

GAMES OF FIXED RANK: A HIERARCHY OF BIMATRIX GAMES

RAVI KANNAN AND THORSTEN THEOBALD

ABSTRACT. We propose and investigate a hierarchy of bimatrix games (A, B) , whose (entry-wise) sum of the pay-off matrices of the two players is of rank k , where k is a constant. We will say the rank of such a game is k . For every fixed k , the class of rank k -games strictly generalizes the class of zero-sum games, but is a very special case of general bimatrix games. We study both the expressive power and the algorithmic behavior of these games. Specifically, we show that even for $k = 1$ the set of Nash equilibria of these games can consist of an arbitrarily large number of connected components. While the question of exact polynomial time algorithms to find a Nash equilibrium remains open for games of fixed rank, we present polynomial time algorithms for finding an ε -approximation.

1. INTRODUCTION

Models of non-cooperative game theory serve to analyze situations of strategic interactions, and the concept of an *equilibrium* plays a prominent role. For some basic models of game theory, such as the *zero-sum games* introduced by von Neumann and Morgenstern [19], the natural equilibrium concepts are rather well understood. Zero-sum games can be described by a single $m \times n$ -matrix A . These games always possess an equilibrium, and the set of all equilibria (which is geometrically a polyhedral set and thus in particular connected) can be computed efficiently using linear programming (see, e.g., [6]).

Nash investigated the model of *bimatrix games* (A, B) (and more generally N -player games) [17, 18], in which the gain of one player does not necessarily agree with the loss of the other player, thus adding much expressive power to the model of zero-sum games. By Nash's results any bimatrix game has at least one equilibrium ("Nash equilibrium"), and it is well-known that under weak non-degeneracy assumptions the number of Nash equilibria in a given bimatrix game is finite and odd (see [12]). However, many questions concerning the combinatorics and the computation of Nash equilibria in general bimatrix games are still open today, thus demonstrating a wide gap between the difficulty in understanding zero-sum games and general bimatrix games. Among the open questions are the following outstanding ones.

Expressive power. The number of Nash equilibria of an $m \times n$ -bimatrix game does not only depend on m and n , but also on the entries of the payoff matrices A and B . The

Key words and phrases. Bimatrix game, Nash equilibrium, fixed rank, approximation.

Part of this work was done while the second author was a Feodor Lynen fellow of the Alexander von Humboldt Foundation at Yale University. A conference version appeared at the Symposium on Discrete Algorithms 2007.

maximum number of Nash equilibria of an $m \times n$ -game is not known, only lower and upper bounds are available; see Section 3.

Computation of equilibria. A central computational question is to understand the best ways to find one respectively all Nash equilibria. In the common computer science viewpoint, the running time of an algorithm is measured in terms of the length of the input, and a central aspect for judging the quality of an algorithm is whether it runs in polynomial time in the input size or not. In 2004, Stengel and Savani [24] have shown that the well-known Lemke-Howson algorithm [12] for finding a Nash equilibrium is not a polynomial-time algorithm.

Within the computer science community, the question whether an equilibrium can be computed in polynomial time at all has been named in 2001 by Papadimitriou to be the most concrete open question on the boundary of the class of polynomial-time solvable problems [21]. Recently, Chen and Deng [3] gave a complexity-theoretical argument that such an algorithm is not to be expected by showing that the problem of finding a Nash equilibrium in a bimatrix game belongs to the complexity class of PPAD-complete problems introduced in [20]. Moreover, together with Teng, they showed that even the problem of computing a $1/n^{\Theta(1)}$ -approximate solution remains PPAD-complete [4]. With regard to positive approximation results, Kontogiannis, Panagopoulou, and Spirakis have provided an algorithm for computing a $\frac{3}{4}$ -approximate Nash equilibrium [11]. For quasi-polynomial time approximation algorithms see Lipton, Markakis, and Mehta [13].

As a consequence of this large gap between understanding zero-sum games and general bimatrix games, it will be of interest to impose restrictions on bimatrix games which while preserving expressive power of the games may admit simple polynomial time algorithms. Lipton et al. [13] investigated games where both payoff matrices A, B are of fixed rank k . They showed that in this restricted model a Nash equilibrium can be found in polynomial time. However, for a fixed rank k , the expressive power of that model is limited; in particular, most zero-sum games do not belong to that class.

In this paper, we propose and investigate a related model based on low-rank restrictions, but which is a strict superset of the model of zero-sum games. The viewpoint we start with is that in a zero-sum game, the sum of the payoff matrices $C := A + B \in \mathbb{R}^{m \times n}$ is the zero matrix, which for our purposes we consider as a matrix of rank 0. In a general bimatrix game the rank of C can take any value up to $\min\{m, n\}$. Here, we consider the hierarchy given by the class of games in which we restrict C to be of rank at most k for some given k . We call these games *rank k -games*.

We show that the expressive power of fixed rank-games is significantly larger than that of zero-sum games. In order to provide this separation, we exhibit a sequence of $d \times d$ -games of rank 1 whose number of connected components of equilibria exceeds any given constant. Our lower bound for the maximal number of Nash equilibria of a $d \times d$ -game is linear in d . This bound is not tight.

Although the problem of finding a Nash equilibrium in a game of fixed rank is a very special case of the problem of finding a Nash equilibrium in an arbitrary bimatrix game, we do not know if there exists an exact polynomial time algorithm for this problem. Note

that the problem strictly generalizes linear programming (see, e.g., [6, Ch. 13.2] for the equivalence of linear programming and zero-sum games).

However, we provide approximation results for two approximation models. Hereby, we concentrate on deterministic approximation algorithms. Firstly, we propose a model of ε -approximation for rank k -games which is a stronger approximation model than the one used in [4, 11] (see Section 2.2). Using existing results from quadratic optimization, we show that we can approximate Nash equilibria of constant rank-games in polynomial time, with an error relative to a natural upper bound on the “maximum loss” of the game (as defined in Section 4.1).

Secondly, we present a polynomial time approximation algorithm for *relative* approximation (with respect to the payoffs in an equilibrium) provided that the matrix C has a nonnegative decomposition.

