Algorithmica (2010) 57: 653–667
DOI 10.1007/s00453-008-9227-6

Well Supported Approximate Equilibria in Bimatrix Games

Spyros C. Kontogiannis · Paul G. Spirakis

Received: 19 July 2007 / Accepted: 12 September 2008 / Published online: 15 October 2008

Abstract In view of the apparent intractability of constructing Nash Equilibria (NE in short) in polynomial time, even for bimatrix games, understanding the limitations of the approximability of the problem is an important challenge.

In this work we study the tractability of a notion of approximate equilibria in bimatrix games, called *well supported approximate Nash Equilibria* (SuppNE in short). Roughly speaking, while the typical notion of approximate NE demands that each player gets a payoff at least an additive term less than the best possible payoff, in a SuppNE each player is assumed to adopt with positive probability only *approximate pure best responses* to the opponent's strategy.

As a first step, we demonstrate the existence of SuppNE with small supports and at the same time good quality. This is a simple corollary of Althöfer's Approximation Lemma, and implies a *subexponential time* algorithm for constructing SuppNE of arbitrary (constant) precision.

We then propose algorithms for constructing SuppNE in win lose and normalized bimatrix games (i.e., whose payoff matrices take values from {0, 1} and [0, 1] respectively). Our methodology for attacking the problem is based on the solvability of

This work was partially supported by the 6th Framework Programme under contract 001907 (DELIS). Preliminary version of parts of this work appeared in [19, 20].

S.C. Kontogiannis (✉)
Department of Computer Science, University of Ioannina, 45110 Ioannina, Greece
e-mail: kontog@cs.uoi.gr

S.C. Kontogiannis · P.G. Spirakis
Research Academic Computer Technology Institute, N. Kazantzaki Str., Patra University Campus, 26500 Rio-Patra, Greece

S.C. Kontogiannis
e-mail: kontog@cti.gr

P.G. Spirakis
e-mail: spirakis@cti.gr

zero sum bimatrix games (via its connection to linear programming) and provides a 0.5-SuppNE for win lose games and a 0.667-SuppNE for normalized games.

To our knowledge, this paper provides the *first* polynomial time algorithms constructing ε-SuppNE for normalized or win lose bimatrix games, for any nontrivial *constant* $0 \le \varepsilon < 1$, bounded away from 1.

Keywords Bimatrix games · Well supported approximate Nash equilibria

1 Introduction

One of the most appealing concepts in game theory is the notion of a Nash equilibrium: *A collection of strategies for the players from which no player has an incentive to unilaterally deviate from her own strategy*. The extremely nice thing about Nash equilibria is that they always exist in any finite k-person game in normal form [24]. This is one of the most important reasons why Nash equilibria are considered to be the prevailing solution concept for finite games in normal form. The problem is that there can be exponentially many of them, of quite different characteristics, even for bimatrix games. Additionally, we do not know yet how to construct them in subexponential time. Therefore, k-NASH, the problem of computing an arbitrary Nash equilibrium of a finite k-person game in normal form, is a fundamental problem in algorithmic game theory and has been recognized as perhaps one of the most outstanding problems at the boundary of **P** [27]. Its complexity has been a long standing open problem, since the introduction of the pioneering (pivoting) algorithm of Lemke and Howson [22]. Unfortunately, it was recently shown by Savani and von Stengel [28] that this algorithm requires an exponential number of steps; moreover, it is also known that the smoothed complexity of the algorithm is likely to be superpolynomial [8]. It is also quite interesting that many (quite natural) refinements of k-NASH are known to be **NP**-complete problems [10, 16].

A flurry of results in the last two years has proved that k-NASH is indeed *complete problem* for the complexity class **PPAD** (introduced by Papadimitriou [26]), even for four players [12], three players [11], and two players [7]. In particular, the result of Chen and Deng [7], complemented by that of Abbott, Kane and Valiant [1], shows that 2-NASH is **PPAD**-complete even for win lose bimatrix games.

Due to the apparent hardness even of 2-NASH, approximate solutions to Nash equilibria have lately attracted the attention of the research community. There are various notions of approximate Nash equilibria one can study. The most popular one is that in which every player gets a payoff at most some (positive) constant ε less than the maximum possible payoff, (denoted here by ε-ApproxNE). An alternative, maybe less popular, but still quite interesting notion of Nash approximation, requires that each player is allowed to adopt **wpp**[1] only actions that are *approximate best responses* to the opponent's strategy, within an additive term (or, precision) ε (denoted here by ε-SuppNE). ApproxNE seem to be the dominant notion of approximate equilibria in the literature, while SuppNE is a rather new and stricter notion (e.g., see

[1] With positive probability.

[8, 9, 13]). As it will be explained later, SuppNE seem to be harder to construct. On the other hand they might be naturally motivated by the players' selfish behavior: Rather than demanding that a player adopts **wpp** only best responses against the opponent's strategy, we allow them to choose *approximate best responses* (within some *additive* precision parameter). This is in contrast to the notion of ApproxNE, in which the two players have no restriction in what kind of actions they choose to play **wpp**, so long as their payoffs are close to their best response payoffs. We would like to argue in this paper that SuppNE is a quite interesting notion of approximate Nash equilibria, due to both its mathematical challenge and also its additional property that the players are not allowed to adopt **wpp** actions that are indeed meaningless to them.

The present paper is a work trying to shed some light on this new notion of approximate equilibria. We provide a quite simple existence proof (as a simple corollary of Althöfer's Approximation Lemma [3]) of ε-SuppNE with arbitrary (constant) precision $\varepsilon > 0$, with logarithmic (in the numbers of players' actions) support sizes. We also provide (to our knowledge) the *first polynomial time algorithms* for the construction of SuppNE in normalized and win lose bimatrix games, for some constant that is clearly *away* from the trivial bound of 1.

2 Preliminaries

2.1 Mathematical Notation

For any integer $k \in \mathbb{N}$, let $[k] \equiv \{1, 2, \ldots, k\}$. We denote by $M \in F^{m\times n}$ any $m \times n$ matrix whose elements have values in some set F. We also denote by $(A, B) \in (F \times F)^{m\times n}$ any $m \times n$ matrix whose elements are *ordered pairs* of values from F. Equivalently, this structure can be seen as an ordered pair of $m \times n$ matrices $A, B \in F^{m\times n}$. Such a pair of matrices is called a **bimatrix**. A $k \times 1$ matrix is also considered to be an k**-vector**. Vectors are denoted by bold small letters (e.g., $\mathbf{x}, \mathbf{y}$). A vector having a 1 in the i-th position and 0 everywhere else is denoted by $\mathbf{e_i}$. We denote by $\mathbf{1_k}$ ($\mathbf{0_k}$) the k-vector having 1s (0s) in all its coordinates. The $k \times k$ matrix $E = \mathbf{1_k} \cdot \mathbf{1_k}^T \in \{1\}^{k\times k}$ has value 1 in all its elements. For a pair of vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we denote the *component–wise comparison* by $\mathbf{x} \geq \mathbf{y}$: $\forall i \in [n], x_i \geq y_i$. Matrices are denoted by capital letters (e.g., $A, B, C, \ldots$), and bimatrices are denoted by ordered pairs of capital letters (e.g., $(A, B), (R, C), \ldots$). For any $m \times n$ (bi)matrix M, M_j is its j-th column (as an $m \times 1$ vector), M^i is the i-th row (as a (transposed) $1 \times n$ vector) and $M_{i,j}$ is the (i, j)-th element. For any matrix $A \in \mathbb{R}^{m\times n}$, we denote by $A_{\max} \equiv \max_{(i,j)\in[m]\times[n]} A_{i,j}$ and $A_{\min} \equiv \min_{(i,j)\in[m]\times[n]} A_{i,j}$ its maximum and its minimum element, respectively.

