Approximate Plutocratic and Egalitarian Nash Equilibria *

(Extended Abstract)

Artur Czumaj, Michail Fasoulakis, Marcin Jurdziński
Department of Computer Science and Centre for Discrete Mathematics and its Applications (DIMAP)
University of Warwick
{A.Czumaj, M.Fasoulakis, M.Jurdzinski}@warwick.ac.uk

ABSTRACT

We pose the problem of computing approximate Nash equilibria in bimatrix games with two simultaneous criteria of optimization: minimization of the incentives to deviate from a strategy profile and maximization of a measure of quality of the strategy profile. We consider two natural measures of quality: the maximum and the minimum of the payoffs of the two players. Maximizing the former yields *plutocratic Nash equilibria*, and maximizing the latter yields *egalitarian Nash equilibria*. We give polynomial-time algorithms that compute ε-Nash equilibria for $\varepsilon \geq \frac{3-\sqrt{5}}{2} \approx 0.382$, and that *approximate the quality of plutocratic and egalitarian Nash equilibria* to various degrees.

1. INTRODUCTION

The fundamental problem in game theory is to study properties of equilibria in noncooperative games. By the classical result of Nash [10], we know that every finite game has at least one Nash equilibrium. However, in many natural scenarios one is not only interested in finding or characterizing any equilibrium, but in fact, one wants to find equilibria that have some desirable properties, one wants to find "best" or "fairest" equilibria. While we understand quite well various properties of the set of all Nash equilibria, if we consider also the quality of a solution sought and aim at characterizing the "best" Nash equilibria, then our knowledge is limited.

The situation is even more challenging if we take into account the computational complexity of the task. For example, even in two-player games, the problem of finding an arbitrary Nash equilibrium is known to be PPAD-complete and the problem of finding a Nash equilibrium with almost any additional constraint is NP-hard [3, 6], 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, or the maximum egalitarian payoffs, or the maximum plutocratic payoffs, obtained in an exact Nash equilibrium, even in symmetric two-player games [3, Corollaries 6–8].

In recent years, the computational hardness of finding (exact) Nash equilibria led to an increasing interest in the study of *approximate Nash equilibria* (ε-Nash equilibria) that are strategy profiles in which each player has the incentive to deviate to another strategy limited by $\varepsilon > 0$. In this paper we study the fundamental problem of *computing approximate Nash equilibria* in bimatrix games, for which a certain measure of quality is close to the best value of that measure among all Nash equilibria in the game. We naturally focus on values of ε for which polynomial-time algorithms are already known for computing (approximate) ε-Nash equilibria, and we aim to develop polynomial-time algorithms for computing such approximate equilibria which are also "approximately best."

There have been some investigations aiming to efficiently find approximate Nash equilibrium that also approximate well the *social welfare* (the total payoff of the players), see Section 1.3. However in this paper, we consider equilibria focusing on other objectives: instead of aiming to maximize the social welfare, or, in short, expected payoff of the players, we will study the scenarios where one wants to maximize the smallest of the payoffs, or that one wants to maximize the largest of the payoffs. That is, we consider two natural settings to describe the notion of the "best Nash equilibrium":

- *plutocratic Nash equilibrium*—that *maximizes the expected maximum* of the payoffs of the players;
- *egalitarian Nash equilibrium*—that *maximizes the expected minimum* of the payoffs of the players.

1.1 Notation

Consider a bimatrix game (R, C) played by the *row* and the *column player*. Each player has n pure strategies at her disposal and the payoff matrix is in $[0,1]^{n\times n}$. If the row player plays a strategy i and the column player plays a strategy j, then the row player receives payoff R_{ij} and the column player receives payoff C_{ij}. The players can randomize over their pure strategies, and a *mixed strategy* x is a probability distribution over the set of all pure strategies. If the row player plays a mixed strategy x and the column player plays a mixed strategy y, then the *expected payoff of the row player* is $x^T Ry$ and the *expected payoff of the column player* is $x^T Cy$.

A *Nash equilibrium* (NE) is a strategy profile (x^*, y^*) such that: x^* is a best response to y^* (i.e., $\forall i\ {x^*}^T Ry^* \geq e_i^T Ry^*$) and y^* is a best response to x^* (i.e., $\forall i\ {x^*}^T Cy^* \geq {x^*}^T Ce_i$), where e_i is a column vector with all 0s except for a 1 in the i^{th} coordinate.

