

Approximate Nash Equilibria with Near Optimal Social Welfare *

Artur Czumaj, Michail Fasoulakis, Marcin Jurdziński

Department of Computer Science
 Centre for Discrete Mathematics and its Applications (DIMAP)
 University of Warwick, United Kingdom

Abstract

It is known that Nash equilibria and approximate Nash equilibria not necessarily optimize social optima of bimatrix games. In this paper, we show that for every fixed $\varepsilon > 0$, every bimatrix game (with values in $[0, 1]$) has an ε -approximate Nash equilibrium with the total payoff of the players at least a constant factor, $(1 - \sqrt{1 - \varepsilon})^2$, of the optimum. Furthermore, our result can be made algorithmic in the following sense: for every fixed $0 \leq \varepsilon^* < \varepsilon$, if we can find an ε^* -approximate Nash equilibrium in polynomial time, then we can find in polynomial time an ε -approximate Nash equilibrium with the total payoff of the players at least a constant factor of the optimum.

Our analysis is especially tight in the case when $\varepsilon \geq \frac{1}{2}$. In this case, we show that for any bimatrix game there is an ε -approximate Nash equilibrium with constant size support whose social welfare is at least $2\sqrt{\varepsilon} - \varepsilon \geq 0.914$ times the optimal social welfare. Furthermore, we demonstrate that our bound for the social welfare is tight, that is, for every $\varepsilon \geq \frac{1}{2}$ there is a bimatrix game for which every ε -approximate Nash equilibrium has social welfare at most $2\sqrt{\varepsilon} - \varepsilon$ times the optimal social welfare.

1 Introduction

The problem of finding good equilibria in noncooperative games and understanding their properties is a central problem in modern game theory. After Nash [Nash, 1951] proved that every finite game has at least one equilibrium (so-called Nash equilibrium), the natural question arose whether we can find one efficiently. After several years of extensive research, this study has culminated in a proof that finding a Nash equilibrium is PPAD-complete even for two-players normal form games [Chen *et al.*, 2009] (see also [Daskalakis *et al.*, 2009a]), making the task of finding an *approximate*

Nash equilibrium one of the central questions in the area of equilibrium computation.

Since scaling the payoffs by any positive factor, and applying any additive constant, results in an equilibrium-equivalent game, one typically considers games with all payoffs normalized to be in the interval $[0, 1]$. Then, we say a set of mixed strategies is an ε -approximate Nash equilibrium, if each player has only at most ε incentive to defect. The PPAD-hardness of finding a Nash equilibrium can be extended to provide a PPAD-hardness of designing a fully polynomial-time approximation scheme for this problem [Chen *et al.*, 2009]. In contrast, [Lipton *et al.*, 2003], based on [Althofer, 1994], showed that for every $\varepsilon > 0$, one can find an ε -approximate Nash equilibrium in quasi-polynomial-time $n^{O(\varepsilon^{-2} \log n)}$ by examining all supports of size $O(\varepsilon^{-2} \log n)$. This work prompted a series of papers [Bosse *et al.*, 2010; Daskalakis *et al.*, 2007; 2009b; Kontogiannis *et al.*, 2006; Tsaknakis and Spirakis, 2008] giving polynomial-time algorithms to find an ε -approximate Nash equilibrium for decreasing values of ε , culminating with the state of the art result by [Tsaknakis and Spirakis, 2008], which finds in polynomial time a 0.3393-approximate Nash equilibrium of a bimatrix game. However, the question whether there is a polynomial-time approximation scheme (which could run in time $n^{O(f(1/\varepsilon))}$) still remains one of the central open questions in the area of equilibria computations.

While the Nash theorem [Nash, 1951] ensures that every finite two-player game has at least one Nash equilibrium, typical games possess many equilibria and it is natural to seek those equilibria that are more desirable than others. One natural measure of the most desirable equilibria is to maximize its social welfare, that is, the sum of players' payoffs. Unlike the problem of finding a Nash equilibrium, which is known to be PPAD-complete, finding a Nash equilibrium with maximal social welfare is known to be NP-hard [Gilboa and Zemel, 1989; Conitzer and Sandholm, 2008], and thus, it is likely to be computationally even more difficult. In fact, it is even NP-hard to approximate (to any positive ratio) the maximum social welfare obtained in an exact Nash equilibrium, even in symmetric 2-player games [Conitzer and Sandholm, 2008, Corollary 6]. Therefore, it is natural to ask the question of computational complexity of finding an ε -approximate Nash equilibrium that approximates well the optimal social welfare. The mentioned above quasi-polynomial-time algorithm

*Research partially supported by the Centre for Discrete Mathematics and its Applications (DIMAP) and by EPSRC award EP/D063191/1. Email: {A.Czumaj, M.Fasoulakis, M.Jurdzinski}@warwick.ac.uk

by [Lipton *et al.*, 2003] not only finds an ε -approximate Nash equilibrium, but also the social welfare of the equilibrium found is an ε -approximation of the social welfare in any Nash equilibrium. In other words, in quasi-polynomial-time we can find an arbitrarily good approximate Nash equilibrium with social welfare near to the best Nash equilibrium. Although this result raised a hope that it may be possible to extend it to design a polynomial-time algorithm, there are strong hardness results known now. Hazan and Krauthgamer [Hazan and Krauthgamer, 2011] show that for a fixed small ε , finding an ε -approximate Nash equilibrium in a two-player game whose social welfare is off by at most ε from best Nash equilibrium is as hard as finding a hidden clique of size $O(\log n)$ in the random graph $G_{n,1/2}$ (see also [Austrin *et al.*, 2013; Minder and Vilenchik, 2009]). These hardness results have been further strengthened by Braverman *et al.* [Braverman *et al.*, 2015], who showed that assuming the deterministic Exponential Time Hypothesis (that any deterministic algorithm for 3SAT requires $2^{\Omega(n)}$ time), there is a constant $\varepsilon > 0$ such that any algorithm for finding an ε -approximate Nash equilibrium whose social welfare is at least $(1 - \varepsilon)$ times the optimal social welfare of a Nash equilibrium of the game, requires $2^{\Omega(\log n)}$ time. The above results demonstrate that it is very unlikely to obtain a polynomial-time approximation scheme that for every positive constants ε and ε' would construct in polynomial time an ε -approximate Nash equilibrium whose social welfare is at least $(1 - \varepsilon')$ times the optimal social welfare of a Nash equilibrium of the game. We note that for large ε , a stronger (optimal) result is possible: Austrin *et al.* [Austrin *et al.*, 2013, Theorem 1.3] gave a polynomial-time algorithm that finds a $\frac{1}{2}$ -approximate Nash equilibrium whose social welfare is as good as that of any Nash equilibrium.