For any integer $k \geq 1$, we denote by $\Delta_k \equiv \{\mathbf{z} \in \mathbb{R}^k : \mathbf{z} \geq \mathbf{0};\ (\mathbf{1_k})^T \mathbf{z} = 1\}$ the $(k-1)$-simplex. For any point $\mathbf{z} \in \Delta_k$, its **support** *supp*$(\mathbf{z})$ is the set of coordinates with positive value: *supp*$(\mathbf{z}) \equiv \{i \in [k] : z_i > 0\}$. For an arbitrary logical expression $\mathcal{E}$, we denote by $\mathbb{P}\{\mathcal{E}\}$ the probability of this expression being true, while $\mathbb{I}_{\{\mathcal{E}\}}$ is the indicator variable of whether this expression is true or false. For any random variable x, $\mathbb{E}\{x\}$ is its expected value (with respect to some given probability measure).

2.2 Game Theoretic Definitions and Notation

An $m \times n$ **bimatrix game** $\langle A, B\rangle$ is a 2-person game in normal form, that is determined by the bimatrix $(A, B) \in (\mathbb{R} \times \mathbb{R})^{m\times n}$ as follows: The first player (called the **row player**) has an m-element *action set* $[m]$, and the second player (called the **column player**) has an n-element *action set* $[n]$. Each row (column) of the bimatrix corresponds to a different action of the row (column) player. The row and the column player's payoffs are determined by the $m \times n$ real matrices A and B respectively. In the special case that the payoff matrices have only rational entries, we refer to a **rational bimatrix game**. If both payoff matrices belong to $[0, 1]^{m\times n}$ then we have a $[0, 1]$**-bimatrix** (aka **normalized**) **game**. When all elements of the bimatrix belong to $\{0, 1\} \times \{0, 1\}$, then we have a $\{0, 1\}$**-bimatrix** (aka **win lose**) **game**. A win lose game having (for some integer $\lambda \ge 1$) at most λ $(1, 0)$-elements per row and at most λ number $(0, 1)$-element per column of the bimatrix, is called λ**-sparse**. A bimatrix game $\langle A, B\rangle$ is called **zero sum**, if it happens that $B = -A$. In that case the game is solvable in polynomial time, since the two players' optimization problems form a primal–dual linear programming pair. In all cases of bimatrix games we assume wlog[2] that $2 \le m \le n$.

Any probability distribution on the action set $[m]$ of the row player, i.e., any point $\mathbf{x} \in \Delta_m$, is a **mixed strategy** for her. The row player then determines her action independently from the column player, according to $\mathbf{x}$. Similarly, any point $\mathbf{y} \in \Delta_n$ is a mixed strategy for the column player. Each extreme point $\mathbf{e_i} \in \Delta_m$ ($\mathbf{e_j} \in \Delta_n$) that enforces the use of the i-th row (j-th column) by the row (column) player, is called a **pure strategy** for her. Any element $(\mathbf{x}, \mathbf{y}) \in \Delta_m \times \Delta_n$ is a (mixed in general) **strategy profile** for the two players. We now define the notions of *approximate (pure) best responses*, that will help us simplify the forthcoming definitions:

Definition 1 (Approximate Best Response) Fix arbitrary constant $\varepsilon > 0$. Given that the column player adopts a strategy $\mathbf{y} \in \Delta_n$ and the payoff matrix of the row player is A, the sets of ε**-approximate (pure) best responses** are:

$$BR(\varepsilon, A, \mathbf{y}) \equiv \left\{\mathbf{x} \in \Delta_m : \mathbf{x}^T A\mathbf{y} \ge \mathbf{z}^T A\mathbf{y} - \varepsilon, \forall \mathbf{z} \in \Delta_m\right\}$$

$$PBR(\varepsilon, A, \mathbf{y}) \equiv \left\{i \in [m] : A^i\mathbf{y} \ge A^r\mathbf{y} - \varepsilon, \forall r \in [m]\right\}$$

The sets of approximate (pure) best responses of the column player against a strategy $\mathbf{x} \in \Delta_m$ of the row player, given that the column player adopts the payoff matrix B, are defined in a similar fashion:

$$BR(\varepsilon, B^T, \mathbf{x}) \equiv \left\{\mathbf{y} \in \Delta_n : \mathbf{y}^T B^T\mathbf{x} \ge \mathbf{z}^T B^T\mathbf{x} - \varepsilon, \forall \mathbf{z} \in \Delta_n\right\}$$

$$PBR(\varepsilon, B^T, \mathbf{x}) \equiv \left\{j \in [n] : B_j^T\mathbf{x} \ge B_r^T\mathbf{x} - \varepsilon, \forall r \in [n]\right\}$$

[2] Without loss of generality.

The notion of Nash equilibria was introduced by John Nash [24]. We give here the definition wrt[3] bimatrix games:

Definition 2 (Nash Equilibrium) For any bimatrix game $\langle A, B\rangle$, a profile $(\mathbf{x}, \mathbf{y}) \in \Delta_m \times \Delta_n$ is a **Nash Equilibrium** (*NE* in short) iff both players adopt 0-approximate best responses against the opponent: $\mathbf{x} \in BR(0, A, \mathbf{y})$ and $\mathbf{y} \in BR(0, B^T, \mathbf{x})$. Equivalently, $(\mathbf{x}, \mathbf{y}) \in \Delta_m \times \Delta_n$ is a NE of $\langle A, B\rangle$ iff both players adopt wpp 0-approximate only *pure best responses* against the opponent: $supp(\mathbf{x}) \subseteq PBR(0, A, \mathbf{y})$ and $supp(\mathbf{y}) \subseteq PBR(0, B^T, \mathbf{x})$. The set of profiles that are NE of $\langle A, B\rangle$ is denoted by $NE(A, B)$.

Due to the apparent difficulty in computing NE for arbitrary bimatrix games, the recent trend is to look for approximate equilibria. Two definitions of approximate equilibria that concern this paper are the following:

Definition 3 (Approximate Nash Equilibria) For any positive number $\varepsilon > 0$ and any bimatrix game $\langle A, B\rangle$, a profile $(\mathbf{x}, \mathbf{y}) \in \Delta_m \times \Delta_n$ is:

– An ε**-approximate Nash Equilibrium** (ε-ApproxNE in short) iff each player chooses an ε-approximate best response against the opponent:

$$\left[\mathbf{x} \in BR(\varepsilon, A, \mathbf{y})\right] \wedge \left[\mathbf{y} \in BR(\varepsilon, B^T, \mathbf{x})\right].$$

– An ε**-well-supported Nash Equilibrium** (ε-SuppNE in short) iff each player assigns positive probability only to ε-approximate *pure best responses* against the strategy of the opponent:

$$\begin{cases} \forall i \in [m], & x_i > 0 \Rightarrow i \in PBR(\varepsilon, A, \mathbf{y}) \\ \forall j \in [n], & y_j > 0 \Rightarrow j \in PBR(\varepsilon, B^T, \mathbf{x}) \end{cases}$$

To see the difference between the two notions of approximate equilibria, consider the well known Matching Pennies game, defined by the bimatrix:

$$(A, B) = \begin{bmatrix} (1,0) & (0,1) \\ (0,1) & (1,0) \end{bmatrix}$$

Consider the profile $(\mathbf{e_1}, \frac{1}{2} \cdot (\mathbf{e_1} + \mathbf{e_2}))$. It is easy to observe that this is actually a 0.5-ApproxNE for the two players, but it is also a (worst possible) 1-SuppNE, since the column player assigns positive probability mass to column 1, by which she gets a payoff of zero, although the pure best response (column 2) assures a payoff of one, given the row player's adopted strategy $\mathbf{e_1}$.

