Players' Effects under Limited Independence

Ronen Gradwohl * Omer Reingold † Ariel Yadin ‡ Amir Yehudayoff §

Abstract

In a function that takes its inputs from various players, the effect of a player measures the variation he can cause in the expectation of that function. In this paper we prove a tight upper bound on the number of players with large effect, a bound that holds even when the players' inputs are only known to be pairwise independent. We also study the effect of a set of players, and show that there always exists a "small" set that, when eliminated, leaves every small set with little effect. Finally, we ask whether there always exists a player with positive effect. We answer this question differently in various scenarios, depending on the properties of the function and the distribution of players' inputs. More specifically, we show that if the function is non-monotone or the distribution is only known to be pairwise independent, then it is possible that all players have 0 effect. If the distribution is pairwise independent with minimal support, on the other hand, then there must exist a player with "large" effect.

1 Introduction

A general recurring theme in the analysis of games is the juxtaposition of two distinct sources of players' motivation to act strategically: The first is the myopic maximization of their own immediate gain, and the second is a consideration of the effect their behavior has on other players and possibly on a collective outcome. This theme surfaces in many settings. For example, in a repeated game with incomplete information in which players both maximize their utilities and learn others' preferences, players must strike a balance between playing to obtain immediate gain and playing to learn (or teach). Another example is in the allocation of a public good: the amount a player declares he is willing to contribute affects both his immediate wealth, but also might have an impact on whether or not the good is eventually allocated.

*Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: ronen.gradwohl@weizmann.ac.il.

†Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: omer.reingold@weizmann.ac.il. Research supported by US-Israel Binational Science Foundation Grants 2002246 and 2006060.

‡Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: ariel.yadin@weizmann.ac.il.

§Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: amir.yehudayoff@weizmann.ac.il. Research supported by a grant from the Israel Ministry of Science (IMOS) - Eshkol Fellowship.

In such settings, players who have little impact on others or on a collective outcome more or less ignore the second source of motivation for their strategic behavior, and can be shown to act myopically. This has been demonstrated more precisely in numerous examples: The provision of a public good [12, 1], repeated games [13, 2], and mechanisms for choosing equilibria in private information economies [3].

There are many ways one can quantify the impact of a player on others or on a collective outcome. Two notable notions that have been widely studied in the economics and computer science literatures respectively are the *effect* and the *influence* of players. We illuminate the distinction between the two notions via the example of voting.

Suppose there are n players, and let a function f be a voting scheme between two candidates, 0 and 1. Each player i has a type $t_i \in \{0, 1\}$ denoting his preferred candidate, and it is common knowledge that this type is drawn from a distribution X_i. Each player submits a vote for his type t_i to f, and given these n binary inputs f outputs the name of one of the candidates. The effect of a player is then the amount of variation he can cause in the expectation of f, where the expectation is over the distributions of other players' types. In other words, it is the difference in the expectation of the function conditioned on the player submitting 0 and the expectation conditioned on the player submitting 1. This is an a priori notion – a player's effect is measured without assuming any knowledge of the other players' votes but only their distribution. The influence of a player, on the other hand, is defined as the probability (over the distribution of players' types) that a specific player casts the deciding vote. This is an a posteriori notion, since it measures the impact of a player on the outcome after others have already submitted their true types to f.

Previous Work The notion of influence was introduced by Ben-Or and Linial [5]. The seminal paper in this line of work is that of Kahn, Kalai and Linial [10] (henceforth KKL), in which they showed that in any voting scheme, if the players' signals are independent then there always exists a player with "large" influence. Following this paper, there has been much work studying the notion of influence (see Kalai and Safra [11] for a survey).

The notion of effect was studied by Mailath and Postlewaite [12] and by Fudenberg, Levine, and Pesendorfer [7]. The results most similar to ours are those of Al-Najjar and Smorodinsky [1], who gave tight bounds on the number of players with large effect. One of their assumptions is that the players' signals are independent (or at least independent conditional on some outside information). Their methods do not apply to general distributions.

Haggstrom, Kalai and Mossel [9] studied the notion of effect for general distributions, and showed that there is complete aggregation of information (this means that a small tendency towards one outcome for each player gives a strong tendency in the general outcome) under certain conditions related to the distribution and the effects.

Our Results We answer questions similar to those of [1] and [10] in the context of distributions that are only weakly independent. As in [1], we give a tight bound on the number of players with large effect. The novelty of our bound is that we assume only minimal independence of players' signals. More precisely, our bound holds even when the players are only pairwise independent.

Note that in general, all players may have maximal effect. For example, suppose players' signals are fully correlated, so that with probability 1/2 they all output 0 and with probability 1/2 they all output 1. If the function outputs the majority vote, then the effect of every player is 1.

Additionally, we study the effects of coalitions of players. We show that a small set of players can be eliminated, leaving only coalitions with small effect.

We also ask whether a KKL-type theorem holds for effect – that is, does there always exist some player with large effect? We have three results here: First, we observe that if the function is not monotone, then it is possible that all effects are 0, even in the fully independent case. Second, we show that there exists a pairwise independent distribution and a monotone function such that all players' effects are again 0. Also, we use similar ideas to show that there exists a pairwise independent distribution and a monotone function such that all players have influence 0; i.e. a KKL-type theorem does not hold for influence either, in the case of pairwise independence. Finally, we give a positive result: we show that if the distribution is pairwise independent and of minimal support size, then there exists a player with large effect.

Organization The rest of this paper is organized as follows. Section 2 begins with some formal definitions, and then proceeds with formal statements of all our results. Section 3 contains the proofs of our theorems bounding the number of players and sets with large effect, and Section 4 contains the proofs of our results on KKL-type theorems.

2 Definitions and Main Results

2.1 Definitions

Let $n \in \mathbb{N}$, and let S be a finite set (whose size may depend on n). Let $f : S^n \mapsto [-1, 1]$ be some function, let $X_1, \ldots, X_n$ be n random variables (which are not necessarily independent) taking values in S, and denote by $X = (X_1, \ldots, X_n)$. We think of each X_i as player i's random variable, and denote the set of players as $[n] = \{1, \ldots, n\}$.

Definition 2.1 (Effect) *For $i \in [n]$, denote*

$$\mathcal{E}_i(f, X) = \max_{s,s' \in S} \left| \mathrm{E}_X[f|X_i = s] - \mathrm{E}_X[f|X_i = s'] \right|$$

the effect of player X_i in f with respect to X. For $\alpha \in \mathbb{R}$, we say that X_i has effect *larger than α in f with respect to X if*

$$\mathcal{E}_i(f, X) > \alpha.$$

Denote by $K(f, X, \alpha)$ the number of players with effect larger than α in f with respect to X.

