

Inapproximability Results for Constrained Approximate Nash Equilibria

Argyrios Deligkas^a, John Fearnley^b, Rahul Savani^b

^a*Technion-Israel Institute of Technology*

^b*University of Liverpool*

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 exponential-time hypothesis (ETH) is true, then solving $(\frac{1}{8} - O(\delta))$ -NE $O(\delta)$ -SW for an $n \times n$ bimatrix game requires $n^{\tilde{\Omega}(\log n)}$ time. 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 ETH holds, where 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.

Keywords: Approximate Nash equilibrium, constrained equilibrium, quasi-polynomial time, lower bound, Exponential Time Hypothesis.

[☆]The authors were supported by EPSRC grant EP/L011018/1. A short version of this paper appeared at WINE 2016.

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 players' 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 [1], while finding a constrained Nash equilibrium turns out to be even harder. Gilboa and Zemel [2] 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 [3] extended the list of NP-complete problems of [2] and furthermore proved inapproximability results for some of them. Recently, Garg et al. [4] and Bilo and Mavronicolas [5, 6] extended these results to many player games and provided ETR-completeness results for them.

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 [7, 8, 9, 10, 11, 12, 13]. 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 [13], 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 [14]. They proved that there is always an ϵ -NE with support of logarithmic size, and then they use 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 Rubinfeld implies that we cannot do better than a QPTAS like the LMM algorithm [15]: assuming an exponential time hypothesis 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 some of them, which we discuss further in Section 3. The question then arises whether one can do better.

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 [16] and Austrin, Braverman and Chlamtac [17] 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 $O(\log n)$ in the random graph $G_{n,1/2}$. This was further strengthened by Braverman, Ko, and Weinstein [18] 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 ETH there is a small constant ϵ such that any algorithm

for $O(\epsilon)$ -NE $O(\epsilon)$ -SW¹ requires $n^{\text{poly}(\epsilon)\log(n)^{1-o(1)}}$ time². We shall refer to this as the BKW result.

It is worth noting that the Rubinstein’s hardness result [15] almost makes this result redundant. If one is willing to accept that PETH is true, which is a stronger assumption 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 ϵ -NE δ -SW. We show that, assuming ETH, that there exists a small constant δ^* such that for all $\delta < \delta^*$ the problem $\left(\frac{1-4g\delta}{8}\right)$ -NE $\left(\frac{g\delta}{4}\right)$ -SW requires $n^{\tilde{\Omega}(\log n)}$ time³, where $g = \frac{1}{138}$.

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, and asymptotically our ϵ approaches $1/8$. Moreover, since δ is some small constant that is certainly less than 1, our lower bound applies to all ϵ -NE with $\epsilon \leq \frac{1-4g}{8} \approx 0.1214$. By contrast, the BKW result applies to ϵ -NE δ -SW where ϵ is a tiny unspecified constant.

Recall that we have argued that the BKW result is almost made redundant by the subsequent result of Rubinstein. We note that 0.1214 is orders of magnitude larger than the inapproximability bound given by Rubinstein’s hardness result, and so our lower bound 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

¹While the proof in [18] produces a lower bound for 0.8-NE $(1 - O(\epsilon))$ -SW, this is in a game with maximum payoff $O(1/\epsilon)$. Therefore, when the payoffs in this game are rescaled to $[0, 1]$, the resulting lower bound only applies to ϵ -NE ϵ -SW.

²Although the paper claims that they obtain a $n^{\tilde{O}(\log n)}$ lower bound, the proof reduces from the *low error* result from [19] (cf. Theorem 36 in [20]), which gives only the weaker lower bound of $n^{\text{poly}(\epsilon)\log(n)^{1-o(1)}}$.

³Here $\tilde{\Omega}(\log n)$ means $\Omega\left(\frac{\log n}{(\log \log n)^c}\right)$ for some constant c .

lower bound. The exponent in their lower bound of $n^{\text{poly}(\epsilon) \log(n)^{1-o(1)}}$ is logarithmic only in the limit, while the exponent in our lower bound of $n^{\tilde{\Omega}(\log n)}$ is logarithmic up to polyloglog factors.

The second set of results in this paper show that, once we have our lower bound on the problem 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 ETH is true, then for every $\epsilon < \frac{1}{8}$ finding an ϵ -NE with the given property in an $n \times n$ bimatrix game requires $n^{\tilde{\Omega}(\log n)}$ time.

One final point is that, in our lower bound against ϵ -NE δ -SW, our value for ϵ is comparatively large (approaching $1/8$), while our value for δ is only a small constant. While it would obviously be more satisfying for both ϵ and δ to be large constants, we note that the value of $1/8$ for ϵ directly leads to hardness for all of the problems in the second set of results for all $\epsilon < 1/8$.

Techniques. At a high level, the proof of our first result is similar in spirit to the proof of the BKW result. They 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 [19]. A free game is played between two players named Merlin₁ and Merlin₂, and a referee named Arthur. The BKW result creates a bimatrix game that simulates the free game, where the two players take the role of Merlin₁ and Merlin₂, while Arthur is simulated using a zero-sum game.

Our result will also be proved by producing a bimatrix game that simulates a free game. However, there are a number of key differences that allow us to prove the stronger lower bounds described above. The first key difference is that we use a different zero-sum game to simulate Arthur. Our zero-sum game is inspired by one used by Althöfer [21]. The advantage of this construction is that it is capable of ensuring that both players use distributions that are very close to uniform in all approximate Nash equilibria, which in turn gives us a

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: Restricted support $S \subseteq [n]$	Is there an ϵ -NE (\mathbf{x}, \mathbf{y}) with $\text{supp}(\mathbf{x}) \subseteq S$?
Problem 3: Two ϵ -NE $d \in (0, 1]$ apart in Total Variation (TV) distance	Are there two ϵ -NE with TV distance $\geq d$?
Problem 4: Small largest probability $p \in (0, 1)$	Is there an ϵ -NE (\mathbf{x}, \mathbf{y}) with $\max_i x_i \leq p$?
Problem 5: 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 6: 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 7: Large total support size $k \in [n]$	Is there an ϵ -WSNE (\mathbf{x}, \mathbf{y}) such that $ \text{supp}(\mathbf{x}) + \text{supp}(\mathbf{y}) \geq 2k$?
Problem 8: Large smallest support size $k \in [n]$	Is there an ϵ -WSNE (\mathbf{x}, \mathbf{y}) such that $\min\{ \text{supp}(\mathbf{x}) , \text{supp}(\mathbf{y}) \} \geq k$?
Problem 9: Large support size $k \in [n]$	Is there an ϵ -WSNE (\mathbf{x}, \mathbf{y}) such that $ \text{supp}(\mathbf{x}) \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 game (R, C) and a quality of approximation $\epsilon \in (0, 1)$. Problems 1 - 6 relate to ϵ -NE, and Problems 7 - 10 relate to ϵ -WSNE.

very accurate simulation of Arthur.

The downside of this zero-sum game is that it requires 2^n rows to force the column player to mix close to uniformly over n rows. Arthur is required to pick two *questions* uniformly from a set of possible questions. The free games provided by Aaronson, Impagliazzo, and Moshkovitz have question sets of linear size, so if we reduce directly from these games, we would end up with an exponentially sized bimatrix game. The conference version of this paper [22] resolved the issue by using a sub-sampling lemma, also proved by Aaronson, Impagliazzo, and Moshkovitz, that produces a free game with logarithmically sized question sets. This allowed us to produce polynomially sized bimatrix games, but at the cost of needing randomization to implement the reduction, and so the result depended on the randomized version of the ETH.

In this version, we show that we are able to assume only the ETH by using a stronger result that was discovered by Babichenko, Papadimitriou, and Rubinfeld [23]. Their results imply that approximating the value of a free game requires quasipolynomial time even when the size of the question sets is logarithmic in the game size. They do not explicitly formulate this result, but it is clearly implied by their techniques. For the sake of completeness, we provide an exposition of their ideas in Section 4.

The second main difference between our result and the BKW result is that we use a different starting point. The BKW result uses the PCP theorem of Moshkovitz and Raz [24], which provides a completeness/soundness gap of $1 - \delta$ vs δ for arbitrarily small constant δ in the label cover problem. The use of this powerful PCP theorem is necessary, as their proof relies on the large completeness/soundness gap produced by that theorem. This choice of PCP theorem directly impacts the running time lower bound that they produce, as the $(\log n)^{1-o(1)}$ term in the exponent arises from the blowup of $n^{1+o(1)}$ from the PCP theorem.

In contrast to this, our stronger simulation of Arthur allows us to use the PCP theorem of Dinur [25] as our starting point. This PCP theorem only involves a blowup of $n \text{ polylog}(n)$, which directly leads to the improved $\tilde{\Omega}(\log n)$

exponent in our lower bound. The improved blowup of the PCP theorem comes at the cost of providing a completeness/soundness gap of only 1 vs $1 - \delta$ for a small constant δ for 3SAT, which is a smaller gap than the one provided by the Moshkovitz-Raz PCP theorem, but our simulation is strong enough to deal with this.

