Approximate Nash Equilibria of Imitation Games: Algorithms and Complexity

Aniket Murhekar
University of Illinois, Urbana-Champaign
aniket2@illinois.edu

Ruta Mehta*
University of Illinois, Urbana-Champaign
rutameht@illinois.edu

ABSTRACT

A two-player finite game is represented by two payoff matrices (A, B), one for each player. *Imitation games* are a subclass of two-player games in which B is the identity matrix, implying that the second player gets a positive payoff only if she "imitates" the first. Given that the problem of computing a Nash equilibrium (NE) is known to be provably hard, even to approximate, we ask if it is any easier for imitation games.

We show that much like the general case, for any $c > 0$, computing a $\frac{1}{n^c}$-approximate NE of imitation games remains PPAD-hard, where n is the number of moves available to the players. On the other hand, we design a polynomial-time algorithm to find ϵ-approximate NE for any given constant $\epsilon > 0$ (PTAS). The former result also rules out the smooth complexity being in P, unless PPAD $\subset$ RP.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory**; **Exact and approximate computation of equilibria**;

KEYWORDS

Two-player games; Imitation games; Nash equilibrium; PTAS; PPAD-hardness

ACM Reference Format:
Aniket Murhekar and Ruta Mehta. 2020. Approximate Nash Equilibria of Imitation Games: Algorithms and Complexity. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 8 pages.

1 INTRODUCTION

Nash equilibrium is arguably one of the most fundamental solution concepts in game theory [31]. It is a state in which no individual can gain by deviating unilaterally. In the previous two decades or more, the field of algorithmic game theory has extensively studied the computability of Nash equilibrium in various games, especially in two-player finite games [7, 16, 33]. Such a game can be represented by two payoff matrices (A, B), one for each player, where a play can be thought of as the first player choosing a row and the second choosing a column.

Computing a Nash equilibrium (NE) of a general two-player game was shown to be PPAD-complete by a series of remarkable results in 2006 [7, 13, 33]; PPAD is a complexity class introduced in [33]. Even computing ϵ-approximate NE (ϵ-NE) for $\epsilon = \frac{1}{poly(n)}$ remains PPAD-complete [7], where n is the number of rows/columns in A and B; at an ϵ-NE no player can achieve more than ϵ additive gain by deviating unilaterally. On the other hand, for a constant ϵ, a quasi-polynomial-time algorithm to find ϵ-NE is known since 2003 [25], but there has been no improvement on this front since then. Recently, this result was shown to be optimal assuming the *exponential time hypothesis* for PPAD [34]. In the light of these negative results, various subclasses of two-player games, like win-lose games, sparse games and constant-rank games have been analyzed both for exact and approximate NE [1, 8, 10, 20] (see Section 1.1 for a detailed discussion).

In this paper we study the complexity of finding an (approximate) NE for one such subclass called *imitation games*. In such a game [28] one of the players, say the second player, is an *imitator*. The imitator gets a payoff of 1 only when she "imitates" the strategy of the other player, and 0 otherwise, and thus her payoff matrix B is an identity matrix. Imitation games are interesting because the symmetric NE of a symmetric bimatrix game are in one-to-one correspondence with the NE strategies of the imitator in an imitation game ([11, 27]). They have also been employed to study the complexity of various computational problems, like providing an alternate proof of the Kakutani fixed point theorem that is brief and elementary [27], relating the Lemke-Howson and Lemke paths's algorithm [28], and other problems on equilibria of two player games (*e.g.,* [11, 17, 29, 30]).

The problem of finding an exact NE in imitation games is PPAD-complete since the same problem on symmetric games reduces to it. However, to the best of our knowledge, the complexity of finding an approximate NE remains unknown. In this paper we obtain the following set of results concerning imitation games: settling the complexity of approximate NE for imitation games (and in doing so, symmetric games), and the smoothed complexity. We also obtain results for a stronger notion of approximation, called approximate *well-supported* Nash equilibrium (wsNE). At an ϵ-wsNE players play a pure strategy with positive probability only if it gives maximum payoff within an additive ϵ.

Our contributions.

- We design a polynomial-time algorithm to find an ϵ-approximate-well-supported NE for a constant $\epsilon > 0$ (PTAS), that runs in time $n^{O(1/\epsilon)}$ poly$(\mathcal{L})$, where $\mathcal{L}$ is the bit-size of the input (see Section 3).
- We show PPAD-hardness for the problem of finding a $\frac{1}{n^c}$-approximate-well-supported NE, and thereby also for $\frac{1}{n^c}$-approximate NE, for any $c > 0$. This hardness result rules out any FPTAS for this problem unless PPAD $\subset$ P (see Section 4).

*The work was supported by NSF CAREER grant CCF 1750436.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

In showing the above, we also prove that computing a symmetric $\frac{1}{n^c}$-approximate-well-supported NE of a symmetric game is also PPAD-hard, for any $c > 0$.

- Towards beyond worst-case complexity, we infer that the above PPAD-hardness result together with a result of [7] rules out the smoothed complexity being in P unless PPAD $\subset$ RP.

1.1 Related work

The Lemke-Howson algorithm [24] is the oldest known algorithm to find an exact Nash equilibrium in general two-player games, and is also the only non-enumerative algorithm for the problem. However it was shown to take exponentially many steps in the worst-case [35]. Efficient algorithms were obtained for special cases, like zero-sum games where $B = -A$ [38], when rank of A or B is a constant [19, 25], or when $rank(A + B) = 1$ [1].

The complexity of finding NE was shown to be PPAD-complete, even for $1/\text{poly}(n)$ approximation [7, 16, 33], that is, an FPTAS for this problem is unlikely unless PPAD $\subset$ P. This was followed by a number of results showing PPAD-hardness for important subclasses: exact NE in constant-rank games [30], exact as well as approximate NE in sparse games ([8]), win-lose games ([10]), and most recently sparse win-lose games ([26]). The hardness of several related decision problems about Nash equilibria in symmetric win-lose bimatrix games were considered in [5]. On the other hand efficient algorithms were obtained to find approximate NE for subclasses like low rank games [2, 20] (FPTAS), and when $(A + B)$ is sparse [4] (PTAS).

Towards constant approximation a $n^{O(\log n/\epsilon^2)}$-time algorithm is known for ϵ-NE [25], and is the best possible assuming exponential time hypothesis for PPAD [34]. While [14] showed existence of $\frac{1}{2}$-NE with support size at most two, [22] gave an efficient algorithm to find $\frac{3}{4}$-NE, and more generally $\frac{2+\lambda}{4}$-NE, where λ is the minimum expected payoff to any player at any Nash equilibrium. There have been several other approaches to compute an ϵ-NE for *constant* ϵ, see for e.g. [6, 15, 37]), with $\epsilon = 0.3393$ being the best so far. Computing ϵ-NE in subclasses has also been studied, relying on the properties of the payoff matrices. See for example [21] for a polynomial time algorithm to compute a $(\frac{1}{3}+\delta)$-NE for a symmetric game, and [37] for a polynomial time algorithm to compute a $\frac{1}{2}$-NE in win-lose games.

Turning to approximate-well-supported Nash equilibrium, [14] showed that computing 5/6-wsNE is possible in polynomial time, assuming a graph theoretic conjecture. A polynomial time algorithm to compute a ϵ-wsNE where ϵ is just above 0.6619 was shown in [18]. For special cases, [23] provided polynomial time algorithms (based on the solvability of zero sum games) for constructing a $\frac{1}{2}$-wsNE for win-lose games and $\frac{2}{3}$-wsNE for normalized games. For symmetric games, [12] provided a linear programming approach to compute a $(\frac{1}{2} + \delta)$-wsNE, for an arbitrarily small constant $\delta > 0$, in polynomial time.

