

Enumerating the Nash Equilibria of Rank-1 Games

Thorsten Theobald

ABSTRACT. A bimatrix game (A, B) is called a game of rank k if the rank of the matrix $A + B$ is at most k . We consider the problem of enumerating the Nash equilibria in (non-degenerate) games of rank 1. In particular, we show that even for games of rank 1 not all equilibria can be reached by a Lemke–Howson path and present a parametric simplex-type algorithm for enumerating all Nash equilibria of a non-degenerate game of rank 1.

1. Introduction

Bimatrix games belong to the fundamental concepts of game theory. A bimatrix game is given by two payoff matrices (A, B) , and by Nash's results any bimatrix game has at least one equilibrium [18, 19]. The problem of computing a Nash equilibrium (named by Papadimitriou in 2001 to be the most concrete open question on the boundary of the complexity class \mathcal{P} [20]) has received much attention in the last years. Chen and Deng recently showed that the problem is \mathcal{PPAD} -complete [3] and (together with Teng [4]) that the problem of computing a $1/n^{\Theta(1)}$ -approximate Nash equilibrium remains \mathcal{PPAD} -complete. Thus it is unlikely that an efficient algorithm exists.

The corresponding *enumeration problem* asks to enumerate all equilibria (in the finite case) or all the extreme equilibria (in the degenerate case where an infinite number of Nash equilibria is possible). The enumeration problem is similar to (but more difficult than) enumerating the vertices of a polyhedron given as the intersection of half-spaces. In the latter vertex enumeration problem, the Upper Bound Theorem gives a tight estimate for the maximal number of vertices, but the analogous problem of determining the maximum number of Nash equilibria of a bimatrix game is an open problem (see [23, 24]).

For the special case of zero-sum games, the set of Nash equilibria defines a polyhedral set in the strategy space, and the problem of computing the set of all Nash equilibria is equivalent to linear programming (see [6, Chapter 13.2]). Hence, for a non-degenerate zero-sum game, the set of Nash equilibria consists of a single point, and thus the enumeration problem becomes trivial.

2000 *Mathematics Subject Classification*. Primary 91A05; Secondary 90C20.
This is the final form of the paper.

Recently, Kannan and Theobald [9] have introduced a hierarchy of bimatrix games in which the matrix $A + B$ is restricted to be of rank at most k , for some fixed constant k . For any fixed $k \geq 1$, this class strictly generalizes the class of zero-sum-games. Of course, the case $k = 1$ is of particular importance; it is the smallest extension of zero-sum games in the hierarchy.

In contrast to zero-sum games, non-degenerate rank- k games (for any fixed $k \geq 1$) can have an arbitrarily large number of equilibria. In particular, this provides a sharp separation between the class of rank-1 games and the class of zero-sum games. From the computational viewpoint, Nash equilibria can be efficiently approximated in games of fixed rank, but the question of exact polynomial time computability is open even for games of rank 1.

In this paper, we consider the enumeration problem for games of rank 1. Similar to the situation above, the rank condition provides additional structure which can be exploited. An initial question is whether any equilibrium can be reached by a Lemke–Howson path (as defined formally in Section 3). For arbitrary bimatrix games Aggarwal has shown that in general not all equilibria can be reached by a Lemke–Howson path [1]. By providing an example of a rank-1 game for which not all equilibria can be reached in this way, we strengthen Aggarwal’s result and thus answer our initial question in the negative.

As main contribution of the paper, we propose a parametric simplex-type algorithm for enumerating the Nash equilibria of rank-1 games. This algorithm is based on the techniques of Konno and Kuno who have investigated linear multiplicative programs ([10, 11]). Our problem can be seen as an enumeration problem of generalized linear multiplicative programming. Moreover, the situation of games provides additional combinatorial structure which can be exploited.

The paper is structured as follows. In Section 2, we introduce the basic concepts of bimatrix games as well as of rank- k games and review existing work on the enumeration of Nash equilibria. In Section 3, we show that not all equilibria can be reached by a Lemke–Howson path in a rank-1 game. Then, in Section 4, we present the parametric simplex-type algorithm for enumerating the Nash equilibria of a non-degenerate rank-1 game; finally, we explain how to modify the algorithm so as to cover degenerate situations as well.

2. Preliminaries

2.1. Bimatrix games. We consider an $m \times n$ -bimatrix game with payoff matrices $A, B \in \mathbb{R}^{m \times n}$. 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

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

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 (see [24]). In the case of a non-degenerate game the set of Nash equilibria consists of finitely many isolated points. We remark that there exist various other definitions of degeneracy in the literature which are equivalent to that notion (see [24, Theorem 2.10]).

2.2. Earlier work on enumeration of equilibria. The classical Lemke–Howson algorithm serves to find one Nash equilibrium in a bimatrix game ([12, 24]). We discuss this algorithm and its (negative) relation to enumeration in more detail in Section 3.

From the viewpoint of computational complexity, the problem of counting the number of Nash equilibria in a bimatrix game is $\#\mathcal{P}$ -hard [5]. Hence, the enumeration problem is $\#\mathcal{P}$ -hard.

The general idea of existing approaches for the enumeration of Nash equilibria is to transform the problem into a problem of polyhedral computation. For each game we define a pair of polyhedra in which each Nash equilibrium of the game corresponds to a *complementary vertex pair*.

