An Improved Quasi-Polynomial Algorithm for Approximate Well-Supported Nash Equilibria

Michail Fasoulakis
Institute of Computer Science,
Foundation for Research
and Technology-Hellas
(ICS-FORTH), Greece
mfasoul@ics.forth.gr

Evangelos Markakis
Department of Informatics,
Athens University of
Economics and Business, Greece
markakis@aueb.gr

Abstract

We focus on the problem of *computing approximate Nash equilibria* in *bimatrix games*. In particular, we consider the notion of *approximate well-supported equilibria*, which is one of the standard approaches for approximating equilibria. It is already known that one can compute an ε*-well-supported Nash equilibrium* in time $n^{\mathcal{O}(\log n/\varepsilon^2)}$, for any $\varepsilon > 0$, in games with n *pure strategies* per player. Such a running time is referred to as *quasi-polynomial*. Regarding faster algorithms, it has remained an open problem for many years if we can have better running times for small values of the approximation parameter, and it is only known that we can compute in polynomial-time a 0.6528*-well-supported Nash equilibrium*. In this paper, we investigate further this question and propose a much better *quasi-polynomial time* algorithm that computes a $(1/2 + \varepsilon)$*-well-supported Nash equilibrium* in time $n^{\mathcal{O}(\log \log n^{\frac{1}{\varepsilon}}/\varepsilon^2)}$, for any $\varepsilon > 0$. Our algorithm is based on appropriately combining sampling arguments, support enumeration, and solutions to systems of linear inequalities.

Keywords

Algorithmic game theory, bimatrix games, Nash equilibria, approximate Nash equilibria.

1 Introduction

We revisit a fundamental by now problem in algorithmic game theory: *computing approximate Nash equilibria* in *bimatrix games*. Right from the outset of the interplay between economics and computation, a couple of decades ago, one of the driving research questions has been to determine the *complexity of computing Nash equilibria* in *non-cooperative games*, see Chapter 2 of (Nisan et al. 2007). A *Nash equilibrium* is a strategy profile where no player has any incentive to deviate, and finding equilibria is important, among others, for predicting strategic behavior, and also for understanding the difficulty of reaching stable points in such competitive situations. Eventually, the problem turned out to be much challenging, even for two players, and ever since the first negative results for finding exact equilibria were established in (Daskalakis, Goldberg, and Papadimitriou 2009; Chen, Deng, and Teng 2009), the research community has focused on *computing approximate Nash equilibria*. Although this is a relaxation of the initial algorithmic question, approximate equilibria still form an appealing concept and can be quite useful in practical settings, see e.g. the recent work of (McCarthy et al. 2018).

Even for *approximate Nash equilibria*, efficient algorithms have been hard to obtain for all interesting notions of approximation. The more standard and popular concept of approximation, dictates that no player can gain much by deviating to a different strategy. This is referred to as an ε*-Nash equilibrium*, where ε is the maximum gain allowed by a deviation. So far, it is known that we can compute such approximate equilibria in time $n^{O(\log n/\varepsilon^2)}$, which is a *quasi-polynomial time algorithm (quasi-PTAS)*, for any constant $\varepsilon > 0$ (Lipton, Markakis, and Mehta 2003). If we insist on polynomial time, then efficient algorithms exist only for $\varepsilon \geq 0.3393$ (Tsaknakis and Spirakis 2008). A strengthening of this concept is that of ε*-well-supported Nash equilibrium*, for which the picture is even worse. An ε*-well-supported Nash equilibrium* is a strategy profile in which the expected payoff of any *pure strategy* played with positive probability (i.e., in the *support* of the *mixed strategy* of some player) is at most ε less than the *best-response payoff* against the strategy of the other player. We can still have a *quasi-polynomial time algorithm* for this stronger notion (Kontogiannis and Spirakis 2010), for any constant $\varepsilon > 0$, but when it comes to polynomial time, the best known algorithm works only for $\varepsilon \geq 0.6528$ (Czumaj et al. 2018). Hence, a major open question is whether we can have polynomial time algorithms for smaller values of ε.

Our focus in this work is on the notion of *well-supported equilibria*. Note that since any ε*-well-supported Nash equilibrium* is also an ε*-Nash equilibrium* (the other direction does not always hold), any algorithmic result directly implies the same approximation for ε*-Nash equilibria*. Given the difficulty of going from *quasi-polynomial* to simply polynomial running times, we take an intermediate step and pose the following question: *Is it possible to beat the* 0.6528*-approximation by quasi-polynomial algorithms of better* *running time?* In other words, if we allow non-polynomial time algorithms, can we achieve a better running time than $n^{O(\log n)}$, for values of ε below 0.6528?

1.1 Our Contribution

We provide a positive answer to the above question, and our main result is a new and faster algorithm for finding a $(1/2+\varepsilon)$-*well-supported Nash equilibrium*, for any constant $\varepsilon > 0$ in 2-player games. The complexity of the algorithm is $n^{\mathcal{O}(\log\log n^{\frac{1}{\varepsilon}}/\varepsilon^2)}$, where n is the number of available *pure strategies* to the players. Although this is still not polynomial, the exponent in the running time is significantly improved (from $\mathcal{O}(\log n)$ to $\mathcal{O}(\log\log n)$) and asymptotically closer to a constant than the previously known results. Our result is based on combining sampling arguments and support enumeration with solving carefully constructed linear systems. The starting point to design the algorithm involves sampling followed by enumerating possible strategy profiles, as in some of the previous works (Lipton, Markakis, and Mehta 2003; Kontogiannis and Spirakis 2010). What differs from previous algorithms is that our analysis allows for an enumeration of candidate profiles over a much reduced search space. To achieve this, our main technical insight is captured by Theorem 2 in Section 3, establishing that for one of the players it suffices to focus on *mixed strategies* of support size $\mathcal{O}(\log\log n)$. On top of that, the strategy of the other player can then be found by solving an appropriate system of linear inequalities within each enumeration step.

