

Empirical Distribution of Equilibrium Play and Its Testing Application

Yakov Babichenko*, Siddharth Barman†, Ron Peretz‡

Abstract

We show that in an n -player m -action strategic form game, we can obtain an approximate equilibrium by sampling any mixed-action equilibrium a small number of times. We study three notions of equilibrium: Nash, correlated and coarse correlated. For each one of them we obtain upper and lower bounds on the asymptotic (where $\max(m, n) \rightarrow \infty$) worst-case number of samples required for the empirical frequency of the sampled action profiles to form an approximate equilibrium with probability close to one.

These bounds imply that using a small number of samples we can test whether or not players are playing according to an approximate equilibrium, even in games where n and m are large. In addition, our results include a substantial improvement over the previously known upper bounds on the existence of a small-support approximate equilibrium in games with many players. For all three notions of equilibrium, we show the existence of approximate equilibrium with support size polylogarithmic in n and m , whereas the best previously known results were polynomial in n [8, 6, 7].

1 Introduction

Consider a setting in which agents implement an underlying mixed strategy during multiple plays of the same normal-form game. Usually in such a setup, the mixed strategy is not known to an outside observer. Rather, the observer sees the pure actions realized by the agents during the play, i.e., she observes independent and

*Center for the Mathematics of Information, California Institute of Technology. E-mail: babich@caltech.edu.

†Center for the Mathematics of Information, California Institute of Technology. E-mail: barman@caltech.edu.

‡Department of Mathematics, London School of Economics. E-mail: ronprtz@gmail.com

identically distributed samples from the mixed strategy. This framework motivates the following fundamental question. How many samples from an equilibrium (Nash, correlated, or coarse correlated) play are required to ensure that the induced empirical distribution forms an approximate equilibrium (again, Nash, correlated, or coarse correlated)? The main objective of this paper is to show that even in large games (i.e., games with a large number of players and/or actions) an extremely small number of samples (with high probability) generate an approximate equilibrium. This result has several useful interpretations.

1. Testing whether or not players are playing according to an equilibrium. In many strategic settings, it is important to test whether players are playing according to an equilibrium or not, but experimental data is limited and costly. Many examples of such scenarios can be found in experimental economics. In such contexts it is desirable to have tests that are reliable and require a small number of data points. Another case wherein this testing exercise is relevant is when the same game is played multiple times in independent environments. We observe a limited number of outcomes/data and our goal is to analyze, through the data, whether the agents are implementing an equilibrium or not. Our results (Theorems 2,5 and Corollary 4) show we can accomplish this testing goal even with a small dataset (i.e., few samples) that consists of i.i.d. action profiles drawn from the underlying mixed strategy. Moreover the results show that the test can be performed via a direct algorithm: We simply need to check whether or not the empirical distribution of the observed data is an approximate equilibrium.

2. Existence of simple approximate equilibrium. The existence of approximate equilibrium with support size polynomial in the number of players has been established in prior work (see [1], [12], and [8]). In particular, these results show that in every normal-form game there exists an approximate equilibrium that is simply the uniform distribution over a small set of action profiles. This implies that even if the game is large it still contains a simple approximate equilibrium. Our result extends this line of work by substantially improving the bounds on the support size of such simple approximate equilibrium. In particular, Corollaries 1 and 3 along with Theorem 4 show that in every game there exists an approximate equilibrium of support size polylogarithmic in n (the number of players) and m (the number of actions per player), see Table 2.

3. Short-Run Stability of Equilibrium in Repeated Games with Bounded Rationality. The premise that players do not know their opponents' utility functions is a central construct in the study of *uncoupled* learning in repeated games (see [7] and [16]). A reasonable assumption along these lines is that players do not know the mixed strategy of their opponents. This gives us a repeated-game model in which

every player learns her opponents' mixed strategies by observing the actions they played in each repetition of the game.¹ In this framework, a natural way of learning opponents' mixed strategy is through the *empirical distribution* (i.e., by approximating opponents' mixed strategy by the empirical distribution of observed action profiles). Such a learning process is considered in the classical fictitious play [15] and in many other recent results, e.g. the regret-testing dynamics (see [4] and [5]).

Overall, we get the following fundamental question in settings where players learn their opponents' mixed strategy through pure-action realizations of the opponents: How fast does the empirical distribution of equilibrium play form an approximate equilibrium? Say after multiple iterations the empirical distribution does not form an approximate equilibrium, then it is likely that an impatient player (who is uncertain about her opponents' strategies) will infer from the observed samples that her current mixed strategy is not a best reply to her opponents' strategies and deviate to some other strategy (this exact setup has been considered in [4] and [5]). Our results imply that (with high probability) even in large games, and even for small number of iterations, such a situation will not occur, i.e., the equilibria are stable in the short run. In particular, we provide almost tight bounds for the above question (see Table 1). The testing results of this paper shows that after small number of iterations, each player can learn whether his current strategy is approximately a best reply to opponents' strategies.

4. Population games Sampling occurs in life very naturally. In biology [13], members of a specie come in different types: sex, size, color, etc. Every newborn has a type which is sampled according to some distribution prescribed by its specie. An ecological system can be modeled as a game between species. When individuals roam the world their fitness (payoff) depends on how their own type interacts with the types of the other individuals; actually, it is only the empirical frequency of the others' types that matters. The more complex situation is modeled by population games (see [10], for example). The abstract model describes the relations between the species as a game in strategic form. It is understood that the model is only an abstraction for a real game being played between individuals. Evolutionary game theory studies the dynamics of the system, the forces that might drive the system into (or away from) an equilibrium. We consider here a different point of view that can result as an interpretation of our work. Assuming the system is already in equilibrium and one might ask what is the minimal number of individuals (the size of the ecosystem) needed in order for it to persist. In other words, take a stable ecosystem and let it shrink until there are so few individuals that the equilibrium

¹Note that this assumption holds in most repeated-game models (i.e., perfect monitoring)

cannot persist. What is that critical point at which the equilibrium is disrupted? Another question we address is how many samples one must take in order to find an (approximate) equilibrium assuming that the system is in equilibrium but not knowing what it is. In real life we know neither the payoff function nor the true strategies of the species. We only observe a few samples of each strategy. We provide bounds on the number of samples one needs to consider in order to understand the equilibrium. The bounds depend only on the number of species and the number of types, but not on the payoff function or the true strategy of the species.

We say that a mixed strategy is k -uniform if it is the uniform distribution over no more than k action profiles. From a game theoretic perspective, a k -uniform strategy is just a pure strategy in a population game with population of size k . Population games are a special case of semi-anonymous games for which Kalai [11] showed there exists a pure equilibrium, when k is large enough. The minimal number of players needed for such an equilibrium to emerge was studied by Azrieli and Shmaya in the even more general framework of Lipschitz games [2]. In this respect the study of k -uniform equilibrium in strategic form games can be viewed as a special pure equilibrium in special class of Lipschitz games. It should be noted that the results we obtain for that special case of population games are stronger than what one could hope to obtain for any Lipschitz game.

1.1 Informal Statement of the Results

We consider large normal-form games with n players and m actions per player (here at least one of the numbers n or m is large). Let x be an equilibrium of the game (Nash, correlated, or coarse correlated), which is a distribution over the action profiles. We observe k i.i.d. samples from x .

For the case of Nash equilibrium, since we know that the players are playing according to a product distribution, the correct notion for the *empirical distribution of play* is the product distribution where each player plays the empirical distribution of the her own realizations. For the cases of correlated and coarse correlated equilibrium the correct notion for the *empirical distribution of play* is simply the empirical distribution of the realizations.

