Inapproximability Results for Approximate Nash Equilibria*

Argyrios Deligkas, John Fearnley, and Rahul Savani

Department of Computer Science, University of Liverpool, UK

Abstract. We study the problem of finding approximate Nash equilibria that satisfy certain conditions, such as providing good social welfare. In particular, we study the problem ϵ-NE δ-SW: find an ϵ-approximate Nash equilibrium (ϵ-NE) that is within δ of the best social welfare achievable by an ϵ-NE. Our main result is that, if the randomized exponential-time hypothesis (RETH) is true, then solving $\left(\frac{1}{8} - \mathrm{O}(\delta)\right)$-NE $\mathrm{O}(\delta)$-SW for an $n \times n$ bimatrix game requires $n^{\widetilde{\Omega}(\delta^{\Lambda} \log n)}$ time, where Λ is a constant. Building on this result, we show similar conditional running time lower bounds on a number of decision problems for approximate Nash equilibria that do not involve social welfare, including maximizing or minimizing a certain player's payoff, or finding approximate equilibria contained in a given pair of supports. We show quasi-polynomial lower bounds for these problems assuming that RETH holds, and these lower bounds apply to ϵ-Nash equilibria for all $\epsilon < \frac{1}{8}$. The hardness of these other decision problems has so far only been studied in the context of exact equilibria.

1 Introduction

One of the most fundamental problems in game theory is to find a Nash equilibrium of a game. Often, we are not interested in finding any Nash equilibrium, but instead we want to find one that also satisfies certain constraints. For example, we may want to find a Nash equilibrium that provides high *social welfare*, which is the sum of the player's payoffs.

In this paper we study such problems for *bimatrix games*, which are two-player strategic-form games. Unfortunately, for bimatrix games, it is known that these problems are hard. Finding any Nash equilibrium of a bimatrix game is PPAD-complete [10], while finding a constrained Nash equilibrium turns out to be even harder. Gilboa and Zemel [16] studied several decision problems related to Nash equilibria. They proved that it is NP-complete to decide whether there exist Nash equilibria in bimatrix games with some "desirable" properties, such as high social welfare. Conitzer and Sandholm [7] extended the list of NP-complete problems of [16] and furthermore proved inapproximability results for some of them. Recently, Garg et al. [15] and Bilo and Mavronicolas [4] extended these results to many player games and provided ETR-completeness results for them.

* The authors were supported by EPSRC grant EP/L011018/1. The full version of this paper, with complete proofs, is available at http://arxiv.org/abs/1608.03574.

Approximate equilibria. Due to the apparent hardness of finding exact Nash equilibria, focus has shifted to *approximate* equilibria. There are two natural notions of approximate equilibrium, both of which will be studied in this paper. An *ϵ-approximate Nash equilibrium* (ϵ-NE) requires that each player has an expected payoff that is within ϵ of their best response payoff. An *ϵ-well-supported Nash equilibrium* (ϵ-WSNE) requires that both players only play strategies whose payoff is within ϵ of the best response payoff. Every ϵ-WSNE is an ϵ-NE but the converse does not hold, so a WSNE is a more restrictive notion.

There has been a long line of work on finding approximate equilibria [5, 8, 11–13, 18, 22]. Since we use an additive notion of approximation, it is common to rescale the game so that the payoffs lie in $[0, 1]$, which allows different algorithms to be compared. The state of the art for polynomial-time algorithms is the following. There is a polynomial-time algorithm that computes an 0.3393-NE [22], and a polynomial-time algorithm that computes a 0.6528-WSNE [8].

There is also a *quasi-polynomial time approximation scheme* (QPTAS) for finding approximate Nash equilibria. The algorithm of Lipton, Markakis, and Mehta finds an ϵ-NE in $n^{O(\frac{\log n}{\epsilon^2})}$ time [19]. They proved that there is always an ϵ-NE with logarithmic support, and then uses a brute-force search over all possible candidates to find one. We will refer to their algorithm as the LMM algorithm.

A recent breakthrough of Rubinstein implies that we cannot do better than a QPTAS like the LMM algorithm [21]: assuming the ETH for PPAD (PETH), there is a small constant, ϵ^*, such that for $\epsilon < \epsilon^*$, every algorithm for finding an ϵ-NE requires quasi-polynomial time. Briefly, PETH is the conjecture that EndOfTheLine, the canonical PPAD-complete problem, cannot be solved faster than exponential time.

Constrained approximate Nash equilibria. While deciding whether a game has an exact Nash equilibrium that satisfies certain constraints is NP-hard for most interesting constraints, this is not the case for approximate equilibria, because the LMM algorithm can be adapted to provide a QPTAS for them. The question then arises whether these results are tight.

Let the problem ϵ-NE δ-SW be the problem of finding an ϵ-NE whose social welfare is within δ of the best social welfare that can be achieved by an ϵ-NE. Hazan and Krauthgamer [17] and Austrin, Braverman and Chlamtac [3] proved that there is a small but constant ϵ such that ϵ-NE ϵ-SW is at least as hard as finding a hidden clique of size $\mathrm{O}(\log n)$ in the random graph $G_{n,1/2}$. This was further strengthened by Braverman, Ko, and Weinstein [6] who showed a lower bound based on the *exponential-time hypothesis* (ETH), which is the conjecture that any deterministic algorithm for 3SAT requires $2^{\Omega(n)}$ time. More precisely, they showed that under the ETH there is a small constant ϵ such that any algorithm for $\mathrm{O}(\epsilon)$-NE $\mathrm{O}(\epsilon)$-SW[1] requires $n^{\mathrm{poly}(\epsilon)\log(n)^{1-o(1)}}$ time[2]. We shall refer to this as the BKW result.

It is worth noting that the Rubinstein's hardness result [21] almost makes this result redundant. If one is willing to accept that PETH is true, which is a stronger conjecture than ETH, then Rubinstein's result says that for small ϵ we require quasi-polynomial time to find *any* ϵ-NE, which obviously implies that the same lower bound applies to ϵ-NE δ-SW for any δ.

Our results. Our first result is a lower bound for the problem of finding ϵ-NE δ-SW. The *randomized ETH* (RETH) is the conjecture that any randomized algorithm for 3SAT requires $2^{\Omega(n)}$ time. We show that, assuming RETH, there exists a small constant Δ such that for all $\delta \in [1/n, \Delta]$ the problem $\left(\frac{1-4g\cdot\delta}{8}\right)$-NE $\left(\frac{g\cdot\delta}{4}\right)$-SW requires $n^{\widetilde{\Omega}(\delta^{\Lambda}\log n)}$ time[3], where $g = \frac{1}{138}$, and Λ is a constant.

