informs
http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 69, No. 2, March–April 2021, pp. 613–631
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Fast Algorithms for Rank-1 Bimatrix Games

Bharat Adsul,[a] Jugal Garg,[b] Ruta Mehta,[c] Milind Sohoni,[a] Bernhard von Stengel[d]

[a] Department of Computer Science and Engineering, Indian Institute of Technology Bombay, Powai, Mumbai 400 076, India; [b] Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana–Champaign, Urbana, Illinois 61801; [c] Department of Computer Science, University of Illinois at Urbana–Champaign, Urbana, Illinois 61801; [d] Department of Mathematics, London School of Economics, London WC2A 2AE, United Kingdom

Contact: adsul@cse.iitb.ac.in (BA); jugal@illinois.edu (JG); rutamehta@cs.illinois.edu (RM); sohoni@cse.iitb.ac.in (MS); b.von-stengel@lse.ac.uk, https://orcid.org/0000-0002-3488-8322 (BvonS)

Received: January 15, 2019
Revised: August 21, 2019
Accepted: November 21, 2019
Published Online in Articles in Advance: June 4, 2020

Subject Classifications: games/group decisions: noncooperative; analysis of algorithms: computational complexity
Area of Review: Methods

https://doi.org/10.1287/opre.2020.1981

Abstract. The rank of a bimatrix game is the matrix rank of the sum of the two payoff matrices. This paper comprehensively analyzes games of rank one and shows the following: (1) For a game of rank r, the set of its Nash equilibria is the intersection of a generically one-dimensional set of equilibria of parameterized games of rank $r-1$ with a hyperplane. (2) One equilibrium of a rank-1 game can be found in polynomial time. (3) All equilibria of a rank-1 game can be found by following a piecewise linear path. In contrast, such a path-following method finds only one equilibrium of a bimatrix game. (4) The number of equilibria of a rank-1 game may be exponential. (5) There is a homeomorphism between the space of bimatrix games and their equilibrium correspondence that preserves rank. It is a variation of the homeomorphism used for the concept of strategic stability of an equilibrium component.

Funding: This material is based upon work supported by the National Science Foundation [Grant CRII 1755619] (J. Garg) and a CAREER award [Computing and Communication Foundations 1750436] (R. Mehta).

Keywords: bimatrix game • Nash equilibrium • rank-1 game • polynomial-time algorithm • homeomorphism

1. Introduction

Noncooperative games are basic economic models. The main concept to analyze them is Nash equilibrium, which recommends to each player a (typically randomized) strategy that is optimal for that player if the other players follow their recommendations. In order to give such a recommendation, a Nash equilibrium (NE) must be found by some method (including any adjustment process). For larger games, this requires computer algorithms. We consider bimatrix games, which are two-player games in strategic form. The algorithm by Lemke and Howson (1964) finds one equilibrium of a bimatrix game. Finding all equilibria is feasible only for small games because of the exponential number of mixed strategies that typically need to be checked for the equilibrium property (Avis et al. 2010).

Kannan and Theobald (2010) introduced a hierarchy of bimatrix games based on the matrix rank of the sum of the two payoff matrices. Games of rank 0 are zero-sum games, which can be solved by linear programming. This paper comprehensively studies games of rank 1. Rank-1 games are economically more interesting than zero-sum games, by allowing a "multiplicative" interaction in addition to an arbitrary zero-sum component (discussed further in Section 10). We will show that, like general bimatrix games, they can have exponentially many disjoint equilibria. On the other hand, as our main results show, they are *computationally tractable*: One equilibrium of a rank-1 game can be found fast (in polynomial time), and finding all equilibria takes comparable time to finding a single equilibrium of a general bimatrix game. Large rank-1 games are therefore attractive as detailed models of interaction, on a similar scale to, but more general than, zero-sum games. Rank-1 bimatrix games and their computational analysis should therefore become a new tool in economic modeling.

The computational complexity (required running time) of computing a Nash equilibrium of a game has received substantial interest in the last two decades. A computational problem is considered tractable if it can be solved in polynomial time. Savani and von Stengel (2006) showed that the algorithm by Lemke and Howson (1964) may have exponential running time. (Their examples require carefully constructed matrices, comparable to linear programs where the simplex algorithm, which otherwise works well in practice, has exponential running time (see Klee and Minty 1972).) The path-following Lemke–Howson algorithm implies that finding an equilibrium of a bimatrix game belongs to the complexity class

(Polynomial Parity Argument in a Directed graph) defined by Papadimitriou (1994, p. 502). PPAD describes certain computational problems where the existence of a solution is known and the problem is to find one explicit solution. (In contrast, the better known complexity class NP applies to decision problems, which are problems that have a "yes" or "no" answer.) Other problems in PPAD include the computation of an approximate Brouwer fixed point, related problems in economics such as market equilibria (Vazirani and Yannakakis 2011), and the computation of an approximate Nash equilibrium of a game with many players. (In games with three or more players, unlike in two-player games, the mixed strategy probabilities in a Nash equilibrium may be irrational num-bers. A suitable concept for such games is approximate Nash equilibrium, and finding an exact Nash equilibrium is an even harder computational problem (see Etessami and Yannakakis 2010).) A celebrated result is that all problems in PPAD can be reduced to finding a Nash equilibrium in a bimatrix game, which makes this problem *PPAD-complete* (Chen and Deng 2006, Chen et al. 2009, Daskalakis et al. 2009). No polynomial-time algorithm for finding a Nash equilibrium of a general bimatrix game is known.

Kannan and Theobald (2010) describe an algorithm to find ε-approximate Nash equilibria in games of fixed rank, with running time that is polynomial in $1/\varepsilon$ and the input length, but exponential in the rank. In the present paper, we prove that an *exact* Nash equilibrium of a rank-1 game can be found in polynomial time. However, we also show that a rank-1 game may have exponentially many equilibria. Moreover, games of higher fixed rank r are PPAD-hard and thus as computationally difficult as general bimatrix games; this has been shown by Mehta (2018) for $r \geq 3$ and is claimed to hold for $r = 2$ (Chen and Paparas 2019). In the context of the "rank" hierarchy, rank-1 games are therefore the most complex type of games that are expected to be computationally tractable.

Section 2 states the notation and preliminary results used in this paper and compares our approach with the work of Theobald (2009). In Section 3, we show that the set of equilibria of a game of rank r is the intersection of a hyperplane with a set of equilibria of parameterized games of rank $r-1$. When $r = 1$, these are parameterized zero-sum games whose equilibria are the solutions to a parameterized linear program (LP). In order to deal with possibly degenerate games that are awkward to handle with pivoting methods, we recall relevant results from Adler and Monteiro (1992) in Section 4. The intersection with the hyperplane gives rise to a polynomial-time binary search for one equilibrium of a rank-1 game, explained in Section 5. In Section 6, we describe completely the set of all Nash equilibria of a rank-1 game and outline a corresponding equilibrium enumeration method.

Section 7 describes an example (which may be useful to consult in between) that illustrates our main results and a second example that shows that binary search fails in general for games of rank 2 or higher. A construction of rank-1 games with exponentially many equilibria is shown in Section 8. In Section 9, we describe a variant of the structure theorem of Kohlberg and Mertens (1986) (KM), which is important for the concept of strategic stability of an equilibrium component. We introduce a new homeomorphism between the space of bimatrix games and its equilibrium correspondence. This homeomorphism preserves the sum of the payoff matrices and hence the rank of the games. In the concluding Section 10, we present a tentative example of an economic model based on rank-1 games and note some open questions.

A preliminary version of our work was published in the *Symposium on Theory of Computing* (Adsul et al. 2011), and the result of Section 8 in von Stengel (2012). The mathematical development in the present paper is almost entirely new in all parts.

2. Bimatrix Games and Best Responses

In this section, we state our notation for bimatrix games and recall the "complementarity" characterization of Nash equilibria in terms of suitable polyhedra. We also briefly compare our approach with Theobald (2009).

We use the following notation. The transpose of a matrix C is written $C^\top$. All vectors are column vectors, so if $x \in \mathbb{R}^m$, then x is an $m \times 1$ matrix and $x^\top$ is the corresponding row vector in $\mathbb{R}^{1\times m}$. In matrix products, scalars are treated like 1×1 matrices. Let $\mathbb{0}$ and $\mathbb{1}$ be vectors with all components equal to 0 and 1, respectively, their dimension depending on the context. Inequalities like $x \geq \mathbb{0}$ hold for all components. The components of a vector $x \in \mathbb{R}^m$ are $x_1, \ldots, x_m$.

For $c \in \mathbb{R}^k$ and $\gamma \in \mathbb{R}$, a *hyperplane* is of the form $\{z \in \mathbb{R}^k \mid c^\top z = \gamma\}$ and a *half space* of the form $\{z \in \mathbb{R}^k \mid c^\top z \leq \gamma\}$. A *polyhedron* is an intersection of finitely many half spaces and called a *polytope* if it is bounded. A *face* of a polyhedron P is of the form $P \cap \{z \in \mathbb{R}^k \mid c^\top z = \gamma\}$, where $P \subseteq \{z \in \mathbb{R}^k \mid c^\top z \leq \gamma\}$. It can be shown that any face of P can be obtained by turning some of the inequalities that define P into equalities (Schrijver 1986, section 8.3). If a face of P consists of a single point, it is called a *vertex* of P. If $S \subseteq X \times Y$ for sets S, X, Y, then $\{x \in X \mid (x, y) \in S \text{ for some } y \in Y\}$ is called the *projection* of S on X, also written as $\{x \mid (x, y) \in S\}$.

A *bimatrix game* is a pair (A, B) of $m \times n$ matrices with rows as pure strategies of player 1 and columns as pure strategies of player 2. The players simultaneously choose their pure strategies, with the corresponding entry of A as payoff to player 1 and of B to player 2. The sets X and Y of *mixed* (that is, randomized) strategies of player 1 and player 2 are given by

$$X = \{x \in \mathbb{R}^m \mid x \geq \mathbb{0},\ \mathbb{1}^\top x = 1\},$$
$$Y = \{y \in \mathbb{R}^n \mid y \geq \mathbb{0},\ \mathbb{1}^\top y = 1\}. \tag{1}$$

For the mixed strategy pair $(x, y) \in X \times Y$, the expected payoffs to the two players are $x^\top Ay$ and $x^\top By$, respectively. A *best response* x of player 1 against y maximizes the player's expected payoff $x^\top Ay$, and a best response y of player 2 against x maximizes the player's expected payoff $x^\top By$. An NE is a pair of mutual best responses.

Consider mixed strategies $x \in X$ and $y \in Y$. If x is a best response to y, then its expected payoff $x^\top Ay$ is clearly at least the payoff $(Ay)_i$ for any pure strategy i of player 1. Moreover, x is a best response to y if and only if any pure strategy i in the *support* of x (that is, where $x_i > 0$) is a pure best response to y (Nash 1951). The following lemma, from Mangasarian (1964), states this *best-response condition* in terms of suitable polyhedra.

Lemma 1. *Let (A, B) be an $m \times n$ bimatrix game. Consider the polyhedra*

$$\overline{P} = \{(x, v) \in X \times \mathbb{R} \mid B^\top x \leq \mathbb{1}v\},$$
$$\overline{Q} = \{(y, u) \in Y \times \mathbb{R} \mid Ay \leq \mathbb{1}u\}. \tag{2}$$

Let $(x, y) \in X \times Y$. Then x is a best response to y if and only if $(y, u) \in \overline{Q}$ and for all rows i

$$x_i = 0 \quad \text{or} \quad (Ay)_i = u \qquad (1 \leq i \leq m), \tag{3}$$

and y is a best response to x if and only if $(x, v) \in \overline{P}$ and for all columns j

$$y_j = 0 \quad \text{or} \quad (B^\top x)_j = v \qquad (1 \leq j \leq n). \tag{4}$$

If both conditions hold, then u and v are the unique payoffs to player 1 *and* 2 *in the Nash equilibrium (x, y).*

A bimatrix game is *degenerate* if there is a mixed strategy that has more pure best responses than the size of its support (von Stengel 2002). A degenerate game may have infinite sets of equilibria. They can be described by suitable faces of $\overline{P}$ and $\overline{Q}$, as explained further in Section 6. Our analysis applies to general games that may be degenerate.

The object of study of our paper is bimatrix games of fixed *rank*, introduced by Kannan and Theobald (2010). They generalize zero-sum games, which are games of rank zero.

Definition 1. The *rank* of a bimatrix game (A, B) is the matrix rank of $A + B$.

For comparison of our approach with Theobald (2009), we consider a quadratic program, from Mangasarian and Stone (1964), that captures the NE of (A, B).

Lemma 2. *The strategy pair (x, y) is a Nash equilibrium of (A, B) if and only if it is a solution to*

$$\underset{x,y,u,v}{\text{maximize}}\ x^\top(A + B)y - u - v$$
$$\text{subject to} \quad (x, v) \in \overline{P}, \quad (y, u) \in \overline{Q}. \tag{5}$$

The optimum value of (5) *is zero, with $u = x^\top Ay$ and $v = x^\top By$.*

Proof. Consider any solution to (5). Then v is at least the best-response payoff of player 2 against x because $(x, v) \in \overline{P}$, and u is at least the best-response payoff of player 1 against y because $(y, u) \in \overline{Q}$. Hence, $x^\top(A + B)y - u - v \leq 0$. Furthermore, (3) and (4) imply that $x^\top(A + B)y - u - v$ is zero if and only if (x, y) is an NE, in which case $u = x^\top Ay$ and $v = x^\top By$. □

The quadratic program (5) shows the importance of the rank of the matrix $A + B$. For zero-sum games, the rank of $A + B$ is zero and (5) is a linear program, a well-known fact (Dantzig 1963). For a rank-1 game (A, B) with $A + B = ab^\top$, the bilinear term $x^\top(A + B)y$ in the objective function becomes the product $(x^\top a)(b^\top y)$ of two linear terms. The resulting optimization problem is called a *linear multiplicative program*. Solving a general linear multiplicative program is NP-hard (Matsui 1996).

Consider a rank-1 game (A, B), where $A + B = ab^\top$. Similar to parametric simplex methods for solving linear multiplicative programs (Konno et al. 1991), Theobald (2009) describes an algorithm to enumerate all equilibria of (A, B). For a real parameter ξ, he considers the parameterized LP:

$$\underset{x,y,u,v}{\text{maximize}}\ x^\top a\xi - u - v$$
$$\text{subject to} \quad (x, v) \in \overline{P}, \quad (y, u) \in \overline{Q}, \quad b^\top y = \xi. \tag{6}$$

In any solution to (6), $x^\top a\xi = x^\top ab^\top y = x^\top(A + B)y$. Hence, by Lemma 2, any optimal solution to (6) is an equilibrium of (A, B) if and only its optimum is zero. Moreover, $b^\top y = \xi$ implies that ξ is a convex combination of the components $b_1, \ldots, b_n$ of b, so that one can restrict ξ to the interval $[\min\{b_1, \ldots, b_n\}, \max\{b_1, \ldots, b_n\}]$. By partitioning this interval into segments, where (6) uses the same basic variables, Theobald (2009) obtains an enumeration of all NEs of (A, B).

