

Biasing Boolean Functions and Collective Coin-Flipping Protocols over Arbitrary Product Distributions

Yuval Filmus 

Computer Science Department, Technion, Haifa, Israel
<http://www.cs.toronto.edu/~yuvalf/>
yuvalfi@cs.technion.ac.il

Lianna Hambardzumyan

School of Computer Science, McGill University, Montreal, QC, Canada
lianna.hambardzumyan@mail.mcgill.ca

Hamed Hatami 

School of Computer Science, McGill University, Montreal, QC, Canada
<https://www.cs.mcgill.ca/~hatami/>
hatami@cs.mcgill.ca

Pooya Hatami 

Department of Computer Science, UT Austin, Austin, TX, USA
<https://pooyahatami.org/>
pooyahat@gmail.com

David Zuckerman

Department of Computer Science, UT Austin, Austin, TX, USA
<http://www.cs.utexas.edu/~diz/>
diz@cs.utexas.edu

Abstract

The seminal result of Kahn, Kalai and Linial shows that a coalition of $O(\frac{n}{\log n})$ players can bias the outcome of *any* Boolean function $\{0, 1\}^n \rightarrow \{0, 1\}$ with respect to the uniform measure. We extend their result to arbitrary product measures on $\{0, 1\}^n$, by combining their argument with a completely different argument that handles very biased input bits.

We view this result as a step towards proving a conjecture of Friedgut, which states that Boolean functions on the continuous cube $[0, 1]^n$ (or, equivalently, on $\{1, \dots, n\}^n$) can be biased using coalitions of $o(n)$ players. This is the first step taken in this direction since Friedgut proposed the conjecture in 2004.

Russell, Saks and Zuckerman extended the result of Kahn, Kalai and Linial to multi-round protocols, showing that when the number of rounds is $o(\log^* n)$, a coalition of $o(n)$ players can bias the outcome with respect to the uniform measure. We extend this result as well to arbitrary product measures on $\{0, 1\}^n$.

The argument of Russell et al. relies on the fact that a coalition of $o(n)$ players can boost the expectation of any Boolean function from ϵ to $1 - \epsilon$ with respect to the uniform measure. This fails for general product distributions, as the example of the AND function with respect to $\mu_{1-1/n}$ shows. Instead, we use a novel boosting argument alongside a generalization of our first result to arbitrary finite ranges.

2012 ACM Subject Classification Theory of computation

Keywords and phrases Boolean function analysis, coin flipping

Digital Object Identifier 10.4230/LIPIcs.ICALP.2019.58

Category Track A: Algorithms, Complexity and Games



© Yuval Filmus, Lianna Hambardzumyan, Hamed Hatami, Pooya Hatami, and David Zuckerman;
licensed under Creative Commons License CC-BY

46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).
Editors: Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi;
Article No. 58; pp. 58:1–58:13



Leibniz International Proceedings in Informatics
LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Related Version A full version of the paper is available at <https://eccc.weizmann.ac.il/report/2019/029/>.

Funding Yuval Filmus: Taub Fellow – supported by the Taub Foundations. The research was funded by ISF grant 1337/16.

Hamed Hatami: Supported by an NSERC grant.

Pooya Hatami: Supported by a Simons Investigator Award (#409864, David Zuckerman).

David Zuckerman: Supported by NSF Grant CCF-1705028 and a Simons Investigator Award (#409864).

Acknowledgements Part of the work on this paper was done while the first three authors were at the Simons Institute for the Theory of Computing at Berkeley, CA, USA.

1 Introduction

How can distributed processors collectively flip a somewhat fair coin if some processors may try to bias the outcome? In the *Collective Coin-Flipping Problem*, a classical problem in distributed computing, n processors wish to generate a single common random bit, even in the presence of faulty and possibly malicious processors. Collective coin-flipping protocols can be used to expedite *Byzantine Agreement* [6] and are closely related to *Leader Election Protocols* [7]. The problem has been considered in several scenarios, depending on the assumptions made on the type of the communication between the processors, the kind and number of faults, and the power of the adversary [6, 3, 7, 2].

A Boolean function $f: \{0, 1\}^n \rightarrow \{0, 1\}$, where $\{0, 1\}^n$ is endowed with a product measure μ , naturally corresponds to a *single round collective coin-flipping protocol* in the *perfect information* model introduced by Ben-Or and Linial [2], where n players each broadcast a bit according to a private distribution, and at the end, the output of the protocol is the value of f on the broadcast string. An interesting and important concept in the design of collective coin-flipping protocols is *resilience* against coalitions of a significant number of players who attempt to influence the output of the protocol towards a particular value.

A *coalition* is a subset S of players that have a particular desired value $b \in \{0, 1\}$ in mind, and if possible, broadcast bits that set the output of the protocol to b . We study the model where the coalition is allowed *rushing*: the corrupt players may wait until all the other players broadcast their bits before deciding on what bit to broadcast. In other words, they *succeed* on $x \sim \mu$ if it is possible to modify x only on the coordinates in S to obtain a string y with $f(y) = b$; they *fail* if the value of f is already determined to be *not* equal to b by the bits broadcast by the players outside the coalition. The success of such a coalition can be easily quantified as the probability that the coalition succeeds on a random $x \sim \mu$.

Fix a parameter $\epsilon > 0$. A protocol f is said to be ϵ -*resilient against coalitions of ℓ players* if no coalition of size at most ℓ succeeds with probability at least $1 - \epsilon$. How resilient can a function be against large coalitions? Over the uniform distribution, perhaps the most natural candidate for a highly resilient function is the majority function, which can be easily seen to be resilient against $\Omega(\sqrt{n})$ size coalitions. However, somewhat surprisingly, it turns out that plain democracy is not the most effective way to be immune against the influence of coalitions. Indeed, Ajtai and Linial [1] gave a *randomized* construction of a Boolean function that is resilient against coalitions of size $\Omega(n/\log^2 n)$, significantly better than the $\Omega(\sqrt{n})$ bound of the majority function. More recently, Chattopadhyay and Zuckerman [5] gave an *explicit* construction of a highly resilient function over the uniform measure. This was a key ingredient in their breakthrough work that introduced explicit two-source extractors

for polylogarithmic min-entropy. Subsequently, Meka [11] gave an explicit construction of a monotone depth three Boolean function that is as resilient as the randomized construction of Ajtai and Linial.