It is worth noting that, if one assumes the gap-ETH [26] instead of the ETH, then our results give an $\Omega(\log n)$ exponent in the lower bound, rather than $\tilde{\Omega}(\log n)$. The gap-ETH conjecture states that it requires 2^{cn} time, for some constant c , to distinguish between 3SAT instances where all of the clauses can be satisfied, and where at most a $1 - \delta$ fraction of the clauses can be satisfied, for some constant δ . The gap-ETH would be implied by the ETH if someone was able to devise a PCP theorem with linear blowup. The reason that the gap-ETH improves our result is that the exponent $\tilde{\Omega}(\log n)$ comes directly from the super-linear n polylog(n) blowup of Dinur's PCP theorem.

One final point of comparison is the size of the payoffs used in our simulation. The zero-sum games that we use have payoffs in the range $(-4, 4)$, which directly leads to the $\frac{1-4g-\delta}{8}$ bound on the quality of approximation. In contrast to this, the zero-sum games used by the BKW result have payoffs of size $O(\frac{1}{\epsilon})$, which ultimately means that their lower bound only applies to the problem ϵ -NE ϵ -SW.

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 Jurziński [27, 28]. In [27], 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 factor $\left(1 - \sqrt{\frac{1-\epsilon'}{1-\epsilon}}\right)^2$ of the best social welfare. They also give further results for the case where $\epsilon > \frac{1}{2}$. In [28] they derived polynomial-time algorithms that compute ϵ -NE for $\epsilon \geq \frac{3-\sqrt{5}}{2}$ that approximate the quality of plutocratic and egalitarian Nash equilibria to various degrees.

2. Preliminaries

Throughout the paper, we use $[n]$ to denote the set of integers $\{1, 2, \dots, n\}$. 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 [29] 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: $\max_{i \in [n]} ((R \cdot \mathbf{y})_i) - \mathbf{x}^T \cdot R \cdot \mathbf{y}$. 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. We say that a pure strategy $j \in [n]$ of the row player is an ϵ -best-response against \mathbf{y} if:

$$\max_{i \in [n]} ((R \cdot \mathbf{y})_i) - (R \cdot \mathbf{y})_j \leq \epsilon.$$

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: $\max_{i \in [n]} ((R \cdot \mathbf{y})_i) - \min_{i \in \text{supp}(\mathbf{x})} ((R \cdot \mathbf{y})_i)$. 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 5, 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 $X \times Y$, and V is a verification function of the form $V : X \times Y \times A \times B \rightarrow \{0, 1\}$.

The game is a co-operative and played between two players, who are called Merlin₁ and Merlin₂, 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 Merlin₁ and y to Merlin₂. Crucially, Merlin₁ does not know the question sent to Merlin₂ and vice versa. Having received x , Merlin₁ then chooses an answer from A and sends it back to Arthur. Merlin₂ 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 Merlin₁ is a function $a : X \rightarrow A$ that gives an answer for every possible question, and likewise a strategy for Merlin₂ is a function $b : Y \rightarrow B$. 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 uniform⁴ distribution \mathcal{U} over $X \times Y$. In particular, this means that there is no correlation between the question sent to Merlin₁ and the question sent to Merlin₂. We are interested in the problem of approximating the value of a free game within an additive error of δ .

⁴More generally, a free game is a two-prover game in which there is no correlation between the questions that are sent to the two players. However, Aaronson et al. [19] point out that this is equivalent to assuming that the distribution is uniform.

FREEGAME $_{\delta}$

Input: A free game \mathcal{T} and a constant $\delta > 0$.

Output: A value p such that $|\omega(\mathcal{T}) - p| \leq \delta$.

3. Quasi-polynomial Upper Bounds

The algorithm of Lipton, Markakis, and Mehta finds an ϵ -NE in $n^{O(\frac{\log n}{\epsilon^2})}$ time [14]. In fact, it turns out that the same technique also provides a quasi-polynomial time approximation scheme for some of the constrained problems that we study in this paper. While this follows in a straightforward way from the techniques used by LMM, this fact does not seem to have been explicitly formulated before in the literature, so we give a short exposition here.

The key idea behind the LMM algorithm is to take a strategy profile (\mathbf{x}, \mathbf{y}) , and to sample a small number of strategies from both \mathbf{x} and \mathbf{y} . Specifically, k independent random samples are taken from \mathbf{x} , and k independent random samples are taken from \mathbf{y} . The strategies \mathbf{x}_s and \mathbf{y}_s are then formed by playing a uniform distribution over the sampled strategies from the respective distributions. These strategies are called *k-uniform*, since they consist of a uniform distribution over k samples.

The key observation of LMM is that when $k \in O(\log(n)/\epsilon^2)$, there is a non-zero probability that both of the following hold for all i .

$$|(R \cdot \mathbf{y})_i - (R \cdot \mathbf{y}_s)_i| \leq \epsilon, \quad |(\mathbf{x}^T \cdot C)_i - (\mathbf{x}_s^T \cdot C)_i| \leq \epsilon.$$

In other words, the payoff of any pure strategy moves by at most ϵ between (\mathbf{x}, \mathbf{y}) and $(\mathbf{x}_s, \mathbf{y}_s)$. So, if one chooses (\mathbf{x}, \mathbf{y}) to be a Nash equilibrium of the game, then this provides a proof via the probabilistic method that there exists a k -uniform approximate Nash equilibrium with k chosen as above. Specifically, this proves the existence of a 2ϵ -NE, since the payoff to each player can decrease by at most ϵ , and the best response of each player can increase by at most ϵ .

The LMM algorithm simply performs a brute-force search over all $n^k = n^{O(\frac{\log n}{\epsilon^2})}$ possible k -uniform profiles in order to find this 2ϵ -NE.

This technique also gives results for constrained approximate Nash equilibria. This is because the sampling technique can be applied to any strategy profile. For example, if we take (\mathbf{x}, \mathbf{y}) to be the ϵ -NE with social welfare s , then with positive probability we will have that $(\mathbf{x}_s, \mathbf{y}_s)$ is a 3ϵ -NE with social welfare at least $s - 2\epsilon$ (since both player's payoffs can decrease by ϵ). This leads to a QPTAS for the following problem: find a 3ϵ -NE that is within 2ϵ of the best social welfare that is achievable by an ϵ -NE. Note that the best social welfare achievable by a 3ϵ -NE may be significantly greater than the best social welfare achievable by an ϵ -NE, and the algorithm approximates the latter rather than the former.

The same technique can be applied to Problems 2, 3, and 5 to obtain similar QPTASs. However, the other problems that we study either ask for certain properties regarding the probabilities themselves, or for large support sizes, neither of which can be addressed through the sub-sampling technique. It is an interesting open question as to whether there exists a QPTAS for these problems.

4. Hardness of approximating free games

The *exponential time hypothesis (ETH)* is the conjecture that any deterministic algorithm for solving 3SAT requires $2^{\Omega(n)}$ time. Aaronson, Impagliazzo, and Moshkovitz have shown that, if ETH holds, then there exists a small constant $\epsilon > 0$ such that approximating the value of a free game within an additive error of ϵ requires quasi-polynomial time. However, their result is not suitable for our purposes, because it produces a free game in which the question and answer sets have the same size, and to prove our result, we will require that the question sets have logarithmic size when compared to the answer sets.

The conference version of this paper [22] solved this issue by using a sub-sampling lemma, also proved by Aaronson, Impagliazzo, and Moshkovitz, which

shows that if we randomly choose logarithmically many questions from the original game, the value of the resulting sub-game is close the value of the original. However, this comes at the cost of needing randomness in the reduction, and so our result depended on the truth of the *randomized* ETH, which is a stronger assumption.

In this exposition, we will instead use a technique of Babichenko, Papadimitriou, and Rubinfeld [23], which allows us to produce a free game with a logarithmic size question set in a deterministic way. The result that we need is a clear consequence of their ideas, but is not explicitly formulated in their paper. For the sake of completeness, in the rest of this section we provide our own exposition of their ideas.

The PCP theorem. The starting point of the result will be a 3SAT instance ϕ . We say that the size of a formula ϕ is the number of variables and clauses in the formula. We define $\text{SAT}(\phi) \in [0, 1]$ to be the maximum fraction of clauses that can be satisfied in ϕ . The first step is to apply a PCP theorem.

Theorem 1 (Dinur’s PCP Theorem [25]). *Given any 3SAT instance ϕ of size n , and a constant ϵ in the range $0 < \epsilon < \frac{1}{8}$, we can produce in polynomial time a 3SAT instance ψ where:*

- *The size of ψ is $n \cdot \text{polylog}(n)$.*
- *Every clause of ψ contains exactly 3 variables and every variable is contained in at most d clauses, where d is a constant.*
- *If $\text{SAT}(\phi) = 1$, then $\text{SAT}(\psi) = 1$.*
- *If $\text{SAT}(\phi) < 1$, then $\text{SAT}(\psi) < 1 - \epsilon$.*

After applying the PCP theorem given above, we then directly construct a free game. Observe that a 3SAT formula can be viewed as a bipartite graph in which the vertices are variables and clauses, and there is an edge between a variable x_i and a clause C_j if and only if x_i appears in C_j . In particular,

the 3SAT formulas produced by Theorem 1 correspond to bipartite graphs with constant degree, since each clause has degree at most 3, and each variable has degree at most d .

