Games of fixed rank: A hierarchy of bimatrix games*

Ravi Kannan† Thorsten Theobald‡

Abstract

We propose and investigate bimatrix games, 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 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 can provide a deterministic polynomial time algorithm for finding an ε-approximation (whose running time is polynomial in $\frac{1}{\varepsilon}$) as well as a randomized polynomial time approximation algorithm (whose running time is similar), but which offers the possibility of finding an exact solution in polynomial time if a conjecture is valid. The latter algorithm is based on a new application of random sampling methods to quadratic optimization problems of fixed rank.

1 Introduction

Models of non-cooperative game theory serve to analyze situations of strategic interactions. Driven by current developments in auction theory as well as in equilibria models for the internet, the basic model of a *Nash equilibrium* has recently attracted much attention (see for example the survey by Papadimitriou [19] or the recent papers [1, 3, 7, 20, 22]).

In [18], von Neumann and Morgenstern introduced the model of *zero-sum games*, which are described by a single $m \times n$-matrix A. These games always possess an equilibrium, and the set of all equilibria (which is 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) [16, 17], 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. Concerning the question whether an equilibrium can be computed in polynomial time (named by Papadimitriou to be the most concrete open question on the boundary of **P** [19]), Chen and Deng recently showed that the problem is **PPAD**-complete [3] and (together with Teng [4]) that the problem of computing a $1/n^{\Theta(1)}$-approximate Nash equilibrium remains **PPAD**-complete. 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].

Thus, it will be of interest to impose restrictions on bimatrix games which while preserving expressive power of the games may admit polynomial time algorithms or polynomial time approximation schemes. Recently, 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*.

Our contributions. 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.

*Part of this work was done while the second author was a Feodor Lynen fellow of the Alexander von Humboldt Foundation at Yale University.

†Dept. of Computer Science, Yale University.

‡Dept. of Mathematics, Technische Universität Berlin.

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).

From the algorithmic point of view, we provide approximation results for two approximation models. 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 results from quadratic optimization, we show that we can deterministically 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). The running time of this algorithm is polynomial in $\frac{1}{\varepsilon}$.

Combining ideas of random sampling and quadratic optimization, we then provide a randomized approximation algorithm for certain quadratic optimization problems, which yields a randomized approximation algorithm for the Nash problem. The running time of the randomized algorithm is similar to the deterministic algorithm, but it has the possibility of finding an exact solution in polynomial time if a conjecture is valid.

Finally, we present a polynomial time 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

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

be the sets of mixed strategies of the two players, and let $\overline{\mathcal{S}}_1 = \{x \in \mathbb{R}^m : \sum_{i=1}^m x_i = 1\}$ and $\overline{\mathcal{S}}_2 = \{y \in \mathbb{R}^n : \sum_{j=1}^n y_j = 1\}$ 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 $(\overline{x}, \overline{y})$ is a *Nash equilibrium* if

$$\overline{x}^T A \overline{y} \;\geq\; x^T A \overline{y} \quad \text{and} \quad \overline{x}^T B \overline{y} \;\geq\; \overline{x}^T B y \tag{2.1}$$

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

$$\begin{aligned} \overline{x}^T A \overline{y} &= \max_{1 \leq i \leq m} A^{(i)} \overline{y} \\ \text{and } \overline{x}^T B \overline{y} &= \max_{1 \leq j \leq n} \overline{x}^T B_{(j)}\,. \end{aligned} \tag{2.2}$$

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 simple way, that is,

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

where f is a simple function. If $f : \{1, \ldots, m\} \times \{1, \ldots, n\} \to \mathbb{Z}$ is an additive function, $f(i,j) = u_i + v_j$ with constants $u_1, \ldots, u_m, v_1, \ldots, v_n$, 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 $a'_{ij} \;=\; a_{ij} - v_j$, $b'_{ij} \;=\; b_{ij} - u_i$. That is, A' results from A by subtracting the column vector $(v_j, \ldots, v_j)^T$ to the j-th column $(1 \leq j \leq n)$ and B' results from B by subtracting the row vector $(u_i, \ldots, u_i)$ to the i-th row $(1 \leq i \leq m)$. Now

$$\begin{aligned} \overline{x}^T A' \overline{y} - x^T A' \overline{y} &= \overline{x}^T A \overline{y} - \sum_{j=1}^{n} v_j \overline{y}_j - x^T A \overline{y} + \sum_{j=1}^{n} v_j \overline{y}_j \\ &= \overline{x}^T A \overline{y} - x^T A \overline{y} \end{aligned}$$

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 [8].