2. PRELIMINARIES

We consider an $m \times n$ -bimatrix game with payoff matrices $A, B \in \mathbb{Z}^{m \times n}$. Let

$$\mathcal{S}_1 = \left\{ x \in \mathbb{R}^m : \sum_{i=1}^m x_i = 1, x \geq 0 \right\} \quad \text{and} \quad \mathcal{S}_2 = \left\{ y \in \mathbb{R}^n : \sum_{j=1}^n y_j = 1, y \geq 0 \right\}$$

be the sets of mixed strategies of the two players, and let $\bar{\mathcal{S}}_1 = \{x \in \mathbb{R}^m : \sum_{i=1}^m x_i = 1\}$ and $\bar{\mathcal{S}}_2 = \{y \in \mathbb{R}^n : \sum_{j=1}^n y_j = 1\}$ denote the underlying affine subspaces. The first player (the row player) plays $x \in \mathcal{S}_1$ and the second player (the column player) plays $y \in \mathcal{S}_2$. The payoffs for player 1 and player 2 are $x^T A y$ and $x^T B y$, respectively.

Let $C^{(i)}$ denote the i -th row of a matrix C (as a row vector), and let $C_{(j)}$ denote the j -th column of C (as a column vector). A pair of mixed strategies (\bar{x}, \bar{y}) is a *Nash equilibrium* if

$$(2.1) \quad \bar{x}^T A \bar{y} \geq x^T A \bar{y} \quad \text{and} \quad \bar{x}^T B \bar{y} \geq \bar{x}^T B y$$

for all mixed strategies x, y . Equivalently, (\bar{x}, \bar{y}) is a Nash equilibrium if and only if

$$(2.2) \quad \bar{x}^T A \bar{y} = \max_{1 \leq i \leq m} A^{(i)} \bar{y} \quad \text{and} \quad \bar{x}^T B \bar{y} = \max_{1 \leq j \leq n} \bar{x}^T B_{(j)}.$$

2.1. Economic interpretation of low-rank games. If $A + B = 0$ then the game is called a zero-sum game. The economic interpretation of a zero-sum game is “What is good for player 1 is bad for player 2”. In order to describe game-theoretic situations which are close to that behavior, we consider a model where $a_{ij} + b_{ij}$ is a function which depends on i and j in a simple way, that is,

$$a_{ij} + b_{ij} = f(i, j)$$

where f is a simple function. If $f : \{1, \dots, m\} \times \{1, \dots, n\} \rightarrow \mathbb{Z}$ is an additive function, $f(i, j) = u_i + v_j$ with constants $u_1, \dots, u_m, v_1, \dots, v_n$, then there is an *equivalent* zero-sum game, i.e., a game having the same set of Nash equilibria. Namely, define the payoff matrices A' and B' by

$$a'_{ij} = a_{ij} - v_j, \quad b'_{ij} = b_{ij} - u_i.$$

That is, A' results from A by subtracting the column vector $(v_j, \dots, v_j)^T$ to the j -th column ($1 \leq j \leq n$) and B' results from B by subtracting the row vector (u_i, \dots, u_i) to the i -th row ($1 \leq i \leq m$). Now

$$\bar{x}^T A' \bar{y} - x^T A' \bar{y} = \bar{x}^T A \bar{y} - \sum_{j=1}^n v_j \bar{y}_j - x^T A \bar{y} + \sum_{j=1}^n v_j \bar{y}_j = \bar{x}^T A \bar{y} - x^T A \bar{y}$$

and a similar relation w.r.t. B holds. So the zero-sum game (A', B') has the same Nash equilibria as (A, B) . We remark that the case $v_j = 0$ yields the row-constant games introduced in [9].

If f is a multiplication function, $f(i, j) = u_i v_j$ with constants $u_1, \dots, u_m, v_1, \dots, v_n$, this is a rank 1-game. If f is a sum of k multiplication functions, this is a game of rank at most k .

Example 2.1. The following instantiation of the famous prisoners' dilemma is of rank 1. Here, the first and second pure strategy of each player refer to “don't confess” and “confess”, respectively. A payoff $-x$ represents an imprisonment for x years.

$$A = \begin{pmatrix} -4 & -12 \\ 0 & -9 \end{pmatrix}, \quad B = \begin{pmatrix} -4 & 0 \\ -12 & -9 \end{pmatrix}.$$

Rank-1 games also occur under the term “multiplication games” in the paper [2] by Bulow and Levin.

2.2. Approximate Nash equilibria. Whenever efficient algorithms for a certain problem are not available, a natural question is whether an *approximate* solution can be found (see, e.g., [29] as a general reference for approximation algorithms). Therefore we will also consider approximate equilibria. To define them, suppose x is not necessarily an optimal strategy for player 1 given that player 2 has played y . Then the “loss” for player 1 (from optimum) is $\max_i A^{(i)} y - x^T A y$. Similarly, if y is not optimal for player 2 given that the first player has played x , the loss for player 2 would be $\max_j x^T B_{(j)} - x^T B y$. We will mainly use the total of the two losses – i.e.,

$$\ell(x, y) = \max_i A^{(i)} y + \max_j x^T B_{(j)} - x^T (A + B) y$$

as a measure of how much (x, y) is off from equilibrium. For a matrix $X \in \mathbb{R}^{m \times n}$ let $|X| = \max_{1 \leq i \leq m, 1 \leq j \leq n} |x_{ij}|$.

Definition 2.2. Let $\varepsilon \geq 0$.

(i) A pair (x, y) of mixed strategies is a *weak ε -approximate equilibrium* if

$$(2.3) \quad \max_i A^{(i)} y - x^T A y \leq \varepsilon |A| \quad \text{and} \quad \max_j x^T B_{(j)} - x^T B y \leq \varepsilon |B|.$$

(ii) A pair (x, y) of mixed strategies is an *ε -approximate equilibrium* if

$$(2.4) \quad \ell(x, y) \leq \varepsilon |A + B|.$$

Note that if (x, y) is an ε -approximate equilibrium then it is also a weak ε -approximate equilibrium, but because of possible cancellations in the entries of $A + B$ a similar statement in the converse direction does not hold. Consequently, the term $|A + B|$ on the right hand side provides a stronger approximation model compared to the term $|A| + |B|$. Our approximation result holds for that stronger model, while in the papers [4, 11] the weak approximation model is used. Also observe that $|A + B|$ is an upper bound for the term $x^T(A + B)y$.