In this article, we are mainly interested in the limitations of resilience. The most classical theorem in this direction is due to Kahn, Kalai, and Linial [10], who proved that, for the uniform distribution, no Boolean function is resilient against coalitions of size $\omega(n/\log n)$. Closing the gap between this bound and the $\Omega(n/\log^2 n)$ construction of Ajtai and Linial remains a longstanding open problem.

Starting with the work of Ben-Or and Linial [2], researchers have studied two natural ways to generalize the discussed protocols: First, allow players to broadcast longer messages, and second, allow many rounds. In this paper, we mostly focus on the latter generalization. In the multi-round setting, the voting procedure that is described above is repeated r times: at every round, first the players who are not in the coalition broadcast their random messages, and then the players in the coalition decide and broadcast their messages in an adversarial manner. When the players are sending single-bit messages, the outcome is decided by a function $f: (\{0, 1\}^n)^r \rightarrow \{0, 1\}$.

The most efficient known protocols are due to Russell and Zuckerman [13] and to Feige [8]. In the case where players are allowed to send longer messages, they constructed $\log^* n + O(1)$ round protocols resilient against coalitions of size βn for any $\beta < 1/2$. In the case when players are allowed to broadcast single bit messages, their protocols use $(1 - o(1)) \log n$ rounds, and are still resilient against coalitions of size βn for any $\beta < 1/2$. For a discussion of various models and known upper and lower bounds, see a survey of Dodis [7].

In the multi-round setting, the players in the coalition have the disadvantage that they will not see the future-round votes of the other players before voting in the current round. Thus, it becomes significantly more difficult to prove limitations on resilience as r grows, and naturally the known bounds are weaker. Russell, Saks and Zuckerman [12], building upon the work of Kahn et al. [10], showed that over the uniform measure, no Boolean function $f: (\{0, 1\}^n)^r \rightarrow \{0, 1\}$ is ϵ -resilient against coalitions of size $\omega_\epsilon \left(\frac{r^2 n}{\log^{(2r-1)} n} \right)$, where $\log^{(2r-1)} n$ is an iterated logarithm. It follows as a simple corollary that $\Omega(\log^* n)$ rounds are necessary in order for a protocol to be resilient against coalitions of size $\Omega(n)$.

The purpose of this paper is to generalize the above results from the uniform distribution to arbitrary product distributions on the Boolean cube.

A moment of reflection reveals that there are major differences between the uniform distribution and the general case, and indeed, prior to this work, it was not clear to us whether similar results were true for general product distributions. We will elaborate on this later, but for now, we only mention that the coordinates x_i that are not highly biased, i.e. $t \leq \Pr[x_i = 1] \leq 1 - t$ for some t that is not too small, can be handled using the same argument as in Kahn et al. [10]. Similarly, the argument of Russell et al. [12] can be used to analyze these coordinates in the multi-round setting. However, the highly biased coordinates behave very differently, and to handle those, we need to take an entirely new approach, and employ a new set of ideas. Indeed, our proofs for the highly biased case have almost no resemblance to those in previous works.

Our first theorem concerns single round protocols. By combining the argument of Kahn, Kalai and Linial with an argument geared towards biased coordinates, we are able to show that these protocols can always be influenced towards a single value, with coalitions which are only slightly worse than those guaranteed by the KKL theorem.

► **Theorem 1.** *Over any product distribution μ , there is no function $f: \{0,1\}^n \rightarrow \{0,1\}$ that is ϵ -resilient against coalitions of size $\omega_\epsilon\left(\frac{n \log \log n}{\log n}\right)$.*

(In contrast, the KKL theorem shows the impossibility of ϵ -resilience against coalitions of size $\omega_\epsilon\left(\frac{n}{\log n}\right)$.)

Next, we prove an impossibility result for resilience in the multi-round setting over arbitrary product distributions. This was posed as an open problem by Russell et al. [12]. Here we face several new challenges. Generalizing our argument for the biased coordinates to the multi-round setting is far from straightforward, and combining it with the argument of Russell et al. [12] for the unbiased coordinates also requires new ideas.

► **Theorem 2.** *Let n and $r < n$ be given. Over any product distribution μ over $(\{0,1\}^n)^r$, there is no r -round coin-flipping protocol $f: (\{0,1\}^n)^r \rightarrow \{0,1\}$ that is ϵ -resilient against coalitions of size $\omega_\epsilon\left(\frac{n(\log^* n)^2}{\log^{(4r)} n}\right)$.*

As a result, over any product distribution μ , $\Omega(\log^* n)$ rounds are necessary in order for a protocol to be resilient against coalitions of size $\Omega(n)$.

Influences. The notion of resilience of a Boolean function is related to the influences of variables and coalitions of variables. For a Boolean function $f: \{0,1\}^n \rightarrow \{0,1\}$ over a product probability measure μ , the *influence* of the k -th variable is defined as

$$I_k(f) := \Pr_{x \sim \mu} [f \text{ is not constant on } B_k(x)],$$

where $B_k(x) := \{y \in \{0,1\}^n : y_j = x_j \text{ for all } j \neq k\}$.

The influence of the k -th variable *towards a value* $b \in \{0,1\}$ is defined as

$$I_k^b(f) := \Pr_{x \sim \mu} [b \in f(B_k(x))].$$

Similarly, the *influence of a coalition* $S \subseteq [n]$ towards a value $b \in \{0,1\}$ is defined as

$$I_S^b(f) := \Pr_{x \sim \mu} [b \in f(B_S(x))],$$

where $B_S(x) := \{y \in \{0,1\}^n : y_j = x_j \text{ for all } j \notin S\}$.

Equivalently, $I_S^b(f)$ is the probability that a random $x \sim \mu$ can be modified on its S variables such that the output of f becomes b .

A function f is *not* ϵ -resilient against coalitions of size ℓ if and only if there exists a set S of size at most ℓ and a value b such that $I_S^b(f) \geq 1 - \epsilon$.

The seminal work of Kahn, Kalai and Linial introduced discrete Fourier-analytic techniques to the study of influences. Their main theorem, known as the KKL inequality, states that over the uniform measure, every unbiased Boolean function $f: \{0,1\}^n \rightarrow \{0,1\}$ has an influential variable. Formally, there exists k such that $I_k(f) \geq \Omega\left(\frac{\alpha \log n}{n}\right)$ when $\alpha \leq \mathbb{E}[f(x)] \leq 1 - \alpha$. Let $b \in \{0,1\}$ satisfy $\Pr[f(x) = b] \geq \epsilon$. Then repeated applications of the KKL inequality imply the existence of a set S with $|S| = O_\epsilon\left(\frac{n}{\log n}\right)$ such that $I_S^b(f) \geq 1 - \epsilon$. In particular, there are no $\omega_\epsilon(n/\log n)$ -resilient functions over the uniform distribution.