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

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

2.2 Approximate Nash equilibria. We 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

$\max_j x^T B_{(j)} - x^T By$. We will 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 $|X| = \max_{1\le i\le m, 1\le j\le n} |x_{ij}|$.

DEFINITION 2.1. *For $\varepsilon \ge 0$, a pair (x,y) of mixed strategies is an* ε-approximate equilibrium *if*

$$\ell(x,y) \;\le\; \varepsilon|A+B|\,. \tag{2.3}$$

Note that the term $|A+B|$ on the right hand side provides a stronger approximation model compared to the term $|A|+|B|$. The latter one is essentially the approximation model used in the papers [4, 11]. 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 2.1, in Section 4.3 we will also consider a notion of "relative" approximation.

LEMMA 2.1. *Suppose $(\overline{x},\overline{y})$ is an ε-approximate equilibrium. Then*

$$x^T A\overline{y} + \overline{x}^T By - \overline{x}^T(A+B)\overline{y} \;\le\; \varepsilon|A+B| \tag{2.4}$$

for any other mixed strategies x,y. Also, conversely, if a pair of mixed strategies $(\overline{x},\overline{y})$ satisfies (2.4) *then it is an ε-approximate equilibrium.*

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

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 2.2. *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 2.1. *Let (A',B') be an ε-approximation of the game (A,B) and $\varepsilon < 1$. If $(\overline{x},\overline{y})$ is a Nash equilibrium of the game (A,B), then $(\overline{x},\overline{y})$ is a 2ε-approximate Nash equilibrium for the game (A',B').*

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

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

We can apply the Singular Value Decomposition (SVD) to approximate the matrix C by a matrix of some given rank k. The approximation factor in Theorem 2.1 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. Here, 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 $\operatorname{supp} y := \{j \,:\, y_j \neq 0\}$ and if the same holds true for the best pure responses of player 2 (see [24]).

If $d \le 4$, then a non-degenerate $d\times d$-game can have at most 2^d-1 Nash equilibria, and this bound is tight (see [10, 15]). For $d \ge 5$, determining the maximal number of equilibria of a non-degenerate $d\times d$-game is an open problem (see [23]). 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 [21]).

The best known lower bound was given by von Stengel [23], who showed that for even d there exists a non-degenerate $d\times d$-game having

$$\tau(d) := f(d/2) + f(d/2-1) - 1 \tag{3.5}$$

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 3.1. *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 $(\overline{x}, \overline{y})$ there exists a Nash equilibrium $(\overline{x}, y')$ 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 game (A, B) is at most $\binom{d}{k+1}^2$.

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 show 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 3.2. *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.*

Proof of Theorem 3.2 (Sketch). 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 $(1 \le i \le d)$ are Nash equilibria. For convenience of notation, we omit the index d in the notation of the game. Let $A, B \in \mathbb{R}^{d \times d}$ be defined by

$$(3.6) \quad a_{ij} = 2ij - i^2 + j^2 , \qquad b_{ij} = 2ij + i^2 - j^2 .$$

Since $A + B = (4ij)_{i,j}$, the matrix $A + B$ is of rank 1.

By an explicit analysis of this construction, it can be shown that the game is non-degenerate and that 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, $1 \le i \le d$, or $x = y = \frac{1}{2}(e_i + e_{i+1})$ for some $i \in \{1, \dots, d-1\}$. Hence, there are $2d - 1$ Nash equilibria. □

The following questions remain unsolved.

OPEN PROBLEM 3.1. 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?

Combining Theorem 3.2 for rank 1-games with von Stengel's result, we obtain the following lower bound for rank k-games.

COROLLARY 3.1. *For odd $d \ge 3$ and $k \le d$, 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 (3.5). 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 3.2. Then the game (A, B) is of rank k and has at least $\tau(k-1) \cdot (2(d-k)+1)$ equilibria.

4 Approximation algorithms

4.1 Deterministic ε-approximation of Nash equilibria of low rank games. For general bimatrix games, no polynomial time algorithm for ε-approximating a Nash equilibrium is known. In a related model to ours, [13] has provided the first subexponential algorithm for finding an approximate equilibrium.