Smoothed analysis is a beyond-worst-case analysis technique which was introduced in [36]. It seeks to show that worst-case instances are sparse and scattered. That is, the smoothed complexity of a problem is in P, if any instance can be solved in polynomial time after subjecting it to independent random perturbations. Using PPAD-hardness for computing $1/poly(n)$-NE, [7] shown that unless PPAD $\subset$ RP, it is unlikely that smoothed complexity of computing a NE is polynomial. Towards the average case, [3] considered random two-player games where all payoffs are i.i.d. random variables in $[0, 1]$ following either the normal or the uniform distribution. They show that with probability at least $1 - O(1/\log n)$, there exists a Nash equilibrium with support of size two. Using this observation, they present a $O(m^2 n \log\log n + n^2 m \log\log m)$-expected time Las Vegas algorithm for finding a Nash equilibrium in such games. It was shown by [32] that in random bimatrix games, where each player's payoffs are bounded and independent random variables with common expectations, the completely mixed uniform strategy profile is an $\tilde{O}(\frac{1}{\sqrt{n}})$-NE with high probability.

The computational complexity of finding Nash equilibria in imitation games has not been studied to the best of our knowledge.

2 PRELIMINARIES

Let $[m] = \{1, 2, \ldots, m\}$ for any $m \in \mathbb{N}$. For $a, b \in \mathbb{R}$, the interval $[a, b]$ is the set $\{x : a \le x \le b\}$, and (a, b) is the set $[a, b] \setminus \{a, b\}$. A $m \times n$ matrix M with entries from set S is denoted as $M \in S^{m\times n}$, and its entries are denoted with the corresponding lowercase letter indexed by the row and column numbers. That is, for an $m \times n$ matrix M, its $(i, j)^{th}$ entry is denoted by $m_{ij} \in S$, where $i \in [m]$ and $j \in [n]$. For a constant c, $M + c$ and cM are the matrices M' and M'' of dimensions $m \times n$ given by $m'_{ij} = m_{ij} + c$ and $m''_{ij} = c \cdot m_{ij}$, respectively, for all $i \in [m], j \in [n]$. We denote by I an identity matrix, whose dimension will be clear from the context. A vector $\mathbf{x}$ is a $m \times 1$ matrix whose i^{th} entry is denoted by $\mathbf{x}_i$. The *support* of a vector $\mathbf{x}$ denoted by $\text{supp}(\mathbf{x})$ is the set of indices with positive value, that is, $\text{supp}(\mathbf{x}) = \{i \in [m] : \mathbf{x}_i > 0\}$. Denote by Δ_m the set of all probability vectors of dimension m. Formally,

$$\Delta_m = \{\mathbf{x} \, : \, \forall i \in [m] \; \mathbf{x}_i \ge 0, \text{ and } \sum_{i=1}^{m} \mathbf{x}_i = 1\}$$

A vector $\mathbf{x} \in \Delta_m$ is said to be *uniform* if for all $i \in [m]$, $\mathbf{x}_i > 0 \implies \mathbf{x}_i = 1/|\text{supp}(\mathbf{x})|$. A vector $\mathbf{x} \in \Delta_m$ is said to be *fully uniform* if for all $i \in [m]$, $\mathbf{x}_i = 1/m$.

A *bimatrix game* or a two player game consists of two players, the *row player* and the *column player*. The row player has a m pure strategies, denoted by the set $[m]$ and the column player has n pure strategies, denoted by $[n]$. The game is specified by two $m\times n$ *payoff matrices* A, B whose entries are reals. If the row player chooses a strategy $i \in [m]$ and the column player chooses a strategy $j \in [n]$, then they receive payoffs equal to a_{ij} and b_{ij} respectively. The players can randomize over their pure strategies, giving rise to a mixed strategy. Formally, a mixed strategy for the row player (resp. column player) is a probability vector $\mathbf{x} \in \Delta_m$ (resp. $\mathbf{y} \in \Delta_n$). Any $(\mathbf{x}, \mathbf{y}) \in \Delta_m \times \Delta_n$ is called a strategy profile. For a strategy profile $(\mathbf{x}, \mathbf{y})$, the *expected payoff* of the row player is $\mathbf{x}^T A\mathbf{y}$ and that of the column player is $\mathbf{x}^T B\mathbf{y}$.

Nash's celebrated theorem, when applied to bimatrix games, states there always exists a strategy profile so that neither player can increase her payoff by unilaterally deviating from the strategy profile. Such a strategy profile is called a *Nash Equilibrium* (NE, for short) ([31]).

Definition 2.1. (Nash Equilibrium) Let (A, B) be a bimatrix game where $A, B \in [0,1]^{m\times n}$. A strategy profile $(\mathbf{x}^*, \mathbf{y}^*) \in \Delta_m \times \Delta_n$ is a Nash equilibrium of (A, B), if for all $\mathbf{x} \in \Delta_m$ and for all $\mathbf{y} \in \Delta_n$, it holds that:

$$(\mathbf{x}^*)^T A\mathbf{y}^* \geq \mathbf{x}^T A\mathbf{y}^* \text{ and } (\mathbf{x}^*)^T B\mathbf{y}^* \geq (\mathbf{x}^*)^T B\mathbf{y}$$

Note that at the Nash equilibrium a player will give positive probability only to pure strategies that give her the maximum payoff against the strategy of the other player. Mathematically, $(\mathbf{x}^*, \mathbf{y}^*)$ is a Nash equilibrium if and only if for all $i \in [m]$ and $j \in [n]$:

$$\begin{aligned} \mathbf{x}_i^* > 0 &\implies (A\mathbf{y}^*)_i = \max_{k\in[m]} (A\mathbf{y}^*)_k \\ \mathbf{y}_j^* > 0 &\implies ((\mathbf{x}^*)^T B)_j = \max_{k\in[n]} ((\mathbf{x}^*)^T B)_k \end{aligned} \tag{1}$$

Observe that the Nash equilibria of a bimatrix game are invariant under scaling by positive constants, that is, the set of NEs of the game (A, B) is the same as the set of NEs of the game $(\alpha A, \beta B)$, for $\alpha, \beta > 0$. The NEs also remain invariant under shifting, that is, the set of NEs of the game (A, B) is the same as the set of NEs of the game $(A + \alpha, B + \beta)$, for any α, β. Thus, it is standard practice to normalize the matrices and assume that all the entries belong to $[0, 1]$.

As it is hard to compute exact Nash equilibria, a natural notion to consider is that of *approximate* equilibria. For $\epsilon > 0$, an *ϵ-approximate Nash Equilibrium* (ϵ-NE for short) is a strategy profile in which neither player has an incentive of more than ϵ of deviating unilaterally.