1.2 Further Related Work

For the simpler class of *zero-sum games*, the works of (Althöfer 1994; Lipton and Young 1994) were the first to demonstrate the existence of approximate *maxmin* strategies with logarithmic *support*. This was extended for general 2-player *normal-form games*, yielding *quasi-polynomial* algorithms for computing both ε-*Nash equilibria* and ε-*well-supported Nash equilibria* in (Lipton, Markakis, and Mehta 2003) and (Kontogiannis and Spirakis 2010), respectively. Similar results for games with a higher number of players have also been obtained, e.g., in (Lipton, Markakis, and Mehta 2003; Babichenko, Barman, and Peretz 2017; Hémon, de Rougemont, and Santha 2008; Czumaj, Fasoulakis, and Jurdziński 2017a).

Regarding polynomial time algorithms for computing additive ε-*Nash equilibria* in bimatrix games, there was a series of works achieving different values of approximation. Namely, for $\varepsilon = 3/4$ (Kontogiannis, Panagopoulou, and Spirakis 2009), for $\varepsilon = 1/2$ (Daskalakis, Mehta, and Papadimitriou 2009), for $\varepsilon \approx 0.38$ (Daskalakis, Mehta, and Papadimitriou 2007; Czumaj, Fasoulakis, and Jurdziński 2017b; Czumaj et al. 2018), for $\varepsilon \approx 0.36$ (Bosse, Byrka, and Markakis 2010), and finally for $\varepsilon = 0.3393$ (Tsaknakis and Spirakis 2008), which is also the currently best known approximation. Furthermore, for *symmetric bimatrix games* there is an algorithm for computing $(1/3+\varepsilon)$-*Nash equilibria* (Kontogiannis and Spirakis 2011), for any $\varepsilon > 0$.

On the other hand, for the stronger notion of approximation of additive ε-*well-supported Nash equilibria*, the picture is more limited. There were polynomial time algorithms for $\varepsilon = 2/3$ (Kontogiannis and Spirakis 2010), for $\varepsilon \approx 0.6619$ (Fearnley et al. 2016), and finally for $\varepsilon = 0.6528$ (Czumaj et al. 2018). For the specific classes of *win-lose* and *symmetric* games, there are polynomial time algorithms achieving an approximation of $1/2$ (Kontogiannis and Spirakis 2010), and $(1/2+\varepsilon)$, for any constant $\varepsilon > 0$ (Czumaj, Fasoulakis, and Jurdziński 2014), respectively. Furthermore, in *bimatrix games* where both payoff matrices are symmetric, a $1/2$-*well-supported Nash equilibrium* can be found in polynomial time (Czumaj, Fasoulakis, and Jurdziński 2017b). Another interesting special case is that of sparse payoff matrices, for which a PTAS has been recently obtained (Barman 2018). Finally, in (Anbalagan et al. 2013; 2015) the authors studied the size of the *support* required at *approximate well-supported Nash equilibria*.

Negative results have also been established, as to what we can hope to achieve algorithmically. The results of (Babichenko, Papadimitriou, and Rubinstein 2016; Rubinstein 2016) show that under certain assumptions, one cannot hope for better than *quasi-polynomial time* algorithms for approximate equilibria. Also, stronger notions of approximation have been proved to be hard in (Etessami and Yannakakis 2010), e.g., if we ask to be geometrically close to an exact equilibrium. Moreover, there are negative results for *approximate Nash equilibria* subject to further constraints (Deligkas, Fearnley, and Savani 2018), e.g., with regard to the social welfare they can guarantee.

Regarding empirical behavior, there have been numerous works, spanning several decades, on heuristic algorithms and experimental comparisons for exact Nash equilibria. These are algorithms that usually have worst-case exponential time, but may behave very well in practice or for special classes of games. This line of works originates with the celebrated *Lemke-Howson algorithm* (Lemke and Howson 1964), and for more recent works, see among others (Bhat and Leyton-Brown 2004; Thompson, Leung, and Leyton-Brown 2011) for the class of action-graph games and (Porter, Nudelman, and Shoham 2008) for the *support enumeration method*. More recently, there have also been experimental evaluations for methods that compute approximate equilibria, as reported in (Tsaknakis, Spirakis, and Kanoulas 2008; Kontogiannis and Spirakis 2011; Fearnley, Igwe, and Savani 2015), highlighting the need for creating new families of testbeds for such algorithms.

2 Definitions

We consider finite, 2-player, *bimatrix games* defined by a pair of matrices, $(R, C) \in [0,1]^{n\times n}$, where R is the *payoff matrix* of the row player and C is the *payoff matrix* of the column player. We assume the matrices are normalized so that every entry lies in the interval $[0,1]$. We also assume each player has n *pure strategies* at her disposal, and we let $[n] = \{1, 2, \ldots, n\}$ denote the set of strategies. This is without loss of generality, since in the case that the players do not have the same number of available *pure strategies*, we can simply add dummy strategies to equalize them. If the row player plays the *pure strategy* i, and the column player plays the *pure strategy* j, then R_{ij} and C_{ij} are the payoffs derived for the row player and the column player respectively.