To understand this result, let us compare it to the BKW result. First, observe that as δ gets smaller, the ϵ in our ϵ-NE gets larger, whereas the approximate Nash equilibria in the BKW result get smaller. Asymptotically, our ϵ approaches 1/8. Moreover, since $\delta \leq 1$, our lower bound applies to all ϵ-NE with $\epsilon \leq \frac{1-4g}{8} \approx 0.1214$. This is orders of magnitude larger than the inapproximability bound given by Rubinstein's hardness result, and so is not made redundant by that result. In short, our hardness result is about the hardness of obtaining good social welfare, rather than the hardness of simply finding an approximate equilibrium.

Secondly, when compared to the BKW result, we obtain a slightly better lower bound. The exponent in their lower bound is logarithmic only in the limit, while ours is always logarithmic. In particular, we obtain quasi-polynomial lower bounds whenever δ is constant.

Finally, our result uses a stronger conjecture when compared to the BKW result. While they assume ETH, our result requires that we assume RETH. This is a stronger conjecture, since even if ETH is true, there may exist randomized sub-exponential algorithms for 3SAT. This means that our result is ultimately incomparable to the BKW result: we obtain a lower bound for larger ϵ, and we have a better lower bound on the running time, but we do so by assuming a stronger conjecture.

To prove our result, we reduce from the problem of approximating the value of a *free game*. Aaronson, Impagliazzo, and Moshkovitz showed quasi-polynomial lower bounds for this problem assuming ETH [1]. In fact, they give two different lower bounds: the *high error* result shows a quasi-polynomial lower bound for determining whether the value of the game is 1 or $1-\delta$ for small δ, while the *low error* result shows a weaker almost-quasi-polynomial lower bound on

[1] While the proof in [6] produces a lower bound for 0.8-NE $(1-\mathrm{O}(\epsilon))$-SW, this is in a game with maximum payoff $\mathrm{O}(1/\epsilon)$. Therefore, when the payoffs in this game are rescaled to $[0,1]$, the resulting lower bound only applies to ϵ-NE ϵ-SW.

[2] Although the paper claims that they obtain a $n^{\widetilde{\mathrm{O}}(\log n)}$ lower bound, the proof reduces from the *low error* result from [1] (cf. Theorem 36 in [2]), which gives only the weaker lower bound of $n^{\mathrm{poly}(\epsilon)\log(n)^{1-o(1)}}$.

[3] Here $\widetilde{\Omega}(\log n)$ means $\Omega(\frac{\log n}{(\log\log n)^c})$ for some constant c.

determining whether the value of the game is 1 or δ for small δ. The BKW result was proved via a reduction from the low error case, while our result uses the high error case. We reduce the free game to a bimatrix game, and prove that in any ϵ-NE of the game, the players must simulate the free game well enough so that we can determine whether the value of the free game is 1 or $1 - \delta$. Our reduction is substantially different from the BKW reduction: we use a sub-sampling result for free games to reduce the number of questions in the free game, and then we use a different method to force the players to simulate the free game.

Once we have our lower bound on the problem of finding ϵ-NE δ-SW, we use it to prove lower bounds for other problems regarding constrained approximate NEs and WSNEs. Table 1 gives a list of the problems that we consider. For each one, we provide a reduction from ϵ-NE δ-SW to that problem. Ultimately, we prove that if RETH is true, then for every $\epsilon < \frac{1}{8}$ finding an ϵ-NE with the given property requires $n^{\widetilde{\Omega}(\log n)}$ time.

Problem description	**Problem definition**
Problem 1: Large payoffs $u \in (0, 1]$	Is there an ϵ-NE $(\mathbf{x}, \mathbf{y})$ such that $\min(\mathbf{x}^T R\mathbf{y}, \mathbf{x}^T C\mathbf{y}) \geq u$?
Problem 2: Small total payoff $v \in [0, 2)$	Is there an ϵ-NE $(\mathbf{x}, \mathbf{y})$ such that $\mathbf{x}^T R\mathbf{y} + \mathbf{x}^T C\mathbf{y} \leq v$?
Problem 3: Small payoff $u \in [0, 1)$	Is there an ϵ-NE $(\mathbf{x}, \mathbf{y})$ such that $\mathbf{x}^T R\mathbf{y} \leq u$?
Problem 4: Restricted support $S \subset [n]$	Is there an ϵ-NE $(\mathbf{x}, \mathbf{y})$ with $\text{supp}(\mathbf{x}) \subseteq S$?
Problem 5: Two ϵ-NE $d \in (0, 1]$ apart in Total Variation (TV) distance	Are there two ϵ-NE with TV distance $\geq d$?
Problem 6: Small largest probability $p \in (0, 1)$	Is there an ϵ-NE $(\mathbf{x}, \mathbf{y})$ with $\max_i \mathbf{x}_i \leq p$?
Problem 7: Large total support size $k \in [n]$	Is there an ϵ-WSNE $(\mathbf{x}, \mathbf{y})$ such that $\lvert\text{supp}(\mathbf{x})\rvert + \lvert\text{supp}(\mathbf{y})\rvert \geq 2k$?
Problem 8: Large smallest support size $k \in [n]$	Is there an ϵ-WSNE $(\mathbf{x}, \mathbf{y})$ such that $\min\{\lvert\text{supp}(\mathbf{x})\rvert, \lvert\text{supp}(\mathbf{y})\rvert\} \geq k$?
Problem 9: Large support size $k \in [n]$	Is there an ϵ-WSNE $(\mathbf{x}, \mathbf{y})$ such that $\lvert\text{supp}(\mathbf{x})\rvert \geq k$?
Problem 10: Restricted support $S_R \subseteq [n]$	Is there an ϵ-WSNE $(\mathbf{x}, \mathbf{y})$ with $S_R \subseteq \text{supp}(\mathbf{x})$?

Table 1: The decision problems that we consider. All of them take as input a bimatrix games and a quality of approximation $\epsilon \in (0, 1)$. Problems 1 - 6 relate to ϵ-NE, and Problems 7 - 10 relate to ϵ-WSNE.

Other related work. The only positive result for finding ϵ-NE with good social welfare that we are aware of was given by Czumaj, Fasoulakis, and Jurdziński [9]. They showed that if there is a polynomial-time algorithm for finding an ϵ-NE, then for all $\epsilon' > \epsilon$ there is also a polynomial-time algorithm for finding an ϵ'-NE that is within a constant multiplicative approximation of the best social welfare. They also give further results for the case where $\epsilon > \frac{1}{2}$.

2 Preliminaries

Throughout the paper, we use $[n]$ to denote the set of integers $\{1, 2, \ldots, n\}$. An $n \times n$ bimatrix game is a pair (R, C) of two $n \times n$ matrices: R gives payoffs for the *row* player and C gives the payoffs for the *column* player.