Clearly any NE is both a 0-ApproxNE and a 0-SuppNE. It is also straightforward to observe that every ε-SuppNE is also an ε-ApproxNE, but *not necessarily vice versa*, as was shown in the previous example. Indeed, the only thing we currently know towards this direction, is that from an arbitrary $\frac{\varepsilon^2}{8n}$-ApproxNE one can

[3] With respect to.

construct an ε-SuppNE in polynomial time [8]. It is also a folklore observation that both ε-ApproxNE and ε-SuppNE are not affected by shifting of the game by some real constant. We shall refine this observation later (see Lemma 2), to show that the addition of arbitrary row vector $\mathbf{r}^T$ to all the rows of A and the addition of arbitrary column vector $\mathbf{c}$ to all the columns of B, does not affect the SuppNE of the game.

Remark Note that both notions of approximate equilibria are defined wrt an *additive* error term ε. Although (exact) NE are known not to be affected by any positive scaling, it is important to mention that approximate notions of NE are indeed affected. Therefore, from now on we adopt the commonly used assumption in the literature (e.g., [7, 8, 13, 21, 23]) that, when referring to ε-ApproxNE or ε-SuppNE, the bimatrix game is considered to be a $[0, 1]$-bimatrix game. This is mainly done for sake of comparison of the results on approximate equilibria.

Of particular importance are the uniform points of the $(k-1)$-simplex Δ_k, considered for example in [3] wrt empirical probability distributions, and in [23] as strategies of players in bimatrix games, or even as approximation points of non-convex quadratic programs [5, 25]:

Definition 4 (Uniform Profiles) A point $\mathbf{x} \in \Delta_r$ is called a k**-uniform strategy** iff it assigns to each action a probability mass that is some multiple of $\frac{1}{k}$: $\mathbf{x} \in \Delta_r \cap \{0, \frac{1}{k}, \frac{2}{k}, \ldots, \frac{k-1}{k}, 1\}^r \equiv \Delta_r(k)$. In the special case that the only possibility for an action is to get either zero probability or $\frac{1}{k}$, we refer to a **strict** k**-uniform strategy**. We denote the space of strict k-uniform strategies by $\hat{\Delta}_r(k) \equiv \Delta_r \cap \{0, \frac{1}{k}\}^r$. A profile $(\mathbf{x}, \mathbf{y}) \in \Delta_m \times \Delta_n$ for which $\mathbf{x}$ is a (strict) k-uniform strategy and $\mathbf{y}$ is a (strict) ℓ-uniform strategy, is called a **(strict)** (k, ℓ)**-uniform** profile.

We shall finally denote by k-NASH the problem of constructing an arbitrary NE for a finite k-player game in normal form.

3 Related Work and Contribution

The computability of NE in bimatrix games has been a long standing open problem for many years. The most popular algorithm for computing NE in these games, is the algorithm of Lemke and Howson [22], which is an adaptation of Lemke's algorithm for finding solutions (if such exist) for arbitrary instances of the Linear Complementarity Problem (LCP). Unfortunately, it has been recently proved by Savani and von Stengel [28] that this pivoting algorithm may require an exponential number of steps before finding a NE, no matter which starting point is chosen. Moreover, it is well known that various (quite natural) restrictions of the problem (e.g., uniqueness, bounds on support sizes, etc.) lead to **NP**-hard problems [10, 16].

A very recent series of research papers within the last two years deal with the complexity of k-NASH. Initially [12, 17] introduced a novel reduction technique and proved that 4-NASH is **PPAD**-complete. Consequently this result was extended to 3-player games [11]. Surprisingly, Chen and Deng [7] proved the same complexity result even 2-NASH. In view of all these hardness results for the k-NASH, understanding the limitations of the (in)approximability of the problem is quite important. To our knowledge, the first result that provides ε-ApproxNE within *subexponential* time, is the work of Lipton et al. [23]. In particular, for any *constant* $\varepsilon > 0$, they prove the existence of an ε-ApproxNE for arbitrary $n \times n$ bimatrix games, which additionally is a uniform profile that has supports of size at most $\lceil \frac{\log n}{\varepsilon^2} \rceil$. This leads to a rather simple subexponential algorithm for constructing ε-ApproxNE for [0, 1]-bimatrix games, simply by checking all possible profiles with support sizes at most $\lceil \frac{\log n}{\varepsilon^2} \rceil$ for each strategy. This still remains the fastest strategy to date, for the general problem of providing ε-ApproxNE for any *constant* $\varepsilon > 0$.

With respect to certain classes of bimatrix games, [2] proved that there is a polynomial time algorithm for finding a NE in a *planar* win lose 2-player game. As for the tractability of a *Fully Polynomial Time Approximation Scheme*(FPTAS) for NE, [8] proved that providing a FPTAS for 2-NASH is also **PPAD**-complete. Namely, they proved that unless **PPAD** $\subseteq$ **P**, there is no algorithm that constructs ε-ApproxNE in time $poly(n, 1/\varepsilon)$, for any $\varepsilon = n^{-\Theta(1)}$. Moreover, they proved that unless **PPAD** $\subseteq$ **RP**, there is no algorithm that constructs a NE in time $poly(n, 1/\sigma)$, where σ is the size of the deviation of the elements of the bimatrix. This latter result essentially states that even the smoothed complexity of the algorithm of Lemke and Howson is not polynomial.

Wrt constant approximations, most of the research has focused on the notion of ApproxNE. Namely, the first two results [13, 21] recently made progress in the direction of providing the first ε-ApproxNE and ε-SuppNE for [0, 1]-bimatrix games and some *constant* $1 > \varepsilon > 0$. In particular, [13] gave a nice and simple $\frac{1}{2}$-ApproxNE for [0, 1]-bimatrix games, involving only two strategies per player. [13] made also a quite interesting connection of the problem of constructing $\frac{1+\varepsilon}{2}$-SuppNE in an arbitrary [0, 1]-bimatrix game, to that of constructing ε-SuppNE for a properly chosen win lose game of the same size. As for [21], based on linear programming techniques, they provided a $\frac{3}{4}$-ApproxNE, as well as a parameterized $\frac{2+\lambda}{4}$-ApproxNE for arbitrary [0, 1]-bimatrix games, where λ is the minimum payoff of a player at a NE of the game. Consequently, [18] provided a PTAS for ApproxNE in bimatrix games in which the sum of the two payoff matrices has fixed rank. Very recently [6, 14, 29] made progress wrt ApproxNE for constant $\varepsilon > 0$, while [15] proved that in order to get an ε-ApproxNE for some $\varepsilon < \frac{1}{2}$, one *must* allow at least logarithmic support sizes. A similar result, but for $\varepsilon < \frac{1}{4}$ had also been proved in [3]. To date, the state-of-art result [29] gives a polynomial time algorithm constructing a 0.3393-ApproxNE for normalized bimatrix games, and a 0.25-ApproxNE for win lose games.