The first step is to apply the following lemma, which allows us to partition the vertices of this bipartite graph. The lemma and proof are essentially identical to [23, Lemma 6], although we generalise the formulation slightly, because the original lemma requires that the two sides of the graph have exactly the same number of nodes and that the graph is d -regular.

Lemma 1 ([23]). *Let (V, E) be a bipartite graph with $|V| = n$, where $V = U \cup W$ are the two sides of the graph, and where each node has degree at most d . Suppose that U and W both have a constant fraction of the vertices, and hence $|U| = c_1 \cdot n$ and $|W| = c_2 \cdot n$ for some constants $c_1 < 1$ and $c_2 < 1$. We can efficiently find a partition $S_1, S_2, \dots, S_{\sqrt{n}}$ of U and a partition $T_1, T_2, \dots, T_{\sqrt{n}}$ of W such that each set has size at most $2\sqrt{n}$, and for all i and j we have*

$$|(S_i \times T_j) \cap E| \leq 2d^2.$$

Proof 1. *The algorithm is as follows. First we arbitrarily split U into \sqrt{n} many sets $S_1, S_2, \dots, S_{\sqrt{n}}$, and so each set S_i has size $c_1\sqrt{n} < 2\sqrt{n}$. Then we iteratively construct the partition of W into sets $T_1, T_2, \dots, T_{\sqrt{n}}$ in the following way. We initialize each set T_j to be the empty set. In each iteration, we pick a vertex of $w \in W$ that has not already been assigned to a set. We find a set T_j such that $|T_j| \leq 2 \cdot \sqrt{n}$, and such that for all i we have $|(S_i \times T_j) \cap E| \leq 2 \cdot d^2$. We assign w to T_j and repeat.*

Obviously, for the algorithm to be correct, we must prove that for each vertex w that is considered, there does exist a set T_j that satisfies the required constraints. For this, we rely on the following two properties.

- *The average number of vertices in a set T_j is at most $c_2\sqrt{n} < \sqrt{n}$, and so by Markov's inequality strictly less than half the sets can have size more than $2\sqrt{n}$, and so we lose strictly less than half the sets T_j to the size constraint.*

- Since each vertex has degree at most d , the graph has at most dn edges, and so the average number of edges between each pair of sets S_i and T_j is $dn/(\sqrt{n} \cdot \sqrt{n}) = d$. Again, using Markov's inequality we can conclude that there are at most $1/2d$ pairs of sets S_i and T_j that have more than $2d^2$ edges between them. Hence, even in the worst case, we can lose at most $1/2d$ sets T_j to the edge constraints.

So, we lose strictly less than half of the sets to the size constraints, and $1/2d \leq 1/2$ of the sets to the edge constraints. Hence, by the union bound, we have shown that there is at least one set T_j that satisfies both constraints simultaneously. \square

A free game. Note that Lemma 1 can be applied to the 3SAT formula that arises from Dinur's PCP theorem, because the number of variables and number of constraints are both a constant fraction of the number of nodes in the associated bipartite graph, and because each vertex has either degree d or degree 3. We use this to construct the following free game, which is highly reminiscent of the clause variable game given by Aaronson, Impagliazzo, and Moshkovitz [19].

Definition 2. Given a 3SAT formula ϕ of size n , we define a free game \mathcal{F}_ϕ in the following way.

1. Arthur begins by applying Dinur's PCP theorem to ϕ to obtain a formula ψ of size $N = n \text{ polylog}(n)$, and then uses Lemma 1 to split the variables of ψ into sets $S_1, S_2, \dots, S_{\sqrt{N}}$ and the clauses of ψ into sets $T_1, T_2, \dots, T_{\sqrt{N}}$.
2. Arthur picks an index i uniformly at random from $[\sqrt{N}]$, and independently an index j uniformly at random from $[\sqrt{N}]$. He sends S_i to Merlin₁ and T_j to Merlin₂.
3. Merlin₁ responds by giving a truth assignment to every variable in S_i , and Merlin₂ responds by giving a truth assignment to every variable that is involved with a clause in T_j .
4. Arthur awards the Merlins payoff 1 if and only if both of the following conditions hold.

- Merlin₂ returns an assignment that satisfies all clauses in T_j .
- For every variable v that appears in S_i and some clause of T_j , the assignment to v given by Merlin₁ agrees with the assignment to v given by Merlin₂. Note that this condition is always satisfied when S_i and T_j share no variables.

Arthur awards payoff 0 otherwise.

If n is the size of ϕ , then when we write \mathcal{F}_ϕ down as a free game $(X, Y, A, B, \mathcal{D}, V)$, the number of questions in the sets X and Y is $\sqrt{n \text{ polylog}(n)}$, and the number of answers in A and B is $2^{2\sqrt{n \text{ polylog}(n)}}$, where the extra $\text{polylog}(n)$ factor arises due to the application of the PCP theorem.

The following lemma shows that if ϕ is unsatisfiable, then the value of this free game is bounded away from 1. Again, the ideas used to prove this lemma are clearly evident in the work of Babichenko, Papadimitriou, and Rubinfeld [23], but this lemma was not explicitly formulated in their paper.

Lemma 2 ([23]). *If ϕ is satisfiable then $\omega(\mathcal{F}_\phi) = 1$. If ϕ is unsatisfiable then $\omega(\mathcal{F}_\phi) \leq 1 - \epsilon/2d$.*

Proof 2. *The case where $\text{SAT}(\phi) = 1$ is straightforward. Since there exists a satisfying assignment for ϕ , there also exists a satisfying assignment for ψ . If the two Merlins play according to this satisfying assignment, then they obviously achieve an expected payoff of 1.*

For the other claim, first observe that we can assume that both Merlins play deterministic strategies, since the game is co-operative, and therefore nothing can be gained through randomization. So, let s_1 be a strategy for Merlin₁. Observe that since $S_1, S_2, \dots, S_{\sqrt{N}}$ partition the variables of ψ , we have that s_1 yields an assignment to the variables of ψ .

Let us fix an arbitrary deterministic strategy s_1 for Merlin₁. We have that the payoff to Merlin₂ for an individual question T_j can be computed as follows:

- For every set S_i for which there are no edges between the variables in S_i and T_j , Merlin₂ gets payoff 1 “for free.”

- Otherwise, Merlin₂ gets payoff 1 only if the assignments to the clauses agree with the assignment implied by s_1 .

From this, we can see that when Merlin₁ plays s_1 , Merlin₂ can maximize his payoff by playing the strategy that agrees everywhere with the assignment chosen by Merlin₁. So let s_2 denote this strategy.

Since ϕ is unsatisfiable, the PCP theorem tells us that $\text{SAT}(\psi) < 1 - \epsilon$. Thus, there are at least ϵdN clauses that are not satisfied when s_1 is played against s_2 . Since Lemma 1 ensures that the maximum number of edges between two sets is $2d^2$, there must therefore be at least $\epsilon dN/2d^2 = \epsilon N/2d$ pairs of sets that give payoff 0 to the Merlins under s_1 and s_2 . Since there are exactly N pairs of sets in total, this means that the expected payoff to the Merlins is bounded by $1 - \epsilon/2d$. \square

Finally, we can formulate the lower bound that we will use in this paper. The proof is the same as the one given in [19], but we use the free game \mathcal{F}_ϕ , rather than the construction originally given in that paper.

Theorem 2 ([23]). *Assuming ETH, there is a small constant δ below which the problem FREEGAME_δ cannot be solved faster than $M^{\tilde{\Omega}(\log M)}$, where M denotes the size of the free game, even when the question sets have size $\log M$.*

Proof 3. *Lemma 2 implies that if we can approximate the value of \mathcal{F}_ϕ with an additive error of less than $\epsilon/2d$, then we can solve the satisfiability problem for ϕ .*

Assume, for the sake of contradiction, that there exists an algorithm that can solve the FREEGAME_δ problem in time $M^{o\left(\frac{\log M}{(\log \log M)^\epsilon}\right)}$ for some constant c that will be fixed later. Observe that the free game \mathcal{F}_ϕ has size $M = O(2\sqrt{n \text{polylog}(n)})$, and so the hypothesized algorithm would run in time:

$$\exp\left(o\left(\frac{\log^2 M}{(\log \log(M))^c}\right)\right) = \exp\left(o\left(\frac{n \text{polylog}(n)}{(\log(\sqrt{n} \text{polylog}(n)))^c}\right)\right).$$

If we set c to be greater than the degree of the polynomial in the $\text{polylog}(n)$ from the numerator, then we can conclude that the running time would be $2^{o(n)}$, which would violate the ETH. \square

5. Hardness of approximating social welfare

Overview. In this section, we study the following *social welfare* problem for a bimatrix game $\mathcal{G} = (R, C)$. The *social welfare* of a strategy profile (\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 bimatrix 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 (from now on we will drop the subscript ϕ , since the exact construction of \mathcal{F} is not relevant to us.) 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 $\text{BSW}(\mathcal{G}, \epsilon) = 2$.
- **(Soundness)** If $\omega(\mathcal{F}) < 1 - \delta$, then the unscaled $\text{BSW}(\mathcal{G}, \epsilon) < 2(1 - g \cdot \delta)$.