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* 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 we 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 4.1. *A pair of mixed strategies* $z = \binom{x}{y} \in \mathcal{S}_1 \times \mathcal{S}_2$ *is a Nash equilibrium if and only if there exists an* $s \in \mathbb{R}$ *such that*

$$\begin{aligned} z^TQz &\geq s \\ s &\geq \left(A^{(i)} \,|\, B^T_{(j)}\right) z \quad \textit{for all} \quad i \in \{1,\dots,m\}, \\ & \qquad\qquad\qquad\qquad\quad j \in \{1,\dots,n\}. \end{aligned}$$

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

$$\begin{aligned} &\min s - z^TQz \\ &\quad s \;\geq\; \left(A^{(i)} \,|\, B^T_{(j)}\right) z \quad \forall i \in \{1,\dots,m\}, \\ &\qquad\qquad\qquad\qquad\quad j \in \{1,\dots,n\}, \\ &\quad z \;\in\; \mathcal{S}_1 \times \mathcal{S}_2 \,. \end{aligned} \tag{4.7}$$

Vavasis has shown the following polynomial approximation result for quadratic optimization problems with compact polyhedral feasible set [25, 26].

PROPOSITION 4.1. *Let* $\min\{\frac{1}{2}x^TQx + q^Tx \,:\, Ax \leq b\}$ *be a quadratic optimization problem with compact support set* $\{x \in \mathbb{R}^n \,:\, Ax \leq b\}$*, and let the rank* k *of* Q *be a fixed constant. If* x^* *and* $x^{\#}$ *denote points minimizing and maximizing the objective function* $f(x) := \frac{1}{2}x^TQx + q^Tx$ *in the feasible region, respectively, then one can find in time* $\mathrm{poly}(\mathcal{L}, 1/\varepsilon)$ *a point* $x^\diamondsuit$ *satisfying*

$$f(x^\diamondsuit) - f(x^*) \;\leq\; \varepsilon(f(x^{\#}) - f(x^*))\,,$$

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

PROOF OF THEOREM 4.1. The feasible region of the quadratic program (4.7) is unbounded. Since the value of z^TQz 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 (4.7), 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 4.1, we can compute in polynomial time an ε-approximation $(z^\diamondsuit, s^\diamondsuit)$ with $z^\diamondsuit = (x^\diamondsuit, y^\diamondsuit)$ of QP$'$. Since the optimal value of QP$'$ is 0, we have

$$s^\diamondsuit - (z^\diamondsuit)^TQz^\diamondsuit \;=\; f(z^\diamondsuit, s^\diamondsuit) \;\leq\; \varepsilon f(z^{\#}, s^{\#}) \;\leq\; \varepsilon|A+B|\,.$$

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

REMARK 4.1. The proof in [25] 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^TQy is bilinear, we can also directly apply an LDU^T factorization on the matrix of the bilinear form.

4.2 Randomized approximation. By Section 4.1, the problem of finding a Nash equilibrium in a bimatrix game of fixed rank can be reduced to a quadratic optimization problem

$$\min_{x \in P} \frac{1}{2}x^TQx + q^Tx \tag{4.8}$$

with compact support set $P = \{x \in \mathbb{R}^n \,:\, Ax \leq b\}$ where the rank k of Q is a fixed constant. Moreover, we can assume that the quadratic form $h(x) := \frac{1}{2}x^TQx$ only depends on the first k variables $x_1, \dots, x_k$. Let $f(x) := h(x) + q^Tx$, $f_* = \min_{x\in P} f(x)$ and $f^* = \max_{x \in P} f(x)$.

We provide a randomized approximation algorithm for this class of problems, which yields a randomized approximation algorithm for a Nash equilibrium. The performance of the randomized algorithm is similar to the deterministic one, but we formulate a plausible conjecture whose proof will make the running time polynomial in $\log(1/\varepsilon)$, thus enabling us to find an exact equilibrium.

Let $\pi : \mathbb{R}^n \to \mathbb{R}^k$ be the projection of $\mathbb{R}^n$ onto the first k variables. We can formulate (4.8) as an optimization problem over the projected polytope $\pi(P)$. Let $\varphi : \pi(P) \to \mathbb{R}$ be defined by

$$\varphi(\hat{x}) \;=\; \min\{q^Tx \,:\, x = (\hat{x}, \overline{x}),\, \overline{x} \in \mathbb{R}^{n-k},\, Ax \leq b\}\,.$$

[It is easy to see that φ is a convex function.] Then our optimization problem is equivalent to

$$\min\{h(\hat{x}) + \varphi(\hat{x}) \,:\, \hat{x} \in \pi(P)\}\,. \tag{4.9}$$

We set $\hat{f}(\hat{x}) = h(\hat{x}) + \varphi(\hat{x})$ and quantify the fraction of the volume of that subset of $\pi(P)$ whose points have objective value close to f_*. For this, we consider the segments from some minimizer $\hat{x}^* \in \mathbb{R}^k$ of (4.9) to the boundary points of $\pi(P)$, and we will show that on each of these segments a certain fraction will lead to sufficiently small function values of the objective function in (4.9). We start by considering the situation in P.

