

Empirical Distribution of Equilibrium Play and Its Testing Application

Ron Peretz

(London School of Economics)

Schedule: December 24-2014 14:15-15:30

Place: Bld. 72, room 465, Dept. of Economics, BGU

Abstract:

We show that in any n -player m -action normal-form game, we can obtain an approximate equilibrium by sampling any mixed-action equilibrium a small number of times. We study three types of equilibria: Nash, correlated and coarse correlated. For each one of them we obtain upper and lower bounds on the number of samples required for the empirical distribution over 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 substantially improve previously known upper bounds on the support size of approximate equilibria in games with many players. In particular, for all the three types of equilibria we show the existence of approximate equilibrium with support size polylogarithmic in n and m , whereas the previously best-known upper bounds were polynomial in n [16, 11, 15].

Empirical Distribution of Equilibrium Play and Its Testing Application

Yakov Babichenko* Siddharth Barman[†] Ron Peretz[‡]

Abstract

We show that in any n -player m -action normal-form game, we can obtain an approximate equilibrium by sampling any mixed-action equilibrium a small number of times. We study three types of equilibria: Nash, correlated and coarse correlated. For each one of them we obtain upper and lower bounds on the number of samples required for the empirical distribution over 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 substantially improve previously known upper bounds on the support size of approximate equilibria in games with many players. In particular, for all the three types of equilibria we show the existence of approximate equilibrium with support size polylogarithmic in n and m , whereas the previously best-known upper bounds were polynomial in n [16, 11, 15].

1 Introduction

Equilibria are central solution concepts in game theory. They are widely used as models to explain the observed behavior of self-interested entities like human players in strategic settings. Arguably the most prominent examples of such notions

*Faculty of Industrial Engineering, Technion - Israel Institute of Technology, yakovbab@tx.technion.ac.il.

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

[‡]Department of Mathematics, London School of Economics. E-mail: r.perez@lse.ac.uk.

of rationality are Nash equilibrium [21], correlated equilibrium [2], and coarse correlated equilibrium [14]. At a high level, these concepts denote distributions over players’ action profiles where no player can benefit, in expectation, by unilateral deviation. Though, in most empirical applications of game theory these underlying distributions are not explicitly known to an observer. Rather, what one observes is the behavior of the players, i.e., players’ realized actions.

This naturally leads us to consider a setting in which players implement an underlying distribution—i.e., a mixed strategy—during multiple plays of the same game. Here, the mixed strategy is not known to an outside observer. Rather, the observer sees the (pure) actions selected by the players during the play; in other words, she observes independent and identically distributed (i.i.d.) samples from the mixed strategy. This framework captures typical (empirical) applications of game theory, and entails 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 equilibrium¹ (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 generate an approximate equilibrium with high probability. This result has several useful interpretations.

1. Testing whether 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. Such scenarios are studied throughout 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. Our results (Theorems 2, 5 and Corollary 4) show that we can accomplish this testing task 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], [19], and [16]). In particular, these results

¹An approximate equilibrium, with approximation parameter $\varepsilon > 0$, is a distribution over action profiles at which no player has more than an ε incentive to deviate.

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; such approximate equilibria are referred to as simple approximate equilibria. Our result extends this line of work by substantially improving the upper bounds on the support size of 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. Population games. Sampling occurs in life very naturally. In biology [25], for example, members of a species come in different traits: sex, size, color, etc. Every newborn has a trait which is sampled according to some distribution prescribed by its species. An ecological system is abstractly modeled as a game, where the players are species, the strategies are their traits, and mixed strategies represent distributions of traits among individuals of a species. This abstraction relies on the assumption that the size of the population is large enough to represent the mixed strategy of the species faithfully. The present paper provides estimates on what should be considered as “large enough” for such strategic reasoning. The above situation is modeled in detail by a *population game*. Here, each individual is modeled as a player, where members of the same species are identical players. It is assumed that individuals can’t change their trait (sex, color, etc.); therefore it is reasonable to look at (approximate) pure equilibria. The diversity of traits is explained by assuming some source of randomization in the creation of individuals; therefore the special interest in sampled equilibria.