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

5.1. *The construction*

We use \mathcal{F} 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.

II			
I	/		
		C	D_2
		R	$-D_2$
		$-D_1$	0
		D_1	0

- The game (R, C) is built from \mathcal{F} 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 Althöfer [21]. Let H be the set of all functions of the form $f : Y \rightarrow \{0, 1\}$ such that $f(y) = 1$ for exactly half⁵ 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 .

⁵If $|Y|$ 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.

- The game $(0, 0)$ is a game in which both players have zero matrices.

Observe that the size of (R, C) is the same as the size of \mathcal{F} . The game $(D_1, -D_1)$ has the same number of columns as C , and the number of rows is at most $2^{|Y|} \leq 2^{O(\log |\mathcal{F}|)} = |\mathcal{F}|^{O(1)}$, where we are crucially using the fact that Theorem 2 allows us to assume that the size of Y is $O(\log |\mathcal{F}|)$. By the same reasoning, the number of columns in $(-D_2, D_2)$ is at most $|\mathcal{F}|^{O(1)}$. Thus, the size of \mathcal{G} is $|\mathcal{F}|^{O(1)}$, and so this reduction is polynomial size.

5.2. Completeness

To prove completeness, it suffices to show that, if $\omega(\mathcal{F}) = 1$, then there exists a $(1 - 4g \cdot \delta)$ -UNE of \mathcal{G} that has social welfare 2. To do this, assume that $\omega(\mathcal{F}) = 1$, and take a pair of optimal strategies (s_1, s_2) for \mathcal{F} 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 harder 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 3. *If $\omega(\mathcal{F}) = 1$, then there exists a $(1 - 4g \cdot \delta)$ -UNE (\mathbf{x}, \mathbf{y}) of \mathcal{G} with $\text{SW}(\mathbf{x}, \mathbf{y}) = 2$.*

Proof 4. *Let $(s_1, s_2) \in S_1 \times S_2$ be a pair of optimal strategies for Merlin₁ and Merlin₂ in \mathcal{F} . For each $(x, a) \in X \times A$ and each $(y, b) \in Y \times B$, we define*

$$\mathbf{x}(x, a) = \begin{cases} \frac{1}{|X|} & \text{if } s_1(x) = a. \\ 0 & \text{otherwise.} \end{cases} \quad \mathbf{y}(y, b) = \begin{cases} \frac{1}{|Y|} & \text{if } s_2(y) = b. \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, by construction, we have that the payoff to the row player under (\mathbf{x}, \mathbf{y}) is equal to $p(\mathcal{F}, s_1, s_2) = 1$, and therefore (\mathbf{x}, \mathbf{y}) has social welfare 2.

On the other hand, we must prove that (\mathbf{x}, \mathbf{y}) is a $(1 - 4g \cdot \delta)$ -UNE. To do so, we will show that neither player has a deviation that increases their payoff

by more than $(1 - 4g \cdot \delta)$. We will show this for the row player; the proof for the column player is symmetric. There are two types of row to consider.

- First suppose that r is a row in the sub-game (R, C) . We claim that the payoff of r is at most 1. This is because the maximum payoff in R is 1, while the maximum payoff in $-D_2$ is 0. Since the row player already obtains payoff 1 in (\mathbf{x}, \mathbf{y}) , row r cannot be a profitable deviation.
- Next suppose that r is a row in the sub-game $(D_1, -D_1)$. Since we have $\sum_{b \in B} \mathbf{y}(y, b) = \frac{1}{|Y|}$ for every question y , we have that all rows in D_1 have the same payoff. This payoff is

$$\frac{1}{2} \cdot \left(\frac{4}{1 + 4 \cdot \delta} \right) = \frac{2}{1 + 4g \cdot \delta} = 2 - \frac{8g \cdot \delta}{1 + 4g \cdot \delta}.$$

Since $\delta \leq 1$ and $g \leq \frac{1}{4}$ we have

$$\frac{8}{1 + 4g \cdot \delta} \geq \frac{8}{1 + 4g} \geq 4.$$

Thus, we have shown that the payoff of r is at most $2 - 4g \cdot \delta$. Thus the row player's regret is at most $1 - 4g \cdot \delta$. \square

5.3. Soundness

We now suppose that $\omega(\mathcal{F}) < 1 - \delta/2$, 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 the subgame (R, C) .

Lemma 4. *If $\text{SW}(\mathbf{x}, \mathbf{y}) > 2 - 2g \cdot \delta$, 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) .

Proof 5. *We will prove the lemma for \mathbf{x} ; the proof for \mathbf{y} is entirely symmetric. For the sake of contradiction, suppose that \mathbf{x} places strictly less than $(1 - g \cdot \delta)$*

probability on rows in (R, C) . Observe that every subgame of \mathcal{G} other than (R, C) is a zero-sum game. Thus, any probability assigned to these sub-games contributes nothing to the social welfare. On the other hand, the payoffs in (R, C) are at most 1. So, even if the column player places all probability on columns in C , the social welfare $\text{SW}(\mathbf{x}, \mathbf{y})$ will be strictly less than $2 \cdot (1 - g \cdot \delta) + g \cdot \delta \cdot 0 = 2 - 2g \cdot \delta$, a contradiction. \square

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 4, 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} , 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 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 $\text{SW}(\mathbf{x}, \mathbf{y})$. 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 Merlin₁ 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 Merlin₂ is defined symmetrically.

We will use $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$ as an intermediary between \mathcal{G} and \mathcal{F} 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} . Since we have a bound on the payoff of any pair of strategies in \mathcal{F} , 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 5. *We have $|p_r(\mathcal{G}, \mathbf{x}, \mathbf{y}) - p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})| \leq 4g \cdot \delta$.*

Proof 6. *By construction, $p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})$ is equal to the payoff that the row player obtains from the subgame (R, C) , and so we have $p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}}) \leq p_r(\mathcal{G}, \mathbf{x}, \mathbf{y})$. On the other hand, since the row player places at most $g \cdot \delta$ probability on rows not in (R, C) , and since these rows have payoff at most $\frac{4}{1+4g \cdot \delta} < 4$, we have $p_r(\mathcal{G}, \mathbf{x}, \mathbf{y}) \leq p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}}) + 4g \cdot \delta$. \square*

Relating $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$ to \mathcal{F} . 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 6. *We have $\|\mathbf{u}_Y - \mathbf{y}'\|_1 < 16g \cdot \delta$ and $\|\mathbf{u}_X - \mathbf{x}'\|_1 < 16g \cdot \delta$.*

We begin by proving an auxiliary lemma. This is similar to Lemma 3 in [30], which was published as [31] without that lemma.

Lemma 7. *If $\|\mathbf{u}_Y - \mathbf{y}'\|_1 \geq c$ then there exists a set $B \subseteq Y$ of size $|Y|/2$ such that*

$$\sum_{i \in B} \mathbf{y}'_i > \frac{1}{2} + \frac{c}{4} - 2g \cdot \delta.$$

Proof 7. *We first define $\mathbf{d} = \mathbf{y}' - \mathbf{u}_Y$, and then we partition Y as follows*

$$U = \{y \in Y : \mathbf{d}_y > \frac{1}{|Y|}\},$$

$$L = \{y \in Y : \mathbf{d}_y \leq \frac{1}{|Y|}\}.$$

Since $\|\mathbf{y}'\|_1 \geq 1 - g \cdot \delta$ and $\|\mathbf{u}\|_1 = 1$, we have that

$$\sum_{y \in U} \mathbf{d}_y \geq c/2 - g \cdot \delta,$$

$$\sum_{y \in L} \mathbf{d}_y \leq -c/2 + g \cdot \delta.$$

We will prove that there exists a set $B \subseteq Y$ of size $|Y|/2$ such that $\sum_{y \in B} \mathbf{d}_y \geq c/4 - g \cdot \delta$.

We have two cases to consider, depending on the size of U .