For a game with $A - B \neq 0$, a pair of strategies is an exact equilibrium if and only if it is a 0-approximate equilibrium. Besides the notion of “absolute” approximation in Definition 2.2, in Section 4.2 we will also consider a notion of “relative” approximation.

Lemma 2.3. *Suppose (\bar{x}, \bar{y}) is an ε -approximate equilibrium. Then*

$$(2.5) \quad x^T A \bar{y} + \bar{x}^T B y - \bar{x}^T (A + B) \bar{y} \leq \varepsilon |A + B| \text{ for any other mixed strategies } x, y.$$

Also, conversely, if a pair of mixed strategies (\bar{x}, \bar{y}) satisfies (2.5) then it is an ε -approximate equilibrium.

Proof. The proof follows from the equivalence of the statements (2.1) and (2.2). \square

2.3. Approximation of games by low-rank games. If the matrix $C = A + B$ of a bimatrix game is “close” to a game with rank k , then the game can be approximated by a rank k -game (A', B') in such a way that the Nash equilibria of the original game (A, B) remain approximate Nash equilibria in the game (A', B') .

Definition 2.4. Let (A, B) be an $m \times n$ -game and $C = A + B$. If a matrix $C' \in \mathbb{R}^{m \times n}$ satisfies $|C - C'| < \varepsilon |A + B|$ then the game (A', B') with $A' = A + \frac{1}{2}(C' - C)$, $B' = B + \frac{1}{2}(C' - C)$ ε -approximates (A, B) .

Note that $A' + B' = C'$.

Under the perturbation of the game, Nash equilibria of the original game are approximate equilibria of the perturbed game (cf. [5, Lemma 2]).

Theorem 2.5. *Let (A', B') be an ε -approximation of the game (A, B) and $\varepsilon < 1$. If (\bar{x}, \bar{y}) is a Nash equilibrium of the game (A, B) , then (\bar{x}, \bar{y}) is a 2ε -approximate Nash equilibrium for the game (A', B') .*

Proof. The loss $\ell'(\bar{x}, \bar{y})$ for (\bar{x}, \bar{y}) with respect to the perturbed game (A', B') satisfies

$$\begin{aligned} \ell'(\bar{x}, \bar{y}) &\leq \max_i (A' - A)^{(i)} \bar{y} + \max_j \bar{x}^T (B' - B)_{(j)} - \bar{x}^T (C' - C) \bar{y} \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} + \varepsilon = 2\varepsilon. \end{aligned}$$

\square

We can apply the Singular Value Decomposition (SVD) to approximate the matrix C by a matrix of some given rank k . That is, for some $C \in \mathbb{R}^{m \times n}$ and $k \in \mathbb{N}$, we want to find the matrix with rank k which is closest to C . If $C = UDV^T$ with an orthogonal matrix $U \in \mathbb{R}^{m \times m}$, an orthogonal matrix $V \in \mathbb{R}^{n \times n}$ and a matrix $D \in \mathbb{R}^{m \times n}$ with $D = \text{diag}(\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n)$, then the product UDV^T is called the *Singular Value*

Decomposition of C . The diagonal entries of D are the *singular values* of C . It is well-known that every matrix has a singular value decomposition (see, e.g., [8]):

Proposition 2.6. *Let $m \geq n$ and $C \in \mathbb{R}^{m \times n}$ be of rank greater than k . Further let $C = UDV^T$ be the singular value decomposition of C . Denoting $D' = \text{diag}(\sigma_1, \dots, \sigma_k, 0, \dots, 0)$, then the matrix $C' = UD'V^T$ is the rank k -matrix which approximates C best in the Frobenius norm.*

Approximating the matrix C of a game by a matrix of some given rank k , the approximation quality in Theorem 2.5 is then a function of the singular values of C .

3. THE EXPRESSIVE POWER OF LOW-RANK GAMES

3.1. The combinatorics of Nash equilibria. One measure for the expressive power of a game-theoretic model is the number of Nash equilibria it can have (depending on the number of strategies m, n). For simplicity, we will concentrate on the case $d := m = n$. If the Nash equilibria are not isolated, then we might count the number of connected components, but we will mainly concentrate on non-degenerate games in which there exist only a finite number of Nash equilibria.

Note that the usual definition of a non-degenerate game is slightly stronger than just requiring isolated Nash equilibria (see the discussion in [26]).

Definition 3.1. A bimatrix game is called *non-degenerate* if the number of the pure best responses of player 1 to a mixed strategy y of player 2 never exceeds the cardinality of the support $\text{supp } y := \{j : y_j \neq 0\}$ and if the same holds true for the best pure responses of player 2.

If $d \leq 4$, then a non-degenerate $d \times d$ -game can have at most $2^d - 1$ Nash equilibria, and this bound is tight (see [10, 16]). For $d \geq 5$, determining the maximal number of equilibria of a non-degenerate $d \times d$ -game is an open problem (see [25]). Based on McMullen's Upper Bound Theorem for polytopes, Keiding [10] gave an upper bound of $\Phi_{d,2d} - 1$, where

$$\Phi_{d,k} := \begin{cases} \frac{k}{k-\frac{d}{2}} \binom{k-\frac{d}{2}}{k-d} & \text{if } d \text{ even,} \\ 2 \binom{k-\frac{d+1}{2}}{k-d} & \text{if } d \text{ odd.} \end{cases}$$

A simple class of configurations which yields an exponential lower bound of $2^d - 1$ is the game where the payoff matrices of both players are the identity matrix I_d (see [23]).

The best known lower bound was given by von Stengel [25], who showed that for even d there exists a non-degenerate $d \times d$ -game having

$$(3.1) \quad \tau(d) := f(d/2) + f(d/2 - 1) - 1$$

Nash equilibria, where $f(n) := \sum_{k=0}^n \binom{n+k}{k} \binom{n}{k}$. Asymptotically, τ grows as $\tau(d) \sim 0.949 \frac{(1+\sqrt{2})^d}{\sqrt{d}}$.