As for SuppNE, [13] proposed an algorithm, which, under a quite interesting graph theoretic conjecture, constructs in polynomial time a non-trivial SuppNE. Unfortunately, the status of this conjecture is still unknown (it is false for some small instances of graphs). Consequently, [19] followed a graph theoretic approach for constructing in polynomial time a SuppNE. It was proved that every win lose bimatrix game either contains a PNE, or has a $(1 - \frac{2}{g})$-SuppNE which is constructed in polynomial time, where $g \geq 4$ is the girth of the Nash Dynamics graph. This result was then extended to a $(1 - \frac{1}{g})$-SuppNE for any [0, 1]-bimatrix game. Unfortunately, the quality of this SuppNE is not necessarily bounded away from the trivial bound of 1, since it depends on the size of the girth in the Nash Dynamics graph.

Concerning random [0, 1]-bimatrix games, the work of Bárány, Vempala and Vetta [4] considers the case where all the cells of the payoff matrices are (either uniform, or normal) **iid**[4] random variables in [0, 1]. They analyze a simple Las Vegas algorithm for finding an exact NE in such a game, by brute force on the support sizes, starting from smaller ones. The running time of their algorithm is $\mathcal{O}(m^2 n \log\log n + n^2 m \log\log m)$, **whp**.[5] Kontogiannis and Spirakis [19] propose a random model that is slightly more general than that of [4], and it is proved there that the strict uniform full mix $(\mathbf{1}\frac{1}{m}, \mathbf{1}\frac{1}{n})$ is an ε-SuppNE **whp**, for any $\varepsilon = \Omega(\sqrt{\log m/m})$. The proposed solution is thus an o(1)-SuppNE which is trivial to construct.

3.1 Our Contribution and Roadmap

We initially prove a result similar to that of [23], but for SuppNE this time (Sect. 4). In particular, we prove that there is a (wlog strict) uniform profile that is also an ε-SuppNE, with at most logarithmic support sizes. This directly yields a trivial $n^{\mathcal{O}(\log n/\varepsilon^2)}$ time algorithm (based on exhaustive search of supports of small size) for constructing ε-SuppNE as well, for any constant $\varepsilon > 0$. The proof of this argument is an extremely simple application of Althöfer's Approximation Lemma [3].

We then present (cf. Sect. 5) a line of attack for constructing SuppNE in bimatrix games, based on the solvability of Linear Programming. In Sect. 5.1 we construct a 0.5-SuppNE for arbitrary win lose games (Sect. 5.1). To our knowledge, this is the first constant SuppNE for arbitrary win lose games. Essentially, our technique is to split *evenly* the divergence from a properly chosen zero sum game, between the two players. Then we solve this zero sum game in polynomial time, using its direct connection to Linear Programming. The computed (exact) NE of the zero sum game we consider, is indeed proved to be also a 0.5-SuppNE for the initial win lose game.

Consequently (cf. Sect. 5.2) we propose a polynomial time algorithm for constructing a 0.667-SuppNE for any [0, 1]-bimatrix game. Again we make only one call to an LP solver.

4 Existence of Uniform SuppNE

The existence of uniform ε-ApproxNE with small support sizes is already known from [23]. In this section we report a similar result for SuppNE, which is a simple corollary of Althöfer's *Approximation Lemma* [3]:

Theorem 1 (Approximation Lemma [3]) *Assume C is any $m \times n$ matrix over the real numbers, with $0 \le C_{i,j} \le 1$, $\forall (i,j) \in [m] \times [n]$. Let $\mathbf{p} \in \Delta_m$ be any m-probability vector. Fix arbitrary positive constant $\varepsilon > 0$. Then, there exists another probability vector $\hat{\mathbf{p}} \in \Delta_m$ with $|supp(\hat{\mathbf{p}})| \le k \equiv \lceil \frac{log(2n)}{2\varepsilon^2} \rceil$, such that $|\mathbf{p}^T C_j - \hat{\mathbf{p}}^T C_j| \le \varepsilon$, $\forall j \in [n]$. Moreover, $\hat{\mathbf{p}} \in \Delta_r(k)$.*

[4] Independent, identically distributed.

[5] With high probability, i.e., with probability $1 - m^{-c}$, for some constant $c > 0$.

The following simple observation is straightforward to prove and will be quite useful in our discussion:

Claim For any real matrix $C \in \mathbb{R}^{m\times n}$ and any probability distribution $\mathbf{p} \in \Delta_m$, the empirical distribution $\hat{\mathbf{p}} \in \Delta_m$ produced by the Approximation Lemma assigns positive probabilities only to rows whose indices belong to $supp(\mathbf{p})$.

We now demonstrate how the Approximation Lemma, along with the previous claim, guarantees the existence of a uniform profile which is also a (2ε)-SuppNE with support sizes at most $\lceil \frac{\log(2n)}{2\varepsilon^2} \rceil$, for any *constant* $\varepsilon > 0$:

Theorem 2 *Fix any positive constant $\varepsilon > 0$ and an arbitrary $[0, 1]$-bimatrix game $\langle A, B\rangle$. There is at least one (k, ℓ)-uniform profile which is also a (2ε)-SuppNE for this game, where $k \le \lceil \frac{\log(2n)}{2\varepsilon^2} \rceil$ and $\ell \le \lceil \frac{\log(2m)}{2\varepsilon^2} \rceil$.*

Proof Assume any profile $(\mathbf{p}, \mathbf{q}) \in NE(A, B)$, whose existence is guaranteed for any finite game in normal form [24]. Due to the Approximation Lemma, there exists profile $(\hat{\mathbf{p}}, \hat{\mathbf{q}}) \in \Delta_m(k) \times \Delta_n(\ell)$ such that: (i) $|supp(\hat{\mathbf{p}})| \le k \equiv \lceil \log(2n)/(2\varepsilon^2)\rceil$ and $|\mathbf{p}^T B_j - \hat{\mathbf{p}}^T B_j| \le \varepsilon,\ \forall j \in [n]$. (ii) $|supp(\hat{\mathbf{q}})| \le \ell \equiv \lceil \log(2m)/(2\varepsilon^2)\rceil$, and $|A^i \mathbf{q} - A^i \hat{\mathbf{q}}| \le \varepsilon,\ \forall i \in [m]$. Therefore (also exploiting the Nash Property of $(\mathbf{p}, \mathbf{q})$ and the claim that $supp(\hat{\mathbf{p}}) \subseteq supp(\mathbf{p})$) we have:

$$\begin{aligned} \forall i \in [m], \quad \hat{p}_i > 0 &\overset{\text{/* Sampling */}}{\Longrightarrow} \quad p_i > 0 \overset{\text{/* Nash Prop. */}}{\Longrightarrow} A^i \mathbf{q} \ge A^r \mathbf{q}, \quad \forall r \in [m] \\ &\overset{\text{/* Approx. Lemma */}}{\Longrightarrow} \quad A^i \hat{\mathbf{q}} \ge A^r \hat{\mathbf{q}} - 2\varepsilon, \quad \forall r \in [m] \end{aligned}$$

The argument for the column player is identical. Therefore, we conclude that $(\hat{\mathbf{p}}, \hat{\mathbf{q}})$ is a (k, ℓ)-uniform profile that is also a (2ε)-SuppNE for $\langle A, B\rangle$. □