The empirical distribution of play is a random variable, in this paper we pose the question: for which values of k will the empirical distribution of play be an approximate equilibrium with probability close to 1? We provide almost tight bounds for the this question, which are summarized in the table below.

Moreover, the opposite is also true. If players are playing according to a distribution that is *not* an approximate equilibrium then for the same values of k (as

Equilibrium	Upper Bound	Lower bound
Nash	$k \leq O(\log m + \log n)$ Theorem 1	$k \geq \Omega(\log m + \log n)$ Examples 2 and 1
Correlated	$k \leq O(m \log m + \log n)$ Theorem 3	$k \geq \Omega(m + \log n)$ Examples 3 and 1
Coarse Correlated	$k \leq O(\log m + \log n)$ Theorem 7	$k \geq \Omega(\log m + \log n)$ Examples 2 and 1

Table 1: Bounds on the number of samples that are required for the empirical distribution of play to form an approximate equilibrium with probability close to 1.

in the upper bound column in Table 1), with probability close to 1, the empirical distribution of play will not form an approximate equilibrium. Therefore, we can test whether players are playing according to an approximate equilibrium using k samples (see Theorems 2 and Corollaries 5 and 4). These results suggests that even in games with very large number of players or very large number of actions there exists efficient tests to determine whether players are playing according to Nash equilibrium or coarse correlated equilibrium. Correlated equilibrium on the other hand is a slightly more complicated notion in this respect. We accomplish this by proving that there does exist a test for approximate correlated equilibrium that uses less than $\Omega(\sqrt{m})$ samples (see Theorem 6). This is in contrast to Nash equilibrium and coarse-correlated equilibrium that require only $O(\log m)$ samples.

Actually, even the fact that the empirical distribution of play forms an approximate equilibrium with *positive* probability is interesting in its own right. The fact that it occurs with positive probability proves existence of an approximate equilibrium with small *support size*. The support size has different meanings in the case of Nash equilibrium and the cases of correlated and coarse correlated equilibrium. For Nash equilibrium the support size is the maximum number of actions that any single player uses (where the maximum is taken over all players). For correlated and coarse correlated equilibrium, the support size is the number of *action profiles* that are used in the approximate equilibrium.

Small-support approximate Nash equilibrium has been previously studied in [1], [12], and [8]. Althofer [1] studied the problem for two-player m -action zero-sum games and established an $O(\log m)$ bound on the size of the support. Lipton, Markakis, and Mehta [12] studied the same question of general n -player m -actions games and achieved a bound of $O(n^2 \log m)$. Hemon, Rougemont, and Santha [8] improved the bound of [12] to $O(n \log m)$. Our result implies existence of approx-

imate Nash equilibrium with support size $O(\log m + \log n)$ (see Corollary 1). This result gives us an algorithm for computing approximate Nash equilibrium with running time $N^{\log \log N}$ where N is the input size (the size of the game), for games where $m = \text{poly}(n)$ (i.e., the number of actions is not significantly larger than the number of players), see Corollary 2. To the best of our knowledge the previously known best algorithm for this problem has a running time of $O(N^{\log N})$ (see [14]).

Small-support correlated equilibrium was studied by Germano and Lugosi in [6] for *exact* correlated equilibrium, where they achieved a bound of $O(nm^2)$ on the size of the support. Applying the technique from [6] we can obtain a bound of $O(nm)$ on the size of the support of exact coarse correlated equilibrium. The results in this paper prove that for *approximate* correlated equilibrium and *approximate* coarse correlated equilibrium the size of the support can be significantly reduced to $O(\log m(\log m + \log n))$ and $O(\log m + \log n)$ respectively.

Our bounds on the support size of approximate equilibria relative to the previously known results are summarized in following table.

Table 2: Support size of approximate equilibria.

Equilibrium	Our Results	Previous Bounds
Nash	$O(\log m + \log n)$ Corollary 1	$O(n \log m)$ [8]
Correlated	$O(\log m(\log m + \log n))$ Theorem 4	Exact: $O(nm^2)$ [6]
Coarse Correlated	$O(\log m + \log n)$ Corollary 3	Exact: $O(nm)$ [6]

2 Notations

We consider n -player games with m actions per player,² such a game will be called *n-player m-action game*. We use the following standard notation. The set of players is $[n] = \{1, 2, \dots, n\}$. The set of actions of each player is $A_i = [m] = \{1, 2, \dots, m\}$. The set of strategy profiles is $A = [m]^n$. The set of probability distributions over a set B is denoted by $\Delta(B)$. Therefore, $\Delta(A)$ is the set of probability distributions over the action profiles, and $\prod_{i \in [n]} \Delta(A_i)$ is the set of *product distributions*. For a

²All the results in the paper hold also for the case where each player has a different number of actions (i.e., player i has m_i actions). For simplicity of notations only, we assume throughout that all players have the same number of actions m .

vector $v = (v_j)_{j \in [n]}$ we denote by $v_{-i} := (v_j)_{j \neq i, j \in [n]}$ the vector that does not contain the i 'th coordinate. The payoff function of player i is $u_i : A \rightarrow [0, 1]$ and it can be extended to $u_i : \Delta(A) \rightarrow [0, 1]$ by $u_i(x) := \mathbb{E}_{a \sim x}[u_i(a)]$.

Definition 1. A distribution over B will be called k -uniform if it is the uniform distribution over a size- k multiset of elements from B . Equivalently, $x \in \Delta(B)$ is k -uniform iff $x(b) = \frac{c_b}{k}$ for every $b \in B$ where $c_b \in \mathbb{N}$.

3 Approximate Nash Equilibrium

Definition 2. A product distribution $x = (x_i)_{i \in [n]}$ is an ε -Nash equilibrium if no player can gain more than ε by deviating to another strategy. Formally, $u_i(x) \geq u_i(a_i, x_{-i}) - \varepsilon$ for every $i \in [n]$ and every $a_i \in A_i$.

In case of $\varepsilon = 0$ we will say that x is an *exact Nash equilibrium*, or simply *Nash equilibrium* for short.

Throughout the paper we will consider ε to be a *constant*, and we will derive asymptotic results for games where at least one of the parameters m or n goes to infinity.

Assume that players are playing according to a product distribution $x = (x_i)_{i \in [n]}$. We observe k i.i.d. *samples* from x that are denoted by $(a(t))_{t \in [k]}$ where $a(t) \in A$. Since we assumed that the players are playing according to a product distribution, the correct interpretation of the observed data is as follows. We denote by s_i^k the *empirical distribution of player i* defined to be the empirical distribution of the samples $(a_i(t))_{t \in [k]}$. Namely, $s_i^k(a_i) = \frac{1}{k} |\{t : a_i(t) = a_i\}|$. The *product empirical distribution of play* is the product distribution $\prod_i s_i^k$.

The following theorem states that if players are playing according to Nash equilibrium then the product empirical distribution of play (which is a random variable) is an ε -Nash equilibrium after $k = O(\log n + \log m)$ samples, with probability close to 1.

Theorem 1. For every $\varepsilon, \alpha > 0$, every n -player m -action game and every Nash equilibrium of the game x , the product empirical distribution of play $(s_i^k)_{i \in [n]}$ is an ε -Nash equilibrium with probability greater than $1 - \alpha$ for

$$k > \frac{8(\ln m + \ln n - \ln \alpha - \ln \varepsilon + \ln 8)}{\varepsilon^2} = O(\log m + \log n).$$

We emphasize the *logarithmic* dependence of the number of samples k on the probability of error α , which means that in order to reduce the probability of error

by a factor of two we should increase the number of samples only by a constant $(\frac{8 \ln 2}{\varepsilon^2})$.

A proof of Theorem 1 is given Section 3.3.

The bound $O(\log m + \log n)$ is tight (up to a constant factor), see Examples 1 and 2 in Section 6.