Definition 2.2 (pivotal player) *For $p, \alpha \in \mathbb{R}$, we say that X_i is* (p, α)-pivotal *in f with respect to X if*

$$\Pr\left[\left| \mathrm{E}_X[f|X_i] - \mathrm{E}_X[f] \right| > \alpha \right] > p.$$

Denote by $K(f, X, p, \alpha)$ the number of (p, α)-pivotal players in f with respect to X.

While the effect measures the *maximal* variation a player can cause in a function's expectation, the pivotalness measures both the variation and the probability with which it occurs. The two are closely related: If some player i is (p, α)-pivotal, and p is smaller than the minimal probability of any element in S (according to X_i), then player i has effect larger than α. Conversely, if player i has effect larger than α, then i is $(p, \alpha/2)$-pivotal for every p less than the minimal probability of any element in S.

Finally, we need one more notion. For a random variable $X = (X_1, \ldots, X_n)$ and a set $T \subseteq [n]$, let $X_T = (X_i)_{i \in T}$ be the projection of X onto the variables in T.

Definition 2.3 (Pivotal set of players) *For $p, \alpha \in \mathbb{R}$, we say that $T \subseteq [n]$ is* (p, α)-pivotal *in f with respect to the random variable $X = (X_1, \ldots, X_n)$ if*

$$\Pr\left[\left|\mathrm{E}_X[f|X_T] - \mathrm{E}_X[f]\right| > \alpha\right] > p.$$

The notions of effect, pivotal players, and pivotal sets of players are relevant even for variables that are not fully independent. We relax the assumption of full independence as follows.

Definition 2.4 (k-wise independence) *The random variables $X_1, \ldots, X_n$ are* k-wise independent *if for any subset $T \subset [n]$, $|T| \leq k$, the random variables $\{X_i : i \in T\}$ are independent.*

We also state here the precise definition of influence.

Definition 2.5 (Influence) *Let $f : \{0,1\}^n \to \{0,1\}$ be a function, and let X be a random variable obtaining values in $\{0,1\}^n$. The* influence *of the i'th player is*

$$I_i(f, X) = \Pr\left[f(X) \neq f(X \oplus e_i)\right],$$

where e_i is the vector with 1 at the i'th index and 0 elsewhere, and $\oplus$ is bitwise XOR.

2.2 Bounds on the Number of Pivotal Players

The following theorem bounds the number of pivotal players.

Theorem 2.6 *Let $n \in \mathbb{N}$, and let S be a finite set. Let $f : S^n \mapsto [-1, 1]$ be some function. Let $X_1, \ldots, X_n$ be n pairwise independent random variables taking values in S, and denote $X = (X_1, \ldots, X_n)$. Then for every positive $\alpha, p \in \mathbb{R}$,*

$$K(f, X, p, \alpha) < \frac{8}{p\alpha^2}.$$

We note that the above theorem was known for the case of fully independent random variables (see [1]). However, we consider the much more general case of pairwise independence. We note also that for the case where all the players vote according to the same coin flip of bias $0 < p < 1$ (i.e.

the signals are 1-wise independent and identical), and f is boolean such that $f(0,\ldots,0)=0$ and $f(1,\ldots,1)=1$,

$$K(f,X,1,1)=n,$$

so the conclusion of the theorem does not hold.

We now turn our attention to the effect of a set of players – that is, how much can a set of players cause variation from the expectation of a function. A first observation is the following: if some player i is (p,α)-pivotal, then any set of players that contains i is at least $(p\alpha/4,\alpha/2)$-pivotal.

However, we still prove a rather strong statement, Theorem 2.7. The theorem is a generalization of a theorem of Gradwohl and Reingold [8] to the case of variables that are not fully independent. Roughly, the theorem states that, for any natural number m, it is possible to eliminate a set T of $O(m)$ players, so that **any** set of size at most m that does not intersect T will have little effect. There is a requirement on the distribution of the variables – in particular, the statement holds as long as the variables are $2m$-wise independent.

Theorem 2.7 *Fix some natural number m. Then for any set of $2m$-wise independent random variables $X_1,\ldots,X_n$, any $0<\alpha<1$, $0<p<1$, and any function f, the following holds: there exists a set $C\subseteq[n]$ of size $|C|\leq 8m/p\alpha^2$, such that for **all** $T\subseteq[n]\setminus C$ of size $|T|\leq m$, the set T is **not** (p,α)-pivotal.*

2.3 An Example

The following example, a function we call Maj_p, is from [1], and shows that our bound on the number of pivotal players in Theorem 2.6 is tight. Let $n\in\mathbb{N}$, let $0<p<1$, and let $S=\{0,1,\perp\}$. For each player $i\in[n]$, suppose $\Pr[X_i=0]=\Pr[X_i=1]=p/2$, and $\Pr[X_i=\perp]=1-p$. Let $f:S^n\mapsto\{0,1\}$ be the majority function over all players that did not output $\perp$. In case of a tie, the function outputs 1.

We say that a player *participates* if he does not output $\perp$. Then every player participates with probability p, and is influential if the remaining players who did not output $\perp$ are split evenly between 0's and 1's. If the number of participating players is roughly pn (which holds with very high probability), then the player is influential with probability roughly $1/\sqrt{pn}$. Thus, every one of the n players is approximately $(p,1/\sqrt{pn})$-pivotal. Setting $\alpha\approx 1/\sqrt{pn}$, we get that $n\approx\frac{1}{p\alpha^2}$, which is the number of (p,α)-pivotal players. Note that we can vary the value of α in this example by picking some natural number $k\leq n$ such that $\alpha\approx 1/\sqrt{pk}$. The function we consider is then the function above, but only over some arbitrary set of k players. In this case, those k players will all be (p,α)-pivotal.

2.4 An Application

As mentioned in the Introduction, bounds on players' effects are useful in many frameworks. For many of these frameworks, weakening the assumption on players' independence is interesting. In this section we give an example in which Theorem 2.7 can be used to improve a known result.