Another justification for seeking such sampling-based pure equilibria was given by Kalai [18], who viewed them as *ex-post* equilibria. That is, the players play a mixed strategy equilibrium, and even after their strategies are realized they have no incentive to modify them.

Our results can be interpreted as bounds on the minimal size of the population that ensures an ex-post equilibrium in a population game. Kalai [18] worked in the more general framework of semi-anonymous games. 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 [3]. It should be noted that the results we obtain for the special case of population games are stronger than what one could hope to obtain for arbitrary Lipschitz games.

4. 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 [15] and [27]). 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 that they play in each repetition of the game.² In this framework, a natural way of learning opponents' mixed strategy is through the empirical distribution (i.e., approximating opponents' mixed strategy by the empirical distribution of the observed action profiles). Such a learning process has been considered in the classical fictitious play literature [23], and in many other recent results; e.g., the regret-testing dynamics (see [10] and [12]).

Overall, we get the following key question in settings where players learn their opponents' mixed strategy through pure-action samples played by the opponents: How fast does the empirical distribution of equilibrium play forms an approximate equilibrium?

The answer to this question sheds light on the short-run stability of equilibria. If the empirical distribution does not form an approximate equilibrium after multiple iterations (i.e., after multiple plays), 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 [10] and [12]. Our results imply that, even in large games, after a small number of iterations with high probability such a situation will not occur. I.e., the equilibria are stable in the short run.

Specifically, we provide almost tight bounds for the above question (see Table 1). These bounds in particular show that after a small number of iterations, each player can learn whether her current strategy is approximately a best reply to her opponents' strategies.

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 the players are playing according to a product distribution, the correct notion for the *empirical distribution of play* is the product of the empirical distributions of the actions played by each player. For the cases of correlated and coarse correlated equilibrium the correct notion for the

²Note that this assumption holds in most repeated-game models, i.e., perfect monitoring.

empirical distribution of play is simply the empirical distribution of the sampled action profiles.

The empirical distribution of play is a random variable. In this paper we address the following question: how large should the number of samples k be to guarantee that the empirical distribution of play forms an approximate equilibrium with probability close to 1? We provide almost tight bounds for this question, which are summarized in the table below.

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

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

Moreover, we show that if players are playing according to a distribution that is *not* an approximate equilibrium, then for the same values of k (as 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 a very large number of players or a very large number of actions there exists efficient tests to determine whether players are playing according to a Nash equilibrium or a coarse correlated equilibrium. Correlated equilibrium on the other hand is a slightly more complicated notion in this respect. We establish this by proving that there does exists 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.

The fact that the empirical distribution of play forms an approximate equilibrium with merely positive probability is interesting in its own right. The fact that this event occurs with positive probability establishes, via the probabilistic method, the existence of an approximate equilibrium with small *support size*. Note that “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 are played with positive probability by any single player. For correlated and coarse correlated equilibrium, the support size is the number of *action profiles* that have a non-zero probability of being played in the approximate equilibrium.

Small-support approximate Nash equilibria have been previously studied in [1], [19], and [16]. Althöfer [1] along with Lipton and Young [20] studied the problem for two-player m -action zero-sum games and established a $\Theta(\log m)$ bound on the support size. Lipton, Markakis, and Mehta [19] studied the same question of general n -player m -actions games and achieved an upper bound of $O(n^2 \log m)$. Hémon et al. [16] improved the bound of [19] to $O(n \log m)$. Our result implies the existence of an approximate Nash equilibrium with support size $O(\log m + \log n)$ (see Corollary 1). This result gives us an algorithm for computing an approximate Nash equilibrium with running time $O(N^{\log \log N})$ in games where $m = \text{poly}(n)$ (i.e., the number of actions is not significantly larger than the number of players). Here N is the input size of the game; see Corollary 2. Note that the running time of the best previously known algorithm for this problem is $O(N^{\log N})$ (see [22]). Hence, the current paper provides an exponential improvement upon prior work in the computation of approximate Nash equilibrium in large games.