3.1 Existence of Simple Approximate Nash Equilibrium

Note that if s^k is an ε -Nash equilibrium with positive probability then *there exists* a sequence of samples of size k that forms an ε -Nash equilibrium (the probabilistic method). Note also that the empirical distribution of every player i is a k -uniform distribution (see Definition 1). This simple observation implies the following corollary from Theorem 1.

Corollary 1. For every ε , every n -player m -action admits an ε -Nash equilibrium where each player plays a k -uniform mixed action for

$$k > \frac{8}{\varepsilon^2}(\ln m + \ln n - \ln \varepsilon + \ln 8) = O(\log m + \log n).$$

One consequence of this corollary is existence of an approximate equilibrium where each player uses only small number of actions in his own mixed strategy (at most $O(\log m + \log n)$). Another consequence, which might also be useful, is the simplicity of the probabilistic structure of the mixed strategy of each player. To see it, consider for example the case of n -players 2-actions games. In terms of the support size Corollary 1 is meaningful. But fix

$$k = \left\lceil \frac{8}{\varepsilon^2}(\ln n + \ln \varepsilon + \ln 16) \right\rceil = O(\log n),$$

then Corollary 1 states that there exists an ε -Nash equilibrium where each player i uses mixed strategy of the form $(\frac{c_i}{k}, 1 - \frac{c_i}{k})$ where $c_i \in \mathbb{N}$.

The fact that such simple approximate Nash equilibrium exists allows us to find approximate Nash equilibrium simply by exhaustively search over all the possible n -tuples of k -uniform strategy. Although the algorithm is naive, to the best of our knowledge this is the best know upper bound for this problem.

Corollary 2. There exists an algorithm for computing ε -Nash equilibrium in every n -players m -actions game with running time m^{nk} for $k = O(\log m + \log n)$.

The running time of an algorithm is usually compared to the input size. In n -players m -actions games the input size is $N = nm^n$.

For *all* games we have

$$m^{nk} = \text{poly}(N^k) \leq \text{poly}(N^{n \log m}) = \text{poly}(N^{\log N})$$

which implies that the running time of the exhaustive search algorithm is at most $N^{\log N}$. This bound coincides with the best known upper bound for computing approximate Nash equilibrium (see [14]).

For the class of games where $m = \text{poly}(n)$ (e.g., n -player games with constant number of actions or n -players n^2 -actions games) the bound of Corollary 2 improves the bound of $N^{\log N}$ to $N^{\log \log N}$:

$$m^{nk} = \text{poly}(N^k) = \text{poly}(N^{\log n}) = \text{poly}(N^{\log \log N}).$$

3.2 Testing Approximate Nash Equilibrium Play

We consider the following settings. We want to test whether players are playing according to an approximate Nash equilibrium or not. We assume that we cannot observe the exact mixed strategies of the players. We can observe samples of the mixed strategies. We focus on the question, how many samples are required to perform this test?

Ideally, we would like to design a test that outputs the answer **YES**, with probability close to 1, if the players are playing according to δ -Nash equilibrium, and it returns **NO**, with probability close to 1, otherwise. It is easy to see that such a test does not exist. The problem arises at the “border”. Consider a distribution x that is a δ -Nash equilibrium, but every arbitrary small neighborhood of x contains a distribution that is not a δ -Nash equilibrium (it is easy to see that such a distribution x always exists). Then the test should distinguish between x and distributions that are arbitrary close to x , which is impossible. Therefore, we weaken our requirements from a test.

Definition 3. A function $T : A^k \rightarrow \{\mathbf{YES}, \mathbf{NO}\}$ is said to be an ε -test that uses k samples and has error probability α for δ -Nash equilibrium if for every product distribution $x = (x_i)_{i \in [n]}$ we have

- $\mathbb{P}(T((a(t))_{t \in [k]}) = \mathbf{YES}) \geq 1 - \alpha$ for every x that is a δ -Nash equilibrium.
- $\mathbb{P}(T((a(t))_{t \in [k]}) = \mathbf{NO}) \geq 1 - \alpha$ for every x that is not a $(\delta + \varepsilon)$ -Nash equilibrium.

In other words, we require that the test returns the correct answer, with probability close to 1, for all distributions that are δ -Nash equilibrium and for all distributions that are not $(\delta + \varepsilon)$ -Nash equilibrium. We allow the test to return any

answer for the case where the distribution is a $(\delta + \varepsilon)$ -Nash equilibrium but not a δ -Nash equilibrium.

The following theorem states that using $O(\log n + \log m)$ samples we can test whether players are playing according to an approximate Nash equilibrium. Moreover, the test is very natural, we need to simply check whether the empirical distribution of play is an approximate Nash equilibrium or not.

Theorem 2. Let

$$T((a(t))_{t \in [k]}) = \begin{cases} \mathbf{YES} & \text{if } (s_i^k)_{i \in [n]} \text{ is a } (\delta + \frac{\varepsilon}{2})\text{-Nash equilibrium} \\ \mathbf{NO} & \text{otherwise.} \end{cases}$$

Then T is an ε -test that has error probability α for δ -Nash equilibrium with

$$k > \frac{72}{\varepsilon^2} (\ln(m+1) + \ln n - \ln \alpha - \ln \varepsilon + \ln 24) = O(\log m + \log n)$$

Note that the number of samples is independent of δ . Section 3.3 contains a proof of the theorem.

3.3 Proofs

The proofs of Theorems 1 and 2 are based on the following Proposition.

Proposition 1. For every n -player m -action game, every player $i \in [n]$, every action $a_i \in A_i = [m]$, and every product distribution of the opponents $x_{-i} = (x_j)_{j \neq i}$ we have

$$\mathbb{P}(|u_i(a_i, s_{-i}^k) - u_i(a_i, x_{-i})| \geq \varepsilon) \leq \frac{4e^{-\frac{\varepsilon^2}{2}k}}{\varepsilon}.$$

In other words, this proposition states that with probability that is exponentially (in k) close to 1, player i is almost indifferent between the case where his opponents are playing the original distribution x_{-i} or the product empirical distribution s_{-i}^k .

We emphasize that this Proposition is the main technical jump relative to the related literature [12] and [8]. Beyond this proposition, our techniques will be similar to [12] and [8].

Proof of Proposition 1. Assume without loss of generality that $i = 1$ and $a_i = 1$. We begin by rewriting the payoff of player 1. For every $l \in [k]$, we can write

$$u_1(1, s_{-1}^k) = \frac{1}{k^{n-1}} \sum_{j_1, j_2, \dots, j_n \in [k]} u_1(1, a_2(j_2 + l), a_3(j_3 + l), \dots, a_n(j_n + l))$$

where the indexes $j_i + l$ are taken modulo k . If we take the average over all possible l we have

$$u_1(1, s_{-1}^k) = \frac{1}{k^{n-1}} \sum_{j_1, j_2, \dots, j_n \in [k]} \frac{1}{k} \sum_{l \in [k]} u_1(1, a_2(j_2 + l), a_3(j_3 + l), \dots, a_n(j_n + l)). \quad (1)$$

For every initial profile of indices $j_* = (j_2, j_3, \dots, j_n) \in [k]^{n-1}$ and every $l \in [k]$, we denote $a_{-1}(j_* + l) := (a_2(j_2 + l), a_3(j_3 + l), \dots, a_n(j_n + l)) \in A_{-1}$, and we define the random variable

$$d(j_*) := \begin{cases} 0 & \text{if } \left| \frac{1}{k} \sum_{l \in [k]} u_1(1, a_{-1}(j_* + l)) - u_1(1, x_{-1}) \right| \leq \frac{\varepsilon}{2} \\ 1 & \text{otherwise.} \end{cases} \quad (2)$$