Our example is an extension to a theorem of Gradwohl and Reingold [8]. They show that in any large Bayesian game in which players' types are independent, all Nash equilibria are partial exposure (PE) ex post Nash. This is a weakening of the standard ex post property, and it roughly means that in any Nash equilibrium, a player's Nash equilibrium strategy remains almost optimal (up to a small additive loss) even if he observes the realized types and chosen actions of some *subset* of the other players (as opposed to *all* players in ex post Nash).

A very rough idea of the proof is the following. Fix some player i, and suppose u_i is his utility function. If player i observes the types and actions of some players, then his expected utility conditional on those observed values might change, and he may wish to change his strategy. However, if the players whose values he observed have little effect on u_i, the for any action of player i, the expected utility does not change much even conditioned on these observed values. Thus, a deviation from the Nash equilibrium strategy can not increase the payoff of player i by much. This implies that his original (Nash) strategy remains almost optimal. The theorem follows by the result of Al-Najjar and Smorodinsky [1] – there can not be many players with a large effect on u_i.

Using Theorem 2.7 instead of the result of [1] we get that the theorem of [8] holds even if players' types are only known to be k-wise independent (for a suitable choice of k that depends on the size of the exposed set). Other then the use of Theorem 2.7 instead of the result of [1], the proof is identical to that of [8].

2.5 General KKL-Type Results

The questions raised in this section are motivated by the celebrated result of Kahn, Kalai, and Linial [10]. Roughly, their result states that for every balanced Boolean function (a function whose expectation is 1/2) on $\{0,1\}^n$ and fully independent inputs there exists a player with influence at least $\Omega(\log(n)/n)$.

The notions of influence and effect are closely related in several specific cases. Here we ask whether a KKL-type theorem holds with regard to effect. More specifically, we ask the following question:

Does there always exist a player with large effect?

We show that the answer to this question depends strongly on the underlying distribution of the players and the properties of the function. We also extend one of our negative results to the original notion of influence, and show that a KKL-type theorem does not hold for general distributions, even for monotone functions (see Section 4.2).

2.5.1 Full Independence

We first consider the question stated above for the case in which the players' signals are fully independent, and in which the function is monotone. Recall that a Boolean function on $\{0,1\}$ is monotone if $f(a_1,\ldots,a_n) \geq f(b_1,\ldots,b_n$ whenever $a_i \geq b_i$ for all i. For such functions, the notions of influence and effect are equivalent [9]. This means that for fully independent players

and balanced monotone functions, the original KKL theorem states that there exists a player with effect $\Omega(\log(n)/n)$.

How about non-monotone functions? It is possible to transform a non-monotone function into a monotone one in such a way that the influences do not increase [10]. Hence, a lower bound on the influence of a player for all monotone functions gives a similar bound for all functions. However, such a transformation does not exist for effects. Consider, for example, the PARITY function, and suppose each player independently outputs a bit generated by a fair coin toss. This function outputs 1 if the sum of the players' values is odd, and 0 otherwise. Then for PARITY all influences are 1, but all effects are 0. In particular, a KKL-type theorem for effect does not hold for non-monotone functions – there is no general non-trivial lower bound on the effect of a player in non-monotone functions, even in the case of full independence.

Since in general there are no non-trivial lower bounds, we must make some assumptions about the types of functions we consider. One "natural" class of functions consists of monotone functions. This is a fairly large class, and in fact most work on the influence of players focuses on monotone functions. Additionally, monotonicity is often closely related to (ex post) incentive compatibility. For example, suppose the function takes players' bids, and outputs the probability of some public good being allocated. Then monotonicity implies that as players increase their bids, the probability of the good being allocated only increases. Another example is voting, in which it seems natural to require that if a player switches his vote to some alternative then that alternative should have an increased chance of being selected.

2.5.2 Pairwise Independence – Negative Results

For monotone functions with full independence, the original KKL theorem implies that there exists a player with large effect (see Section 2.5.1). We also saw that without monotonicity, this does not hold. In this section we ask the following question: for monotone functions, does a player with large effect necessarily exist even if the inputs are only pairwise independent?

In Section 4.1 we show that there exists a balanced monotone function f and a pairwise independent distribution D over $\{0,1\}^n$ such that

$$\mathcal{E}_i(f, D) = 0$$

for all $i \in [n]$. This implies that there is no non-trivial lower bound on the effect of a player for pairwise independent distributions, even for monotone functions.

Furthermore, in Section 4.2 we extend these results to show that there exists a balanced monotone function g and a pairwise independent distribution D' such that

$$I_i(g, D') = 0$$

for all $i \in [n]$.

2.5.3 Pairwise Independence – Positive Result

In Section 2.5.2, we argued that, for the case of pairwise independence, monotonicity does not suffice in order for some KKL-type theorem to hold. In this section, we show that a KKL-type

theorem does hold in a restricted special case. Roughly, we show that for all pairwise independent distributions with minimal support size, there is a player with effect at least $1/\sqrt{n}$ (for any balanced function).

Theorem 2.8 *Let $n+1 = 2^k$. Let μ be a pairwise independent distribution on $(X_1, \ldots, X_n) \in \{0,1\}^n$ such that $\Pr[X_i = 1] = 1/2$ for all $i \in [n]$, and μ has support of size $n+1$. Let $f : \{0,1\}^n \mapsto \{0,1\}$. Then*

$$\sum_{i \in [n]} (\mathcal{E}_i(f))^2 = \frac{\mathrm{Var}[f]}{4}.$$

A function f is balanced if its expectation is $1/2$, which also implies that its variance is $1/4$. For such functions, the above theorem states that the sum of squares of effects is at least $1/16$, and so there exists a player with effect at least $1/(4\sqrt{n})$.

We now discuss the premise of the above theorem, namely pairwise independent distributions on $\{0,1\}^n$ with support size $n+1$. First, we note that there are such distributions – in the sequel (Section 4.1) we will construct two such examples, μ and $\overline{\mu}$. Second, any pairwise independent distribution on $\{0,1\}^n$ has support of size *at least* $n+1$. Finally, such a distribution with support of size *exactly* $n+1$ must be uniform on its support – see Benjamini, Gurel-Gurevich, and Peled [4].

We also note that the distribution D from Section 4.1 that serves as our counter-example to any KKL-type theorem for effect has support of size $2(n+1)$. It seems a small difference in support size can make a significant difference: For any pairwise independent distribution with support size $n+1$ and balanced function, there exists a player with effect roughly $1/\sqrt{n}$. On the other hand, a convex sum of two such distributions yields a distribution for which there exists a balanced function in which all the effects are 0.