Support-size upper bounds for *exact* correlated equilibria were studied by Germano and Lugosi in [11]. Specifically, they showed that every n -player m -action game admits a correlated equilibrium with support size $O(nm^2)$. Applying the technique of [11] we can obtain a bound of $O(nm)$ on the support size 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 and previously known results are summarized in Table 2.

2 Notation and Preliminaries

We consider n -player m -action games, i.e., games with n -player and m actions per player.³ Write $N := nm^n$ to denote the input size of such games; i.e., the size of a list

³All 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 ease of presentation, we assume throughout that all players have the same number of actions m .

Approximate Equilibrium	Our Results	Previous Bounds
Nash	$O(\log m + \log n)$ Corollary 1	$O(n \log m)$ [16]
Correlated	$O(\log m(\log m + \log n))$ Theorem 4	Exact: $O(nm^2)$ [11]
Coarse Correlated	$O(\log m + \log n)$ Corollary 3	Exact: $O(nm)$ [11]

Table 2: Support size of approximate equilibria.

that enumerates players' utilities at every pure action profile. 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 all probability distributions over a set B is denoted by $\Delta(B)$. Therefore, $\Delta(A)$ is the set of all probability distributions over the action profiles, and $\prod_{i \in [n]} \Delta(A_i)$ is the set of all *product distributions*. For a 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 payoffs of the players are normalized between 0 and 1. Specifically, the payoff function of player i is denoted by $u_i : A \rightarrow [0, 1]$ 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 a set B is 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 $c_b \in \mathbb{N}$. The set of all k -uniform distributions on B is denoted $\Delta_k(B)$.

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 $a_i \in A_i$.

When $\varepsilon = 0$ we say that x is an *exact Nash equilibrium*, or simply *Nash equilibrium*.

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

Assume that the 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 $(a(t))_{t \in [k]}$, where

each $a(t) \in A$. Since we assume 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 $(a_i(t))_{t \in [k]}$. Formally, $s_i^k(a_i) := \frac{1}{k} |\{t : a_i(t) = a_i\}|$. 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 a 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. Let x be a Nash equilibrium of an n -player m -action game and parameters $\varepsilon, \alpha \in (0, 1)$. Then, the product empirical distribution of play $(s_i^k)_{i \in [n]}$ defined over k i.i.d. samples from x is an ε -Nash equilibrium with probability greater than $1 - \alpha$, for every

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

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 in Section 3.3. We note that 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 multiset of k samples that forms an ε -Nash equilibrium; this claim follows from 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. Here, we say that an approximate Nash equilibrium $\prod_i s_i$ is k -uniform if every s_i is a k -uniform distribution.

Corollary 1. Every n -player m -action game admits an k -uniform ε -Nash equilibrium for every

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

This corollary guarantees the existence of an approximate equilibrium (in every n -player m -action game) where each player uses only a small number of actions in

her mixed strategy (at most $O(\log m + \log n)$). Another useful consequence, is the simplicity of the probabilistic structure of the mixed strategy of each player. To see it, consider, for example, the case of n -player 2-action games. Here, corollary 1 with

$$k = \left\lceil \frac{8}{\varepsilon^2} (\ln n + \ln \varepsilon + \ln 16) \right\rceil + 1 = O(\log n),$$

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

The fact that such a simple approximate Nash equilibrium exists allows us to find an approximate Nash equilibrium just by exhaustively searching over all the possible n -tuples of k -uniform strategies. Although the algorithm is simple, to the best of our knowledge it provides the best-known running-time bound for this problem.

Corollary 2. Given an n -player m -action normal-form game and constant $\varepsilon > 0$, there exists an algorithm that computes an ε -Nash equilibrium of the game in time $\min\{m^{nk}, k^{nm}\}$, with $k = O(\log m + \log n)$.