If the ranks of A and B are bounded by a fixed constant, then the number of Nash equilibria is bounded polynomially in d :

Theorem 3.2. *For any $d \times d$ -bimatrix game (A, B) in which the ranks of both A and B are bounded by a fixed constant k , the number of connected components of the Nash equilibria is bounded by $\binom{d}{k+1}^2$.*

In particular, for a non-degenerate game the number of Nash equilibria is at most $\binom{d}{k+1}^2$, i.e., that number is bounded polynomially in d .

Proof. Let A and B be of rank at most k . The column space of Ay has dimension at most k . By applying Carathéodory's Theorem on the columns of Ay , it was shown in [13, Theorem 4] that for every Nash equilibrium (\bar{x}, \bar{y}) there exists a Nash equilibrium (\bar{x}, y') in which the second player plays at most $k + 1$ pure strategies with positive probability. The same argument can be used to bound the number of pure strategies which are used by player 1. It follows from that argument that there exists a continuous path from the original Nash equilibrium to the Nash equilibrium with small support.

Since for a given support of the equilibria, the set of equilibria with that support is a polyhedral set, the number of connected components of the equilibria of the game (A, B) is at most $\binom{d}{k+1}^2$. \square

Now we show that the expressive power of fixed rank-games is significantly higher than the expressive power of zero-sum games. In order to achieve this, we prove that the number of Nash equilibria of a rank 1-game can exceed any given constant and give a linear lower bound.

Theorem 3.3. *For any $d \in \mathbb{N}$ there exists a non-degenerate $d \times d$ -game of rank 1 with at least $2d - 1$ many Nash equilibria.*

The following questions remain unsolved.

Open problem 3.4. Is the maximal number of Nash equilibria for non-degenerate $d \times d$ -games of rank k smaller than the maximal number of Nash equilibria of non-degenerate $d \times d$ -games of arbitrary rank? Is the maximal number of Nash equilibria for non-degenerate $d \times d$ -games of rank k polynomially bounded in d ?

In order to prove Theorem 3.3, we use the following representation of Nash equilibria introduced by Mangasarian [15].

Definition 3.5. For an $m \times n$ -bimatrix game (A, B) , the polyhedra \bar{P} and \bar{Q} are defined by

$$(3.2) \quad \bar{P} = \{(\bar{x}, v) \in \mathbb{R}^m \times \mathbb{R} : \underbrace{\bar{x} \geq 0}_{\text{inequalities } 1, \dots, m}, \underbrace{\bar{x}^T B \leq \mathbf{1}^T v}_{\text{inequalities } m+1, \dots, m+n}, \mathbf{1}^T \bar{x} = 1\},$$

$$(3.3) \quad \bar{Q} = \{(\bar{y}, u) \in \mathbb{R}^n \times \mathbb{R} : \underbrace{A\bar{y} \leq \mathbf{1}u}_{\text{inequalities } 1, \dots, m}, \underbrace{\bar{y} \geq 0}_{\text{inequalities } m+1, \dots, m+n}, \mathbf{1}^T \bar{y} = 1\}.$$

A pair of mixed strategies $(\bar{x}, \bar{y}) \in \mathcal{S}_1 \times \mathcal{S}_2$ is a Nash equilibrium if and only if there exist $u, v \in \mathbb{R}$ such that $(\bar{x}, v) \in \bar{P}$, $(\bar{y}, u) \in \bar{Q}$ and for all $i \in \{1, \dots, m+n\}$, the i -th inequality of \bar{P} or \bar{Q} is binding. Here, u and v represent the payoffs of player 1 and player 2, respectively. For $i \in \{1, \dots, m\}$ we call the inequality $x_i \geq 0$ the i -th

nonnegativity inequality of \overline{P} , and for $j \in \{1, \dots, n\}$ we call the inequality $\overline{x}^T B_{(j)} \leq v$ the j -th best response inequality of \overline{P} . And analogously for \overline{Q} .

3.2. A class of low-rank games with arbitrarily many Nash equilibria. We construct a sequence (A_d, B_d) of $d \times d$ -games of rank 1 in which all pairs (i, i) of pure strategies ($1 \leq i \leq d$) are Nash equilibria. For convenience of notation, we will omit the index d in the notation of the game. In order to achieve the desired properties, we enforce that for every $i \in \{1, \dots, d\}$ the element a_{ii} is the maximal element in the i -th column of A and the element b_{ii} is the maximal element in the i -th row of B .

Let us begin with an auxiliary sequence of games $(\overline{A}, \overline{B})$. Let $\overline{A}, \overline{B} \in \mathbb{R}^{d \times d}$ be defined by

$$(3.4) \quad \overline{a}_{ij} = \overline{b}_{ij} = -(i - j)^2.$$

Then for every $i \in \{1, \dots, d\}$ the element \overline{a}_{ii} is the largest element in the i -th column of \overline{A} , and the element \overline{b}_{ii} is the largest element in the i -th row of \overline{B} . Expanding (3.4) shows that both \overline{A} and \overline{B} can be written as the sum of three rank 1-matrices; since $\overline{A} = \overline{B}$, it follows that the game $(\overline{A}, \overline{B})$ is a rank 3-game.

In order to transform $(\overline{A}, \overline{B})$ into a rank 1-game, we observe that adding a constant column vector to a column of A or adding a constant row vector to a row of B does not change the set of Nash equilibria. For $j \in \{1, \dots, d\}$, we add the constant vector $(2j^2, \dots, 2j^2)^T$ to the j -th column of \overline{A} , and for $i \in \{1, \dots, d\}$ we add the constant vector $(2i^2, \dots, 2i^2)$ to the i -th row of \overline{B} . Let $A, B \in \mathbb{R}^{d \times d}$ be the resulting matrices, i.e.,

$$(3.5) \quad a_{ij} = 2ij - i^2 + j^2, \quad b_{ij} = 2ij + i^2 - j^2.$$

Since $A + B = (4ij)_{i,j}$, the matrix $A + B$ is of rank 1. Note that the game (A, B) is symmetric, i.e., $A = B^T$.