Our approach is somewhat similar, with a parameter λ and the equality $x^\top a = \lambda$. However, we consider a different LP that is parameterized by λ and involves only the payoff matrix A and the vector b used in $A + B = ab^\top$. That LP, given in (19) below, has x as primal and y as dual variables, whereas in (6) both x and y are primal with less closely related constraints. We consider the hyperplane defined by $x^\top a = \lambda$ *separately* from the parameterized LP. The intersection of the hyperplane with the solutions to the parameterized

LP defines the equilibria of the rank-1 game. This structural insight can be used both for finding an exact NE in polynomial time by binary search (see Section 5) and for enumerating all equilibria (see Section 6). As a topic for further research, it may be interesting if this approach can be extended to more general linear multiplicative programs.

3. Rank Reduction

The central result of this short section is Theorem 1. It states that the set of Nash equilibria of a game of rank r is the intersection of a set $\mathcal{N}$ of equilibria of parameterized games of rank $r-1$ with a suitable hyperplane. In subsequent sections, we show how to exploit this property algorithmically when $r = 1$.

The following lemma states the well-known fact that the equilibria of a bimatrix game are unchanged when subtracting a separate constant b_j from each column j of the row player's payoff matrix. Call two games *strategically equivalent* if they have the same best responses to mixed strategies.

Lemma 3. *If $b \in \mathbb{R}^n$, then the $m \times n$ game (A, B) is strategically equivalent to the game $(A - \mathbb{1}b^\top, B)$.*

Proof. This holds by Lemma 1, because the best-response payoff u to player 1 in the game (A, B) changes to $u - b^\top y$ in $(A - \mathbb{1}b^\top, B)$: Clearly, $Ay \le \mathbb{1}u$ is equivalent to $(A - \mathbb{1}b^\top)y \le \mathbb{1}(u - b^\top y)$, and $(Ay)_i = u$ is equivalent to $((A - \mathbb{1}b^\top)y)_i = u - b^\top y$. □

Lemma 4. *An $m \times n$ bimatrix game of positive rank r can be written as $(A, C + ab^\top)$ for suitable $a \in \mathbb{R}^m$, $b \in \mathbb{R}^n$, and a game (A, C) of rank $r - 1$.*

Proof. An $m \times n$ matrix is of rank at most r if and only if it can be written as the sum of r rank-1 matrices, that is, as $a_1 b_1^\top + \cdots + a_r b_r^\top$ for suitable $a_q \in \mathbb{R}^m$ and $b_q \in \mathbb{R}^n$ for $1 \le q \le r$. This is easily seen by writing the jth column of the matrix as $\sum_{q=1}^r a_q b_{qj}$ and letting $b_q^\top = (b_{q1}, \ldots, b_{qn})$ (see also Wardlaw 2005). Suppose (A, B) is of rank r, with $A + B = \sum_{q=1}^r a_q b_q^\top$ and therefore $B = -A + \sum_{q=1}^r a_q b_q^\top$. Let $C = -A + \sum_{q=1}^{r-1} a_q b_q^\top$ and $a = a_r$, $b = b_r$, so that $B = C + ab^\top$; obviously, $A + C$ is of rank $r - 1$. □

The following is a simple but central lemma.

Lemma 5. *Let $A, C \in \mathbb{R}^{m\times n}$, $x \in X$, $y \in Y$, $a \in \mathbb{R}^m$, $b \in \mathbb{R}^n$, $\lambda \in \mathbb{R}$. The following are equivalent:*

a. *(x, y) is an equilibrium of $(A, C + ab^\top)$,*
b. *(x, y) is an equilibrium of $(A, C + \mathbb{1}\lambda b^\top)$ and $x^\top a = \lambda$, and*
c. *(x, y) is an equilibrium of $(A - \mathbb{1}\lambda b^\top, C + \mathbb{1}\lambda b^\top)$ and $x^\top a = \lambda$.*

Proof. The equivalence of (a) and (b) holds because the players get in both games the same expected payoffs for their pure strategies: this is immediate for player 1; if $x^\top a = \lambda$, then the column payoffs are given by

$$x^\top(C + ab^\top) = x^\top C + \lambda b^\top = x^\top C + x^\top \mathbb{1}\lambda b^\top = x^\top(C + \mathbb{1}\lambda b^\top). \tag{7}$$

The games in (b) and (c) are strategically equivalent by Lemma 3. □

Consider a game (A, B) of positive rank r, where $B = C + ab^\top$ so that (A, C) is a game of rank $r - 1$ according to Lemma 4. Then the game $(A - \mathbb{1}\lambda b^\top, C + \mathbb{1}\lambda b^\top)$ in Lemma 5(c) has the same sum $A + C$ of its payoff matrices and hence also rank $r - 1$, for any choice of the parameter λ. Let $\mathcal{N}$ be the set of Nash equilibria together with λ of these parameterized games,

$$\mathcal{N} = \{(\lambda, x, y) \in \mathbb{R} \times X \times Y \mid (x, y) \text{ is a NE of } (A - \mathbb{1}\lambda b^\top, C + \mathbb{1}\lambda b^\top)\}, \tag{8}$$

where by Lemma 5(b),

$$\mathcal{N} = \{(\lambda, x, y) \in \mathbb{R} \times X \times Y \mid (x, y) \text{ is a NE of } (A, C + \mathbb{1}\lambda b^\top)\}. \tag{9}$$

These considerations imply the following main result of this section.

Theorem 1. *Given a bimatrix game $(A, C + ab^\top)$, its set of Nash equilibria is exactly the projection on $X \times Y$ of the intersection of $\mathcal{N}$ and the hyperplane H defined by*

$$H = \{(\lambda, x, y) \in \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^n \mid x^\top a = \lambda\}. \tag{10}$$

Theorem 1 asserts that for any rank-r game of the form $(A, C + ab^\top)$, every Nash equilibrium of the game is captured by the set $\mathcal{N}$ in (8) of games of rank $r - 1$, which are parameterized by λ, intersected with the hyperplane H in (10). Can this rank reduction be leveraged to get an efficient algorithm to find a Nash equilibrium for a game of arbitrary constant rank? As will be discussed in Section 7, this does not work in general. However, it does work for rank-1 games.

4. Parameterized Linear Programs

Our aim is to describe the equilibria of rank-1 games $(A, -A + ab^\top)$ using the rank reduction of the previous section. For this, we consider the set $\mathcal{N}$ in (9) for $C = -A$,

$$\mathcal{N} = \{(\lambda, x, y) \in \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^n \mid (x, y) \text{ is a NE of } (A, -A + \mathbb{1}\lambda b^\top)\}, \tag{11}$$

where by (8),

$$\mathcal{N} = \{(\lambda, x, y) \in \mathbb{R} \times \mathbb{R}^m \times \mathbb{R}^n \mid (x, y) \text{ is a NE of } (A - \mathbb{1}\lambda b^\top, -A + \mathbb{1}\lambda b^\top)\}, \tag{12}$$

which is the set of equilibria of zero-sum games parameterized by λ. These correspond to the solutions of a parameterized LP. In this section, we review the structure of such parameterized LPs with a particular view toward nongeneric cases and polynomial-time algorithms as studied by Adler and Monteiro (1992). In essence, such parameterized LPs have finitely many special values of the parameter λ called *breakpoints*. These separate the set $\mathcal{N}$ into a connected sequence of polyhedral *segments* (which generically are line segments). They are described in Theorem 3 in the next section, where we will present a polynomial-time algorithm for finding one equilibrium of a rank-1 game. In the subsequent section, we present another algorithm for finding all equilibria.

We assume familiarity with notions of linear programming, such as LP duality and complementary slackness (see, for example, Schrijver 1986). The following well-known lemma (Dantzig 1963) states that the equilibria of a zero-sum game are the primal and dual solutions to an LP.

Lemma 6. *Consider an $m \times n$ zero-sum game $(M, -M)$. In any equilibrium (x, y) of this game, y is a minmax strategy of player 2, which is a solution to the LP with variables y in $\mathbb{R}^n$ and u in $\mathbb{R}$:*

$$\begin{aligned} &\underset{y,u}{\text{maximize}} \quad u \\ &\text{subject to} \quad My + \mathbb{1}u \le \mathbb{0}, \quad y \in Y, \end{aligned} \tag{13}$$

and x is a maxmin strategy of player 1, which is a solution to the dual LP to (13). For the optimal value of u in (13), the maxmin payoff to player 1 and minmax cost to player 2 and hence the value of the game is $-u$.

Proof. The dual LP to (13) has variables $x \in \mathbb{R}^m$ and $v \in \mathbb{R}$ and states

$$\begin{aligned} &\underset{x,v}{\text{minimize}} \quad v \\ &\text{subject to} \quad x^\top M + v\mathbb{1}^\top \ge \mathbb{0}^\top, \quad x \in X. \end{aligned} \tag{14}$$

Both LPs are feasible (with sufficiently small u and large v). Let (y, u) be an optimal solution to (13) and (x, v) to (14). Then $u = v$ by LP duality, and (13) and (14) state $My \le \mathbb{1}(-u)$, that is, player 2 pays no more than $-u$ for any row, and $x^\top M \ge (-v)\mathbb{1}^\top$, that is, player 1 gets at least $-v$ in every column, where $-u = -v$, which is therefore the value of the game.

With the dual constraints written as $x^\top(-M) \le v\mathbb{1}^\top$, the complementary slackness conditions between the primal and the dual are exactly the Nash equilibrium conditions (3) and (4) of Lemma 1 (except for the changed sign of u so that we do not have to write $x \in X$ in (14) as $-\mathbb{1}^\top x = -1$ and $x \ge \mathbb{0}$). Hence, (x, y) is a Nash equilibrium. □

Applied to $M = A - \mathbb{1}\lambda b^\top$ in (12), the LP (13) in Lemma 6 says

$$\begin{aligned} &\underset{y,u}{\text{maximize}} \quad u \\ &\text{subject to} \quad \left(A - \mathbb{1}\lambda b^\top\right)y + \mathbb{1}u \le \mathbb{0}, \ y \in Y. \end{aligned} \tag{15}$$

In (15), the matrix A is parameterized. The substitution $u = \lambda b^\top y + t$ gives the equivalent LP, where only the objective function is parameterized:

$$\begin{aligned} &\underset{y,t}{\text{maximize}} \quad \lambda b^\top y + t \\ &\text{subject to} \quad Ay + \mathbb{1}t \le \mathbb{0}, \quad y \in Y. \end{aligned} \tag{16}$$

This is a standard parameterized linear programming problem. We stay close to the notation of Adler and Monteiro (1992) who consider a primal LP with minimization subject to equality constraints, variables x, and a parameterized right-hand side, of which (16) is the dual, a maximization problem subject to inequalities, with variables y, and a parameterized objective function. We write (16) as

$$\begin{aligned} D_\lambda: \quad &\underset{y,t}{\text{maximize}} \quad \lambda b^\top y + t \\ &\text{subject to} \quad (y, t) \in D, \end{aligned} \tag{17}$$

with the fixed polyhedron

$$D = \{ (y, t) \in \mathbb{R}^n \times \mathbb{R} \mid \begin{array}{ll} Ay + \mathbb{1}t & \le \mathbb{0} \\ \mathbb{1}^\top y & = 1 \\ y & \ge \mathbb{0} \ \}. \end{array} \tag{18}$$

The LP D_λ is the dual of the following LP P_λ with a parameterized right-hand side, where we use slack variables $s \in \mathbb{R}^n$ to express the inequality $A^\top x + \mathbb{1}v \ge b\lambda$ as an equality, in line with Adler and Monteiro (1992):

$$\begin{aligned} P_\lambda: \quad &\underset{x,v,s}{\text{minimize}} \quad v \\ &\text{subject to} \quad A^\top x + \mathbb{1}v - s = b\lambda \\ &\qquad\qquad\quad \mathbb{1}^\top x = 1 \\ &\qquad\qquad\quad x, \quad s \ge \mathbb{0}. \end{aligned} \tag{19}$$

For optimal solutions (y, t) to D_λ and (x, v, s) to P_λ, we have $\lambda b^\top y + t = v$. The next lemma (essentially a corollary to Lemma 5 and Lemma 6) states that $-t$ and v can be interpreted as the players' payoffs for the games in Lemma 5 (a) and (b), and asserts that t, v, s are uniquely determined by (λ, x, y) (that is, a point on $\mathcal{N}$).

Lemma 7. *Let $\lambda \in \mathbb{R}$. Then (x, y) is an equilibrium of the game $(A, -A + \mathbb{1}\lambda b^\top)$ if and only if (y, t) is an optimal solution to D_λ in (17) for some t, which is uniquely determined by y, and (x, v, s) is an optimal solution to P_λ in (19) for some v and s, which are uniquely determined by λ and x. The equilibrium payoffs are $-t$ to player 1 and v*

to player 2. If $x^\top a = \lambda$, these are also the payoffs in the game $(A, -A + ab^\top)$, and (x, y) is an equilibrium of that game.

Proof. By Lemma 5 with $C = -A$, the game $(A, -A + ab^\top)$ has the same equilibria (x, y) and, by (7), payoffs as the game $(A, -A + \mathbb{1}\lambda b^\top)$ if $x^\top a = \lambda$. Consider any optimal solutions (y, t) to D_λ and (x, v, s) to P_λ. Then $Ay + \mathbb{1}t \le \mathbb{0}$ states for each row i of A the inequality $(Ay)_i \le -t$. Complementary slackness, equivalent to LP optimality, states that $(Ay)_i = -t$ whenever $x_i > 0$. This is the equilibrium condition in (3) that states that x is a best response to y. Because it holds for at least one i, it uniquely determines $-t$, which is the equilibrium payoff to player 1 in the above games.

Similarly, the constraint $s = A^\top x - b\lambda + \mathbb{1}v$ in (19) means that s is determined by (x, λ, v), and states $s_j = (A^\top x - b\lambda)_j + v \ge 0$ for all j, or equivalently $((-A^\top + b\lambda\mathbb{1}^\top)x)_j \le v$. Complementary slackness, equivalent to LP optimality, states that this inequality is tight whenever $y_j > 0$. This is the condition (4) that states that y is a best response to x in the game $(A, -A + \mathbb{1}\lambda b^\top)$, and it uniquely determines v as the equilibrium payoff to player 2. □

Primal-dual pairs P_λ, D_λ of LPs with a parameter λ have been studied since Gass and Saaty (1955). The next result is well known, which we show following Jansen et al. (1997).

Lemma 8. *For $\lambda \in \mathbb{R}$, let $\phi(\lambda)$ be the optimum value of P_λ and hence of D_λ. Then $\phi : \mathbb{R} \to \mathbb{R}$ is the pointwise maximum of a finite number of affine functions on $\mathbb{R}$ and therefore piecewise linear and convex.*

Proof. The optimum of D_λ exists for any λ and is taken at a vertex of the polyhedron D in (18). Let V be the set of vertices of D, which is finite. Hence,

$$\phi(\lambda) = \max\{\lambda(b^\top y) + t \mid (y, t) \in V\}, \tag{20}$$

where for each of the finitely many (y, t) in V the function $\lambda \mapsto \lambda(b^\top y) + t$ is affine. Hence, ϕ is the pointwise maximum of a finite number of affine functions as claimed. The epigraph of ϕ given by $E = \{(\lambda, \theta) \mid \theta \ge \phi(\lambda)\}$ is the intersection of the convex epigraphs of these affine functions, so E is convex and ϕ is a convex function. □

By (20), the function $\phi(\lambda)$ is the "upper envelope" of the affine functions $\lambda \mapsto \lambda(b^\top y) + t$ defined by the vertices (y, t) of D. A *breakpoint* is any λ^* so that $\phi(\lambda)$ has different left and right derivatives when λ approaches λ^* from below or above, denoted by $\phi'_-(\lambda^*)$ and $\phi'_+(\lambda^*)$, respectively.