Each player has n *pure* strategies. To play the game, both players simultaneously select a pure strategy: the row player selects a row $i \in [n]$, and the column player selects a column $j \in [n]$. The row player then receives payoff $R_{i,j}$, and the column player receives payoff $C_{i,j}$.

A *mixed strategy* is a probability distribution over $[n]$. We denote a mixed strategy for the row player as a vector $\mathbf{x}$ of length n, such that $\mathbf{x}_i$ is the probability that the row player assigns to pure strategy i. A mixed strategy of the column player is a vector $\mathbf{y}$ of length n, with the same interpretation. If $\mathbf{x}$ and $\mathbf{y}$ are mixed strategies for the row and the column player, respectively, then we call $(\mathbf{x}, \mathbf{y})$ a *mixed strategy profile.* The expected payoff for the row player under strategy profile $(\mathbf{x}, \mathbf{y})$ is given by $\mathbf{x}^T R\mathbf{y}$ and for the column player by $\mathbf{x}^T C\mathbf{y}$. We denote the *support* of a strategy $\mathbf{x}$ as $\text{supp}(\mathbf{x})$, which gives the set of pure strategies i such that $\mathbf{x}_i > 0$.

Nash equilibria. Let $\mathbf{y}$ be a mixed strategy for the column player. The set of *pure best responses* against $\mathbf{y}$ for the row player is the set of pure strategies that maximize the payoff against $\mathbf{y}$. More formally, a pure strategy $i \in [n]$ is a best response against $\mathbf{y}$ if, for all pure strategies $i' \in [n]$ we have: $\sum_{j \in [n]} \mathbf{y}_j \cdot R_{i,j} \geq \sum_{j \in [n]} \mathbf{y}_j \cdot R_{i',j}$. Column player best responses are defined analogously.

A mixed strategy profile $(\mathbf{x}, \mathbf{y})$ is a *mixed Nash equilibrium* if every pure strategy in $\text{supp}(\mathbf{x})$ is a best response against $\mathbf{y}$, and every pure strategy in $\text{supp}(\mathbf{y})$ is a best response against $\mathbf{x}$. Nash [20] showed that every bimatrix game has a mixed Nash equilibrium. Observe that in a Nash equilibrium, each player's expected payoff is equal to their best response payoff.

Approximate Equilibria. There are two commonly studied notions of approximate equilibrium, and we consider both of them in this paper. The first notion is that of an *ϵ-approximate Nash equilibrium* (ϵ-NE), which weakens the requirement that a player's expected payoff should be equal to their best response payoff. Formally, given a strategy profile $(\mathbf{x}, \mathbf{y})$, we define the *regret* suffered by the row player to be the difference between the best response payoff and the actual payoff: $\max_{i \in [n]} \big((R \cdot y)_i\big) - \mathbf{x}^T \cdot R \cdot \mathbf{y}$. Regret for the column player is defined analogously. We have that $(\mathbf{x}, \mathbf{y})$ is an ϵ-NE if and only if both players have regret less than or equal to ϵ.

The other notion is that of an ϵ-approximate-well-supported equilibrium (ϵ-WSNE), which weakens the requirement that players only place probability on best response strategies. Given a strategy profile $(\mathbf{x}, \mathbf{y})$ and a pure strategy $j \in [n]$, we say that j is an ϵ-best-response for the row player if: $\max_{i\in[n]} \big((R \cdot y)_i\big) - (R \cdot y)_j \leq \epsilon$. An ϵ-WSNE requires that both players only place probability on ϵ-best-responses. Formally, the row player's *pure strategy regret* under $(\mathbf{x}, \mathbf{y})$ is defined to be: $\max_{i\in[n]} \big((R\cdot y)_i\big) - \min_{i\in \mathrm{supp}(\mathbf{x})} \big((R\cdot y)_i\big)$. Pure strategy regret for the column player is defined analogously. A strategy profile $(\mathbf{x}, \mathbf{y})$ is an ϵ-WSNE if both players have pure strategy regret less than or equal to ϵ.

Since approximate Nash equilibria use an additive notion of approximation, it is standard practice to rescale the input game so that all payoffs lie in the range $[0, 1]$, which allows us to compare different results on this topic. For the most part, we follow this convention. However, for our result in Section 3, we will construct a game whose payoffs do not lie in $[0, 1]$. In order to simplify the proof, we will prove results about approximate Nash equilibria in the unscaled game, and then rescale the game to $[0, 1]$ at the very end. To avoid confusion, we will refer to an ϵ-approximate Nash equilibrium in this game as an ϵ-UNE, to mark that it is an additive approximation in an unscaled game.

Two-prover games. A two-prover game is defined as follows.

Definition 1 (Two-prover game). *A two-prover game $\mathcal{T}$ is defined by a tuple $(X, Y, A, B, \mathcal{D}, V)$ where X and Y are finite sets of* questions*, A and B are finite sets of* answers*, $\mathcal{D}$ is a probability distribution defined over $\in X \times Y$, and V is a* verification function *of the form $V : X \times Y \times A \times B \to \{0, 1\}$.*

The game is a co-operative game played between two players, who are called Merlin$_1$ and Merlin$_2$, and an adjudicator called Arthur. At the start of the game, Arthur chooses a question pair $(x, y) \in X \times Y$ randomly according to $\mathcal{D}$. He then sends x to Merlin$_1$ and y to Merlin$_2$. Crucially, Merlin$_1$ does not know the question sent to Merlin$_2$ and vice versa. Having received x, Merlin$_1$ then chooses an answer from A and sends it back to Arthur. Merlin$_2$ similarly picks an answer from B and returns it to Arthur. Arthur then computes $p = V(x, y, a, b)$ and awards payoff p to both players. The size of the game, denoted $|\mathcal{T}| = |X \times Y \times A \times B|$ is the total number of entries needed to represent V as a table.

A *strategy* for Merlin$_1$ is a function $a : X \to A$ that gives an answer for every possible question, and likewise a strategy for Merlin$_2$ is a function $b : Y \to B$. We define S_i to be the set of all strategies for Merlin$_i$. The *payoff* of the game under a pair of strategies $(s_1, s_2) \in S_1 \times S_2$ is denoted as $p(\mathcal{T}, s_1, s_2) = E_{(x,y)\sim\mathcal{D}}[V(x, y, s_1(x), s_2(y))]$.

The *value* of the game, denoted $\omega(\mathcal{T})$, is the maximum expected payoff to the Merlins when they play optimally: $\omega(\mathcal{T}) = \max_{s_1\in S_1} \max_{s_2\in S_2} p(\mathcal{T}, s_1, s_2)$.