The players are also allowed to play a *mixed strategy*, which is simply a probability distribution on the set of *pure strategies*. We will denote a *mixed strategy* by a column vector, so that for a *mixed strategy* $x = (x(1), x(2), \ldots, x(n))^T$, the probability $x(i)$ corresponds to the probability of playing the i-th *pure strategy*. Obviously, the probabilities need to satisfy that $\sum_{i\in[n]} x(i) = 1$. *Pure strategies* can also be written as *mixed strategies*, in vector format, and we will denote by e_i the vector corresponding to the i-th pure strategy, i.e., the column vector whose i-th coordinate equals 1 and all the rest are equal to 0. If x is a *mixed strategy* of the row player and y is a *mixed strategy* of the column player, then the expected payoff of the row player equals

$$u_R(x, y) = \sum_{i\in[n]} \sum_{j\in[n]} x(i)y(j)R_{ij} = x^T Ry.$$

Similarly, the expected payoff of the column player can be written as $u_C(x, y) = x^T Cy$.

An important notion both in our work and in many previous related works is the *support* of a *mixed strategy* x, denoted by $supp(x)$. This is simply the set of *pure strategies* that are played with positive probability under the mixed strategy x:

$$supp(x) = \{i \in [n] : x(i) > 0\}.$$

Most algorithms on *approximating Nash equilibria* are based on finding strategies with small *support*. Furthermore, the following form of simplified, small support *mixed strategies* is crucial for our algorithm.

Definition 1. *A mixed strategy x is a k-uniform strategy if and only if $x(i) \in \{0, \frac{1}{k}, \frac{2}{k}, \ldots, 1\}$, for any $i \in [n]$.*

Hence, for a k-*uniform* strategy x, it holds that $|supp(x)| \leq k$, and as an example, the uniform distribution over all strategies is the n-*uniform* strategy $(1/n, 1/n, \ldots, 1/n)^T$. When looking for strategies with small *support* in the sequel, we will focus on k-*uniform* strategies for appropriate values of k. We continue now with defining the relevant equilibrium notions.

Definition 2. *A strategy profile (x^*, y^*) is a Nash equilibrium if and only if, for any $i \in [n]$,*

$$x^{*T} Ry^* \geq e_i^T Ry^*, \text{ and}$$
$$x^{*T} Cy^* \geq x^{*T} Ce_i.$$

In other words the players mutually play *best-response strategies* and no player has an incentive to deviate. Note that in the definition, it suffices to check only deviations to pure strategies (a profitable deviation by a mixed strategy implies there exists a deviation by a pure strategy). Perhaps the most natural way to relax this requirement is the notion of an ε-*Nash equilibrium*.

Definition 3. *A strategy profile (x^*, y^*) is an ε-Nash equilibrium if and only if, for any $i \in [n]$,*

$$x^{*T} Ry^* + \varepsilon \geq e_i^T Ry^*, \text{ and}$$
$$x^{*T} Cy^* + \varepsilon \geq x^{*T} Ce_i.$$

In other words, no player can gain more than ε by deviating to another strategy.

Finally, to define the stronger notion of approximation that we will work with, we will say that a strategy is an ε-*best-response strategy* to the other player's action, if the expected payoff of playing this strategy is at most ε less than the payoff of the *best-response strategy* against the other player's action.

Definition 4. *A strategy profile (x^*, y^*) is an ε-well-supported Nash equilibrium if and only if, for any $i \in [n]$, for any $k \in supp(x^*)$, and for any $\ell \in supp(y^*)$,*

$$e_k^T Ry^* + \varepsilon \geq e_i^T Ry^*, \text{ and}$$
$$x^{*T} Ce_\ell + \varepsilon \geq x^{*T} Ce_i.$$

Hence, in an ε-well-supported Nash equilibrium, any strategy in the support of x^* is an ε-*best-response strategy* against y^* and vice versa. It is easy to see that Definition 4 implies Definition 3, thus an ε-*well-supported Nash equilibrium* is also an ε-*Nash equilibrium*.

3 Computing approximate well-supported Nash equilibria

As already mentioned, the currently best approximation by polynomial time algorithms achieves a 0.6528-*well-supported Nash equilibrium*. Furthermore, in the work of (Czumaj, Fasoulakis, and Jurdziński 2014), the authors provided a polynomial time algorithm for computing a $(1/2 + \varepsilon)$-*well-supported Nash equilibrium*, for any constant $\varepsilon > 0$, but under the condition that there exists an exact *Nash equilibrium*, where the payoffs for both players are simultaneously no greater or greater than $1/2$. Specifically, as it was pointed out in section 3.3 of (Czumaj, Fasoulakis, and Jurdziński 2014), in the case when both values of a *Nash equilibrium* are no greater than 1/2, a $1/2$-*well-supported Nash equilibrium* can be found by using a linear program in polynomial time. On ther other hand, in the case when both values are greater than 1/2, a $(1/2 + \varepsilon)$-*well-supported Nash equilibrium* can be found with support enumeration of strategies with small support and systems of linear constraints in polynomial time. We state their result in the following theorem.