For any LP L, say, let *OptFace*(L) be the face of the domain of L where its optimum is attained. For any λ, we denote *OptFace*(D_λ) by $Y(\lambda)$, that is,

$$Y(\lambda) = \{(y, t) \in D \mid \lambda b^\top y + t = \phi(\lambda)\}. \tag{21}$$

Then the left and right derivatives of ϕ at λ are characterized as follows (obvious from (20), also proposition 2.4 of Adler and Monteiro 1992):

$$\begin{aligned} \phi'_-(\lambda) &= \min\{b^\top y \mid (y, t) \in Y(\lambda)\}, \\ \phi'_+(\lambda) &= \max\{b^\top y \mid (y, t) \in Y(\lambda)\}, \end{aligned} \tag{22}$$

which are the optima of the two LPs

$$\begin{aligned} SL^{\min}(\lambda):\ &\underset{y,t}{\text{minimize}} && b^\top y \\ &\text{subject to} && (y, t) \in Y(\lambda), \\ SL^{\max}(\lambda):\ &\underset{y,t}{\text{maximize}} && b^\top y \\ &\text{subject to} && (y, t) \in Y(\lambda). \end{aligned} \tag{23}$$

That is, λ^* is a breakpoint if and only if $\phi'_-(\lambda^*) < \phi'_+(\lambda^*)$. Clearly, in that case there are at least two vertices (y, t) and $(\hat{y}, \hat{t})$ of D that define two different affine functions $\lambda \mapsto \lambda(b^\top y) + t$ and $\lambda \mapsto \lambda(b^\top \hat{y}) + \hat{t}$ that meet at $\lambda = \lambda^*$ to define the maximum $\phi(\lambda^*)$ in (20). These are also vertices of $Y(\lambda^*)$, which is then a higher-dimensional face (such as an edge) of D. The following central observation shows that the breakpoints give all the information about the optimal faces $Y(\lambda)$ of D_λ for any λ between these breakpoints.

Theorem 2 (Adler and Monteiro 1992, theorem 4.1). *Let $\lambda_1, \ldots, \lambda_K$ be the breakpoints, in increasing order, for the parameterized LPs P_λ and D_λ, and let $\lambda_0 = -\infty$ and $\lambda_{K+1} = \infty$. For $0 \le k \le K$, consider any $\lambda'_k \in (\lambda_k, \lambda_{k+1})$. Then $Y(\lambda'_k) = OptFace(SL^{\max}(\lambda_k))$ for $1 \le k \le K$, and $Y(\lambda'_k) = OptFace(SL^{\min}(\lambda_{k+1}))$ for $0 \le k \le K - 1$.*

For finding the solutions to P_λ as a function of λ, the *nondegenerate* case is straightforward, where $Y(\lambda)$ is a vertex of D_λ unless λ is a breakpoint, in which case $Y(\lambda)$ is an edge of D_λ. Then these vertices uniquely describe the pieces of the piecewise linear function $\phi(\lambda)$ and can be traversed by a parameterized simplex algorithm (Gass and Saaty 1955). An example is shown in the right diagram of Figure 4 below with the Constraints (44) for $Ay + \mathbb{1}t \le \mathbb{0}$ in D, with the additional constraints $0 \le y_2 \le 1$ to represent $y \in Y$, and objective function $\lambda b^\top y + t$ given by $\lambda(1 - 2y_2) + t$. The three linear parts of $\phi(\lambda)$ are

$$\phi(\lambda) = \begin{cases} -\lambda - 1 & \text{for } \lambda \le -\frac{1}{2} \\ -\frac{1}{2} & \text{for } -\frac{1}{2} \le \lambda \le \frac{1}{2} \\ \lambda - 1 & \text{for } \frac{1}{2} \le \lambda, \end{cases} \tag{24}$$

which correspond to the optimal vertices (y_2, t) of D given by $(1, -1)$, $(\frac{1}{2}, -\frac{1}{2})$, and $(0, -1)$. The two breakpoints are $\lambda_1 = -\frac{1}{2}$ and $\lambda_2 = \frac{1}{2}$, which correspond to the two edges of D.

In the degenerate case, one typically does not get polynomial-time algorithms by considering vertices and corresponding basic solutions to the LP P_λ as in a parameterized simplex algorithm. Instead of partitioning the variables of P_λ into basic and nonbasic variables, Adler and Monteiro (1992, p. 164) consider "optimal partitions"; we use here only the partition part that replaces the nonbasic variables, which we denote by $M(\lambda) \cup N(\lambda)$ in (26) below (called $N(\lambda)$ in Adler and Monteiro 1992). This is the set of variables of the dual LP D_λ that may be strictly positive in an optimal solution, which represent the "true inequalities" of $Y(\lambda)$.

Definition 2. For some A, b, C, d, suppose that the constraints in x

$$Ax \le b, \qquad Cx = d \tag{25}$$

are feasible. Then any row i of $Ax \le b$ so that $(b - Ax)_i > 0$ for some feasible x is called a *true inequality* of (25).

If there are solutions x and $\hat{x}$ to (25) so that $(b - Ax)_i > 0$ and $(b - A\hat{x})_j > 0$, then both inequalities are true for $x\frac{1}{2} + \hat{x}\frac{1}{2}$, so there is a unique largest set of true inequalities with some feasible solution where all these strict inequalities hold simultaneously. These define the relative interior of the polyhedron defined by (25).

Let $A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^n$. Let $M(\lambda) \cup N(\lambda)$ be the set of true inequalities of the optimal face $Y(\lambda)$ of D_λ in (17), that is,

$$\begin{aligned} M(\lambda) &= \{\, i \in \{1, \ldots, m\} \mid (Ay)_i + t < 0 \\ &\qquad \text{for some } (y, t) \in Y(\lambda)\,\}, \\ N(\lambda) &= \{\, j \in \{1, \ldots, n\} \mid y_j > 0 \\ &\qquad \text{for some } (y, t) \in Y(\lambda)\,\}. \end{aligned} \tag{26}$$

Any nontrue inequality of $Y(\lambda)$ is always tight, that is, $(Ay)_i + t = 0$ if $i \notin M(\lambda)$ and $y_j = 0$ if $j \notin N(\lambda)$. It can be shown that for such i and j, there are optimal solutions (x, v, s) to P_λ, where $x_i > 0$ and $s_j > 0$, so these are the true inequalities of $OptFace(P_\lambda)$. This is also known as "strict complementary slackness" (Schrijver 1986, p. 95, condition (36)). Consider the polyhedron P of the constraints for P_λ in (19), where λ is allowed to vary,

$$P = \big\{(\lambda, x, v, s) \in \mathbb{R} \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^n \mid A^\top x + \mathbb{1}v - s = b\lambda, \\ x \in X,\ s \ge \mathbb{0}\big\}. \tag{27}$$

The following lemma considers the face of P defined by the equations $x_i = 0$ for $i \in M(\lambda)$ and $s_j = 0$ for $j \in N(\lambda)$, which are necessary and sufficient for a feasible solution to P_λ to be optimal. This is immediate from the standard complementary slackness condition.

Lemma 9. *Let $A \in \mathbb{R}^{m\times n}$ and $b \in \mathbb{R}^n$. For $M \subseteq \{1, \ldots, m\}$ and $N \subseteq \{1, \ldots, n\}$, with $x_M = (x_i)_{i\in M}$ and $s_N = (s_j)_{j\in N}$, define*

$$P(M, N) = \{(\lambda, x, v, s) \in P \mid x_M = \mathbb{0},\ s_N = \mathbb{0}\}. \tag{28}$$

Then any feasible solution (x, v, s) to P_λ is optimal if and only if $(\lambda, x, v, s) \in P(M(\lambda), N(\lambda))$.

Crucially, according to Theorem 2, for any λ in an open interval $(\lambda_k, \lambda_{k+1})$ (for $0 \le k \le K$), the optimal face $Y(\lambda)$ is constant in λ. Hence, for all $\lambda \in (\lambda_k, \lambda_{k+1})$, the true inequalities $(M(\lambda), N(\lambda))$ of $Y(\lambda)$ are equal to some fixed (M, N), and for the points (λ, x, v, s) in $P(M, N)$ the value of λ can be any real in the *closed* interval $[\lambda_k, \lambda_{k+1}]$. Namely, with the LPs

$$\begin{aligned} BR^{\max}(M,N):\quad &\underset{\lambda,x,v,s}{\text{maximize}}\ \lambda \\ &\text{subject to}\ (\lambda, x, v, s) \in P(M, N), \\ BR^{\min}(M,N):\quad &\underset{\lambda,x,v,s}{\text{minimize}}\ \lambda \\ &\text{subject to}\ (\lambda, x, v, s) \in P(M, N), \end{aligned} \tag{29}$$

the following holds.

Lemma 10. *Consider $\lambda_0, \lambda_1, \ldots, \lambda_K, \lambda_{K+1}$ and $\lambda'_k \in (\lambda_k, \lambda_{k+1})$ for $0 \le k \le K$ as in Theorem 2. Let $M^k = M(\lambda'_k)$ and $N^k = N(\lambda'_k)$ (which do not depend on the choice of λ'_k). Then for $1 \le k \le K$,*

a. *the breakpoint λ_k is the optimum of the LP $BR^{\max}(M^{k-1}, N^{k-1})$ and of the LP $BR^{\min}(M^k, N^k)$; and*

b. *if $(\lambda, x, v, s) \in P(M(\lambda_k), N(\lambda_k))$, then $\lambda = \lambda_k$.*

Proof. See Adler and Monteiro (1992), p. 171, for (a) and theorem 3.1(a) and lemma 3.1(b) for (b). □

Lemma 10(a) implies that for any λ in the open interval $(\lambda_k, \lambda_{k+1})$, for $1 \le k \le K - 1$, the endpoints of the closed interval $[\lambda_k, \lambda_{k+1}]$ are given by the minimum and maximum of λ for $(\lambda, x, v, s) \in P(M, N)$, where $M = M(\lambda)$ and $N = N(\lambda)$. Lemma 10(b) and Lemma 9 imply that if λ is itself a breakpoint, then $P(M, N) = \{\lambda\} \times OptFace(P_\lambda)$.

As we will describe in detail in the next section, Theorem 2 and Lemma 10 lead to a description of the set of optimal solutions to P_λ and D_λ for all λ with the help of the breakpoints $\lambda_1, \ldots, \lambda_K$ in the form of $2K + 1$ polyhedral *segments* (which are lines in the nondegenerate case). Any solution (x, v, s) to P_λ is optimal if and only if (λ, x, v, s) belongs to $P(M(\lambda), N(\lambda))$, which is a face of P, and any solution to D_λ is optimal if and only if it belongs to $Y(\lambda)$, which is a face of D. For λ between two breakpoints, these faces do not change (but x typically varies with λ), and their Cartesian product defines $K + 1$ of the segments. If λ is equal to a breakpoint, the set $P(M(\lambda), N(\lambda))$ is a subset of the two adjoining faces $P(M(\lambda'), N(\lambda'))$ for λ' near λ, whereas $Y(\lambda)$ is a superset of the adjoining faces $Y(\lambda')$, as described in Theorem 2. This defines the other K segments. Using this we will give a precise description of the set $\mathcal{N}$ in Theorem 3 below.

Adler and Monteiro (1992) describe how to generate the breakpoints of P_λ, D_λ in polynomial time per breakpoint, with a polynomial-time algorithm applied to the LPs (17), (23), (29), which we will adapt to our purpose. (However, the number of breakpoints may be exponential, see Murty (1980).) The true inequalities in Definition 2 can also be found with an LP, according to the following lemma (proposition 4.1 of Adler and Monteiro (1992)), due to Freund et al. (1985); for an alternative polynomial-time algorithm see Mehrotra and Ye (1993).

Lemma 11. *For A, b, C, d, and the Constraints (25), consider the LP*

$$\begin{aligned} \underset{x,u,\alpha}{\text{maximize}} \quad & \mathbb{1}^\top u \\ \text{subject to} \quad & Ax + u - b\alpha \le \mathbb{0}, \\ & Cx - d\alpha = 0, \\ & \mathbb{0} \le u \le \mathbb{1}, \\ & \alpha \ge 1. \end{aligned} \tag{30}$$

Then (25) is feasible if and only if (30) is feasible and bounded, and any optimal solution (x^*, u^*, α^*) to (30) satisfies $u_i^* = 1$ (and $u_i^* = 0$ otherwise) if and only if i is a true inequality of (25). For such an optimal solution (x^*, u^*, α^*) to (30), $x = x^*(1/\alpha^*)$ is a solution to (25), where $(b - Ax)_i > 0$ for all true inequalities i.

Proof. If the LP (30) is feasible, then it is also bounded because $u \le \mathbb{1}$. Let I be the set of true inequalities of (25), that is, $(b - Ax)_i = \varepsilon_i > 0$ for $i \in I$ for some x with $Cx = d$. Choose $\alpha^* \ge 1$ so that $\alpha^* \ge 1/\varepsilon_i$ for all $i \in I$. Then $(b\alpha^* - A(x\alpha^*))_i = (b - Ax)_i\alpha^* = \varepsilon_i\alpha^* \ge 1$ for $i \in I$. Hence, $x^* = x\alpha^*$ and u^* defined by $u_i^* = 1$ if $i \in I$, and $u_i^* = 0$ otherwise, give a feasible solution (x^*, u^*, α^*) to the LP (30). This solution is also optimal because any solution $(\hat{x}, \hat{u}, \hat{\alpha})$ to (30) where $\hat{u}_i > 0$ would give a solution $x = \hat{x}(1/\hat{\alpha})$ to (25) with $(b - A\hat{x})_i > 0$ and thus $i \in I$, so for any feasible solution (x, u, α) to (30), we have $u_i = 0$ whenever $i \notin I$. This proves the claim. □

5. Finding One Equilibrium of a Rank-1 Game by Binary Search

We use the results of the previous section to present a polynomial-time algorithm for finding one equilibrium of a rank-1 game $(A, -A + ab^\top)$, using binary search for a suitable value of the parameter λ in Theorem 1. The search maintains a pair of successively closer parameter values and corresponding equilibria of the game $(A, -A + \mathbb{1}\lambda b^\top)$ that are on opposite sides of the hyperplane H in (10). Generically, the set $\mathcal{N}$ in (11) is a piecewise linear path that has to intersect H between these two parameter values. In general, the *segments* of that "path" are products of certain faces of the polyhedra D in (17) and P in (27) described in Theorem 2 and Lemma 10 using the breakpoints $\lambda_1, \ldots, \lambda_K$ of the LPs P_λ and D_λ.