Definition 4.1. *Let P be a polytope and $f : P \to \mathbb{R}$ continuous. We say that f is of* bounded σ-variation *if there exists some δ_1 such that for some minimal point x^* to $\min\{f(x) : x \in P\}$, for all $x \in bd(P)$ and for all $\lambda \in [0, \delta_1]$ the function f satisfies*

$$f((1-\lambda)x^* + \lambda x) - f(x^*) \;\le\; \sigma\,. \tag{4.10}$$

Let ε be given. If f is of bounded $(\varepsilon(f^* - f_*))$-variation then there is a polynomial fraction of "good volume", i.e., there is a sufficiently large subset in which the points are relative ε-approximations. In the following we show that the quadratic optimization problem (4.9) is always of bounded $(\varepsilon(f^*-f_*))$-variation (i.e., for every $\varepsilon \in (0,1]$); in this case δ_1 will be a function of ε.

Lemma 4.2. *For every $\varepsilon \in (0,1]$, the quadratic optimization problem (4.9) is of bounded $(\varepsilon(f^* - f_*))$-variation with $\delta_1 := \varepsilon/4$.*

In order to prove this, we can assume without loss of generality that a minimal point $\hat{x}^*$ of $\hat{f}$ is located in the origin. Let $\hat{x} \in \mathrm{bd}(\pi(P))$; by definition of $\pi(P)$, there exist $\overline{x^*}$ and $\overline{x}$ such that $x^* = (\hat{x^*}, \overline{x^*}) \in P$ and $x = (\hat{x}, \overline{x}) \in \mathrm{bd}(P)$ with $f(x) = \hat{f}(\hat{x})$. On the segment from x^* to x, the function f can be regarded as a univariate quadratic function. By parameterizing the line from x^* to a point $x \in \mathrm{bd}(P)$ by $t \in [0,1]$, we can assume that $h = at^2 + bt + c$. Moreover, we can assume that $c = 0$.

Thus we want to show the following lemma.

Lemma 4.3. *Let $h : \mathbb{R} \to \mathbb{R}$, $h(t) = at^2 + bt$, $\alpha = \max_{t\in[0,1]} h(t) - \min_{t\in[0,1]} h(t)$ and 0 be a global minimal point of h. Then for all $\varepsilon \in [0,1]$ and $\delta_1 := \varepsilon/4$ we have $h(t) \le \varepsilon\alpha$ for all $t \in [0,\delta_1]$.*

Proof. If $a > 0$ then h is convex. Hence, $h(0) = 0$ and $h(1) = a + b$ imply that for $t \in [0,1]$ we have $h(t) \le (a+b)t \le \alpha t \le \alpha\delta_1 \le \alpha\varepsilon$.

Now consider the case $a < 0$, and let $t_m = -\frac{b}{2a}$ be the maximal point of h. Since for $t_m \le 0$ we do not have a global minimum at 0, we can assume $t_m > 0$, i.e., $b < 0$. For $t_m < \frac{1}{2}$, we do not have a global minimum at 0 either. Hence, it suffices to consider the case $t_m \ge \frac{1}{2}$. Since the range of h in the interval $[0, 1/2]$ is at most α, and since $\int_0^{1/2} 2\alpha\, dt = \alpha$, the slope of h in 0 is at most 2α. Hence, for $\delta_1 \le \varepsilon/4$ the claim follows.

If we consider the segment in the projection, i.e., the segment from $\hat{x^*}$ to $\hat{x}$, then by definition of $\hat{f}$ we have $\hat{f}(\pi(x)) \le f(x)$, i.e., the function value w.r.t. $\hat{f}$ is not larger than $f(x)$. Hence, on every segment we have at least an $\frac{\varepsilon}{4}$-fraction of good function values.