1.1 Approximate Nash equilibria with near optimal social welfare

In this paper we take a more pragmatic approach and focus on the analysis of the social welfare in ε -approximate Nash equilibria in a two-player game for a fixed ε , for the regime when we know that we can find an ε -approximate Nash equilibrium. Our goal is more general than that presented in earlier works, like e.g., in [Austrin *et al.*, 2013; Braverman *et al.*, 2015; Hazan and Krauthgamer, 2011; Minder and Vilenchik, 2009]; it is not to compare the social welfare of an ε -approximate Nash equilibrium to that of any Nash equilibrium, but rather to *compare it with the optimal social welfare*.

It is known that a Nash equilibrium can be arbitrarily far from the optimal social welfare in a bimatrix game. A simple example describing this situation is a *prisoners' dilemma* game:

	C	D
C	$(\frac{2}{3}, \frac{2}{3})$	$(0, 1)$
D	$(1, 0)$	(δ, δ)

Assuming that $\delta < \frac{2}{3}$, the optimal social welfare is achieved by the strategy profile (C, C) with total payoff of $\frac{4}{3}$, but the unique Nash equilibrium is the strategy profile (D, D) with total payoff of 2δ . Thus, by taking δ arbitrarily small, we

can make the social welfare of a Nash equilibrium arbitrarily far from the optimal social welfare of a game.

The central question studied in this paper is if we allow the players up to ε loss to deviate from the best response strategy, whether we can find a stable strategy profile (an ε -approximate Nash equilibrium) that guarantees the players a value close to the social optimum?

We note that, to the best of our knowledge, the known polynomial-time algorithms to construct an ε -approximate Nash equilibrium for a constant $\varepsilon > 0$, do not guarantee any welfare for the ε -approximate Nash equilibrium and they return an ε -approximate Nash equilibrium strategy profile which can be arbitrarily far from the optimal social welfare (see, e.g., [Bosse *et al.*, 2010; Daskalakis, 2013; Daskalakis *et al.*, 2007; 2009b; Kontogiannis *et al.*, 2006; Tsaknakis and Spirakis, 2008] for more details).

1.2 New contributions

In this paper we provide several results showing that for every bimatrix game, for every $\varepsilon > 0$, there is always an ε -approximate Nash equilibrium with near optimal social welfare, at least a constant fraction the optimal social welfare. Our analysis shows that by considering an appropriate mixture of the optimal strategies and exact or approximate Nash equilibria, one can find the desired approximate Nash equilibrium with near optimal social welfare.

We begin with the case when $\varepsilon \geq \frac{1}{2}$, the case for which it is known that there is always an ε -approximate Nash equilibrium with constant size support (cf. [Daskalakis *et al.*, 2009b]). We show that in that case we can find an ε -approximate Nash equilibrium with constant size support whose social welfare is at least $2\sqrt{\varepsilon} - \varepsilon \geq 0.914$ times the optimal social welfare. Furthermore, we demonstrate that our bound for the social welfare is *tight*.

Theorem 1 *For every $\varepsilon \geq \frac{1}{2}$, we can construct in polynomial time an ε -approximate Nash equilibrium (and with constant size support) whose social value is at least $2\sqrt{\varepsilon} - \varepsilon$ times the optimal social welfare. Furthermore, there is a bimatrix game for which for every $\varepsilon \geq \frac{1}{2}$, every ε -approximate Nash equilibrium has social welfare no more than $2\sqrt{\varepsilon} - \varepsilon$ times the optimal social welfare.*

In particular, we can construct in polynomial time a $\frac{1}{2}$ -approximate Nash equilibrium whose social welfare is at least $\frac{2\sqrt{2}-1}{2} \approx 0.914$ times the optimal social welfare.

As a byproduct of our approach, we also obtain a stronger result for the class of win-lose bimatrix games and show that for any $\varepsilon \in [\frac{1}{2}, 1]$, for any win-lose bimatrix game with values in $\{0, 1\}$, we can find in polynomial time an ε -approximate Nash equilibrium with optimal social welfare (Theorem 5).

The case $\varepsilon < \frac{1}{2}$ is more challenging and while we do not have a tight bound for the social welfare in this case, we can still construct an ε -approximate Nash equilibrium with social welfare that is at least κ_ε times the optimum, for some positive constant κ_ε . One challenge in the case $\varepsilon < \frac{1}{2}$ stems from the fact that there are bimatrix games with no ε -approximate Nash equilibrium with constant support (cf. [Feder *et al.*, 2007]), which requires us to use a different approach than that in Theorem 1 to deal with this case. Using as a starting point

ε^* -approximate Nash equilibria with arbitrary social welfare and $\varepsilon^* < \varepsilon$, we modify them to obtain an ε -approximate Nash equilibrium with high social welfare to get the following.

Theorem 2 For every fixed positive $\varepsilon < \frac{1}{2}$ there is a positive constant $\kappa_\varepsilon = (1 - \sqrt{1 - \varepsilon})^2$, such that every bimatrix game has an ε -approximate Nash equilibrium with social welfare at least κ_ε times the optimal social welfare.

Our construction is algorithmic and gives the following.

Theorem 3 Let ε^* be such that there is a polynomial time algorithm for finding an ε^* -approximate Nash equilibrium of a bimatrix game. Then for every fixed positive $\varepsilon > \varepsilon^*$, there is a positive constant $\zeta_{\varepsilon, \varepsilon^*} = (1 - \sqrt{\frac{1 - \varepsilon}{1 - \varepsilon^*}})^2$, such that for every bimatrix game one can find in polynomial time an ε -approximate Nash equilibrium with social welfare at least $\zeta_{\varepsilon, \varepsilon^*}$ times the optimal social welfare.