Theorem 1. *(Czumaj, Fasoulakis, and Jurdziński 2014). For bimatrix games where there exists a Nash equilibrium such that the payoffs for the two players are either both greater or both no greater than 1/2, there is a polynomial time algorithm for computing a $(1/2 + \varepsilon)$-well-supported Nash equilibrium, for any constant $\varepsilon > 0$.*

Thus, to obtain a polynomial time $(1/2 + \varepsilon)$-approximation, it suffices to handle the remaining case, i.e., games where in all exact *Nash equilibria*, the payoff of one player is no greater than 1/2 and the payoff of the other player is greater than 1/2. The main insight of our work is the following theorem, which establishes that for such games, we can focus on strategies with much smaller support for one of the two players.

Theorem 2. *Let $(R, C) \in [0,1]^{n \times n}$ be a bimatrix game, such that there exists a Nash equilibrium (x^*, y^*) where the payoff of the row player satisfies that $u_R = x^{*T} R y^* \in (\frac{1}{2}, 1]$, and the payoff of the column player satisfies that $u_C = x^{*T} C y^* \in [0, \frac{1}{2}]$. Then, for any $\varepsilon > 0$, there exists a strategy profile $(\hat{x}, \hat{y})$, such that,*

- *$\hat{x}$ is a k-uniform strategy with $k = \ln(n)/\varepsilon^2$, and $\hat{y}$ is an ℓ-uniform strategy with $\ell = \ln\ln(n^{\frac{1}{\varepsilon}})/\varepsilon^2$,*
- *$supp(\hat{x}) \subseteq supp(x^*)$, and $supp(\hat{y}) \subseteq supp(y^*)$,*
- *$(\hat{x}, \hat{y})$ is a $(1/2+\varepsilon)$-well-supported Nash equilibrium,*
- *for any $i \in supp(\hat{x})$, $e_i^T R\hat{y} \geq u_R - \varepsilon > 1/2 - \varepsilon$,*
- *for any $j \in [n]$, $\hat{x}^T C e_j \leq u_C + \varepsilon \leq 1/2 + \varepsilon$.*

We postpone the proof of Theorem 2 till Section 4, but we elaborate further now on some consequences. First, an analogous statement is true in the reverse case where $u_R \in [0, \frac{1}{2}]$, and $u_C \in (\frac{1}{2}, 1]$. Second, the most important aspect of the theorem is the fact that $|supp(\hat{y})| \leq \ln\ln(n^{\frac{1}{\varepsilon}})/\varepsilon^2$. Although Theorem 2 is purely an existential statement, it allows for a brute-force enumeration of such strategies in time that is much better than $n^{O(\log n)}$. Finally, one still needs to worry that the support of $\hat{x}$ may be larger, but as we will see, this can be taken care of by avoiding any enumeration for this player, and solving instead a linear system of inequalities.

Algorithm 1

Input: A bimatrix game $(R, C) \in [0,1]^{n \times n}$ and a parameter $\varepsilon \in (0, 1]$.

1. Run the algorithm of (Czumaj, Fasoulakis, and Jurdziński 2014). If you find a $(1/2+\varepsilon)$-well-supported Nash equilibrium, then STOP.
2. Enumerate all the multisets of size $(\ln\ln n^{\frac{1}{\varepsilon}})/\varepsilon^2$ among the pure strategies of the column player.
   - Every multiset corresponds to an ℓ-uniform strategy $\hat{y}$ for the column player, with $\ell = (\ln\ln n^{\frac{1}{\varepsilon}})/\varepsilon^2$. Given $\hat{y}$, identify the set $A \subseteq [n]$ of the pure strategies i of the row player such that $e_i^T R\hat{y} > 1/2 - \varepsilon$.
   - For this set A, solve the linear system (1), described below. If you find a solution $\hat{x}$, then STOP and return the profile $(\hat{x}, \hat{y})$.
3. Enumerate all the multisets of size $(\ln\ln n^{\frac{1}{\varepsilon}})/\varepsilon^2$ for the row player.
   - As before, every such multiset corresponds to a strategy $\hat{x}$ of the row player with support at most $(\ln\ln n^{\frac{1}{\varepsilon}})/\varepsilon^2$. Given $\hat{x}$, find the set $B \subseteq [n]$ of the pure strategies j of the column player such that $\hat{x}^T C e_j > 1/2 - \varepsilon$.
   - For this set B, solve the linear system (2). If you find a solution $\hat{y}$, then STOP and return the profile $(\hat{x}, \hat{y})$.

Linear System (1)

$$\begin{aligned}
&\hat{x}(i) \geq 0, \quad \forall i \in A,\\
&\hat{x}(j) = 0, \quad \forall j \notin A,\\
&\sum_{i=1}^{n} \hat{x}(i) = 1,\\
&\hat{x}^T C e_j \leq 1/2 + \varepsilon, \text{ for any j} \in [n].
\end{aligned}$$

Linear System (2)

$$\begin{aligned}
&\hat{y}(j) \geq 0, \quad \forall j \in B,\\
&\hat{y}(i) = 0, \quad \forall i \notin B,\\
&\sum_{j=1}^{n} \hat{y}(j) = 1,\\
&e_i^T R\hat{y} \leq 1/2 + \varepsilon, \text{ for any } i \in [n].
\end{aligned}$$

3.1 The algorithm

We are now ready to state the algorithm for *computing* $(1/2+\varepsilon)$*-well-supported Nash equilibria*, for any constant $\varepsilon > 0$. Since we cannot compute an exact *Nash equilibrium*, we first try to see if the input game falls within the cases of Theorem 1, by running the algorithm of (Czumaj, Fasoulakis, and Jurdziński 2014). If this does not succeed, we then make use of Theorem 2, and we explore both possible cases: $u_C \in [0, \frac{1}{2}]$, $u_R \in (\frac{1}{2}, 1]$, and $u_R \in [0, \frac{1}{2}]$, $u_C \in (\frac{1}{2}, 1]$.