The following Lemma allows us to assume wlog that there is actually a *strict* (k, ℓ)-uniform profile that is also a (2ε)-SuppNE for $\langle A, B\rangle$:

Lemma 1 *Fix arbitrary constant $\varepsilon > 0$, $k = \lceil \frac{\log(2n)}{2\varepsilon^2} \rceil$ and $\ell = \lceil \frac{\log(2m)}{2\varepsilon^2} \rceil$. For any $m \times n$ $[0, 1]$-bimatrix game $\langle A, B\rangle$, there is some $(km) \times (\ell n)$ $[0, 1]$-bimatrix game $\langle A', B'\rangle$ that is polynomial-time equivalent with $\langle A, B\rangle$ with regard to (k, ℓ)-uniform ε-SuppNE. That is, there are polynomial-time computable maps $F_I : \Delta_m \mapsto \Delta_{km}$ and $F_{II} : \Delta_n \mapsto \Delta_{\ell n}$ such that for any profile $(\mathbf{p}, \mathbf{q}) \in \Delta_m \times \Delta_n$ that is also an ε-SuppNE, there is a uniquely defined profile $(\tilde{\mathbf{p}}, \tilde{\mathbf{q}}) = (F_I(\mathbf{p}), F_{II}(\mathbf{q})) \in \Delta_{km} \times \Delta_{\ell n}$, that is also an ε-SuppNE of $\langle A', B'\rangle$. Conversely, there are polynomial-time computable maps $H_I : \Delta_{km} \mapsto \Delta_m$ and $H_{II} : \Delta_{\ell n} \mapsto \Delta_n$ such that for any ε-SuppNE $(\tilde{\mathbf{p}}, \tilde{\mathbf{q}}) \in \Delta_{km} \times \Delta_{\ell n}$ of $\langle A, B\rangle$, there is a unique profile $(\hat{\mathbf{p}}, \hat{\mathbf{q}}) = (H_I(\tilde{\mathbf{p}}), H_{II}(\tilde{\mathbf{q}})) \in \Delta_m \times \Delta_n$ that is an ε-SuppNE of $\langle A, B\rangle$. Finally, the proposed mappings assign (k, ℓ)-uniform profiles of $\langle A, B\rangle$ to strict (k, ℓ)-uniform profiles of $\langle A', B'\rangle$ and vice versa.*

Proof For convenience, we first consider the intermediate $m \times (\ell n)$ bimatrix $(A'', B'') \equiv [(A, B), \ldots, (A, B)]$ (by multiplying the columns of (A, B) ℓ times, and consequently we construct the $(km) \times (\ell n)$ bimatrix (A', B') by multiplying k times each row of (A'', B''). We now consider the following mapping of strategies $\mathbf{p} \in \Delta_m$ to strategies $\tilde{\mathbf{p}} = F_I(\mathbf{p}) \in \Delta_{k\cdot m}$: $\forall i \in [m]$,

if $p_i = \frac{k_i}{k} \in \{\frac{1}{k}, \ldots, \frac{k-1}{k}, 1\}$
then $\tilde{p}_i = \tilde{p}_{m+i} = \cdots = \tilde{p}_{m(k_i-1)+i} = \frac{1}{k}$; $\tilde{p}_{mk_i+i} = \cdots = \tilde{p}_{m(k-1)+i} = 0$
else $\tilde{p}_i = p_i$; $\tilde{p}_{m+i} = \cdots = \tilde{p}_{m(k-1)+i} = 0$

Similarly, for the column player we determine a mapping of strategies $\mathbf{q} \in \Delta_n$ to strategies $\tilde{\mathbf{q}} = F_{II}(\mathbf{q}) \in \Delta_{\ell\cdot n}$: $\forall j \in [n]$,

if $q_j = \frac{\ell_j}{\ell} \in \{\frac{1}{\ell}, \ldots, \frac{\ell-1}{\ell}, 1\}$
then $\tilde{q}_j = \tilde{q}_{n+j} = \cdots = \tilde{q}_{n(\ell_j-1)+j} = \frac{1}{\ell}$; $\tilde{q}_{n\ell_j+j} = \cdots = \tilde{q}_{n(\ell-1)+j} = 0$
else $\tilde{q}_j = q_j$; $\tilde{q}_{n+j} = \cdots = \tilde{q}_{n(\ell-1)+j} = 0$

The inverse mappings are simpler to define: $\forall \tilde{\mathbf{p}} \in \Delta_{km}$, we define $\hat{\mathbf{p}} = H_I(\tilde{\mathbf{p}}) \in \Delta_m$ as follows: $\forall i \in [m]$, $\hat{p}_i = \sum_{r=0}^{k-1} \tilde{p}_{rm+i}$. Similarly, $\forall \tilde{\mathbf{q}} \in \Delta_{\ell n}$, we define $\hat{\mathbf{q}} = H_{II}(\tilde{\mathbf{q}}) \in \Delta_n$ as follows: $\forall j \in [n]$, $\hat{q}_j = \sum_{r=0}^{\ell-1} \tilde{q}_{rn+j}$.

It is clear from the definition of these transformations, that any (k, ℓ)-uniform profile for $\langle A, B\rangle$ is mapped to a unique strict (k, ℓ)-uniform profile for $\langle A', B'\rangle$, and vice versa. It is also straightforward to check that the NE approximability of profiles is preserved by this set of transformations. Thus, we conclude that $(\tilde{\mathbf{p}}, \tilde{\mathbf{q}})$ is also an ε-SuppNE of $\langle A', B'\rangle$.

Conversely, if we start from an arbitrary ε-SuppNE $(\tilde{\mathbf{p}}, \tilde{\mathbf{q}})$ of $\langle A', B'\rangle$, an almost identical reasoning will lead us to the conclusion that its inverse map $(\hat{\mathbf{p}}, \hat{\mathbf{q}}) = (H_I(\hat{\mathbf{p}}), H_{II}(\hat{\mathbf{q}})) \in \Delta_m \times \Delta_n$ is an ε-SuppNE of $\langle A, B\rangle$. The reason is that $\hat{p}_i > 0 \Rightarrow \exists\, 0 \le r \le k-1 : \tilde{p}_{rm+i} > 0$. The calculations are quite similar to the ones above. □

Therefore, in our quest for SuppNE in [0, 1]-bimatrix games $\langle A, B\rangle$, we can assume wlog the existence of a strict (k, ℓ)-uniform profile with $k, \ell \le \lceil \frac{\log(2n)}{2\varepsilon^2} \rceil$, that is also (2ε)-SuppNE, due to Lemma 1 and Theorem 2.

5 A Linear Programming Approach for Constructing SuppNE

We shall now exploit the tractability of zero sum games due to their connection to linear programming, in order to provide a 0.5-SuppNE for arbitrary win lose games and a 0.667-SuppNE for any normalized bimatrix game.

5.1 Construction of a 0.5-SuppNE for Win Lose Games

In this subsection we provide a 0.5-SuppNE for win lose games, which directly translates to a 0.75-SuppNE for arbitrary normalized games, if one exploits the nice observation of [13]. But first we remark that additive transformations (i.e., shift operations) have no effect on well supported equilibria:

Lemma 2 *Fix arbitrary* $[0,1]$-*bimatrix game* $\langle A, B\rangle$ *and any real matrices* $R, C \in \mathbb{R}^{m\times n}$, *such that* $\forall i \in [m]$, $R^i = \mathbf{r}^T \in \mathbb{R}^n$ *and* $\forall j \in [n]$, $C_j = \mathbf{c} \in \mathbb{R}^m$. *Then,* $\forall 1 > \varepsilon > 0$, $\forall(\mathbf{x}, \mathbf{y}) \in \Delta_m \times \Delta_n$, *if* $(\mathbf{x}, \mathbf{y})$ *is an* ε-*SuppNE for* $\langle A, B\rangle$ *then it is also an* ε-*SuppNE for* $\langle A + R, B + C\rangle$.