Definition 2.2. (ϵ-approximate Nash Equilibrium) Let (A, B) be a bimatrix game where $A, B \in [0,1]^{m\times n}$. For an arbitrary $\epsilon > 0$, a strategy profile $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \in \Delta_m \times \Delta_n$ is an ϵ-approximate Nash equilibrium if:

$$\forall \mathbf{x} \in \Delta_m : \tilde{\mathbf{x}}^T A\tilde{\mathbf{y}} \geq \mathbf{x}^T A\tilde{\mathbf{y}} - \epsilon$$
$$\forall \mathbf{y} \in \Delta_n : \tilde{\mathbf{x}}^T B\tilde{\mathbf{y}} \geq \tilde{\mathbf{x}}^T B\mathbf{y} - \epsilon$$

A stronger notion of approximation of a Nash equilibrium is the *ϵ-approximate-well-supported Nash equilibrium* (ϵ-wsNE for short), in which neither player has an incentive of more than ϵ to unilaterally deviate from any of the pure strategies used in her mixed strategy.

Definition 2.3. (ϵ-approximate well-supported Nash Equilibrium) Let (A, B) be a bimatrix game where $A, B \in [0,1]^{m\times n}$. For an arbitrary $\epsilon > 0$, a strategy profile $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in \Delta_m \times \Delta_n$ is an ϵ-well-supported Nash equilibrium if:

$$\forall i \in [m] : \bar{\mathbf{x}}_i > 0 \implies (A\bar{\mathbf{y}})_i \geq \max_{k\in[m]} (A\bar{\mathbf{y}})_k - \epsilon$$
$$\forall j \in [n] : \bar{\mathbf{y}}_j > 0 \implies (\bar{\mathbf{x}}^T B)_j \geq \max_{k\in[n]} (\bar{\mathbf{x}}^T B)_k - \epsilon$$

It is easy to see that every ϵ-wsNE is also ϵ-NE, but not vice versa. However as is observed in [9], the two approximate notions of Nash equilibrium are polynomially equivalent:

Lemma 2.4. *([9]) From every $\epsilon^2/8$-approximate Nash equilibrium of a bimatrix game, we can compute in polynomial time an ϵ-approximate-well-supported Nash equilibrium of the same game.*

Symmetric bimatrix games are a subclass of bimatrix games in which both players have the same set of pure strategies, and the payoffs depend only on the strategies chosen and not the players who play them, that is, $B = A^T$. Nash ([31]) showed that every symmetric game has a *symmetric Nash equilibrium* $(\mathbf{y}^*, \mathbf{y}^*)$.

An *imitation game* ([28]) is a bimatrix game in which the column player is an *imitator*, that is, she gets a payoff of 1 only when she picks the same strategy as the row player, otherwise her payoff is 0. Thus, the payoff matrix of the imitator is the identity matrix, that is, $B = I$.

Definition 2.5. (Imitation game, I-equilibrium) An imitiation game is a bimatrix game (A, I), where $A \in [0,1]^{n\times n}$. An *I-equilibrium* of an imitation game is a mixed strategy $\mathbf{y}$ for the imitator such that $\text{supp}(\mathbf{y}) \subseteq \text{argmax}_{k\in[n]} (A\mathbf{y})_k$.

The symmetric Nash equilibria of any symmetric game (A, A^T) are in one-to-one correspondence with the I-equilibria of the imitation game (A, I). Thus any efficient algorithm computing Nash equilibria of imitation games can be used to efficiently compute symmetric Nash equilibria of symmetric games. The following properties about Nash equilibria of imitation games are well-known (and appear in different forms in [28], [17] and [29]).

Lemma 2.6. *Let $A \in [0,1]^{n\times n}$ be a payoff matrix and let $\mathbf{y} \in \Delta_n$ be a mixed strategy. Then $(\mathbf{y},\mathbf{y})$ is a symmetric NE of (A, A^T) if and only if $\mathbf{y}$ is an I-equilibrium of (A, I).*

Proof. Observe that from equation 1, $(\mathbf{y},\mathbf{y})$ is a symmetric NE of (A, A^T) if and only if for all $i \in [n] : \mathbf{y}_i > 0 \implies (A\mathbf{y})_i = \max_{k\in[n]} (A\mathbf{y})_k$, which holds if and only if $i \in \text{supp}(\mathbf{y}) \implies i \in \text{argmax}_{k\in[n]} (A\mathbf{y})_k$, which is true if and only if $\mathbf{y}$ is an I-equilibrium of (A, I). □

Lemma 2.7. *For any Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*) \in \Delta_n \times \Delta_n$ of an imitation game (A, I) where $A \in [0,1]^{n\times n}$, $\text{supp}(\mathbf{y}^*) \subseteq \text{supp}(\mathbf{x}^*)$.*

Proof. Let $(\mathbf{x}^*, \mathbf{y}^*)$ be a Nash equilibrium of an imitation game (A, I). From equation 1, for all $i \in [n]$, $\mathbf{y}_i^* > 0 \implies ((\mathbf{x}^*)^T I)_i = \max_{k\in[n]} ((\mathbf{x}^*)^T I)_k > 0$. Thus, $i \in \text{supp}(\mathbf{y}^*) \implies i \in \text{supp}(\mathbf{x}^*)$, and hence $\text{supp}(\mathbf{y}^*) \subseteq \text{supp}(\mathbf{x}^*)$. □

Next we observe that imitation games always have a Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ where $\mathbf{x}^*$ is uniform. As we shall see in Section 3, this fact will be useful in constructing a PTAS for computing an approximate-well-supported Nash equilibrium in an imitation game.

Lemma 2.8. *For any imitation game (A, I) where $A \in [0,1]^{n\times n}$, there exists a Nash equilibrium $(\hat{\mathbf{x}}, \hat{\mathbf{y}}) \in \Delta_n \times \Delta_n$ where $\hat{\mathbf{x}}$ is uniform.*

Proof. By Nash's theorem ([31]), we know that there exists at least one Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*) \in \Delta_n \times \Delta_n$ of (A, I). From Lemma 2.7, if for some $i \in [n]$, $\mathbf{y}_i^* > 0$, then $\mathbf{x}_i^* > 0$. Together with equation 1, we have for all $i \in [n]$:

$$\mathbf{y}_i^* > 0 \implies \mathbf{x}_i^* > 0 \implies (A\mathbf{y}^*)_i = \max_{k\in[n]} (A\mathbf{y}^*)_k \tag{2}$$

Consider a mixed strategy $\hat{\mathbf{x}}$ for the row player given by $\hat{\mathbf{x}}_i = 1/|\text{supp}(\mathbf{y}^*)| \iff \mathbf{y}_i^* > 0$. Clearly $\hat{\mathbf{x}}$ is a uniform vector in Δ_n.

We also have that for all $i \in [n]$:

$$\mathbf{y}^*_i > 0 \iff \hat{\mathbf{x}}_i = \max_{k\in[n]} \hat{\mathbf{x}}_k > 0 \quad (3)$$

Set $\hat{\mathbf{y}} = \mathbf{y}^*$. Now equations 2 and 3 together with equation 1 imply that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a Nash equilibrium of (A, I) where $\hat{\mathbf{x}}$ is uniform. □

3 POLYNOMIAL-TIME ALGORITHM FOR CONSTANT APPROXIMATE NE

We now present a polynomial-time approximation scheme (PTAS) for the problem of computing a well-supported approximate Nash Equilibrium of an imitation game. Let (A, I) be an imitation game where $A \in [0,1]^{n\times n}$ is the payoff matrix of the row player and I, the $n \times n$ identity matrix is the payoff matrix of the column player. Given a constant $\epsilon \in (0,1)$, we will show how to compute an ϵ-approximate well-supported Nash Equilibrium $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ in $n^{O(1/\epsilon)}\,\mathrm{poly}(\mathcal{L})$ time, where $\mathcal{L}$ is the bit-size of the input, that is, the sum of the bit-sizes of the n^2 entries of A.