Recall that $N = nm^n$ denotes the input size of n -player m -action games. 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 $N^{O(\log N)}$. This bound coincides with the best-known upper bound for computing approximate Nash equilibrium (see [22]).

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

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

For the class of games where the number of actions is constant ($m = O(1)$), the bound of Corollary 2 further improves to $N^{\log \log \log N}$:

$$k^{nm} = N^{\frac{nm \log k}{\log N}} = \text{poly}(N^{\log k}) = \text{poly}(N^{\log \log \log N}).$$

3.2 Testing Approximate Nash Equilibrium

Our goal is to test whether players are playing according to an approximate Nash equilibrium. The assumption is that the exact mixed strategies of the players are unobserved. Instead, we observe i.i.d. samples from the underlying mixed strategies.

Our focus is on the following 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 “boundary”. 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 using a finite number of samples, which is impossible. Therefore, we weaken our requirements from a test by providing a slack of ε .

Definition 3. Given a number of samples $k \in \mathbb{Z}_+$, a function $T : A^k \rightarrow \{\mathbf{YES}, \mathbf{NO}\}$ is said to be an ε -test that has error probability α for δ -Nash equilibrium if for every product distribution $x = (x_i)_{i \in [n]} \in \prod_i \Delta(A_i)$ 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 when 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, when the number of samples

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

Note that the number of samples is independent of δ . Section 3.3 contains a proof of Theorem 2.

3.3 Proofs

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

Lemma 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 lemma 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 lemma is a key technical contribution of this work that proves a novel concentration inequality for product distributions. Even though we apply the sampling method as in [19] and [16], the fact that we use Lemma 1 instead of some standard concentration inequality essentially enables us to significantly improve upon the previously-best-know bounds. The proof of Lemma 1 is brought in the appendix.

Proof of Theorem 1. The proof uses similar idea to [19] or [16]. Lemma 1 and the choice of k guarantee 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. Lemma 1 and the choice of k guarantee that

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

for every player i and every action $a_i \in [m]$. In addition, Hoeffding's inequality (see [17]) guarantees that for a given x_{-i} we have

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

for every player $i \in [n]$. Note that there are $n(m+1)$ inequalities of the form (1) and (2). Therefore, the union bound implies that with probability greater than $1 - \alpha$ the following $n(m+1)$ inequalities hold:

$$\begin{aligned} |u_i(a_i, s_{-i}^k) - u_i(a_i, x_{-i})| &\leq \frac{\varepsilon}{6} & \forall i \in [n], \forall a_i \in [m]. \\ |u_i(s_i^k, x_{-i}) - u_i(x_i, x_{-i})| &\leq \frac{\varepsilon}{6} & \forall i \in [n]. \end{aligned} \quad (3)$$

Throughout the proof we will assume that all the inequalities in (3) 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, the choice of k guarantees that all the inequalities in (3) 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, i.e., a setting in which players followed their mixed strategies independently. We will now consider complementary notions of equilibria that address settings in which players' actions are correlated. Specifically, this section is focused on correlated equilibria and the next section is on coarse correlated equilibria.

A typical interpretation 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 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 the 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*, 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* from x that are denoted by $(a(t))_{t \in [k]}$ where $a(t) \in A$. Write s^k for 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. Let x be a correlated equilibrium of an n -player m -action game and parameters $\varepsilon, \alpha \in (0, 1)$. Then, the empirical distribution s^k defined over k i.i.d. samples from x is an ε -correlated equilibrium with probability greater than $1 - \alpha$ for every

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

The bounds of Theorem 3 are almost tight. Specifically, Example 1 in Section 6 demonstrates 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 appears in Section 4.3.