We give a complete description of $\mathcal{N}$ in terms of these faces of P and D, which we project to $\mathbb{R} \times X$ (for the possible values of (λ, x)) and Y. Namely, consider $\lambda_0, \lambda_1, \ldots, \lambda_K, \lambda_{K+1}$ and $\lambda'_k \in (\lambda_k, \lambda_{k+1})$ for $0 \le k \le K$ as in Theorem 2. For $0 \le k \le K$, define

$$X'_k = \{(\lambda, x) \mid (\lambda, x, v, s) \in P(M(\lambda'_k), N(\lambda'_k))\}. \tag{31}$$

Note that for any $(\lambda, x, v, s) \in P(M(\lambda'), N(\lambda'))$ (for any $\lambda' \in \mathbb{R}$), the components v and s are uniquely determined by (λ, x) by Lemma 7. Similarly, let

$$Y'_k = \{y \mid (y, t) \in Y(\lambda'_k)\}, \tag{32}$$

where again t in (y, t) is uniquely determined by y. Recall that the choice of $\lambda'_k \in (\lambda_k, \lambda_{k+1})$ does not matter for the definitions of X'_k and Y'_k. The polyhedra $X'_k \times Y'_k$ for $0 \le k \le K$ (which for $k = 0$ and $k = K + 1$ are infinite, otherwise bounded) represent $K + 1$ of the segments that constitute $\mathcal{N}$ between any two breakpoints λ_k and λ_{k+1}. They are successively connected by K further segments, which are polytopes $X_k \times Y_k$ that correspond to the breakpoints themselves. These are for $1 \le k \le K$ defined by

$$X_k = \{(\lambda, x) \mid (\lambda, x, v, s) \in P(M(\lambda_k), N(\lambda_k))\} \tag{33}$$

and

$$Y_k = \{y \mid (y, t) \in Y(\lambda_k)\}. \tag{34}$$

Theorem 3. *The set $\mathcal{N}$ in (11) is given by*

$$\mathcal{N} = (X'_0 \times Y'_0) \;\cup\; \bigcup_{k=1}^{K} ((X_k \times Y_k) \cup (X'_k \times Y'_k)), \tag{35}$$

where for $1 \le k \le K$ we have

$$Y_k \supseteq Y'_{k-1} \cup Y'_k \tag{36}$$

and

$$X_k \subseteq X'_{k-1} \cap X'_k. \tag{37}$$

Proof. This follows from Lemma 7, Lemma 9, and Theorem 2. By Theorem 2, $Y(\lambda'_k)$ is the optimal face of $SL^{\max}(\lambda_k)$, which is a subset of $Y(\lambda_k)$. Hence, $Y'_k \subseteq Y_k$, and similarly $Y'_{k-1} \subseteq Y_k$, which implies (36). In addition, we have $M(\lambda'_k) \subseteq M(\lambda_k)$ and $N(\lambda'_k) \subseteq N(\lambda_k)$ and thus $X_k \subseteq X'_k$ because of the additional tight constraints in $P(M(\lambda_k), N(\lambda_k))$. Similarly, $X_k \subseteq X'_{k-1}$. This shows (37). □

The preceding characterization of $\mathcal{N}$ is used in the following lemma.

Lemma 12. *Let $\underline{\lambda} \le \overline{\lambda}$ and $\underline{x}, \overline{x} \in X$ and $\underline{y}, \overline{y} \in Y$ so that for $\mathcal{N}$ in (11)*

$$(\underline{\lambda}, \underline{x}, \underline{y}) \in \mathcal{N}, \quad \underline{\lambda} \le \underline{x}^\top a, \quad (\overline{\lambda}, \overline{x}, \overline{y}) \in \mathcal{N}, \quad \overline{x}^\top a \le \overline{\lambda}. \tag{38}$$

Then $x^\top a = \lambda$ for some $(\lambda, x, y) \in \mathcal{N}$ with $\lambda \in [\underline{\lambda}, \overline{\lambda}]$.

Proof. Consider the largest λ^* so that $\lambda^* \in [\underline{\lambda}, \overline{\lambda}]$ and there are x^*, y^* with $(\lambda^*, x^*, y^*) \in \mathcal{N}$ and $\lambda^* \le x^{*\top} a$, which exists since $\underline{\lambda}$ fulfills this property and $\mathcal{N}$ is closed by Theorem 3.

If $\lambda^* = \overline{\lambda}$ then both $(\lambda^*, \overline{x})$ and (λ^*, x^*) belong to the same set X_k or X'_k which is convex, where since $\overline{x}^\top a \le \lambda^*$ and $\lambda^* \le x^{*\top} a$ we have $x^\top a = \lambda^*$ for a suitable convex combination x of $\overline{x}$ and x^*, and $(\lambda^*, x, y^*) \in \mathcal{N}$, as claimed.

Hence, we can assume $\lambda^* < \overline{\lambda}$. Suppose λ^* is a breakpoint λ_k, so that $(\lambda^*, x^*) \in X_k$. Consider $\lambda' \in (\lambda_k, \min\{\lambda_{k+1}, \overline{\lambda}\})$ and $(\lambda', x', y') \in X'_k \times Y'_k$ where $\lambda' > x'^\top a$ by maximality of λ^*. By (37), we have $(\lambda^*, x^*) \in X'_k$ and hence $(\lambda^*, x^*, y') \in X'_k \times Y'_k$. Because $\lambda^* \le x^{*\top} a$ and $\lambda' > x'^\top a$, a suitable convex combination (λ, x, y') of (λ^*, x^*, y') and (λ', x', y') belongs to $\mathcal{N}$ and fulfills $\lambda = x^\top a$ as claimed (in fact, $(\lambda, x, y') = (\lambda^*, x^*, y')$ does by maximality of λ^*). If λ^* is not a breakpoint, we directly have $(\lambda^*, x^*, y^*) \in X'_k \times Y'_k$ for some k and can choose $(\lambda', x', y^*) \in X'_k \times Y'_k$ with $\lambda^* < \lambda' \le \overline{\lambda}$ and apply the same argument. □

The binary search algorithm will maintain (38) as an *invariant* while halving the length of the interval $[\underline{\lambda}, \overline{\lambda}]$ in each iteration.

Lemma 12 ensures that the interval contains some λ with $(\lambda, x, y) \in \mathcal{N}$ and $x^\top a = \lambda$ (which is not true when applied to games of higher rank, as shown in the example in Figure 5 below). Let $\lambda' = (\underline{\lambda} + \overline{\lambda})/2$ and let x' be the strategy of player 1 in an equilibrium (x', y') of the game $(A, -A + \mathbb{1}\lambda' b^\top)$, which is found as a solution (x', v', s') to $P_{\lambda'}$. If $\lambda' \le x'^\top a$, it is natural to proceed with $\underline{\lambda}$ set to λ' (written as $\underline{\lambda} \leftarrow \lambda'$), otherwise with $\overline{\lambda} \leftarrow \lambda'$. The binary search should terminate once $\underline{\lambda}$ and $\overline{\lambda}$ are in the same interval $[\lambda_k, \lambda_{k+1}]$ between two breakpoints, with the desired equilibrium found in $(X'_k \times Y'_k) \cap H$.

However, this straightforward approach has the following problems:

i. The search may converge to an equilibrium (x, y) with $x^\top a = \lambda$, where λ is a breakpoint λ_k, so that $\underline{\lambda}$ and $\overline{\lambda}$ are always in different intervals $(\lambda_{k-1}, \lambda_k]$ and $[\lambda_k, \lambda_{k+1})$ and the described termination condition fails;

ii. The number of digits to describe $\underline{\lambda}$ and $\overline{\lambda}$ may pile up, which slows down solving $P_{\lambda'}$.

We address these problems as follows. First, we identify with $M = M(\lambda')$, $N = N(\lambda')$ the face $P(M, N)$ of P that contains (λ', x', v', s'). We then check if that face contains some (λ, x, v, s) with $x^\top a = \lambda$. Depending on whether $\lambda' \le x'^\top a$ or $x'^\top a \le \lambda'$, this is achieved by one of the following variations of the LPs in (29) (these variations will also be used for the enumeration of all equilibria in Section 6):

$$
\begin{aligned}
P^{\max}(M,N,a,\lambda'):\quad & \underset{\lambda,x,v,s}{\text{maximize}} \quad \lambda - x^\top a \\
& \text{subject to} \quad (\lambda, x, v, s) \in P(M,N), \\
& \qquad\qquad x^\top a \ge \lambda \ge \lambda', \\
P^{\min}(M,N,a,\lambda'):\quad & \underset{\lambda,x,v,s}{\text{minimize}} \quad \lambda - x^\top a \\
& \text{subject to} \quad (\lambda, x, v, s) \in P(M,N), \\
& \qquad\qquad x^\top a \le \lambda \le \lambda'.
\end{aligned}
\tag{39}
$$

Figure 1 illustrates $P^{\max}(M, N, a, \lambda')$ where $\lambda' < x'^\top a$, and λ' is between two breakpoints λ_{k-1} and λ'_k (but λ' could also be a breakpoint itself), so that $P(M, N)$ is projected to X'_{k-1}. Here the optimal solution x' to $P_{\lambda'}$ is not unique but always fulfills $\lambda' < x'^\top a$. Moreover, $X'_{k-1} \times Y'_{k-1}$ and H intersect. In the left diagram in Figure 1, $P(M, N)$ is not just a line segment but a higher-dimensional polytope. It contains some (λ, x, v, s) and $(\lambda, \hat{x}, \hat{v}, \hat{s})$ with $x^\top a < \lambda < \hat{x}^\top a$, for example, for $\lambda = \hat{\lambda}$ but not for $\lambda = \lambda'$ nor $\lambda = \lambda_k$. In the right diagram of Figure 1, we always have $\lambda < x^\top a$, and $P^{\max}(M, N, a, \lambda')$ attains its optimum λ^* at λ' because for the corresponding (x^*, λ^*), shown as a dot, $\lambda^* - x^{*\top} a$ is least negative. Here, the solution $\lambda^* = \lambda_k$ would be more useful for proceeding because it is the next breakpoint. We will introduce an extra computation step to achieve this, as we discuss further below.

The next lemma states that the appropriate LP in (39) identifies if there is an equilibrium (x, y) of the game $(A, -A + \mathbb{1}\lambda b^\top)$ with $x^\top a = \lambda$ for some λ between λ' and the next breakpoint λ_k.

Lemma 13. *Let λ_k be a breakpoint of P_λ and D_λ as in Theorem 2, $1 \le k \le K$. Let $\lambda' \in \mathbb{R}$ and let (x', v', s') be an optimal solution to $P_{\lambda'}$, and let $(M, N) = (M(\lambda'), N(\lambda'))$ as in (26).*

a. *Suppose $\lambda' \in (\lambda_{k-1}, \lambda_k]$ and $\lambda' \le x'^\top a$. Let $(\lambda^*, x^*, v^*, s^*)$ be an optimal solution to $P^{\max}(M, N, a, \lambda')$. Then $\lambda^* \in [\lambda', \lambda_k]$, and the game $(A, -A + \mathbb{1}\lambda b^\top)$ has an equilibrium (x, y) with $x^\top a = \lambda$ for some $\lambda \in [\lambda', \lambda_k]$ if and only if this holds for $\lambda = \lambda^*$ and $x = x^*$.*

b. *Suppose $\lambda' \in [\lambda_k, \lambda_{k+1})$ and $\lambda' \ge x'^\top a$. Let $(\lambda^*, x^*, v^*, s^*)$ be an optimal solution to $P^{\min}(M, N, a, \lambda')$. Then $\lambda^* \in [\lambda_k, \lambda']$, and the game $(A, -A + \mathbb{1}\lambda b^\top)$ has an equilibrium (x, y) with $x^\top a = \lambda$ for some $\lambda \in [\lambda_k, \lambda']$ if and only if this holds for $\lambda = \lambda^*$ and $x = x^*$.*

Proof. We prove (a), where (b) is entirely analogous. By Lemma 9, (λ', x', v', s') is feasible for $P^{\max}(M, N, a, \lambda')$. Clearly $\lambda' \le \lambda^*$, and Lemma 10 implies $\lambda^* \le \lambda_k$. Because $\lambda \le x^\top a$ for any feasible solution (λ, x, v, s), the objective function $\lambda - x^\top a$ is nonpositive and zero and

Figure 1. (Color online) Illustration of $P^{\max}(M, N, a, \lambda')$ in (39) for $\lambda' \in (\lambda_{k-1}, \lambda_k)$, with $M = M(\lambda')$, $N = N(\lambda')$, and $P(M, N)$ as a Polytope (Left) or Line Segment (Right)

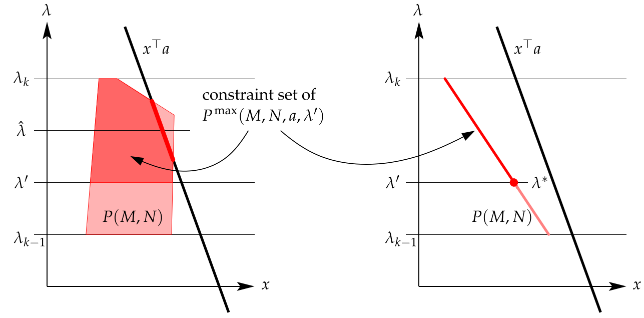

hence optimal if and only if $\lambda = x^\top a$, in which case x is part of the described equilibrium (x, y). □

We now describe the BinSearch algorithm in Figure 2, where we will return to the LPs in (39). The conditions $x^\top a = \lambda$ and $x \in X$ mean that λ is a convex combination of the components $a_1, \ldots, a_m$ of a, so that we can initialize $\underline{\lambda}$ and $\overline{\lambda}$ as their minimum and maximum in line 3 of the algorithm. The main loop of the algorithm is between lines 4 and 18. The candidate value for λ (called λ' in the above explanations) is the midpoint between $\underline{\lambda}$ and $\overline{\lambda}$ in line 5. Line 6 computes some optimal solution (x, v, s) of the LP P_λ in (19), where the dual LP D_λ in (17) is typically solved alongside P_λ. The optimum $\phi(\lambda)$ of P_λ and D_λ determines the optimal face $Y(\lambda)$ of D_λ in (21). The true inequalities M, N of $Y(\lambda)$ in line 7 are determined according to (26), for example with the help of the LP in Lemma 11.