Recall that an ϵ-wsNE of an imitation game (A, I) is a mixed strategy profile $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in \Delta_n \times \Delta_n$ such that for all $i \in [n]$ and for all $j \in [n]$:

$$\begin{aligned} \bar{\mathbf{x}}_i > 0 &\implies (A\bar{\mathbf{y}})_i \geq \max_{k\in[n]} (A\bar{\mathbf{y}})_k - \epsilon \\ \bar{\mathbf{y}}_j > 0 &\implies \bar{\mathbf{x}}_j \geq \max_{k\in[n]} \bar{\mathbf{x}}_k - \epsilon \end{aligned} \quad (4)$$

We assume $\epsilon \in (0,1)$ is a constant given to us in binary. Let $\ell = \lceil \frac{1}{\epsilon} \rceil$. Since $\ell \geq \frac{1}{\epsilon}$, any $1/\ell$-wsNE is also an ϵ-wsNE. From Lemma 2.8, we know that there exists a NE $(\mathbf{x}^*, \mathbf{y}^*)$ of (A, I) where $\mathbf{x}^*$ is uniform, that is, $\mathbf{x}^*_i > 0 \implies \mathbf{x}^*_i = \frac{1}{|\mathrm{supp}(\mathbf{x}^*)|}$.

We separately analyze the problem depending on the size of the support of the row player's strategy in any Nash equilibrium. In Section 3.1 we discuss the case where there exists a Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ where $\mathbf{x}^*$ is uniform and has support of size less than ℓ. In Section 3.2, we discuss the case where in every Nash equilibrium $(\mathbf{x}^*, \mathbf{y}^*)$ with $\mathbf{x}^*$ uniform, the support of $\mathbf{x}^*$ is of size at least ℓ. Our algorithm, presented in Section 3.3 finds a $\frac{1}{\ell}$-approximate-well-supported Nash equilibrium by solving a finite set of linear programs, which are presented in the next two sections, of which one is guaranteed to be feasible. Using the solution to this feasible program we recover the desired ϵ-approximate well-supported Nash equilibrium of the imitation game (A, I).

3.1 Support less than ℓ

Let S be a subset of $[n]$ of cardinality m. Consider the following linear program $LP_1(S)$ with variables $(\Pi, \mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_n), \mathbf{y} = (\mathbf{y}_1, \ldots, \mathbf{y}_n))$:

$$\underline{\mathrm{LP}_1(S)}$$

$$\begin{aligned} &\forall i \in S : \Pi = (A\mathbf{y})_i \\ &\forall i \notin S : \Pi \geq (A\mathbf{y})_i \\ &\forall i \in S : \mathbf{x}_i = 1/m \\ &\forall i \notin S : \mathbf{x}_i = 0 \\ &\forall j \notin S : \mathbf{y}_j = 0 \\ &\sum_{j=1}^{n} \mathbf{y}_j = 1 \end{aligned}$$

Proposition 3.1. *The imitation game (A, I) has a Nash equilibrium $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ where $\boldsymbol{x}^*$ is uniform and has a support of size less than ℓ if and only if there is a set $S \subseteq [n]$ of size less than ℓ such that $LP_1(S)$ is feasible. Further any $(\boldsymbol{x}, \boldsymbol{y})$ in its feasible region is a Nash equilibrium.*

Proof. ($\Rightarrow$) Let $(\mathbf{x}^*, \mathbf{y}^*)$ be a Nash equilibrium of (A, I) where $\mathbf{x}^*$ is uniform and has a support of size $m < \ell$. Then consider the linear program $LP_1(S)$ where we set $S = \mathrm{supp}(\mathbf{x}^*)$. We claim that $(\Pi, \mathbf{x}^*, \mathbf{y}^*)$ lies in the feasible region of $LP_1(S)$, where $\Pi = \max_{k\in[n]}(A\mathbf{y}^*)_k$. This is true because:

- Since $(\mathbf{x}^*, \mathbf{y}^*)$ is a NE, by equation 1, $\mathbf{x}^*_i > 0 \implies (A\mathbf{y}^*)_i = \max_{k\in[n]}(A\mathbf{y}^*)_k$, thus for all $i \in S$, $\Pi = (A\mathbf{y}^*)_i$, and for all $i \notin S$, $\Pi \geq (A\mathbf{y}^*)_i$.
- Since $(\mathbf{x}^*, \mathbf{y}^*)$ is a NE of an imitation game, by Lemma 2.7, we have that $\mathrm{supp}(\mathbf{y}^*) \subseteq \mathrm{supp}(\mathbf{x}^*)$, equivalently $\mathbf{y}^*_j = 0$ for $j \notin S$.

($\Leftarrow$) Suppose on the other hand there is set $S \subseteq [n]$ of cardinality $m < \ell$ such that $LP_1(S)$ is feasible. Let $(\mathbf{x}, \mathbf{y})$ be any point in its feasible region. Then we have for all $i \in [n]$:

- $\mathbf{x}_i > 0 \implies i \in S \implies \Pi = (A\mathbf{y})_i = \max_{k\in[n]}(A\mathbf{y})_k$
- $\mathbf{y}_i > 0 \implies i \in S \implies \mathbf{x}_i = 1/m = \max_{k\in[n]} \mathbf{x}_k$

Thus by equation 1, $(\mathbf{x}, \mathbf{y})$ is a Nash equilibrium of (A, I) where $\mathbf{x}$ is a uniform vector with a support of size less than ℓ. □

3.2 Support at least ℓ

Suppose every NE $(\mathbf{x}^*, \mathbf{y}^*)$ of (A, I) where $\mathbf{x}^*$ is uniform has a support of size at least ℓ. For a set $S \subseteq [n]$, with $|S| = \ell$, consider the following linear program with variables $(\Pi, \mathbf{y} = (\mathbf{y}_1, \ldots, \mathbf{y}_n))$:

$$\underline{\mathrm{LP}_2(S)}$$

$$\begin{aligned} &\forall i \in S : \Pi = (A\mathbf{y})_i \\ &\forall i \notin S : \Pi \geq (A\mathbf{y})_i \\ &\forall j : \mathbf{y}_j \geq 0 \\ &\sum_{j=1}^{n} \mathbf{y}_j = 1 \end{aligned}$$

Proposition 3.2. *If every Nash equilibrium $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ of the imitation game (A, I) where $\boldsymbol{x}^*$ is uniform is such that $|\mathrm{supp}(\boldsymbol{x}^*)| \geq \ell$, then there exists a set $S \subseteq [n]$ of size exactly ℓ such that $LP_2(S)$ is feasible. Further for every $(\Pi, \bar{\mathbf{y}})$ in its feasible region, there exists a uniform $\bar{\mathbf{x}} \in \Delta_n$ such that $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is a $\frac{1}{\ell}$-approximate well-supported Nash equilibrium.*

Proof. Let $(\mathbf{x}^*, \mathbf{y}^*)$ be some Nash equilibrium of the imitiation game (A, I) where $\mathbf{x}^*$ is uniform, which we know exist thanks to Lemma 2.8. We further assume that $|\mathrm{supp}(\mathbf{x}^*)| \geq \ell$. Let S be any ℓ-element subset of $\mathrm{supp}(\mathbf{x}^*)$. Then $LP_2(S)$ is feasible because the point $(\Pi, \mathbf{y}^*)$ lies in its feasible region, where $\Pi = \max_{k\in[n]}(A\mathbf{y}^*)_k$. This is true, since we have for all $i \in [n]$ if $i \in S$, then $\mathbf{x}^*_i > 0$, which in turn implies from equation 1 that $(A\mathbf{y}^*)_i = \max_{k\in[n]}(A\mathbf{y}^*)_k = \Pi$.