4.1 Existence of Simple Approximate Correlated Equilibrium

We should emphasize again that the support of a correlated equilibrium is the number of *action profiles* in the support of the equilibrium. Also, note that in an n -player m -action game the number of action profiles is m^n . If we use the existence of small-support approximate Nash equilibrium (which is also an approximate correlated equilibrium) we obtain the existence of an 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 we obtain an approximate correlated equilibrium support size poly-logarithmic in m , instead of a polynomial? Example 3 demonstrates that if we sample from an *arbitrary* correlated equilibrium then we cannot. But, if we sample from a *specific* approximate correlated equilibrium, then poly-logarithmic number of samples will be 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 support size poly-logarithmic in 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

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

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

4.2 Testing Approximate Correlated Equilibrium Play

As in 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.⁴

Definition 5. Given a number of samples $k \in \mathbb{Z}_+$, 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 $x \in \Delta(A)$ 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 quite natural, we need to simply check whether the empirical distribution of play is an approximate correlated equilibrium or not.

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, when the number of samples

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

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 in Section 6 demonstrate that the natural test that is presented in the theorem cannot use less than $\Omega(m)$ 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.

⁴We note again, that it is impossible to design such a test for distinguishing between δ -correlated equilibrium and not a δ -correlated equilibrium. This follows via arguments similar to the ones in Section 3.2.

4.3 Proofs

Proof of Theorem 3. Note that $R_f^i(a)$ where $a \sim x$ 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. \square

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

$$\Pr(\mathbb{E}_{a \sim s^k}[R_f^i(a)] \geq \varepsilon) \leq e^{-\frac{\varepsilon^2}{8}k}. \quad (5)$$

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

$$e^{-\frac{\varepsilon^2}{8}k} < \frac{1}{nm^b}. \quad (6)$$

Therefore, using inequality (5) 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 implies, via the probabilistic method, that such a k -uniform correlated equilibrium exists. \square

Proof of Theorem 5. Hoeffding's inequality and the choice of k guarantee 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 rules is nm^m , therefore with probability of at least $1 - \alpha$ we have

$$|\mathbb{E}_{a \sim x}[R_f^i(a)] - \mathbb{E}_{a \sim s^k}[R_f^i(a)]| < \frac{\varepsilon}{2}, \quad (7)$$

for all players i and all switching rules $f : A_i \rightarrow A_i$. Throughout the rest of the paper, we assume that inequality (7) is satisfied for every switching rule.

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. □

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

We consider the same game as in Example 3 in Section 6: consider a two-players $2m$ -actions zero-sum game in which 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

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

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 choose the same dummy action.

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:

$$y_b((r_1, d_1), (r_2, d_2)) = \begin{cases} \frac{1}{m} & \text{if } d_1 = d_2 = d \text{ and } (r_1, r_2) = b_d, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Loosely speaking, the distribution y_b picks for every $d \in [m]$ a single action (r_i, r_j) for both players and puts a measure of $\frac{1}{m}$ on it. This is in 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 $\frac{1}{2}$ -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 the recommendations of the mediator), 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. Let x be a coarse correlated equilibrium of an n -player m -action game and parameters $\varepsilon, \alpha \in (0, 1)$. Then, the empirical distribution s^k defined over k i.i.d. samples from x is an ε -coarse correlated equilibrium with probability greater than $1 - \alpha$ for every

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

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 that establishes 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

$$k > \frac{2}{\varepsilon^2}(\ln m + \ln n) = O(\log m + \log n). \quad (9)$$

This bound is tight, see Examples 1 and 2.

Analogous to Theorem 5, we have the following result for testing approximate coarse correlated equilibrium play.

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

The following example of Althöfer [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 a 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 $r_i \in [2]$ 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

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

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 choose 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., [9], Lemma 4). So, with probability that is exponentially close to 1, in the resulting uniform distribution at

least one player may increase its 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 probability exponentially small in m .

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, Theorem 1 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 central limit theorem, 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 latter 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 of the players can deduce their opponent's strategy from the recommendation, which cannot occur in an equilibrium.

7.2 Hypothesis Testing

The present paper proposes simple intuitive tests to whether players are following some (approximate) equilibrium or behaving far from any equilibrium. The notion of distance that we use, ε -equilibrium, is given in terms of the game incentives. This is in contrast to the general-purpose hypothesis-testing literature that refers to distances between distributions, such as the total variation norm.