Given Theorem 2, the following statement is the main conclusion of our work.

Corollary 1. *For any bimatrix game $(R, C) \in [0,1]^{n \times n}$, there is an algorithm for computing a $(1/2+\varepsilon)$-well-supported Nash equilibrium, for any constant $\varepsilon > 0$, in time $n^{\mathcal{O}(\log\log n^{\frac{1}{\varepsilon}}/\varepsilon^2)}$.*

Proof. We know that every game has at least one exact *Nash equilibrium* by Nash's theorem (Nash 1951). Fix one such equilibrium and let u_R be the payoff of the row player and u_C be the payoff of the column player respectively. Note that we cannot know in advance any information about u_R and u_C, and hence we need to explore all possible cases. If both payoff values are no greater than 1/2 or both greater than 1/2, then step 1 of our algorithm will terminate with a $(1/2+\varepsilon)$*-well-supported Nash equilibrium* in polynomial-time, given the result of (Czumaj, Fasoulakis, and Jurdziński 2014). Suppose now that $u_R \in (1/2, 1]$ and $u_C \in [0, 1/2]$, which corresponds to step 2 of our algorithm. By Theorem 2, we know there exists a strategy profile that has the properties we need, along with additional properties on the payoffs (conditions 4 and 5 of Theorem 2). It is precisely this strategy profile that we attempt to find in step 2 of the algorithm. Since we know its existence, we then construct the set A as the subset of pure strategies that satisfy condition 4 of Theorem 2 and we give this as input to the linear system. Hence, we have a guarantee that there will be at least one occasion in the enumeration process where the linear system will have a solution and the profile returned will satisfy the desired approximation. Here we need to clarify that our algorithm may return a strategy profile $(\hat{x}, \hat{y})$ that is a $(1/2 + \varepsilon)$-*well-supported equilibrium* but different from the strategy profile described by Theorem 2. This is of no concern, since our goal is still achieved.

In the case that $u_C \in (1/2, 1]$ and $u_R \in [0, 1/2]$, then we can apply completely analogous arguments as in step 2 of the algorithm. Hence, in all cases, the algorithm returns a $(1/2 + \varepsilon)$-*well-supported Nash equilibrium*.

To find the complexity of the algorithm we need to find first the number of all possible multisets of size $(\ln \ln n^{\frac{1}{\varepsilon}})/\varepsilon^2$, which is equal to

$$\binom{n + (\ln \ln n^{\frac{1}{\varepsilon}})/\varepsilon^2 - 1}{(\ln \ln n^{\frac{1}{\varepsilon}})/\varepsilon^2} = n^{\mathcal{O}((\log \log n^{\frac{1}{\varepsilon}})/\varepsilon^2)}.$$

In worst case, we need to examine these multisets both for the row and the column player. For any multiset under examination, the algorithm performs polynomially many steps. In particular, note that for every multiset, we need to solve a system of linear inequalities. But since this can be done in polynomial time, the total complexity of the algorithm is $n^{\mathcal{O}((\log \log n^{\frac{1}{\varepsilon}})/\varepsilon^2)}$, as claimed. □

4 Proving Theorem 2

For the proof of Theorem 2, we will give a probabilistic argument inspired by previous works such as (Althöfer 1994; Lipton and Young 1994; Lipton, Markakis, and Mehta 2003; Czumaj, Fasoulakis, and Jurdziński 2014). In particular, given a *Nash equilibrium* profile (x^*, y^*), satisfying the properties stated in the theorem, we will perform sampling with replacement on the equilibrium strategies x^* and y^*, creating two empirical distributions $\hat{x}$ and $\hat{y}$. Note that Theorem 2 is purely an existential result, we do not need to compute the *Nash equilibrium* from which we do our sampling.

We will analyze the empirical distributions we create by using the well-known *Hoeffding's inequality*, which is stated next.

Lemma 1 (Hoeffding's inequality (Hoeffding 1963))**.** *Let $Y_1, Y_2, \ldots, Y_\ell$, be ℓ independent random variables in the interval $[0,1]$, with $\overline{Y} = \frac{1}{\ell}\sum_{i=1}^{\ell} Y_i$, and $E[\overline{Y}]$ be the expectation of of the random variable $\overline{Y}$. Then, for any $\varepsilon > 0$, it holds that $\mathbf{Pr}\big[|\overline{Y} - E[\overline{Y}]| \geq \varepsilon\big] \leq 2e^{-2\varepsilon^2 \ell}$.*