The above argument shows that unless f is already very biased towards 0 or 1, one can pick any $b \in \{0,1\}$ and find a small coalition S that can bias f towards b . However, this is no longer true if we consider general product distributions.

► **Example 3.** Consider the p -biased distribution μ_p^n over $\{0, 1\}^n$, i.e. each coordinate is 1 with probability p . Set $p = 1/n$ and let f be the OR function $\bigvee_{i=1}^n x_i$. Obviously, $\mathbb{E}[f] = 1 - (1 - p)^n \approx 1 - \frac{1}{e}$, and yet for every S with $|S| = o(n)$, we have $I_S^0(f) = 1 - (1 - p)^{n-|S|} \approx 1 - \frac{1}{e}$. In other words, despite the fact that the expected value of the function is bounded away from both 0 and 1, no small coalition can influence the output of the function towards 0. However, this is not a counterexample to Theorem 1 because any set S with $|S| = 1$ satisfies $I_S^1(f) = 1$, and thus the function is not even 1-resilient.

As the above example illustrates, part of the difficulty of generalizing the coalition theorem of KKL is to figure out which $b \in \{0, 1\}$ to bias towards.

Using the notation $I_S^b(f)$, Theorem 1 can be restated as follows.

► **Theorem 4** (Theorem 1 reformulated). *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}$ be a function over a product distribution μ . There exists a set S of size $O_\epsilon(\frac{n \log \log n}{\log n})$ such that $I_S^b(f) \geq 1 - \epsilon$ for some $b \in \{0, 1\}$.*

► **Remark 5.** To simplify the statement, in Theorem 4, we did not explicitly state the dependence of $|S|$ on ϵ . Our proof yields the bound $|S| = O(\frac{\log(1/\epsilon)n}{\epsilon \log n} + \frac{n \log \log n}{\epsilon \log n})$.

Continuous cube and a conjecture of Friedgut. The Bernoulli distribution on $\{0, 1\}$ with parameter p can be embedded in the continuous interval $[0, 1]$ via the measure-preserving map $\sigma: [0, 1] \rightarrow \{0, 1\}$ defined as $\sigma(x) = 1$ if and only if $x \geq 1 - p$. By taking the product of these maps, for every product probability measure μ on $\{0, 1\}^n$, we obtain a measure-preserving map $\sigma_\mu: [0, 1]^n \rightarrow \{0, 1\}^n$. As a result, every function $f: (\{0, 1\}^n, \mu) \rightarrow \{0, 1\}$ naturally corresponds to a function $\bar{f}: [0, 1]^n \rightarrow \{0, 1\}$ defined by $\bar{f} = f \circ \sigma_\mu$. Note that $I_S^b[\bar{f}] = I_S^b[f]$, for every $S \subseteq [n]$ and $b \in \{0, 1\}$. Thus, a more general setting for studying resilience is the set of measurable functions $f: [0, 1]^n \rightarrow \{0, 1\}$. Indeed, Bourgain et al [4] proved a generalization of the KKL inequality, but erroneously claimed that as a corollary, if $\epsilon \leq \mathbb{E}[f]$, then $I_S^1[f] \geq 1 - \epsilon$ for a set S of size $|S| = o_\epsilon(n)$. Interestingly, Example 3, which was introduced in the same paper to demonstrate that the proof of the KKL inequality breaks down for the continuous cube, is also a counterexample to this false claim. Friedgut [9] pointed out this error, and suggested the following tantalizing conjecture to replace the false statement¹.

► **Conjecture 6** ([9]). *Let $f: [0, 1]^n \rightarrow \{0, 1\}$ be a measurable function. There exists a set S of size $o_\epsilon(n)$ such that $I_S^b(f) \geq 1 - \epsilon$ for some $b \in \{0, 1\}$.*

A standard compression argument shows that it suffices to prove this conjecture for increasing functions, and indeed the original form of the conjecture is stated for increasing functions. Furthermore, by discretization, the statement can be further reduced to functions $f: \{1, \dots, n^2\}^n \rightarrow \{0, 1\}$, where the domain is endowed with the uniform measure. Note that this form of the conjecture corresponds to resilience of one-round collective coin-flipping protocols where each player is allowed to send $\log n$ -bit messages.

The above discussion show that, qualitatively, Conjecture 6 is a generalization of Theorem 4, and thus our theorem can be considered as a step towards resolving Friedgut's conjecture. However, our techniques and ideas seem to fall short of proving the full conjecture.

¹ Nati Linial told the last author about this error and conjecture years earlier, but as far as we know this is the first published account.

Beyond the Boolean range. As we discussed above, the coalition theorem of KKL says that if $\mathbb{E}[f(x) = b] \geq \epsilon$ then there exists a small coalition S such that $I_S^b(f) \geq 1 - \epsilon$. Now consider a function $h: \{0, 1\}^n \rightarrow \mathcal{R}$ over the uniform distribution, where $\mathcal{R}$ is a constant size set. Pick any $b \in \mathcal{R}$ with $\Pr[h(x) = b] \geq \epsilon$. We can apply the KKL theorem to the function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ defined as $f(x) = 1$ if and only if $h(x) = b$, and conclude that there is a coalition of size $O_\epsilon(n/\log n)$ with $I_S^b(f) \geq 1 - \epsilon$. This shows that over the uniform distribution, the general range $\mathcal{R}$ easily reduces to the Boolean range.

Unfortunately, the above reduction cannot be carried for general product distributions, for in Theorem 4, the final outcome b is dictated to us by the function. To illustrate the problem, consider a function $h: \{0, 1\}^n \rightarrow \{0, 1, 2\}$ and a general product distribution μ . By bundling $\{1, 2\}$ into a single value and applying Theorem 4, we can conclude that there exists a small coalition S such that either it biases the outcome of the function towards 0, or it biases the outcome towards being in $\{1, 2\}$. If it is the former case, then we are done, but in the latter case, it is not clear how to proceed.