Putting our problem in a more general context, gives rise to a question of independent interest.

Question 1. Consider the set of all probability measures on d elements, $\Delta(d)$. Given a polytope $P \subset \Delta(d)$ and $\varepsilon > 0$, how many independent samples from an unknown distribution $q \in \Delta(d)$ are needed to ascertain with high probability that either

- H_1 : q is in P , or
- H_2 : the total variation distance $\text{dist}(q, P) > \varepsilon$?

Upper bounds for Question 1 translate automatically to upper bounds for our correlated and coarse-correlated equilibrium testing; this is because the game payoff function is 1-Lipschitz in the total-variation norm.

It should be noted that the upper bounds obtained in this paper are stronger than what one could hope to deduce from total variation estimates. A special case of Question 1, when P consists of a single point, is known as the “multinomial goodness of fit” problem in hypothesis-testing literature [24, 8], and as the “identity testing” problem in the computer science literature [13, 6, 4, 5, 7, 26]. For this special case the answer to Question 1 is $\Theta(\sqrt{d})$. Note that this special case corresponds to exact equilibrium testing ($\delta = 0$) in games with a unique equilibrium. Here the dimension is $d = m^n$; therefore an upper bound of $m^{\frac{n}{2}}$ is obtained. Even in this particular case, the total variation bound is generally weaker than what we obtained. The only exception is the case of correlated equilibrium in two-player games, where our upper bound is $O(m \log m)$, and the total variation result suggest the possibility of an $O(m)$ upper bound.

Lower bounds for Question 1 do not directly translate to equilibrium testing, because being close in payoffs does not mean being close in total variation distance. For example, a full support equilibrium can be approximated (in terms of payoffs) by mixed strategies of much smaller support.

7.3 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 poly-logarithmic 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.

Question 2. Let $k = k(n, m, \epsilon)$ be the smallest number such that every n -player m -action game admits a k -uniform ϵ -equilibrium. Fix $\epsilon > 0$ and $m \in \mathbb{N}$ (i.e., we refer to them as constants), and denote $f(n) = f(n, m, \epsilon)$. What is the asymptotic behavior of $k(n)$? In particular, does $\lim_{n \rightarrow \infty} k(n) = \infty$? The question remains open for all three equilibrium types: Nash, correlated, and coarse correlated.

Acknowledgements

The authors would like to thank Federico Echenique and Tuğkan Batu for their constructive suggestions and comments.

References

- [1] Ingo Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199:339–355, 1994.
- [2] Robert J Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.
- [3] Yaron Azrieli and Eran Shmaya. Lipschitz games. *Mathematics of Operations Research*, 38(2):350–357, 2013.
- [4] Tuğkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D Smith, and Patrick White. Testing closeness of discrete distributions. *Journal of the ACM (JACM)*, 60(1):4, 2013.
- [5] Tuğkan Batu, Eldar Fischer, Lance Fortnow, Ravi Kumar, Ronitt Rubinfeld, and Patrick White. Testing random variables for independence and identity. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 442–451. IEEE, 2001.
- [6] Tuğkan Batu, Lance Fortnow, Ronitt Rubinfeld, Warren D Smith, and Patrick White. Testing that distributions are close. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 259–269. IEEE, 2000.
- [7] Siu-On Chan, Ilias Diakonikolas, Paul Valiant, and Gregory Valiant. Optimal algorithms for testing closeness of discrete distributions. In *SODA*, pages 1193–1203. SIAM, 2014.
- [8] Noel Cressie and Timothy RC Read. Multinomial goodness-of-fit tests. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 440–464, 1984.
- [9] Martín Farach-Colton and Miguel A Mosteiro. Sensor network gossiping or how to break the broadcast lower bound. In *Algorithms and Computation*, pages 232–243. Springer, 2007.
- [10] Dean P Foster, H Peyton Young, et al. Regret testing: learning to play nash equilibrium without knowing you have an opponent. *Theoretical Economics*, 2006.
- [11] Fabrizio Germano and Gábor Lugosi. Existence of sparsely supported correlated equilibria. *Economic Theory*, 32(3):575–578, 2007.