Lines 8–12, and symmetrically 13–17, use the LPs in (39). In order to match the notation in the discussion before Lemma 13, let $\lambda' = \lambda$. Consider the case $\lambda' \le x^\top a$, handled in lines 8–12. Line 9 invokes the LP $P^{\max}(M, N, a, \lambda')$. By Lemma 13, the optimum $(\lambda^*, x^*, v^*, s^*)$ to this LP will find the desired equilibrium with $\lambda^* = x^{*\top} a$ if there is one for some λ^* up to the next breakpoint λ_k, that is, for $\lambda^* \in [\lambda', \lambda_k]$. Suppose this is not the case, that is, $\lambda^* < x^{*\top} a$ and the optimum $\lambda^* - x^{*\top} a$ of $P^{\max}(M, N, a, \lambda')$ is negative. By Lemma 13, in this case the next breakpoint λ_k does *not* define an equilibrium, so that

Figure 2. The BinSearch Algorithm for Finding one Nash Equilibrium of a Rank-1 Game $(A, -A + ab^\top)$

BinSearch

1 Input: $A \in \mathbb{R}^{m\times n}$, $a \in \mathbb{R}^m$, $b \in \mathbb{R}^n$
2 Output: one Nash equilibrium of the game $(A, -A + ab^\top)$
3 $\underline{\lambda} \leftarrow \min\{a_1, \ldots, a_m\}$, $\overline{\lambda} \leftarrow \max\{a_1, \ldots, a_m\}$
4 **repeat**
5 $\quad\lambda \leftarrow (\underline{\lambda} + \overline{\lambda})/2$
6 $\quad(x, v, s) \leftarrow$ solution of P_λ
7 $\quad(M, N) \leftarrow (M(\lambda), N(\lambda))$
8 $\quad$**if** $\lambda \le x^\top a$ **then**
9 $\quad\quad(\lambda^*, x^*, v^*, s^*) \leftarrow$ solution of $P^{\max}(M, N, a, \lambda)$
10 $\quad\quad$**if** $\lambda^* < x^{*\top} a$ **then**
11 $\quad\quad\quad(\lambda^*, x^*, v^*, s^*) \leftarrow$ solution of $BR^{\max}(M, N)$
12 $\quad\quad\underline{\lambda} \leftarrow \lambda^*$
13 $\quad$**else** [know: $x^\top a < \lambda$]
14 $\quad\quad(\lambda^*, x^*, v^*, s^*) \leftarrow$ solution of $P^{\min}(M, N, a, \lambda)$
15 $\quad\quad$**if** $x^{*\top} a < \lambda^*$ **then**
16 $\quad\quad\quad(\lambda^*, x^*, v^*, s^*) \leftarrow$ solution of $BR^{\min}(M, N)$
17 $\quad\quad\overline{\lambda} \leftarrow \lambda^*$
18 **until** $x^{*\top} a = \lambda^*$
19 $(y^*, t^*) \leftarrow$ solution of D_{λ^*}
20 **output** (x^*, y^*)

problem (i) above does not occur. However, as shown in the right diagram in Figure 1, this may result in $\lambda^* = \lambda'$. We could simply continue with $\underline{\lambda} \leftarrow \lambda^*$ as in line 12, but if $\lambda^* = \lambda'$, this increases the description size of $\underline{\lambda}$, which we would like to keep bounded to avoid problem (ii) (the description size of λ probably increases only by one bit per main iteration, but it is useful to keep it independent of the number of iterations both for the computation and for the analysis). In line 10, the condition $\lambda^* < x^{*\top} a$ recognizes that the current segment of $\mathcal{N}$ contains no equilibrium, and then $BR^{\max}(M, N)$ in line 11 computes λ^* as the next breakpoint λ_k according to Lemma 10(a); the LP in line 11 can be solved by starting from the current solution to $P^{\max}(M, N, a, \lambda')$. The left diagram in Figure 1 shows that we cannot simply replace the objective function $\lambda - x^\top a$ of $P^{\max}(M, N, a, \lambda')$ by λ: Although this would compute the next breakpoint λ_k, it may overlook that the current segment of $\mathcal{N}$ defined by $P(M, N)$ intersects the hyperplane H; this could possibly miss the equilibrium altogether, for example, if $\overline{\lambda} = \hat{\lambda}$ as shown in the diagram (in particular if $\overline{\lambda}$ still has its initial value, which is not checked in the algorithm as to whether it produces an equilibrium).

In summary, lines 8–11 find λ^* and x^* so that either (a) $x^{*\top} a = \lambda^*$ or (b) $\lambda^* < x^{*\top} a$ and λ^* is a breakpoint and $(\underline{\lambda} + \overline{\lambda})/2 = \lambda \le \lambda^* < \overline{\lambda}$, which implies $\overline{\lambda} - \lambda^* \le (\overline{\lambda} - \underline{\lambda})/2$. The next value of $\underline{\lambda}$ is set to λ^* in line 12. In case (a), the loop terminates in line 18. In case (b), the loop continues, and in the next iteration the difference $\overline{\lambda} - \underline{\lambda}$ has shrunk by at least one half. The analogous statements hold for lines 13–17. The following theorem states the correctness and polynomial running time of the algorithm.

Theorem 4. *Algorithm* BinSearch *(Figure 2) finds one equilibrium of the rank-1 game $(A, -A + ab^\top)$. Assume that the entries of A, a, b are rational numbers with combined bit length L and that LPs are solved with polynomial-time solvers that return extreme LP solutions obtained from linear equations derived from A, a, b. Then* BinSearch *runs in polynomial time in L.*

Proof. During the main loop, the invariant (38) is preserved, and the length of the interval $[\underline{\lambda}, \overline{\lambda}]$ shrinks by at least a factor of two per iteration. By Lemma 12, a solution $(\lambda, x, y) \in \mathcal{N}$ with $x^\top a = \lambda$ and $\lambda \in [\underline{\lambda}, \overline{\lambda}]$ is guaranteed to exist. The termination condition $x^{*\top} a = \lambda^*$ in line 18 holds once λ reaches a segment of $\mathcal{N}$ that intersects H, which is identified with one of the LPs in line 9 or 14 by Lemma 13. Because the length of the search interval $[\underline{\lambda}, \overline{\lambda}]$ shrinks by at least half in each iteration, the search interval eventually contains at most one breakpoint λ_k. If there is no breakpoint in $[\underline{\lambda}, \overline{\lambda}]$, then $(M(\underline{\lambda}), N(\underline{\lambda})) = (M(\overline{\lambda}), N(\overline{\lambda})) = (M(\lambda), N(\lambda))$ for $\lambda = (\underline{\lambda} + \overline{\lambda})/2$. Hence, a solution $(\lambda^*, x^*, v^*, s^*)$ to

$P^{\max}(M(\lambda), N(\lambda), a, \lambda)$ or to $P^{\min}(M(\lambda), N(\lambda), a, \lambda)$ determines an equilibrium (x^*, y^*) of $(A, -A + ab^\top)$ by Lemma 13 and Lemma 5. This holds also if there is a single breakpoint λ_k in $[\underline{\lambda}, \overline{\lambda}]$. Hence, as claimed, the algorithm computes an equilibrium (x^*, y^*) of $(A, -A + ab^\top)$.

The number of overall iterations is polynomial for the following reason. Any breakpoint λ is part of a vertex (λ, x, v, s) of P by Lemma 10(a). This vertex is a solution to a linear system of equations where each component (such as λ) is a fraction with an integer determinant obtained from A, b in the denominator (which has a polynomial of bits) and distinct fractions for different breakpoints λ. Hence, any two breakpoints have minimum distance $1/2^{p(L)}$ for some polynomial p (see also Schrijver 1986, section 10.2). Therefore, there will be at most $O(p(L))$ binary search iterations until the search interval contains at most one breakpoint and the search terminates.

Each iteration of the algorithm solves three or four LPs. The first is P_λ in line 6. Using the optimum $\phi(\lambda)$ of that LP, in line 7 the true inequalities in (26) of $Y(\lambda)$ in (21) are found with another LP as in Lemma 11. The third LP is either $P^{\max}(M, N, a, \lambda)$ in line 9 or $P^{\min}(M, N, a, \lambda)$ in line 14. The fourth LP is either $BR^{\max}(M, N)$ or $BR^{\min}(M, N)$ in line 11 or 16, respectively, (which just relaxes the extra constraints of the previous LP in (39) and has a different objective function). In all cases, the output λ^* is described in terms of A, a, b and found in polynomial time in the bit size L, and λ^* itself has polynomial bit size (Schrijver 1986, corollary 10.2a(iii)). In the next iteration, λ^* determines with the constant arithmetic expression in line 5 the next parameter λ for P_λ in line 6 and for (M, N) in line 7 so that the bit size of λ remains polynomial in L. Hence, each main iteration takes polynomial time, and the overall running time is polynomial. □

In practice, as observed in Adler and Monteiro (1992, section 5), in the nondegenerate case the segments of $\mathcal{N}$ are line segments. Then the LP in line 9 or 14 is solved starting from the current solution to P_λ in line 6 with a single pivot, and similarly the next LP in line 11 or 16.

6. Enumerating All Equilibria of a Rank-1 Game

In this section, we show how to obtain a complete description of all Nash equilibria of a rank-1 game with the help of Theorem 1 and Theorem 3.

A degenerate bimatrix game may have infinite sets of Nash equilibria. They can be described via *maximal Nash subsets* (Jansen 1981), called "sub-solutions" by Nash (1951, p. 290). A Nash subset for (A, B) is a nonempty product set $S \times T$, where $S \subseteq X$ and $T \subseteq Y$, so that every (x, y) in $S \times T$ is an equilibrium of (A, B); in other words, any two equilibrium strategies $x \in S$ and $y \in T$ are "exchangeable." Using the "best response polyhedra" $\overline{P}$ and $\overline{Q}$ in (2), it can be shown that any maximal Nash subset $S \times T$ is a polytope, with S as a suitable face of $\overline{P}$ projected to X, and T as a suitable face of $\overline{Q}$ projected to Y (Avis et al. 2010). These faces are defined by converting some inequalities in (2) to equations, which have to fulfill the equilibrium conditions (3) and (4). The usual output for enumerating all equilibria consists of listing all maximal Nash subsets $S \times T$ via the vertices of S and T. These are vertices of $\overline{P}$ and $\overline{Q}$, respectively, (projected to X and Y) that define the "extreme" Nash equilibria of (A, B), with maximal Nash subsets obtained as maximally exchangeable sets of such vertices (Avis et al. 2010, proposition 4). Maximal Nash subsets may intersect, in which case their vertex sets intersect. In a nondegenerate game, all maximal Nash subsets are singletons.

For a rank-1 game $(A, -A + ab^\top)$, its set of Nash equilibria is $\mathcal{N} \cap H$ projected to $X \times Y$ by Theorem 1, with $\mathcal{N}$ in (11) and H in (10). By (35), $\mathcal{N}$ is the union of polyhedra, whose nonempty intersections with H give almost directly the maximal Nash subsets.

Theorem 5. *Let $(A, -A + ab^\top)$ be a rank-1 bimatrix game, and let $\lambda_0, \lambda_1, \ldots, \lambda_K, \lambda_{K+1}$, and $\lambda'_k \in (\lambda_k, \lambda_{k+1})$ for $0 \le k \le K$ as in Theorem 2. With (31), (32), (33), and (34), let*

$$\begin{aligned} S_k &= \{ x \mid (\lambda, x) \in X_k,\ x^\top a = \lambda \} && (1 \le k \le K), \\ L_k &= \{ \lambda \mid (\lambda, x) \in X_k,\ x^\top a = \lambda \} && (1 \le k \le K), \\ S'_k &= \{ x \mid (\lambda, x) \in X'_k,\ x^\top a = \lambda \} && (0 \le k \le K), \\ L'_k &= \{ \lambda \mid (\lambda, x) \in X'_k,\ x^\top a = \lambda \} && (0 \le k \le K). \end{aligned} \tag{40}$$

Then the maximal Nash subsets of $(A, -A + ab^\top)$ are the sets $S_k \times Y_k$ if $S_k \neq \emptyset$, and $S'_k \times Y'_k$ if $S'_k \neq \emptyset$ and L'_k is not equal to $\{\lambda_k\}$ or $\{\lambda_{k+1}\}$.

Proof. Each set S_k is the projection of $(X_k \times Y_k) \cap H$ on X, and S'_k is the projection of $(X'_k \times Y'_k) \cap H$ on X, with L_k and L'_k containing the corresponding set of λ's. Hence, by Theorem 3, if $S_k \neq \emptyset$, then $S_k \times Y_k$ is a Nash subset, and if $S'_k \neq \emptyset$, then $S'_k \times Y'_k$ is a Nash subset, and the union of these is the set of all equilibria, which is the projection of $\mathcal{N} \cap H$ on $X \times Y$ by Theorem 1. The only question is which of these Nash subsets are inclusion-maximal. By corollary 3.2 of Adler and Monteiro (1992), $Y_k \cap Y_{k+1} = Y'_k$, where Y_k and Y_{k+1} contain Y'_k properly, $Y_k \cap Y_\ell = \emptyset$ whenever $|k - \ell| \ge 2$, and $Y'_k \cap Y'_\ell = \emptyset$ whenever $k \neq \ell$, and Lemma 10 implies $L_k = \{\lambda_k\} = L_{k-1} \cap L_k$. So the only possible inclusions are that $S'_k \times Y'_k$ is a subset of $S_k \times Y_k$ or of $S_{k+1} \times Y_{k+1}$. Suppose $x \in S'_k$, that is, $(\lambda, x) \in X'_k$ and $x^\top a = \lambda$. If this implies $\lambda = \lambda_k$, then $L'_k = \{\lambda_k\}$. By Lemma 9, this means x is part of an optimal solution (x, v, s) to P_{λ_k} and hence $x \in S_k$, which shows the proper inclusion $S'_k \times Y'_k \subset S_k \times Y_k$ because $Y'_k \subset Y_k$. Similarly, $L'_k = \{\lambda_{k+1}\}$ implies $S'_k \times Y'_k \subset S_k \times Y_{k+1}$. These are the only possible inclusions because if $x \in S'_k$ with $(\lambda, x) \in X'_k$ so that $x^\top a = \lambda \notin \{\lambda_k, \lambda_{k+1}\}$, we clearly cannot have $x \in S_k$, say, where $x^\top a = \lambda_k$.

This proves the theorem. We also note that the described sets S_k and S'_k are defined in terms of the game $(A, -A + ab^\top)$ independently of the parameter λ. Namely, the condition $x^\top a = a^\top x = \lambda$ implies that the polyhedron $\overline{P}$ in (2) for $B = -A + ab^\top$ is given by

$$\begin{aligned} \overline{P} &= \{(x, v) \in X \times \mathbb{R} \mid (-A + ab^\top)^\top x \le \mathbb{1} v\} \\ &= \{(x, v) \in X \times \mathbb{R} \mid -A^\top x + b\lambda \le \mathbb{1} v\}, \end{aligned} \tag{41}$$

so S_k and S'_k are projections of certain faces of $\overline{P}$. □

A suitable algorithm that enumerates all Nash equilibria can be adapted from the algorithm by Adler and Monteiro (1992) that proceeds from breakpoint to breakpoint using Theorem 2. The corresponding segments of $\mathcal{N}$ can then be checked for nonempty intersections with H, which are then output as maximal Nash subsets if they meet the conditions of Theorem 5.

We give an outline of this algorithm. Suppose λ is equal to a breakpoint λ_k. Then Y_k in (34) is the projection of $Y(\lambda_k) = OptFace(D_{\lambda_k})$, and X_k in (33) is the projection of $OptFace(P_{\lambda_k})$ by Lemma 10(b) and Lemma 9. If $(X_k \times Y) \cap H$ is not empty, its projection to $X \times Y$ is a maximal Nash subset $S_k \times Y_k$. Start from some $(\lambda, x) \in X_k$. If $\lambda = x^\top a$, then $x \in S_k$, which is a suitable starting point for the vertex enumeration of the polytope S_k, for example with the program *lrs* (Avis 2000). If $\lambda < x^\top a$ or $\lambda > x^\top a$, then the condition $(X_k \times Y) \cap H \neq \emptyset$ is checked with one of the LPs in (39) by Lemma 13, which then have optimal value zero, with optimum $(\lambda^*, x^*, v^*, s^*)$; then $\lambda^* = x^{*\top} a$, and $x^* \in S_k$ is a new starting point to enumerate the vertices of S_k.