We can conclude that there is a polynomial fraction of "good volume":

Corollary 4.1. *For a quadratic optimization problem of the form (4.8), let π denote the projection on the first k variables. Then there exists a subset $A \subset \pi(P)$ of volume $\ge (\varepsilon/4)^k \operatorname{vol}_k(\pi(P))$ such that for all $\hat{x} \in A$*

$$\hat{f}(\hat{x}) - f_* \;\le\; \varepsilon(f^* - f_*)\,.$$

Proof. Let x^* be an arbitrary globally optimal point. On every segment in $\pi(P)$ starting in $\pi(x^*)$ to a boundary point on $\pi(P)$, a fraction of $\varepsilon/4$ has function values as desired. This proves the claim.

We note that from the proofs of the two preceding Lemmas, it actually follows that on the ray from x^* to x, at least for half the length of the ray, the function $h(\cdot)$ is increasing. So, if we were to start from this half of the ray and do gradient descent (now just on that ray), we would arrive at x^*. We conjecture that this generalizes to gradient descent in the whole body :

Conjecture 1. *Let K be the set of points x^0 in $\pi(P)$ such that gradient descent for minimizing $\hat{f}$ on $\pi(P)$ starting at x^0 ends at x^*. Then $\operatorname{vol}(K)/\operatorname{vol}(\pi(P))$ is at least a constant.*

The idea of our randomized algorithm is to sample points in the projected polytope $\pi(P)$. Since the sampling statements are phrased in terms of oracles of convex bodies, we start by explaining how to construct a membership or a separation oracle for the projection $\pi(P)$. Given a membership oracle for P and a point $x \in P$, we can easily provide a point in $\pi(P)$ (namely, $\hat{x} = \pi(x) \in \pi(P)$) and construct a polynomial time membership oracle for $\pi(P)$: Given $x \in \mathbb{R}^n$, we have $x \in \pi(P)$ iff the linear program $\min\{\bar{x} \in \mathbb{R}^{n-k} : x = (\hat{x}, \overline{x}),\, Ax \le b\}$ has a solution.

The following lemma shows how to construct a separation oracle in polynomial time:

Lemma 4.4. *Let $P = \{x \in \mathbb{R}^n : Ax \le b\}$ be a rational polytope given as an intersection of halfspaces. Then there exists a polynomial time algorithm for providing a strong separation oracle for $\pi(P)$.*

Proof. If $P = \{x \in \mathbb{R}^n : Ax \leq b\}$ with $A = (A_1\, A_2) \in \mathbb{R}^{m\times k} \times \mathbb{R}^{m\times(n-k)}$, we can easily provide a polynomial time separation oracle for $\pi(P)$. Given some $\hat{z} \in \mathbb{R}^k$, we can test whether $\hat{z} \in \pi(P)$ by checking feasibility of the linear program in $\overline{x}$

$$\max 0^T\overline{x} \quad \text{s.t. } A_2\overline{x} \ \leq \ b - A_1\hat{z}\,.$$

If $\hat{z}$ is not feasible for $\pi(P)$ then the dual program

$$\min(b - A_1\hat{z})^T u \quad \text{s.t. } A_2^T u = 0,\ u \geq 0$$

is unbounded, i.e., there exists some $u^* \in \mathbb{R}^m$ with $(b - A_1\hat{z})u^* = 1$, $A_2^T u^* = 0$, $u^* \geq 0$. Now it can easily be seen that the hyperplane $H = \{\hat{x} \in \mathbb{R}^k : (u^*)^T A_1\hat{x} = (u^*)^T b\}$ yields the desired separating hyperplane.

In order to sample the points in $\pi(P)$, we use the following sampling theorem of Kannan, Lovász, and Simonovits [9, Theorem 2.2]. Here, for two probability measures p, q on a convex body K, the *total variation distance* between p and q is defined by $\sup\{|p(A) - q(A)| : A \subset K\}$.

PROPOSITION 4.2. *For a convex body $K \subset \mathbb{R}^n$ given by a separation oracle and $r, R \in \mathbb{Q}$ with $r\mathbb{B}^n \subset K \subset R\mathbb{B}^n$, an integer $N > 0$ and $\varepsilon > 0$, we can generate a set of N random points $v_1, \ldots, v_N$ in K that are*

1. *almost uniform in the sense that the distribution of each one is at most ε away from the uniform in total variation distance, and*
2. *almost (pairwise) independent in the sense that for every $1 \leq i < j \leq N$ and every two measurable sets A and B of K,*

$$\big|\text{Prob}(v_i \in A, v_j \in B) - \text{Prob}(v_i \in A)\text{Prob}(v_j \in B)\big| \ \leq \ \varepsilon\,.$$

The running time of the algorithm is polynomial in N, n, $\log\frac{1}{\varepsilon}$ and in $\log\frac{R}{r}$.