We also obtain further algorithmic results improving the bounds for the social welfare above in several special cases for $\varepsilon < \frac{1}{2}$. For example, in the case when the optimal social welfare is at least $\frac{2-3\varepsilon}{1-\varepsilon}$, then in Theorem 8 we design a polynomial-time algorithm that finds an ε -approximate Nash equilibrium with constant support size and with social welfare at least $(1 - \frac{\varepsilon(1-\varepsilon)}{2-3\varepsilon}) \geq 0.5$ times the optimum social welfare. For this case we will prove that if the optimum social welfare is less than $\frac{2-3\varepsilon}{1-\varepsilon}$, we need logarithmic support in order to create an ε -Nash equilibrium.

We will prove Theorem 1 in Section 3 and Theorems 2 and 3 in Section 4.

2 Preliminaries

Consider a two-player normal form game with n strategies in the disposal of every player and let (R, C) be the payoff matrices in $[0, 1]^{n \times n}$ of the row player and the column player respectively. If the row player plays the strategy i and the column player plays the strategy j then the row player's payoff is R_{ij} and the column player's payoff is C_{ij} . A mixed strategy $x \in [0, 1]^n$ is a column vector that describes a probability distribution on the n pure strategies of a player; a support of a mixed strategy x is the set of the pure strategies i such that $x_i > 0$. Note that if the row player plays a mixed strategy x and the column player plays a mixed strategy y the expected payoff of the row player is $x^T R y$ and the expected payoff of the column player is $x^T C y$. The social welfare is the total payoff of both players, i.e., it is $\text{cost} = x^T R y + x^T C y = x^T (R + C) y$.

A Nash equilibrium is a strategy profile (x^*, y^*) such that

$$\begin{aligned} x^{*T} R y^* &\geq e_i^T R y^* && \text{for every } i = 1, \dots, n, \\ x^{*T} C y^* &\geq x^{*T} C e_i && \text{for every } i = 1, \dots, n, \end{aligned}$$

where $e_i \in [0, 1]^n$ is the column vector with 1 in its coordinate i and 0 elsewhere.

For any $\varepsilon \geq 0$, an ε -approximate Nash equilibrium is any strategy profile (x^*, y^*) such that

$$\begin{aligned} x^{*T} R y^* + \varepsilon &\geq e_i^T R y^* && \text{for every } i = 1, \dots, n, \\ x^{*T} C y^* + \varepsilon &\geq x^{*T} C e_i && \text{for every } i = 1, \dots, n. \end{aligned}$$

Note that a 0-Nash equilibrium is a (exact) Nash equilibrium.

Throughout the paper, we let (i, j) to denote the pure strategy profile that maximizes the sum of the payoffs of the two players (utilitarian objective). We define opt to be the optimal social welfare, that is,

$$\forall x, y \in [0, 1]^n \text{opt} = R_{ij} + C_{ij} \geq x^T (R + C) y. \quad (1)$$

(Note that i and j can be trivially found in $O(n^2)$ time.)

We define the pure strategy r of the row player as the best response strategy of the row player to the strategy j of the column player and the pure strategy c of the column player as the best response strategy of the column player to the strategy i of the row player. The optimality of the profile (i, j) yields:

$$\text{opt} = R_{ij} + C_{ij} \geq R_{ic} + C_{ic}, \quad (2)$$

$$\text{opt} = R_{ij} + C_{ij} \geq R_{rj} + C_{rj}. \quad (3)$$

The central goal of this paper is for a fixed $\varepsilon \in [0, 1]$, to find an ε -approximate Nash equilibrium strategy profile (x^*, y^*) whose social welfare cost is as close to opt as possible.

In our analysis, we will consider several cases depending on the values of $R_{rj} - R_{ij}$ and $C_{ic} - C_{ij}$:

- $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$,
- $R_{rj} - R_{ij} \geq \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$ (and the symmetric case $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} \geq \varepsilon$),
- $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$ (and the symmetric case $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} < \varepsilon$).

We will use the fact that in the first case, when $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$, the strategy profile (i, j) (which can be found in polynomial time) is an ε -approximate Nash equilibrium, and since it has the optimal social welfare, in this case we can find an optimal solution by choosing strategy (i, j) . Thus, our main task will be to find a good algorithm to construct an ε -approximate Nash equilibrium in the other cases.

In our analysis, we will separately consider two regimes: one when $\varepsilon \geq \frac{1}{2}$ and one when $\varepsilon < \frac{1}{2}$.

3 Approximation with $\varepsilon \geq \frac{1}{2}$

We begin with the scenario when $\varepsilon \geq \frac{1}{2}$, proving Theorem 1. We will show in Section 3.1 that if $\varepsilon \geq \frac{1}{2}$, then one can find an ε -approximate Nash equilibrium with constant size support that has an almost optimal social welfare, at least $2\sqrt{\varepsilon} - \varepsilon \geq 0.914$ times the optimal social welfare. We will also prove that our bound is tight for any $\varepsilon \geq \frac{1}{2}$, by showing in Section 3.2 explicit bimatrix games for which every ε -approximate Nash equilibrium has social welfare no more than $2\sqrt{\varepsilon} - \varepsilon$ times the optimal social welfare.

Let us recall that (i, j) is the pure strategy profile that maximizes the sum of the payoffs of the two players, and hence $\text{opt} = R_{ij} + C_{ij}$ (cf. (1)). Let us recall that r is the pure strategy of the row player that is the best response strategy of the row player to the strategy j of the column player and that c is the pure strategy of the column player that is the best response strategy of the column player to the strategy i of the row player. We will now consider several cases depending on the values of $R_{rj} - R_{ij}$ and $C_{ic} - C_{ij}$.