By the definition of $d(j_*)$, we have

$$d(j_*) + \frac{\varepsilon}{2} \geq \left| \frac{1}{k} \sum_{l \in [k]} u_1(1, a_{-1}(j_* + l)) - u_1(1, x_{-1}) \right|. \quad (3)$$

Note also that for any fixed j_* the random action profiles $a_{-1}(j_* + 1), a_{-1}(j_* + 2), \dots, a_{-1}(j_* + k)$ are independent. Therefore by Hoeffding's inequality (see [9]) we have

$$\mathbb{E}[d(j_*)] \leq 2e^{-\frac{\varepsilon^2}{2}k}. \quad (4)$$

Using representation (1) of the payoffs and inequalities (3) and (4), we get

$$\begin{aligned} & \mathbb{P}(|u_i(1, s_{-1}^k) - u_i(1, x_{-1})| \geq \varepsilon) = \\ & = \mathbb{P} \left(\left| \frac{1}{k^{n-1}} \sum_{j_* \in [k]^{n-1}} \frac{1}{k} \sum_{l \in [k]} u_1(1, a_{-1}(j_* + l)) - u_1(1, x_{-1}) \right| \geq \varepsilon \right) \\ & \leq \mathbb{P} \left(\frac{1}{k^{n-1}} \sum_{j_* \in [k]^{n-1}} \left| \frac{1}{k} \sum_{l \in [k]} u_1(1, a_{-1}(j_* + l)) - u_1(1, x_{-1}) \right| \geq \varepsilon \right) \\ & \leq \mathbb{P} \left(\frac{1}{k^{n-1}} \sum_{j_* \in [k]^{n-1}} d(j_*) \geq \frac{\varepsilon}{2} \right) \leq \frac{4e^{-\frac{\varepsilon^2}{2}k}}{\varepsilon} \end{aligned} \quad (5)$$

where the last inequality follows from Markov's inequality. □

Proof of Theorem 1. The proof uses similar idea to [12] or [8]. Proposition 1 and the choice of k guarantees that

$$\mathbb{P}(|u_i(a_i, s_{-i}^k) - u_i(a_i, x_{-i})| \geq \frac{\varepsilon}{2}) \leq \frac{8e^{-\frac{\varepsilon^2}{8}k}}{\varepsilon} < \frac{\alpha}{mn}$$

for every player i and every action $a_i \in [m]$. Using the union bound, we get that with probability greater than $1 - \alpha$ we have $|u_i(a_i, s_{-i}^k) - u_i(a_i, x_{-i})| < \frac{\varepsilon}{2}$ for *all* players $i \in [n]$ and *all* actions $a_i \in [m]$. In such a case $(s_i^k)_{i \in [n]}$ is an ε -Nash equilibrium because:

$$\begin{aligned} u_i(a_i, s_{-i}^k) &\leq u_i(a_i, x_{-i}) + \frac{\varepsilon}{2} \leq \sum_{a'_i \in A_i} s_i^k(a'_i) u_i(a'_i, x_{-i}) + \frac{\varepsilon}{2} \\ &\leq \sum_{a'_i \in A_i} s_i^k(a'_i) u_i(a'_i, s_{-i}^k) + \varepsilon = u_i(s_i^k, s_{-i}^k) + \varepsilon, \end{aligned}$$

where the second inequality holds because all the strategies in the support of s_i^k are in the support of x_i , which contains only best replies to x_{-i} . \square

Proof of Theorem 2. Proposition 1 and the choice of k guarantees that

$$\mathbb{P}(|u_i(a_i, s_{-i}^k) - u_i(a_i, x_{-i})| \geq \frac{\varepsilon}{3}) \leq \frac{24e^{-\frac{\varepsilon^2}{72}k}}{\varepsilon} < \frac{\alpha}{(m+1)n} \quad (6)$$

for every player i and every action $a_i \in [m]$. In addition, Hoeffding's inequality (see [9]) guarantees that for a given x_{-i} we have

$$\mathbb{P}(|u_i(s_i^k, x_{-i}) - u_i(x_i, x_{-i})| \geq \frac{\varepsilon}{6}) \leq 2e^{-\frac{\varepsilon^2}{18}k} < \frac{24e^{-\frac{\varepsilon^2}{72}k}}{\varepsilon} < \frac{\alpha}{(m+1)n} \quad (7)$$

for every player $i \in [n]$. Note that there are $m(n+1)$ inequality of the form (6) and (7). Therefore, the union bound implies that with probability greater than $1 - \alpha$ the following $m(n+1)$ inequalities hold:

$$\begin{aligned} |u_i(a_i, s_{-i}^k) - u_i(a_i, x_{-i})| &\leq \frac{\varepsilon}{6} \quad \forall i \in [n], \forall a_i \in [m]. \\ |u_i(s_i^k, x_{-i}) - u_i(x_i, x_{-i})| &\leq \frac{\varepsilon}{6} \quad \forall i \in [n]. \end{aligned} \quad (8)$$

Throughout the proof we will assume that all the inequalities in (8) are satisfied.

If $(x_i)_{i \in [n]}$ is a δ -Nash equilibrium then $(s_i^k)_{i \in [n]}$ is a $(\delta + \frac{\varepsilon}{2})$ -Nash equilibrium because

$$\begin{aligned} u_i(a_i, s_{-i}^k) &\leq u_i(a_i, x_{-i}) + \frac{\varepsilon}{6} \leq u_i(x_i, x_{-i}) + \delta + \frac{\varepsilon}{6} \leq u_i(s_i^k, x_{-i}) + \delta + \frac{\varepsilon}{3} \\ &= \sum_{a_i \in A_i} s_i^k(a_i) u_i(a_i, x_{-i}) + \delta + \frac{\varepsilon}{3} \leq \sum_{a_i \in A_i} s_i^k(a_i) u_i(a_i, s_{-i}^k) + \delta + \frac{\varepsilon}{2} = u_i(s_i^k, s_{-i}^k) + \delta + \frac{\varepsilon}{2}. \end{aligned}$$

On the other hand, if $(x_i)_{i \in [n]}$ is not a $(\delta + \varepsilon)$ -Nash equilibrium, then there exists a player i and an action a_i^* such that $u_i(a_i^*, x_{-i}) > u_i(x_i, x_{-i}) + \delta + \varepsilon$. In such a case $(s_i^k)_{i \in [n]}$ is not a $(\delta + \frac{\varepsilon}{2})$ -Nash equilibrium because

$$\begin{aligned} u_i(a_i^*, s_{-i}^k) &\geq u_i(a_i^*, x_{-i}) - \frac{\varepsilon}{6} > u_i(x_i, x_{-i}) + \delta + \frac{5\varepsilon}{6} \geq u_i(s_i^k, x_{-i}) + \delta + \frac{4\varepsilon}{6} \\ &= \sum_{a_i \in A_i} s_i^k(a_i) u_i(a_i, x_{-i}) + \delta + \frac{2\varepsilon}{3} \geq \sum_{a_i \in A_i} s_i^k(a_i) u_i(a_i, s_{-i}^k) + \delta + \frac{\varepsilon}{2} = u_i(s_i^k, s_{-i}^k) + \delta + \frac{\varepsilon}{2}. \end{aligned}$$

Summarizing, by the choice of k we guarantee that all the inequalities in (8) will be satisfied with probability of at least $1 - \alpha$. If those inequalities are satisfied then we have the following:

- For every product distribution x that is a δ -Nash equilibrium, the product empirical distribution is a $(\delta + \frac{\varepsilon}{2})$ -Nash equilibrium. Hence, for δ -Nash equilibria, the given test T returns the correct answer **YES**.