Now suppose $LP_2(S)$ is feasible for some subset S of $[n]$ containing exactly ℓ elements. Let $(\Pi, \bar{\mathbf{y}})$ be a point in its feasible region. Clearly, $\Pi = \max_{k\in[n]}(A\bar{\mathbf{y}})_k$. Let $\bar{\mathbf{x}} \in \Delta_n$ be given by $\bar{\mathbf{x}}_i = \frac{1}{\ell}$ if $i \in S$, and 0 otherwise. Note that $\max_{k\in[n]} \bar{\mathbf{x}}_k = \frac{1}{\ell}$. Then $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is a $\frac{1}{\ell}$-approximate well-supported Nash equilibrium of (A, I) since it holds that:

- for all $i \in [n]$, $\bar{\mathbf{x}}_i > 0 \implies i \in S \implies \Pi = (A\bar{\mathbf{y}})_i \implies (A\bar{\mathbf{y}})_i \geq \max_{k\in[n]}(A\bar{\mathbf{x}})_k - \frac{1}{\ell}$
- for all $i \in [n]$, $\bar{\mathbf{x}}_i \geq 0$. Thus $\bar{\mathbf{y}}_i > 0 \implies \bar{\mathbf{x}}_i \geq \max_{k\in[n]} \bar{\mathbf{x}}_k - \frac{1}{\ell} = 0$ is also true.

Thus from Definition 2.3, it follows that $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ is a $\frac{1}{\ell}$-approximate well-supported Nash equilibrium, and thus also a ϵ-wsNE. □

3.3 PTAS for imitation games

Given an imitation game (A, I) and a constant $\epsilon > 0$, the following algorithm finds a ϵ-approximate-well-supported Nash equilibrium.

Algorithm 1 PTAS for Imitation games

1: Compute $\ell = \lceil \frac{1}{\epsilon} \rceil$.
2: Iterate over all subsets S of $[n]$ of size less than ℓ and check if $LP_1(S)$ is feasible. If yes, output any point in its feasible region.
3: If not, iterate over all subsets S of $[n]$ of size ℓ and check if $LP_2(S)$ is feasible. Use Proposition 3.2 to output a $\frac{1}{\ell}$-wsNE.

Theorem 3.3. *Given an imitation game (A, I), where $A \in [0, 1]^{n\times n}$, and a constant $\epsilon > 0$, Algorithm 1 computes an ϵ-approximate-well-supported Nash equilibrium of (A, I) in time $n^{O(1/\epsilon)}\text{poly}(\mathcal{L})$, where $\mathcal{L}$ is the bit size of the matrix A.*

Proof. *Correctness.* Due to Propositions 3.1 and 3.2, at least one of the linear programs examined in Steps 2 or 3 of Algorithm 1 will be feasible. If the algorithm succeeds in Step 2, then it outputs an exact NE of the imitation game due to Proposition 3.1; and if not, it outputs a ϵ-wsNE due to Proposition 3.2 in Step 3.

Complexity. In step 2, Algorithm 1 iterates over all subsets of $[n]$ of size less than ℓ, which are $\binom{n}{1} + \binom{n}{2} + \cdots + \binom{n}{\ell-1} \leq (n+1)^\ell$ in number. Checking if an LP is feasible takes polynomial time in $\mathcal{L}$, the bit size of the input A. Thus step 2 of the algorithm takes time at most:

$$\sum_{i=1}^{\ell} \binom{n}{i}\text{poly}(\mathcal{L}) \leq (n+1)^\ell \text{poly}(\mathcal{L}) = n^{O(1/\epsilon)}\text{poly}(\mathcal{L})$$

In step 3, Algorithm 1 iterates over all subsets S of $[n]$ of size ℓ and checks if the corresponding linear program $\text{LP}_2(S)$ is feasible. This takes time at most:

$$\binom{n}{\ell}\text{poly}(\mathcal{L}) \leq n^\ell\text{poly}(\mathcal{L}) = n^{O(1/\epsilon)}\text{poly}(\mathcal{L})$$

Thus, Algorithm 1 runs in time $n^{O(1/\epsilon)}\text{poly}(\mathcal{L})$, and computes an ϵ-approximate-well-supported Nash equilibrium of the imitation game (A, I). □

Having presented a *polyomial time approximation scheme* (PTAS), we now ask if there is a *fully polynomial time approximation scheme* (FPTAS) for the problem of computing an approximate Nash equilibrium of an imitation game. The results of the next section show that an FPTAS is unlikely.

4 HARDNESS OF $1/n^{\Theta(1)}$ -APPROXIMATION

It was shown in [7] that the problem of computing an ϵ-approximate-well-supported Nash equilibrium of a bimatrix game is PPAD-hard for $\epsilon = \frac{1}{n^c}$, for any $c > 0$. In this section we show that a similar hardness result holds for imitation games as well. We do this by first showing that it remains hard to compute a $\frac{1}{n^c}$-approximate-well-supported symmetric Nash equilibrium of symmetric games, for any $c > 0$. Then we show that any polynomial-time algorithm that computes a $\frac{1}{n^c}$-approximate-well-supported Nash equilibrium of an imitation game (A, I) can be used to compute a $\frac{1}{n^c}$-approximate-well-supported Nash equilibrium of a symmetric game (A, A^T) in polynomial time, for any $c \geq 1$, showing PPAD-hardness. We then extend the result to show that computing an $\frac{1}{n^{1/c}}$-wsNE of imitation games is PPAD-hard as well, for integers $c \geq 1$. Therefore this rules out an FPTAS for computing approximate Nash equilibria of imitation games, unless PPAD ⊂ P.

Lemma 4.1. *For any $c > 0$, the problem of computing a symmetric $\frac{1}{n^c}$-approximate-well-supported Nash equilibrium of a symmetric game is* PPAD*-hard.*

Proof. Let (A, B) be any bimatrix game where $A, B \in [0, 1]^{n\times n}$. Consider the symmetric game (C, C^T), where C is the following $2n \times 2n$ matrix, where $m = 6$.

$$C = \begin{bmatrix} O & A+m \\ B^T+m & O \end{bmatrix}$$

where O is a zero matrix of appropriate dimensions, and $m > 0$. Let $(\bar{\mathbf{z}}, \bar{\mathbf{z}})$ be a symmetric ϵ-approximate-well-supported Nash equilibrium of (C, C^T), where $0 < \epsilon < 1$. Let $\mathbf{x}, \mathbf{y}$ be such that for all $i \in [n]$: $\mathbf{x}_i = \bar{\mathbf{z}}_i$ and $\mathbf{y}_i = \bar{\mathbf{z}}_{n+i}$. Let $X = \sum_{i\in[n]} \mathbf{x}_i$ and $Y = \sum_{j\in[n]} \mathbf{y}_j$. Since $\bar{\mathbf{z}} \in \Delta_{2n}$, $X + Y = 1$. Assume without loss of generality that $X \geq 1/2$. We have from Definition 2.3 that for all $i \in [n]$:

$$\bar{\mathbf{z}}_i > 0 \implies (C\bar{\mathbf{z}})_i \geq \max_{k\in[2n]}(C\bar{\mathbf{z}})_k - \epsilon \tag{5}$$