Let us first note that it is impossible to have $R_{rj} - R_{ij} \geq \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$, or to have $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} \geq \varepsilon$ (since these cases are symmetric, we will focus only on the first one). To show that we cannot have $R_{rj} - R_{ij} \geq \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$, we first observe that these inequalities yield:

$$R_{ij} \leq R_{rj} - \varepsilon \leq 1 - \varepsilon \text{ and } C_{ij} < C_{ic} - \varepsilon < 1 - \varepsilon . \quad (4)$$

Next, $R_{rj} - R_{ij} \geq \varepsilon$ together with (3) yield $R_{ij} + C_{ij} \geq R_{rj} + C_{rj} \geq R_{ij} + \varepsilon + C_{rj}$, what implies $C_{ij} \geq \varepsilon$. Similarly, $C_{ic} - C_{ij} > \varepsilon$ and (2) give $R_{ij} + C_{ij} \geq R_{ic} + C_{ic} > R_{ic} + C_{ij} + \varepsilon$, and hence $R_{ij} > \varepsilon$. Now, however, we observe that with the assumption $\varepsilon \geq \frac{1}{2}$, the inequalities above form a contradiction, and therefore this case cannot happen.

Since we cannot have either of the cases $R_{rj} - R_{ij} \geq \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$, or $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} \geq \varepsilon$, we only have to consider one of the following three scenarios: (1) $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$, (2) $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$, (3) $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} < \varepsilon$.

We will now consider these cases, depending on the values of $R_{rj} - R_{ij}$ and $C_{ic} - C_{ij}$:

- (1) If $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$, then we know that the strategy profile (i, j) is an ε -approximate Nash equilibrium with the optimal social welfare.
- (2) If $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$, then we note that

$$C_{ic} > C_{ij} + \varepsilon \geq \max\{C_{ij}, \varepsilon\} , \quad (5)$$

and that (2) yields $R_{ij} - R_{ic} \geq C_{ic} - C_{ij} > \varepsilon$.

Next, we prove a key lemma describing an ε -approximate Nash equilibrium in our setting.

Lemma 4 *Let $\varepsilon \in [\frac{1}{2}, 1]$, $R_{rj} - R_{ij} < \varepsilon$, and $C_{ic} - C_{ij} > \varepsilon$. Let $\mathbf{p} = \frac{\varepsilon}{C_{ic} - C_{ij}}$. The strategy profile $(i, \mathbf{p}j + (1 - \mathbf{p})c)$, where \mathbf{p} is the probability for the column player to play strategy j and $(1 - \mathbf{p})$ is the probability of playing strategy c respectively, is an ε -approximate Nash equilibrium.*

Proof. Let us first notice that \mathbf{p} is well defined with $0 < \mathbf{p} \leq 1$ since $0 < \varepsilon < C_{ic} - C_{ij}$.

Let b be the best response strategy of the row player to the strategy $\mathbf{p}j + (1 - \mathbf{p})c$ of the column player. If the row player plays strategy i , her incentive to deviate is:

$$\begin{aligned} & \mathbf{p}R_{bj} + (1 - \mathbf{p})R_{bc} - \mathbf{p}R_{ij} - (1 - \mathbf{p})R_{ic} \\ & \leq \mathbf{p}R_{rj} + (1 - \mathbf{p}) - \mathbf{p}R_{ij} - (1 - \mathbf{p})R_{ic} \\ & \leq \mathbf{p} + (1 - \mathbf{p}) - \mathbf{p}(R_{ij} - R_{ic}) \\ & = 1 - \frac{\varepsilon \cdot (R_{ij} - R_{ic})}{C_{ic} - C_{ij}} \leq 1 - \varepsilon \leq \varepsilon . \end{aligned}$$

The first inequality follows from $R_{bj} \leq R_{rj}$ and $R_{bc} \leq 1$, the second one because of the fact that $R_{rj} \leq 1$ and $R_{ic} \geq 0$, the third one because $R_{ij} + C_{ij} \geq R_{ic} + C_{ic}$, and the final one follows from the fact that $\varepsilon \geq \frac{1}{2}$.

On the other hand, the incentive to deviate for the column player when the row player plays i is $C_{ic} - \mathbf{p}C_{ij} - (1 - \mathbf{p})C_{ic} = \varepsilon$. Hence the strategy profile $(i, \mathbf{p}j + (1 - \mathbf{p})c)$ is an ε -approximate Nash equilibrium. \square

- (3) $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} < \varepsilon$ is symmetric to case (2).

3.1 Upper bound in Theorem 1

We now prove that ε -APPROXIMATE NASH (R, C, ε) presented below, returns an ε -approximate Nash equilibrium with social welfare at least $(2\sqrt{\varepsilon} - \varepsilon)\text{opt}$. By the arguments above, we only have to consider the following scenarios: (1) $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$, (2) $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$, (3) $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} < \varepsilon$.

ε -APPROXIMATE NASH (R, C, ε)

- Find i, j such that $R_{ij} + C_{ij}$ is maximized.
- Find r, c such that R_{rj} is maximized and C_{ic} is maximized.
- If $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$, then return strategy profile (i, j) .
- If $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$, then set $\mathbf{p} = \frac{\varepsilon}{C_{ic} - C_{ij}}$ and return strategy profile $(i, \mathbf{p}j + (1 - \mathbf{p})c)$.
- If $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} < \varepsilon$, then set $\mathbf{p} = \frac{\varepsilon}{R_{rj} - R_{ij}}$ and return strategy profile $(\mathbf{p}i + (1 - \mathbf{p})r, j)$.

Let us recall that if $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$, then the strategy (i, j) is an ε -approximate Nash equilibrium with social welfare opt , and therefore the algorithm will return an optimum solution that is an ε -approximate Nash equilibrium. Therefore, we only have to consider scenarios (2) and (3). Since these scenarios are symmetric, we focus only on scenario (2), when $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$: we prove that the strategy profile $(i, \mathbf{p}j + (1 - \mathbf{p})c)$ with $\mathbf{p} = \frac{\varepsilon}{C_{ic} - C_{ij}}$ has social welfare at least $(2\sqrt{\varepsilon} - \varepsilon)\text{opt}$.