Lemma 3.6. *For any mixed strategy $x \in \mathcal{S}_1$ there are at most two pure best responses for player 2. And for any mixed strategy $y \in \mathcal{S}_2$ there are at most two pure best responses for player 1.*

Proof. Let y be a mixed strategy of player 2 with support $J := \{j_1, \dots, j_k\}$. We assume that there exists a 3-element subset $I = \{i_1, i_2, i_3\} \subset \{1, \dots, d\}$ such that

$$(3.6) \quad (Ay)_{i_1} = (Ay)_{i_2} = (Ay)_{i_3} \geq (Ay)_i \text{ for all } i \notin I.$$

The equations in (3.6) imply that for all distinct $i, i' \in I$ we have

$$\sum_{j \in J} (2ij - i^2 + j^2) y_j = \sum_{j \in J} (2i'j - i'^2 + j^2) y_j,$$

which, using $\sum_{j \in J} y_j = 1$, is equivalent to $2(i - i') \sum_{j \in J} j y_j = (i^2 - i'^2)$. Hence, $2 \sum_{j \in J} j y_j = (i + i')$. The left hand side of this equation is independent of i . Therefore there cannot be more than two indices in I such that this equation is satisfied for all pairs of these indices.

The proof of the other statement is symmetric. \square

Lemma 3.7. *Each of the two polyhedra \overline{P} and \overline{Q} has $\frac{d}{6}(d^2 + 5)$ vertices, which come in two classes:*

- (1) *There exists a $j \in \{1, \dots, d\}$ such that the best response inequality of \overline{Q} with index j is binding and all nonnegativity inequalities of \overline{Q} but the one with index j are binding (d vertices).*
- (2) *There exist $j_1, j_2 \in \{1, \dots, d\}$, $j_1 < j_2$ and $i \in \{j_1, \dots, j_2 - 1\}$ such that the best response inequalities with indices i and $i + 1$ are binding and all nonnegativity inequalities except those with indices j_1, j_2 are binding (altogether $\sum_{k=1}^{d-1} k(d - k)$ vertices).*

And similarly for \overline{P} .

Proof. We consider the polyhedron \overline{Q} . By Lemma 3.6, at most two best response inequalities can be binding at a vertex of \overline{Q} .

If there is a single binding best response inequality, say, with index i , then, at a vertex v , at least $d - 1$ of the nonnegativity inequalities must be binding, and therefore there exists a single index j such that y_j is nonzero; hence $y_j = 1$. Now the condition $v \in \overline{Q}$ implies $a_{ij} \geq a_{ij'}$ for all $j' \in \{1, \dots, d\}$, and it suffices to observe that for a fixed j the value a_{ij} is maximized for $i = j$, and this defines indeed a vertex.

Now assume that there are two binding best response inequalities i_1 and i_2 with $i_1 < i_2$. Then there are at most two nonzero components of y , say y_{j_1} and y_{j_2} . We can assume that $j_1 \neq j_2$ since otherwise we are in the situation discussed before.

We claim that i_1 and i_2 are neighboring indices. Otherwise there would exist an i' with $i_1 < i' < i_2$. Now, similar to the calculations in the proof of Lemma 3.6, the property $i' + i_2 > i_1 + i_2$ implies that $2(j_1 y_{j_1} + j_2 y_{j_2}) = (i_1 + i_2) < (i' + i_2)$ and therefore

$$(Ay)_{i'} > (Ay)_{i_1} = (Ay)_{i_2}.$$

This contradicts $v \in \overline{Q}$.

Now let $i_2 = i_1 + 1$. Computing the solutions for y_{j_1} and y_{j_2} of the equations

$$\begin{aligned} 2j_1 y_{j_1} + 2j_2 y_{j_2} &= i_1 + i_2, \\ y_{j_1} + y_{j_2} &= 1 \end{aligned}$$

yields

$$y_{j_1} = \frac{2j_2 - (i_1 + i_2)}{2(j_2 - j_1)}, \quad y_{j_2} = \frac{(i_1 + i_2) - 2j_1}{2(j_2 - j_1)},$$

which in connection with $y \geq 0$ shows $j_1 \leq i_1$ and $j_2 > i_1$.

It remains to show that the stated pairs indeed define vertices. In order to prove this, we have to show that for $i' < i_1$ or $i' > i_2$ we obtain $(Ay)_{i'} < (Ay)_{i_1}$, which follows in the same way as in the case $i_1 < i' < i_2$ that was discussed before.

Now summing up over all the possibilities proves the stated number. \square

Corollary 3.8. *A pair of mixed strategies (x, y) is a Nash equilibrium of the game (A, B) if and only if $x = y = e_i$ for some unit vector e_i , $1 \leq i \leq d$, or $x = y = \frac{1}{2}(e_i + e_{i+1})$ for some $i \in \{1, \dots, d - 1\}$.*

Proof. By the characterization of the vertices in Lemma 3.7, the Nash equilibria come in two classes. If for some $i \in \{1, \dots, d\}$ both players play the i -th pure strategy, then this gives a Nash equilibrium. Moreover, for every $i \in \{1, \dots, d-1\}$, if both players only use the i -th and the $(i+1)$ -th pure strategy, there exists a Nash equilibrium. It is easy to check that in this situation, both players play both of their pure strategies with probability $\frac{1}{2}$. \square

Combining Theorem 3.3 for rank 1-games with von Stengel's result, we obtain the following lower bound for rank k -games.

Corollary 3.9. *For odd $d \geq 3$ and $k \leq d$, there exists a $d \times d$ -game of rank k with at least $\tau(k-1) \cdot (2(d-k)+1)$ Nash equilibria, where τ is defined as in (3.1). For fixed k , this sequence converges to ∞ as d tends to ∞ .*

Proof. We construct a $d \times d$ -game (A, B) of rank k with

$$A = \left(\begin{array}{c|c} A' & 0 \\ \hline 0 & A'' \end{array} \right) \quad \text{and} \quad B = \left(\begin{array}{c|c} B' & 0 \\ \hline 0 & B'' \end{array} \right)$$

where $A', B' \in \mathbb{R}^{k-1} \times \mathbb{R}^{k-1}$ define a $(k-1) \times (k-1)$ -game with $\tau(k-1)$ equilibria, which exists by von Stengel's construction. Moreover, let $A'', B'' \in \mathbb{R}^{d-k+1} \times \mathbb{R}^{d-k+1}$ define a $(d-k+1) \times (d-k+1)$ -game of rank 1 with $2(d-k+1) - 1$ equilibria based on the construction in Theorem 3.3. Then the game (A, B) is of rank k and has at least $\tau(k-1) \cdot (2(d-k)+1)$ equilibria. \square