Free games. A two-prover game is called a *free game* if the probability distribution $\mathcal{D}$ is the uniform distribution $\mathcal{U}$ over $X \times Y$. In particular, this means that there is no correlation between the question sent to Merlin$_1$ and the question sent to Merlin$_2$. We are interested in the problem of approximating the value of a free game within an additive error of δ.

FreeGame_δ
Input: A free game $\mathcal{T}$ and a constant $\delta > 0$.
Output: A value p such that $\mid \omega(\mathcal{T}) - p \mid \leq \delta$.

The *exponential time hypothesis (ETH)* is the conjecture that any deterministic algorithm for solving 3SAT requires $2^{\Omega(n)}$ time. The *randomized exponential time hypothesis (RETH)* is the same hypothesis, but for randomized algorithms. Aaronson, Impagliazzo, and Moshkovitz have shown that, if ETH holds, then we have the following inapproximability result [1].

Theorem 2 (Theorem 32 in [2]). *If the ETH holds, then there exists a constant Δ such that for all $\delta \in [1/n, \Delta]$ the problem* FreeGame_δ *cannot be solved faster than* $n^{\frac{\tilde{O}(\log n)}{\delta}}$.

This theorem was proved by providing a family of games such that, each game $\mathcal{F}$ had either $\omega(\mathcal{F}) = 1$, or $\omega(\mathcal{F}) < 1 - \delta$, and showing that it is hard to decide which of these is the case. Theorem 2 produces a free game where the size of the question sets X and Y is proportional to the size of the answer sets A and B. For our proof we would like the size of X and Y to be *logarithmic* in the size of A and B. Fortunately, this can be achieved by applying the following sub-sampling result from the same paper. Since our results will rely on this sub-sampling lemma, our lower bounds will depend on RETH, rather than ETH.

Lemma 3 (Corollary 46 in [2]). *Given a free game $\mathcal{F} = (X, Y, A, B, \mathcal{U}, V)$ and $\epsilon > 0$, we can randomly select a free game $\mathcal{F}' = (X', Y', A, B, \mathcal{U}, V)$ such that $|X| = |Y| = 2 \cdot \epsilon^{-\Lambda} \cdot \log(|A| + |B|)$ for some constant Λ such that, with high probability, we have $|\omega(\mathcal{F}) - \omega(\mathcal{F}')| \leq \epsilon$.*

3 Hardness of approximating social welfare

Overview. In this section, we study the following *social welfare* problem. The *social welfare* of a pair of strategies $(\mathbf{x}, \mathbf{y})$ is denoted by $\text{SW}(\mathbf{x}, \mathbf{y})$ and is defined to be $\mathbf{x}^T R\mathbf{y} + \mathbf{x}^T C\mathbf{y}$. Given an $\epsilon \geq 0$, we define the set of all ϵ equilibria as $E^\epsilon = \{(\mathbf{x}, \mathbf{y}) \ : \ (\mathbf{x}, \mathbf{y}) \text{ is an } \epsilon\text{-NE}\}$. Then, we define the *best social welfare* achievable by an ϵ-NE in $\mathcal{G}$ as $\text{BSW}(\mathcal{G}, \epsilon) = \max\{\text{SW}(\mathbf{x}, \mathbf{y}) \ : \ (\mathbf{x}, \mathbf{y}) \in E^\epsilon\}$. Using these definitions we now define the main problem that we consider:

ϵ-NE δ-SW
Input: A game $\mathcal{G}$, and two constants $\epsilon, \delta > 0$.
Output: An ϵ-NE $(\mathbf{x}, \mathbf{y})$ s.t. $\text{SW}(\mathbf{x}, \mathbf{y})$ is within δ of $\text{BSW}(\mathcal{G}, \epsilon)$.

We show a lower bound for this problem by reducing from FreeGame_δ. Let $\mathcal{F}$ be a free game of size n from the family of free games that were used to prove Theorem 2. We have that either $\omega(\mathcal{F}) = 1$ or $\omega(\mathcal{F}) < 1 - \delta$ for some fixed constant δ, and that it is hard to determine which of these is the case. We will construct a game $\mathcal{G}$ such that for $\epsilon = 1 - 4g \cdot \delta$, where $g < \frac{5}{12}$ is a fixed constant that we will define at the end of the proof, we have the following properties.

- (**Completeness**) If $\omega(\mathcal{F}) = 1$, then the unscaled $\mathrm{BSW}(\mathcal{G}, \epsilon) = 2$.
- (**Soundness**) If $\omega(\mathcal{F}) < 1 - \delta$, then the unscaled $\mathrm{BSW}(\mathcal{G}, \epsilon) < 2(1 - g \cdot \delta)$.

This will allow us to prove our lower bound using Theorem 2.

3.1 The construction

The first step of the proof is to apply Lemma 3 to $\mathcal{F}$ with $\epsilon = \delta/2$ to produce a free game $\mathcal{F}_{\mathrm{s}} = (X, Y, A, B, \mathcal{U}, V)$ that will be fixed for the rest of this section. Since the question sets in $\mathcal{F}$ have size $O(|\mathcal{F}|)$, we have that the question sets X and Y in $\mathcal{F}_{\mathrm{s}}$ have size $\log(|\mathcal{F}|)$. Furthermore, with high probability, it is hard to decide whether $\omega(\mathcal{F}_{\mathrm{s}}) = 1$ or $\omega(\mathcal{F}_{\mathrm{s}}) = 1 - \delta/2$. Next, we use $\mathcal{F}_{\mathrm{s}}$ to construct a bimatrix game, which we will denote as $\mathcal{G}$ throughout the rest of this section. The game is built out of four subgames, which are arranged and defined as follows.

I \ II		
	R , C	$-D_2$, D_2
	D_1 , $-D_1$	0 , 0

- The game (R, C) is built from $\mathcal{F}_{\mathrm{s}}$ in the following way. Each row of the game corresponds to a pair $(x, a) \in X \times A$ and each column corresponds to a pair $(y, b) \in Y \times B$. Since all free games are cooperative, the payoff for each strategy pair $(x, a), (y, b)$ is defined to be $R_{(x,a),(y,b)} = C_{(x,a),(y,b)} = V(x, y, a(x), b(y))$.
- The game $(D_1, -D_1)$ is a zero-sum game. The game is a slightly modified version of a game devised by Feder, Nazerzadeh, and Saberi [14]. Let H be the set of all functions of the form $f : Y \to \{0, 1\}$ such that $f(y) = 1$ for exactly half[4] of the elements $y \in Y$. The game has $|Y \times B|$ columns and $|H|$ rows. For all $f \in H$ and all $(y, b) \in Y$ the payoffs are

$$(D_1)_{f,(y,b)} = \begin{cases} \frac{4}{1+4g\cdot\delta} & \text{if } f(y) = 1, \\ 0 & \text{otherwise.} \end{cases}$$

- The game $(-D_2, D_2)$ is built in the same way as the game $(D_1, -D_1)$, but with the roles of the players swapped. That is, each column of $(-D_2, D_2)$ corresponds to a function that picks half of the elements of X.
- The game $(0, 0)$ is a game in which both players have zero matrices.