3 Proofs: The Number of Pivotal Players and Sets

We begin with some preliminary definitions. In Section 3.2 we then discuss a weaker result than our main theorem because its proof is instructive in that it contains many of the ideas of the main theorem. The more general case will be proven in Section 3.3.

3.1 Preliminaries

We need some preliminary definitions. Let X be a random variable taking values in $\{0,1\}^k$. For two functions $g, h : \{0,1\}^k \mapsto [-1,1]$, define the inner product (with respect to X)

$$\langle g, h \rangle = \sum_{x \in \{0,1\}^k} \Pr[X = x] \cdot g(x) \cdot h(x)$$

(for simplicity of notation, we omit the dependency on X from the notation, and will make sure it is clear from the context). Define the norm of g to be

$$\|g\| = \sqrt{\langle g, g \rangle}.$$

3.2 Binary Independent Inputs

The theorem we prove here is the following:

Proposition 3.1 *Let $n \in \mathbb{N}$, and let $f : \{0,1\}^n \mapsto [-1,1]$. If $X_1, \ldots, X_n$ are n independent random variables such that $\Pr[X_i = 1] = \Pr[X_i = 0] = 1/2$ for all $i \in [n]$, then for every positive $\alpha \in \mathbb{R}$,*

$$K(f, X, \alpha) < \frac{4}{\alpha^2},$$

where $X = (X_1, \ldots, X_n)$.

Proof: Let $k = K(f, X, \alpha)$, and without loss of generality assume that the first k variables are the ones with effect at least α. Define the function $f' : \{0,1\}^k \mapsto [-1,1]$ as

$$f'(x_1, \ldots, x_k) = \mathrm{E}_X[f | X_1 = x_1, \ldots, X_k = x_k].$$

Also denote $X' = (X_1, \ldots, X_k)$. For every $b \in \{0,1\}$ and for every $i \in [k]$,

$$\mathrm{E}_{X'}[f' | X_i = b] = \mathrm{E}_{X'}[\mathrm{E}_X[f | X_1 = x_1, \ldots, X_k = x_k] | X_i = b] = \mathrm{E}_X[f | X_i = b],$$

and thus player i has effect larger than α in f' if and only if player i has effect larger than α in f. Furthermore, assume without loss of generality that for all $i \in [k]$, we have $\mathrm{E}_X[f | X_i = 0] > \mathrm{E}_X[f | X_i = 1]$, and so

$$\mathrm{E}_X[f | X_i = 0] - \mathrm{E}_X[f | X_i = 1] > \alpha$$

(otherwise consider the function $g(x_1, \ldots, x_n) = f(x_1, \ldots, x_{i-1}, 1 - x_i, \ldots, x_n)$, and note that, due to the independence and symmetry of the random variables, this does not alter the effects of the players).

For $i \in [k]$, define the functions $b_i : \{0,1\}^k \mapsto [-1,1]$ as

$$\forall\ x = (x_1, \ldots, x_k) \in \{0,1\}^k \quad b_i(x) = (-1)^{x_i}.$$

The functions b_i are useful because

$$\left\langle b_i, f' \right\rangle = 2^{-k} \sum_{x : x_i = 0} f'(x) - 2^{-k} \sum_{x : x_i = 1} f'(x) = \frac{1}{2}\mathrm{E}_X[f' | X_i = 0] - \frac{1}{2}\mathrm{E}_X[f' | X_i = 1] > \frac{\alpha}{2}$$

by assumption on the effects of $X_1, \ldots, X_k$. Also for $i, j \in [k]$, since $X_1, \ldots, X_n$ are independent,

$$\langle b_i, b_j \rangle = \Pr_{X'}[X_i = X_j] - \Pr_{X'}[X_i \neq X_j] = \begin{cases} 1 & i = j \\ 0 & i \neq j, \end{cases}$$

and so

$$\|b_1 + \ldots + b_k\| = \sqrt{k}.$$

Now, on the one hand,

$$\left\langle b_1 + \ldots + b_k, f' \right\rangle > \frac{k\alpha}{2}.$$

On the other hand,

$$\langle b_1 + \ldots + b_k, f' \rangle \leq \|b_1 + \ldots + b_k\| \cdot \|f'\| \leq \sqrt{k},$$

by Cauchy-Schwarz and since $\|f'\| \leq 1$. Combining the two inequalities yields

$$k < \frac{4}{\alpha^2}.$$

■

3.3 More General $\{0,1\}^n$ Case

To prove Theorem 2.6, we first assume that the variables X_i are binary, albeit with skewed probabilities. Later we reduce the general case to such variables.

We first prove a bound on the sum of effects of a subset of players.

Lemma 3.2 *Let $n \in \mathbb{N}$ and $f : \{0,1\}^n \mapsto [-1,1]$, and consider pairwise independent binary random variables $X_1, \ldots, X_n$ with $\Pr[X_i = 0] = q$, where $0 < q < 1$, for all $i \in [n]$. Then for all $k \in [n]$,*

$$\sum_{i \in [k]} \mathcal{E}_i(f, X) \leq \sqrt{\frac{2k}{p}},$$

where $p = \min\{q, 1-q\}$.

Proof: Let $z \in \{0,1\}$ be such that $\Pr[X_i = z] = p$. For each $i \in [k]$, fix $a_i \in \{-1,1\}$ such that

$$a_i \cdot (\mathrm{E}_X[f|X_i = z] - \mathrm{E}_X[f|X_i = 1-z]) \geq 0.$$

Also define the functions $b_i : \{0,1\}^k \mapsto [-1,1]$ as

$$b_i(x) = \begin{cases} a_i & x_i = z \\ -a_i \cdot \frac{p}{1-p} & x_i = 1-z \end{cases}$$

for all $x = (x_1, \ldots, x_n) \in \{0,1\}^n$. Again consider the quantity

$$\begin{aligned} \langle b_i, f \rangle &= a_i \sum_{x : x_i = z} \Pr[X = x] f(x) - a_i \frac{p}{1-p} \sum_{x : x_i = 1-z} \Pr[X = x] f(x) \\ &= a_i \Pr[X_i = z] \cdot \mathrm{E}[f|X_i = z] - a_i \frac{p}{1-p} \cdot \Pr[X_i = 1-z] \cdot \mathrm{E}[f|X_i = 1-z] \\ &= a_i p \left(\mathrm{E}[f|X_i = z] - \mathrm{E}[f|X_i = 1-z]\right) \\ &= p \cdot \mathcal{E}_i(f, X). \end{aligned}$$

By Cauchy-Schwarz and since $\|f\| \leq 1$,

$$p \cdot \sum_{i \in [k]} \mathcal{E}_i(f, X) = \langle b_1 + \ldots + b_k, f \rangle \leq \|b_1 + \ldots + b_k\| \cdot \|f\| \leq \|b_1 + \ldots + b_k\|.$$

We now bound $\|b_1 + \ldots + b_k\|$.