Proof The proof is rather simple and therefore it is left as an exercise. □

Our next theorem tries to construct the "right" zero sum game that would stand between the two extreme zero sum games $\langle R, -R\rangle$ and $\langle -C, C\rangle$, wrt an arbitrary win lose bimatrix game $\langle R, C\rangle$.

Theorem 3 *For arbitrary win lose bimatrix game* $\langle A, B\rangle$, *there is a polynomial time constructible profile that is a* 0.5-*SuppNE of the game.*

Proof Consider arbitrary win lose game $\langle A, B\rangle \in \{(0,0), (0,1), (1,0)\}^{m\times n}$. We have excluded the $(1,1)$-elements because, as we already know, these are trivial PNE of the game. We transform the bimatrix (A, B) into a bimatrix (R, C) by subtracting $1/2$ from all the possible payoffs in the bimatrix: $R = A - \frac{1}{2}E$ and $C = B - \frac{1}{2}E$, where $E = \mathbf{1}\cdot\mathbf{1}^T$. We already know that this transformation does not affect the quality of a SuppNE (cf. Lemma 2).

We observe that the row player would never accept a payoff less than the one achieved by the (exact) Nash equilibrium $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ of the (zero sum) game $\langle R, -R\rangle$. This is because strategy $\hat{\mathbf{x}}$ is a maximin strategy for the row player, and thus the row player can achieve a payoff of at least $\hat{V}_I \equiv \hat{\mathbf{x}}^T R\hat{\mathbf{y}}$ by adopting $\hat{\mathbf{x}}$, for any possible column that the column player chooses **wpp**. Similarly, the column player would never accept a profile $(\mathbf{x}, \mathbf{y})$ with payoff for her less than $\tilde{V}_{II} \equiv \tilde{\mathbf{x}}^T C\tilde{\mathbf{y}}$, where $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ is the (exact) NE of the zero sum game $\langle -C, C\rangle$. So, we already know that any 0-SuppNE for $\langle R, C\rangle$ should assure payoffs at least $\hat{V}_I$ and at least $\tilde{V}_{II}$ for the row and the column player respectively. Clearly, $(\hat{\mathbf{x}}, \tilde{\mathbf{y}})$ is a $\max\{\frac{1}{2} - \hat{V}_I, \frac{1}{2} - \tilde{V}_{II}\}$-ApproxNE of the game, but we cannot assure that it is a nontrivial SuppNE of $\langle R, C\rangle$. Nevertheless, inspired by this observation, we attempt to set up the right zero sum game that is somehow connected to $\langle R, C\rangle$, whose (exact) NE would provide a good SuppNE for $\langle R, C\rangle$. Therefore, we consider an arbitrary zero sum game $\langle D, -D\rangle$, for which it holds that $D = R + X \Leftrightarrow X = D - R$ and $-D = C + Y \Leftrightarrow Y = -(D + C)$ for some $m \times n$ bimatrix (X, Y). Let again $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in NE(D, -D)$. Then we have (by Nash property):

$$(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in NE(D, -D) = NE(R + X, C + Y) \Leftrightarrow \forall i, r \in [m],\ \ \forall j, s \in [n],$$

$$\begin{cases} \bar{x}_i > 0 \\ \bar{y}_j > 0 \end{cases} \Rightarrow \begin{cases} (R + X)^i\bar{\mathbf{y}} \ge (R + X)^r\bar{\mathbf{y}} \\ (C + Y)_j^T\bar{\mathbf{x}} \ge (C + Y)_s^T\bar{\mathbf{x}} \end{cases} \Rightarrow \begin{cases} R^i\bar{\mathbf{y}} \ge R^r\bar{\mathbf{y}} - [X^i - X^r]\bar{\mathbf{y}} \\ C_j^T\bar{\mathbf{x}} \ge C_s^T\bar{\mathbf{x}} - [Y_j - Y_s]^T\bar{\mathbf{x}} \end{cases}$$

Since $D = R + X = -(-D) = -(C + Y) \Leftrightarrow -Z \equiv R + C = -(X + Y)$, we can simply set $X = Y = \frac{1}{2}Z$, and then we conclude that:

$$(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in NE(D, -D) \Leftrightarrow \begin{cases} \forall i, r \in [m], & \bar{x}_i > 0 \Rightarrow R^i\bar{\mathbf{y}} \ge R^r\bar{\mathbf{y}} - \frac{1}{2}\cdot[Z^i - Z^r]\bar{\mathbf{y}} \\ \forall j, s \in [n], & \bar{y}_j > 0 \Rightarrow C_j^T\bar{\mathbf{x}} \ge C_s^T\bar{\mathbf{x}} - \frac{1}{2}\cdot[Z_j - Z_s]^T\bar{\mathbf{x}} \end{cases}$$

Observe now that, since $R, C \in \{(-\frac{1}{2}, -\frac{1}{2}), (-\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, -\frac{1}{2})\}^{m\times n}$, any row of $Z = -(R+C)$ is a vector in $\{0,1\}^n$, and any column of Z is a vector in $\{0,1\}^m$. But it holds that $\forall \hat{\mathbf{z}}, \tilde{\mathbf{z}} \in \{0,1\}^k, \forall \mathbf{w} \in \Delta_k, (\hat{\mathbf{z}} - \tilde{\mathbf{z}})^T \mathbf{w} \le \mathbf{1}^T \mathbf{w} = 1$. So we conclude that $\forall i, r \in [m], \forall \mathbf{y} \in \Delta_n, [Z^i - Z^r]\mathbf{y} \le \mathbf{1}^T \mathbf{y} = 1$, and $\forall j, s \in [n], \forall \mathbf{x} \in \Delta_m, [Z_j - Z_s]^T \mathbf{x} \le \mathbf{1}^T \mathbf{x} = 1$. Therefore we conclude that:

$$(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in NE\left(R + \frac{1}{2}Z, C + \frac{1}{2}Z\right) \Rightarrow \begin{cases} \forall i, r \in [m], & \bar{x}_i > 0 \Rightarrow R^i \bar{\mathbf{y}} \ge R^r \bar{\mathbf{y}} - \frac{1}{2} \\ \forall j, s \in [n], & \bar{y}_j > 0 \Rightarrow C_j^T \bar{\mathbf{x}} \ge C_s^T \bar{\mathbf{x}} - \frac{1}{2} \end{cases}$$

$$\Rightarrow (\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in \tfrac{1}{2}\text{-SuppNE}(R, C).$$

□

5.2 SuppNE for Normalized Bimatrix Games

Given our result on win lose games, applying a lemma of Daskalakis et al. [13, Lemma 4.6] for constructing $\frac{1+\varepsilon}{2}$-SuppNE of a $[0,1]$-bimatrix game $\langle A, B\rangle$ by any ε-SuppNE of a properly chosen win lose game of the same size, we could directly generalize our result to SuppNE for any $[0,1]$-bimatrix game:

Corollary 1 *For any* $[0,1]$*-bimatrix game* $\langle R, C\rangle$*, there is a* 0.75*-SuppNE that can be computed in polynomial time.*

The question is whether we can do better than that. Indeed we can, if we modify the rationale of the proof of Theorem 3. This way we shall get a parameterized approximation for $[0,1]$-bimatrix games. The next theorem demonstrates this parameterized method.