[4] If $|Y|$ is not even, then we can create a new free game in which each question in $|Y|$ appears twice. This will not change the value of the free game.

Observe that the size of (R, C) is the same as the size of $\mathcal{F}_{\mathrm{s}}$, which is at most $|\mathcal{F}|$. The game $(D_1, -D_1)$ has the same number of columns as C, and the number of rows is at most $2^{|Y|} \leq 2^{2\delta^{-\Lambda} \log |\mathcal{F}|} = |\mathcal{F}|^{2\delta^{-\Lambda}}$, where Λ is the constant obtained from Lemma 3. By the same reasoning, the number of columns in $(-D_2, D_2)$ is at most $|\mathcal{F}|^{2\delta^{-\Lambda}}$. Thus, the size of $\mathcal{G}$ is $|\mathcal{F}|^{O(\delta^{-\Lambda})}$, and in particular, for every constant $\delta > 0$, this reduction is polynomial.

3.2 Completeness

To prove completeness, it suffices to show that, if $\omega(\mathcal{F}_{\mathrm{s}}) = 1$, then there exists a $(1-4g\cdot\delta)$-UNE of $\mathcal{G}$ that has social welfare 2. To do this, assume that $\omega(\mathcal{F}_{\mathrm{s}}) = 1$, and take a pair of optimal strategies (s_1, s_2) for $\mathcal{F}_{\mathrm{s}}$ and turn them into strategies for the players in $\mathcal{G}$. More precisely, the row player will place probability $\frac{1}{|X|}$ on each answer chosen by s_1, and the column player will place probability $\frac{1}{|Y|}$ on each answer chosen by s_2. By construction, this gives both players payoff 1, and hence the social welfare is 2. The hard part is to show that this is an approximate equilibrium, and in particular, that neither player can gain by playing a strategy in $(D_1, -D_1)$ or $(-D_2, D_2)$. We prove this in the following lemma.

Lemma 4. *If $\omega(\mathcal{F}_s) = 1$, then there exists a $(1 - 4g \cdot \delta)$-UNE $(\mathbf{x}, \mathbf{y})$ of $\mathcal{G}$ with* $\mathrm{SW}(\mathbf{x}, \mathbf{y}) = 2$.

3.3 Soundness

We now suppose that $\omega(\mathcal{F}_{\mathrm{s}}) < 1-\delta/2$, and we will prove that all $(1-4g\cdot\delta)$-UNE provide social welfare at most $2 - 2g \cdot \delta$. Throughout this subsection, we will fix $(\mathbf{x}, \mathbf{y})$ to be a $(1 - 4g \cdot \delta)$-UNE of $\mathcal{G}$. We begin by making a simple observation about the amount of probability that is placed on (R, C).

Lemma 5. *If* $\mathrm{SW}(\mathbf{x}, \mathbf{y}) > 2-2g\cdot\delta$*, then* $\mathbf{x}$ *places at least* $(1-g\cdot\delta)$ *probability on rows in* (R, C)*, and* $\mathbf{y}$ *places at least* $(1 - g \cdot \delta)$ *probability on columns in* (R, C).

So, for the rest of this subsection, we can assume that both $\mathbf{x}$ and $\mathbf{y}$ place at least $1 - g \cdot \delta$ probability on the subgame (R, C). We will ultimately show that, if this is the case, then both players have payoff at most $1 - \frac{1}{2} \cdot \delta + mg \cdot \delta$ for some constant m that will be derived during the proof. Choosing $g = 1/(2m+2)$ then ensures that both players have payoff at most $1 - g \cdot \delta$, and therefore that the social welfare is at most $2 - 2g \cdot \delta$.

A two-prover game. We use $(\mathbf{x}, \mathbf{y})$ to create a two-prover game. First, we define two distributions that capture the marginal probability that a question is played by $\mathbf{x}$ or $\mathbf{y}$. Formally, we define a distribution $\mathbf{x}'$ over X and a distribution $\mathbf{y}'$ over Y such that for all $x \in X$ and $y \in Y$ we have $\mathbf{x}'(x) = \sum_{a\in A} \mathbf{x}(x, a)$, and $\mathbf{y}'(y) = \sum_{b\in B} \mathbf{y}(y, b)$. By Lemma 5, we can assume that $\|\mathbf{x}'\|_1 \geq 1 - g \cdot \delta$ and $\|\mathbf{y}'\|_1 \geq 1 - g \cdot \delta$.

Our two-prover game will have the same question sets, answer sets, and verification function as $\mathcal{F}_{\mathrm{s}}$, but a different distribution over the question sets. Let $\mathcal{T}_{(\mathbf{x},\mathbf{y})} = (X, Y, A, B, \mathcal{D}, V)$, where $\mathcal{D}$ is the product of $\mathbf{x}'$ and $\mathbf{y}'$. Note that we have cheated slightly here, since $\mathcal{D}$ is not actually a probability distribution.

If $\|\mathcal{D}\|_1 = c < 1$, then we can think of this as Arthur having a $1 - c$ probability of not sending any questions to the Merlins and awarding them payoff 0.

The strategies $\mathbf{x}$ and $\mathbf{y}$ can also be used to give a us a strategy for the Merlins in $\mathcal{T}_{(\mathbf{x},\mathbf{y})}$. Without loss of generality, we can assume that for each question $x \in X$ there is exactly one answer $a \in A$ such that $\mathbf{x}(x, a) > 0$, because if there are two answers a_1 and a_2 such that $\mathbf{x}(x, a_1) > 0$ and $\mathbf{x}(x, a_2) > 0$, then we can shift all probability onto the answer with (weakly) higher payoff, and (weakly) improve the payoff to the row player. Since (R, C) is cooperative, this can only improve the payoff of the columns in (R, C), and since the row player does not move probability between questions, the payoff of the columns in $(-D_2, D_2)$ does not change either. Thus, after shifting, we arrive at a $(1 - 4g \cdot \delta)$-UNE of $\mathcal{G}$ whose social welfare is at least as good as $\mathrm{SW}(\mathbf{x}, \mathbf{y})$. Similarly, we can assume that for each question $y \in Y$ there is exactly one answer $b \in B$ such that $\mathbf{y}(y, b) > 0$.

So, we can define a strategy $s_{\mathbf{x}}$ for Merlin$_1$ in the following way. For each question $x \in X$, the strategy $s_{\mathbf{x}}$ selects the unique answer $a \in A$ such that $\mathbf{x}(x, a) > 0$. The strategy $s_{\mathbf{y}}$ for Merlin$_2$ is defined symmetrically.