First we bound $\langle b_i, b_i \rangle$.

$$\begin{aligned}\langle b_i, b_i \rangle &= \Pr[X_i = z] + \frac{p^2}{(1-p)^2} \Pr[X_i = 1 - z] \\ &= p + \frac{p^2}{1-p}.\end{aligned}$$

Now we bound $\langle b_i, b_j \rangle$ for $i \neq j$ using the pairwise independence of X_i and X_j.

$$\begin{aligned}\langle b_i, b_j \rangle &= a_i a_j \Pr[X_i = X_j = z] + a_i a_j \frac{p^2}{(1-p)^2} \Pr[X_i = X_j = 1 - z] - a_i a_j \frac{p}{1-p} \Pr[X_i \neq X_j] \\ &= a_i a_j \left(p^2 + p^2 - 2p^2\right) \\ &= 0.\end{aligned}$$

Thus,

$$\|b_1 + \ldots + b_k\|^2 \leq 2pk,$$

since $p \leq 1/2$. The lemma follows. ■

We use the previous lemma to get the following bound on the number of players with large effect.

Lemma 3.3 *Let $n \in \mathbb{N}$ and $f : \{0,1\}^n \mapsto [-1,1]$, and consider pairwise independent binary random variables $X_1, \ldots, X_n$ with $\Pr[X_i = 0] = q$, where $0 < q < 1$, for all $i \in [n]$. Then for every positive $\alpha \in \mathbb{R}$,*

$$K(f, X, \alpha) < \frac{2}{p\alpha^2},$$

where $p = \min\{q, 1-q\}$.

Proof: As before, let $k = K(f, X, \alpha)$, and without loss of generality assume that the first k variables are the ones with effect at least α. Using Lemma 3.2,

$$\alpha \cdot k < \sum_{i \in [k]} \mathcal{E}_i(f, X) \leq \sqrt{\frac{2k}{p}},$$

which implies

$$k < \frac{2}{p\alpha^2}.$$

■

3.4 Reducing the General Case to the Boolean Case

We now wish to reduce the general case to the case in which the random variables are binary and have skewed marginals.

Lemma 3.4 *Let $n \in \mathbb{N}$, and let S be a finite set. Let $f : S^n \mapsto [-1,1]$ be some function. Let $X_1, \ldots, X_n$ be n pairwise independent random variables taking values in S, and denote $X = (X_1, \ldots, X_n)$. Let $\alpha > 0$ and $0 \leq p \leq 1$. Then there exist an integer $k \in \mathbb{N}$, a function $g : \{0,1\}^k \mapsto [-1,1]$ and pairwise independent binary random variables $Y_1, \ldots, Y_k$ such that*

- *For every $i \in [k]$, $\Pr[Y_i = 0] = \frac{p}{2}$, and*
- *$K(f, X, p, \alpha) \leq 2 \cdot K(g, Y, \alpha) = 2k$,*

where $Y = (Y_1, \ldots, Y_k)$.

Proof: Let $I \subseteq [n]$ be the set of (p,α)-pivotal players in f, and suppose that for at least $|I|/2$ of $i \in I$,

$$\Pr[\mathrm{E}[f|X_i] - \mathrm{E}[f] > \alpha] > \frac{p}{2}. \tag{1}$$

If this does not hold, then simply consider the function $f' = 1 - f$, for which (1) will hold. Denote by I^+ the set of indices for which (1) holds. Without loss of generality, assume $I^+ = \{1, \ldots, k\}$, and note that $k \geq K(f, X, p, \alpha)/2$.

For every $i \in I^+$, denote

$$p_i = \Pr[\mathrm{E}[f|X_i] - \mathrm{E}[f] > \alpha] > \frac{p}{2}.$$

Define $Y_i = Y_i(x_i)$ as follows:

- If $\mathrm{E}[f|X_i = x_i] - \mathrm{E}[f] > \alpha$, then with probability $\frac{p}{2p_i}$ set $Y_i = 0$ (independently of all other random variables).
- Otherwise, set $Y_i = 1$.

Thus, for every $i \in I^+$, we have $\Pr[Y_i = 0] = \frac{p}{2}$. Furthermore, since $X_1, \ldots, X_n$ are pairwise independent, $Y_1, \ldots, Y_k$ are pairwise independent. All that remains is to define a function g in which all of the Y_i's will have large effect.

To this end, define the function $g : \{0,1\}^k \mapsto [-1,1]$ as

$$g(y) = \mathrm{E}\left[f \middle| Y_1 = y_1, \ldots, Y_k = y_k\right]$$

for all $y = (y_1, \ldots, y_k) \in \{0,1\}^k$.

Now, for every $i \in [k]$ and $z \in \{0,1\}$,

$$\begin{aligned}
\mathrm{E}[f|Y_i = z] &= \sum_y \mathrm{E}[f|Y = y] \cdot \Pr[Y = y|Y_i = z] \\
&= \sum_y g(y) \cdot \Pr[Y = y|Y_i = z] \\
&= \mathrm{E}[g|Y_i = z].
\end{aligned} \tag{2}$$

An additional claim we need is that $\mathrm{E}[f|Y_i = 0] > \mathrm{E}[f] + \alpha$. Denote by

$$T = \{t \in \mathrm{supp}(X_i) : \mathrm{E}[f|X_i = t] > \mathrm{E}[f] + \alpha\}.$$

For any $t \in T$,

$$\begin{aligned}\mathrm{E}[f|Y_i = 0, X_i = t] &= \sum_{x:x_i=t} f(x) \cdot \frac{p}{2p_i} \cdot \frac{\Pr[X = x]}{\Pr[X_i = t, Y_i = 0]} \\ &= \mathrm{E}[f|X_i = t] \cdot \frac{p}{2p_i} \cdot \frac{1}{\Pr[Y_i = 0|X_i = t]} \\ &= \mathrm{E}[f|X_i = t].\end{aligned}$$

Hence,

$$\mathrm{E}[f|Y_i = 0] = \sum_{t \in T} \mathrm{E}[f|X_i = t] \Pr[X_i = t|Y_i = 0] > \mathrm{E}[f] + \alpha.$$