Proof of Theorem 2. Fix $\varepsilon > 0$, and consider a game that has a *Nash equilibrium* profile (x^*, y^*), for which the expected payoffs of the two players satisfy that $u_R = x^{*T} R y^* \in (\frac{1}{2}, 1]$ and $u_C = x^{*T} C y^* \in [0, \frac{1}{2}]$. We sample with replacement on the equilibrium strategies x^* and y^*, creating two empirical distributions $\hat{x}$ and $\hat{y}$, respectively. Specifically, for the row player we sample ℓ times, with $\ell = \ln(n)/\varepsilon^2$, from the *pure strategies* in the *support* of the *mixed strategy* x^*, according to the probability distribution of x^*. This corresponds to sampling a multiset among the *support* of x^*, since each *pure strategy* may be sampled multiple times. Hence, the sampling process yields an ℓ-*uniform mixed strategy* $\hat{x}$ (the empirical distribution), where for each $i \in [n]$, the probability $\hat{x}(i)$ is the fraction of the number of appearances of *pure strategy* i in the sampling, divided by ℓ. For the column player we do a similar procedure creating a k-*uniform strategy profile* $\hat{y}$, with $k = \ln(\ln(n^{\frac{1}{\varepsilon}}))/\varepsilon^2$, by sampling from the distribution y^*. This way, we satisfy the first condition of Theorem 2.

We take a closer look now at distribution $\hat{y}$. For each sample $t \in [k]$, let J_t be a random variable, whose value equals the index of the column that we sampled. Since we sample from y^*, we have that $\mathbf{Pr}\big[J_t = j\big] = y^*(j)$. The distribution $\hat{y}$ that we create from y^* can be written as $\hat{y} = (\hat{y}(1), \hat{y}(2), \ldots, \hat{y}(n))^T$ with

$$\hat{y}(j) = \frac{1}{k}\sum_{t=1}^{k} \mathbb{1}_{\{J_t = j\}}.$$

By the way that $\hat{y}$ is created from y^*, it trivially follows that $supp(\hat{y}) \subseteq supp(y^*)$ (and similarly $supp(\hat{x}) \subseteq supp(x^*)$), thus satisfying the second condition of Theorem 2. We now focus on establishing the last three conditions of Theorem 2 and consider the random variable $e_i^T R \hat{y}$, for any $i \in [n]$, which expresses the expected payoff of the row player when she chooses her i-th row while the column player plays $\hat{y}$. We can write this as: $e_i^T R \hat{y} = \frac{1}{k}\sum_{t=1}^{k} R_{i,J_t}$. To derive the expectation of this expression, note that

$$E[R_{i,J_t}] = \sum_{j=1}^{n} \mathbf{Pr}\big[J_t = j\big] R_{ij} = \sum_{j=1}^{n} y^*(j) R_{ij} = e_i^T R y^*.$$

This directly implies that the expectation of $e_i^T R \hat{y}$ is equal to $e_i^T R y^*$. Thus, we can apply *Hoeffding's inequality*[1] and establish that for any $i \in [n]$:

$$\begin{aligned}\mathbf{Pr}\big[e_i^T R \hat{y} - e_i^T R y^* \leq -\varepsilon\big] &= \mathbf{Pr}\big[-e_i^T R \hat{y} + e_i^T R y^* \geq \varepsilon\big] \\ &\leq e^{-2\varepsilon^2 k} = \frac{\varepsilon^2}{\ln^2(n)}. \quad (1)\end{aligned}$$

Similarly, for the strategy $\hat{x}$, and for any $j \in [n]$, we can use *Hoeffding's inequality* and obtain that

$$\mathbf{Pr}\big[|\hat{x}^T C e_j - x^{*T} C e_j| \geq \varepsilon\big] \leq 2e^{-2\varepsilon^2 \ell} = \frac{2}{n^2}. \quad (2)$$

Moving on, we provide an auxiliary observation that we

[1] To derive inequality (1), we use a simpler variant of *Hoeffding's inequality*, since we do not need to bound the absolute value of the difference.

will use shortly. In particular, we have that

$$\begin{aligned}
&\mathbf{Pr}\big[\exists i \in supp(\hat{x}) \text{ s.t. } e_i^T R\hat{y} \le e_i^T Ry^* - \varepsilon\big] \\
&= \mathbf{Pr}\big[\exists i \in supp(\hat{x}) \text{ s.t. } e_i^T R\hat{y} \le u_R - \varepsilon\big] \\
&\le \sum_{i \in supp(\hat{x})} \mathbf{Pr}\big[e_i^T R\hat{y} \le u_R - \varepsilon\big] \\
&\le \frac{\ln(n)}{\varepsilon^2} \cdot \frac{\varepsilon^2}{\ln^2(n)} = \frac{1}{\ln(n)}. \qquad (3)
\end{aligned}$$

To see the above implications, recall the fact that $u_R = e_i^T Ry^*$, for any $i \in supp(\hat{x}) \subseteq supp(x^*)$ (by the definition of *Nash equilibrium*). The first inequality holds by the *Union bound* and the second inequality holds by (1) and the fact that $|supp(\hat{x})| \le \ln(n)/\varepsilon^2$.