We will use $\mathcal{T}_{(\mathbf{x},\mathbf{y})}$ as an intermediary between $\mathcal{G}$ and $\mathcal{F}_{\mathrm{s}}$ by showing that the payoff of $(\mathbf{x}, \mathbf{y})$ in $\mathcal{G}$ is close to the payoff of $(s_{\mathbf{x}}, s_{\mathbf{y}})$ in $\mathcal{T}_{(\mathbf{x},\mathbf{y})}$, and that the payoff of $(s_{\mathbf{x}}, s_{\mathbf{y}})$ in $\mathcal{T}_{(\mathbf{x},\mathbf{y})}$ is close to the payoff of $(s_{\mathbf{x}}, s_{\mathbf{y}})$ in $\mathcal{F}_{\mathrm{s}}$. Since we have a bound on the payoff of any pair of strategies in $\mathcal{F}_{\mathrm{s}}$, this will ultimately allow us to bound the payoff to both players when $(\mathbf{x}, \mathbf{y})$ is played in $\mathcal{G}$.

Relating $\mathcal{G}$ to $\mathcal{T}_{(\mathbf{x},\mathbf{y})}$. For notational convenience, let us define $p_r(\mathcal{G}, \mathbf{x}, \mathbf{y})$ and $p_c(\mathcal{G}, \mathbf{x}, \mathbf{y})$ to be the payoff to the row player and column player, respectively, when $(\mathbf{x}, \mathbf{y})$ is played in $\mathcal{G}$. We begin by showing that the difference between $p_r(\mathcal{G}, \mathbf{x}, \mathbf{y})$ and $p(\mathcal{T}_{(\mathbf{x},\mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})$ is small. Once again we prove this for the payoff of the row player, but the analogous result also holds for the column player.

Lemma 6. *We have* $|p_r(\mathcal{G}, \mathbf{x}, \mathbf{y}) - p(\mathcal{T}_{(\mathbf{x},\mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})| \leq 4g \cdot \delta$.

Relating $\mathcal{T}_{(\mathbf{x},\mathbf{y})}$ to $\mathcal{F}_{\mathbf{s}}$. First we show that if $(\mathbf{x}, \mathbf{y})$ is indeed a $(1 - 4g \cdot \delta)$-UNE, then $\mathbf{x}'$ and $\mathbf{y}'$ must be close to uniform over the questions. We prove this for $\mathbf{y}'$, but the proof can equally well be applied to $\mathbf{x}'$. The idea is that, if $\mathbf{y}'$ is sufficiently far from uniform, then there is set $B \subseteq Y$ of $|Y|/2$ columns where $\mathbf{y}'$ places significantly more than 0.5 probability. This, in turn, means that the row of $(D_1, -D_1)$ that corresponds to B, will have payoff at least 2, while the payoff of $(\mathbf{x}, \mathbf{y})$ can be at most $1 + 3g \cdot \delta$, and so $(\mathbf{x}, \mathbf{y})$ would not be a $(1 - 4g \cdot \delta)$-UNE. We formalise this idea in the following lemma. Define $\mathbf{u}_X$ to be the uniform distribution over X, and $\mathbf{u}_Y$ to be the uniform distribution over Y.

Lemma 7. *We have* $\|\mathbf{u}_Y - \mathbf{y}'\|_1 < 16g \cdot \delta$ *and* $\|\mathbf{u}_X - \mathbf{x}'\|_1 < 16g \cdot \delta$.

With Lemma 7 at hand, we can now prove that the difference between $p(\mathcal{T}_{(\mathbf{x},\mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})$ and $p(\mathcal{F}_{\mathrm{s}}, s_{\mathbf{x}}, s_{\mathbf{y}})$ must be small. This is because the question distribution $\mathcal{D}$ used in $\mathcal{T}_{(\mathbf{x},\mathbf{y})}$is a product of two distributions that are close to uniform, while the question distribution $\mathcal{U}$ used in $\mathcal{F}_{\mathrm{s}}$ is a product of two uniform distributions. In the following lemma, we show that if we transform $\mathcal{D}$ into $\mathcal{U}$, then we do not change the payoff of $(s_{\mathbf{x}}, s_{\mathbf{y}})$ very much.

Lemma 8. *We have* $|p(\mathcal{T}_{(\mathbf{x},\mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}}) - p(\mathcal{F}_s, s_{\mathbf{x}}, s_{\mathbf{y}})| \le 64g \cdot \delta$.

Completing the soundness proof. The following lemma uses the bounds derived in Lemmas 6 and 8, along with a suitable setting for g, to bound the payoff of both players when $(\mathbf{x}, \mathbf{y})$ is played in $\mathcal{G}$.

Lemma 9. *If* $g = \frac{1}{138}$*, then both players have payoff at most* $1 - g \cdot \delta$ *when* $(\mathbf{x}, \mathbf{y})$ *is played in* $\mathcal{G}$.

Hence, we have proved that $\mathrm{SW}(\mathbf{x}, \mathbf{y}) \le 2 - 2g \cdot \delta$.

3.4 The result

We can now state the theorem that we have proved in this section. We first rescale the game so that it lies in $[0, 1]$. The maximum payoff in $\mathcal{G}$ is $\frac{4}{1+4g\cdot\delta} \le 4$, and the minimum payoff is $-\frac{4}{1+4g\cdot\delta} \ge -4$. To rescale this game, we add 4 to all the payoffs, and then divide by 8. Let us refer to the scaled game as $\mathcal{G}_s$. Observe that an ϵ-UNE in $\mathcal{G}$ is a $\frac{\epsilon}{8}$-NE in $\mathcal{G}_s$ since adding a constant to all payoffs does not change the approximation guarantee, but dividing all payoffs by a constant does change the approximation guarantee. So, we have the following theorem.

Theorem 10. *If RETH holds, then there exists a constant* Δ *such that for all* $\delta \in [1/n, \Delta]$ *the problem* $(\frac{1-4g\cdot\delta}{8})$*-NE* $\frac{g}{4} \cdot \delta$*-SW, where* $g = \frac{1}{138}$*, cannot be solved faster than* $n^{\widetilde{O}(\delta^{\Lambda} \log n)}$*, for some fixed constant* Λ.

4 Hardness results for other decision problems

In this section we study a range of decision problems associated with approximate equilibria. Most are known to be NP-complete for the case of exact Nash equilibria [16, 7]. Table 1 shows all of the decision problems that we consider. For each problem, the input includes a bimatrix game and a quality of approximation $\epsilon \in (0, 1)$. We consider decision problems related to both ϵ-NE and ϵ-WSNE. Since ϵ-NE is a weaker concept than ϵ-WSNE, the hardness results for ϵ-NE imply the same hardness for ϵ-WSNE. We consider problems for ϵ-WNSE only where the corresponding problem for ϵ-NE is trivial. For example, observe that approximate ϵ-NE with large support is a trivial problem, since we can always add a tiny amount of probability to each pure strategy without changing our expected payoff very much.