- [12] Fabrizio Germano and Gabor Lugosi. Global nash convergence of foster and young’s regret testing. *Games and Economic Behavior*, 60(1):135–154, 2007.
- [13] Oded Goldreich and Dana Ron. On testing expansion in bounded-degree graphs. In *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, pages 68–75. Springer, 2011.
- [14] James Hannan. Approximation to bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.
- [15] Sergiu Hart, Andreu Mas-Colell, and Yakov Babichenko. *Simple adaptive strategies: from regret-matching to uncoupled dynamics*, volume 4. World Scientific Publishing Company, 2013.
- [16] Sébastien Hémon, Michel de Rougemont, and Miklos Santha. Approximate nash equilibria for multi-player games. In *Algorithmic Game Theory*, pages 267–278. Springer, 2008.
- [17] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963.
- [18] Ehud Kalai. Large robust games. *Econometrica*, 72(6):1631–1665, 2004.
- [19] Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 36–41. ACM, 2003.
- [20] Richard J Lipton and Neal E Young. Simple strategies for large zero-sum games with applications to complexity theory. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 734–740. ACM, 1994.
- [21] John Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, 1951.
- [22] Noam Nisan. Approximate nash, Blog: Turing’s invisible hand: Computation, economics and game theory. <http://agtb.wordpress.com/2009/06/07/approximate-nash/>, 2009.
- [23] Julia Robinson. An iterative method of solving a game. *Annals of mathematics*, pages 296–301, 1951.

- [24] Minoru Siotani, Yasunori Fujikoshi, et al. Asymptotic approximations for the distributions of multinomial goodness-of-fit statistics. *Hiroshima mathematical journal*, 14(1):115–124, 1984.
- [25] J Maynard Smith and GR Price. The logic of animal conflict. *Nature*, 246:15, 1973.
- [26] Gregory Valiant, Paul Valiant, Arnab Bhattacharyya, Eldar Fischer, and Ronitt Rubinfeld. Instance-by-instance optimal identity testing. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 20, page 111, 2013.
- [27] H Peyton Young. *Strategic learning and its limits*. Oxford University Press, 2004.

A Concentration Inequality for Product Distributions

Here we generalize the classic Hoeffding’s inequality [17] to product probability spaces. In order to state our generalization we need the following definition.

Definition 7. Let (Ω, μ) be a discrete probability space. The k -sample approximation of μ is the random k -uniform distribution $\mu^{(k)} \in \Delta_k(\Omega)$ given by taking the average of k i.i.d. samples from μ .

Formally, one can implement $\mu^{(k)}$ by taking k independent random variables x_1, \dots, x_k assuming values in Ω with distribution μ . By identifying Ω with $\Delta_1(\Omega)$, $\mu^{(k)}$ is given as the $\Delta_k(\Omega)$ -valued random variable

$$\mu^{(k)} = \frac{1}{k} \sum_{i=1}^k x_i.$$

The classic Hoeffding’s Inequality can be stated as follows:

Theorem 8 (Hoeffding 1963). Let (Ω, μ) be a (discrete) probability space. For every $\varepsilon > 0$, $k \in \mathbb{N}$, and $f : \Omega \rightarrow [0, 1]$

$$\mathbb{P}(|\mathbb{E}_{\mu^{(k)}}[f] - \mathbb{E}_{\mu}[f]| > \varepsilon) \leq 2e^{-\frac{\varepsilon^2}{2}k}.$$

When (Ω, μ) is the product of n probability spaces, we obtain a similar inequality in which the approximation is achieved by a product measure and the error term is independent of n .