We know that except for the x 's that belong to a small-measure set $\mathcal{E}$, the coalition can modify x in such a way that the outcome is in $\{1, 2\}$. Now at first glance, it might seem that by applying Theorem 4 again, we can find another coalition T that can modify x further to refine the outcome to a single value $b \in \{1, 2\}$, and thus conclude that for most x 's the alliance $S \cup T$ can influence the outcome of the function towards b . Unfortunately, this is actually not the case. One reason is that S and T might intersect, and suggest conflicting modifications to x . Even if S and T are disjoint, the proof doesn't work: denoting by x' the vector obtained from $x \sim \mu$ after modification by S , we no longer have $x' \sim \mu$, and so there is no guarantee that on most inputs T can be applied successfully. In other words, $\Pr[x' \in \mathcal{E}]$ need not be small.

The above discussion shows that one cannot deduce the general case via the simple reduction that was outlined above for the uniform measure, but surely, as cumbersome as it may be, one can go over the proof and generalize every step from $\{0, 1\}$ to $\{0, 1, 2\}$ by making small notational adjustments. This turns out not to be the case either! The proof of Theorem 4, rather unexpectedly, relies on the assumption that the function takes only two values. Indeed, to generalize the result to larger ranges, we had to introduce new ideas, and in particular a strengthening of Theorem 4 (see Theorem 15 below) that provides stronger control over the set $\mathcal{E}$ described above.

► **Theorem 7** (Single round, general range). *Let $\mathcal{R}$ be a constant size set, and $f: \{0, 1\}^n \rightarrow \mathcal{R}$ be a function over a product distribution μ . There exists a set S of size $O_\epsilon(\frac{n \log \log n}{\log n})$ such that $I_S^b(f) \geq 1 - \epsilon$ for some $b \in \mathcal{R}$.*

► **Remark 8.** At the heart of the proof of Theorem 7 there is an intermediate result, Theorem 15, which states that if all coordinates are biased, say $\Pr[x_i = 1] < \alpha$, then a random coalition of size $O(\log^3 |\mathcal{R}| \log \log |\mathcal{R}| \cdot \alpha n)$ biases the outcome with high probability. This intermediate result is an essential ingredient in the proof of our result on the multi-round setting, Theorem 2. For this application, it was crucial to obtain a bound which depends only polylogarithmically in $|\mathcal{R}|$.

Even though Theorem 4 is a special case of Theorem 7, we prove them separately, as Theorem 4 can be proven using a shorter and simpler proof.

Paper organization. We prove Theorem 1, which shows that all single-round protocols can be biased using coalitions of size $o(n)$, in Section 2. We prove Theorem 7, which generalizes the preceding result to arbitrary finite domains, in Section 3. We prove our main result, Theorem 2, which shows the multi-round protocols can be biased, in Section 4.

Due to space constraints, some of the proofs are available only in the full version of the paper, which is attached as an appendix.

2 Single Round Case: Proof of Theorem 1

In this section we prove Theorem 1, showing that, under any product distribution, there exists a small coalition which can bias the output of the function towards one of the outputs.

Note that in order to prove Theorem 1, without loss of generality, we can assume that $\Pr_{x \sim \mu}[x_i = 1] \leq \frac{1}{2}$ for every $i \in [n]$, as otherwise we can simply change the role of 0 and 1 for the i -th coordinate. In light of this observation, the coordinates can be divided into two sets: the small bias coordinates, satisfying $\Pr_{x \sim \mu}[x_i = 1] \in (\alpha_0, \frac{1}{2}]$, and the highly biased coordinates, satisfying $\Pr_{x \sim \mu}[x_i = 1] \leq \alpha_0$, where α_0 is a threshold that is chosen to be $\alpha_0 = \frac{1}{\log n}$.

Indeed, we first consider the case where all the coordinates are of the same type:

- *Small bias case:* $\Pr_{x \sim \mu}[x_i = 1] \in (\alpha_0, \frac{1}{2}]$ for every $i \in [n]$.
- *Large bias case:* $\Pr_{x \sim \mu}[x_i = 1] \leq \alpha_0$ for every $i \in [n]$.

We handle the large bias case in Section 2.1, which is the novel part of the proof. The small bias case is handled in Section 2.2 via a reduction to the previous work of Russell et al. [12]. Finally, in Section 2.3 we show how to combine the two cases to handle any product distribution μ , thus completing the proof of Theorem 1.

2.1 Large Bias Case

We will sometimes identify the subsets of $[n]$ with elements of $\{0, 1\}^n$. For example, $S \sim \mu$ would mean that $S = \text{supp}(x)$, where x is sampled according to μ . We construct the coalitions from a certain boosted form of μ .

► **Definition 9** (Boosted distribution). *For a positive integer t , we denote by $\mu^{(t)}$ the distribution of $x^1 \vee \dots \vee x^t$, where $x^1, \dots, x^t$ are i.i.d. random variables distributed according to μ .*

The large bias case of Theorem 1 follows from the following general proposition, that holds for distributions that are not necessarily product distributions.

► **Proposition 10.** *Consider $f: (\{0, 1\}^n, \mu) \rightarrow \{0, 1\}$, where μ is an arbitrary probability measure, and let $S \sim \mu^{(k)}$, where $k \approx \frac{10 \log \frac{1}{\epsilon}}{\epsilon}$. For some $b \in \{0, 1\}$, we have $\Pr_S[I_S^b[f] > 1 - \epsilon] > 1 - \epsilon$.*

Note that Proposition 10 implies (via a straightforward concentration bound) that in the large bias case, there exists a random coalition of expected size at most $k\alpha_0 n$ such that $\Pr_S[I_S^b[f] > 1 - \epsilon] > 1 - \epsilon$. As it will become apparent later, for the application to the multiround setting, it is important that in Proposition 10 the set S is chosen randomly from a distribution that does not depend on f .

Proposition 10 is a direct consequence of the following lemma, as for the Boolean range $\{0, 1\}$, either Condition I holds for $b = 0$ or Condition II holds for $b = 1$. This, however, is not true for larger $\mathcal{R}$.

► **Lemma 11** (Key Lemma for Single Round). *Consider $f: (\{0, 1\}^n, \mu) \rightarrow \mathcal{R}$, where μ is an arbitrary probability measure. Let $x, y \sim \mu$, $S \sim \mu^{(k)}$, where $k \approx \frac{10 \log \frac{1}{\epsilon}}{\epsilon}$. For $b \in \mathcal{R}$, either of*

- *Condition I:* $\Pr_x[\Pr_y[f(x \vee y) = b] \geq 1 - \epsilon] > \epsilon/2$, or
- *Condition II:* $\Pr_x[\Pr_y[f(x \vee y) = b] \geq \epsilon] \geq 1 - \epsilon/2$,

implies $\Pr_S[I_S^b[f] > 1 - \epsilon] > 1 - \epsilon$.