Definition 2.1. For an $m \times n$ -bimatrix game (A, B) , define the polyhedra P and Q by

$$(2.3) \quad P = \left\{ (x, \pi_2) \in \mathbb{R}^m \times \mathbb{R} : \underbrace{x \geq 0}_{\substack{\text{inequalities} \\ 1, \dots, m}}, \underbrace{x^T B \leq \mathbf{1}^T \pi_2, \mathbf{1}^T x = 1}_{\substack{\text{inequal's} \\ m+1, \dots, m+n}} \right\},$$

$$(2.4) \quad Q = \left\{ (y, \pi_1) \in \mathbb{R}^n \times \mathbb{R} : \underbrace{Ay \leq \mathbf{1} \pi_1}_{\substack{\text{inequalities} \\ 1, \dots, m}}, \underbrace{y \geq 0}_{\substack{\text{inequal's} \\ m+1, \dots, m+n}}, \mathbf{1}^T y = 1 \right\},$$

where $\mathbf{1}$ is the all-1 vector.

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 $\pi_1, \pi_2 \in \mathbb{R}$ such that $(\bar{x}, \pi_2) \in P$, $(\bar{y}, \pi_1) \in Q$ and for all $i \in \{1, \dots, m+n\}$, the i th inequality of P or Q is binding. Here, π_1 and π_2 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 P , and for $j \in \{1, \dots, n\}$ we call the inequality $\bar{x}^T B_{(j)} \leq \pi_2$ the j th *best response inequality* of P . And analogously for Q .

Example 2.2 (Taken from [22]). The bimatrix game with payoff matrices

$$A = \begin{pmatrix} 2 & 1 & 5 \\ 3 & 0 & 4 \end{pmatrix}, \quad B = \begin{pmatrix} 7 & 8 & 1 \\ 2 & 1 & 6t \end{pmatrix}$$

has 3 Nash equilibria:

$$\left((1, 0)^T, (0, 1, 0)^T \right), \quad \left(\left(\frac{1}{2}, \frac{1}{2} \right)^T, \left(\frac{1}{2}, \frac{1}{2}, 0 \right)^T \right), \quad \left(\left(\frac{2}{5}, \frac{3}{5} \right)^T, \left(\frac{1}{2}, 0, \frac{1}{2} \right)^T \right).$$

The polyhedra P and Q are shown in Figure 1.




FIGURE 1. In the example, P is a two-dimensional polyhedron embedded in $\overline{\mathcal{S}}_1$, and Q is a three-dimensional polyhedron embedded in $\overline{\mathcal{S}}_2$. The vertical direction corresponds to the variables π_2 and π_1 , respectively.

The situation becomes more involved when one also wants to handle degenerate games. For an arbitrary bimatrix game, the set of all Nash equilibria is in general a non-convex subset, but can be represented as the union of a finite number of polytopes (called *maximal Nash subsets*) in the strategy space (see [15]). Hence, in the degenerate situation, the task is to enumerate all vertices of every maximal Nash subset of a bimatrix game. Every equilibrium point is a convex combination of some of these *extreme equilibria* (cf. [8, 26]).

These concepts can be used to provide algorithms for enumerating Nash equilibria of bimatrix games. The earliest ones can be found in Vorob'ev [25] and Mangasarian [13]; for later developments see Muhamediev [17], Winkels [26], and Audet, Hansen, Jaumard, and Sauvard [2]. The latter paper also contains a detailed historical account on these algorithms.

2.3. Games of fixed rank. For a given constant $k \geq 0$, a bimatrix game is called a *game of rank k* if the matrix $A + B$ has rank at most k .

Kannan and Theobald have shown that for any fixed $k \geq 1$, the number of Nash equilibria in a non-degenerate game of rank k cannot be bounded by a function in terms of k [9]. In particular, the case $k = 1$ stands in sharp contrast to the case $k = 0$ of zero-sum games; there, the number of Nash equilibria in the finite situation is always 1. For the case $d := m = n$, the best known lower bound for the maximal number of Nash equilibria of rank-1 games is linear in d :

Proposition 2.3. *For any $d \in \mathbb{N}$ there exists a non-degenerate $d \times d$ -game of rank 1 with at least $2d - 1$ Nash equilibria.*

A construction achieving this number is given by the $d \times d$ -game (A, B) with

$$(2.5) \quad a_{ij} = 2ij - i^2 + j^2, \quad b_{ij} = 2ij + i^2 - j^2$$

(see [9]). Since $A + B = (4ij)_{i,j}$, the matrix $A + B$ is of rank 1.

It is not known whether in games of fixed rank a Nash equilibrium can be computed in polynomial time. In [9], the following approximation result was shown. Here, a pair (x, y) of mixed strategies is called an ε -approximate equilibrium if

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

where $\ell(x, y)$ denotes the sum of the losses of the players,

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

and $|\cdot|$ denotes the maximum absolute value of the entries of a matrix.

Proposition 2.4. *Let k be a fixed constant and $\varepsilon > 0$. In a game of rank k , 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.*

There are several operations on a bimatrix game (A, B) which do not change the set of Nash equilibria:

- (1) adding multiples of the all-1 vector to a given column of A or a given row of B ;
- (2) positive scaling of a given column of A or a given row of B .

Since these operations can change the rank of the game, the following useful consequence is obtained.