The next segment to be tested for its intersection with H is $X'_k \times Y'_k$ in (31) and (32). For that purpose it is not necessary to find some $\lambda' \in (\lambda_k, \lambda_{k+1})$, because $Y(\lambda') = OptFace(SL^{\max}(\lambda_k))$ by Theorem 2, and the true inequalities $M \cup N$ of that face are found by Lemma 11, so that one obtains X'_k as the projection of $P(M, N)$. Moreover, we have $x \in X_k \subseteq X'_k$. If $\lambda = x^\top a$, then x is also a starting point for the enumeration of the vertices of S'_k, which gives the Nash subset $S'_k \times Y'_k$ (which is, however, not maximal if $S'_k \subseteq S_k$, see Theorem 5). If $\lambda < x^\top a$, then we solve $P^{\max}(M, N, a, \lambda_k)$ in (39) to find out if H intersects the current segment $X'_k \times Y'_k$, and similarly $P^{\min}(M, N, a, \lambda_k)$ if $\lambda > x^\top a$. Finally, the next breakpoint λ_{k+1} is found as the solution to $BR^{\max}(M, N)$ in (29) by Lemma 10(a).

For initialization and termination of this algorithm, we use that the possible values of λ can be restricted to $[\underline{\alpha}, \overline{\alpha}]$ with $\underline{\alpha}$ and $\overline{\alpha}$ as minimum and maximum of $\{a_1, \ldots, a_m\}$. The initialization is $\lambda = \underline{\alpha}$, which is decided to be a breakpoint or not as described after (23). The constraint $\lambda \le \overline{\alpha}$ is added to the step of finding the next breakpoint, which terminates the algorithm when it is found to hold as equality.

This algorithm, based on Theorem 5, for enumerating all Nash equilibria of a rank-1 game has the following noteworthy features. First, it works for all games (degenerate or not), and its characterization of maximal Nash subsets is simpler than for general bimatrix games (Avis et al. 2010) and could even be adapted to easily represent these Nash subsets in terms of their inequalities rather than their vertices (which would be of interest if they are high dimensional). Secondly, the algorithm in effect traverses $\mathcal{N}$, which is generically a path. Rather than by solving a succession of LPs, it can also be implemented by a variant of the algorithm by Lemke (1965) with the additional linear constraints $\lambda \ge x^\top a$ or $\lambda \le x^\top a$, depending on the current sign of $\lambda - x^\top a$. Here, traversing this path gives *all* Nash equilibria, whereas for general bimatrix games Lemke's algorithm (as in von Stengel et al. 2002 or Govindan and Wilson 2003) only finds *one* Nash equilibrium.

7. Two Examples

In this section, we illustrate the results of the previous sections with an example of a rank-1 game. After that we will give an example that shows that binary search will in general not work for a game of rank 2 or higher, even though Lemma 5 suggests the possibility of finding a Nash equilibrium of such a game via a recursive rank reduction.

Consider the following rank-1 game (A, B),

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & -2 \\ -1 & 0 \end{bmatrix}, A + B = \begin{bmatrix} 2 & -2 \\ -1 & 1 \end{bmatrix} = ab^\top, \tag{42}$$

where $a^\top = (2, -1)$ and $b^\top = (1, -1)$. This game has the two pure equilibria $((1, 0), (1, 0))$ and $((0, 1), (0, 1))$, and the mixed equilibrium $((\frac{1}{4}, \frac{3}{4}), (\frac{1}{2}, \frac{1}{2}))$. By Theorem 1(b), these are the equilibria (x, y) of the game $(A, -A + \mathbb{1}\lambda b^\top)$ so that $x^\top a = \lambda$. For $x = (1, 0), (\frac{1}{4}, \frac{3}{4}), (0, 1)$, this means $\lambda = 2, -\frac{1}{4}, -1$.

Figure 3 shows the set $\mathcal{N}$ in (11) where (x, y) is an equilibrium of the parameterized game $(A, -A + \mathbb{1}\lambda b^\top)$, where

$$-A + \mathbb{1}\lambda b^\top = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} + \begin{bmatrix} \lambda & -\lambda \\ \lambda & -\lambda \end{bmatrix}. \tag{43}$$

These equilibria are pure except when $\lambda \in [-\frac{1}{2}, \frac{1}{2}]$, when the unique mixed strategy $(1 - x_2, x_2)$ of player 1 is given by equalizing the column payoffs, $-(1 - x_2) + \lambda = -x_2 - \lambda$, that is, $\lambda = \frac{1}{2} - x_2$. The white dots indicate the intersection of $\mathcal{N}$ with the hyperplane H in (10),

Figure 3. (Color online) The Path $\mathcal{N}$ in (11) for the Game (43), for $x = (1 - x_2, x_2) \in X$ and $y = (1 - y_2, y_2) \in Y$, and the Hyperplane H in (10)

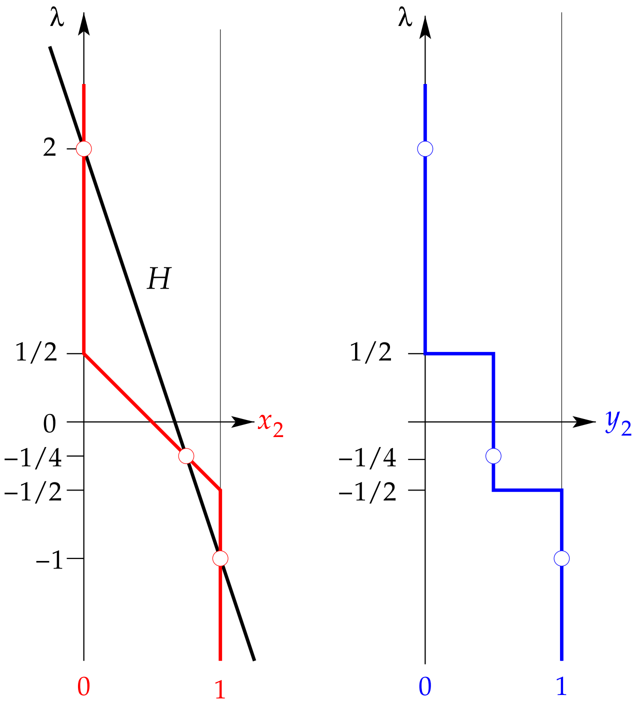

which is defined by the equation $\lambda = x^\top a = 2(1 - x_2) - x_2 = 2 - 3x_2$, and no constraints on y.

Figure 4 shows the domains of the LPs P_λ in (19) and D_λ in (17) for $\lambda = -\frac{1}{4}$. Again we show x in X as $(1 - x_2, x_2)$ and y in Y as $(1 - y_2, y_2)$. The constraints of P_λ are then $1 - x_2 + v \geq \lambda$ and $x_2 + v \geq -\lambda$, which for $\lambda = -\frac{1}{4}$ are $v \geq -\frac{5}{4} + x_2$ and $v \geq \frac{1}{4} - x_2$. The constraints $Ay + \mathbb{1}t \leq \mathbb{0}$ of D_λ are

$$1 - y_2 + t \leq 0 \qquad \text{and} \qquad y_2 + t \leq 0, \tag{44}$$

Figure 4. (Color online) The LP P_λ in (19) and the Polyhedron D in (18) with the Objective Function of the LP D_λ in (17) for $\lambda = -\frac{1}{4}$, for the Game (43)

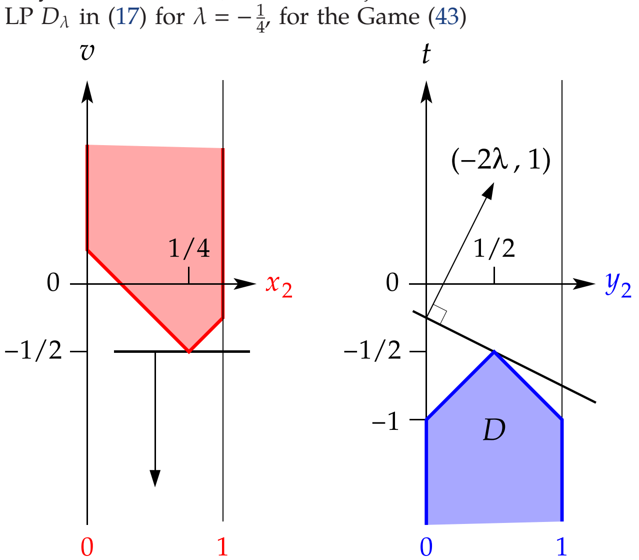

and the objective function $\lambda b^\top y + t$ is $\lambda(1 - y_2 - y_2) + t$, with gradient $(\frac{\partial}{\partial y_2}, \frac{\partial}{\partial t}) = (-2\lambda, 1) = (\frac{1}{2}, 1)$ for $\lambda = -\frac{1}{4}$. For $\lambda > \frac{1}{2}$, the optimum of D_λ is attained at the vertex $(y_1, y_2, t) = (1, 0, -1)$ of D, for $\frac{1}{2} > \lambda > -\frac{1}{2}$ at the vertex $(\frac{1}{2}, \frac{1}{2}, -\frac{1}{2})$, and for $-\frac{1}{2} > \lambda$ at the vertex $(0, 1, -1)$. For $\lambda_2 = \frac{1}{2}$ and $\lambda_1 = -\frac{1}{2}$, the optimal face of D_λ is an edge of D. These are the two breakpoints λ_1 and λ_2 in Theorem 2.

Figure 3 also demonstrates the characterization of the path $\mathcal{N}$ in Theorem 3. The left diagram shows (from left to right) the three pieces X'_2, X'_1, and X'_0, each of which happen to intersect H. In the central diagram, the vertical parts of the path are Y'_2, Y'_1, and Y'_0, and the horizontal parts (for the breakpoints) are Y_2 and Y_1. This corresponds to the following, more elementary game-theoretic explanation. Except when $\lambda = -\frac{1}{2}$ or $\lambda = \frac{1}{2}$, player 2's equilibrium strategy y in the game $(A, -A + \mathbb{1}\lambda b^\top)$ is constant in λ, which holds because player 1's payoff matrix A does not change with λ and y is chosen so as to make player 1 indifferent between the pure strategies in the support of his equilibrium strategy. When $\lambda = -\frac{1}{2}$ or $\lambda = \frac{1}{2}$, the game is degenerate, and player 2's equilibrium strategies form a line segment, which allows the change of support of player 2's equilibrium strategy y.

Our second example shows that the binary search algorithm no longer works for rank-r games with $r > 1$. Consider the following game (A, B) of rank 2:

$$A = \begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix}, \quad C = \begin{bmatrix} 4 & 0 \\ 0 & 0 \end{bmatrix},$$
$$B = C + ab^\top = \begin{bmatrix} 1 & 0 \\ 2 & 0 \end{bmatrix}, \tag{45}$$

where $a^\top = (-3, 2)$ and $b^\top = (1, 0)$. Here, (A, B) is of rank 2 and (A, C) is of rank 1. The only equilibrium of (A, B) is the pure equilibrium $((1, 0), (1, 0))$. The parameterized game $(A, C + \mathbb{1}\lambda b^\top)$ has payoff matrices

$$A = \begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix}, \quad C + \mathbb{1}\lambda b^\top = \begin{bmatrix} 4 + \lambda & 0 \\ \lambda & 0 \end{bmatrix}. \tag{46}$$

It has the following equilibria (x, y) depending on λ, which define the set $\mathcal{N}$ in (9), shown in Figure 5: The pure equilibrium $((1, 0), (1, 0))$ for $\lambda \geq -4$; the pure equilibrium $((0, 1), (0, 1))$ for $\lambda \leq 0$; the mixed equilibrium $((-\frac{\lambda}{4}, 1 + \frac{\lambda}{4}), (\frac{1}{2}, \frac{1}{2}))$ for $-4 < \lambda < 0$, and two further components $((1, 0), (1 - y_2, y_2))$ with $y_2 \in [0, \frac{1}{2}]$ when $\lambda = -4$ and $((0, 1), (1 - y_2, y_2))$ with $y_2 \in [\frac{1}{2}, 1]$ when $\lambda = 0$, where the game in (46) is degenerate. These are multiple disjoint equilibrium components for $-4 \leq \lambda \leq 0$, which cannot happen for a parameterized zero-sum game. As a result, λ may change nonmonotonically along the path $\mathcal{N}$, which in general causes a binary search to fail, as we show next.

We describe a suitably adapted binary search method for this example, where instead of solving parameterized LPs we find equilibria of the parameterized game (46) of

Figure 5. (Color online) The Path $\mathcal{N}$ of Equilibria of the Games in (46) Where the Binary Search Method Fails

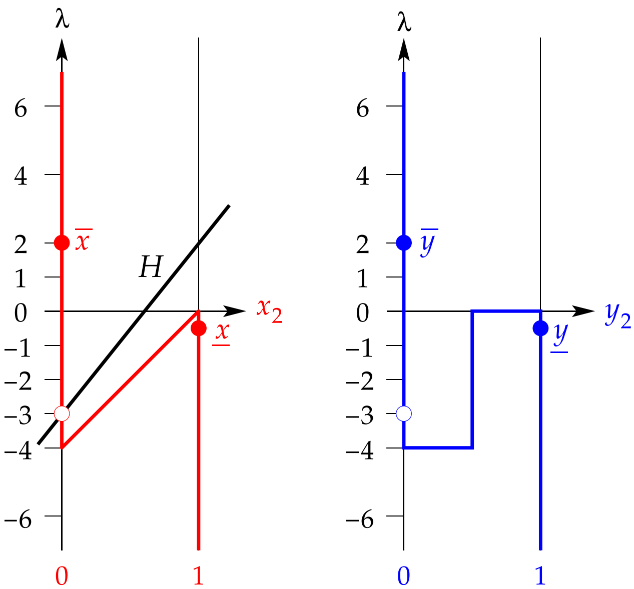

lower rank. The smallest and largest components of a as in line 3 of the BinSearch algorithm are $\underline{\lambda} = -3$ and $\overline{\lambda} = 2$. For $\lambda = \overline{\lambda}$, the only equilibrium of the game in (46) is $(\overline{x}, \overline{y}) = ((1,0),(1,0))$; but for $\lambda = \underline{\lambda}$, there are multiple equilibria, where we choose $(\underline{x}, \underline{y}) = ((0,1),(0,1))$. Then $\underline{\lambda} = -3 < \underline{x}^\top a = 2$ and $\overline{x}^\top a = -3 < \overline{\lambda} = 2$, so we next consider the midpoint $\lambda = (\underline{\lambda} + \overline{\lambda})/2 = -1/2$ as in line 5 of BinSearch and compute a new equilibrium of this parameterized game. Suppose this is again $(x,y) = ((0,1),(0,1))$, so that because $\lambda < x^\top a$ the assignment $(\underline{\lambda}, \underline{x}, \underline{y}) \leftarrow (\lambda, x, y)$ takes place for the binary search to continue. This is the situation shown in Figure 5. At this point, the method will no longer succeed in finding a suitable value of λ because the search interval $[\underline{\lambda}, \overline{\lambda}] = [-\frac{1}{2}, 2]$ no longer contains the only possible value for λ, namely, -3. The problem is that in that interval, the set $\mathcal{N}$ consists of two disconnected parts where $\lambda < x^\top a$ and $\lambda > x^\top a$ are on opposite sides of the hyperplane H, so that $\mathcal{N}$ no longer intersects with H. Hence, even though the values of λ converge, the corresponding equilibria (x,y) on the two sides of H will not converge.