Proof. Let $S = \text{supp}(y^1 \vee \dots \vee y^k)$, where $y^1, \dots, y^k \sim \mu$ are drawn independently. Let the sets X^I and X^{II} denote the following subsets of the input space $\{0, 1\}^n$:

$$X^I = \{x : \Pr_y[f(x \vee y) = b] \geq 1 - \epsilon\},$$

$$X^{II} = \{x : \Pr_y[f(x \vee y) = b] \geq \epsilon\}.$$

If we are in the Type I setting, then $\Pr[X^I] > \epsilon/2$, and so

$$\Pr_S[S \text{ contains some } x \in X^I] \geq 1 - \Pr[y^1, \dots, y^k \notin X^I] \geq 1 - \left(1 - \frac{\epsilon}{2}\right)^k > 1 - \epsilon.$$

Note that if there exists $z \in X^I$ which is a subset of S then for every x , the two elements x and $x \vee z$ can only differ on a subset of S , and thus

$$I_S^b(f) \geq \Pr_x[f(x \vee z) = b] > 1 - \epsilon.$$

Now we turn our attention to Condition II. In this case, we shall prove that $\Pr_S[I_S^b[f] < 1 - \epsilon] \leq \epsilon$. Indeed,

$$\begin{aligned} \Pr_S[I_S^b[f] < 1 - \epsilon] &\leq \Pr_{y^1, \dots, y^k} \left[\Pr_x [\exists i \in [k], f(x \vee y^i) = b] < 1 - \epsilon \right] \\ &= \Pr_{y^1, \dots, y^k} \left[\Pr_x [\forall i \in [k], f(x \vee y^i) \neq b] \geq \epsilon \right]. \end{aligned} \quad (1)$$

To bound the last probability, for $x \in \{0, 1\}^n$ let E_x denote the event that for every $i \in [k]$, $f(x \vee y^i) \neq b$. Then

$$\Pr_x[E_x] \leq \Pr_x[x \notin X^{II}] + \Pr_x[E_x \wedge x \in X^{II}] \leq \frac{\epsilon}{2} + \Pr_x[E_x \mid x \in X^{II}].$$

Plugging this into (1), we get

$$\begin{aligned} \Pr_S[I_S^b[f] < 1 - \epsilon] &\leq \Pr_{y^1, \dots, y^k} [\Pr_x[E_x] \geq \epsilon] \leq \Pr_{y^1, \dots, y^k} \left[\Pr_x[E_x \mid x \in X^{II}] \geq \frac{\epsilon}{2} \right] \leq \\ &\quad \frac{1}{(\epsilon/2)} \Pr_{y^1, \dots, y^k, x} [E_x \mid x \in X^{II}]. \end{aligned}$$

Since $k \approx \frac{10 \log \frac{1}{\epsilon}}{\epsilon}$,

$$\Pr_{x, y^1, \dots, y^k} [E_x \mid x \in X^{II}] \leq (1 - \epsilon)^k \leq \frac{\epsilon^2}{2},$$

showing that

$$\Pr_S[I_S^b[f] < 1 - \epsilon] \leq \frac{1}{(\epsilon/2)} \cdot \frac{\epsilon^2}{2} \leq \epsilon. \quad \blacktriangleleft$$

2.2 Small Bias Case

To handle the small bias case for the sake of proving Theorem 1, one can simply repeat the argument of Kahn et al. [10], i.e. iteratively select influential variables and set them to the value that increases the probability of success. However, for the purposes of our results in the multi-round setting, we will need to prove a stronger result, which states that even if the coalition is selected *randomly*, there is a nontrivial chance of succeeding in influencing the outcome.

► **Lemma 12.** *Let $n \in \mathbb{N}$, $\gamma \in (0, 1/2)$ and $m \leq n$. Let $f: (\{0, 1\}^n, \mu) \rightarrow \{0, 1\}$, where μ is a product distribution such that for all i , $1/n < \alpha \leq \mathbb{E}[x_i] \leq 1/2$. Assume $m > \frac{n \log 1/\alpha}{2\gamma \log n}$. If $\mathbb{E}[f] \geq \gamma$ then*

$$\Pr_{S \subseteq [n]: |S|=m} [I_S^1[f] \geq 1 - \gamma] > \frac{1}{2} \left(\frac{m}{4n \log 1/\alpha} \right)^{2 \frac{80n \log 1/\alpha}{m\gamma}}.$$

Proof. The result is proved in [12] for constant α . The general case follows by representing a μ_p distributed variable as an AND of t variables that are distributed according μ_c , where $c \approx 1/2$ and $p = c^t$. The complete details appear in the full version of the paper. ◀

2.3 Finishing the Proof: Combining the Two Cases

We are ready to finish the proof of Theorem 1. Let $A := \{i : \Pr_{x \sim \mu}[x_i = 1] \in (\alpha_0, \frac{1}{2}]\}$, and recall that $\alpha_0 = \frac{1}{\log n}$. For every $y \in \{0, 1\}^A$, define $f_y: \{0, 1\}^{[n] \setminus A} \rightarrow \{0, 1\}$ as $f_y(z) := f(y, z)$. By Proposition 10, for every $y \in \{0, 1\}^A$, there exists $b := b_y \in \{0, 1\}$ such that

$$\Pr_{S \sim \mu_{[n] \setminus A}^{(k)}} [I_S^b[f_y] > 1 - \frac{\epsilon}{2}] > 1 - \frac{\epsilon}{2},$$

where $k = O\left(\frac{\log(1/\epsilon)}{\epsilon}\right)$. Moreover, since every variable i in $[n] \setminus A$ satisfies $\mathbb{E}[x_i] \leq \alpha_0 = 1/\log n$, Chernoff's bound gives,

$$\Pr_{S \sim \mu_{[n] \setminus A}^{(k)}} \left[|S| \geq \frac{C \log(1/\epsilon)n}{\epsilon \log n} \right] \leq \exp\left(-\Omega\left(\frac{\log(1/\epsilon)n}{\epsilon \log n}\right)\right) \leq \frac{\epsilon}{2},$$

for some constant $C > 0$. Therefore,

$$\Pr_{S \sim \mu_{[n] \setminus A}^{(k)}} \left[I_S^b[f_y] > 1 - \frac{\epsilon}{2} \text{ and } |S| \leq \frac{C \log(1/\epsilon)n}{\epsilon \log n} \right] > 1 - \epsilon.$$