Lemma 2.5. *Let (A, B) be a $d \times d$ -game of rank d . Then there exists a game of rank $d - 1$ with the same set of Nash equilibria.*

PROOF. Assume without loss of generality that $d \geq 2$ and that $C := A + B$ is of rank exactly d . Then there exists a column j of C which is not a multiple of the all-1 vector. Denote by v the column vector obtained from the entries of the j th column of C . Since the affine line $v + \mathbb{R}(1, \dots, 1)^T$ intersects the $(d - 1)$ -dimensional linear subspace defined by the $d - 1$ other columns, there exists some $\lambda \in \mathbb{R}$ such that adding $\lambda(1, \dots, 1)^T$ to the j th column of C yields a matrix of rank at most $d - 1$. Thus, adding $\lambda(1, \dots, 1)^T$ to the j th column of A turns the game (A, B) into a game of rank $d - 1$. \square

3. Lemke–Howson paths

In this section, we recall the classical Lemke–Howson algorithm for finding a Nash equilibrium in a bimatrix game and then show that not even for games of rank 1 all equilibria can be reached by a Lemke–Howson path.

The Lemke–Howson algorithm is a simplex-type algorithm which can be combinatorially described in terms of a graph. In order to define this graph, we start from the polyhedral description of the bimatrix game in terms of P and Q . Consider a pair of strategies $(x, y) \in \mathcal{S}_1 \times \mathcal{S}_2$, and let π_1 and π_2 be the resulting payoffs. We label each of the strategies x and y by the indices of the inequalities in (2.3) and (2.4) that are binding. For a non-degenerate $m \times n$ -game, only the vertices of P have m labels and the vertices of Q have n labels, and there do not exist points in P and Q with more than m or n labels, respectively (see [24, Theorem 2.7]).

We define the graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ as follows. The vertex set V_1 of G_1 consists of the vertices of P , with an additional vertex $0 \in \mathbb{R}^m$ that has all labels in the set $\{1, \dots, m\}$. Two vertices x and x' are connected by an edge if they differ in exactly one label, i.e., if they have $m - 1$ labels in common. Similarly, let G_2 be the graph whose vertex set V_2 consists of the vertices of Q , with an additional vertex $0 \in \mathbb{R}^n$ having all labels in $\{m + 1, \dots, m + n\}$. Two vertices in G_2 are connected if they have $n - 1$ labels in common.

The product graph $G_1 \times G_2$ of G_1 and G_2 is defined by the vertex set $V_1 \times V_2$, and the edges are given by $\{x\} \times \{y, y'\}$ for vertices x of G_1 and edges $\{y, y'\}$ of G_2 , or by $\{x, x'\} \times \{y\}$ for edges $\{x, x'\}$ of G_1 and vertices y of G_2 .

From a combinatorial viewpoint, the Lemke–Howson algorithm can now be described as follows. Fix an $r \in \{1, \dots, m+n\}$. A vertex (x, y) of $G_1 \times G_2$ is called *r-almost completely labeled* if the union of the labels is the set $\{1, \dots, m+n\} \setminus \{r\}$. Since two adjacent vertices x and x' in G_1 have $m-1$ common labels, the edge $\{x, x'\} \times \{y\}$ of $G_1 \times G_2$ is also *r-almost completely labeled* if y has the remaining n labels except r . And similarly for edges $\{x\} \times \{y, y'\}$ of $G_1 \times G_2$. The Lemke–Howson algorithm starts from the artificial equilibrium $(0, 0)$ which has all labels and then follows the unique path where the label r is missing. After finitely many steps, it reaches a Nash equilibrium of the game.

For different choices of r it is possible that we reach different Nash equilibria. This led to the early question in the algorithmic study of games on whether any equilibrium can be reached by a Lemke–Howson path, i.e., by some choice of r . For general bimatrix games it is known that the set of Lemke–Howson paths does not enumerate all Nash equilibria (Aggarwal [1]; see also [24] and the references therein).

Since games of rank 1 are a very special case of general bimatrix games, the question arises whether for a game of rank 1 all equilibria can be reached by a Lemke–Howson path. Here, we strengthen the (un-)reachability result by showing that even for games of rank 1 not all equilibria can be reached. Namely, consider any bimatrix game (A', B') which has two pure and one mixed equilibria. By Lemma 2.5, this game can be transformed into a rank-1 game (A, B) with the same equilibria. For the sake of concreteness, we consider the rank-1 game

$$A = \begin{pmatrix} -28 & -18 \\ -8 & -23 \end{pmatrix}, \quad B = \begin{pmatrix} 10 & 30 \\ 20 & 15 \end{pmatrix},$$

which results from applying Lemma 2.5 onto Aggarwal’s example. The polyhedra P and Q are shown in Figure 2.

The game has three Nash equilibria:

$$((1, 0)^T, (0, 1)^T), \quad \text{with payoffs } (-18, 30);$$




FIGURE 2. The polyhedra P and Q and the labels of the vertices.

$$\begin{aligned} &((0, 1)^T, (1, 0)^T), \quad \text{with payoffs } (-8, 20); \quad \text{and} \\ &((\frac{1}{5}, \frac{4}{5})^T, (\frac{1}{5}, \frac{4}{5})^T), \quad \text{with payoffs } (-20, 18). \end{aligned}$$

With regard to the Lemke–Howson graphs, the locally maximal peaks of the lower hull of P are connected with the 0-vertex (which has labels 1, 2) and the maximal peaks in the lower hull of Q are connected with the 0-vertex (which has labels 3, 4). Considering all the four possible values of r , only the first two equilibria can be obtained via Lemke–Howson paths, and the third equilibrium cannot be obtained. E.g., for the initial missing label $r = 1$ we obtain the following path with labels. Here, the two components correspond to the graphs G_1 and G_2 .