- *First suppose that $|U| > |Y|/2$. If this is the case, then there must exist a set $B \subseteq U$ with $|B| = |U|/2$ and $\sum_{i \in B} \mathbf{d}_i \geq c/4 - g \cdot \delta$. We can then add arbitrary columns from $U \setminus B$ to B in order to make $|B| = |Y|/2$, and since $\mathbf{d}_i > 0$ for all $i \in U$, this cannot decrease $\sum_{i \in B} \mathbf{d}_i$. Thus, we have completed the proof for this case.*

- Now suppose that $|U| \leq |Y|/2$. If this is the case, then there must exist a set $C \subseteq L$ with $|C| = |L|/2$ and $\sum_{i \in C} \mathbf{d}_i \geq -\frac{c}{4} + g \cdot \delta$. So, let $C' \subseteq C$ be an arbitrarily chosen subset such that $|C'| + |U| = |Y|/2$. This is possible since $|L| = |Y| - |U|$ and hence $|L|/2 = |Y|/2 - |U|/2$, which implies that $|L|/2 + |U| > |Y|/2$. Setting $B = C' \cup U$ therefore gives us a set with $|B| = |Y|/2$ such that

$$\begin{aligned} \sum_{i \in B} \mathbf{d}_i &\geq (c/2 - g \cdot \delta) - (c/4 + g \cdot \delta) \\ &= c/4 - 2 \cdot g \cdot \delta . \end{aligned}$$

So we have completed the proof of this case, and the lemma as a whole. \square

We can now proceed with the proof of Lemma 6.

Proof 8 (Proof of Lemma 6). Suppose, for the sake of contradiction that one of these two properties fails. Without loss of generality, let us assume that $\|\mathbf{u}_Y - \mathbf{y}'\|_1 \geq c$. We will show that the row player can gain more than 1 in payoff by deviating to a new strategy, which will show that (\mathbf{x}, \mathbf{y}) is not a 1-UNE, contradicting our assumption that it is a $(1 - 4g \cdot \delta)$ -UNE.

By assumption, \mathbf{x} places at least $1 - g \cdot \delta$ probability on rows in (R, C) . The maximum payoff in R is 1, and the maximum payoff in $-D_2$ is 0. On the one hand, the rows in D_2 give payoff at most $8/(2 + g \cdot \delta) \leq 4$. So the row player's payoff under (\mathbf{x}, \mathbf{y}) is bounded by

$$(1 - g \cdot \delta) \cdot 1 + (g \cdot \delta) \cdot 4 = 1 + 3g \cdot \delta.$$

On the other hand, we can apply Lemma 7 with $c = 16g \cdot \delta$ to find a set $B \subseteq Y$ such that

$$\begin{aligned} \sum_{i \in B} \mathbf{y}'(i) &> \frac{1}{2} + \frac{16g \cdot \delta}{4} - 2g \cdot \delta \\ &= \frac{1}{2} + 2g \cdot \delta \\ &= \frac{1 + 4g \cdot \delta}{2}. \end{aligned}$$

So, let r_B be the row of D_1 that corresponds to B . This row has payoff $\frac{8}{2+g\cdot\delta}$ for every entry in B . So, the payoff of row r_B must be at least

$$\left(\frac{1+4g\cdot\delta}{2}\right) \cdot \left(\frac{4}{1+4g\cdot\delta}\right) = 2.$$

Thus, the row player can deviate to r_B and increase his payoff by at least $1-3g\cdot\delta$, and (\mathbf{x}, \mathbf{y}) is not a $(1-4g\cdot\delta)$ -UNE. \square

With Lemma 6 at hand, we can now prove that the difference between $p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})$ and $p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}})$ 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} 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 $|p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}}) - p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}})| \leq 64g \cdot \delta$.*

Proof 9. *The distribution used in \mathcal{F} is the product of \mathbf{u}_Y and \mathbf{u}_X , while the distribution used in $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$ is the product of \mathbf{y}' and \mathbf{x}' . Furthermore, Lemma 6 tells us that $\|\mathbf{u}_Y - \mathbf{y}'\|_1 < 16g \cdot \delta$ and $\|\mathbf{u}_X - \mathbf{x}'\|_1 < 16g \cdot \delta$. Our approach is to transform \mathbf{u}_X to \mathbf{x}' while bounding the amount that $p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}})$ changes. Once we have this, we can apply the same transformation to \mathbf{u}_Y and \mathbf{y}' .*

Consider the effect of shifting probability from a question $x_1 \in X$ to a different question $x_2 \in X$. Since all entries of V are in $\{0, 1\}$, if we shift q probability from x_1 to x_2 , then $p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}})$ can change by at most $2q$. This bound also holds if we remove probability from x_1 without adding it to x_2 (which we might do since $\|\mathbf{x}\|_1$ may not be 1.) Thus, if we shift probability to transform \mathbf{u}_X into \mathbf{x}' , then we can change $p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}})$ by at most $32g \cdot \delta$.

The same reasoning holds for transforming \mathbf{u}_Y into \mathbf{y}' . This means that we can transform \mathcal{F} to $\mathcal{T}_{(\mathbf{x}, \mathbf{y})}$ while changing the payoff of $(s_{\mathbf{x}}, s_{\mathbf{y}})$ by at most $64g \cdot \delta$, which completes the proof. \square

Completing the soundness proof. The following lemma uses the bounds derived in Lemmas 5 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} .*

Proof 10. *Lemmas 5 and 8 tell us that*

$$\begin{aligned} |p_r(\mathcal{G}, \mathbf{x}, \mathbf{y}) - p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}})| &\leq 4g \cdot \delta, \\ |p(\mathcal{T}_{(\mathbf{x}, \mathbf{y})}, s_{\mathbf{x}}, s_{\mathbf{y}}) - p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}})| &\leq 64g \cdot \delta. \end{aligned}$$

Hence, we have $|p_r(\mathcal{G}, \mathbf{x}, \mathbf{y}) - p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}})| \leq 68g \cdot \delta$. However, we know that $p(\mathcal{F}, s_{\mathbf{x}}, s_{\mathbf{y}}) \leq 1 - \delta/2$. So, if we set $g = \frac{1}{138}$, then we will have that

$$\begin{aligned} p_r(\mathcal{G}, \mathbf{x}, \mathbf{y}) &\leq 1 - \frac{1}{2} \cdot \delta + \frac{68}{138} \cdot \delta \\ &= 1 - \frac{1}{138} \cdot \delta \\ &= 1 - g \cdot \delta. \end{aligned}$$

□

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

5.4. 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 $\frac{4}{1+4g\delta} \leq 4$, and the minimum payoff is $-\frac{4}{1+4g\delta} \geq -4$. 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 3. *If ETH holds, then there exists a constant δ below which the problem $(\frac{1-4g\delta}{8})$ -NE $(\frac{g}{4} \cdot \delta)$ -SW, where $g = \frac{1}{138}$, requires $n^{\tilde{\Omega}(\log n)}$ time.*

Proof 11. *By Lemmas 3 and 9, we have*

- *if $\omega(\mathcal{F}) = 1$ then there exists a $(1 - 4g \cdot \delta)$ -UNE of \mathcal{G} with social welfare $1 + 1 = 2$. In the rescaled game this translates to a $(\frac{1-4g\delta}{8})$ -NE of \mathcal{G}_s with social welfare $\frac{1+4}{8} + \frac{1+4}{8} = \frac{10}{8}$.*

- if $\omega(\mathcal{F}) < 1 - \delta$ then all $(1 - 4g \cdot \delta)$ -UNE of \mathcal{G} have social welfare at most $(1 - g \cdot \delta) + (1 - g \cdot \delta) = 2 - 2g \cdot \delta$. After rescaling, we have that all $(\frac{1-4g \cdot \delta}{8})$ -NE of \mathcal{G}_s have social welfare at most

$$\frac{5 - g \cdot \delta}{8} + \frac{5 - g \cdot \delta}{8} = \frac{10}{8} - \frac{g \cdot \delta}{4}$$

By Theorem 2, assuming ETH we require $|\mathcal{F}|^{\tilde{O}(\log |\mathcal{F}|)}$ time to decide whether the value of \mathcal{F} is 1 or $1 - \delta$ for some small constant δ . Thus, we also require $n^{\tilde{O}(\log |n|)}$ to solve the problem $(\frac{1-4g \cdot \delta}{8})$ -NE $(\frac{g}{4} \cdot \delta)$ -SW. \square

As we observed in the introduction, when δ is small this theorem gives a lower bound for all ϵ -NE where $\epsilon \approx 1/8$. While this is reasonably far from zero, we note that it is unlikely to be best possible, since currently we only have polynomial time algorithms for finding a 0.3393-NE in bimatrix games [13]. The $1/8$ in our result is ultimately derived from the fact that our payoffs lie in the range $(-4, 4)$, and this was necessary in order to force the two players to play uniformly over the question sets. A better simulation of the free game would be needed to improve upon this. It should also be noted that our result only holds when δ , the approximation quality for social welfare, is a very small constant. An interesting question is whether one can show lower bounds against ϵ -NE δ -SW where both ϵ and δ are far from zero.