Our conditional quasi-polynomial lower bounds will hold for all $\epsilon < \frac{1}{8}$. Thus fix $\epsilon^* < \frac{1}{8}$ for the rest of this section. We will appeal to Theorem 10, and thus we compute from ϵ^* the parameters n and δ that we require to apply this theorem. In particular, compute δ^* to solve $\epsilon^* = (\frac{1-4g\cdot\delta^*}{8})$, which comes from Theorem 10, and choose n^* as $\frac{1}{\delta^*}$. Then, for $n > n^*$ and $\delta = \delta^*$ we can apply Theorem 10 to bound the social welfare acheivable if $\omega(\mathcal{F}_s) < 1 - \delta^*$ as $\mathfrak{u} = \frac{6}{8} - \frac{1}{522}\delta^*$. Theorem 10 implies that in order to decide whether the game $\mathcal{G}_s$ possess an ϵ^*-NE that yields social welfare strictly greater than $\mathfrak{u}$ requires $n^{\tilde{O}(\log n)}$ time, where δ no longer appears in the exponent since we have fixed it as a constant δ^* according to our choice of ϵ^*.

The hardness of Problem 1 is a corollary of Theorem 10 when we set $u = \frac{3}{8}$. For the other problems in Table 1, we use $\mathcal{G}_s$ to construct two new games: $\mathcal{G}'$, which adds one row and column to $\mathcal{G}_s$, is used to show hardness of Problems 2 - 9, and $\mathcal{G}''$, which in turn adds one row and column to $\mathcal{G}'$, is used to show hardness of Problem 10. We define $\mathcal{G}'$ and $\mathcal{G}''$ using the constants $\mathfrak{u}$ and ϵ^* fixed above. The game $\mathcal{G}'$ extends $\mathcal{G}_s$ by adding the pure strategy $\mathfrak{i}$ for the row player, and the pure strategy $\mathfrak{j}$ for the column player, with payoffs as shown in Figure 1.

$$\mathcal{G}' = \begin{array}{c|ccc|c|}
 & & & & \mathfrak{j} \\
\cline{2-5}
 & & & & 0, \frac{3}{8} + \epsilon^* \\
\cline{5-5}
 & & \mathcal{G}_s & & \vdots \\
\cline{5-5}
 & & & & 0, \frac{3}{8} + \epsilon^* \\
\cline{2-5}
\mathfrak{i} & \multicolumn{1}{|c|}{\frac{3}{8} + \epsilon^*, 0} & \multicolumn{1}{c|}{\cdots} & \frac{3}{8} + \epsilon^*, 0 & 1, 1 \\
\cline{2-5}
\end{array}$$

Fig. 1: The game $\mathcal{G}'$.

The payoffs for $\mathfrak{i}$ and $\mathfrak{j}$ were chosen so that: If the game $\mathcal{G}_s$ possess an ϵ^*-NE with social welfare $\frac{6}{8}$, then $\mathcal{G}'$ posses at least one ϵ^*-NE where the players do not play the pure strategies $\mathfrak{i}$ and $\mathfrak{j}$; if every ϵ^*-NE of the game $\mathcal{G}_s$ yields social welfare at most $\mathfrak{u}$, then in *every* ϵ^*-NE of $\mathcal{G}'$, the players place almost all of their probability on $\mathfrak{i}$ and $\mathfrak{j}$ respectively. Lemmas 11 and 12 show further properties that hold in the first case but not the second.

Notice that the expected payoff for the row player from the pure strategy $\mathfrak{i}$ is at least $\frac{3}{8} + \epsilon^*$ irrespective of the strategy the column player chooses. The same holds for the column player and the pure strategy $\mathfrak{j}$, i.e., the expected payoff that the column players gets from the pure strategy $\mathfrak{j}$ is at least $\frac{3}{8} + \epsilon^*$ irrespective of the strategy chosen by the row player. In what follows we will use S_R (S_C), or S when it is clear from the context, to denote the set of pure strategies available to the row (column) from the (R, C) part of $\mathcal{G}_s$ that corresponds to different questions in the free game $\mathcal{F}_s$.

First, we derive some properties of the equilibria of $\mathcal{G}'$ when $\mathcal{G}_s$ posses an ϵ^*-NE with social welfare $\frac{6}{8}$.

Lemma 11. *If $\mathcal{G}_s$ posses an ϵ^*-NE $(\mathbf{x}, \mathbf{y})$ with social welfare $\frac{6}{8}$, then $(\mathbf{x}, \mathbf{y})$ is an ϵ^*-WSNE for $\mathcal{G}'$ such that:*

(a) $\mathbf{x}^T R \mathbf{y} = \frac{3}{8}$ *and* $\mathbf{x}^T C \mathbf{y} = \frac{3}{8}$*,*
(b) $\mathrm{supp}(\mathbf{x}) \subseteq S_R$ *and* $\mathrm{supp}(\mathbf{y}) \subseteq S_C$*,*
(c) $|\mathrm{supp}(\mathbf{x})| = |S_R|$ *and* $|\mathrm{supp}(\mathbf{y})| = |S_C|$*,*
(d) $\max_i \mathbf{x}_i \leq \frac{1}{|S_R|}$ *and* $\max_j \mathbf{y}_j \leq \frac{1}{|S_C|}$*.*

Next, we prove certain properties that *all* ϵ^*-NE and ϵ^*-WSNE of $\mathcal{G}'$ possess if every ϵ^*-NE of $\mathcal{G}_s$ yields social welfare at most $\mathfrak{u}$.

Lemma 12. *If every ϵ^*-NE of $\mathcal{G}_s$ yields social welfare at most $\mathfrak{u}$, then in every ϵ^*-NE $(\mathbf{x}, \mathbf{y})$ of $\mathcal{G}'$ it holds that:*

(α) $\mathbf{x}_\mathfrak{i} > 1 - \epsilon^*$ *and* $\mathbf{y}_\mathfrak{j} > 1 - \epsilon^*$*,*

(β) $\mathbf{x}^T R\mathbf{y} > 1 - 2\epsilon^*$ *and* $\mathbf{x}^T C\mathbf{y} > 1 - 2\epsilon^*$*.*

Furthermore, in every ϵ^*-WSNE $(\mathbf{x}, \mathbf{y})$ of $\mathcal{G}'$ it holds that

(γ) $|\mathrm{supp}(\mathbf{x})| = |\mathrm{supp}(\mathbf{y})| = 1$*.*

Observe that the combination of the claims of Lemmas 11 and 12 give the desired hardness results for the Problems 2 - 9. The combination of claim (a) from Lemma 11 with the claim (β) from Lemma 12 gives the hardness result for the Problems 2 and 3; the combination of (b) with (α) gives the hardness for Problems 4 and 5; the combination of (d) with (α) gives the hardness for the Problem 6; and finally that hardness of Problems 7 - 9 follows from the combination of (c) with (γ).