Remark 3.10. Generalizing the construction in (3.4), for a mapping $g : \{1, \dots, d\} \rightarrow \mathbb{R}$ and a polynomial $p = \sum_{i=0}^n a_i x^i$ of degree n , the matrix $C \in \mathbb{R}^{d \times d}$ defined by

$$c_{ij} = p(g(i) - g(j))$$

has rank at most $\frac{1}{2}(n+1)(n+2)$. This follows immediately from applying the Binomial Theorem on $p(g(i) - g(j))$,

$$p(g(i) - g(j)) = \sum_{k=0}^n a_k \sum_{l=0}^k \binom{k}{l} g(i)^l (-g(j))^{k-l},$$

and observing that the rank of C is bounded by the number of terms in this expansion.

4. APPROXIMATION ALGORITHMS

For general bimatrix games, no polynomial time algorithm for ε -approximating a Nash equilibrium is known. In the model of weak approximation, Lipton et. al. [13] have provided the first subexponential algorithm for finding an approximate equilibrium (i.e., an algorithm whose running time is bounded by $2^{O(\sqrt{\mathcal{L}})}$, where \mathcal{L} is the total input length.)

Kontogiannis et. al. [11] showed that the following simple algorithm yields a $\frac{3}{4}$ -approximation algorithm in the weak approximation model. Defining the indices i_1, i_2, j_1, j_2 , by $a_{i_1, j_1} = \max_{i,j} a_{i,j}$ and $b_{i_2, j_2} = \max_{i,j} b_{i,j}$, then player 1 plays his pure strategies i_1, i_2 with probability $\frac{1}{2}$, and player 2 plays his pure strategies j_1, j_2 with probability $\frac{1}{2}$.

4.1. ε -approximating Nash equilibria of low-rank games. Here, we show the following result for our restricted class of bimatrix games.

Theorem 4.1. *Let k be a fixed constant and $\varepsilon > 0$. If $A + B$ is of rank k then an ε -approximate Nash equilibrium can be found in time $\text{poly}(\mathcal{L}, 1/\varepsilon)$, where \mathcal{L} is the bit length of the input.*

Set

$$Q = \left(\begin{array}{c|c} 0 & \frac{1}{2}(A+B) \\ \hline \frac{1}{2}(A^T+B^T) & 0 \end{array} \right) \quad \text{and} \quad z = \begin{pmatrix} x \\ y \end{pmatrix}$$

so that the quadratic form $x^T(A+B)y$ can be written as $\frac{1}{2}z^TQz$ with a symmetric matrix Q . We assume that $A+B$ has rank k for a fixed constant k ; thus Q has rank $2k$. Since the trace of the matrix Q is zero, this matrix is either the zero matrix or an indefinite matrix. Hence, in the case $Q \neq 0$ the quadratic form defined by Q is indefinite.

We use the following straightforward formulation of a Nash equilibrium as a solution of a system of linear and quadratic inequalities.

Lemma 4.2. *A pair of mixed strategies $z = \begin{pmatrix} x \\ y \end{pmatrix} \in \mathcal{S}_1 \times \mathcal{S}_2$ is a Nash equilibrium if and only if there exists an $s \in \mathbb{R}$ such that*

$$\begin{aligned} z^T Q z &\geq s \\ s &\geq (A^{(i)} | B_{(j)}^T) z \quad \text{for all } i \in \{1, \dots, m\}, j \in \{1, \dots, n\}. \end{aligned}$$

Since $z^T Q z \leq s$ in any feasible solution of this optimization problem, we have $z^T Q z = s$ for any feasible solution. Hence, the Nash equilibria are exactly the optimal solutions of the quadratic optimization problem

$$(4.1) \quad \begin{aligned} (\text{QP} :) \quad & \min s - z^T Q z \\ & s \geq (A^{(i)} | B_{(j)}^T) z \quad \text{for all } i \in \{1, \dots, m\}, j \in \{1, \dots, n\}, \\ & z \in \mathcal{S}_1 \times \mathcal{S}_2. \end{aligned}$$

This quadratic optimization problem with objective function of fixed rank can be well approximated. Namely, Vavasis has shown the following polynomial approximation result for quadratic optimization problems with compact polyhedral feasible set [27, 28].

Proposition 4.3. *Let $\min\{\frac{1}{2}x^T Q x + q^T x : Ax \leq b\}$ be a quadratic optimization problem with compact support set $P = \{x \in \mathbb{R}^n : Ax \leq b\}$, and let the rank k of Q be a fixed constant. If x^* and $x^\#$ denote points minimizing and maximizing the objective function $f(x) := \frac{1}{2}x^T Q x + q^T x$ in the feasible region, respectively, then one can find in time $\text{poly}(\mathcal{L}, 1/\varepsilon)$ a point x^\diamond satisfying*

$$f(x^\diamond) - f(x^*) \leq \varepsilon(f(x^\#) - f(x^*)),$$

where \mathcal{L} is the bit length of the quadratic problem. Such a point x^\diamond is called an ε -approximation of the quadratic problem.

To provide some intuition on the ideas underlying this statement, consider the case where the quadratic form f depends only on the first k variables. Let $C = \pi(P) \subset \mathbb{R}^k$ be the projection of the feasible set P onto the first k variables, and consider the optimization problem as an optimization problem over C . We can compute in polynomial time a weak Löwner-John pair for the convex body C in the k -dimensional space (see [14, Theorem 2.4.1]). This is a pair of concentric ellipsoids E_1, E_2 such that $E_1 \subset C \subset E_2$ and E_1 is obtained from E_2 by shrinking each dimension by $1/((k+1)\sqrt{k})$. After an affine transformation, we can assume that the Löwner-John pair for C is given by the pair (S_1, S_2) of spheres centered in the origin with radii 1 and $(k+1)\sqrt{k}$, respectively. These two spheres are contained in the box $B = [-(k+1)\sqrt{k}, (k+1)\sqrt{k}]^k$. Dividing each of these k intervals into $\lceil \frac{k(k+1)}{\sqrt{\varepsilon}} \rceil$ many pieces establishes a grid on B . Approximating the quadratic form f on the each grid cell by its linear Taylor approximation yields the desired approximation result.