6. Hardness results for other decision problems

In this section we study a range of decision problems associated with approximate equilibria. Table 1 shows all of the decision problems that we consider. Most are known to be NP-complete for the case of exact Nash equilibria [2, 3]. For each problem in Table 1, 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 solution concept than ϵ -WSNE, i.e., every ϵ -WSNE is an ϵ -NE, 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 deciding if there is an

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 $\epsilon < \frac{1}{8}$, so let us fix $\epsilon^* < \frac{1}{8}$ for the rest of this section. Using Theorem 3, we compute from ϵ^* the parameters n and δ that we require to apply Theorem 3. In particular, set δ^* to solve $\epsilon^* = (\frac{1-4g\cdot\delta^*}{8})$, and choose n^* as $\frac{1}{\delta^*}$. Then, for $n > n^*$ and $\delta = \delta^*$ we can apply Theorem 3 to bound the social welfare achievable if $\omega(\mathcal{F}) < 1 - \delta^*$ as

$$u := \frac{10}{8} - \frac{1}{522}\delta^*.$$

Theorem 3 implies that in order to decide whether the game \mathcal{G}_s possesses an ϵ^* -NE that yields social welfare strictly greater than u requires $n^{\tilde{O}(\log n)}$ time, where δ no longer appears in the exponent since we have fixed it as the constant δ^* .

Problem 1 asks to decide whether a bimatrix game possesses an ϵ^* -NE where the expected payoff for each player is at least u , where u is an input to the problem. When we set $u = \frac{5}{8}$, the conditional hardness of this problem is an immediate corollary of Theorem 3.

For Problems 2 - 9, we use \mathcal{G}_s to construct a new game \mathcal{G}' , which adds one row i and one column j to \mathcal{G}_s . The payoffs are defined using the constants u and ϵ^* , as shown in Figure 1.

$$\mathcal{G}' =$$

			j	
\mathcal{G}_s			$0, \frac{5}{8} + \epsilon^*$	
			\vdots	
			$0, \frac{5}{8} + \epsilon^*$	
i	$\frac{5}{8} + \epsilon^*, 0$	\dots	$\frac{5}{8} + \epsilon^*, 0$	$1, 1$

Figure 1: The game \mathcal{G}' .

In \mathcal{G}' , the expected payoff for the row player for i is at least $\frac{5}{8} + \epsilon^*$ irrespective of the column player's strategy. Similarly, the expected payoff for j is at least $\frac{5}{8} + \epsilon^*$ irrespective of the row player's strategy. This means that:

- If \mathcal{G}_s possesses an ϵ^* -NE with social welfare $\frac{10}{8}$, then \mathcal{G}' possesses at least one ϵ^* -NE where the players do not play the pure strategies i and j .
- If every ϵ^* -NE of \mathcal{G}_s yields social welfare at most u , then in *every* ϵ^* -NE of \mathcal{G}' , the players place almost all of their probability on i and j respectively. Note that (i, j) is a pure exact Nash equilibrium.

Problem 2 asks whether a bimatrix game possesses an ϵ -NE where the row player plays with positive probability only strategies in a given set S . Let S_R (S_C) denote the set of pure strategies available to the row (column) player from the subgame (R, C) of \mathcal{G}_s . To show the hardness of Problem 2, we will set $S = S_R$.

Recall that \mathcal{G}_s is created from \mathcal{F} . First, we prove in Lemma 10 that if $\omega(\mathcal{F}) = 1$, then \mathcal{G}' possesses an ϵ^* -NE such that the answer to Problem 2 is “Yes”. Note that we actually argue in Lemma 10 about the existence of an ϵ^* -WSNE, since this stronger claim will be useful when we come to deal with Problems 7 - 9.

Lemma 10. *If $\omega(\mathcal{F}) = 1$, then \mathcal{G}' possesses an ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) such that $\text{supp}(\mathbf{x}) \subseteq S_R$. Under (\mathbf{x}, \mathbf{y}) , both players get payoff $\frac{5}{8}$, so $\text{SW}(\mathbf{x}, \mathbf{y}) = \frac{10}{8}$. Moreover, $|\text{supp}(\mathbf{x})| = |X|$ and $\max_i x_i \leq \frac{1}{|X|}$, where X is the question set of Merlin₁ in \mathcal{F} .*

Proof 12. *The proof of Lemma 3 shows that, if $\omega(\mathcal{F}) = 1$, then \mathcal{G} possesses an ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) where the expected payoff for each player is 1 and $\text{supp}(\mathbf{x}) \subseteq S_R$. The reason that (\mathbf{x}, \mathbf{y}) is well supported is that all rows in $\text{supp}(\mathbf{x})$ have equal expected payoff. Moreover, \mathbf{x} is a uniform mixture over a pure strategy set of size $|X|$, where X is the question set of Merlin₁ in \mathcal{F} . Since \mathcal{G}_s is obtained from \mathcal{G} by adding 4 to the payoffs and dividing by 8, (\mathbf{x}, \mathbf{y}) is as an ϵ^* -NE in \mathcal{G}_s where each player has payoff $\frac{5}{8}$. To complete the proof we show that (\mathbf{x}, \mathbf{y}) in an*

ϵ^* -NE for \mathcal{G}' , which is the same as \mathcal{G}_s apart from the additional pure strategies i and j . Since i and j yield payoff $\frac{5}{8} + \epsilon^*$, but not more, the claim holds. \square

Next we prove that if $\omega(\mathcal{F}) < 1 - \delta^*$, then the answer to Problem 2 is “No”.

Lemma 11. *If $\omega(\mathcal{F}) < 1 - \delta^*$, then in every ϵ^* -NE (\mathbf{x}, \mathbf{y}) of \mathcal{G}' it holds that $x_i > 1 - \frac{\epsilon^*}{1 - \epsilon^*}$ and $y_j > 1 - \frac{\epsilon^*}{1 - \epsilon^*}$.*

Proof 13. Let $\mathcal{G}_s := (P, Q)$ and suppose that (\mathbf{x}, \mathbf{y}) is an ϵ^* -NE of \mathcal{G}' . From Theorem 3 we know that if $\omega(\mathcal{F}) < 1 - \delta^*$, then in any ϵ^* -NE of \mathcal{G}_s we have that each player gets payoff at most $\frac{u}{2} < \frac{5}{8}$. Under (\mathbf{x}, \mathbf{y}) in \mathcal{G}' the row player gets payoff

$$\begin{aligned} \mathbf{x}^T P \mathbf{y} &< (1 - x_i) \cdot (1 - y_j) \cdot \frac{5}{8} + x_i \cdot (1 - y_j) \left(\frac{5}{8} + \epsilon^* \right) + x_i \cdot y_j \\ &= x_i \cdot ((1 - y_j) \cdot \epsilon^* + y_j) + (1 - y_j) \cdot \frac{5}{8}. \end{aligned}$$

From the pure strategy i , the row player gets

$$P_i \cdot \mathbf{y} = (1 - y_j) \left(\frac{5}{8} + \epsilon^* \right) + y_j.$$

In order for (\mathbf{x}, \mathbf{y}) to be an ϵ^* -NE it must hold that $\mathbf{x}^T P \mathbf{y} \geq P_i \mathbf{y} - \epsilon^*$. Using the upper bound on $\mathbf{x}^T P \mathbf{y}$ that we just derived, we get:

$$x_i > 1 - \frac{\epsilon^*}{(1 - y_j) \cdot \epsilon^* + y_j}. \quad (1)$$

By symmetry, we also have that the column player must play j with probability:

$$y_j > 1 - \frac{\epsilon^*}{(1 - x_i) \cdot \epsilon^* + x_i}. \quad (2)$$

Recall that in this section ϵ^* is a constant. Observe that the right-hand side of (2) is increasing in x_i , and we can thus use it to replace x_i in (2) as follows:

$$\begin{aligned} y_j &> 1 - \frac{\epsilon^*}{(1 - 1 + \frac{\epsilon^*}{(1 - y_j) \epsilon^* + y_j}) \epsilon^* + 1 - \frac{\epsilon^*}{(1 - y_j) \epsilon^* + y_j}} \\ &= 1 - \frac{\epsilon^*}{\frac{\epsilon^{*2}}{(1 - y_j) \epsilon^* + y_j} + 1 - \frac{\epsilon^*}{(1 - y_j) \epsilon^* + y_j}} \\ &= 1 - \frac{(1 - y_j) \epsilon^{*2} + y_j \epsilon^*}{\epsilon^{*2} + (1 - y_j) \epsilon^* + y_j - \epsilon^*}. \end{aligned}$$

Noting that $(\epsilon^{*2} + (1 - \mathbf{y}_j)\epsilon^* + \mathbf{y}_j - \epsilon^*) \geq 0$, by rearranging we get that

$$\mathbf{y}_j^2(1 - \epsilon^*) + \mathbf{y}_j(2\epsilon^* - 1) > 0.$$

Then, since $\epsilon^* < \frac{1}{8}$, we have $1 - \epsilon^* > 0$, and we get that

$$\mathbf{y}_j > \frac{1 - 2\epsilon^*}{1 - \epsilon^*} = 1 - \frac{\epsilon^*}{1 - \epsilon^*}.$$

By symmetry, we have $\mathbf{x}_i > 1 - \frac{\epsilon^*}{1 - \epsilon^*}$, which completes the proof. \square

Next we recall Problems 3 and 4 and we show that, as for Problem 2, Lemmas 10 and 11 can also be used to immediately show that there are instances of these decision problems where the answer is “Yes” if and only if $\omega(\mathcal{F}) = 1$.