It follows that

$$\mathbb{E}_{S \sim \mu_{[n] \setminus A}^{(k)}} \left[\Pr_{y, b} [I_S^b[f_y] \geq 1 - \frac{\epsilon}{2}] \right] > \frac{1 - \epsilon}{2} \geq \frac{1}{4},$$

assuming without loss of generality that $\epsilon \leq 1/2$. Hence, there exists a fixed $b_0 \in \{0, 1\}$ and a set S , satisfying $|S| \leq \frac{C \log(1/\epsilon)n}{\epsilon \log n}$ and

$$\Pr_y [I_S^{b_0}[f_y] \geq 1 - \frac{\epsilon}{2}] \geq \frac{1}{4}.$$

Now, define $h: \{0, 1\}^n \rightarrow \{0, 1\}$ as $h(y) = 1$ if and only if $I_S^{b_0}[f_{y|A}] \geq 1 - \epsilon/2$. Note that, h depends only on A variables. The above inequality asserts that $\mathbb{E}[h] \geq \frac{1}{4}$. Since, A contains only small bias variables, we may apply Lemma 12. Namely, there is $m = O(\frac{n \log \log n}{\epsilon \log n})$ such that

$$\Pr_{T \subseteq [n]: |T|=m} [I_T^1[h] \geq 1 - \frac{\epsilon}{2}] > 0.$$

Thus, there exists a coalition $T \subseteq A$ of size $O(\frac{n \log \log n}{\epsilon \log n})$ of players that can bias h towards 1. In other words, T can bias y towards cases where S is able to bias f_y towards b_0 . As a result,

$$I_{S \cup T}^{b_0}[f] \geq \left(1 - \frac{\epsilon}{2}\right) \left(1 - \frac{\epsilon}{2}\right) > 1 - \epsilon.$$

Moreover, $|S \cup T| = O\left(\frac{n \log \log n}{\epsilon \log n} + \frac{\log(1/\epsilon)n}{\epsilon \log n}\right)$, as desired.

3 The Larger Range: Proof of Theorem 7

As outlined in the introduction, there are certain obstacles to generalizing Theorem 1 to larger ranges. In particular, the fact that the set $\mathcal{E}$ of all the points on which the coalition fails in Theorem 1 is of small measure does not seem to be a sufficiently strong condition for an induction to go through. We will need to prove a strengthening of Theorem 1 which shows that not only is $\mathcal{E}$ of small measure, but it is also small if it is measured via the boosted distributions introduced in Definition 9. This leads to a more general definition of influence.

► **Definition 13** (Boosted influence towards value). *Let $\mathcal{R}$ be an arbitrary set. For a function $f: \{0, 1\}^n \rightarrow \mathcal{R}$ and $b \in \mathcal{R}$, define $I_S^{b,t}(f) = \Pr_{x \sim \mu^{(t)}}[b \in f(B_S(x))]$.*

Note that $I_S^b(f) = I_S^{b,1}(f)$, as $\mu^{(1)} = \mu$.

The following lemma generalizes Lemma 11, as we spell out in its corollary.

► **Lemma 14.** *Consider $f: \{0, 1\}^n \rightarrow \mathcal{R}$, let $t \in \mathbb{N}$, and let $S \sim \mu^{(k)}$, where $k = \frac{10t}{\delta} \log \frac{t}{\epsilon}$. Let $b \in \mathcal{R}$. We have $\Pr_S[\forall \ell \leq t, I_S^{b,\ell}(f) \geq 1 - \epsilon] \geq 1 - \epsilon$, if any of the following two cases hold:*

- *Case I: For some $s \leq t$, $\Pr_{u \sim \mu^{(s)}}[\Pr_{v \sim \mu^{(t)}}[f(u \vee v) = b] \geq 1 - \epsilon/2] \geq \delta$.*
- *Case II: For every $s \leq t$, $\Pr_{u \sim \mu^{(s)}}[\Pr_{v \sim \mu^{(t)}}[f(u \vee v) = b] \geq \delta] \geq 1 - \frac{\epsilon}{2}$.*

Proof. The complete proof can be found in the full version of the paper. ◀

We can now state the main result of this section. The failure output $\dagger$ allows the inductive proof of Theorem 15, as well as our multi-round result, Theorem 17, to go through, as we explain in Section 4.

► **Theorem 15.** *Let $f: \{0, 1\}^n \rightarrow \{0, 1\}^m \cup \{\dagger\}$, and suppose that $\{0, 1\}^n$ is endowed with a probability measure μ . Let t be a positive integer, and let $S \sim \mu^{(k)}$, where $k = k(m, t, \epsilon) = O(tm^3 \epsilon^{-2} \log \frac{tm}{\epsilon})$. If $\Pr_{\mu^{(\ell)}}[\dagger] < \frac{\epsilon^4}{2^{16}}$ for every $\ell \leq 2t$, then there exists a value $b \in \{0, 1\}^m$ such that $\Pr_S[\forall \ell \leq t, I_S^{b,\ell}(f) \geq 1 - \epsilon] \geq 1 - \epsilon$.*

Proof. The complete proof can be found in the full version of the paper. ◀

Theorem 7 follows from Theorem 15 using an argument very similar to that in Section 2.3, as we show in the full version of the paper.

4 Multi-Round Protocols: Proof of Theorem 2

In this section we will prove Theorem 2, showing that even in the multi-round setting, there are no protocols that are resilient against all coalitions of size $o(n)$. As described in the introduction, here at every round, first the players who are not in the coalition broadcast their random messages, and then the players in the coalition decide and broadcast their messages in an adversarial manner. The outcome is decided by a function $f: (\{0, 1\}^n)^r \rightarrow \{0, 1\}$.