The social welfare of our solution is $\text{cost} = \mathbf{p}(R_{ij} + C_{ij}) + (1 - \mathbf{p})(R_{ic} + C_{ic})$. Let

$$\begin{aligned} \rho & = \frac{\text{opt}}{\text{cost}} = \frac{R_{ij} + C_{ij}}{\mathbf{p}(R_{ij} + C_{ij}) + (1 - \mathbf{p})(R_{ic} + C_{ic})} \\ & \leq \frac{R_{ij} + C_{ij}}{\mathbf{p}(R_{ij} + C_{ij}) + (1 - \mathbf{p})C_{ic}} . \end{aligned} \quad (6)$$

Observe that if we consider the last bound as a function of R_{ij} , we obtain a function of the form $f(x) = \frac{x + \beta}{\mathbf{p}x + \gamma}$, with $0 \leq \mathbf{p} \leq 1$, $\beta = C_{ij}$ and $\gamma = \mathbf{p}C_{ij} + (1 - \mathbf{p})C_{ic}$. Notice further that since by (5), we have $\mathbf{p}C_{ij} + (1 - \mathbf{p})C_{ic} > \mathbf{p}C_{ij} + (1 - \mathbf{p})C_{ij} = C_{ij} \geq 0$, we obtain $\gamma > \beta \geq 0$. Therefore, by considering the derivative $f'(x) = \frac{\gamma - \mathbf{p}\beta}{(\mathbf{p}x + \gamma)^2} > 0$, we observe that f is increasing in x . Thus, the right hand side of (6) takes the maximum value when R_{ij} is maximum, that is, is equal to 1, independently from the other variables. Hence,

$$\begin{aligned} \rho & \leq \frac{1 + C_{ij}}{\mathbf{p} + \mathbf{p}C_{ij} + (1 - \mathbf{p})C_{ic}} \\ & = \frac{-C_{ij}^2 - C_{ij}(1 - C_{ic}) + C_{ic}}{-C_{ij}(C_{ic} - \varepsilon) + C_{ic}(C_{ic} - \varepsilon) + \varepsilon} . \end{aligned} \quad (7)$$

We note that the right hand side of (7) takes maximum when $C_{ic} = C_{ij} + \sqrt{\varepsilon}$, and hence when $\mathbf{p} = \sqrt{\varepsilon}$. If we plug this in (7), then we obtain $\rho \leq \frac{1 + C_{ij}}{2\sqrt{\varepsilon} - \varepsilon + C_{ij}}$. Next, we observe that since $\varepsilon \in [\frac{1}{2}, 1]$ we have $2\sqrt{\varepsilon} - \varepsilon \leq 1$, and hence




Figure 1: Bound for $\rho = \frac{\text{opt}}{\text{cost}}$ as a function of ε , $\varepsilon \geq \frac{1}{2}$. Notice that $\rho(1) \leq 1$ and $\rho(\frac{1}{2}) \leq \frac{2}{2\sqrt{2}-1} \approx 1.094$.

the right hand side of is decreasing and takes the maximum at $C_{ij} = 0$. Therefore $\rho \leq \frac{1}{2\sqrt{\varepsilon}-\varepsilon}$. This completes the proof of the first part (upper bound) of Theorem 1.

Figure 1 depicts the upper bound as a function of ε .

3.2 Lower bound in Theorem 1

We now show the second part of Theorem 1 and for every $\varepsilon \in [\frac{1}{2}, 1]$, we present a game for which the social welfare of every ε -approximate Nash equilibrium is at most $(2\sqrt{\varepsilon} - \varepsilon)\text{opt}$.

Fix ε , $\frac{1}{2} \leq \varepsilon \leq 1$. Consider a bimatrix game with one strategy for the row player, strategy i , and with two strategies for the column player, strategies j and c . Set $R_{ij} = 1$, $C_{ij} = 0$, $R_{ic} = 0$, and $C_{ic} = \sqrt{\varepsilon}$, resulting in the following game:

$$\begin{array}{c|cc} & j & c \\ \hline i & (1, 0) & (0, \sqrt{\varepsilon}) \end{array}$$

The optimal strategy is (i, j) with the social welfare $\text{opt} = 1$. In order to obtain an ε -approximate Nash equilibrium the column player needs to randomize between her strategies, playing strategy j with probability \mathbf{p} and strategy c with probability $(1 - \mathbf{p})$. Then, the strategy profile $(i, \mathbf{p}j + (1 - \mathbf{p})c)$ is an ε -approximate Nash equilibrium if and only if $\sqrt{\varepsilon} \leq (1 - \mathbf{p})\sqrt{\varepsilon} + \varepsilon$. This is equivalent to $\mathbf{p} \leq \sqrt{\varepsilon}$. Conditioned on this, we bound the social welfare of any ε -approximate Nash equilibrium for this game. For any $0 \leq \mathbf{p} \leq \sqrt{\varepsilon}$, if we denote the social welfare of an ε -approximate Nash equilibrium with fixed \mathbf{p} by $\text{cost}_{\mathbf{p}}$, then we obtain, $\text{cost}_{\mathbf{p}} = \mathbf{p} + (1 - \mathbf{p})\sqrt{\varepsilon} \leq \sqrt{\varepsilon} + \sqrt{\varepsilon}(1 - \sqrt{\varepsilon}) = 2\sqrt{\varepsilon} - \varepsilon$. Therefore, since $\text{opt} = 1$, we conclude that for the game defined above, the social welfare of every ε -approximate Nash equilibrium is at most $2\sqrt{\varepsilon} - \varepsilon$ times the optimal social welfare. This completes the proof of the second part (lower bound) of Theorem 1.

3.3 Win-lose games with $\varepsilon \geq \frac{1}{2}$

We note that for the class of win-lose games, one can easily show the following stronger bound.

Theorem 5 For any win-lose bimatrix game with values in $\{0, 1\}$ and any $\varepsilon \in [\frac{1}{2}, 1]$, we can find in polynomial time an ε -approximate Nash equilibrium with optimal social welfare.




Figure 2: Bound for $(1 - \sqrt{1 - \varepsilon})^2$ as a function of ε , as in Theorem 2; $(1 - \sqrt{1 - \varepsilon})^2 \approx 0.0858$ for $\varepsilon = \frac{1}{2}$.