$$\begin{pmatrix} 1, 2 \\ 3, 4 \end{pmatrix} \rightarrow \begin{pmatrix} 2, 4 \\ 3, 4 \end{pmatrix} \rightarrow \begin{pmatrix} 2, 4 \\ 1, 3 \end{pmatrix},$$

where the last vertex pair gives the Nash equilibrium $((1, 0)^T, (0, 1)^T)$. Hence we can conclude:

Theorem 3.1. *There exist games of rank 1 for which not all equilibria can be reached by a Lemke–Howson path.*

Let G' be the graph with vertex set $V_1 \times V_2$ whose edge set is the union (over all r) of r -almost completely labeled edges. In a paper of Shapley [21], the following example attributed to Wilson is given which shows that for games of arbitrary rank the graph G' can even be disconnected. (However, note that in general a disconnected graph does not necessarily imply that there exists some equilibria which cannot be reached from the artificial equilibrium; namely, the almost complementary edges can form cycles). Let the 3×3 -game (A, B) be defined by

$$A = \begin{pmatrix} 0 & 3 & 0 \\ 2 & 2 & 0 \\ 3 & 0 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & 2 & 3 \\ 3 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

In the graph G' , the mixed equilibrium of (A, B) cannot be reached from the artificial equilibrium. Hence, even modifications of the Lemke–Howson-algorithm which are allowed to change the index of the missing label within the graph traversal cannot enumerate all Nash equilibria.




FIGURE 3. The projections of the polyhedra P and Q in Wilson’s examples. The picture contains the labels of the points with three labels. The underlined numbers show the best responses of the regions.

Applying the rank reduction Lemma 2.5 on Wilson's example immediately implies that for 3×3 -games of rank 2, the graph G' can be disconnected. It would be interesting to know if this graph can also be disconnected for the case of rank-1 games. While from the principle viewpoint this question is decidable, currently we do not know the answer.

4. A parametric algorithm for enumerating all Nash equilibria of rank-1 games

In this section, we present a parametric simplex-type enumeration algorithm for games of rank 1. For simplicity we concentrate on the situation where the game is non-degenerate and the polytopes P and Q are in general position. In Section 4.2 we explain how to modify the algorithm to cover degenerate situations as well. An example is presented in Section 4.3.

4.1. Non-degenerate situations. We consider the following well-known characterization of a Nash equilibrium in terms of a quadratic program [14]. For any $(x, \pi_2) \in P$ and $(y, \pi_1) \in Q$, we have

$$x^T(A+B)y \leq \max_{1 \leq i \leq m} A^{(i)}x + \max_{1 \leq j \leq n} x^T B_{(j)} \leq \pi_1 + \pi_2$$

with equality if and only if (x, y) is a Nash equilibrium with payoffs π_1 and π_2 . Consider the quadratic program

$$\begin{aligned} \text{(QP)} \quad & \max x^T(A+B)y - \pi_1 - \pi_2 \\ & (x, \pi_2) \in P, \\ & (y, \pi_1) \in Q. \end{aligned}$$

Hence, we obtain the following characterization of the Nash equilibria in terms of the quadratic program.

- Lemma 4.1.** (1) *For any feasible solution of QP, the objective value is nonpositive.*
 (2) *A pair $(x^*, y^*) \in \mathcal{S}_1 \times \mathcal{S}_2$ is a Nash equilibrium of the bimatrix game (A, B) if and only if there exist $\pi_1^*, \pi_2^* \in \mathbb{R}$ such that $(x^*, y^*, \pi_1^*, \pi_2^*)$ is a feasible solution of QP with objective value 0.*

Note that the inequalities describing the feasible set are linear. Moreover, the objective function only depends on the sum $A+B$ rather than on A or B individually.

If the game is of rank 1, we write $A+B$ in the form $A+B = b \cdot c^T$ with $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$. That is, after a linear variable transformation we want to enumerate the global optima of a function of the form $f+g$ where f is the product of two variables and g is a linear function. The problem of optimizing a product of two variables over a polyhedron is called a *linear multiplicative program* ([10, 11]). In these references, parametric primal-dual simplex-type algorithms were presented to find the optimal value. Based on these techniques, we now present an algorithm which enumerates all equilibria of a rank-1 game. Since we are starting from a game, we will see below that the relevant bases have a special combinatorial structure, where m elements of the basis refer to the polyhedron P and $n-1$ elements refer to the polyhedron Q .

ALGORITHM 1. Enumeration algorithm.

Set $\xi = \xi_{\min}$;
while $\xi \leq \xi_{\max}$ **do**
 Compute the optimal point $z := z(\xi)$ of LP(ξ);
 if $z(\xi)$ has objective value 0 **then**
 Let x, y be the strategies played in $z(\xi)$;
 Output “Nash equilibrium:” x, y ;
 Compute the smallest $\xi' > \xi$ where the optimal basis changes and the new optimal basis;

By setting $\xi := c^T y$, we can write the quadratic program as

$$\begin{aligned}
 (4.1) \quad & \max (x^T b)\xi - \pi_1 - \pi_2, \\
 & (x, \pi_2) \in P, \\
 & (y, \pi_1) \in Q, \\
 & c^T y = \xi.
 \end{aligned}$$

We consider ξ as a parameter to the optimization problem. For a given value to ξ , the problem becomes a linear program which we call LP(ξ). Geometrically, for fixed ξ we are slicing the feasible set polyhedron by a hyperplane $c^T y = \xi$ on which the strategy of the second player satisfies a particular linear condition.