For Problem 10, we define a new game $\mathcal{G}''$ by extending $\mathcal{G}'$. We add the new pure strategy $\mathfrak{i}'$ for the row player and the new pure strategy $\mathfrak{j}'$ for the column player, with payoffs constructed as shown in Figure 2. We prove that if the game $\mathcal{G}_s$ posses an ϵ^*-NE with social welfare $\frac{3}{8}$, then the game $\mathcal{G}''$ possess an ϵ^*-WSNE $(\mathbf{x}, \mathbf{y})$ such that $\mathfrak{i}' \in \mathrm{supp}(\mathbf{x})$. Furthermore, we prove that if all ϵ^*-NE of $\mathcal{G}_S$ yield social welfare at most $\mathfrak{u}$, then for any ϵ^*-WSNE $(\mathbf{x}, \mathbf{y})$ it holds that $\mathfrak{i}' \notin \mathrm{supp}(\mathbf{x})$.

$$
\mathcal{G}'' = \begin{array}{c|ccc|c|}
 & & & & \mathfrak{j}' \\
\hline
 & & & & \frac{3}{8}, \frac{3}{8} \\
 & & \mathcal{G}' & & \vdots \\
 & & & & \frac{3}{8}, \frac{3}{8} \\
\hline
\mathfrak{i}' & \frac{3}{8}, \frac{3}{8} & \cdots & \frac{3}{8}, \frac{3}{8} & 0, 0 \\
\hline
\end{array}
$$

Fig. 2: The game $\mathcal{G}''$.

Lemma 13. *If the game $\mathcal{G}_s$ posses an ϵ^*-NE with social welfare $\frac{6}{8}$, then the game $\mathcal{G}''$ posses an ϵ^*-WSNE $(\mathbf{x}, \mathbf{y})$ such that $\mathfrak{i}' \in \mathrm{supp}(\mathbf{x})$.*

Lemma 14. *If all the ϵ^*-NE of $\mathcal{G}_s$ yield social welfare at most $\mathfrak{u}$, then for any ϵ^*-WSNE $(\mathbf{x}, \mathbf{y})$ of $\mathcal{G}''$ it holds that $\mathfrak{i}' \notin \mathrm{supp}(\mathbf{x})$.*

We now summarize the results of this section in the following theorem. Given the game $\mathcal{G}_s$ we can construct games $\mathcal{G}'$ and $\mathcal{G}''$ such that the answer to the Problems 2 - 10 is "Yes" if $\mathcal{G}_s$ possess an ϵ^*-NE with social welfare $\frac{3}{8}$ and "No" if every ϵ^*-NE of $\mathcal{G}_s$ has social welfare at most $\mathfrak{u}$.

Theorem 15. *Assuming the RETH, any algorithm that solves the Problems 1 - 10 for any constant $\epsilon < \frac{1}{8}$ requires $n^{\tilde{\Omega}(\log n)}$ time.*

References

1. S. Aaronson, R. Impagliazzo, and D. Moshkovitz. AM with multiple merlins. In *Proc. of CCC*, pages 44–55, 2014.
2. S. Aaronson, R. Impagliazzo, and D. Moshkovitz. AM with multiple merlins. *CoRR*, abs/1401.6848, 2014.
3. P. Austrin, M. Braverman, and E. Chlamtac. Inapproximability of NP-complete variants of Nash equilibrium. *Theory of Computing*, 9:117–142, 2013.
4. V. Bilò and M. Mavronicolas. A catalog of $\exists\mathbb{R}$-complete decision problems about Nash equilibria in multi-player games. In *Proc. of STACS*, pages 17:1–17:13, 2016.
5. H. Bosse, J. Byrka, and E. Markakis. New algorithms for approximate Nash equilibria in bimatrix games. *Theoretical Computer Science*, 411(1):164–173, 2010.
6. M. Braverman, Y. Kun-Ko, and O. Weinstein. Approximating the best Nash equilibrium in $n^{o(\log n)}$-time breaks the exponential time hypothesis. In *Proc. of SODA*, pages 970–982, 2015.
7. V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621 – 641, 2008.
8. A. Czumaj, A. Deligkas, M. Fasoulakis, J. Fearnley, M. Jurdziński, and R. Savani. Distributed methods for computing approximate equilibria. In *Proc. of WINE*, 2016.
9. A. Czumaj, M. Fasoulakis, and M. Jurdziński. Approximate Nash equilibria with near optimal social welfare. In *Proc. of IJCAI*, pages 504–510, 2015.
10. C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
11. C. Daskalakis, A. Mehta, and C. H. Papadimitriou. Progress in approximate Nash equilibria. In *Proc. of EC*, pages 355–358, 2007.
12. C. Daskalakis, A. Mehta, and C. H. Papadimitriou. A note on approximate Nash equilibria. *Theoretical Computer Science*, 410(17):1581–1588, 2009.
13. J. Fearnley, P. W. Goldberg, R. Savani, and T. B. Sørensen. Approximate well-supported Nash equilibria below two-thirds. In *Proc. of SAGT*, pages 108–119, 2012.
14. T. Feder, H. Nazerzadeh, and A. Saberi. Approximating Nash equilibria using small-support strategies. In *Proc. of EC*, pages 352–354, 2007.
15. J. Garg, R. Mehta, V. V. Vazirani, and S. Yazdanbod. ETR-completeness for decision versions of multi-player (symmetric) Nash equilibria. In *Proc. of ICALP*, pages 554–566, 2015.
16. I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80 – 93, 1989.
17. E. Hazan and R. Krauthgamer. How hard is it to approximate the best Nash equilibrium? *SIAM J. Comput.*, 40(1):79–91, 2011.
18. S. C. Kontogiannis and P. G. Spirakis. Well supported approximate equilibria in bimatrix games. *Algorithmica*, 57(4):653–667, 2010.
19. R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proc. of EC*, pages 36–41, 2003.
20. J. Nash. Non-cooperative games. *The Annals of Mathematics*, 54(2):286–295, 1951.
21. A. Rubinstein. Settling the complexity of computing approximate two-player Nash equilibria. In *Proc. of FOCS*, pages 258–265, 2016.
22. H. Tsaknakis and P. G. Spirakis. An optimization approach for approximate Nash equilibria. *Internet Mathematics*, 5(4):365–382, 2008.