4 Approximation with $\varepsilon < \frac{1}{2}$

The analysis of the case $\varepsilon < \frac{1}{2}$ is more complicated and our results are not as tight as those for the case $\varepsilon \geq \frac{1}{2}$. One important reason why this case is more challenging is that for $\varepsilon < \frac{1}{2}$, we know that we have to consider large support size of the strategies. This follows from [Feder *et al.*, 2007], who showed that for $\varepsilon < \frac{1}{2}$, to find an ε -approximate-Nash equilibrium the support needs to be of size logarithmic in the number of strategies available to the players.

We begin with a general transformation that takes an arbitrary ε^* -approximate Nash equilibria with arbitrary social welfare and outputs an ε -approximate Nash equilibrium, $\varepsilon^* < \varepsilon$, with social welfare at least a constant fraction the optimal social welfare. This is achieved by considering an appropriate mixture of a strategy profile with the optimal social welfare and an ε^* -approximate Nash equilibrium. We also show that our transformation runs in polynomial time, and thus if there is a polynomial-time algorithm finding an ε^* -approximate Nash equilibrium then our scheme can find in polynomial time an ε -approximate Nash equilibrium, $\varepsilon^* < \varepsilon$, with social welfare at least a constant fraction the optimal social welfare. Next, we will analyze the special case where the social welfare is greater or equal to $\frac{2-3\varepsilon}{1-\varepsilon}$, when we find ε -approximate Nash equilibria with high social welfare.

4.1 Reducing social welfare

As mentioned earlier, if $\varepsilon < \frac{1}{2}$ then we cannot hope to find an ε -approximate Nash equilibrium with constant size support, which is the approach we used in Section 3. However, we will show that using an existing ε^* -approximate Nash equilibrium, $\varepsilon^* < \varepsilon$, with an arbitrary social welfare, we can construct an ε -approximate Nash equilibrium with social welfare that is at least a constant times optimal, to conclude Theorem 2. We begin with the following key lemma.

Lemma 6 Let $0 \leq \varepsilon^* < 1$, and $\varepsilon^* < \varepsilon < 1$. Let (x^*, y^*) be the strategy profile of an ε^* -approximate Nash equilibrium. Then, for $\mathbf{p} = 1 - \sqrt{(1 - \varepsilon)/(1 - \varepsilon^*)}$, the strategy profile $(\mathbf{p}i + (1 - \mathbf{p})x^*, \mathbf{p}j + (1 - \mathbf{p})y^*)$ is an ε -approximate Nash equilibrium with the social welfare $\text{cost} \geq \mathbf{p}^2 \text{opt}$.

Proof. Since (x^*, y^*) is a strategy profile of an ε^* -approximate Nash equilibrium, the maximum incentive to deviate for any player in the strategy profile (x^*, y^*) is ε^* . Therefore, since under strategies (i, j) , (i, y^*) , (x^*, j) no player can improve its payoff by more than 1, we obtain that if the players play the strategy profile $(p_i + (1-p)x^*, p_j + (1-p)y^*)$, then the maximum incentive to deviate for any player is upper bounded by the following:

$$(1-p)^2\varepsilon^* + p^2 + p(1-p) + p(1-p) = 1 - (1-p)^2(1-\varepsilon^*).$$

Hence, to ensure that this strategy is an ε -approximate Nash equilibrium for $0 \leq \varepsilon^* < \varepsilon < 1$, we set $p = 1 - \sqrt{(1-\varepsilon)/(1-\varepsilon^*)}$. It is easy to check that $0 \leq p \leq 1$.

Next, we can bound the social welfare $\text{cost} = p^2(R_{ij} + C_{ij}) + (1-p)^2x^{*T}(R+C)y^* \geq p^2(R_{ij} + C_{ij}) = p^2\text{opt}$. \square

Proof of Theorem 2: We choose $\varepsilon^* = 0$ in Lemma 6 (here we use Nash theorem [Nash, 1951] to guarantee the existence of an exact Nash equilibrium (x^*, y^*)) to ensure that one can use the strategy profile $(p_i + (1-p)x^*, p_j + (1-p)y^*)$ with $p = 1 - \sqrt{1-\varepsilon}$ to obtain an ε -approximate Nash equilibrium with $\text{cost} \geq p^2\text{opt} = (1 - \sqrt{1-\varepsilon})^2\text{opt}$. \square

Our construction above can be trivially transformed into a polynomial time algorithm, assuming that we have at hand a polynomial-time algorithm for finding an ε^* -approximate Nash equilibrium in any bimatrix game. This proves Theorem 3 with $\zeta_{\varepsilon, \varepsilon^*} = (1 - \sqrt{\frac{1-\varepsilon}{1-\varepsilon^*}})^2$. Since the best currently known value for ε^* is 0.3393 [Tsaknakis and Spirakis, 2008], this approach works (currently) only for $\varepsilon > 0.3393$.

4.2 Analysis of the case $\text{opt} \geq \frac{2-3\varepsilon}{1-\varepsilon}$

We consider a special case, when $\text{opt} \geq \frac{2-3\varepsilon}{1-\varepsilon}$, for which we can construct approximate Nash equilibria with high social welfare. We will show in Theorem 8 that there is a good ε -approximate Nash equilibrium that has a constant size support and high social welfare. This result is complemented by Theorem 9 that shows that if $\text{opt} < \frac{2-3\varepsilon}{1-\varepsilon}$, then an ε -Nash equilibrium may require a logarithmic size support.

We begin with the case $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$ (the case $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} < \varepsilon$ is symmetric).

Lemma 7 Let $\varepsilon \in [0, \frac{1}{2})$, $R_{rj} - R_{ij} < \varepsilon$, $C_{ic} - C_{ij} > \varepsilon$, and $p = \frac{1-\varepsilon}{-1+2\varepsilon+2R_{ij}(1-\varepsilon)}$. If $\text{opt} \geq \frac{2-3\varepsilon}{1-\varepsilon}$ then the strategy profile $(i, p_j + (1-p)c)$ is an ε -approximate Nash equilibrium with social welfare greater than $\frac{2-4\varepsilon+\varepsilon^2}{2-3\varepsilon} \cdot \text{opt} = (1 - \frac{\varepsilon(1-\varepsilon)}{2-3\varepsilon}) \cdot \text{opt}$.