We can now prove our approximation result on Nash equilibria.

PROOF OF THEOREM 4.1. The feasible region of the quadratic program (4.1) is unbounded. Since the value of $z^T Q z$ is at most $|A+B|$ for any feasible solution z and since the objective value for a Nash equilibrium is 0, we can add the constraint $s \leq |A+B|$ to (4.1), which makes the feasible region compact. Denote the resulting quadratic optimization problem by QP' and recall that the approximation ratio of the quadratic program depends on the maximum objective value in the feasible region.

By Proposition 4.3, we can compute in polynomial time an ε -approximation (z^\diamond, s^\diamond) with $z^\diamond = (x^\diamond, y^\diamond)$ of QP'. Since the optimal value of QP' is 0, we have

$$s^\diamond - (z^\diamond)^T Q z^\diamond = f(z^\diamond, s^\diamond) \leq \varepsilon f(z^\#, s^\#) \leq \varepsilon |A+B|.$$

Hence, (x^\diamond, y^\diamond) is an ε -approximate Nash equilibrium of the game (A, B) . \square

Remark 4.4. The proof in [27] computes an LDL^T factorization of the matrix Q defining the quadratic form and then constructs a sufficiently fine grid in the fixed-dimensional space. Since the quadratic form $x^T Q y$ is bilinear, we can also directly apply an LDU^T factorization on the matrix of the bilinear form.

4.2. Relative approximation in case of a nonnegative decomposition. The right hand side in Definition 2.2 of an approximate Nash equilibrium depends only on ε and on $|A+B|$. Since different equilibria in the same game can differ strongly in their payoffs, we introduce a notion of *relative approximation* with respect to a Nash payoff which takes into account these differences.

Consider the quadratic problem (4.1). In a Nash equilibrium $(x, y) \in \mathcal{S}_1 \times \mathcal{S}_2$ there exists an $s \in \mathbb{R}$ such that (x, y, s) is a feasible solution to (4.1); in this situation s coincides with the sum of the payoffs of the two players. In the relative approximation, we aim at finding pairs of strategies (x, y) for which there exists an $s \in \mathbb{R}$ such that (x, y, s) is feasible and

$$s - x^T(A+B)y \leq \rho s$$

for some approximation ratio ρ . Using our notion of loss, by observing $s = \max_i A^{(i)}y + \max_j x^T B_{(j)}$ for an optimally chosen s , this means

$$\ell(x, y) \leq \rho(\max_i A^{(i)}y + \max_j x^T B_{(j)}).$$

We provide an efficient approximation algorithm for the case that $C = A + B$ has a known decomposition of the form

$$(4.2) \quad C = \sum_{i=1}^k u^{(i)}(v^{(i)})^T$$

with non-negative vectors $u^{(i)}$ and $v^{(i)}$.

Theorem 4.5. *If C has a known decomposition of the form (4.2) then for any given $\varepsilon > 0$ a relatively approximate Nash equilibrium with approximation ratio $1 - \frac{1}{(1+\varepsilon)^2}$ can be computed in time $\text{poly}(\mathcal{L}, 1/\log(1+\varepsilon))$, where \mathcal{L} is the bit length of the input.*

Let $z_i = x^T \cdot u^{(i)}$, $w_i = (v^{(i)})^T \cdot y$. We put a grid on each of the z_i and on each of the w_i in a geometric progression: denoting by

$$(z_i)_{\min} = \min_{x \in \mathcal{S}_1} x^T \cdot u^{(i)} \quad \text{and} \quad (z_i)_{\max} = \max_{x \in \mathcal{S}_1} x^T \cdot u^{(i)}$$

the minimum and the maximum possible value for z_i , we partition the interval $[(z_i)_{\min}, (z_i)_{\max}]$ into the intervals $[(z_i)_{\min}, (1+\varepsilon)(z_i)_{\min}]$, $[(1+\varepsilon)(z_i)_{\min}, (1+\varepsilon)^2(z_i)_{\min}]$, and so on. And analogously for the w_i .

For every cell we construct a linear program which ‘‘approximates’’ the quadratic program (4.1). Let the intervals of a grid cell be $[\alpha_i, (1+\varepsilon)\alpha_i]$ and $[\beta_i, (1+\varepsilon)\beta_i]$, i.e.,

$$\begin{aligned} \alpha_i &\leq z_i \leq (1+\varepsilon)\alpha_i, \\ \beta_i &\leq w_i \leq (1+\varepsilon)\beta_i. \end{aligned}$$

Then for any pair of strategies $(x, y) \in \mathcal{S}_1 \times \mathcal{S}_2$ falling into that cell, the quadratic form $x^T C y$ satisfies

$$(4.3) \quad \sum_{i=1}^k \alpha_i \beta_i \leq x^T C y \leq (1+\varepsilon)^2 \sum_{i=1}^k \alpha_i \beta_i,$$

where the left inequality uses that all the values in the decomposition are nonnegative. For the grid cell, we consider the linear program

$$\begin{aligned} \min s - \sum_{i=1}^k \alpha_i \beta_i \\ \alpha_i &\leq x^T \cdot u^{(i)} \leq (1+\varepsilon)\alpha_i, \\ \beta_i &\leq (v^{(i)})^T \cdot y \leq (1+\varepsilon)\beta_i, \\ s &\geq \left(A^{(i)} \mid B_{(j)}^T \right) z \quad \text{for all } i \in \{1, \dots, m\}, j \in \{1, \dots, n\}, \\ (x, y) &\in \mathcal{S}_1 \times \mathcal{S}_2, s \in \mathbb{R}. \end{aligned}$$

In at least one of the cells there exists a Nash equilibrium. The linear program corresponding to that cell yields a solution with

$$(4.4) \quad \sum_{i=1}^k \alpha_i \beta_i \leq s \leq (1 + \varepsilon)^2 \left(\sum_{i=1}^k \alpha_i \beta_i \right).$$

Hence, by the left inequality in (4.3) and the right inequality in (4.4) we have

$$x^T C y \geq \sum_{i=1}^k \alpha_i \beta_i \geq \frac{s}{(1 + \varepsilon)^2}.$$

We conclude

$$s - x^T C y \leq s \left(1 - \frac{1}{(1 + \varepsilon)^2} \right),$$

which shows Theorem 4.5.