To be more formal, let $\mu = \mu_1 \times \dots \times \mu_r$ be a product distribution over $\{0, 1\}^{rn} \equiv (\{0, 1\}^n)^r$, where each μ_i is a product distribution over $\{0, 1\}^n$. An (n, r) coin-flipping protocol is simply a map $f: (\{0, 1\}^n)^r \rightarrow \{0, 1\}$. Such a protocol is executed in r rounds. In the presence of a coalition $B \subseteq [n]$ of bad players, the protocol operates as follows. In round i , the players in $[n] \setminus B$ select $\alpha^i \in \{0, 1\}^{[n] \setminus B}$ according to $\mu_i|_{[n] \setminus B}$. Then, the bad players B choose their values depending on $\alpha^1, \dots, \alpha^i$. Formally, an (n, r) -strategy for a set $B \subseteq [n]$ is a sequence $\pi = (\pi_1, \dots, \pi_r)$ of functions where $\pi_i: (\{0, 1\}^{[n] \setminus B})^i \rightarrow \{0, 1\}^B$. The function π_i describes the choice of bits the bad players make in the i -th round based on the broadcasted bits of the good players in the first i rounds.

► **Definition 16.** Let $f: (\{0,1\}^n)^r \rightarrow \{0,1\}$ be an (n,r) coin-flipping protocol, and let μ be a product distribution on $(\{0,1\}^n)^r$. Given a Boolean value $b \in \{0,1\}$, a set $B \subseteq [n]$, and an (n,r) -strategy π for the bad players B ,

- $I_{\pi,B}^b(f)$ is the probability that f outputs b given that the bad players B follow π .
- $I_B^b(f) := \sup_{\pi} \{I_{\pi,B}^b(f)\}$ is the influence of B on f towards b .

Our goal is to show that there exists a coalition B of size $o(n)$ such that $I_B^b(f) \geq 1 - \epsilon$ for some $b \in \{0,1\}$. For the moment, let us assume that we have only two rounds, and let $f(x,y)$ denote the protocol, where $x, y \in \{0,1\}^n$ correspond to the inputs in the first and the second round respectively. Let us also denote $f_x(y) := f(x,y)$.

Russell et al. [12] proof of the uniform case. Pick $b \in \{0,1\}$ such that $\Pr[f(x,y) = b] \geq \frac{1}{2}$. Let $\mathcal{A}$ be the set of all $x \in \{0,1\}^n$ that satisfy $\Pr_y[f_x(y) = b] \geq \frac{1}{4}$, and note that $\Pr_x[x \in \mathcal{A}] \geq \frac{1}{4}$. By Lemma 12 of Russell et al. [12], for every $x \in \mathcal{A}$, a random coalition S can bias f_x towards b , with a probability δ that is not too small. Since S is chosen randomly and independently of x , it follows that there exists a fixed coalition S_0 that can bias f_x for at least a δ fraction of $x \in \mathcal{A}$, and thus for at least a $\frac{\delta}{4}$ fraction of $\{0,1\}^n$. Let $\mathcal{A}' \subseteq \mathcal{A} \subseteq \{0,1\}^n$ denote the set of such x . If $x \in \mathcal{A}'$, the coalition S_0 is able to bias the protocol by only interfering in the second round. The set $\mathcal{A}'$ is of measure at least $\frac{\delta}{4}$, which is not too small. Thus, applying Lemma 12 again, we can find another coalition $T_0 = o(n)$ which can modify most x 's to fall in $\mathcal{A}'$. Now we can form the desired coalition $B = T_0 \cup S_0$: In the first round, the players in T_0 try to modify x into an element in $\mathcal{A}'$, and if they succeed, in the second round, the players in S_0 interfere to change the outcome of the protocol into b . This argument easily generalizes to more rounds.

We point out that it was crucial for the above argument, that the distribution of S in Proposition 10 is independent of f .

What fails for the general product distributions. Consider $f(x,y)$ over $\mu = \mu_1 \times \mu_2$, where μ_1 is *highly biased*, and μ_2 is the *uniform distribution*. Similar to the previous paragraph, we can find a set $\mathcal{A}' \subseteq \{0,1\}^n$, a value $b \in \{0,1\}$, and a small coalition S_0 such that $\Pr_{x \sim \mu_1}[x \in \mathcal{A}] \geq \frac{\delta}{4}$, and moreover for every $x \in \mathcal{A}'$, the coalition S_0 is able to influence f towards b by interfering only in the second round. Now, if we are to follow the argument of Russell et al., we would like to find a set T_0 of players to add to the coalition such that, with high probability, T_0 is able to modify a random $x \sim \mu_1$ into an element in $\mathcal{A}'$. We could then conclude that $B = S_0 \cup T_0$ can bias f towards b .

Unfortunately, Proposition 10, the highly-biased counterpart of Lemma 12, only guarantees the existence of a small coalition T_0 which *either* modifies a random $x \sim \mu$ into being in $\mathcal{A}'$ *or* modifies a random $x \sim \mu$ into *not* being in $\mathcal{A}'$; in the latter case, the coalition T_0 is useless. As Example 3 shows, this is not just a caveat of the proof of the proposition. To be more concrete, suppose μ_1 is the $\frac{1}{n}$ -biased distribution, and $\mathcal{A}'$ consists only of the single element $x = \vec{0}$. Even though $\Pr[x \sim \mathcal{A}] \geq \frac{1}{4}$, there is no coalition of size $o(n)$ which can, with high probability, modify a random $x \sim \mu_1$ into an element in $\mathcal{A}'$. On the other hand, even a single player can modify every x into an element outside $\mathcal{A}'$, but this is not helpful for our purposes, as the elements outside $\mathcal{A}'$ are the elements that S_0 cannot handle.

How to overcome the problem. Consider the same setting as in the previous paragraph. We know that for every x , a random coalition S of size $o(n)$ succeeds in influencing f_x towards one of the outputs, with probability at least δ , where δ is not too small. Instead of

picking one S_0 , we select a collection of coalitions that cover almost all x 's. More precisely, we find $S_1, \dots, S_M$ and $b_1, \dots, b_M$, where $M = O_\delta(1)$, such that apart from a small set of exceptions $\mathcal{E} \subseteq \{0, 1\}^n$, every f_x can be biased towards some b_i using the coalition S_i .

Let $h: \{0, 1\}^n \rightarrow \{1, \dots, M\} \cup \{\dagger\}$ be defined as follows: If $x \in \mathcal{E}$, then $h(x) = \dagger$, and otherwise $h(x)$ is equal to some i such that S_i can bias f_x towards b_i . This brings us to the non-Boolean range case, which was analyzed in Section 3. We can apply Theorem 15 to find a coalition T that can influence h towards one of the values in $j \in \{1, \dots, M\}$. Now $B = T \cup S_j$ will be our desired coalition. With high probability, in the first round the players in T can successfully modify a random element x into an element x' with $h(x') = j$, and then in the second round, the players in S_j can modify x' to bias the outcome towards b_j . This is the main new idea used below to resolve the multi-round setting over arbitrary distributions.