In order to optimize over P, we then use the following algorithm:

Input: An optimization problem of the form (4.8) with h quadratic, g linear and $P = \{x \in \mathbb{R}^n : Ax \leq b\}$, an integer $N > 0$ and $0 < \varepsilon < 1$, $\delta > 0$.

Output: Points $v_1, \ldots, v_N \in P$.

1. Sample points $\hat{v}_1, \ldots, \hat{v}_N \in \pi(P)$ which are almost uniform in $\pi(P)$ (with total variation distance at most ε) and almost independent.

2. For each $j \in \{1, \ldots, N\}$ solve the linear program in the variables $\overline{x} = (x_{k+1}, \ldots, x_n)$

$$\min g(x) \ \text{ s.t. } \ Ax \leq b,\ \hat{x} = \hat{v}_j$$

to obtain points $v_1, \ldots, v_N \in P$.

This algorithm has the following performance guarantee:

THEOREM 4.2. *Given an optimization problem of the form (4.8) with h quadratic, g linear and P a polytope given as the intersection of halfspaces, $N > 0$, $0 < \varepsilon < 1$, and $\delta > 0$, the algorithm finds points $v_1, \ldots, v_N \in P$ with the following properties:*

1. *$\pi(v_1), \ldots, \pi(v_N)$ are almost uniform in $\pi(P)$ (with total variation distance at most δ) and almost independent.*
2. *In the underlying uniform distribution in $\pi(P)$, we have for any fixed $j \in \{1, \ldots, N\}$*

$$\text{Prob}(v_j \textit{ is an } \varepsilon\textit{-approximation of the quadratic problem}) \ \geq \ \left(\frac{\varepsilon}{4}\right)^k.$$

The running time of the algorithm is $\text{poly}(N, k, \mathcal{L}, \log\frac{1}{\delta})$.

Proof. The quality of the approximation follows from the lower bound of the volume in Corollary 4.1.

In order to prove the running time, with regard to Proposition 4.2 we have to show that in polynomial time we can compute radii R and r of concentric balls containing and being contained in $\pi(P)$. Since by Lemma 4.4 we can provide a separation oracle for $\pi(P)$ in polynomial time, we can compute a Löwner-John pair for $\pi(P)$ in polynomial time. The smallest principal axes of the small ellipsoid and the largest principal axes of the large ellipsoid yield the desired pair of balls. By appropriate scaling, we can assume $r = 1$. Hence, with regard to the sampling theorem, $\log\frac{R}{r}$, i.e., the bit length of R/r, is polynomial in the input length.

By choosing $\delta := \varepsilon/3$, and boosting the probability, this yields

COROLLARY 4.2. *Given an optimization problem of the form (4.8) with h quadratic, g linear and P a polytope given as the intersection of halfspaces. For any $\varepsilon > 0$, we can find in time polynomial in $\mathcal{L}$ and $\frac{1}{\varepsilon}$ with probability at least $\frac{1}{2}$ an ε-approximate solution to the quadratic program.*

Consequently, we obtain a randomized approximation algorithm for finding Nash equilibria in games of fixed rank.

4.3 Relative approximation in case of a nonnegative decomposition.

The right hand side in Definition 2.1 of an approximate 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 (4.7). 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 (4.7); 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

$$s - x^T(A+B)y \;\le\; \rho s\,.$$

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

We provide an efficient approximation algorithm for the case that $C = A + B$ has a known decomposition of the form $C = \sum_{i=1}^k u^{(i)}(v^{(i)})^T$ with non-negative vectors $u^{(i)}$ and $v^{(i)}$.