- For every product distribution x that is not a $(\delta + \varepsilon)$ -Nash equilibrium, the product empirical distribution is not a $(\delta + \frac{\varepsilon}{2})$ -Nash equilibrium. Hence, for a distribution that is not a $(\delta + \varepsilon)$ -Nash equilibrium, the given test T returns the correct answer **NO**. \square

4 Approximate Correlated Equilibrium

Section 3 considered the case of product distributions, and hence the players do not have a correlation device. In scenarios where players can correlate their actions the appropriate notions of equilibria are correlated equilibria (considered in this section) and coarse correlated equilibria (considered in the next section).

The high level idea behind the definition of correlated equilibrium is as follows. There exists a *mediator* who samples an action profile $a = (a_i)_{i \in [n]}$ according to a distribution x . Then the mediator (privately) tells to every player i the corresponding action a_i . We will call the drawn action a_i *the recommendation to player i* . A distribution $x \in \Delta(A)$ is an ε -*correlated equilibrium* if no player can gain more than ε by deviating from the recommendation of the mediator. A deviation from mediator's recommendation is described by a *switching rule* $f : A_i \rightarrow A_i$, that corresponds to the case where instead of the recommended action a_i the player chooses to play $f(a_i)$.

Definition 4. For every switching rule $f : A_i \rightarrow A_i$ we denote by $R_f^i(a) := u_i(f(a_i), a_{-i}) - u_i(a_i, a_{-i})$ the regret of player i for not implementing the switching rule f at strategy profile a .

A distribution $x \in \Delta(A)$ is an ε -correlated equilibrium if $\mathbb{E}_{a \sim x}[R_f^i(a)] \leq \varepsilon$ for every player i and every switching rule $f : A_i \rightarrow A_i$.

Unlike the case of product distributions where it was reasonable to consider the *product empirical distribution* (i.e., the product of each player individual empirical distribution), here in the case of general (not necessarily product) distributions we consider the empirical distribution of the sampled profiles. We assume that players are playing according to a distribution $x \in \Delta(A)$. We observe k i.i.d. *samples* of x which are denoted by $(a(t))_{t \in [k]}$ where $a(t) \in A$. We denote by s^k the *empirical distribution of the samples*, specifically $s^k(a) := \frac{1}{k} |\{t \in [k] : a(t) = a\}|$.

We begin with stating the analogue of Theorem 1 for the case of correlated equilibrium.

Theorem 3. For every $\varepsilon, \alpha > 0$, every n -player m -action game and every correlated equilibrium of the game x , the empirical distribution of the samples s^k is an ε -correlated equilibrium with probability greater than $1 - \alpha$ for

$$k > \frac{2}{\varepsilon^2} (m \ln m + \ln n - \ln \alpha) = O(m \log m + \log n).$$

The bounds of Theorem 3 are almost tight. Namely, Example 1 demonstrate that the $\log n$ dependence on n is tight, and Example 3 demonstrates that at least $\Omega(m)$ samples are required in order to form an approximate correlated equilibrium.

The arguments for proving Theorem 3 are more direct than the Nash-equilibrium case (i.e., Theorem 1). A proof of Theorem 3 is given in Section 4.3.

4.1 Existence of Simple Approximate Correlated Equilibrium

We should emphasize again that the support of correlated equilibrium is the number of *action profiles* in the support of the equilibrium. If we use the existence of small-support approximate Nash equilibrium (which is also an approximate correlated equilibrium) we obtain the existence of approximate correlated equilibrium with support of size $O(\log m + \log n)^n$.

By observing that Theorem 3 holds with positive probability we can deduce the existence of approximate correlated equilibrium with support of size $O(m \log m + \log n)$. But can the support size of an approximate correlated equilibrium have a poly-logarithmic dependence on m , instead of a polynomial one? Example 3 demonstrates that if we will sample from an *arbitrary* correlated equilibrium then we cannot. But, if we will sample from a *specific* approximate correlated equilibrium

than poly-logarithmic number of samples are sufficient. It turns out that the specific approximate correlated equilibrium from which we should sample is an equilibrium in which *each player* uses only a small number of her own actions in the support of the equilibrium. Existence of such an approximate correlated equilibrium is proved in Corollary 1 (because every approximate Nash equilibrium is also an approximate correlated equilibrium).

The following theorem shows that there always exists an approximate correlated equilibrium with poly-logarithmic dependence on n and m , moreover the probabilistic structure of the equilibrium is simple: it is a k -uniform distribution (see Definition 1).

Theorem 4. Every n -player m -action game admits a k -uniform ε -correlated equilibrium for every

$$k > \frac{264}{\varepsilon^4} \ln m (\ln m + \ln n - \ln \varepsilon + \ln 16) = O(\log m (\log m + \log n)) \quad (9)$$

A proof of the theorem, which uses the above mentioned ideas, appears in Section 4.3.

4.2 Testing Approximate Correlated Equilibrium Play

Similar to Section 3.2, we would like to design a test that uses k samples to determine whether players are playing according to a δ -correlated equilibrium or according to a distribution that is not a $(\delta + \varepsilon)$ -correlated equilibrium.³

Definition 5. An ε -test with error probability α for δ -correlated equilibrium that uses k -samples, is a function $T : A^k \rightarrow \{\mathbf{YES}, \mathbf{NO}\}$, such that for every distribution x we have

- $\mathbb{P}(T((a(t))_{t \in [k]}) = \mathbf{YES}) \geq 1 - \alpha$ for every x that is a δ -correlated equilibrium.
- $\mathbb{P}(T((a(t))_{t \in [k]}) = \mathbf{NO}) \geq 1 - \alpha$ for every x that is not a $(\delta + \varepsilon)$ -correlated equilibrium.

The following theorem states that using $O(m \log m + \log n)$ samples we can test whether players are playing according to an approximate correlated equilibrium. Moreover, the test is very natural, we should simply check whether the empirical distribution of play is an approximate correlated equilibrium or not.

³We note again, that it is impossible to design such a test for the case of “ δ -correlated equilibrium or not a δ -correlated equilibrium”, by arguments similar to the ones in Section 3.2.

Theorem 5. Let

$$T((a(t))_{t \in [k]}) = \begin{cases} \mathbf{YES} & \text{if } s^k \text{ is a } (\delta + \frac{\varepsilon}{2})\text{-correlated equilibrium} \\ \mathbf{NO} & \text{otherwise.} \end{cases}$$

Then T is an ε -test with an α -error-probability for δ -correlated equilibrium for

$$\frac{8}{\varepsilon^2}(m \ln m + \ln n - \ln \alpha) = O(m \log m + \log n).$$

Section 4.3 contains a proof of this theorem.

An unsatisfactory property of the above test is the polynomial dependence on the number of actions. Example 3 demonstrate that the natural test that is presented in the theorem cannot use less then $\Omega(m)$ samples. Hypothetically, it could be the case that there exists some other test that uses significantly fewer samples. The following theorem states that this is not the case. The number of samples must be polynomial in m , even for the case where ε and α are constants.

Theorem 6. Every $\frac{1}{2}$ -test with an error probability $\frac{1}{4}$ for exact correlated equilibrium for two-player m -action games must use at least $\sqrt{\frac{m}{2}}$ samples.

See Section 4.3 for a proof.