Remark 4.2. In the special case where c is a multiple of the all-1 vector, the hyperplane defined by $c^T y = \xi$ is parallel to the hyperplane defined by $\mathbf{1}^T y = 1$ (which is part of the definition of Q). In that situation, all columns of $A + B$ are identical, and games of this type are called *row-constant games* (see [7]). We will come back to that special case below.

The feasible set of QP is unbounded. However, since $y \in \mathcal{S}_2$ the range $[\xi_{\min}, \xi_{\max}]$ of ξ is bounded, namely

$$(4.2) \quad [\xi_{\min}, \xi_{\max}] = \left[\min_{y \in \mathcal{S}_2} c^T y, \max_{y \in \mathcal{S}_2} c^T y \right] = \left[\min_{1 \leq j \leq n} c_j, \max_{1 \leq j \leq n} c_j \right].$$

Even for fixed ξ , the feasible region of LP(ξ) can be unbounded. But by Lemma 4.1, the objective value of QP is bounded from above and hence also LP(ξ) is bounded from above.

Let $\mathcal{I} := \{1, \dots, m, m+1, \dots, m+n\}$ be the index set of the inequalities of P , and $\mathcal{J} := \{1, \dots, m, m+1, \dots, m+n\}$ be the index set of the inequalities of Q .

We consider a fixed parameter value ξ . If the resulting LP(ξ) is in general position, then the optimal point v of LP(ξ) is unique and v is a vertex of the feasible set polyhedron of LP(ξ). If ξ is a sufficiently generic value (we will come back to this aspect below) then v can be uniquely described in terms of a basis, i.e., by a choice $I \subset \mathcal{I}$ of cardinality m and $J \subset \mathcal{J}$ of cardinality $n-1$.

For a given ξ , let (I, J) denote an optimal basis for the linear program LP(ξ) depending on ξ . The idea of the enumeration algorithm is to keep track on the values of ξ where the optimal basis changes. This yields the enumeration algorithm in the Algorithm 1.

In order to explain how the update steps can be performed efficiently, we analyze the set of the values of the parameter ξ which yield the same optimal basis. For this, we consider the dual program of the parametric LP (4.1).

Let u_i be the dual variable corresponding to the i th inequality of P , $1 \leq i \leq m+n$, and u_{m+n+j} be the dual variable corresponding to the j th inequality of Q , $1 \leq j \leq m+n$. Further let $u_{2(m+n)+1}$ and $u_{2(m+n)+2}$ be the dual variables of the equations $\mathbf{1}^T x = 1$ and $\mathbf{1}^T y = 1$, and $u_{2(m+n)+3}$ be the dual variable of the equation $c^T y = \xi$. The dual variables $u_1, \dots, u_{2(m+n)}$ are sign-restricted while the dual variables $u_{2(m+n)+1}, \dots, u_{2(m+n)+3}$ are not.

We set $z = (x, y, \pi_1, \pi_2)$ and the constants $K = 2(m+n)$ and $N = m+n+2$. Let $M_1 \in \mathbb{R}^{K \times N}$ and $e_1 \in \mathbb{R}^K$ be defined by

$$M_1 = \left(\begin{array}{c|c|c|c} -I & & & \\ B^T & & & \\ \hline & A & -\mathbf{1} & \\ \hline & & -I & \\ \hline & & & -\mathbf{1} \end{array} \right), \quad e_1 = 0.$$

Further, let $M_2 \in \mathbb{R}^{3 \times N}$ and $e_2 \in \mathbb{R}^3$ be defined by

$$M_2 = \left(\begin{array}{c|c|c|c} \mathbf{1}^T & & & \\ \hline & \mathbf{1}^T & & \\ \hline & & c^T & \\ \hline & & & \end{array} \right), \quad e_2 = \begin{pmatrix} 1 \\ 1 \\ \xi \end{pmatrix}.$$

Then the feasible region of the linear program $\text{LP}(\xi)$ is given by $M_1 z \leq e_1$, $M_2 z = e_2$. For fixed ξ , the dual of $\text{LP}(\xi)$ is

$$\begin{aligned} & \min(e_1^T, e_2^T)u \\ & (M_1^T \mid M_2^T)u = (b_1\xi, \dots, b_m\xi, \underbrace{0, \dots, 0}_{n \text{ times}}, -1, -1)^T, \\ & u_1, \dots, u_K \geq 0, \end{aligned}$$

where $u = (u_1, \dots, u_{K+3})$ is the vector of dual variables.

By the complementary slackness conditions for linear programming, the optimal solution of the dual program satisfies

$$(4.3) \quad u_i = 0 \text{ for all } i \notin I \quad \text{and} \quad u_{m+n+j} = 0 \text{ for all } j \notin J.$$

A basis (I, J) of $\text{LP}(\xi)$ is also an optimal basis for all those programs $\text{LP}(\xi')$ for which the point described by (I, J) is feasible and for which there exists a feasible solution to the dual program satisfying the complementarity condition (4.3).