In a similar manner, using (2), we can prove that

$$\begin{aligned}
&\mathbf{Pr}\big[\exists j \in [n] \text{ s.t. } \hat{x}^T Ce_j \ge {x^*}^T Ce_j + \varepsilon\big] \\
&\le \sum_{j \in [n]} \mathbf{Pr}\big[\hat{x}^T Ce_j \ge {x^*}^T Ce_j + \varepsilon\big] \le n\frac{2}{n^2} = \frac{2}{n}. \qquad (4)
\end{aligned}$$

Our goal is to prove that

$$\begin{aligned}
&\mathbf{Pr}\big[e_i^T R\hat{y} \ge u_R - \varepsilon, \text{ for all } i \in supp(\hat{x}), \\
&\qquad \text{and } \hat{x}^T Ce_j \le u_C + \varepsilon, \text{ for all } j \in [n]\big] > 0, \qquad (5)
\end{aligned}$$

and since by the *Nash equilibrium* definition we have that ${x^*}^T Ce_j \le u_C$, for any j, it suffices to prove that

$$\begin{aligned}
&\mathbf{Pr}\big[e_i^T R\hat{y} \ge u_R - \varepsilon, \text{ for all } i \in supp(\hat{x}), \\
&\qquad \text{and } \hat{x}^T Ce_j \le {x^*}^T Ce_j + \varepsilon, \text{ for all } j \in [n]\big] > 0.
\end{aligned}$$

Looking at the complement of this event, we obtain that

$$\begin{aligned}
&\mathbf{Pr}\big[e_i^T R\hat{y} < u_R - \varepsilon, \text{ for some } i \in supp(\hat{x}), \\
&\qquad \text{or } \hat{x}^T Ce_j > {x^*}^T Ce_j + \varepsilon, \text{ for some } j\big] \\
&\le \sum_{I \subseteq supp(x^*)} \Big(\mathbf{Pr}\big[supp(\hat{x}) = I\big] \cdot \\
&\mathbf{Pr}\big[e_i^T R\hat{y} \le u_R - \varepsilon, \text{ for some } i \in I | supp(\hat{x}) = I\big]\Big) \\
&\quad + \sum_{j \in [n]} \mathbf{Pr}\big[\hat{x}^T Ce_j \ge {x^*}^T Ce_j + \varepsilon\big] \\
&\le \frac{1}{\ln(n)} \sum_{I \subseteq supp(x^*)} \mathbf{Pr}\big[supp(\hat{x}) = I\big] + \frac{2}{n} \\
&= \frac{1}{\ln(n)} + \frac{2}{n} < 1.
\end{aligned}$$

The second inequality holds by (3) and (4). It is easy to see that the last inequality holds for any $n \ge 6$. Thus, we have proved that the probability of our desired event is strictly positive, which means that there exists a strategy profile $(\hat{x}, \hat{y})$ such that inequality (5) holds, or in other words there is a strategy profile $(\hat{x}, \hat{y})$ such that for any $i \in supp(\hat{x})$, $e_i^T R\hat{y} \ge u_R - \varepsilon > 1/2 - \varepsilon$, and for any $j \in [n]$, $\hat{x}^T Ce_j \le u_C + \varepsilon \le 1/2 + \varepsilon$, since $u_R > 1/2$ and $u_C \le 1/2$. This means we have satisfied the fourth and the fifth condition of the theorem.

To conclude the proof, we now establish that the strategy profile $(\hat{x}, \hat{y})$ is a $(1/2 + \varepsilon)$-*well-supported Nash equilibrium*. Let i^* be the *best-response strategy* of the row player against the strategy $\hat{y}$ of the column player. Then, we have that, for any $i \in supp(\hat{x})$,

$$e_{i^*}^T R\hat{y} - e_i^T R\hat{y} \le 1 - u_R + \varepsilon < 1/2 + \varepsilon,$$

the first inequality holds by the fourth condition of the theorem, and the second by the fact that $u_R > 1/2$. On the other hand, let j^* be the *best-response strategy* of the column player against the strategy $\hat{x}$ of the row player. We have that, for any j,

$$\hat{x}^T Ce_{j^*} - \hat{x}^T Ce_j \le u_C + \varepsilon \le 1/2 + \varepsilon,$$

the first inequality holds by the fifth condition of the theorem, and the fact that $\hat{x}^T Ce_j \ge 0$, for any j. The second inequality holds since $u_C \le 1/2$. $\square$

5 Conclusions

We have provided an improved algorithm for obtaining a $(1/2+\varepsilon)$-well-supported Nash equilibrium. We believe that given this positive result, even better algorithms may still be possible for this stronger notion of approximation. A major open question that still remains is whether there exists a polynomial time algorithm for finding a $1/2$-*well-supported Nash equilibrium*.

We would also like to explore possible applications of our technique for other notions of approximation. As an example, for ε-*Nash equilibria*, we only know of $n^{O(\log n)}$ algorithms when $\varepsilon < 0.3393$. It would be exciting if we can provide improved running times for low values of ε beyond 0.3393.

6 Acknowledgments

For this work the first author is supported by the Stavros Niarchos-FORTH post-doc fellowship for the project ARCHERS.

References

Althöfer, I. 1994. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and Its Applications* 199:339–355.

Anbalagan, Y.; Norin, S.; Savani, R.; and Vetta, A. 2013. Polylogarithmic supports are required for approximate well-supported Nash equilibria below 2/3. In *Proceedings of the 9th International Conference Web and Internet Economics (WINE'13)*, 15–23.

Anbalagan, Y.; Huang, H.; Lovett, S.; Norin, S.; Vetta, A.; and Wu, H. 2015. Large supports are required for well-supported Nash equilibria. In *Proceedings of the 18th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'15)*, 78–84.

Babichenko, Y.; Barman, S.; and Peretz, R. 2017. Empirical distribution of equilibrium play and its testing application. *Math. Oper. Res.* 42(1):15–29.