This example shows that because of the nonmonotonicity of λ along the path $\mathcal{N}$, there is no equivalent statement to Lemma 12 that would guarantee that a binary search will succeed.

8. Rank-1 Games with Exponentially Many Equilibria

Kannan and Theobald (2010, open problem 9) asked if the number of Nash equilibria of a nondegenerate rank-1 game is polynomially bounded. This is not the case, because our next result shows that this number may be exponential.

Theorem 6. *Let $p > 2$ and let (A,B) be the $n \times n$ bimatrix game with entries of A*

$$a_{ij} = \begin{cases} 2p^{i+j} & \text{if } j > i \\ p^{2i} & \text{if } j = i \\ 0 & \text{if } j < i \end{cases} \tag{47}$$

for $1 \le i,j \le n$, and $B = A^\top$. Then $A+B$ is of rank 1, and (A,B) is a nondegenerate bimatrix game with $2^n - 1$ many Nash equilibria.

Proof. By (47), $A + B = ab^\top$ with the n components of a and b defined by $a_i = p^i$ and $b_j = 2p^j$ for $1 \le i,j \le n$, so $A+B$ is of rank 1.

Let $y \in Y$ with support S. Consider a row i and let $T = \{j \in S \mid j > i\}$. Because A is upper triangular, the expected payoff against y in row i is

$$(Ay)_i = a_{ii}y_i + \sum_{j \in T} a_{ij}y_j. \tag{48}$$

Suppose $i \notin S$. If T is empty, then $(Ay)_i = 0 < (Ay)_1$, otherwise let $t = \min T$ and note that for $j \in T$, we have $a_{ij} = 2p^{i+j} < p^{1+i+j} \le p^{t+j} \le a_{tj}$, so $(Ay)_i < (Ay)_t$. Hence, no row i outside S is a best response to y. Similarly, because the game is symmetric, any column that is a best response to x in X belongs to the support of x. This shows that the game is nondegenerate. Moreover, if (x,y) is an equilibrium of (A,B), then x and y have equal supports.

For any nonempty subset S of $\{1,\ldots,n\}$, we construct a mixed strategy y with support S so that (y,y) is an equilibrium of (A,B). This implies that the game has $2^n - 1$ many equilibria, one for each support set S. The equilibrium condition holds if $(Ay)_i = u$ for $i \in S$ with equilibrium payoff u, because then $(Ay)_i < u$ for $i \notin S$ as shown above. We start with $s = \max S$, where $(Ay)_s = a_{ss}y_s = u$, by fixing u as some positive constant (e.g., $u = 1$), which determines y_s. Once y_i is known for all $i \in S$ (and $y_i = 0$ for $i \notin S$), we scale y and u by multiplication with $1/\mathbb{1}^\top y$ so that y becomes a mixed strategy. Assume that $i \in S$ and $T = \{j \in S \mid j > i\} \ne \emptyset$ and assume that y_k has been found for all k in T so that $(Ay)_k = u$ for all k in T, which is true for $T = \{s\}$. Then, as shown above, $\sum_{j\in T} a_{ij}y_j < \sum_{j \in T} a_{tj}y_j = (Ay)_t = u$ for $t = \min T$, so y_i is determined by $(Ay)_i = u$ in (48), and $y_i > 0$. By induction, this determines y_i for all i in S and after rescaling gives the desired equilibrium strategy y. □

By Theorem 1, the equilibria (x,y) of a rank-1 game are the intersection of the path $\mathcal{N}$ in (11) with the hyperplane H in (10). The exponential number of Nash equilibria of the game in Theorem 6 shows that $\mathcal{N}$ has exponentially many line segments. Murty (1980) describes a parameterized LP with such an exponentially long path of length 2^n. The payoffs for the game in Theorem 6 have been inspired by Murty's example but are not systematically constructed from it, which would be interesting. See von Stengel (2012) for further discussions and related work on the maximal number of Nash equilibria in bimatrix games, such as von Stengel (1999).

9. A Rank-Preserving Structure Theorem

Nash equilibria of games are in general not unique, which has led to a large literature on equilibrium *refinements* (van Damme 1991) that impose additional conditions on equilibria, such as *stability* against small changes in the game parameters, as proposed in the seminal paper by Kohlberg and Mertens (1986). The authors showed that stability has to apply to equilibrium *components*, that is, maximal sets of equilibria that are topologically connected (which for bimatrix games are unions of intersecting maximal Nash subsets, see Section 6). That is, an equilibrium component is stable if every perturbed game has an equilibrium near that component (although possibly in different positions depending on the perturbation, which is why any single equilibrium may fail to be stable). KM proved the existence of stable equilibrium components with the help of a *structure theorem* (Kohlberg and Mertens 1986, theorem 1), which states that the equilibrium correspondence E over the set Γ of strategic-form games with a given number of players and numbers of strategies is homeomorphic to Γ itself.

In this section, we present in Theorem 7 a similar structure theorem with a new homeomorphism for bimatrix games that *preserves rank*. In analogy to Kohlberg and Mertens (1986, appendix B), one consequence of this new structure theorem is the existence of an equilibrium component in a game (A, B) that is stable with respect to small perturbations that preserve the sum $A + B$ of the payoff matrices. This is not interesting for zero-sum games, which always have only one component, but it is for games of higher rank and applies, for example, to perturbations of the matrix A in a rank-1 game given as $(A, -A + ab^\top)$. Furthermore, a number of equilibrium-finding algorithms can be interpreted as following a path on the equilibrium correspondence E via the KM homeomorphism and suitable projections (Wilson 1992, Govindan and Wilson 2003). As a topic for further research, it may be interesting to study our new homeomorphism in this context or, similar to Jansen and Vermeulen (2001), the computation of equilibrium components that are stable with respect to small perturbations that preserve the sum $A + B$ of the payoff matrices.

We first recall the KM homeomorphism. Let Γ be the set of $m \times n$ bimatrix games (A, B) and $E \subseteq \Gamma \times X \times Y$ be its equilibrium correspondence,

$$E = \{(A, B, x, y) \mid (A, B) \in \Gamma,\ (x, y) \text{ is a NE of } (A, B)\}. \tag{49}$$

To distinguish the dimensions of the all-zero and all-one vectors we write them as $\mathbf{0}, \mathbf{1} \in \mathbb{R}^m$ and $\mathbb{0}, \mathbb{1} \in \mathbb{R}^n$. Let a and b be the vectors of row and column averages of A and B,

$$a = A\mathbb{1}\tfrac{1}{n}, \qquad b = B^\top\mathbf{1}\tfrac{1}{m}. \tag{50}$$

Then A and B correspond uniquely to pairs $(\tilde{A}, a)$ and $(\tilde{B}, b)$ with

$$A = \tilde{A} + a\mathbb{1}^\top, \quad B = \tilde{B} + \mathbf{1}b^\top, \quad \tilde{A}\mathbb{1} = \mathbf{0}, \quad \mathbf{1}^\top\tilde{B} = \mathbb{0}^\top, \tag{51}$$

with a and b as in (50). That is, (A, B) is parameterized by a "base game" $(\tilde{A}, \tilde{B})$, where each row of player 1 and each column of player 2 gets payoff zero when the other player randomizes uniformly (as in $\tilde{A}\mathbb{1}\frac{1}{n} = \mathbf{0}$, where the factor $\frac{1}{n}$ does not matter), and a pair of vectors a in $\mathbb{R}^m$ and $b^\top$ with b in $\mathbb{R}^n$ that are added to the rows of $\tilde{A}$ and columns of $\tilde{B}$, respectively, to obtain the correct payoffs.

The KM homeomorphism $\phi : \Gamma \to E$ only changes a and b. It is most easily described by its inverse $\phi^{-1} : E \to \Gamma$ defined by $\phi^{-1}(A, B, x, y) = (C, D)$,

$$C = \tilde{A} + (Ay + x)\mathbb{1}^\top, \qquad D = \tilde{B} + \mathbf{1}(x^\top B + y^\top). \tag{52}$$

That is, (C, D) has the same base game $(\tilde{A}, \tilde{B})$ as (A, B) but different parameters $(Ay + x) \in \mathbb{R}^m$ and $(B^\top x + y) \in \mathbb{R}^n$. The fact that (x, y) is an equilibrium of (A, B) implies that ϕ^{-1} is injective (and therefore ϕ well defined), by the following intuition. Because x is a best response to y, each row of the vector Ay of expected payoffs in the support of x has maximal and equal value u among all components of Ay, by (3). This condition allows us to reconstruct x from the sum $c = Ay + x$, which is used in the definition of C in (52) and which can be obtained from C. Suppose the components c_i of c are heights of m "poles in the water" of which a certain amount x_i is "above the waterline" depending on the "water level" w, where

$$x_i = \max(c_i - w, 0), \tag{53}$$

so $x_i \geq 0$ and if $c_i < w$, then $x_i = 0$. For any $c \in \mathbb{R}^m$, there is a unique choice of $w \in \mathbb{R}$ in (53) so that $\sum_{i=1}^m x_i = 1$ and therefore $x \in X$. Let $p = c - x$, that is, $c = x + p$. Then all components p_i of the vector p fulfill (a) $w = \max_k p_k$, and (b) $x_i > 0$ implies $p_i = w$, as when $p = Ay$ and x is a best response to y. In a similar way, y is a best response to x, and the sum $x^\top B + y^\top$ used to define D in (52) is special because it allows us first to obtain a vector $d \in \mathbb{R}^n$ from D and second to obtain the original $y \in Y$ and $q \in \mathbb{R}^n$ so that $d = q + y$ and $q^\top = x^\top B$. The following lemma states this construction, which we apply afterward to define the KM homeomorphism and will later use again for our new homeomorphism.

Lemma 14. *Given $c \in \mathbb{R}^m$ and $d \in \mathbb{R}^n$, there are unique $x \in X$, $y \in Y$, $p \in \mathbb{R}^m$ and $q \in \mathbb{R}^n$ so that*

$$\begin{aligned} c &= p + x, & d &= q + y, \\ x_i &= 0 \quad \text{or} & p_i &= u = \max_{1\le k\le m} p_k \quad (1 \le i \le m), \\ y_j &= 0 \quad \text{or} & q_j &= v = \max_{1\le l\le n} q_l \quad (1 \le j \le n). \end{aligned} \tag{54}$$

Proof. For $t \in \mathbb{R}$, let $t^+ = \max(t, 0)$, and

$$\begin{aligned} u &= \min\left\{ w \in \mathbb{R} \mid \sum_{i=1}^{m} (c_i - w)^+ \le 1 \right\}, \\ v &= \min\left\{ w \in \mathbb{R} \mid \sum_{j=1}^{n} (d_j - w)^+ \le 1 \right\}, \end{aligned} \tag{55}$$

where u (and similarly v) is the unique lowest "water level" w so that the "heights" of the components c_i of c that are "above the waterline" sum up to (at most) one. Then

$$x_i = (c_i - u)^+ \ (1 \le i \le m), \quad y_j = (d_j - v)^+ \ (1 \le j \le n), \tag{56}$$

and $p = c - x$ and $q = d - y$ fulfill (54), and x, y, p, q are uniquely determined by the conditions $x \in X$, $y \in Y$, and (54). □

The KM homeomorphism $\phi : (C, D) \mapsto (A, B, x, y)$ is then defined as follows.

a. Let $c = C\mathbb{1}\frac{1}{n}$, $d = D^\top \mathbf{1}\frac{1}{m}$, $\tilde{A} = C - c\mathbb{1}^\top$, and $\tilde{B} = D - \mathbf{1}d^\top$.

b. Apply Lemma 14 to get x, y, p, q so that (54) holds.

c. Let $a = c - x - \tilde{A}y$ and $b = d - y - \tilde{B}^\top x$, and define A and B by (51).

Then ϕ is continuous because it is defined by continuous linear mappings and (55) and (56) for (b). We show that $(A, B, x, y) \in E$. We have $Ay = (\tilde{A} + a\mathbb{1}^\top)y = \tilde{A}y + a = \tilde{A}y + c - x - \tilde{A}y = p$ and similarly $x^\top B = x^\top \tilde{B} + b^\top = d^\top - y^\top = q^\top$. Then the conditions (54) are equivalent to the best-response conditions (3) and (4), that is, (x, y) is indeed an equilibrium of (A, B). Moreover, $c = p + x = Ay + x$ and $d = B^\top x + y$, which shows that the (continuous) function $(A, B, x, y) \mapsto (C, D)$ in (52) is indeed the inverse of ϕ (so ϕ is injective) and also that ϕ is surjective, because we can start in (52) from any $(A, B, x, y) \in E$.

The KM homeomorphism does not operate within a subset of games of fixed rank (for example, the zero-sum games). Our new homeomorphism $\psi : \Gamma \to E$ has this property. Consider a fixed matrix $M \in \mathbb{R}^{m\times n}$, the set Γ_M bimatrix games (A, B) with $A + B = M$, and E_M as the equilibrium correspondence E in (49) restricted to these games,

$$\begin{aligned} \Gamma_M &= \{(A, B) \in \Gamma \mid A + B = M\}, \\ E_M &= \{(A, B, x, y) \in E \mid (A, B) \in \Gamma_M\}. \end{aligned} \tag{57}$$

The following theorem states we can restrict ψ to a homeomorphism $\Gamma_M \to E_M$ for any M (for example, the all-zero matrix M). Also, ψ is continuous in M and therefore a homeomorphism $\Gamma \to E$ like the KM homeomorphism.