THEOREM 4.3. *If C has a known nonnegative decomposition 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* $\mathrm{poly}(\mathcal{L}, 1/\log(1+\varepsilon))$, *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 z_i and w_i in a geometric progression: denoting by $(z_i)_{\min} = \min_{x \in \mathcal{S}_1} x^T \cdot u^{(i)}$ and $(z_i)_{\max} = \max_{x \in \mathcal{S}_1} x^T \cdot u^{(i)}$ the minimum and 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 (4.7). Let the intervals of a grid cell be $[\alpha_i, (1+\varepsilon)\alpha_i]$ and $[\beta_i, (1+\varepsilon)\beta_i]$, i.e.,

$$\alpha_i \;\le\; z_i \;\le\; (1+\varepsilon)\alpha_i \quad \text{and} \quad \beta_i \;\le\; w_i \;\le\; (1+\varepsilon)\beta_i\,.$$

Then for any pair of strategies $(x, y) \in \mathcal{S}_1 \times \mathcal{S}_2$ falling into that cell, we have

$$(4.11)\quad \sum_{i=1}^k \alpha_i\beta_i \;\le\; x^T C y \;\le\; (1+\varepsilon)^2 \sum_{i=1}^k \alpha_i\beta_i\,,$$

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

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

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

$$(4.12)\quad \sum_{i=1}^k \alpha_i\beta_i \;\le\; s \;\le\; (1+\varepsilon)^2 \left(\sum_{i=1}^k \alpha_i\beta_i\right).$$

Hence, by the left inequality in (4.11) and the right inequality in (4.12) we have $x^T C y \;\ge\; \sum_{i=1}^k \alpha_i\beta_i \;\ge\; \frac{s}{(1+\varepsilon)^2}$. We conclude $s - x^T C y \;\le\; s\left(1 - \frac{1}{(1+\varepsilon)^2}\right)$, which shows Theorem 4.3.

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 a Nash equilibrium in a game of fixed rank can be found in polynomial time, remain open, and deserve further algorithmic study.

From the viewpoint of algorithmic optimization, we can interpret our randomized algorithm as a sampling-based method to optimize quadratic functions of low rank over a polytope. Current work aims at generalizing sampling-based optimization methods to more general optimization problems with some suitable "low-rank" structure. Moreover, for general optimization lacking a low-rank structure, it seems to be fruitful to combine existing techniques of low-rank optimization with these sampling-based low-rank optimization techniques.

References

[1] I. Barany, S. Vempala, and A. Vetta. Nash equilibria in random games. In *Proc. 46th IEEE Foundations of Computer Science* (Pittsburgh, PA), 2005.

[2] J. Bulow and J. Levin. Matching and price competition. To appear in American Economic Review.

[3] X. Chen and X. Deng. Settling the complexity of 2-player Nash equilibrium. *Electronic Colloquium on Computational Complexity.* Report TR05–140, 2005.

[4] X. Chen and X. Deng. Computing Nash equilibria: Approximation and smoothed complexity. *Electronic Colloquium on Computational Complexity.* Report TR06–23, 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 Sytems (AAMAS)*, Hakodate, Japan, 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), 2006.

[8] 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.

[9] R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures and Algorithms* 11:1–50, 1997.

[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, P.G. Spirakis. Polynomial algorithms for approximating Nash equilibria of bimatrix games. *Electronic Colloquium on Computational Complexity.* Report TR06–81, 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] O.L. Mangasarian. Equilibrium points of bimatrix games. *J. Soc. Industr. Appl. Math.* 12:778–780, 1964.

[15] 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.

[16] J. Nash. Equilibrium points in n-person games. *Proc. Amer. Math. Soc.* 36:48–49, 1950.

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

[18] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton, NJ, 1944.

[19] C.H. Papadimitriou. Algorithms, games and the Internet. In *Proc. 33rd ACM Symp. Theory of Computing*, Chersonissos, Kreta, 749–753, 2001.

[20] C.H. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In *Symp. on Discrete Algorithms* (Vancouver, BC), 82–91, 2005.

[21] T. Quint and M. Shubik. A bound on the number of Nash equilibria in a coordination game. *Economic Letters* 77:323–327, 2002.

[22] 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.

[23] B. von Stengel. New maximal numbers of equilibria in bimatrix games. *Discrete Comput. Geom.* 21:557–568, 1999.

[24] B. von Stengel. Computing equilibria for two-person games. In R.J. Aumann, S. Hart (eds.), *Handbook of Game Theory*, North-Holland, Amsterdam, 2002.

[25] S. Vavasis. Approximation algorithms for indefinite quadratic programming. Technical Report 91-1228, Dept. of Computer Science, Cornell University (Ithaca, NY), 1991.

[26] S. Vavasis. Approximation algorithms for indefinite quadratic programming. *Math. Program.* 57:279–311, 1992.