Therefore, since

$$\mathrm{E}[f] = \mathrm{E}[f|Y_i = 0] \Pr[Y_i = 0] + \mathrm{E}[f|Y_i = 1] \Pr[Y_i = 1],$$

it follows that

$$\mathrm{E}[f|Y_i = 1] < \mathrm{E}[f],$$

which implies

$$\mathrm{E}[f|Y_i = 0] - \mathrm{E}[f|Y_i = 1] > \alpha.$$

Thus, using (2), for all $i \in \{1, \ldots, k\}$,

$$\mathrm{E}[g|Y_i = 0] - \mathrm{E}[g|Y_i = 1] = \mathrm{E}[f|Y_i = 0] - \mathrm{E}[f|Y_i = 1] > \alpha.$$

■

3.5 Proof of Theorem 2.6

Proof: By Lemma 3.4, there exists a function g and distribution Y such that $K(f, X, p, \alpha) \leq 2 \cdot K(g, Y, \alpha)$. Since the distribution Y is such that $\Pr[Y_i = 0] = p/2$, and the Y_i's are pairwise independent, Lemma 3.3 implies that $K(g, Y, \alpha) < 2/(p/2)\alpha^2 = 4/p\alpha^2$. Hence,

$$K(f, X, p, \alpha) < \frac{8}{p\alpha^2}.$$

■

3.6 Proof of Theorem 2.7

Proof: We first bound the possible number of *disjoint* pivotal sets. Let $\{C_i\}_{i=1}^t$ be any maximal collection of sets (i.e. t is maximal) satisfying the following:

- $C_i \subseteq [n]$, $|C_i| \leq m$, for all $i \in [t]$.
- For all $i \neq j$, $C_i \cap C_j = \emptyset$.
- For all i, C_i is (p, α)-pivotal.

Since $\{C_i\}_{i=1}^t$ is a maximal collection of such sets, any other (p, α)-pivotal set of size at most m intersects at least one of the C_i's. We now provide an upper bound on the number t of such sets.

To simplify the exposition, suppose all C_i's are of size m, $C_1 = \{1, \ldots, m\}$, $C_2 = \{m+1, \ldots, 2m\}$, and so on. Now consider the function

$$f'(X'_1, \ldots, X'_t, X_{mt+1}, \ldots, X_n) = f(X_1, \ldots, X_n),$$

where $X'_i = X_{C_i}$. We call all players in f' "meta-players". This function takes the same values as f, except that it considers the inputs of all the players in C_i as the input of one meta-player. Note that f' has the same expectation as f.

The variables $X_1, \ldots, X_n$ are $2m$-wise independent, and so $X'_1, \ldots, X'_t, X_{mt+1}, \ldots, X_n$ are pairwise independent. This holds because any meta-player depends on at most m original players, so any 2 meta-players consist of at most $2m$ players that are all independent. Hence, every 2 meta-players are also independent.

By Theorem 2.6, the number of (p, α)-pivotal players in f' is less than $8/p\alpha^2$. Note that if the set of players C_i is (p, α)-pivotal in f, then the meta-player X'_i is (p, α)-pivotal in f'. Thus, we can conclude that $t < 8/p\alpha^2$.

Let $C = \bigcup_{i \in [t]} C_i$. Now consider some set $T \subseteq [n]$ of size $|T| \leq m$. If T is disjoint from C, then T can **not** be (p, α)-pivotal, since the C_i's are a maximal collection of disjoint pivotal sets.

We are now done: $|C| < 8m/p\alpha^2$ as required, and for any $T \subseteq [n] \setminus C$ of size $|T| \leq m$, T is **not** (p, α)-pivotal. ■

4 Proofs: General KKL-Type Results

4.1 Pairwise Independence – Negative Result

We will now present a balanced monotone function and a pairwise independent distribution such that the effects of all players are 0. This will imply that there is no non-trivial lower bound on the effect of a player for pairwise independent distributions, even for monotone functions.

The rough idea of the construction is that since the support of a pairwise independent distribution can be small, monotonicity does not play much of a role. In the next section, however, we show that if the support of the distribution is very small, some player must have large effect. We begin by describing the distribution D used in the counter-example. D will be the convex sum of two other pairwise independent distributions μ and $\overline{\mu}$.

Assume $n+1=2^k$, and identify the set $\{0,\ldots,n\}$ with $\{0,1\}^k$ (by the binary representation). μ will be the uniform distribution over a set of $n+1$ elements in $\{0,1\}^n$. These $n+1$ elements in the support of μ will be denoted by x^z, where $z=(z_1,\ldots,z_k)$ runs over all vectors in $\{0,1\}^k$. Let $y=(y_1,\ldots,y_k)$ be a nonzero element of $\{0,1\}^k$, or, equivalently, an element of $\{1,\ldots,n\}$. Then the y'th index of x^z is

$$x^z_y = \langle z,y\rangle \stackrel{\text{def}}{=} \sum_{i\in[k]} z_i\cdot y_i \mod 2.$$

μ is the quintessential pairwise independent distribution, with marginals $1/2$.

For example, suppose $n=3$ and $k=2$. In this case μ is the uniform distribution over $n+1=4$ elements from $\{0,1\}^3$. These elements are x^z, where $z=(z_1,z_2)\in\{0,1\}^2$. For example, the element $x^{(0,1)}$ is $(1,0,1)$, since $\langle(0,1),(0,1)\rangle=1$, $\langle(0,1),(1,0)\rangle=0$, and $\langle(0,1),(1,1)\rangle=1$. The other elements are $x^{(1,0)}=(0,1,1)$, $x^{(1,1)}=(1,1,0)$, and $x^{(0,0)}=(0,0,0)$. Note that, except for the 0 vector, every element in the support of μ has 2 ones.

In fact, this is true more generally: The support of μ for any k consists of the $(0,\ldots,0)$ vector and n other vectors, each with $(n+1)/2$ ones and $(n-1)/2$ zeros. Moreover, aside from the $(0,\ldots,0)$ vector, all other vectors are incomparable under the natural partial order on $\{0,1\}^n$.

