the fact that x^{-j} does not make f a constant, and so $f(1, x_{-j})$ must be 0.

Recalling that x_{-j} follows the distribution $\mathbf{A}^{n-1}$, we get

$$\sum_{i \in B_j} \mathbf{I}_i^p(g) \leq 2 \mathop{\mathbb{E}}_{x_{-j} \sim \mathbf{A}^{n-1}} \big[\Pr_{x_j \sim \mathbf{A}}[f(x_j, x_{-j}) = 1] - f(1, x_{-j}) \big] = 2\, \mathbf{J}_j^0(f, \mathbf{A}).$$

The claim now follows by summing over $j \in [n]$. □

4.2 Approximating the uniform distribution

We next choose a so that $\mathbf{A}$ approximates the uniform distribution on $[r]$. It is helpful to note that

$$\Pr_{y \sim \mu_p^\ell}[\phi_a(y) \leq i] = (1-p)^{|\{j \in [\ell] : a_j > i\}|}.$$

This just follows by definition of ϕ_a since the encoded value is at most i exactly when every bit assigned a value greater than i is zero. We choose the number of such bits to make this probability close to i/r. Taking differences for consecutive values of i gives the following.

Lemma 4.3. *There exists $\ell \in \mathbb{N}$ and a nondecreasing sequence $a = (a_1, \ldots, a_\ell) \in [r]^\ell$ such that for all $i \in [r]$, we have that*

$$\left| \Pr_{y \sim \mu_p}[\phi_a(y) = i] - \frac{1}{r} \right| \leq 2p$$

Proof. For $i \in [r]$, let $\chi(i)$ be the unique nonnegative integer such that

$$(1-p)^{\chi(i)+1} < \frac{i}{r} \leq (1-p)^{\chi(i)}. \tag{2}$$

Let a be the nondecreasing sequence containing one copy of 1 and, for each $i \geq 2$, $\chi(i-1) - \chi(i)$ copies of i. We observe that $\chi(r) = 0$ and so, for any $i \in [r]$, the number of terms in a that have value $> i$ is $\chi(i)$. Therefore, $\Pr_{y \sim \mu_p}[\phi_a(y) \leq i] = (1-p)^{\chi(i)}$.

Thus, for $i \geq 2$, we have

$$\begin{aligned} \Pr_{y \sim \mu_p}[\phi_a(y) = i] &= \Pr[\phi_a(y) \leq i] - \Pr[\phi_a(y) \leq i-1] \\ &= (1-p)^{\chi(i)} - (1-p)^{\chi(i-1)}. \end{aligned} \tag{3}$$

To bound Equation (3), we will make use of the following claim, which is proved at the end of this subsection.

Claim 4.4. *The following holds:*

$$\frac{i}{r} \leq (1-p)^{\chi(i)} \leq \frac{i}{r}(1+2p).$$

First, consider the case of $i = 1$. Since ϕ_a takes values in $[r]$, we have

$$\Pr_{y \sim \mu_p}[\phi_a(y) = 1] = (1-p)^{\chi(1)}.$$

Applying Claim 4.4, we obtain

$$\frac{1}{r} \leq \Pr_{y \sim \mu_p}[\phi_a(y) = 1] \leq \frac{1}{r}(1+2p) \leq \frac{1}{r} + 2p,$$

yielding the desired bound for $i = 1$.

For the remainder of the proof, fix $i \geq 2$. First, using Equation (3), we compute that

$$\begin{aligned} \Pr_{y \sim \mu_p}[\phi_a(y) = i] &= (1-p)^{\chi(i)} - (1-p)^{\chi(i-1)} \\ &\leq \frac{i}{r}(1+2p) - \frac{i-1}{r} \\ &\leq \frac{1}{r} + 2p \end{aligned}$$

where for the middle inequality, we applied Claim 4.4.

Similarly, again using Equation (3), we have that

$$\begin{aligned} \Pr_{y \sim \mu_p}[\phi_a(y) = i] &= (1-p)^{\chi(i)} - (1-p)^{\chi(i-1)} \\ &\geq \frac{i}{r} - \frac{i-1}{r}(1+2p) \end{aligned}$$

$$\geq \frac{1}{r} - 2p$$

where for the middle inequality, we again applied Claim 4.4.

Hence, we conclude that

$$\left| \Pr_{y \sim \mu_p} [\phi_a(y) = i] - \frac{1}{r} \right| \leq 2p$$

as desired. Lastly, we prove our remaining claim.

Proof of Claim 4.4. By definition of $\chi(i)$ from Equation (2), our lower bound of $(1-p)^{\chi(i)} \geq \frac{i}{r}$ follows.