Given two probability distributions \mathbf{x} and \mathbf{x}' , the *Total Variation (TV)* distance between them is $\max_i \{|\mathbf{x}_i - \mathbf{x}'_i|\}$. We define the TV distance between two strategy profiles (\mathbf{x}, \mathbf{y}) and $(\mathbf{x}', \mathbf{y}')$ to be the maximum over the TV distance of \mathbf{x} and \mathbf{x}' and the TV distance of \mathbf{y} and \mathbf{y}' . Problem 3 asks whether a bimatrix game possesses two ϵ -NEs with TV distance at least d . In order to apply Lemmas 10 and 11, we will set $d = 1 - \frac{\epsilon^*}{1 - \epsilon^*}$. Then an instance \mathcal{G}' of Problem 3 is “Yes” when $\omega(\mathcal{F}) = 1$ since the ϵ^* -NE (\mathbf{x}, \mathbf{y}) identified in Lemma 10, has TV distance one from the pure exact Nash equilibrium (i, j) . Lemma 11 says that, if $\omega(\mathcal{F}) < 1 - \delta^*$, every ϵ^* -NE (\mathbf{x}, \mathbf{y}) of \mathcal{G}' has $\mathbf{x}_i > 1 - \frac{\epsilon^*}{1 - \epsilon^*}$ and so all ϵ^* -NE are within TV distance $1 - \frac{\epsilon^*}{1 - \epsilon^*}$ of each other.

Problem 4 asks to decide whether there exists an ϵ -NE where the row player does not play any pure strategy with probability more than p . For this problem, we set $p = \frac{1}{|X|}$, where X is the question set for Merlin₁. According Lemma 10, if $\omega(\mathcal{F}) = 1$, then an instance \mathcal{G}' of Problem 4 is a “Yes”. Lemma 11 says that, if $\omega(\mathcal{F}) < 1 - \delta^*$, then for all ϵ^* -NE (\mathbf{x}, \mathbf{y}) of \mathcal{G}' , $\max_i \mathbf{x}_i \geq \mathbf{x}_i > 1 - \frac{\epsilon^*}{1 - \epsilon^*} > \frac{1}{|X|}$.

Problem 5 asks whether a bimatrix game possesses an ϵ -NE with social welfare at most v , and Problem 6 asks whether a bimatrix game possesses an ϵ -NE where the expected payoff of the row player is at most u . We fix $v = \frac{10}{8}$ for Problem 5, and for Problem 6 we fix $u = \frac{5}{8}$. As we have already explained in the proof of Lemma 10, if $\omega(\mathcal{F}) = 1$, then there is an ϵ^* -NE for \mathcal{G}' such that

the expected payoff for each player is $\frac{5}{8}$ and thus the social welfare is $\frac{10}{8}$. So, if $\omega(\mathcal{F}) = 1$, then the answer to Problems 5 and 6 is “Yes”. On the other hand, from the proof of Lemma 11 we know that if $\omega(\mathcal{F}) < 1 - \delta^*$, then in any ϵ^* -NE of \mathcal{G}' both players play the strategies \mathbf{i} and \mathbf{j} with probability at least $1 - \frac{\epsilon^*}{1 - \epsilon^*}$. So, each player gets payoff at least $(1 - \frac{\epsilon^*}{1 - \epsilon^*})^2 > \frac{5}{8}$, since $\epsilon^* < \frac{1}{8}$, from their pure strategies \mathbf{i} and \mathbf{j} . So, if $\omega(\mathcal{F}) < 1 - \delta^*$, then the answer to Problems 5 and 6 is “No”.

Problems 7 - 9 relate to deciding if there exist approximate *well-supported* equilibria with large supports (for ϵ -NE these problems would be trivial). Problem 7 asks whether a bimatrix game possesses an ϵ -WSNE (\mathbf{x}, \mathbf{y}) with $|\text{supp}(\mathbf{x})| + |\text{supp}(\mathbf{y})| \geq 2k$. Problem 8 asks whether a bimatrix game possesses an ϵ -WSNE (\mathbf{x}, \mathbf{y}) with $\min\{|\text{supp}(\mathbf{x})|, |\text{supp}(\mathbf{y})|\} \geq k$. Problem 9 asks whether a bimatrix game possesses an ϵ -WSNE (\mathbf{x}, \mathbf{y}) with $|\text{supp}(\mathbf{x})| \geq k$. Recall that X and Y are the question sets of Merlin₁ and Merlin₂ respectively that were used to define \mathcal{F} and in turn \mathcal{G}_s . We will fix $k = |X| = |Y|$ for all three problems.

If $\omega(\mathcal{F}) = 1$, then Lemma 10 says that there exists an ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) for \mathcal{G}' such that $|\text{supp}(\mathbf{x})| = |\text{supp}(\mathbf{y})| = k$ and thus the answer to Problems 7 - 9 is “Yes”. On the other hand, if $\omega(\mathcal{F}) < 1 - \delta$, then we will prove that there is a unique ϵ^* -WSNE where the row player plays only the pure strategy \mathbf{i} and the column player plays the pure strategy \mathbf{j} .

Lemma 12. *If $\omega(\mathcal{F}) < 1 - \delta^*$, then there is a unique ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) in \mathcal{G}' such that $\mathbf{x}_i = 1$ and $\mathbf{y}_j = 1$.*

Proof 14. *We consider only the case that $\omega(\mathcal{F}) < 1 - \delta^*$. Then Lemma 11 says that in every ϵ^* -NE of \mathcal{G}' the column player plays the pure strategy \mathbf{j} with probability at least $1 - \frac{\epsilon^*}{1 - \epsilon^*}$. Against \mathbf{j} , the row player gets 0 for all pure strategies $i \neq \mathbf{i}$ and 1 for \mathbf{i} . Thus, in any ϵ^* -NE of \mathcal{G}' , for every pure strategy $i \neq \mathbf{i}$, the row player gets at most $\frac{\epsilon^*}{1 - \epsilon^*}$ from every pure strategy i , and the row player gets at least $1 - \frac{\epsilon^*}{1 - \epsilon^*}$ from \mathbf{i} . So, in every ϵ^* -WSNE the row player must play only the pure strategy \mathbf{i} since from every other pure strategy the player suffers regret at least $1 - \frac{2\epsilon^*}{1 - \epsilon^*}$, which is strictly larger than ϵ^* for every $\epsilon^* < \frac{1}{8}$. In*

turn, against \mathbf{i} , every pure strategy $j \neq j$ for the column player yields zero payoff while the strategy j yields payoff 1. So, the unique ϵ^* -WSNE of \mathcal{G}' is $\mathbf{x}_i = 1$ and $\mathbf{y}_j = 1$. \square

Hence, when $\omega(\mathcal{F}) < 1 - \delta^*$ the answer to Problems 7 - 9 is “No”. Thus, we have shown the following:

Theorem 4. *Assuming ETH, any algorithm that solves the Problems 2 - 9 for any constant $\epsilon < \frac{1}{8}$ requires $n^{\tilde{\Omega}(\log n)}$ time.*

Finally, for Problem 10, we define a new game \mathcal{G}'' by extending \mathcal{G}' . We add the new pure strategies i' for the row player and j' for the column player. The payoffs are shown in Figure 2. Recall that Problem 10 asks whether a bimatrix game possesses an ϵ -WSNE such that every strategy from a given set S is played with positive probability.

In order to prove our result we fix $S = i'$. First, we prove that if $\omega(\mathcal{F}) = 1$ then the game \mathcal{G}'' possesses an ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) such that $i' \in \text{supp}(\mathbf{x})$. Then we prove that if $\omega(\mathcal{F}) < 1 - \delta$, then for any ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) it holds that $i' \notin \text{supp}(\mathbf{x})$.

$$\mathcal{G}'' =$$

			j'	
\mathcal{G}'			$\frac{5}{8}, \frac{5}{8}$	
			\vdots	
			$\frac{5}{8}, \frac{5}{8}$	
i'	$\frac{5}{8}, \frac{5}{8}$	\dots	$\frac{5}{8}, \frac{5}{8}$	$0, 0$

Figure 2: The game \mathcal{G}'' .

Lemma 13. *If $\omega(\mathcal{F}) = 1$, then \mathcal{G}'' possesses an ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) such that $i' \in \text{supp}(\mathbf{x})$.*

Proof 15. *Lemma 10 says that if $\omega(\mathcal{F}) = 1$, then \mathcal{G}' possesses an ϵ^* -WSNE $(\mathbf{x}', \mathbf{y}')$ that gives payoff $\frac{5}{8}$ for each player, and \mathbf{x}' is uniform on a set of size $|X|$.*

We construct the required ϵ^* -WSNE of \mathcal{G}'' from $(\mathbf{x}', \mathbf{y}')$ as follows. We add i' to the support of \mathbf{x}' so that \mathbf{x} is a uniform mixture over $\text{supp}(\mathbf{x}') \cup i'$. For the column player, we extend \mathbf{y}' by adding zero probability for j' .