Babichenko, Y.; Papadimitriou, C. H.; and Rubinstein, A. 2016. Can almost everybody be almost happy? In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science (ITCS'16)*, 1–9.

Barman, S. 2018. Approximating Nash equilibria and dense subgraphs via an approximate version of Carathéodory's theorem. *SIAM J. Comput.* 47(3):960–981.

Bhat, N. A. R., and Leyton-Brown, K. 2004. Computing Nash equilibria of action-graph games. In *Proceedings of the 20th Conference in Uncertainty in Artificial Intelligence (UAI'04)*, 35–42.

Bosse, H.; Byrka, J.; and Markakis, E. 2010. New algorithms for approximate Nash equilibria in bimatrix games. *Theoretical Computer Science* 411(1):164–173.

Chen, X.; Deng, X.; and Teng, S.-H. 2009. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM* 56(3):14:1–14:57.

Czumaj, A.; Deligkas, A.; Fasoulakis, M.; Fearnley, J.; Jurdziński, M.; and Savani, R. 2018. Distributed methods for computing approximate equilibria. *Algorithmica*.

Czumaj, A.; Fasoulakis, M.; and Jurdziński, M. 2014. Approximate well-supported Nash equilibria in symmetric bimatrix games. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT'14)*, 244–254.

Czumaj, A.; Fasoulakis, M.; and Jurdziński, M. 2017a. Multi-player approximate Nash equilibria. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS'17)*, 1511–1513.

Czumaj, A.; Fasoulakis, M.; and Jurdziński, M. 2017b. Zero-sum game techniques for approximate Nash equilibria. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS'17)*, 1514–1516.

Daskalakis, C.; Goldberg, P. W.; and Papadimitriou, C. H. 2009. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing* 39(1):195–259.

Daskalakis, C.; Mehta, A.; and Papadimitriou, C. 2007. Progress in approximate Nash equilibria. In *Proceedings of the 8th (ACM) Conference on Electronic Commerce (EC'07)*, 355–358.

Daskalakis, C.; Mehta, A.; and Papadimitriou, C. 2009. A note on approximate Nash equilibria. *Theoretical Computer Science* 410(17):1581–1588.

Deligkas, A.; Fearnley, J.; and Savani, R. 2018. Inapproximability results for constrained approximate Nash equilibria. *Inf. Comput.* 262(Part 1):40–56.

Etessami, K., and Yannakakis, M. 2010. On the complexity of Nash equilibria and other fixed points. *SIAM J. of Computing* 39(6):2531–2597.

Fearnley, J.; Goldberg, P. W.; Savani, R.; and Sørensen, T. B. 2016. Approximate well-supported Nash equilibrium below two-thirds. *Algorithmica* 76(2):297–319.

Fearnley, J.; Igwe, T. P.; and Savani, R. 2015. An empirical study of finding approximate equilibria in bimatrix games. In *Proceedings of the 14th International Symposium on Experimental Algorithms (SEA'15)*, 339–351.

Hémon, S.; de Rougemont, M.; and Santha, M. 2008. Approximate Nash equilibria for multi-player games. In *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT'08)*, 267–278.

Hoeffding, W. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58(301):13–30.

Kontogiannis, S. C., and Spirakis, P. G. 2010. Well supported approximate equilibria in bimatrix games. *Algorithmica* 57(4):653–667.

Kontogiannis, S. C., and Spirakis, P. G. 2011. Approximability of symmetric bimatrix games and related experiments. In *Proceedings of the 10th International Symposium on Experimental Algorithms (SEA'11)*, 1–20.

Kontogiannis, S. C.; Panagopoulou, P. N.; and Spirakis, P. G. 2009. Polynomial algorithms for approximating Nash equilibria of bimatrix games. *Theor. Comput. Sci.* 410(17):1599–1606.

Lemke, C. E., and Howson, J. T. 1964. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics* 12:413–423.

Lipton, R. J., and Young, N. E. 1994. 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 (STOC'94)*, 734–740.

Lipton, R.; Markakis, E.; and Mehta, A. 2003. Playing large games using simple strategies. In *Proceedings of the 4th (ACM) Conf. on Electronic Commerce (EC'03)*, 36–41.

McCarthy, S. M.; Laan, C. M.; Wang, K.; Vayanos, P.; Sinha, A.; and Tambe, M. 2018. The price of usability: Designing operationalizable strategies for security games. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, (IJCAI '18)*, 454–460.

Nash, J. 1951. Non-cooperative games. *Annals of Mathematics* 54(2):286–295.

Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic Game Theory*. Cambridge Univ. Press.

Porter, R.; Nudelman, E.; and Shoham, Y. 2008. Simple search methods for finding a Nash equilibrium. *Games and Economic Behavior* 63(2):642–662.

Rubinstein, A. 2016. Settling the complexity of computing approximate two-player Nash equilibria. In *IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS'16)*, 258–265.

Thompson, D. R. M.; Leung, S.; and Leyton-Brown, K. 2011. Computing Nash equilibria of action-graph games via support enumeration. In *Proceedings of the 7th International Workshop on Internet and Network Economics (WINE'11)*, 338–350.

Tsaknakis, H., and Spirakis, P. G. 2008. An optimization approach for approximate Nash equilibria. *Internet Mathematics* 5(4):365–382.

Tsaknakis, H.; Spirakis, P. G.; and Kanoulas, D. 2008. Performance evaluation of a descent algorithm for bi-matrix games. In *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE'08)*, 222–230.