Theorem 4 *For any* $[0,1]$*-bimatrix game* $\langle R, C\rangle$*, and the matrix* $Z = -(R+C)$*, there is a polynomial-time constructible* $\varepsilon(\delta)$*-SuppNE for any* $0 < \delta < 1$*, where* $\varepsilon(\delta) \le \max\{\delta, 1-\delta\} \cdot (Z_{\max} - Z_{\min})$.

Proof We try to find a zero sum game that lies somehow between $\langle R, -R\rangle$ and $\langle -C, C\rangle$ and indeed provides a guaranteed SuppNE for $\langle R, C\rangle$. Therefore, we fix a constant $\delta \in (0,1)$, to be determined later. Consequently, we consider the matrix $Z = -(R+C)$. The zero sum bimatrix game $\langle R + \delta Z, -(R+\delta Z)\rangle$ is solvable in polynomial time (by use of linear programming). We denote with $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ the (exact) NE of this game. By the definition of NE, the row and the column player assign positive probability mass only to maximizing elements of the vectors $(R+\delta Z)\bar{\mathbf{y}}$ and $(-R-\delta Z)^T \bar{\mathbf{x}}$ respectively. That is: $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in NE(R+\delta Z, -(R+\delta Z)) \Leftrightarrow \forall i, r \in [m],\ \forall j, s \in [n]$,

$$\begin{cases} \bar{x}_i > 0 \\ \bar{y}_j > 0 \end{cases} \Rightarrow \begin{cases} (R+\delta Z)^i \bar{\mathbf{y}} \ge (R+\delta Z)^r \bar{\mathbf{y}} \\ (-R-\delta Z)_j^T \bar{\mathbf{x}} \ge (-R-\delta Z)_s^T \bar{\mathbf{x}} \end{cases}$$

$$\overset{/*\, C+Z=-R\, */}{\Rightarrow} \begin{cases} R^i \bar{\mathbf{y}} + \delta Z^i \bar{\mathbf{y}} \ge R^r \bar{\mathbf{y}} + \delta Z^r \bar{\mathbf{y}} \\ (C+Z-\delta Z)_j^T \bar{\mathbf{x}} \ge (C+Z-\delta Z)_s^T \bar{\mathbf{x}} \end{cases}$$

$$\Rightarrow \qquad \begin{cases} R^i\bar{\mathbf{y}} \geq R^r\bar{\mathbf{y}} - \delta[Z^i - Z^r]\bar{\mathbf{y}} \\ C_j^T\bar{\mathbf{x}} \geq C_s^T\bar{\mathbf{x}} - (1-\delta)(Z_j - Z_s)^T\bar{\mathbf{x}} \end{cases} \Rightarrow \begin{cases} R^i\bar{\mathbf{y}} \geq R^r\bar{\mathbf{y}} - \varepsilon(\delta) \\ C_j^T\bar{\mathbf{x}} \geq C_s^T\bar{\mathbf{x}} - \varepsilon(\delta) \end{cases}$$

where,

$$\begin{aligned} \varepsilon(\delta) &\equiv \max_{i,r\in[m],\, j,s\in[n],\mathbf{x}\in\Delta_m,\mathbf{y}\in\Delta_n} \left\{\delta\cdot\left[Z^i - Z^r\right]\mathbf{y},\, (1-\delta)\cdot\left[Z_j^T - Z_s^T\right]\mathbf{x}\right\} \\ &\leq \max\{\delta, (1-\delta)\}\cdot(Z_{\max} - Z_{\min}) \end{aligned} \tag{1}$$

The last inequality holds since the vectors $\mathbf{x}\in\Delta_m$ and $\mathbf{y}\in\Delta_n$ considered in the definition of $\varepsilon(\delta)$ are probability distributions over the rows and the columns of Z respectively. Obviously, for any $\delta\in[0,1]$ it holds that $(\bar{\mathbf{x}},\bar{\mathbf{y}})$ is an $\varepsilon(\delta)$-SuppNE for $\langle R, C\rangle$. □

We already know that for win lose bimatrix games $\forall i, r\in[m], \forall\mathbf{y}\in\Delta_n, [Z^i - Z^r]\mathbf{y}\leq \mathbf{1}^T\mathbf{y} = 1$. This directly yields the result of Theorem 3, if we simply set $\delta = 0.5$. But let's see what can be said about arbitrary $[0,1]$-bimatrix games:

Theorem 5 *For any* $[0,1]$*-bimatrix game, a* $\frac{2}{3}$*-SuppNE is constructible in polynomial time.*

Proof Our initial steps are in complete analogy as in the proof of Theorem 4. Therefore, we know how to construct in polynomial time an $\varepsilon(0.5)$-SuppNE, where, $\varepsilon(0.5)\leq\frac{Z_{max}-Z_{\min}}{2}$.

Let's assume now that, for some $0<\zeta<1$, we are in the seek of some ζ-SuppNE of $\langle R, C\rangle$. It is clear that the existence of any element $(R,C)_{i,j}\in[1-\zeta,1]\times[1-\zeta,1]$ would indicate a (pure) profile $(\mathbf{e_i},\mathbf{e_j})$ that is already a ζ-SuppNE. Since these are detectable in time $\mathcal{O}(nm)$, we suppose wlog that for each element of the bimatrix (R,C), it holds that $(R_{i,j}<1-\zeta)\vee(C_{i,j}<1-\zeta)$. Now, for Z we observe that $\forall(i,j)\in[m]\times[n]$,

$$\begin{array}{ll} \textbf{if} & 0\leq R_{i,j}, C_{i,j} < 1-\zeta \\ \textbf{then} & -2+2\zeta < Z_{i,j} = -(R_{i,j}+C_{i,j}) \leq 0 \\ \textbf{else} & \left(0\leq R_{i,j} < 1-\zeta\leq C_{i,j}\leq 1\right)\vee\left(0\leq C_{i,j}<1-\zeta\leq R_{i,j}\leq 1\right) \\ & \quad\Longrightarrow -2+\zeta < Z_{i,j} = -(R_{i,j}+C_{i,j}) \leq -1+\zeta \end{array}$$

So, since $0\leq\zeta<1$, we conclude that $Z\in(-2+\zeta,0]^{m\times n}$ and therefore, $\forall i,r\in[m],\forall\mathbf{y}\in\Delta_n, (Z^i - Z^r)\mathbf{y}\leq Z_{\max}-Z_{\min}<2-\zeta$, which implies that $\varepsilon(0.5) = 1-\frac{\zeta}{2}$. Since our approximation is $\max\{\zeta, 1-\frac{\zeta}{2}\}$, our best choice is to set $\zeta^* = \frac{2}{3}$ and we are done. □

6 Conclusions

In this work we have explored the existence of well supported approximate Nash equilibria (SuppNE) with small supports, as well as the tractability of constructing them, both in normalized and win lose bimatrix games.