For the special case of a zero-sum game, the parametric formulation degenerates to the well-known pair of dual linear programs associated with the game. Namely, if $b = c = 0$ then $\xi_{\min} = \xi_{\max} = 0$ and the matrix M_1 becomes

$$M_1 = \left(\begin{array}{c|c|c|c} -I & & & \\ (-A)^T & & & \\ \hline & A & -\mathbf{1} & \\ \hline & & -I & \\ \hline & & & -\mathbf{1} \end{array} \right).$$

For $\xi = 0$, the dual then becomes

$$\begin{aligned} & \min u_{K+1} + u_{K+2} \\ \text{s.t.} & \left(\begin{array}{c|c|c|c|c|c} -I & -A & & & \mathbf{1} & \\ & & A^T & & & \\ & & -\mathbf{1}^T & & & \\ & & & -I & & \\ & & & & & \mathbf{1} \end{array} \right) u = \begin{pmatrix} 0 \\ 0 \\ -1 \\ -1 \end{pmatrix}, \quad u_1, \dots, u_K \geq 0. \end{aligned}$$

Hence we can conclude:

Lemma 4.3. *If the game is a zero-sum game (i.e., $A + B = 0$) then the parametric problem (4.1) is only feasible for $\xi = 0$. In this case, the constraints of the dual LP coincide with the constraints of the primal program under the identifications $x_i = u_{m+n+i}$, $1 \leq i \leq m$, $y_j = u_{m+j}$, $1 \leq j \leq n$, $\pi_1 = u_{2(m+n)+2}$, $\pi_2 = u_{2(m+n)+1}$, and considering u_1, \dots, u_m and u_{2m+n+1}, \dots, u_K as slack variables; the objective functions are additive inverses under these identifications.*

Similarly, for row-constant games the range of ξ -values consists of a single point, and the equilibria of these games can be phrased as linear programs. Namely, if there exist constants $u_1, \dots, u_m \in \mathbb{R}$ with

$$a_{ij} + b_{ij} = u_i \quad \text{for } 1 \leq i \leq m, 1 \leq j \leq n$$

then the zero-sum game (A', B') defined by $b'_{ij} = b_{ij} - u_i$ has the same Nash equilibria as (A, B) .

From now on, let the game (A, B) neither be a zero-sum game nor a row-constant game. Let (I, J) be an optimal basis of $\text{LP}(\xi)$ for some parameter value ξ . Let $\mathcal{B} = I \cup J$, and for a matrix A let $A_{\mathcal{B}}$ be the submatrix of A with rows in \mathcal{B} . By our assumption, the system of linear equations

$$(4.4) \quad (M_1)_{\mathcal{B}} z = (e_1)_{\mathcal{B}},$$

$$(4.5) \quad M_2 z = e_2$$

has a unique solution. Let $z(\xi)$ be the solution point of this system. In order to check whether there exists a dual solution satisfying the complementary slackness conditions, set $\mathcal{B}' = \mathcal{B} \cup \{K+1, K+2, K+3\}$. Compute

$$((M_1)_{\mathcal{B}'}^T \mid M_2^T)^{-1} (b_1 \xi, \dots, b_m \xi, 0, \dots, 0, -1, -1)^T$$

and set all components of u indexed by $\{1, \dots, 2K\} \setminus \mathcal{B}$ to zero to obtain a vector $u = u(\xi)$ with $(M_1^T \mid M_2^T)u(\xi) = (b_1 \xi, \dots, b_m \xi, 0, \dots, 0, -1, -1)^T$ and which satisfies the complementarity conditions.

Lemma 4.4. *Let $z(\xi)$ be defined by solving (4.4) and (4.5) for z and $u(\xi)$ be as described before. The set of ξ such that \mathcal{B} is an optimal basis of $\text{LP}(\xi)$ is given by the two conditions*

$$(4.6) \quad M_1 z(\xi) \leq e_1,$$

$$(4.7) \quad u_1(\xi), \dots, u_K(\xi) \geq 0.$$

PROOF. The first condition is satisfied if and only if $z(\xi)$ is feasible. For a vertex $z(\xi)$ of the feasible set polyhedron the second condition is satisfied if and only if $z(\xi)$ is optimal for $\text{LP}(\xi)$. \square

For a variable ξ , both conditions in (4.6) and (4.7) are linear conditions in ξ . Hence, the range of ξ -values in which both conditions are satisfied defines an interval. Let $[\alpha_1, \alpha_2]$ and $[\beta_1, \beta_2]$ be the intervals defined by (4.6) and (4.7), respectively. Then $[\xi_1, \xi_2] := [\max\{\alpha_1, \beta_1\}, \min\{\alpha_2, \beta_2\}]$ is the interval for ξ in which both conditions are satisfied. Since (I, J) is an optimal basis for some ξ , the interval is nonempty. We distinguish two cases:

CASE $\xi_2 = \alpha_2$. Then for the value $\xi = \xi_2$, there are $m + n$ inequalities, indexed by $\mathcal{B} \cup \{j\}$ for some j , which are binding in the primal program. One of the inequalities i of the current vector $z(\xi)$ becomes violated for $\xi > \xi_2$. We remove the index i from \mathcal{B} and replace it by the index j . Since i was chosen to be the earliest violated inequality, after this dual simplex step the new basis defines a feasible and optimal point for sufficiently small $\xi > \xi_2$.

From the viewpoint of the game, we can distinguish the following subcases corresponding to the set of $m + n$ active inequalities.

- (1) $m + 1$ inequalities for x and $n - 1$ inequalities for y are binding.

Then i and j refer to indices of inequalities for the polyhedron P . If $i \in \{1, \dots, m\}$ then one of the unplayed pure strategies of the first player is now effectively played. If $i \in \{m + 1, \dots, m + n\}$ then one of the previous best pure responses of the second player becomes a suboptimal response.

If $j \in \{1, \dots, m\}$ then one of the played pure strategies of the first player becomes unplayed. If $j \in \{m + 1, \dots, m + n\}$ then one of the previous suboptimal pure responses of the second player becomes a best response.

- (2) m inequalities for x and n inequalities for y are binding.