$\overline{\mu}$ will be the uniform distribution on vectors that complement those of μ: for every $x\in\text{supp}(\mu)$, there is an $\overline{x}\in\text{supp}(\overline{\mu})$ such that $\overline{x}=(1,\ldots,1)\oplus x$, where $\oplus$ is the bit-wise XOR. Formally, $\overline{\mu}$ is the uniform distribution over $n+1$ strings

$$\overline{x^z} \stackrel{\text{def}}{=} (1,\ldots,1)\oplus x^z,$$

where $z=(z_1,\ldots,z_k)\in\{0,1\}^k$. The support of $\overline{\mu}$ complements that of μ: it consists of the $(1,\ldots,1)$ vector and n other vectors, each with $(n-1)/2$ ones and $(n+1)/2$ zeros. Since μ is pairwise independent and has marginals $1/2$, so does $\overline{\mu}$.

Set

$$D=\frac{\mu}{2}+\frac{\overline{\mu}}{2}.$$

Since μ and $\overline{\mu}$ are pairwise independent and have marginals $1/2$, so does D. Except for the $(1,\ldots,1)$ and $(0,\ldots,0)$ vectors, none of the vectors in the support of D are comparable (for $n\geq 7$). This means that every function f with $f(1,\ldots,1)=1$ and $f(0,\ldots,0)=0$ is monotone on the support of D.

Define the function $f:\{0,1\}^n\mapsto\{0,1\}$ as follows. For all $x\in\text{supp}(\mu)$, $f(x)=0$. For all $x\in\text{supp}(\overline{\mu})$, $f(x)=1$. Note that regardless of how f is defined on other inputs, f is monotone on the support of D. Furthermore, f is balanced when the inputs are drawn from D. Finally, it is possible to extend f to all of $\{0,1\}^n$ in such a way that f will remain monotone and balanced.

It remains to show that all the effects of players in f with respect to D are 0. Since f is constant on the support of μ, the effects of all players in f with respect to μ are 0. The same is true for $\overline{\mu}$. Using Lemma 4.1 below, we conclude that the effects of all players in D are 0.

Note that for functions that are constant on some distribution, all effects are trivially 0 with respect to that distribution. Such functions, however, are not balanced. The reason f is interesting is that, with respect to D, the effects of all the players are 0 despite f being balanced.

Lemma 4.1 *Let η_1 and η_2 be two distributions on $(X_1, \ldots, X_n) \in \{0,1\}^n$ with marginals $0 < p < 1$, and let $\eta = q\eta_1 + (1-q)\eta_2$, where $0 \leq q \leq 1$. Then for any $g : \{0,1\}^n \mapsto \{0,1\}$,*

$$\begin{aligned}&\mathrm{E}_\eta[g|X_i = 1] - \mathrm{E}_\eta[g|X_i = 0] \\ &\qquad = q\left(\mathrm{E}_{\eta_1}[g|X_i = 1] - \mathrm{E}_{\eta_1}[g|X_i = 0]\right) + (1-q)\left(\mathrm{E}_{\eta_2}[g|X_i = 1] - \mathrm{E}_{\eta_2}[g|X_i = 0]\right).\end{aligned}$$

Proof:

$$\begin{aligned}&\mathrm{E}_\eta[g|X_i = 1] - \mathrm{E}_\eta[g|X_i = 0] \\ &\qquad = \frac{1}{p}\sum_{x:x_i=1}\eta(x)g(x) - \frac{1}{1-p}\sum_{x:x_i=0}\eta(x)g(x) \\ &\qquad = \frac{1}{p}\sum_{x:x_i=1}\Big(q\eta_1(x) + (1-q)\eta_2(x)\Big)g(x) - \frac{1}{1-p}\sum_{x:x_i=0}\Big(q\eta_1(x) + (1-q)\eta_2(x)\Big)g(x) \\ &\qquad = \frac{1}{p}\sum_{x:x_i=1}q\eta_1(x)g(x) - \frac{1}{1-p}\sum_{x:x_i=0}q\eta_1(x)g(x) \\ &\qquad\quad + \frac{1}{p}\sum_{x:x_i=1}(1-q)\eta_2(x)g(x) - \frac{1}{1-p}\sum_{x:x_i=0}(1-q)\eta_2(x)g(x) \\ &\qquad = q\left(\mathrm{E}_{\eta_1}[g|X_i = 1] - \mathrm{E}_{\eta_1}[g|X_i = 0]\right) + (1-q)\left(\mathrm{E}_{\eta_2}[g|X_i = 1] - \mathrm{E}_{\eta_2}[g|X_i = 0]\right),\end{aligned}$$

where $x \in \{0,1\}^n$ for all sums above. ■

4.2 Pairwise Independence – Negative Result for Influence

Section 4.1 deals with KKL-type theorems for effect. In this section we ask whether a KKL-type theorem holds for influence when the distribution is not fully independent. We show that such a theorem does not hold.

Recall Definition 2.5, and note that when μ is not fully independent the vector $x \oplus e_i$ may not even be in the support of μ. Thus, this may not be the "correct" measure of influence for non-independent distributions (which is one of the reasons for considering effect). The original KKL theorem was proved for the case where μ is a fully-independent distribution (see [10]). Understanding the most general scenario in which KKL holds is an interesting open question. Our example from Section 4.1 shows that a KKL-type theorem for effect does not hold under the assumption that μ is pairwise independent. We now show that a KKL-type theorem does not hold for influence either, assuming only pairwise independence (and monotonicity).

Consider the pairwise independent distribution D from Section 4.1. On the support of D, let f be defined as in Section 4.1. Note that if n is large enough, for any $x, y \in \mathrm{supp}(D) \setminus \{(0,\ldots,0),(1,\ldots,1)\}$ and any $i, j \in [n]$, we have that $x \oplus e_i$ is not comparable to either y or $y \oplus e_j$. Thus, f can be extended to a monotone function on all of $\{0,1\}^n$, such that if $x \in \mathrm{supp}(D)$, then for any $i \in [n]$, $f(x \oplus e_i) = f(x)$. This implies that f is a balanced monotone function and D is a pairwise independent distribution such that

$$I_i(f, D) = 0$$

for all $i \in [n]$.

4.3 Pairwise Independence – Positive Result

In this section we prove our one positive result on KKL-type theorems for effect – Theorem 2.8.

Proof: Since μ is pairwise independent and has support of size $n+1$, μ is uniform on its support (see [4]). Thus, for every $x \in \text{supp}(\mu)$, we have $\mu(x) = 2^{-k}$.

We identify the set $\{0, \ldots, n\}$ with $\{0,1\}^k$ (by the binary representation). There is a bijection between $\{0,1\}^k$ and the support of μ: for every $z \in \{0,1\}^k$ fix a corresponding $x^z \in \text{supp}(\mu)$.