Proof. We first show that p is well defined with $0 \leq p \leq 1$. Since $C_{ic} - C_{ij} > \varepsilon$, we get $C_{ij} < 1 - \varepsilon$. Thus, if $\text{opt} \geq \frac{2-3\varepsilon}{1-\varepsilon}$, then $\text{opt} = R_{ij} + C_{ij}$ yields $R_{ij} \in (\frac{1-\varepsilon-\varepsilon^2}{1-\varepsilon}, 1]$ and $C_{ij} \in [\frac{1-2\varepsilon}{1-\varepsilon}, 1 - \varepsilon)$. Hence, $-1 + 2\varepsilon + 2R_{ij}(1-\varepsilon) > -1 + 2\varepsilon + 2(1-\varepsilon-\varepsilon^2) \geq 1 - \varepsilon$, and thus p is well defined.

Next, we prove that the strategy profile $(i, p_j + (1-p)c)$ is an ε -approximate Nash equilibrium. Let b be the best response strategy of the row player to the strategy profile

$(p_j + (1-p)c)$ of the column player. Then the incentive of the row player to deviate from strategy i is:

$$pR_{bj} + (1-p)R_{bc} - pR_{ij} - (1-p)R_{ic} \leq 1 - pR_{ij},$$

and the incentive of the column player to deviate is:

$$\begin{aligned} C_{ic} - pC_{ij} - (1-p)C_{ic} &= p(C_{ic} - C_{ij}) \\ &\leq p \left(1 - \left(\frac{2-3\varepsilon}{1-\varepsilon} - R_{ij} \right) \right) = p \left(R_{ij} - \frac{1-2\varepsilon}{1-\varepsilon} \right). \end{aligned}$$

Here we use the facts that c is the best response strategy of the column player to the strategy i of the row player, and that $C_{ic} \leq 1$ and $C_{ij} = \text{opt} - R_{ij} \geq \frac{2-3\varepsilon}{1-\varepsilon} - R_{ij}$.

Our choice of p ensures that $p \left(R_{ij} - \frac{1-2\varepsilon}{1-\varepsilon} \right) = 1 - pR_{ij} = \frac{-1+2\varepsilon+R_{ij}(1-\varepsilon)}{-1+2\varepsilon+2R_{ij}(1-\varepsilon)}$, which takes the maximum at $R_{ij} = 1$. Therefore $1 - pR_{ij} \leq \frac{-1+2\varepsilon+(1-\varepsilon)}{-1+2\varepsilon+2(1-\varepsilon)} = \varepsilon$, what implies that the strategy profile $(i, p_j + (1-p)c)$ is an ε -approximate Nash equilibrium. This yields the following lower bound for the social welfare cost of the strategy profile $(i, p_j + (1-p)c)$:

$$\begin{aligned} \frac{\text{opt}}{\text{cost}} &\leq \frac{1 + \frac{1-2\varepsilon}{1-\varepsilon}}{pR_{ij} + \frac{1-2\varepsilon}{1-\varepsilon}} = \frac{1 + \frac{1-2\varepsilon}{1-\varepsilon}}{\frac{R_{ij}(1-\varepsilon)}{-1+2\varepsilon+2R_{ij}(1-\varepsilon)} + \frac{1-2\varepsilon}{1-\varepsilon}} \\ &\leq \frac{1 + \frac{1-2\varepsilon}{1-\varepsilon}}{1 - \varepsilon + \frac{1-2\varepsilon}{1-\varepsilon}} = \frac{2-3\varepsilon}{2-4\varepsilon+\varepsilon^2}. \quad \square \end{aligned}$$

With Lemma 7 at hand, we can prove the following.

Theorem 8 Let $\varepsilon \in [0, \frac{1}{2})$ and $\text{opt} \geq \frac{2-3\varepsilon}{1-\varepsilon}$. Then one can find in polynomial-time an ε -approximate Nash equilibrium with constant support size and with social welfare at least $(1 - \frac{\varepsilon(1-\varepsilon)}{2-3\varepsilon}) \cdot \text{opt} \geq 0.5 \cdot \text{opt}$.

Proof. We consider three cases:

- If $R_{rj} - R_{ij} \leq \varepsilon$ and $C_{ic} - C_{ij} \leq \varepsilon$, then the strategy profile (i, j) is an ε -approximate Nash equilibrium with $\text{cost} = \text{opt}$.
- If $R_{rj} - R_{ij} \geq \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$ (the $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} \geq \varepsilon$ is symmetric), then $\text{opt} = R_{ij} + C_{ij} < (R_{rj} - \varepsilon) + (C_{ic} - \varepsilon) \leq 2(1 - \varepsilon)$. But is impossible if at the same time $\varepsilon < \frac{1}{2}$ and $\text{opt} \geq \frac{2-3\varepsilon}{1-\varepsilon}$, and therefore this case cannot happen.
- Finally, if $R_{rj} - R_{ij} < \varepsilon$ and $C_{ic} - C_{ij} > \varepsilon$ (the case $R_{rj} - R_{ij} > \varepsilon$ and $C_{ic} - C_{ij} < \varepsilon$ is symmetric), then by Lemma 7, the strategy profile $(i, p_j + (1-p)c)$ with $p = \frac{1-\varepsilon}{-1+2\varepsilon+2R_{ij}(1-\varepsilon)}$, is an ε -approximate Nash equilibrium with social welfare $\text{cost} \geq (1 - \frac{\varepsilon(1-\varepsilon)}{2-3\varepsilon}) \cdot \text{opt}$.

The bound $(1 - \frac{\varepsilon(1-\varepsilon)}{2-3\varepsilon})\text{opt} \geq \frac{1}{2}\text{opt}$ follows from the fact that in the interval $\varepsilon \in [0, \frac{1}{2}]$, function $1 - \frac{\varepsilon(1-\varepsilon)}{2-3\varepsilon}$ is non-increasing in ε , and hence it is minimized at $\varepsilon = \frac{1}{2}$ with the value $\frac{1}{2}$.