Then i and j refer to indices of inequalities for the polyhedron Q . If $i \in \{1, \dots, m\}$ then one of the previous best pure responses of the first player becomes a suboptimal response. If $i \in \{m + 1, \dots, m + n\}$ then one of the unplayed pure strategies of the second player is now effectively played.

If $j \in \{1, \dots, m\}$ then one of the previous suboptimal pure responses of the first player becomes a best response. If $j \in \{m + 1, \dots, m + n\}$ then one of the played pure strategies of the second player becomes unplayed.

CASE $\xi_2 = \beta_2$. Then for the value $\xi = \xi_2$ there exists an index $i \in \mathcal{B}$ such that the dual variable y_i becomes zero. For $\xi > \xi_2$ the current vector $z(\xi)$ is no longer optimal. We remove the index i from the basis and perform a simplex step moving along an edge of the polyhedron. Since i was chosen to be the earliest violated optimality condition, this simplex step gives a new optimal basis \mathcal{B}' .

We conclude:

Theorem 4.5. *Let (A, B) be a non-degenerate bimatrix game of rank 1. Then the parametric algorithm enumerates all Nash equilibria. The number of pivot steps of the algorithm is polynomial in the product $f_0(P) \cdot f_0(Q)$, where $f_0(P)$ and $f_0(Q)$ denote the number of vertices of P and Q respectively.*

PROOF. For those parameter values ξ where the basis does not change, the optimal basis is given by an m -element subset $I \subset \mathcal{I}$ corresponding to a vertex of P and by an $(n - 1)$ -element subset $J \subset \mathcal{J}$ that can be extended to an n -element subset characterizing a vertex of Q . \square

As mentioned in Section 2 note that in general not every vertex pair of $P \times Q$ corresponds to a Nash equilibrium.

4.2. Degenerate games. If the game is degenerate then there are two issues. The polyhedra might not be simple and the number of Nash equilibria can become infinite. In order to resolve the first of these points, we have to cope with the same issues as in the case of the simplex algorithm (in particular, the issue of possible cycling when changing a basis). However, with the same techniques as for linear programming (such as symbolic perturbation, lexicographic ordering), these situations can be resolved.

In order to extend the algorithm to degenerate games with an infinite number of equilibria as well, by Section 2.2 the extreme equilibria are sufficient to determine all maximal Nash subsets. Indeed, our method can be modified to find all extreme equilibria even in degenerate cases.

4.3. Example. We consider the rank-1 game from (2.5) for $d = 2$:

$$A = \begin{pmatrix} 2 & 7 \\ 1 & 8 \end{pmatrix}, \quad B = \begin{pmatrix} cc2 & 1 \\ 7 & 8 \end{pmatrix};$$

i.e., $A + B = bc^T$ with $b = (2, 4)^T$, $c = (2, 4)^T$. We have $\min_{y \in S_2} c^T y = 2$ and $\max_{y \in S_2} c^T y = 4$. For the value of $\xi = 2$, the inequalities with indices in $\{2, 3, 5, 8\}$ are binding.

$\xi = c^T y$	objective	binding
2	0	$\{2, 3, 5, 8\}$
$\xi \in (2, \frac{5}{2})$	< 0	$\{2, 3, 5\}$
$\frac{5}{2}$	$-\frac{1}{4}$	$\{2, 3, 4, 5\}$
$\xi \in (\frac{5}{2}, 3)$	< 0	$\{3, 4, 5\}$
3	0	$\{3, 4, 5, 6\}$
$\xi \in (3, \frac{7}{2})$	< 0	$\{3, 4, 6\}$
$\frac{7}{2}$	$-\frac{1}{4}$	$\{1, 3, 4, 6\}$
$\xi \in (\frac{7}{2}, 4)$	< 0	$\{1, 4, 6\}$
4	0	$\{1, 4, 6, 7\}$

For $\xi = \frac{5}{2}$, in the optimal situation we obtain uniquely $y = (\frac{3}{4}, \frac{1}{4})^T$ and $\pi_1 = 13/4$. The optimal values for x and π_2 are not unique; by substituting $x_2 = 1 - x_1$ we can analyze the situation locally around the parameter value $\xi = \frac{5}{2}$ in the (x_1, π_2) -plane (see Figure 4; but note that the x and y -axis are scaled differently). The induced optimization problem is

$$\begin{aligned} \max \quad & (-2x_1 + 4)\xi - \pi_2 - \frac{13}{4} \\ & 0 \leq x_1 \leq 1, \quad G \quad (1, 2) \\ & -5x_1 + 7 \leq \pi_2, \quad G \quad (3) \\ & -7x_1 + 8 \leq \pi_2. \quad (4) \end{aligned}$$

For $\xi = \frac{5}{2}$ the objective function is $-5x_1 - \pi_2 + 27/4$, so that both points $p_1 := (1, 2)^T$ and $p_2 := (\frac{1}{2}, \frac{9}{2})^T$ (as well as all convex combinations) are optimal. The first one comes from the basis $\{2, 3, 5\}$ and the second one from the basis $\{3, 4, 5\}$. For some sufficiently small ε , in the case of $\xi = \frac{5}{2} - \varepsilon$, the first of these bases is optimal, and in the case of $\xi = \frac{5}{2} + \varepsilon$ the second of these bases is optimal.




FIGURE 4. The situation for $\xi = \frac{5}{2}$. Since the x_1 - and π_2 -axes are scaled differently, the outer normal vectors to the lines do not look orthogonal in the figure.

5. Conclusion and outlook

We have presented an enumeration algorithm for the Nash equilibria of non-degenerate rank-1 games which is adapted to the rank-1 structure. Our running time estimate was based on bounding the number of vertices of the polyhedra involved. An open question is in how far one can improve the running time analysis of the algorithm.