First of all we demonstrated the existence of SuppNE of arbitrary (constant) precision, which additionally have logarithmic supports in both strategies. The proof we provide is an extremely simple application of Althöfer's Approximation Lemma, and directly leads to a subexponential (but unfortunately not polynomial) time algorithm for constructing arbitrarily precise SuppNE. We then exploited the connection of zero sum bimatrix games with linear programming, in order to get a 0.5-SuppNE for win lose games, and a 0.667-SuppNE for normalized games. As for the tractability of ApproxNE, it is already known how to construct 0.3393-ApproxNE in polynomial time [29].

The important question, whether there exists a polynomial time approximation scheme (PTAS) for the construction of either ε-ApproxNE or ε-SuppNE, for any *positive constant* $1 > \varepsilon > 0$, still remains open. It would also be interesting to find polynomial time algorithms for constructing ε-SuppNE, for some constant $0 < \varepsilon < 0.5$ for win lose games and $0 < \varepsilon < 0.667$ for the general case. Even for the case of ApproxNE, we do not currently know how to construct ε-ApproxNE for some precision $0 < \varepsilon < 0.3393$ for normalized games, or even $0 < \varepsilon < 0.25$ for win lose games.

References

1. Abbott, T., Kane, D., Valiant, P.: On the complexity of two-player win-lose games. In: Proc. of 46th IEEE Symp. on Found. of Comp. Sci. (FOCS '05), pp. 113–122 (2005)
2. Addario-Berry, L., Olver, N., Vetta, A.: A polynomial time algorithm for finding Nash equilibria in planar win-lose games. J. Graph Algorithms Appl. **11**(1), 309–319 (2007)
3. Althöfer, I.: On sparse approximations to randomized strategies and convex combinations. Linear Algebra Appl. **199**, 339–355 (1994)
4. Bárány, I., Vempala, S., Vetta, A.: Nash equilibria in random games. In: Proc. of 46th IEEE Symp. on Found. of Comp. Sci. (FOCS '05), pp. 123–131. IEEE Comput. Soc., Los Alamitos (2005)
5. Bomze, I., de Klerk, E.: Solving standard quadratic optimization problems via linear, semidefinite and copositive programming. J. Global Optim. **24**, 163–185 (2002)
6. Bosse, H., Byrka, J., Markakis, E.: New algorithms for approximate Nash equilibria in bimatrix games. In: Proc. of 3rd W. on Internet and Net. Econ. (WINE '07). LNCS, vol. 4858, pp. 17–29. Springer, Berlin (2007)
7. Chen, X., Deng, X.: Settling the complexity of 2-player Nash equilibrium. In: Proc. of 47th IEEE Symp. on Found. of Comp. Sci. (FOCS '06), pp. 261–272. IEEE Comput. Soc., Los Alamitos (2006)
8. Chen, X., Deng, X., Teng, S.H.: Computing Nash equilibria: Approximation and smoothed complexity. In: Proc. of 47th IEEE Symp. on Found. of Comp. Sci. (FOCS '06), pp. 603–612. IEEE Comput. Soc., Los Alamitos (2006)
9. Chen, X., Deng, X., Teng, S.H.: Sparse games are hard. In: Proc. of 2nd W. on Internet and Net. Econ. (WINE '06). LNCS, vol. 4286, pp. 262–273. Springer, Berlin (2006)
10. Conitzer, V., Sandholm, T.: Complexity results about Nash equilibria. In: Proc. of 18th Int. Joint Conf. on Art. Intel. (IJCAI '03), pp. 765–771. Morgan Kaufmann, San Mateo (2003)
11. Daskalakis, C., Papadimitriou, C.: Three player games are hard. Technical Report TR05-139, Electr. Coll. on Comp. Compl. (ECCC) (2005)
12. Daskalakis, C., Goldberg, P., Papadimitriou, C.: The complexity of computing a Nash equilibrium. In: Proc. of 38th ACM Symp. on Th. of Comp. (STOC '06), pp. 71–78. Assoc. Comput. Mach., New York (2006)
13. Daskalakis, C., Mehta, A., Papadimitriou, C.: A note on approximate equilibria. In: Proc. of 2nd W. on Internet and Net. Econ. (WINE '06). LNCS, vol. 4286, pp. 297–306. Springer, Berlin (2006)
14. Daskalakis, C., Mehta, A., Papadimitriou, C.: Progress in approximate Nash equilibrium. In Proc. of 8th ACM Conf. on El. Comm. (EC '07), pp. 355–358 (2007)

15. Feder, T., Nazerzadeh, H., Saberi, A.: Approximating Nash equilibria using small–support strategies. In Proc. of 8th ACM Conf. on El. Comm. (EC ’07), pp. 352–354 (2007)
16. Gilboa, I., Zemel, E.: Nash and correlated equilibria: Some complexity considerations. Games Econ. Behav. **1**, 80–93 (1989)
17. Goldberg, P., Papadimitriou, C.: Reducibility among equilibrium problems. In: Proc. of 38th ACM Symp. on Th. of Comp. (STOC ’06), pp. 61–70. Assoc. Comput. Mach., New York (2006)
18. Kannan, R., Theobald, T.: Games of fixed rank: A hierarchy of bimatrix games. In Proc. of 18th ACM-SIAM Symp. on Discr. Alg. (SODA ’07) (2007)
19. Kontogiannis, S., Spirakis, P.: Well supported approximate equilibria in bimatrix games: A graph theoretic approach. In: Proc. of 32nd Int. Symp. on Math. Found. of Comp. Sci. (MFCS ’07). LNCS, vol. 4708, pp. 596–608. Springer, Berlin (2007)
20. Kontogiannis, S., Spirakis, P.: Efficient algorithms for constant well supported approximate equilibria in bimatrix games. In: Proc. of 34th Int. Col. on Aut., Lang. and Progr. (ICALP ’07). LNCS, vol. 4596, pp. 595–606. Springer, Berlin (2007)
21. Kontogiannis, S., Panagopoulou, P., Spirakis, P.: Polynomial algorithms for approximating Nash equilibria in bimatrix games. In: Proc. of 2nd W. on Internet and Net. Econ. (WINE ’06). LNCS, vol. 4286, pp. 286–296. Springer, Berlin (2006)
22. Lemke, C.E., Howson, J.T. Jr.: Equilibrium points of bimatrix games. J. Soc. Ind. Appl. Math. **12**, 413–423 (1964)
23. Lipton, R., Markakis, E., Mehta, A.: Playing large games using simple strategies. In: Proc. of 4th ACM Conf. on El. Comm. (EC ’03), pp. 36–41. Assoc. Comput. Mach., New York (2003)
24. Nash, J.: Noncooperative games. Ann. Math. **54**, 289–295 (1951)
25. Nesterov, Y.: Random walk in a simplex and quadratic optimization over convex polytopes. Technical report, DOCUMENT DE TRAVAIL 2003/71, EUEN/CORE—Center for operations research and econometrics, CORE (2003)
26. Papadimitriou, C.: On the complexity of the parity argument and other inefficient proofs of existence. J. Comput. Syst. Sci. **48**, 498–532 (1994)
27. Papadimitriou, C.: Algorithms, games and the Internet. In Proc. of 33rd ACM Symp. on Th. of Comp. (STOC ’01), pp. 749–753 (2001)
28. Savani, R., von Stengel, B.: Hard-to-solve bimatrix games. Econometrica **74**(2), 397–429 (2006)
29. Tsaknakis, H., Spirakis, P.: An optimization approach for approximate Nash equilibria. In: Proc. of 3rd W. on Internet and Net. Econ. (WINE ’07). LNCS, vol. 4858, pp. 42–56. Springer, Berlin (2007)