An *ε-Nash equilibrium* (ε-NE) is a strategy profile (x^*, y^*) such that: x^* is an ε-best response to y^* (i.e., $\forall i\ {x^*}^T Ry^* \geq e_i^T Ry^* - \varepsilon$) and y^* is an ε-best response to x^* (i.e., $\forall i\ {x^*}^T Cy^* \geq {x^*}^T Ce_i - \varepsilon$).

In this paper, our goal is to find ε-NE that are close to a "best" NE. We consider two very natural versions of a best NE:

A **plutocratic NE** is a NE that maximizes the following:

$$u^*_{max} \stackrel{\text{def}}{=} \max\left\{ \max\{x^T Ry, x^T Cy\} : (x, y) \text{ is a NE} \right\} .$$

An **egalitarian NE** is a NE that maximizes the following:

$$u^*_{eg} \stackrel{\text{def}}{=} \max\left\{ \min\{x^T Ry, x^T Cy\} : (x, y) \text{ is a NE} \right\} .$$

*Research partially supported by the Centre for Discrete Mathematics and its Applications (DIMAP) and by the EPSRC awards EP/D063191/1 and EP/G069034/1.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

We extend these notions to approximate Nash equilibria.

DEFINITION 1. *An* ***(additive)*** ***ρ-plutocratic ε-NE*** *is an ε-NE (x, y) such that $u^*_{max} - \max\{x^T Ry, x^T Cy\} \leq \rho$. If a strategy profile (x, y) is a 0-plutocratic ε-NE then it is a* ***plutocratic ε-NE****.*

DEFINITION 2. *An* ***(additive)*** ***ρ-egalitarian ε-NE*** *is an ε-NE (x, y) such that $u^*_{eg} - \min\{x^T Ry, x^T Cy\} \leq \rho$. A* ***relative ρ-egalitarian ε-NE*** *is an ε-NE (x, y) with $\rho \cdot u^*_{eg} \leq \min\{x^T Ry, x^T Cy\}$.*

1.2 New results

For all $\varepsilon \geq \frac{3-\sqrt{5}}{2} \approx 0.382$, we give polynomial-time algorithms for computing ε-Nash equilibria, which approximate what plutocratic and egalitarian Nash equilibria optimize.

THEOREM 3. *There is a polynomial-time algorithm that computes a plutocratic $\frac{1}{2}$-NE.*

THEOREM 4. *Let $\frac{3-\sqrt{5}}{2} \leq \varepsilon < \frac{1}{2}$. There is a polynomial-time algorithm that computes an additive $\frac{1-2\varepsilon}{1-\varepsilon}$-plutocratic ε-NE.*

THEOREM 5. *There is a polynomial-time algorithm that computes a relative $\frac{1}{2}$-egalitarian $\frac{1}{2}$-NE.*

THEOREM 6. *Let $\frac{3-\sqrt{5}}{2} \leq \varepsilon < \frac{1}{2}$. There is a polynomial-time algorithm that computes an additive $\frac{1}{2}$-egalitarian ε-NE.*

The results are obtained by first showing the existence of such approximate Nash equilibria adapting the method of Daskalakis et al. [5] and by carefully maintaining the tradeoff between the payoffs of the players. Once the existence of an appropriate approximate NE is given, we apply an LP-based approach to find them efficiently.

Observe that our analysis assumes that $\varepsilon \geq \frac{3-\sqrt{5}}{2} \approx 0.382$, which covers almost the entire range for which ε-Nash equilibria are currently known to be computable in polynomial time. Indeed, the smallest value of ε for which a polynomial-time algorithm is known that computes an ε-Nash equilibrium is $\varepsilon \approx 0.3393$ [11].

1.3 Related works

The tradeoff between minimizing incentives to deviate and the quality of the solution found has been studied in the past, though the main focus of the research was on the goal of finding approximate Nash equilibria that approximate the optimal *social welfare*. There have been two strands of these investigations, one comparing the social welfare to the optimal social welfare of the game and one comparing it to the optimal social welfare in any NE.