Theorem 2 is a consequence of the following more elaborate theorem which states that for sufficiently large n , and $r \leq \log^* n/5$, no (n, r) protocol over an arbitrary product distribution is resilient against coalitions of $m = o(n)$ bad players.

► **Theorem 17.** *For every $\epsilon > 0$, and integers $n > 0$, and $r < \log^* n/5$, there exists $\delta = \Omega(\frac{1}{\log(1/\epsilon)^r n})$, and $m = o(n)$ such that the following holds. For every $f: (\{0, 1\}^n)^r \rightarrow \{0, 1\}$ over a product distribution μ , there exists $b \in \{0, 1\}$, such that the corresponding r -round protocol satisfies $\Pr_{S \sim \nu}[I_S^b(f) \geq 1 - \epsilon] \geq \delta_r$, where ν is a distribution on $\binom{[n]}{m}$ that depends only μ but not on f . To be more precise, one can take $m = O_\epsilon\left(\frac{n \cdot r \cdot 4^r}{\log^{(4r)} n}\right) = O_\epsilon\left(\frac{n(\log^* n)^2}{\log^{(4r)} n}\right)$.*

Proof. The complete proof can be found in the full version of the paper. ◀

5 Concluding Remarks and Open Problems

- Perhaps the most interesting next step is proving limitations for resilience of protocols where players may send longer messages. As was discussed below Conjecture 6, it is conjectured that even when the players are allowed to broadcast arbitrarily long messages, only resilience against coalitions of size $o(n)$ is possible. This question has also been studied in the multi-round setting [12, 13, 8]. In this case, if the players are allowed $\log n$ -bit messages, we know of $(\log^* n + O(1))$ -round protocols resilient against coalitions of size $(1/2 - \epsilon)n$ [13, 8]. On the other hand, Russell et al. [12] showed that $\Omega(\log^* n)$ rounds are necessary if we have the added restriction that in the i -th round the players are allowed messages of length $(\log^{(2i-1)} n)^{1-o(1)}$. Strengthening this impossibility result to messages of length $\Omega(\log n)$ is another interesting problem that remains open.
- The key qualitative point of Theorems 1 and 2 is that there always exists a coalition of size $o(n)$ that can bias the outcome of the protocol towards a particular value. Interestingly, we are not aware of a simpler proof of this weaker qualitative statement even in the case of the uniform measure. The proof techniques introduced in this paper for the highly biased coordinates are more combinatorial and probabilistic in nature; however, the less biased coordinates are ultimately handled by the Fourier-analytic proof of [10]. These Fourier analytic arguments are *hard* in nature, in the sense that their purpose is to give effective bounds. It would be interesting to find more intuitive combinatorial proofs for these statements, potentially at the cost of obtaining less effective bounds, or by appealing to *soft analytic tools* such as compactness, at the cost of obtaining no quantitative bounds. We refer the reader to Terence Tao's blog post [14] for a discussion about hard and soft analysis.
- Over the uniform distribution, Kahn et al. [10] proved that there exists no Boolean function that is ϵ -resilient against coalitions of size $\omega_\epsilon\left(\frac{n}{\log n}\right)$. In this work we show that a similar bound of $\omega_\epsilon\left(\frac{n \log \log n}{\log n}\right)$ on resilience holds over arbitrary product distributions.

A natural question is whether the $\log \log n$ in our bound necessary. However, even in the uniform setting there is work left to be done. Here, the best known constructions guarantee resilience against coalitions of size $O(\frac{n}{\log^2 n})$ [11, 1], which is a factor of $\log n$ off from the impossibility result of Kahn, Kalai, and Linial.

References

- 1 Miklós Ajtai and Nathan Linial. The influence of large coalitions. *Combinatorica*, 13(2):129–145, 1993.
- 2 Michael Ben-Or and Nathan Linial. Collective coin flipping. *Advances in Computing Research*, 5:91–115, 1989.
- 3 Michael Ben-Or, Nathan Linial, and Michael Saks. *Collective coin flipping and other models of imperfect randomness*. IBM Thomas J. Watson Research Division, 1989.
- 4 Jean Bourgain, Jeff Kahn, Gil Kalai, Yitzhak Katznelson, and Nathan Linial. The influence of variables in product spaces. *Israel Journal of Mathematics*, 77(1–2):55–64, 1992.
- 5 Eshan Chattopadhyay and David Zuckerman. Explicit two-source extractors and resilient functions. *Annals of Mathematics, to appear*, 2016. Preliminary version in STOC 2016.
- 6 Benny Chor and Cynthia Dwork. Randomization in Byzantine Agreement. *Advances in Computing Research*, 5:443–497, 1989.
- 7 Yevgeniy Dodis. Fault-tolerant leader election and collective coin-flipping in the full information model. *Survey*, 2006.
- 8 Uriel Feige. Noncryptographic Selection Protocols. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*, page 142. IEEE Computer Society, 1999.
- 9 Ehud Friedgut. Influences in Product Spaces: KKL and BKKKL Revisited. *Combinatorics, Probability and Computing*, 13(1):17–29, 2004.
- 10 Jeff Kahn, Gil Kalai, and Nathan Linial. The influence of variables on Boolean functions. In *Proceedings of the 29th annual FOCS*, pages 68–80, 1988.
- 11 Raghu Meka. Explicit resilient functions matching Ajtai-Linial. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1132–1148. SIAM, 2017.
- 12 Alexander Russell, Michael Saks, and David Zuckerman. Lower bounds for leader election and collective coin-flipping in the perfect information model. *SIAM Journal on Computing*, 31(6):1645–1662, 2002.
- 13 Alexander Russell and David Zuckerman. Perfect information leader election in $\log^* n + O(1)$ rounds. *Journal of Computer and System Sciences*, 63(4):612–626, 2001.
- 14 Terence Tao. Soft analysis, hard analysis, and the finite convergence principle. URL: <https://terrytao.wordpress.com/2007/05/23/soft-analysis-hard-analysis-and-the-finite-convergence-principle/>, 2007. Accessed 10 Feb 2019.