All required strategies can be found in polynomial-time. \square

Theorem 8 ensures that if $\text{opt} \geq \frac{2-3\varepsilon}{1-\varepsilon}$ and $\varepsilon < \frac{1}{2}$, then we can create an ε -approximate Nash equilibrium with social welfare greater than or equal to $\frac{1}{2}\text{opt}$, which is a superior upper bound to the general case from Theorem 2.

Lower bound. We can prove also a lower bound that for any $\varepsilon \leq \frac{1}{2}$, if $\text{opt} = \frac{2-3\varepsilon}{1-\varepsilon}$ then for any $\hat{\varepsilon} < \varepsilon$, we may need support of size $\Omega(\log n)$ to construct an $\hat{\varepsilon}$ -Nash equilibrium.

Theorem 9 Let $\varepsilon \leq \frac{1}{2}$. There exists a bimatrix game (R, C) in $[0, 1]^{n \times n}$ for which the maximum sum of the payoffs of the players is $\text{opt} = \frac{2-3\varepsilon}{1-\varepsilon}$, and for any $\hat{\varepsilon} < \varepsilon$, any $\hat{\varepsilon}$ -Nash equilibrium requires logarithmic support.

Proof. Let $k = \log n - 2 \log \log n$. Let (R, C) be the two payoff matrices in $[0, 1]^{n \times n}$ in which every entry is chosen independently at random from the set $\{(1, \frac{1-2\varepsilon}{1-\varepsilon}), (0, 1)\}$. We consider the row player; the case of the column player is analogous. We will show that with high probability, for any k columns in the payoff matrix of the column player, there is at least one row that has all 1s in these k columns.

Fix any set of k columns. The probability that a single row has at least one 0 in these k columns is $1 - 2^{-k}$. Thus, the probability that every row has at least one 0 in these k columns is $(1 - 2^{-k})^n$. Hence, the probability that there is a set of k columns for which all rows have at least one 0 in these k columns is at most $\binom{n}{k}(1 - 2^{-k})^n$. Since our choice of k yields $\binom{n}{k}(1 - 2^{-k})^n \ll 1$, we conclude that with high probability, for every set of k columns there is at least one row that has all 1s in these k columns. Analogous arguments hold for the column player. Let us condition on the two events and assume that for every set of k columns in the payoff matrix of the row player there is a row that has all 1s in these columns, and that for every set of k rows in the payoff matrix of the column player there is a column that has all 1s in these rows.

Let us assume that there is an $\hat{\varepsilon}$ -Nash equilibrium (x^*, y^*) for some $\hat{\varepsilon} < \varepsilon \leq \frac{1}{2}$, with the support of size k . Let $p = \sum_{\ell, m} x_\ell^* y_m^*$, where the sum is over all pairs (ℓ, m) , $1 \leq \ell, m \leq n$, such that $(R_{\ell m}, C_{\ell m}) = (1, \frac{1-2\varepsilon}{1-\varepsilon})$. p is the probability that the players play the strategy profile $(1, \frac{1-2\varepsilon}{1-\varepsilon})$ in the $\hat{\varepsilon}$ -Nash equilibrium, and $1 - p$ is the probability that the players play the strategy profile $(0, 1)$ in the $\hat{\varepsilon}$ -Nash equilibrium. The expected payoff of the row player is p , and the expected payoff of the column player is $p \left(\frac{1-2\varepsilon}{1-\varepsilon} \right) + (1 - p)$.

Since (x^*, y^*) is an $\hat{\varepsilon}$ -Nash equilibrium, $p + \hat{\varepsilon} \geq 1$ for the row player, and thus $p > 1 - \varepsilon$. Hence, the expected payoff of the column player is $p \left(\frac{1-2\varepsilon}{1-\varepsilon} \right) + (1 - p) = 1 - \frac{p\varepsilon}{1-\varepsilon} < 1 - \varepsilon < 1 - \hat{\varepsilon}$. But this contradicts the condition for the column player in the assumption that (x^*, y^*) is an $\hat{\varepsilon}$ -Nash equilibrium. \square

References

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

[Austrin et al., 2013] P. Austrin, M. Braverman, and E. Chlamtac. Inapproximability of NP-complete variants of Nash equilibrium. *Theory of Computing*, 9:117–142, 2013.

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

[Braverman et al., 2015] M. Braverman, Y. K. Ko, and O. Weinstein. Approximating the best Nash equilibrium in $n^{o(\log n)}$ -time breaks the exponential time hypothesis. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 970–982, 2015.

[Chen et al., 2009] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.

[Conitzer and Sandholm, 2008] V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.

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

[Daskalakis et al., 2009a] 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.

[Daskalakis et al., 2009b] C. Daskalakis, A. Mehta, and C. Papadimitriou. A note on approximate Nash equilibria. *Theoretical Computer Science*, 410:1581–1588, 2009.

[Daskalakis, 2013] C. Daskalakis. On the complexity of approximating a Nash equilibrium. *ACM Transactions on Algorithms*, 9(3), June 2013.

[Feder et al., 2007] T. Feder, H. Nazerzadeh, and A. Saberi. Approximating Nash equilibria using small-support strategies. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC)*, pages 352–354, 2007.

[Gilboa and Zemel, 1989] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, March 1989.

[Hazan and Krauthgamer, 2011] E. Hazan and R. Krauthgamer. How hard is it to approximate the best Nash equilibrium? *SIAM Journal on Computing*, 40(1):79–91, 2011.

[Kontogiannis et al., 2006] S. C. Kontogiannis, P. N. Panagopoulou, and P. G. Spirakis. Polynomial algorithms for approximating Nash equilibria of bimatrix games. In *Proceedings of the 2nd Workshop on Internet and Network Economics (WINE)*, pages 286–296, 2006.

[Lipton et al., 2003] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC)*, pages 36–41, 2003.

[Minder and Vilenchik, 2009] L. Minder and D. Vilenchik. Small clique detection and approximate Nash equilibria. In *Proceedings of the 13th International Workshop on Randomization and Computation*, pages 673–685, 2009.

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

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