It is known that the social welfare of a NE can be arbitrarily far from the optimal social welfare in a bimatrix game [4]. Motivated by this result, there have been recent studies aiming to find an approximate NE with the social welfare close to the best possible (optimal social welfare of the game, not necessarily in a NE). It was shown in [4] that for every fixed $\varepsilon > 0$, every bimatrix game has an ε-NE with the social welfare at least a constant factor of the optimum. E.g., for any $\varepsilon \geq \frac{1}{2}$, there is always an ε-NE whose social welfare is at least $2\sqrt{\varepsilon} - \varepsilon \geq 0.914$ times the optimal social welfare, and this bound is tight. Furthermore, these results are algorithmic, and for every fixed $0 \leq \varepsilon^* < \varepsilon$, if one can find an ε^*-NE in polynomial time, then one can find in polynomial time an ε-NE with the social welfare at least a constant factor of the optimum.

There has been also some research aiming to efficiently find an approximate NE that *approximates well the social welfare in any NE*. (Observe that if an ε-NE approximates well the optimal social welfare, as in [4], then it also approximates well the social welfare in any NE.) It has been noted that the quasi-polynomial-time algorithm for approximating NE by Lipton et al. [9] not only finds an ε-NE for arbitrary $\varepsilon > 0$, but also the social welfare of the equilibrium found is an ε-approximation of the social welfare in any NE. In other words, in time $n^{O(\log n/\varepsilon^2)}$ we can find an arbitrarily good approximate NE with social welfare near to the best NE. (Further, we note that it is straightforward to extend the quasi-polynomial-time algorithm from [9] to find an ε-NE whose plutocratic (or egalitarian) payoff is no more than ε smaller than the maximum plutocratic (or egalitarian, respectively) payoff of a NE in the game.)

While this result raised a hope that it may be possible to extend it to design a polynomial-time ε-approximation algorithm, recent hardness results [1, 2, 8] showed that it is unlikely. Braverman et al. [2] showed that assuming the deterministic Exponential Time Hypothesis, there is a constant $\varepsilon > 0$ such that any algorithm for finding an ε-NE whose social welfare is at least $(1 - \varepsilon)$ times the optimal social welfare of a NE of the game, requires $n^{\Omega(\log n)}$ time (see also [1, 8] for related results that assume hardness of finding a large planted clique). These hardness results show that it is unlikely to obtain a polynomial-time approximation scheme that for every constants ε and ε' would construct in polynomial time an ε-NE whose social welfare is at least $(1 - \varepsilon')$ times the optimal social welfare of a NE of the game. (Austrin et al. [1] showed that many similar variants of the bicriteria approximation are similarly hard.)

On the other hand, the results from [4] can be combined with the polynomial-time algorithm finding an ε^*-NE in any bimatrix game with $\varepsilon^* \approx 0.3393$ [11], to obtain a polynomial-time algorithm that for any $\varepsilon > \varepsilon^*$ finds an ε-NE whose social welfare is at least $(1 - \sqrt{\frac{1-\varepsilon}{1-\varepsilon^*}})^2$ times the optimal social welfare of a NE of the game; for $\varepsilon \geq \frac{1}{2}$, the approximation bound can be made at least $2\sqrt{\varepsilon} - \varepsilon \geq 0.914$ times the optimal social welfare.

REFERENCES

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

[2] M. Braverman, Y. K. Ko, and O. Weinstein. Approximating the best Nash equilibrium in $n^{o(\log n)}$-time breaks the exponential time hypothesis. *26th SODA*, pp. 970–982, 2015.

[3] V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *Games Econom. Behav.*, 63:621–641, 2008.

[4] A. Czumaj, M. Fasoulakis, and M. Jurdziński. Approximate Nash equilibria with near optimal social welfare. *Proc. 24th IJCAI*, pp. 504–510, 2015.

[5] C. Daskalakis, A. Mehta, and C. Papadimitriou. A note on approximate Nash equilibria. *Theoret. Comput. Sci.*, 410(17):1581–1588, 2009.

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

[7] A. Greenwald and K. Hall. Correlated Q-learning. *Proc. 20th ICML*, pp. 242–249, 2003.

[8] E. Hazan, R. Krauthgamer. How hard is it to approximate the best Nash equilibrium? *SIAM J. Comput.*, 40:79–91, 2011.

[9] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. *Proc. 4th EC*, pp. 36–41, 2003.

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

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