We have for all $i \in [n]$, $(C\bar{\mathbf{z}})_i = (A\mathbf{y})_i + mY$ and $(C\bar{\mathbf{z}})_{n+i} = (B^T\mathbf{x})_i + mX$. Since $X \geq 1/2$, there exists $i \in [n]$ such that $\mathbf{x}_i > 0$. Then we have that for any $j \in [n]$:

$$(A\mathbf{y})_i + mY \geq (B^T\mathbf{x})_j + mX - \epsilon$$

This gives:

$$Y \geq X - \frac{\epsilon}{m} + \frac{(B^T\mathbf{x})_j - (A\mathbf{y})_i}{m}$$

Since entries of A, B are from $[0, 1]$, $(B^T\mathbf{x})_j - (A\mathbf{y})_i \geq -1$. Thus for $m = 6$,

$$Y \geq \frac{1}{2} - \frac{\epsilon}{m} - \frac{1}{m} \geq \frac{m-2}{2m} = \frac{1}{3}$$

Now consider $\bar{\mathbf{x}}, \bar{\mathbf{y}} \in \Delta_n$ such that for all $i \in [n]$, $\bar{\mathbf{x}}_i = \frac{\mathbf{x}_i}{X}$ and $\bar{\mathbf{y}}_i = \frac{\mathbf{y}_i}{Y}$. Since $(\bar{\mathbf{z}}, \bar{\mathbf{z}})$ is an ϵ-wsNE of (C, C^T), it follows from equation 5 that for all $i \in [n]$:

$$\mathbf{x}_i > 0 \implies (A\mathbf{y})_i + mY \geq \max_{k\in[n]}(A\mathbf{y})_k + mY - \epsilon \text{ , thus}$$

$$\bar{\mathbf{x}}_i > 0 \implies (A\bar{\mathbf{y}})_i \geq \max_{k\in[n]}(A\bar{\mathbf{y}})_k - \frac{\epsilon}{Y} \geq \max_{k\in[n]}(A\bar{\mathbf{y}})_k - 3\epsilon$$

Similarly from equation 5 we have for all $i \in [n]$:

$$\mathbf{y}_i > 0 \implies (B^T\mathbf{x})_i + mX \geq \max_{k\in[n]}(B^T\mathbf{x})_k + mX - \epsilon \text{ , thus}$$

$$\bar{\mathbf{y}}_i > 0 \implies (B^T\bar{\mathbf{x}})_i \geq \max_{k\in[n]}(B^T\bar{\mathbf{x}})_k - \frac{\epsilon}{X} \geq \max_{k\in[n]}(B^T\bar{\mathbf{x}})_k - 2\epsilon$$

Thus from Definition 2.3, $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in \Delta_n \times \Delta_n$ is a 3ϵ-wsNE of (A, B). The entries of the matrix C are in $[0, 7]$. By noting that scaling the entries of the payoff matrices by the same constant causes the approximation factor ϵ to change only by a constant multiplicatively, we can observe that a symmetric ϵ-wsNE of (C, C^T) is also a symmetric $\frac{\epsilon}{7}$-wsNE of (D, D^T), where $D = \frac{1}{7}C$ is a matrix with entries in $[0, 1]$. Thus in fact from any symmetric ϵ-wsNE of the symmetric game (D, D^T), we can construct a 21ϵ-wsNE of the general bimatrix game (A, B). Since we know from [7] that for any $c > 0$, computing an $\frac{1}{n^c}$-aproximate Nash equilibrium of a general bimatrix game is PPAD-hard, we conclude because of the above reduction that the problem of computing a symmetric $\frac{1}{n^c}$-wsNE of a symmetric game is PPAD-hard as well. □

We now show our first hardness result for imitation games:

Theorem 4.2. *For $c \geq 1$, the problem of computing an $1/n^c$-approximate-well-supported Nash equilibrium of an imitation game (A, I) is* PPAD*-hard.*

Proof. Let (A, I) be an imitation game where A is an $n \times n$ matrix with entries from $[0, 1]$. Fix $c \geq 1$. We first observe that for every strategy profile $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ that is a $1/n^c$-approximate well-supported NE of (A, I), the strategy profile $(\bar{\mathbf{y}}, \bar{\mathbf{y}})$ is a $1/n^c$-approximate-well-supported NE of (A, A^T).

Let $\epsilon = 1/n^c$. By definition of ϵ-approximate well-supported NE, we have for all $i \in [n]$:

$$\bar{\mathbf{x}}_i > 0 \implies (A\bar{\mathbf{y}})_i \geq \max_k (A\bar{\mathbf{y}})_k - \epsilon \tag{6}$$

$$\bar{\mathbf{y}}_i > 0 \implies \bar{\mathbf{x}}_i \geq \max_k \bar{\mathbf{x}}_k - \epsilon \tag{7}$$

Since $\bar{\mathbf{x}} \in \Delta_n$, $\max_k \bar{\mathbf{x}}_k \geq 1/n$. If $\max_k \bar{\mathbf{x}}_k = 1/n$, then in fact for each $i \in [n]$, $\bar{\mathbf{x}}_i = 1/n > 0$. On the other hand suppose $\max_k \bar{\mathbf{x}}_k > 1/n$. Since $\epsilon \leq 1/n$, from equation 7 we have that if $\bar{\mathbf{y}}_i > 0$ then $\bar{\mathbf{x}}_i \geq \max_k \bar{\mathbf{x}}_k - \epsilon > 0$. Thus, in either case whenever $\bar{\mathbf{y}}_i > 0$, it holds that $\bar{\mathbf{x}}_i > 0$. Thus from equations 6 and 7 we have for all $i \in [n]$:

$$\bar{\mathbf{y}}_i > 0 \implies \bar{\mathbf{x}}_i > 0 \implies (A\bar{\mathbf{y}})_i \geq \max_k (A\bar{\mathbf{y}})_k - \epsilon$$

Thus, $(\bar{\mathbf{y}}, \bar{\mathbf{y}})$ is a symmetric $1/n^c$-approximate-well-supported symmetric NE of (A, A^T). Therefore, the problem of computing a symmetric $1/n^c$-approximate-well-supported Nash equilibrium of the symmetric game (A, A^T) reduces to the problem of computing a $1/n^c$-approximate well-supported Nash equilibrium of an imitation game (A, I). Since we know from Lemma 4.1 that the former is PPAD-hard, the theorem follows. □

We now show that the hardness extends to the problem of computing a $\frac{1}{n^{1/c}}$-wsNE of an imitation game, for $c \geq 1$.