4.3 Proofs

Proof of Theorem 3. Note that $R_f^i(a)$ where $x \sim a$ is a random variable that assumes values in $[-1, 1]$, and $\mathbb{E}_{a \sim s^k}[R_f^i(a)] = \frac{1}{k}R_f^i(a(t))$ is the average regret on the samples. Since x is a correlated equilibrium we know that $\mathbb{E}_{a \sim x}[R_f^i(a)] \leq 0$. Therefore by Hoeffding's inequality and the choice of k we have

$$\mathbb{P}(\mathbb{E}_{a \sim s^k}[R_f^i(a)] \geq \varepsilon) \leq e^{-\frac{\varepsilon^2}{2}k} \leq \frac{\alpha}{nm^m}$$

For every player i , there are m^m switching rules of the form $f : A_i \rightarrow A_i$. Hence, summing across n players, we get that the total number of different switching rules is nm^m . Therefore, the union bound implies that with probability greater than $1 - \alpha$ we have $\mathbb{E}_{a \sim s^k}[R_f^i(a)] < \varepsilon$. Hence, with probability at least $1 - \alpha$, the empirical distribution s^k is an ε -correlated equilibrium. □

Proof of Theorem 4. By Corollary 1, there exists an $\frac{\varepsilon}{2}$ -Nash equilibrium x where every player i uses at most $b = \lceil \frac{32}{\varepsilon^2}(\ln n + \ln m - \ln \varepsilon + \ln 16) \rceil$ actions from A_i . We denote the set of player's i actions that are played with positive probability in x by

B_i , where $|B_i| \leq b$. Let us implement the sampling method for the distribution x which is an $\frac{\varepsilon}{2}$ -Nash equilibrium, and therefore, also an $\frac{\varepsilon}{2}$ -correlated equilibrium.

Since $\mathbb{E}_{a \sim x}[R_f^i(a)] \leq \frac{\varepsilon}{2}$ by Hoeffding's inequality we have

$$\Pr(\mathbb{E}_{a \sim s^k}[R_f^i(a)] \geq \varepsilon) \leq e^{-\frac{\varepsilon^2}{8}k}. \quad (10)$$

Note that s^k is an ε -correlated equilibrium iff $\mathbb{E}_{a \sim s^k}[R_f^i(a)] \leq \varepsilon$ for every switching rule $f : B_i \rightarrow A_i$ (note that the number of such switching rules is at most m^b for every player). In other words, we can consider only switching rules $f : B_i \rightarrow A_i$ instead of $f : A_i \rightarrow A_i$, because all the recommendations to player i will be from the set B_i .

The choice of k guarantees that

$$e^{-\frac{\varepsilon^2}{8}k} < \frac{1}{nm^b}. \quad (11)$$

Therefore, using inequality (10) and the union bound, we get that with positive probability $\mathbb{E}_{a \sim s^k}[R_f^i(a)] \leq \varepsilon$ is satisfied for every $f : B_i \rightarrow A_i$, which implies that s^k is an ε -correlated equilibrium. This obviously implies that such a k -uniform correlated equilibrium exists. \square

Proof of Theorem 5. Hoeffding's inequality and the choice of k guarantees that

$$\mathbb{P}(|\mathbb{E}_{a \sim x}[R_f^i(a)] - \mathbb{E}_{a \sim s^k}[R_f^i(a)]| \geq \frac{\varepsilon}{2}) \leq 2e^{-\frac{\varepsilon}{8}k} < \frac{\alpha}{nm^m}.$$

The total number of switching rule is nm^m , therefore with probability of at least $1 - \alpha$ we have

$$|\mathbb{E}_{a \sim x}[R_f^i(a)] - \mathbb{E}_{a \sim s^k}[R_f^i(a)]| < \frac{\varepsilon}{2} \quad (12)$$

for all players i and all switching rules $f : A_i \rightarrow A_i$. Throughout, we will assume that inequalities (12) are satisfied.

If x is a δ -correlated equilibrium then

$$\mathbb{E}_{a \sim s^k}[R_f^i(a)] \leq \mathbb{E}_{a \sim x}[R_f^i(a)] + \frac{\varepsilon}{2} \leq \delta + \frac{\varepsilon}{2},$$

which means that s^k is an $(\delta + \frac{\varepsilon}{2})$ -correlated equilibrium.

If x is not a $(\delta + \varepsilon)$ -correlated equilibrium then there exists a player i and a switching rule f^* such that $\mathbb{E}_{a \sim x}[R_{f^*}^i(a)] > \delta + \varepsilon$. So,

$$\mathbb{E}_{a \sim s^k}[R_{f^*}^i(a)] \geq \mathbb{E}_{a \sim x}[R_{f^*}^i(a)] - \frac{\varepsilon}{2} > \delta + \frac{\varepsilon}{2},$$

which means that s^k is not a $(\delta + \frac{\varepsilon}{2})$ -correlated equilibrium. \square

Proof of Theorem 6. Instead of proving that in 2-players m -actions games every test must use $k = \sqrt{\frac{m}{2}}$ samples, we will prove the equivalent statement that in 2-players $(2m)$ -actions games every test must use $k = \sqrt{m}$ samples.

We consider the game that is described in Example 3. Let x be the correlated equilibrium that is considered in Example 3:

$$x((r_1, d_1), (r_2, d_2)) = \begin{cases} \frac{1}{4m} & \text{if } d_1 = d_2 \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

Let $b := (b_d)_{d \in [m]}$ be a vector of size m , where each coordinate b_d is a pair $b_d \in \{(1, 1), (1, 2), (2, 1), (2, 2)\}$. We define the distributions y_b (we have 4^m such distributions) by:

$$y_b((r_1, d_1), (r_2, d_2)) = \begin{cases} \frac{1}{m} & \text{if } d_1 = d_2 = d \text{ and } (r_1, r_2) = b_d \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

Simply speaking, the distribution y_b picks for every $d \in [m]$ single action (r_i, r_j) for both players and puts a measure of $\frac{1}{m}$ on it. This is in a contrast to x which puts an equal measure of $\frac{1}{4m}$ on all four actions (r_i, r_j) .

Let ω be the event $\omega := \{((r(t), d(t)))_{t \in [k]} : d(t) \neq d(t') \text{ for } t \neq t'\}$; i.e., all the samples have different values of d . Note that $\mathbb{P}_x(\omega) = \mathbb{P}_{y_b}(\omega)$ for every b , because the event ω depends only on the samples of d , and both x and y_b have the uniform distribution over the values of d .

We claim that if $k = \lfloor \sqrt{m} \rfloor$ then $\mathbb{P}_x(\omega) = \mathbb{P}_{y_b}(\omega) > \frac{1}{2}$ (for every b).

The t th sample will have the same value d as one of the previous with probability of at most $\frac{t-1}{m}$. Using the union bound we get that

$$1 - \mathbb{P}_x(\omega) \leq \frac{0}{m} + \frac{1}{m} + \frac{2}{m} + \dots + \frac{\lfloor \sqrt{m} \rfloor - 1}{m} \leq \frac{(\sqrt{m} - 1)\sqrt{m}}{2m} < \frac{1}{2}.$$

A test with error-probability $\frac{1}{4}$ should return with probability $\frac{3}{4}$ the answer **YES** for the correlated equilibrium x , and it should return the answer **NO** with probability $\frac{3}{4}$ for all the distributions y_b which are not a $\frac{1}{2}$ -correlated equilibria. In particular, if we first draw the distribution from which we sample (according to some probability distribution), and then sample from the chosen distribution, the probability of an error of the test should be less than $\frac{1}{4}$ (because for each one of the distributions that we draw the probability of error is less than $\frac{1}{4}$). Let us draw the distribution from which we sample as follows. The distribution x is chosen with probability $\frac{1}{2}$, and each one of the distributions y_b is chosen with probability $\frac{1}{2 \cdot 4^m}$. If the sequence of samples is $(r(t), d(t))_{t \in [k]} \in \omega$ then, by the symmetry

of the distributions $\{y_b\}_b$, the probability that it is sampled from x is equal to the probability that it is sampled from one of the distributions y_b . Therefore, for sequences of samples in ω the test makes an error with probability of at least $\frac{1}{2}$, and sequence of samples is in ω with probability of at least $\frac{1}{2}$. Therefore, the probability of an error is at least $\frac{1}{4}$. □

5 Approximate Coarse Correlated Equilibrium

The case of coarse correlated equilibrium is the simplest one. Here we present the results for coarse correlated equilibria without the proofs, since they are quite similar to the proofs of the correlated-equilibrium results presented in Section 4.