Against \mathbf{y} , pure strategies in $\text{supp}(\mathbf{x}')$ give payoff $\frac{5}{8}$, pure strategy i in \mathcal{G}' yields payoff $\frac{5}{8} + \epsilon^*$, and i' gives payoff $\frac{5}{8}$. Thus, since $(\mathbf{x}', \mathbf{y}')$ is an ϵ^* -WSNE of \mathcal{G}' , \mathbf{x} has pure regret at most ϵ^* against \mathbf{y} , as required. What remains is to show that the pure regret of \mathbf{y} is no more than ϵ^* against \mathbf{x} . Recall that, in \mathcal{G}' , against \mathbf{x}' , the payoff of each pure strategy in $\text{supp}(\mathbf{y}')$ is $\frac{5}{8}$. Now consider \mathcal{G}'' . Since, against i' , the column player gets $\frac{5}{8}$ for all $j \in \text{supp}(\mathbf{y})$, the column player still gets $\frac{5}{8}$ against \mathbf{x} for all $j \in \text{supp}(\mathbf{y})$. Moreover, against \mathbf{x} , the payoff of j' is $\frac{|X|}{|X|+1} \cdot \frac{5}{8} < \frac{5}{8}$. Thus, since $(\mathbf{x}', \mathbf{y}')$ is an ϵ^* -WSNE of \mathcal{G}' , we have that (\mathbf{x}, \mathbf{y}) is an ϵ^* -WSNE of \mathcal{G}'' with $i' \in \text{supp}(\mathbf{x})$, which completes the proof. \square

Lemma 14. *If $\omega(\mathcal{F}) < 1 - \delta^*$, then for any ϵ^* -WSNE (\mathbf{x}, \mathbf{y}) of \mathcal{G}'' it holds that $i' \notin \text{supp}(\mathbf{x})$.*

Proof 16. *We prove that the unique ϵ^* -WSNE of \mathcal{G}'' is the pure profile (i, j) . Using exactly the same arguments as in the proof of Lemma 11 we can prove that if $\omega(\mathcal{F}) < 1 - \delta^*$, then in any ϵ^* -NE of \mathcal{G}'' it holds that $x_i > 1 - \frac{\epsilon^*}{1-\epsilon^*}$ and $y_j > 1 - \frac{\epsilon^*}{1-\epsilon^*}$. Then, using exactly the same arguments as in Lemma 12 we can get that the pure strategy j for the column player yields payoff at least $1 - \frac{\epsilon^*}{1-\epsilon^*}$ while any other pure strategy, including j' , yields payoff at most $\frac{\epsilon^*}{1-\epsilon^*}$. Hence, in any ϵ^* -WSNE of \mathcal{G}'' the column player must play only the pure strategy j . Then, in order to be in an ϵ^* -WSNE the row player must play the pure strategy i . Our claim follows. \square*

The combination of Lemmas 13 and 14 gives the following theorem.

Theorem 5. *Assuming the ETH, any algorithm that solves the Problem 10 for any constant $\epsilon < \frac{1}{8}$ requires $n^{\tilde{\Omega}(\log n)}$ time.*

Acknowledgements. We would like to thank Aviad Rubinstein for alerting us to the existence of Theorem 2.

References

- [1] C. Daskalakis, P. W. Goldberg, C. H. Papadimitriou, The complexity of computing a Nash equilibrium, *SIAM Journal on Computing* 39 (1) (2009) 195–259.
- [2] I. Gilboa, E. Zemel, Nash and correlated equilibria: Some complexity considerations, *Games and Economic Behavior* 1 (1) (1989) 80 – 93.
- [3] V. Conitzer, T. Sandholm, New complexity results about Nash equilibria, *Games and Economic Behavior* 63 (2) (2008) 621 – 641.
- [4] J. Garg, R. Mehta, V. V. Vazirani, S. Yazdanbod, ETR-completeness for decision versions of multi-player (symmetric) Nash equilibria, in: *Proc. of ICALP*, 2015, pp. 554–566.
- [5] V. Bilò, M. Mavronicolas, A catalog of $\exists\mathbb{R}$ -complete decision problems about Nash equilibria in multi-player games, in: *Proc. of STACS*, 2016, pp. 17:1–17:13.
- [6] V. Bilò, M. Mavronicolas, $\exists\mathbb{R}$ -complete decision problems about symmetric Nash equilibria in symmetric multi-player games, in: *Proc. of STACS*, 2017, pp. 13:1–13:14.
- [7] H. Bosse, J. Byrka, E. Markakis, New algorithms for approximate Nash equilibria in bimatrix games, *Theoretical Computer Science* 411 (1) (2010) 164–173.
- [8] A. Czumaj, A. Deligkas, M. Fasoulakis, J. Fearnley, M. Jurdziński, R. Savani, Distributed methods for computing approximate equilibria, in: *Proc. of WINE*, 2016, pp. 15–28.
- [9] C. Daskalakis, A. Mehta, C. H. Papadimitriou, Progress in approximate Nash equilibria, in: *Proc. of EC*, 2007, pp. 355–358.
- [10] C. Daskalakis, A. Mehta, C. H. Papadimitriou, A note on approximate Nash equilibria, *Theoretical Computer Science* 410 (17) (2009) 1581–1588.

- [11] J. Fearnley, P. W. Goldberg, R. Savani, T. B. Sørensen, Approximate well-supported Nash equilibria below two-thirds, in: Proc. of SAGT, 2012, pp. 108–119.
- [12] S. C. Kontogiannis, P. G. Spirakis, Well supported approximate equilibria in bimatrix games, *Algorithmica* 57 (4) (2010) 653–667.
- [13] H. Tsaknakis, P. G. Spirakis, An optimization approach for approximate Nash equilibria, *Internet Mathematics* 5 (4) (2008) 365–382.
- [14] R. J. Lipton, E. Markakis, A. Mehta, Playing large games using simple strategies, in: Proc. of EC, 2003, pp. 36–41.
- [15] A. Rubinfeld, Settling the complexity of computing approximate two-player Nash equilibria, in: Proc. of FOCS, 2016, pp. 258–265.
- [16] E. Hazan, R. Krauthgamer, How hard is it to approximate the best Nash equilibrium?, *SIAM J. Comput.* 40 (1) (2011) 79–91.
- [17] P. Austrin, M. Braverman, E. Chlamtac, Inapproximability of NP-complete variants of Nash equilibrium, *Theory of Computing* 9 (2013) 117–142.
- [18] M. Braverman, Y. Kun-Ko, O. Weinstein, Approximating the best Nash equilibrium in $n^{o(\log n)}$ -time breaks the exponential time hypothesis, in: Proc. of SODA, 2015, pp. 970–982.
- [19] S. Aaronson, R. Impagliazzo, D. Moshkovitz, AM with multiple merlins, in: Proc. of CCC, 2014, pp. 44–55.
- [20] S. Aaronson, R. Impagliazzo, D. Moshkovitz, AM with multiple merlins, CoRR abs/1401.6848.
- [21] I. Althöfer, On sparse approximations to randomized strategies and convex combinations, *Linear Algebra and its Applications* 199 (1994) 339–355.
- [22] A. Deligkas, J. Fearnley, R. Savani, Inapproximability results for approximate Nash equilibria, in: Proc. of WINE, 2016, pp. 29–43.

- [23] Y. Babichenko, C. H. Papadimitriou, A. Rubinstein, Can almost everybody be almost happy?, in: Proc. of ITCS, 2016, pp. 1–9.
- [24] D. Moshkovitz, R. Raz, Two-query PCP with subconstant error, J. ACM 57 (5) (2010) 29:1–29:29.
- [25] I. Dinur, The PCP theorem by gap amplification, J. ACM 54 (3) (2007) Article No. 12.
- [26] I. Dinur, Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover, Electronic Colloquium on Computational Complexity (ECCC) 23 (2016) 128.
- [27] A. Czumaj, M. Fasoulakis, M. Jurdziński, Approximate Nash equilibria with near optimal social welfare, in: Proc. of IJCAI, 2015, pp. 504–510.
- [28] A. Czumaj, M. Fasoulakis, M. Jurdziński, Approximate plutocratic and egalitarian Nash equilibria: (extended abstract), in: Proc. of AAMAS, 2016, pp. 1409–1410.
- [29] J. Nash, Non-cooperative games, The Annals of Mathematics 54 (2) (1951) 286–295.
- [30] C. Daskalakis, C. H. Papadimitriou, On oblivious PTAS’s for Nash equilibrium, CoRR abs/1102.2280. [arXiv:1102.2280](https://arxiv.org/abs/1102.2280).
- [31] C. Daskalakis, C. H. Papadimitriou, On oblivious PTAS’s for Nash equilibrium, in: Proc. of STOC, 2009, pp. 75–84.