Theorem 4.3. *For $c \geq 1$, the problem of computing a $\frac{1}{n^{1/c}}$-approximate-well-supported Nash equilibrium of an imitation game (A, I) is PPAD-hard.*

Proof. Let (A, I) be an imitation game, where $A \in [0, 1]^n$. Fix an integer $c \geq 1$. We construct an $m \times m$ matrix A', where $m = (2n)^c$, given by:

$$A' = \begin{bmatrix} \frac{1}{2}A + \frac{1}{2} & H \\ O & O \end{bmatrix}$$

where H is an $(m - n) \times (m - n)$ matrix with every entry $\frac{1}{2}$, and O denotes zero matrices of appropriate size. Since every entry of A is in $[0, 1]$, every non-zero entry of A' is at least $\frac{1}{2}$ and at most 1. Let $(\mathbf{x}', \mathbf{y}')$ be an ϵ'-wsNE of the imitation game (A', I), where $\epsilon' = \frac{1}{m^{1/c}}$. Thus for any $i \in [m]$:

$$\mathbf{x}'_i > 0 \implies (A'\mathbf{y}')_i \geq \max_{k\in[m]}(A'\mathbf{y}')_k - \epsilon' \tag{8}$$

and for any $j \in [m]$:

$$\mathbf{y}'_j > 0 \implies \mathbf{x}'_j \geq \max_{k\in[m]} \mathbf{x}'_k - \epsilon' \tag{9}$$

Note that for any $i \in [n]$, $(A'\mathbf{y}')_i \geq \frac{1}{2}$, and for any $i \notin [n]$, $(A'\mathbf{y}')_i = 0$. Thus by the contrapositive of Equation 8, we get that for all $i \notin [n]$, $\mathbf{x}'_i = 0$. Thus $\mathrm{supp}(\mathbf{x}') \subseteq [n]$. Similarly note that since for all $j \notin [n]$, $\mathbf{x}'_j = 0$, it follows from the contrapositive of Equation 9 that $\mathbf{y}'_j = 0$. Thus $\mathrm{supp}(\mathbf{y}') \subseteq [n]$.

Now we define vectors $\mathbf{x} \in \Delta_n$ and $\mathbf{y} \in \Delta_n$ given by $\mathbf{x}_i = \mathbf{x}'_i$ and $\mathbf{y}_i = \mathbf{y}'_i$, for all $i \in [n]$. Observe that for $i \in [n]$:

$$(A'\mathbf{y}')_i = \sum_{j=1}^{m} a'_{ij}\mathbf{y}'_j = \sum_{j=1}^{n}\left(\frac{1}{2}a_{ij} + \frac{1}{2}\right)\mathbf{y}_j = \frac{1}{2}(A\mathbf{y})_i + \frac{1}{2}$$

With $\epsilon' = \frac{1}{m^{1/c}} = \frac{1}{2n}$, we have from Equation 8 for all $i \in [n]$:

$$\mathbf{x}_i > 0 \implies \frac{1}{2}(A\mathbf{y})_i + \frac{1}{2} \geq \max_{k\in[n]} \frac{1}{2}(A\mathbf{y})_k + \frac{1}{2} - \frac{1}{2n}$$

Equivalently, for all $i \in [n]$:

$$\mathbf{x}_i > 0 \implies (A\mathbf{y})_i \geq \max_{k\in[n]}(A\mathbf{y})_k - \frac{1}{n}$$

Similarly, from Equation 9 we have for all $j \in [n]$:

$$\mathbf{y}_j > 0 \implies \mathbf{x}_j \geq \max_{k\in[n]} \mathbf{x}_k - \frac{1}{2n} \geq \max_{k\in[n]} \mathbf{x}_k - \frac{1}{n}$$

This in fact shows that $(\mathbf{x}, \mathbf{y})$ is an $\frac{1}{n}$-wsNE of the imitation game (A, I). Thus any algorithm that computes an $\frac{1}{m^{1/c}}$-wsNE of (A', I), where $A' \in [0, 1]^{m\times m}$, can be used to compute an $\frac{1}{n}$-wsNE of (A, I), where $A \in [0, 1]^{n\times n}$. Since the latter problem is PPAD-hard due to Theorem 4.2, the former problem must also be PPAD-hard. □

We summarize Theorems 4.2 and 4.3:

Theorem 4.4. *For any $c > 0$, the problem of computing a $\frac{1}{n^c}$-approximate-well-supported Nash equilibrium of an imitation game (A, I) is PPAD-hard.*

Recall from Lemma 2.4 that the two notions of approximate Nash equilibria are polyomially equivalent. Thus we have:

Corollary 4.5. *For any $c > 0$, the problem of computing a $1/n^c$-approximate Nash equilibrium of a imitation game is* PPAD*-hard.*

This implies that a *fully polynomial time approximation scheme*, that is, an algorithm which runs in time polynomial in n and $1/\epsilon$, for the problem of computing an ϵ-approximate-well-supported Nash equilibrium of an imitation game is unlikely, unless PPAD $\subset$ P.

This hardness result also rules out the smoothed complexity of computing an approximate NE in imitation games being in P, as was shown in [7] for general bimatrix games:

Corollary 4.6. *It is unlikely that the problem of computing a Nash equilibrium of an imitation game is in smoothed polynomial time, under uniform perturbations, unless* PPAD $\subset$ RP.

Since an FPTAS is unlikely and so is obtaining smoothed complexity in P, we can ask if the average case is any easier. Indeed, a result of [32] applied to *random* imitation games, where the payoffs (in $[0, 1]$) of the row player are chosen independently and randomly from the same distribution, shows that with high probability, the fully-uniform strategy profile is an $\tilde{O}(\frac{1}{\sqrt{n}})$-approximate Nash equilibrium. Note that no assumptions are made on the probability distribution itself.

Theorem 4.7 ([32]). *Consider an imitation game (A, I) where $A \in [0, 1]^{n \times n}$, in which the entries of A are chosen independently at random from the same distribution. Then with probability at least $1 - \frac{1}{n}$, the fully uniform strategy profile is an ϵ-approximate Nash equilibrium, where $\epsilon = O\left(\sqrt{\frac{\ln n}{n}}\right)$.*

5 DISCUSSION

In this paper we studied the complexity of finding approximate Nash equilibria in imitation games. In general two-player games, the problem of computing an ϵ-approximate NE, for constant $\epsilon > 0$, is known to admit a quasi-polynomial-time algorithm, which is in fact optimal assuming the exponential-time-hypothesis for PPAD [34]. In contrast, we showed that for imitation games this problem can be solved in polynomial time due to our polynomial-time approximation scheme (PTAS) presented in Section 3.

On the other hand we showed that when $\frac{1}{\text{poly}(n)}$-approximate NE are considered, the problem remains PPAD-hard just like the case of general two-player games. We in fact showed that computing a $\frac{1}{n^c}$-approximate NE is PPAD-hard, for any $c > 0$. In showing this result we also showed PPAD-hardness of finding a $\frac{1}{n^c}$-approximate-well-supported NE in both symmetric and imitation games, for any $c > 0$. While the above results rule out smoothed complexity of the problem being in P (unless PPAD $\subset$ RP), in the average case, quite like general games, the fully uniform strategy is with high probability an $\tilde{O}(1/\sqrt{n})$-approximate NE of an imitation game.