The key difference between coarse correlated equilibrium and correlated equilibrium is that in coarse correlated equilibrium every player is allowed to deviate to one fixed pure action (irrespective of mediators recommendation), instead of allowing the player to deviate to different actions for different recommendations.

Definition 6. For every pure action $j \in A_i$ we denote by $R_j^i(a) := u_i(j, a_{-i}) - u_i(a_i, a_{-i})$ the regret of player i for not choosing the action j at strategy profile a .

A distribution $x \in \Delta(A)$ is an ε -coarse correlated equilibrium if $\mathbb{E}_{a \sim x}[R_j^i(a)] \leq \varepsilon$ for every player i and every action $j \in A_i$.

Theorem 7. For every $\varepsilon, \alpha > 0$, every n -player m -action game and every coarse correlated equilibrium of the game x , the empirical distribution of the samples s^k is an ε coarse correlated equilibrium with probability greater than $1 - \alpha$ for

$$k > \frac{2}{\varepsilon^2}(\ln m + \ln n - \ln \alpha) = O(\log m + \log n).$$

We can establish this theorem using the same ideas as in the proof of Theorem 3. The only difference is that here, instead of nm^m inequalities (one for every R_f^i , where $f : A_i \rightarrow A_i$), we need to satisfy only nm inequalities, one for every R_j^i .

This theorem gives us the following result regarding the existence of simple approximate coarse correlated equilibrium:

Corollary 3. Every n -player m -action game admits a k -uniform ε -coarse correlated equilibrium for every

$$k > \frac{2}{\varepsilon^2}(\ln m + \ln n) = O(\log m + \log n). \tag{15}$$

This bound is tight, see Examples 1 and 2.

Regarding the simplicity of testing approximate coarse correlated equilibrium play, we have the following result which is analogous to Theorem 5.

Corollary 4. Let

$$T((a(t))_{t \in [k]}) = \begin{cases} \mathbf{YES} & \text{if } s^k \text{ is a } (\delta + \frac{\varepsilon}{2})\text{-coarse correlated equilibrium} \\ \mathbf{NO} & \text{otherwise.} \end{cases}$$

Then T is an ε -test with error probability α for δ -coarse correlated equilibrium for

$$k > \frac{8}{\varepsilon^2}(\ln m + \ln n - \ln \alpha) = O(\log m + \log n).$$

The theorem follows from a proof similar to the one for Theorem 5.

6 Lower Bounds

In this section we present lower bounds for the number of samples from an equilibrium that are required in order that the empirical distribution of play will be an approximate equilibrium (with high probability).

The following example demonstrates that $\Omega(\log n)$ samples are required for all cases: Nash equilibrium, correlated equilibrium, and coarse-correlated equilibrium.

Example 1. Consider the following $2n$ -players two-actions game. We have n pairs of players $(p_i^1, p_i^2)_{i \in [n]}$. Player p_i^j is playing matching-pennies with his partner p_i^{3-j} (the actions of the pair (p_i^1, p_i^2) have no influence on the payoffs of other pairs). Consider the Nash equilibrium where each player is playing $(\frac{1}{2}, \frac{1}{2})$ (which is also a correlated equilibrium and a coarse-correlated equilibrium). If the number of samples is $k \leq \frac{\log n}{2}$ then the probability that player's p_i^j empirical distribution of play will be a pure strategy is

$$2 \left(\frac{1}{2} \right)^{\frac{\log n}{2}} \geq \frac{1}{\sqrt{n}}.$$

Therefore the probability that no player will have a pure-strategy empirical distribution is at most

$$\left(1 - \frac{1}{\sqrt{n}} \right)^{2n} \approx e^{-2\sqrt{n}}$$

Note that the requirement that no player will have a pure-strategy empirical distribution is necessary for the empirical distribution of play to be $\frac{1}{2}$ -coarse correlated

equilibrium (and therefore it is necessary also for $\frac{1}{2}$ -correlated equilibrium and $\frac{1}{2}$ -Nash equilibrium). So the probability that the empirical distribution of play will form a $\frac{1}{2}$ -equilibrium is exponentially small in n .

The following example of Althofer [1] demonstrates that $\Omega(\log m)$ samples are required for all cases: Nash equilibrium, correlated equilibrium and coarse-correlated equilibrium (actually for the correlated equilibrium case Example 3 will demonstrate much stronger result).

Example 2. Let $m = 4^b$ for $b \in \mathbb{N}$, and consider the following two-players m -actions zero-sum game.

Player 1 picks an element $i \in [2b]$ (player 1 has $2b < m$ actions).

Player 2 picks a subset of $S_j \subset [2b]$ such that $|S_j| = b$ (player 2 has $\binom{2b}{b} < m$ actions).

The payoffs are defined by

$$u_2(i, S_j) = -u_1(i, S_j) = \begin{cases} 1 & \text{if } i \in S_j, \\ 0 & \text{otherwise.} \end{cases}$$

Player 1 can guarantee to pay at most $\frac{1}{2}$ by playing the uniform distribution. If in the support of the distribution (which might be correlated) player 1 plays at most b different actions then player 2 has a pure strategy that will yield a payoff of 1. Therefore in every $\frac{1}{4}$ -equilibrium (Nash, correlated or coarse-correlated) player 1 should play at least $b + 1$ different strategies. Therefore, in order that the empirical distribution will be a $\frac{1}{4}$ -equilibrium the number of samples must be greater than $b = \frac{\log m}{2}$.

The following example demonstrates that $\Omega(m)$ samples are required for the case of correlated equilibrium.

Example 3. Consider the following two-players $2m$ -actions zero-sum game. The players are playing matching-pennies, but in addition to player's i "real" action $r_1 \in [2]$ player i also chooses a "dummy" action $d_i \in [m]$ which does not influence the payoff. Formally, the payoff functions of the players are defined by

$$u_1((r_1, d_1), (r_2, d_2)) = -u_2((r_1, d_1), (r_2, d_2)) = \begin{cases} 1 & \text{if } r_1 = r_2, \\ 0 & \text{otherwise.} \end{cases}$$

Consider the correlated equilibrium x where $x((r_1, d), (r_2, d)) = \frac{1}{4m}$ for every $d \in [m]$ and every $r_1, r_2 \in [2]$. In other words, x is the correlated equilibrium

where beyond the actual $(\frac{1}{2}, \frac{1}{2})$ play of the real matching-pennies, the players always chooses the same dummy action.

If the number of samples is $k = m$, then for any $d \in [m]$ the probability that it is picked exactly once during the sampling is $m \frac{1}{m} \cdot (1 - \frac{1}{m})^{m-1} \approx \frac{1}{e}$. If a certain d was picked exactly once then both players can deduce from d which action their opponent will play. Note that the expected number of $d \in [m]$ that are sampled exactly once is $\frac{m}{e}$. Moreover, the probability that the number of exactly-once-sampled d 's will be smaller than $\frac{m}{2e}$ is exponentially small in m (see , e.g., [3], Lemma 4). So, with probability that is exponentially close to 1, in the resulting uniform distribution at least one player may increase it's payoff by at least $\frac{1}{4e}$ by reacting optimally to the opponent's known strategy in all cases where she got the recommendation (r_i, d) where d was chosen exactly once. Therefore the empirical distribution of samples is an $\frac{1}{4e}$ -correlated equilibrium with exponentially small (in m) probability.