Let y be a nonzero element of $\{0,1\}^k$, or, equivalently, an element of $\{1, \ldots, n\}$. Denote by χ_y the map from $\{0,1\}^k$ to $\{1,-1\}$ defined as

$$\forall\ z \in \{0,1\}^k\ \ ,\ \ \chi_y(z) = (-1)^{x^z[y]},$$

where $x^z[y]$ is the y'th index of x^z. Also denote $\chi_{(0,\ldots,0)}$ the map defined as

$$\forall\ z \in \{0,1\}^k\ \ ,\ \ \chi_{(0,\ldots,0)}(z) = 1.$$

We will consider the vector space of maps from $\{0,1\}^k$ to $\mathbb{R}$. We will now show that the set of maps $\{\chi_y\}_{y \in \{0,1\}^k}$ form an orthonormal basis for this vector space with respect to the inner product

$$\langle g, g' \rangle = \sum_{z \in \{0,1\}^k} 2^{-k} g(z) g'(z).$$

For all $y \neq y'$ nonzero elements in $\{0,1\}^k$, we have

$$\begin{aligned}
\langle \chi_y, \chi_{y'} \rangle &= \sum_{z \in \{0,1\}^k} 2^{-k} \chi_y(z) \chi_{y'}(z) \\
&= \sum_{z \in \{0,1\}^k} 2^{-k} (-1)^{x^z[y]} (-1)^{x^z[y']} \\
&= \Pr[X_y = X_{y'}] - \Pr[X_y \neq X_{y'}] \\
&= 0,
\end{aligned}$$

where the last equality follows since μ is pairwise independent with marginals $1/2$ (and we think of y and y' as elements of $[n]$). Furthermore,

$$\langle \chi_y, \chi_{(0,\ldots,0)} \rangle = \sum_{z \in \{0,1\}^k} 2^{-k} \chi_y(z) = \Pr[X_y = 0] - \Pr[X_y = 1] = 0,$$

where the last equality follows since the marginals are $1/2$. Finally, for every $y \in \{0,1\}^k$, we have

$$\langle \chi_y, \chi_y \rangle = \sum_{z \in \{0,1\}^k} 2^{-k} = 1.$$

Thus, the set of maps $\{\chi_y\}_{y \in \{0,1\}^k}$ form an orthonormal basis.

We can also think of f as a map from $\{0,1\}^k$ to $\mathbb{R}$ as follows:

$$\forall\ z \in \{0,1\}^k\ \ ,\ \ \ f(z) = f(x^z).$$

Thus, we can write

$$f = \sum_{y \in \{0,1\}^k} \widehat{f}(y)\chi_y,$$

where

$$\widehat{f}(y) = \langle f, \chi_y \rangle$$

($\widehat{f}(\cdot)$ is called the Fourier transform of f). By Parseval's equality (see c.f. [6]),

$$\sum_{z \in \{0,1\}^k} |f(z)|^2 = 2^k \sum_{y \in \{0,1\}^k} |\widehat{f}(y)|^2.$$

We will now show that $|\widehat{f}(y)|$ is twice the effect of the y'th player. For a nonzero $y \in \{0,1\}^k$, since the marginals are $1/2$,

$$\begin{aligned}
\widehat{f}(y) &= \sum_{z \in \{0,1\}^k} 2^{-k} f(z)(-1)^{x^z[y]} \\
&= \sum_{z : x^z[y]=0} 2^{-k} f(z) - \sum_{z : x^z[y]=1} 2^{-k} f(z) \\
&= 2\left(\mathrm{E}[f|X_y = 0] - \mathrm{E}[f|X_y = 1]\right).
\end{aligned}$$

In addition,

$$\widehat{f}\big((0,\ldots,0)\big) = \sum_{z \in \{0,1\}^k} 2^{-k} f(z) = \mathrm{E}[f].$$

Thus,

$$\begin{aligned}
4 \sum_{i \in [n]} |\mathcal{E}_i(f)|^2 &= \sum_{y \in \{0,1\}^k} |\widehat{f}(y)|^2 - |\widehat{f}\big((0,\ldots,0)\big)|^2 \\
&= 2^{-k} \sum_{z \in \{0,1\}^k} |f(z)|^2 - |\widehat{f}\big((0,\ldots,0)\big)|^2 \\
&= \mathrm{E}[f^2] - (\mathrm{E}[f])^2 \\
&= \mathrm{Var}[f].
\end{aligned}$$

■

References

[1] N. I. Al-Najjar and R. Smorodinsky. Pivotal players and the characterization of influence. *Journal of Economic Theory* **92** (2000): 318–342.

[2] N. I. Al-Najjar and R. Smorodinsky. Large non-anonymous repeated games. *Games and Economic Behavior* **37** (2001): 26–39.

[3] N. I. Al-Najjar and R. Smorodinsky. The efficiency of competitive mechanisms under private information. *Journal of Economic Theory* **137** (2007): 383–403.

[4] I. Benjamini, O. Gurel-Gurevich, and R. Peled. On k-wise independent distributions and boolean functions. Submitted.

[5] M. Ben-Or and N. Linial. Collective coin flipping. In S. Micali (editor), *Randomness and Computation* (1989): 91–115.

[6] J. Doob. Measure Theory. New York: Springer-Verlag, 1994

[7] D. Fudenberg, D. Levine, and W. Pesendorfer. When are non-anonymous players negligible? *Journal of Economic Theory* **79** (1998): 46–71.

[8] R. Gradwohl and O. Reingold. Partial exposure in large games. Submitted.

[9] O. Haggstrom, G. Kalai and E. Mossel. A law of large numbers for weighted majority. *Advances in Applied Mathematics* **37** (2006), no. 1: 112–123.

[10] J. Kahn, G. Kalai and N. Linial. The influence of variables on boolean functions. *Proceedings of the Symposium on the Foundations of Computer Science* (1988): 68–80.

[11] G. Kalai and S. Safra. Threshold phenomena and influence. *Computational Complexity and Statistical Physics*, Oxford University Press, Oxford (2005).

[12] G. J. Mailath and A. Postlewaite. Asymmetric information bargaining problems with many agents, *Review of Economic Studies* **57** (1990): 351–167.

[13] H. Sabourian. Anonymous repeated games with a large number of players and random outcomes. *Journal of Economic Theory* **51** (1990): 92–110.