A widely open question is to develop enumeration algorithms for games of rank k (for fixed $k \in \mathbb{N}$) exploiting the low-rank structure.

Acknowledgments. Thanks to the reviewers for very helpful comments and corrections.

References

1. V. Aggarwal, *On the generation of all equilibrium points for bimatrix games through the Lemke–Howson algorithm*, Math. Programming **4** (1973), 233–234.
2. C. Audet, P. Hansen, B. Jaumard, and G. Savard, *Enumeration of all extreme equilibria of bimatrix games*, SIAM J. Sci. Comput. **23** (2001), no. 1, 323–338.
3. X. Chen and X. Deng, *Settling the complexity of 2-player Nash equilibrium*, 47th Annual IEEE Symposium on Foundations of Computer Science (Berkeley, CA, 2006), IEEE, Los Alamitos, CA, 2006, pp. 261–272.
4. X. Chen, X. Deng, and S.-H. Teng, *Computing Nash equilibria: approximation and smoothed complexity*, 47th Annual IEEE Symposium on Foundations of Computer Science (Berkeley, CA, 2006), IEEE, Los Alamitos, CA, 2006, pp. 603–612.
5. V. Conitzer and T. Sandholm, *Complexity results about Nash equilibria*, 18th International Joint Conference on Artificial Intelligence (Acapulco, 2003), Morgan Kaufmann Publ., San Francisco, CA, 2003, pp. 765–771.
6. G. B. Dantzig, *Linear programming and extensions*, Princeton Univ. Press, Princeton, NJ, 1963.
7. K. Isaacson and C. B. Millham, *On a class of Nash-solvable bimatrix games and some related Nash subsets*, Naval Res. Logist. Quart. **27** (1980), no. 3, 407–412.
8. M. J. M. Jansen, *Maximal Nash subsets for bimatrix games*, Naval Res. Logist. Quart. **28** (1981), no. 1, 147–152.
9. R. Kannan and T. Theobald, *Games of fixed rank: a hierarchy of bimatrix games*, Proceedings of 18th ACM–SIAM Symposium on Discrete Algorithms (New Orleans, LA, 2007), ACM, New York; SIAM, Philadelphia, PA, 2007, pp. 1124–1132.

10. H. Konno and T. Kuno, *Linear multiplicative programming*, Math. Programming **56** (1992), no. 1, Ser. A, 51–64.
11. H. Konno, P. T. Thach, and H. Tuy, *Optimization on low rank nonconvex structures*, Nonconvex Optim. Appl., vol. 15, Kluwer Acad. Publ. Dordrecht, 1997.
12. C. E. Lemke and J. T. Howson, Jr., *Equilibrium points of bimatrix games*, J. Soc. Indust. Appl. Math. **12** (1964), 413–423.
13. O. L. Mangasarian, *Equilibrium points of bimatrix games*, J. Soc. Indust. Appl. Math. **12** (1964), 778–780.
14. O. L. Mangasarian and H. Stone, *Two-person nonzero-sum games and quadratic programming*, J. Math. Anal. Appl. **9** (1964), 348–355.
15. C. B. Millham, *On Nash subsets of bimatrix games*, Naval Res. Logist. Quart. **21** (1974), 307–317.
16. H. Mills, *Equilibrium points in finite games*, J. Soc. Indust. Appl. Math. **8** (1960), 397–402.
17. B. M. Muhamediev, *Solution of the bilinear programming problem and the determination of all equilibrium situations in bimatrix games*, Zh. Vychisl. Mat. i Mat. Fiz. **18** (1978), no. 2, 351–359 (Russian).
18. J. Nash, *Equilibrium points in n -person games*, Proc. Nat. Acad. Sci. U.S.A. **36** (1950), 48–49.
19. ———, *Non-cooperative games*, Ann. of Math. (2) **54** (1951), 286–295.
20. C. Papadimitriou, *Algorithms, games and the internet*, Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing (Chersonissos, 2001), ACM, New York, 2001, pp. 749–753.
21. L. S. Shapley, *A note on the Lemke–Howson algorithm*, Pivoting and extensions, Math. Programming Stud. (1974), no. 1, 175–189.
22. T. Theobald, *Geometrie und Kombinatorik von Nash-Gleichgewichten*, Math. Semesterber. **52** (2005), no. 2, 221–239.
23. B. von Stengel, *New maximal numbers of equilibria in bimatrix games*, Discrete Comput. Geom. **21** (1999), no. 4, 557–568.
24. ———, *Computing equilibria for two-person games*, Handbook of Game Theory with Economics Applications, Vol. 3 (R. J. Aumann and S. Hart, eds.), North-Holland, Amsterdam, 2002, pp. 1723–1759.
25. N. N. Vorobev, *Equilibrium points in bimatrix games*, Theory Probab. Appl. **3** (1958), 318–331.
26. H.-M. Winkels, *An algorithm to determine all equilibrium points of a bimatrix game*, Game Theory and Related Topics (Proc. Sem., Bonn and Hagen, 1978) (O. Moeschlin and D. Pallaschke, eds.), North-Holland, Amsterdam–New York, 1979, pp. 137–148.

INSTITUT FÜR MATHEMATIK, JOHANN WOLFGANG GOETHE-UNIVERSITÄT, POSTFACH 111932,
FACH 187, 60054 FRANKFURT AM MAIN, GERMANY
E-mail address: `theobald@math.uni-frankfurt.de`