5. CONCLUSION AND FUTURE RESEARCH

We have introduced the model of games of fixed rank and presented various combinatorial and algorithmic results on games of fixed rank. Both from the viewpoint of game theory and from the viewpoint of generalizations of linear programming, we think that this model has much to offer and suggest further investigation.

From the viewpoint of game theory, it provides a flexible hierarchy between zero-sum games and general bimatrix games. As mentioned above, some fundamental questions, such as the question whether an exact Nash equilibrium in a game of fixed rank can be found in polynomial time, remain open, and deserve further algorithmic study.

From the computational viewpoint, besides the deterministic algorithms, a central issue is to understand randomized approximation algorithms for games of fixed rank. Current work aims at generalizing these optimization techniques for low-rank games to more general optimization problems with some suitable “low-rank” structure.

REFERENCES

- [1] I. Bárány, S. Vempala, and A. Vetta. Nash equilibria in random games. In *Proc. 46th IEEE Foundations of Computer Science* (Pittsburgh, PA), 123–131, 2005.
- [2] J. Bulow and J. Levin. Matching and price competition. *American Economic Review* 96:652–668, 2006.
- [3] X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *Proc. 47th IEEE Foundations of Computer Science* (Berkeley, CA), 261–272, 2006.
- [4] X. Chen, X. Deng, and S.-H. Teng. Computing Nash equilibria: Approximation and smoothed complexity. In *Proc. 47th IEEE Foundations of Computer Science* (Berkeley, CA), 603–612, 2006.
- [5] V. Conitzer and T. Sandholm. A technique for reducing normal-form games to compute a Nash equilibrium. *Proc. Autonomous Agents and Multi Agent Systems (AAMAS)*, Hakodate, Japan, 537–544, 2006.
- [6] G.B. Dantzig. *Linear Programming and Extensions*. Princeton Univ. Press, Princeton, NJ, 1963.
- [7] C. Daskalakis, P.W. Goldberg, and C.H. Papadimitriou. The complexity of computing a Nash equilibrium. *Proc. 38th ACM Symp. Theory of Computing*, (Seattle, WA), 71–78, 2006.
- [8] R.A. Horn and C.R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

- [9] K. Isaacson and C.B. Millham. On a class of Nash-solvable bimatrix games and some related Nash subsets. *Naval. Res. Logist. Quarterly* 23:311–319, 1980.
- [10] H. Keiding. On the maximal number of Nash equilibria in an $n \times n$ bimatrix game. *Games Econom. Behavior* 21:148–160, 1997.
- [11] S.C. Kontogiannis, P.N. Panagopoulou, and P.G. Spirakis. Polynomial algorithms for approximating Nash equilibria of bimatrix games. In *Proc. Workshop on Internet and Network Economics* (Patra, Greece), Lecture Notes in Computer Science, Vol. 4286, 286–296, Springer-Verlag, 2006.
- [12] C.E. Lemke and J.T. Howson. Equilibrium points of bimatrix games. *J. Soc. Indust. Appl. Math.* 12:413–423, 1964.
- [13] R.J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proc. ACM Conf. on Electronic Commerce* (San Diego, CA), 36–41, 2003.
- [14] L. Lovász. *An Algorithmic Theory of Numbers, Graphs, and Convexity*. SIAM, Philadelphia, PA, 1986.
- [15] O.L. Mangasarian. Equilibrium points of bimatrix games. *J. Soc. Industr. Appl. Math.* 12:778–780, 1964.
- [16] A. McLennan and I.-U. Park. Generic 4×4 two person games have at most 15 Nash equilibria. *Games Econom. Behavior* 26:111–130, 1997.
- [17] J. Nash. Equilibrium points in n -person games. *Proc. Amer. Math. Soc.* 36:48–49, 1950.
- [18] J. Nash. Non-cooperative games. *Ann. Math.* 54:286–295, 1951.
- [19] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.
- [20] C.H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.* 48:498–532.
- [21] C.H. Papadimitriou. Algorithms, games and the Internet. In *Proc. 33rd ACM Symp. Theory of Computing*, Chersonissos, Kreta, 749–753, 2001.
- [22] C.H. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In *Symp. on Discrete Algorithms* (Vancouver, BC), 82–91, 2005.
- [23] T. Quint and M. Shubik. A bound on the number of Nash equilibria in a coordination game. *Economic Letters* 77:323–327, 2002.
- [24] R. Savani and B. von Stengel. Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In *Proc. 45th IEEE Foundations of Computer Science* (Rome), 258–257, 2004.
- [25] B. von Stengel. New maximal numbers of equilibria in bimatrix games. *Discrete Comput. Geom.* 21:557–568, 1999.
- [26] B. von Stengel. Computing equilibria for two-person games. In R.J. Aumann, S. Hart (eds.), *Handbook of Game Theory*, North-Holland, Amsterdam, 2002.
- [27] S. Vavasis. Approximation algorithms for indefinite quadratic programming. Technical Report 91-1228, Dept. of Computer Science, Cornell University (Ithaca, NY), 1991.
- [28] S. Vavasis. Approximation algorithms for indefinite quadratic programming. *Math. Program.* 57:279–311, 1992.
- [29] V.V. Vazirani. *Approximation Algorithms*. Springer-Verlag, Berlin, 2001.

R. KANNAN: DEPT. OF COMPUTER SCIENCE, YALE UNIVERSITY, P.O. BOX 208285, NEW HAVEN, CT 06520–8285, USA

E-mail address: kannan@cs.yale.edu

T. THEOBALD: FACHBEREICH INFORMATIK UND MATHEMATIK, J.W. GOETHE-UNIVERSITÄT, D-60054 FRANKFURT AM MAIN, GERMANY

E-mail address: theobald@math.uni-frankfurt.de