The focus of the paper was on the dependence of the number of samples on m and n . However, Theorems 1,3, and 7 proves also a dependence on ε . For the case of Nash equilibrium, Theorem1 proves a bound of $O(\frac{1}{\varepsilon^2} \log(\frac{1}{\varepsilon}))$. Theorems 3 and 7 proves a bound of $O(\frac{1}{\varepsilon^2})$ for correlated and coarse-correlated equilibrium. The following example demonstrates that those bounds are tight (except for the case of Nash equilibrium where is it almost tight).

Example 4. Consider the matching-pennies game, with the unique Nash equilibrium $((\frac{1}{2}, \frac{1}{2})(\frac{1}{2}, \frac{1}{2}))$. A necessary condition for the empirical distribution of play to form an ε -equilibrium (Nash correlated or coarse-correlated) is that the empirical distribution of player 1 should be $(p, 1 - p)$ where $p \in [\frac{1}{2} - \varepsilon, \frac{1}{2} + \varepsilon]$. By the strong law of large numbers, after k samples, with constant probability the deviation from the expectation ($p = \frac{1}{2}$) is at least $\frac{1}{\sqrt{k}}$. Therefore, if we draw k samples for $k < \frac{1}{\varepsilon^2}$, then with positive probability the deviation from $\frac{1}{2}$ will be at least $\frac{1}{\sqrt{k}} > \varepsilon$.

7 Discussion

7.1 Sampling from One Type of Equilibrium to Achieve Another

In this paper we considered three types of equilibria: Nash, correlated, and coarse correlated. Our high level approach was to sample from an equilibrium of a particular type to generate an approximate equilibrium of the same type. We can modify this approach a bit and, in principle, ask the following question: How many samples

from an equilibrium of a particular type are required to generate an approximate equilibrium of a different type?

Note that the notion of coarse correlated equilibrium is a generalization of correlated equilibrium, and the later generalizes Nash equilibrium. In general, we cannot hope to get a more refined notion of equilibrium by sampling from a more general one. But, hypothetically, it might be the case that fewer samples from a refined equilibrium type are sufficient for generating an approximate equilibrium of a more general type.

First we observe that $\Omega(\log n + \log m)$ samples are necessarily required to generate a coarse correlated equilibrium, even if the samples are drawn from a Nash equilibrium. This follows from the lower bound of Example 1 (wherein we actually sample from a Nash equilibrium) and Example 2 (in which the counting argument holds *irrespective* of the initial distribution).

So the remaining question is, can $o(m)$ samples from a Nash equilibrium generate an approximate correlated equilibrium? In other words, can we overcome the $\Omega(m)$ sampling lower bound established in Example 3? The answer to this question is *no*. In particular, consider the same game as in Example 3, but now draw m samples from the Nash equilibrium where both players are playing the uniform distribution over their $2m$ actions. We say that a pair (d_1, d_2) *appears exactly once* if the pair (d_1, d_2) appears in the sample (i.e., one of the samples is $((r_1, d_1), (r_2, d_2))$ for some $r_1, r_2 \in [2]$) and, for $i = 1, 2$, d_i appears exactly once among all the samples.

For every pair (d_1, d_2) , the probability that it appears exactly once is equal to $m \frac{1}{m^2} (1 - \frac{2m-1}{m^2})^m \approx \frac{1}{e^2 m}$. Therefore, the expected number of recommendation pairs, (d_1, d_2) , that appear exactly once is $\frac{m}{e^2}$. Hence, among the m samples a significant fraction of pairs appear exactly once. Note that if a recommendation appears exactly once, then both the players can deduce their opponent's strategy from the recommendation, which cannot occur at an equilibrium.

7.2 Future Work

This paper establishes tight bounds on the rate of convergence of the empirical distribution (of equilibrium play) to an approximate equilibrium. These bounds imply the existence of *small-support* approximate equilibria. But, whether our polylogarithmic upper bounds on support size are tight remains an open question. Note that the $\log m$ lower bound developed in Example 2 applies to support size as well. In particular, Example 2 establishes that there does not exist an ε -equilibrium (Nash, correlated, or coarse correlated) with support size smaller than $\log m$. However, to the best of our knowledge, lower bounds on support size of approximate equilibrium

(Nash, correlated, or coarse correlated) that depends on n have not been established.

Open Question: Let $k = k(n, m, \epsilon)$ be the smallest number such that every n -player m -action game admits a k -uniform ϵ -equilibrium. Fixed $\epsilon > 0$ and $m \in \mathbb{N}$ (i.e., we refer to them as constants). What is the asymptotic behavior of $\lim_{n \rightarrow \infty} k(n)$? In particular, does $\lim_{n \rightarrow \infty} k(n) = \infty$? The question remains open for all three equilibrium types: Nash, correlated, and coarse correlated.

Finally, other directions that warrant further study include Bayesian games. In that setting, we can consider additional solution concepts such as communication equilibria.

Acknowledgements

The authors would like to thank Federico Echenique for his constructive suggestions and comments.

References

- [1] Althofer, I. (1994) “On Sparse Approximations to Randomized Strategies and Convex Combinations,” *Linear Algebra and Its Applications* 199, 339-355.
- [2] Azrieli, Y. and Shmaya, E. (2013) “Lipschitz Games,” *Mathematics of Operations Research*, forthcoming.
- [3] Farach-Colton, M. and Mosteiro, M.A. (2007) “Sensor Network Gossiping or How to Break the Broadcast Lower Bound,” *Algorithms and Computation*, 4835, 232-243.
- [4] Foster, D.P. and Young, H.P., (2006) “Regret-testing: Learning to play Nash equilibrium without knowing you have an opponent”, *Theoretical Economics* 1, 341-367.
- [5] Germano, F. and Lugosi, G. (2007) “Global Nash convergence of Foster and Young’s regret-testing,” *Games and Economic Behavior* 60, 135-154.
- [6] Germano, F. and Lugosi, G. (2007) “Existence of Sparsely Supported Correlated Equilibria,” *Economic Theory*, 32, 575-578.
- [7] Hart, S. and Mas-Colell, A. (2013) *Simple Adaptive Strategies: From Regret-Matching to Uncoupled Dynamics*, World Scientific Publishing.

- [8] Hémon, S., Rougemont, M., and Santha, M. (2008) “Approximate Nash Equilibria for Multy-player Games,” *Algorithmic Game Theory, Lecture Notes in Computer Science* 4997, 267–278.
- [9] Hoeffding, W. (1963) “Probability Inequalities for Sums of Bounded Random Variables,” *Journal of the American Statistical Association* 58, 13–30.
- [10] Hofbauer J., Sigmund, K., (1998) *Evolutionary Games and Population Dynamics*, Cambridge University Press.
- [11] Kalai, E. (2004) “Large Robust Games,” *Econometrica* 72, 1631–1665.
- [12] Lipton, R. J., Markakis, E., and Mehta, A. (2003) “Playing Large Games Using Simple Strategies,” *Proceedings of the 4th ACM Conference on Electronic Commerce*, pp. 36–41.
- [13] Maynard Smith, J., Price, G. R., (1973) “The Logic of Animal Conflict,” *Nature* 246, 15 – 18.
- [14] Nisan, N. (2009) “Approximate Nash,” blog: Turing’s Invisible Hand: Computation, Economics and Game Theory. <http://agtb.wordpress.com/2009/06/07/approximate-nash/>.
- [15] Robinson, J. (1951) “An Iterative Method of Solving a Game,” *Annals of Mathematics*, 52, 296–301.
- [16] Young, H.P., (2004) *Strategic Learning and Its Limits*, Oxford UK: Oxford University Press.