Similarly, from Equation (2), we compute our upper bound as

$$\begin{aligned} (1-p)^{\chi(i)} &\leq \frac{i}{r} \cdot \frac{1}{1-p} \\ &= \frac{i}{r} \cdot \left(1 + \frac{p}{1-p}\right) \\ &\leq \frac{i}{r} \cdot (1 + 2p)\,, \end{aligned}$$

where the last inequality uses the fact that $p \in (0, 1/2)$. □

□

4.3 Completing the proof

Fix a sequence a that satisfies the properties in Lemma 4.3. We first record the resulting bounds on the distance from uniform.

Lemma 4.5. *The following holds regarding* $\mathbf{A}$.

1. $$\|\mathbf{A} - \mathbf{U}_r\|_\infty \leq 2p.$$

2. *For all* $i \geq 1$,
$$\left|\mathbf{A}^i - \mathbf{U}_r^i\right| \leq pri.$$

Proof of Lemma 4.5. The first part is immediate from Lemma 4.3. For the second part, we proceed by induction on i. For $i = 1$, by definition of statistical distance and the first part,

$$|\mathbf{A} - \mathbf{U}_r| = \frac{1}{2} \sum_{a \in [r]} \left| \Pr[\mathbf{A} = a] - \frac{1}{r} \right| \leq pr.$$

For $i \geq 2$, applying the triangle inequality, we obtain

$$\begin{aligned} \left|\mathbf{A}^i - \mathbf{U}_r^i\right| &\leq \left|\mathbf{A}^{i-1} \times \mathbf{A} - \mathbf{U}_r^{i-1} \times \mathbf{A}\right| + \left|\mathbf{U}_r^{i-1} \times \mathbf{A} - \mathbf{U}_r^{i-1} \times \mathbf{U}_r\right| \\ &= \left|\mathbf{A}^{i-1} - \mathbf{U}_r^{i-1}\right| + |\mathbf{A} - \mathbf{U}_r| \\ &\leq pr(i-1) + pr = pri, \end{aligned}$$

where the equality holds because taking a product with the same distribution preserves statistical distance, and the last inequality follows from the induction hypothesis and the case $i = 1$. □

We next relate the influence of f towards 0 under $\mathbf{A}$ to its influence under the uniform distribution.

Claim 4.6. *It holds that*

$$\mathbf{J}^0(f, \mathbf{A}) \leq \mathbf{J}^0(f) + 2pr \cdot n^2.$$

Proof of Claim 4.6. By monotonicity of f, $\mathbf{P}_i^0(f, \mathbf{A}) = \Pr_{x \sim \mathbf{A}^n}[f(1, x_{-i}) = 0]$, and the same equation holds under the uniform distribution as well. Applying Claim 2.5 and Lemma 4.5 to the indicators of $f(1, x_{-i}) = 0$ and $f(x) = 0$, respectively, we obtain

$$\mathbf{P}_i^0(f, \mathbf{A}) \leq \mathbf{P}_i^0(f) + prn \qquad \text{and} \qquad \Pr_{x \sim \mathbf{A}^n}[f(x) = 0] \geq \Pr_{x \sim \mathbf{U}_n}[f(x) = 0] - prn.$$

Therefore,

$$\begin{aligned}
\mathbf{J}^0(f, \mathbf{A}) &= \sum_i \left(\mathbf{P}_i^0(f, \mathbf{A}) - \Pr_{x \sim \mathbf{A}^n}[f(x) = 0] \right) \\
&\leq \sum_i \left(\mathbf{P}_i^0(f) + prn - \left(\Pr_{x \sim \mathbf{U}_n}[f(x) = 0] - prn \right) \right) \\
&= \mathbf{J}^0(f) + 2pr \cdot n^2,
\end{aligned}$$

as desired. □

With the above bounds, we are ready to prove Lemma 3.5.

Proof of Lemma 3.5. We prove each of these results one after the other.

1. Let $y_1 \leq y_2 \in \{0,1\}^m$, and let $x_1, x_2 \in [r]^n$ be their corresponding encoded values. By definition of ϕ_a, changing an input bit from 0 to 1 cannot decrease the encoded value of its block. Thus, $x_1 \leq x_2$. Since f is monotone, we obtain $g(y_1) = f(x_1) \leq f(x_2) = g(y_2)$, as desired.

2. By definition of g, we have that $\mathbb{E}_{y \sim \mu_p}[g(y)] = \mathbb{E}_{x \sim \mathbf{A}^n}[f(x)]$. Hence, it suffices to show that

$$\left| \mathop{\mathbb{E}}_{x \sim \mathbf{A}^n}[f(x)] - \mathop{\mathbb{E}}_{x \sim \mathbf{U}_n}[f(x)] \right| \leq prn.$$

   Applying the data processing inequality (Claim 2.5), it suffices to show that $|\mathbf{A}^n - \mathbf{U}_n| \leq prn$. This last claim directly follows from Lemma 4.5.