REFERENCES

[1] Bharat Adsul, Jugal Garg, Ruta Mehta, and Milind Sohoni. 2011. Rank-1 Bimatrix Games: A Homeomorphism and a Polynomial Time Algorithm. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing (STOC '11)*. ACM, New York, NY, USA, 195–204. https://doi.org/10.1145/1993636.1993664

[2] Noga Alon, Troy Lee, Adi Shraibman, and Santosh Vempala. 2013. The Approximate Rank of a Matrix and Its Algorithmic Applications: Approximate Rank. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing (STOC '13)*. ACM, New York, NY, USA, 675–684. https://doi.org/10.1145/2488608.2488694

[3] Imre Bárány, Santosh Vempala, and Adrian Vetta. 2007. Nash Equilibria in Random Games. *Random Struct. Algorithms* 31, 4 (Dec. 2007), 391–405. https://doi.org/10.1002/rsa.v31:4

[4] Siddharth Barman. 2015. Approximating Nash Equilibria and Dense Bipartite Subgraphs via an Approximate Version of Caratheodory's Theorem. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing (STOC '15)*. ACM, New York, NY, USA, 361–369. https://doi.org/10.1145/2746539.2746566

[5] Vittorio Bilò and Marios Mavronicolas. 2019. The Complexity of Computational Problems about Nash Equilibria in Symmetric Win-Lose Games. *CoRR* abs/1907.10468 (2019). arXiv:1907.10468 http://arxiv.org/abs/1907.10468

[6] Hartwig Bosse, Jaroslaw Byrka, and Evangelos Markakis. 2010. New algorithms for approximate Nash equilibria in bimatrix games. *Theoretical Computer Science* 411, 1 (2010), 164 – 173. https://doi.org/10.1016/j.tcs.2009.09.023

[7] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2006. Computing Nash Equilibria: Approximation and Smoothed Complexity. *CoRR* abs/cs/0602043 (2006). arXiv:cs/0602043 http://arxiv.org/abs/cs/0602043

[8] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2006. Sparse Games Are Hard. In *WINE*.

[9] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2009. Settling the Complexity of Computing Two-Player Nash Equilibria. *J. ACM* 56, 3, Article Article 14 (May 2009), 57 pages. https://doi.org/10.1145/1516512.1516516

[10] Xi Chen and Shang hua Teng. 2007. The approximation complexity of win-lose games. In *In Proceedings of SODA*.

[11] Bruno Codenotti and Daniel Štefankovič. 2005. On the computational complexity of Nash equilibria for (0,1) bimatrix games. *Inform. Process. Lett.* 94, 3 (2005), 145 – 150. https://doi.org/10.1016/j.ipl.2005.01.010

[12] Artur Czumaj, Michail Fasoulakis, and Marcin Jurdziński. 2014. Approximate Well-Supported Nash Equilibria in Symmetric Bimatrix Games. In *Algorithmic Game Theory*, Ron Lavi (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 244–254.

[13] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. 2006. The Complexity of Computing a Nash Equilibrium. In *Proceedings of the Thirty-eighth Annual ACM Symposium on Theory of Computing (STOC '06)*. ACM, New York, NY, USA, 71–78. https://doi.org/10.1145/1132516.1132527

[14] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. 2006. A Note on Approximate Nash Equilibria. In *Internet and Network Economics*, Paul Spirakis, Marios Mavronicolas, and Spyros Kontogiannis (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 297–306.

[15] Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. 2007. Progress in Approximate Nash Equilibria. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC '07)*. ACM, New York, NY, USA, 355–358. https://doi.org/10.1145/1250910.1250962

[16] Konstantinos Daskalakis and Christos Papadimitriou. 2005. Three-Player Games Are Hard. *Electronic Colloquium on Computational Complexity (ECCC)* (01 2005).

[17] Vladimir Estivill-Castro and Mahdi Parsa. 2009. Computing Nash Equilibria Gets Harder: New Results Show Hardness Even for Parameterized Complexity. In *CATS*.

[18] John Fearnley, Paul W. Goldberg, Rahul Savani, and Troels Bjerre Sørensen. 2012. Approximate Well-Supported Nash Equilibria Below Two-Thirds. In *Algorithmic Game Theory*, Maria Serna (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 108–119.

[19] Jugal Garg, Albert Xin Jiang, and Ruta Mehta. 2011. Bilinear Games: Polynomial Time Algorithms for Rank Based Subclasses. In *Internet and Network Economics*, Ning Chen, Edith Elkind, and Elias Koutsoupias (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 399–407.

[20] Ravi Kannan and Thorsten Theobald. 2007. Games of Fixed Rank: A Hierarchy of Bimatrix Games. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '07)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1124–1132. http://dl.acm.org/citation.cfm?id=1283383.1283504

[21] Spyros Kontogiannis and Paul Spirakis. 2011. Approximability of Symmetric Bimatrix Games and Related Experiments. In *Experimental Algorithms*, Panos M. Pardalos and Steffen Rebennack (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–20.

[22] Spyros C. Kontogiannis, Panagiota N. Panagopoulou, and Paul G. Spirakis. 2006. Polynomial Algorithms for Approximating Nash Equilibria of Bimatrix Games. In *Internet and Network Economics*, Paul Spirakis, Marios Mavronicolas, and Spyros Kontogiannis (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 286–296.

[23] Spyros C. Kontogiannis and Paul G. Spirakis. 2010. Well Supported Approximate Equilibria in Bimatrix Games. *Algorithmica* 57, 4 (01 Aug 2010), 653–667. https://doi.org/10.1007/s00453-008-9227-6

[24] C. E. Lemke and J. T. Howson, Jr. 1964. Equilibrium Points of Bimatrix Games. *J. Soc. Indust. Appl. Math.* 12, 2 (1964), 413–423. https://doi.org/10.1137/0112033 arXiv:https://doi.org/10.1137/0112033

[25] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. 2003. Playing Large Games Using Simple Strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC '03)*. ACM, New York, NY, USA, 36–41.

[26] Zhengyang Liu and Ying Sheng. 2018. On the Approximation of Nash Equilibria in Sparse Win-Lose Games. (2018). https://aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16311

[27] Andrew McLennan and Rabee Tourky. 2005. From Imitation Games to Kakutani. (01 2005).

[28] Andrew McLennan and Rabee Tourky. 2010. Imitation games and computation. *Games and Economic Behavior* 70, 1 (2010), 4 – 11. https://doi.org/10.1016/j.geb.2009.08.003 Special Issue In Honor of Ehud Kalai.

[29] Andrew McLennan and Rabee Tourky. 2010. Simple complexity from imitation games. *Games and Economic Behavior* 68, 2 (2010), 683 – 688. https://doi.org/10.1016/j.geb.2009.10.003

[30] Ruta Mehta. 2014. Constant Rank Bimatrix Games Are PPAD-hard. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing (STOC '14)*. ACM, New York, NY, USA, 545–554. https://doi.org/10.1145/2591796.2591835

[31] John Nash. 1951. Non-cooperative games. *Annals of mathematics* (1951), 286–295.

[32] Panagiota N. Panagopoulou and Paul G. Spirakis. 2011. Random Bimatrix Games Are Asymptotically Easy to Solve (A Simple Proof). In *Algorithmic Game Theory*, Giuseppe Persiano (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 190–199.

[33] Christos H. Papadimitriou. 1994. On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence. *J. Comput. Syst. Sci.* 48, 3 (June 1994), 498–532. https://doi.org/10.1016/S0022-0000(05)80063-7

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

[35] Rahul Savani and Bernhard von Stengel. 2006. Hard-to-Solve Bimatrix Games. *Econometrica* 74, 2 (2006), 397–429. http://www.jstor.org/stable/3598806

[36] Daniel A. Spielman and Shang-Hua Teng. 2004. Smoothed Analysis of Algorithms: Why the Simplex Algorithm Usually Takes Polynomial Time. *J. ACM* 51, 3 (May 2004), 385–463. https://doi.org/10.1145/990308.990310

[37] Haralampos Tsaknakis and Paul G. Spirakis. 2007. An Optimization Approach for Approximate Nash Equilibria. In *Internet and Network Economics*, Xiaotie Deng and Fan Chung Graham (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 42–56.

[38] J. v. Neumann. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100, 1 (dec 1928), 295–320. https://doi.org/10.1007/bf01448847