Theorem 9. Let $(\Omega_1, \mu_1), \dots, (\Omega_n, \mu_n)$ be discrete probability spaces. Consider the product space $(\Omega = \prod_i \Omega_i, \mu = \prod_i \mu_i)$. For every $\varepsilon > 0$, $k \in \mathbb{N}$, and $f : \Omega \rightarrow [0, 1]$

$$\mathbb{P}(|\mathbb{E}_{\prod_i \mu_i^{(k)}}[f] - \mathbb{E}_\mu[f]| > \varepsilon) \leq \frac{4e^{-\frac{\varepsilon^2}{8}k}}{\varepsilon}.$$

Lemma 1 follows immediately from Theorem 9.

Proof of Lemma 1. The strategy profile s_{-i}^k is just the k -sampling approximation of x_{-i} . \square

Proof of Theorem 9. For every $i \in [n]$, let us implement each $\mu_i^{(k)}$ as the average of the k i.i.d. random variables x_1^i, \dots, x_k^i prescribed by the definition of k -sampling approximation. We begin by rewriting $\mathbb{E}_\mu[f]$. For every $l \in [k]$, we can write

$$\mathbb{E}_\mu[f] = \frac{1}{k^n} \sum_{j_1, j_2, \dots, j_n \in [k]} \mathbb{E}[f(x_{j_1+l}^1, \dots, x_{j_n+l}^n)],$$

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

$$\mathbb{E}_\mu[f] = \frac{1}{k^n} \sum_{j_1, j_2, \dots, j_n \in [k]} \frac{1}{k} \sum_{l \in [k]} \mathbb{E}[f(x_{j_1+l}^1, \dots, x_{j_n+l}^n)]. \quad (10)$$

For every initial multi-index $j_* = (j_1, j_2, \dots, j_n) \in [k]^n$ and every $l \in [k]$, we denote $x_{j_*+l} := (x_{j_1+l}^1, \dots, x_{j_n+l}^n) \in \Omega$, and we define the random variable

$$d(j_*) := \begin{cases} 0 & \text{if } \left| \frac{1}{k} \sum_{l \in [k]} \mathbb{E}[f(x_{j_*+l})] - \mathbb{E}_\mu[f] \right| \leq \frac{\varepsilon}{2}, \\ 1 & \text{otherwise.} \end{cases} \quad (11)$$

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

$$d(j_*) + \frac{\varepsilon}{2} \geq \left| \frac{1}{k} \sum_{l \in [k]} \mathbb{E}[f(x_{j_*+l})] - \mathbb{E}_\mu[f] \right|. \quad (12)$$

Note also that for any fixed j_* , the random variables $x_{j_*+1}, x_{j_*+2}, \dots, x_{j_*+k}$ are independent with distribution μ ; therefore their average implements $\mu^{(k)}$; therefore, by Hoeffding's inequality, we have

$$\mathbb{E}[d(j_*)] \leq 2e^{-\frac{\varepsilon^2}{8}k}. \quad (13)$$

Using representation (10) of $\mathbb{E}_\mu[f]$ and inequalities (12) and (13), we get

$$\begin{aligned}
& \mathbb{P}(|\mathbb{E}_{\prod_i(\mu_i^{(k)})}[f] - \mathbb{E}_\mu[f]| \geq \varepsilon) = \\
& = \mathbb{P}\left(\left|\frac{1}{k^n} \sum_{j_* \in [k]^n} \frac{1}{k} \sum_{l \in [k]} \mathbb{E}[f(x_{j_*+l})] - \mathbb{E}_\mu[f]\right| \geq \varepsilon\right) \\
& \leq \mathbb{P}\left(\frac{1}{k^n} \sum_{j_* \in [k]^n} \left|\frac{1}{k} \sum_{l \in [k]} \mathbb{E}[f(x_{j_*+l})] - \mathbb{E}_\mu[f]\right| \geq \varepsilon\right) \\
& \leq \mathbb{P}\left(\frac{1}{k^n} \sum_{j_* \in [k]^n} d(j_*) \geq \frac{\varepsilon}{2}\right) \leq \frac{4e^{-\frac{\varepsilon^2}{8}k}}{\varepsilon},
\end{aligned} \tag{14}$$

where the last inequality follows from Markov's inequality. \square