3. This directly follows from Claim 4.6 and Claim 4.2.

4. Let $S \subseteq [m]$ be arbitrary such that $\mathbf{P}_S^1(g, \mu_p) \geq 1 - \varepsilon$. For $i \in [n]$, let $B_i = \{(i-1) \cdot \ell + 1, \ldots, i \cdot \ell\}$. We define our set T as follows: $T = \{i \in [n] : B_i \cap S \neq \emptyset\}$. Then, $|T| \leq |S|$. We claim that

$$\mathbf{P}_T^1(f, \mathbf{A}) = \Pr_{x \sim \mathbf{A}^n}[f(r^T, x_{-T}) = 1] \geq 1 - \varepsilon.$$

   First, let us see how our desired claim follows from this. From Lemma 4.5, we have that $|\mathbf{A}^n - \mathbf{U}_r^n| \leq prn$. Applying Claim 2.5 to the function obtained by fixing the coordinates in T to r, we have that

$$\mathbf{P}_T^1(f) = \Pr_{x \sim \mathbf{U}_r^n}[f(r^T, x_{-T}) = 1] \geq \mathbf{P}_T^1(f, \mathbf{A}) - prn \geq 1 - \varepsilon - prn$$

as desired.

We now prove our remaining claim that $\mathbf{P}^1_T(f, \mathbf{A}) \geq 1-\varepsilon$. Let $S' = \bigcup_{i\in T} B_i$. Setting all bits in each block indexed by T to 1 produces encoded values at most r. By monotonicity of f, fixing the original coordinates in T to r can only increase the output. Thus, $\mathbf{P}^1_T(f, \mathbf{A}) \geq \mathbf{P}^1_{S'}(g, \mu_p)$. By the definitions of T and S', we have $S \subseteq S'$. Therefore,

$$\mathbf{P}^1_T(f, \mathbf{A}) \geq \mathbf{P}^1_{S'}(g, \mu_p) \geq \mathbf{P}^1_S(g, \mu_p) \geq 1-\varepsilon.$$

□

Acknowledgements M.G. would like to thank David Zuckerman for helpful conversations. Part of the work was done while the authors were visiting the Simons Institute for Theory of Computing, UC Berkeley. We would like to thank the Simons Institute for the stimulating research environment and the hospitality.

AI use statement. The authors used ChatGPT Pro extensively for exploring ideas, and proving several technical parts. The authors take full responsibility for the correctness of all results in the paper.

References

[AL93] Miklós Ajtai and Nathan Linial. “The influence of large coalitions”. In: *Comb.* 13.2 (1993), pp. 129–145. DOI: 10.1007/BF01303199 (cit. on p. 2).

[BL85] Michael Ben-Or and Nathan Linial. “Collective Coin Flipping, Robust Voting Schemes and Minima of Banzhaf Values”. In: *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985*. IEEE Computer Society, 1985, pp. 408–416. DOI: 10.1109/SFCS.1985.15 (cit. on pp. 3, 6).

[BKKKL92] Jean Bourgain, Jeff Kahn, Gil Kalai, Yitzhak Katznelson, and Nathan Linial. “The influence of variables in product spaces”. In: *Israel Journal of Mathematics* 77.1 (1992), pp. 55–64 (cit. on pp. 1, 6).

[Dod06] Yevgeniy Dodis. *Fault-tolerant leader election and collective coin-flipping in the full information model*. 2006 (cit. on p. 3).

[FHHHZ19] Yuval Filmus, Lianna Hambardzumyan, Hamed Hatami, Pooya Hatami, and David Zuckerman. “Biasing Boolean Functions and Collective Coin-Flipping Protocols over Arbitrary Product Distributions”. In: *46th International Colloquium on Automata, Languages, and Programming (ICALP)*. Vol. 132. LIPIcs. 2019, 58:1–58:13. DOI: 10.4230/LIPICS.ICALP.2019.58 (cit. on p. 2).

[Fri04] Ehud Friedgut. “Influences in product spaces: KKL and BKKKL revisited”. In: *Combinatorics, Probability and Computing* 13.1 (2004), pp. 17–29 (cit. on p. 2).

[Hat12] Hamed Hatami. “A structure theorem for Boolean functions with small total influences”. In: *Annals of Mathematics* (2012), pp. 509–533. DOI: 10.4007/annals.2012.176.1.9 (cit. on pp. 3, 5, 9).

[KKL88] Jeff Kahn, Gil Kalai, and Nathan Linial. “The Influence of Variables on Boolean Functions (Extended Abstract)”. In: *29th Annual Symposium on Foundations of Computer Science, White Plains, New York, USA, 24-26 October 1988*. IEEE Computer Society, 1988, pp. 68–80. DOI: 10.1109/SFCS.1988.21923 (cit. on pp. 1, 6).

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il