Theorem 7. *Let $M \in \mathbb{R}^{m\times n}$. There is a homeomorphism $\psi : \Gamma_M \to E_M$, $(C, D) \mapsto (A, B, x, y)$, that is, $A + B = M$ for all $(C, D) \in \Gamma_M$.*

Proof. We will use a new parameterization of any matrix A in $\mathbb{R}^{m\times n}$, which corresponds uniquely to a quadruple $(\hat{A}, \gamma, a, b)$ with $\hat{A} \in \mathbb{R}^{m\times n}$, $\gamma \in \mathbb{R}$, $a \in \mathbb{R}^m$, and $b \in \mathbb{R}^n$ according to

$$A = \hat{A} + \mathbf{1}\gamma\mathbb{1}^\top + a\mathbb{1}^\top + \mathbf{1}b^\top \tag{58}$$

so that

$$\mathbf{1}^\top \hat{A} = \mathbb{0}^\top, \qquad \hat{A}\mathbb{1} = \mathbf{0}, \qquad \mathbf{1}^\top a = 0, \qquad b^\top \mathbb{1} = 0. \tag{59}$$

It is easy to see that $\hat{A}$, γ, a, and b are uniquely given by A, (58), and

$$\gamma = \tfrac{1}{m}\mathbf{1}^\top A\mathbb{1}\tfrac{1}{n}, \quad a = A\mathbb{1}\tfrac{1}{n} - \mathbf{1}\gamma, \quad b^\top = \tfrac{1}{m}\mathbf{1}^\top A - \gamma\mathbb{1}^\top. \tag{60}$$

The homeomorphism $\psi : \Gamma_M \to E_M$, $(C, D) \mapsto (A, B, x, y)$ uses this parameterization of C and only changes the vectors a and b and maintains the sum M of the payoff matrices, that is, $A + B = C + D = M$. Like for the KM homeomorphism, we first describe its inverse ψ^{-1}, which maps (A, B, x, y) in E_M to (C, D) in Γ_M. Let $A + B = M$ and (x, y) be an equilibrium of (A, B). Let A be represented as in (58) so that (59) holds, and let

$$C = \hat{A} + \mathbf{1}\gamma\mathbb{1}^\top + c\mathbb{1}^\top + \mathbf{1}d^\top, \tag{61}$$

with c and d given by

$$c = \rho(Ay + x), \qquad d = \sigma(B^\top x + y), \tag{62}$$

where $\rho : \mathbb{R}^m \to \mathbb{R}^m$ and $\sigma : \mathbb{R}^n \to \mathbb{R}^n$ are the linear projections on the hyperplane through the origin with normal vector $\mathbf{1}$, respectively $\mathbb{1}$,

$$\rho(x) = x - \mathbf{1}\left(\tfrac{1}{m}\mathbf{1}^\top x\right), \qquad \sigma(y) = y - \mathbb{1}\left(\tfrac{1}{n}\mathbb{1}^\top y\right), \tag{63}$$

which achieves $\mathbf{1}^\top \rho(x) = 0$ and $\mathbb{1}^\top \sigma(y) = 0$ for any $x \in \mathbb{R}^m$ and $y \in \mathbb{R}^n$, as required for a parameterization of the payoff matrix C like it is done for A in (59). With C thus encoded, we let $D = M - C$.

The homeomorphism $\psi : (C, D) \mapsto (A, B, x, y)$ itself is obtained as follows. Let $(C, D) \in \Gamma_M$. Similar to (58), we represent C by (61), whereas in (60)

$$\gamma = \tfrac{1}{m}\mathbf{1}^\top C\mathbb{1}\tfrac{1}{n}, \quad c = C\mathbb{1}\tfrac{1}{n} - \mathbf{1}\gamma, \quad d^\top = \tfrac{1}{m}\mathbf{1}^\top C - \gamma\mathbb{1}^\top, \tag{64}$$

which implies

$$\mathbf{1}^\top \hat{A} = \mathbb{0}^\top, \quad \hat{A}\mathbb{1} = \mathbf{0}, \quad \mathbf{1}^\top c = 0, \quad d^\top \mathbb{1} = 0. \tag{65}$$

Given c and d, we determine $x \in X$, $y \in Y$, $p \in \mathbb{R}^m$, and $q \in \mathbb{R}^n$ by Lemma 14 so that (54) holds. Then, let

$$a = c - \rho(\hat{A}y + x), \qquad b = \sigma((M - \hat{A})^\top x + y) - d, \quad (66)$$

so that a and b fulfill (59), define A by (58), and let $B = M - A$. Like ϕ before, ψ is defined by linear maps and the continuous operations in (55) and (56) and is therefore continuous.

We show that $\psi(C, D) = (A, B, x, y) \in E_M$. Because $A + B = M$, we only need to show the equilibrium property. Using (58), $\mathbb{1}^\top y = 1$, (66), $c = p + x$, and the definition of ρ in (63),

$$\begin{aligned}
Ay &= \hat{A}y + \mathbf{1}\gamma\mathbb{1}^\top y + a\mathbb{1}^\top y + \mathbf{1}b^\top y \\
&= \hat{A}y + \mathbf{1}\gamma + a + \mathbf{1}b^\top y \\
&= \hat{A}y + \mathbf{1}\gamma + c - \rho(\hat{A}y + x) + \mathbf{1}b^\top y \\
&= \hat{A}y + \mathbf{1}\gamma + p + x - (\hat{A}y + x) + \mathbf{1}\left(\frac{1}{m}\mathbf{1}^\top\left(\hat{A}y + x\right)\right) \\
&\quad + \mathbf{1}b^\top y \\
&= p + \mathbf{1}\left(\gamma + \frac{1}{m}\mathbf{1}^\top\left(\hat{A}y + x\right) + b^\top y\right) \\
&= p + \mathbf{1}\alpha \qquad (67)
\end{aligned}$$

for some $\alpha \in \mathbb{R}$, which means that $(Ay)_i = p_i + \alpha$ for $1 \le i \le m$ and therefore by (54) the best-response condition (3) holds (which is unaffected by a constant shift), that is, x is a best response to y. Similarly, using $\mathbf{1}^\top x = 1$, (66), the definition of σ in (63), and $d = q + y$,

$$\begin{aligned}
B^\top x &= (M - A)^\top x \\
&= (M - \hat{A} - \mathbf{1}\gamma\mathbb{1}^\top - a\mathbb{1}^\top - \mathbf{1}b^\top)^\top x \\
&= (M - \hat{A})^\top x - \mathbb{1}\gamma\mathbf{1}^\top x - \mathbb{1}a^\top x - b\mathbf{1}^\top x \\
&= (M - \hat{A})^\top x - \mathbb{1}\gamma - \mathbb{1}a^\top x - b \\
&= (M - \hat{A})^\top x - \mathbb{1}\gamma - \mathbb{1}a^\top x - \sigma((M - \hat{A})^\top x + y) + d \\
&= -\mathbb{1}\gamma - \mathbb{1}a^\top x - y + \mathbb{1}\tfrac{1}{n}\mathbb{1}^\top((M - \hat{A})^\top x + y) + q + y \\
&= \mathbb{1}\beta + q \qquad (68)
\end{aligned}$$

for some $\beta \in \mathbb{R}$, which means that $(B^\top x)_j = q_j + \beta$ for $1 \le j \le n$ and therefore by (54) the best-response condition (4) holds, that is, y is a best response to x. Hence, (x, y) is indeed an equilibrium of (A, B).

To show that ψ has the inverse described in (61) and (62), note that ρ and σ in (63) are linear and $\rho(\mathbf{1}) = \mathbf{0}$ and $\sigma(\mathbb{1}) = \mathbb{0}$. Therefore, for $\psi(C, D) = (A, B, x, y)$ with C as in (61), we have by (67) and (68) and because $\mathbf{1}^\top c = \mathbf{0}$ and $\mathbb{1}^\top d = \mathbb{0}$,

$$\begin{aligned}
\rho(Ay + x) &= \rho(p + \mathbf{1}\alpha + x) = \rho(p + x) = \rho(c) = c, \\
\sigma(B^\top x + y) &= \sigma(\mathbb{1}\beta + q + y) = \sigma(q + y) = \sigma(d) = d, \quad (69)
\end{aligned}$$

that is, ψ has indeed the (continuous) inverse described in (62) and ψ is both injective and surjective. This shows that ψ is indeed a homeomorphsim from Γ_M to E_M. □

10. Conclusions

We conclude with some open questions. Our analysis shows that rank-1 games are computationally easy to analyze: One Nash equilibrium can be found in polynomial time, and enumerating all equilibria can be performed by following a piecewise linear path, similar to finding a single Nash equilibrium of a bimatrix game (which is in general a PPAD-hard problem).

As described in Section 6, the path of solutions to the parameterized LP consists in general of polyhedral segments whose intersections with the hyperplane H define the sets of Nash equilibria of the rank-1 game. This setup suggests the application of *smoothed analysis* as pioneered by Spielman and Teng (2004) for the "shadow vertex algorithm" for parameterized LPs. This analysis has been subsequently improved and simplified; for recent developments, see Dadush and Huiberts (2018). In smoothed analysis, the LP data are perturbed by some moderate Gaussian noise which cancels "pathological" cases that lead to exponential worst-case examples, like the game constructed in Section 8. Applied to our parameterized LP, it would imply that in expectation there is a polynomial number of segments in Theorem 3. If this holds, the number of Nash equilibria is similarly polynomially bounded by Theorem 5 (the Nash subsets are all single equilibria because the perturbed game is generic and therefore nondegenerate with probability one). However, the standard framework of smoothed analysis (as in, e.g., Dadush and Huiberts 2018) assumes that the LP constraints are of the form $Ax \le \mathbb{1}$, which is not the case for the LP (16) that we consider, so combining this with our approach requires a careful study that we leave for future work. For a general bimatrix game, finding one equilibrium is PPAD-hard even under smoothed analysis (Chen et al. 2009). However, it is not known if a perturbed game may have exponentially long Lemke–Howson paths; the long paths in Savani and von Stengel (2006) do not persist because of exponential size differences in the payoffs.

In Section 8, we described rank-1 games with exponentially many equilibria (also with exponential size differences in the payoffs). This raises the following question: Can all equilibria of a rank-1 game be computed in running time that is polynomial in the size of the input *and output*? Such an algorithm is called "output efficient." For example, the algorithm by Adler and Monteiro (1992) that computes all segments of a parameterized LP is output efficient. We have extended this algorithm in Section 6. For general bimatrix games, an output-efficient algorithm that finds all Nash equilibria would imply P = NP because it is NP-hard to decide if a game has more than one Nash equilibrium (Gilboa and Zemel 1989). Our binary search algorithm gives no information about the existence of a second equilibrium, so it is conceivable that finding a second Nash equilibrium of a rank-1 game is also NP-hard. The existence of an output-efficient algorithm to find all Nash equilibria of a rank-1 game is an open question.

General bimatrix games are computationally difficult, but rank-1 games are computationally easy. One should therefore investigate *economic applications* of large rank-1 games, also as approximate economic models that can serve as fast-solvable benchmarks. As a possible starting point, we describe here a simple "trade game," which suggests that rank-1 games are much more versatile and economically interesting than zero-sum games. Let player 1 be a seller of a product who can choose possible *quality levels* a_i for $i = 1, \ldots, m$, and let player 2 be a buyer who can decide on possible *quantity levels* b_j for $j = 1, \ldots, n$ that the buyer buys from the seller. A *price* p_{ij} that is paid from buyer to seller can be chosen arbitrarily for each i and j. Suppose there are further parameters α, β, γ_j, and δ_i so that the payoffs to the players are

$$\begin{aligned} &\text{payoff to player 1:} \quad && p_{ij} - \alpha a_i b_j + \gamma_j \\ &\text{payoff to player 2:} \quad && -p_{ij} + \beta a_i b_j + \delta_i. \end{aligned} \tag{70}$$

We further assume that $\beta > \alpha > 0$, which reflects that high quality is costly to produce for player 1 and beneficial for player 2, with $\beta - \alpha$ representing the benefits from trade. The additional parameter γ_j (increasing with b_j) is an additional benefit to player 1 for higher amounts of sold quantities and similarly δ_i to player 2 for higher quality. Neither γ_j nor δ_i affect the players' best responses and can therefore be assumed to be zero. This gives a strategically equivalent game whose sums of payoffs are $(\beta - \alpha)a_i b_j$ and therefore of rank one. Because rank-1 games can be analyzed very fast, this "trade game" can be studied for large values of m and n, and in particular for its possibly many price levels. The concrete economic interpretation of such games and their equilibria remains to be investigated. Bulow and Levin (2006, p. 654) consider a "multiplication game," which is a matching game between n workers and n firms where the suitability of a worker for a firm is described by a matrix of rank one. However, it is a game with $2n$ players, not two players.

Acknowledgments

The authors thank two anonymous referees for their detailed comments which helped improve the manuscript. Heinrich Nax suggested the "trade game" (70) in Section 10.

References

Adler I, Monteiro RDC (1992) A geometric view of parametric linear programming. *Algorithmica* 8(1–6):161–176.

Adsul B, Garg J, Mehta R, Sohoni M (2011) Rank-1 bimatrix games: a homeomorphism and a polynomial time algorithm. *Proc. 43rd Annual ACM Sympos. Theory Comput. (STOC)* (Association for Computing Machinery, New York), 195–204.

Avis D (2000) lrs: A revised implementation of the reverse search vertex enumeration algorithm. Kalai G, Ziegler G, eds. *Polytopes–Combinatorics and Computation* (Birkhäuser Verlag, Basel, Switzerland), 177–198.

Avis D, Rosenberg GD, Savani R, von Stengel B (2010) Enumeration of Nash equilibria for two-player games. *Econom. Theory* 42(1):9–37.

Bulow J, Levin J (2006) Matching and price competition. *Amer. Econom. Rev.* 96(3):652–668.

Chen X, Deng X (2006) Settling the complexity of two-player Nash equilibrium. *47th Annual IEEE Sympos. Foundations Comput. Sci. (FOCS)* (IEEE Computer Society, Los Alamitos, CA), 261–272.

Chen X, Deng X, Teng SH (2009) Settling the complexity of computing two-player Nash equilibria. *J. ACM* 56(3):Article 14.

Chen X, Paparas D (2019) Rank-2 bimatrix games are PPAD-hard. August 10, personal communication by email from Xi Chen to Ruta Mehta.

Dadush D, Huiberts S (2018) A friendly smoothed analysis of the simplex method. *Proc. 50th Annual ACM Sympos. Theory Comput. (STOC)* (Association for Computing Machinery, New York), 390–403.

Dantzig GB (1963) *Linear Programming and Extensions* (Princeton University Press, Princeton, NJ).

Daskalakis C, Goldberg PW, Papadimitriou CH (2009) The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39(1):195–259.

Etessami K, Yannakakis M (2010) On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.* 39(6):2531–2597.

Freund RM, Roundy R, Todd MJ (1985) Identifying the set of always-active constraints in a system of linear inequalities by a single linear program. Sloan Working Paper No. 1674-85, Massachusetts Institute of Technology, Cambridge, MA.

Gass S, Saaty T (1955) The computational algorithm for the parametric objective function. *Naval Res. Logist. Quart.* 2(1–2):39–45.

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

Govindan S, Wilson R (2003) A global Newton method to compute Nash equilibria. *J. Econom. Theory* 110(1):65–86.

Jansen B, De Jong J, Roos C, Terlaky T (1997) Sensitivity analysis in linear programming: Just be careful! *Eur. J. Oper. Res.* 101(1):15–28.

Jansen M, Vermeulen D (2001) On the computation of stable sets and strictly perfect equilibria. *Econom. Theory* 17(2):325–344.

Jansen MJM (1981) Maximal Nash subsets for bimatrix games. *Naval Res. Logist. Quart.* 28(1):147–152.

Kannan R, Theobald T (2010) Games of fixed rank: a hierarchy of bimatrix games. *Econom. Theory* 42(1):157–173.

Klee V, Minty GJ (1972) How good is the simplex algorithm? Shisha O, ed. *Inequalities, III* (Academic Press, New York), 159–175.

Kohlberg E, Mertens JF (1986) On the strategic stability of equilibria. *Econometrica* 54(5):1003–1037.

Konno H, Yajima Y, Matsui T (1991) Parametric simplex algorithms for solving a special class of nonconvex minimization problems. *J. Global Optim.* 1(1):65–81.

Lemke CE (1965) Bimatrix equilibrium points and mathematical programming. *Management Sci.* 11(7):681–689.

Lemke CE, Howson JT Jr (1964) Equilibrium points of bimatrix games. *J. Soc. Indust. Appl. Math.* 12(2):413–423.

Mangasarian OL (1964) Equilibrium points of bimatrix games. *J. Soc. Indust. Appl. Math.* 12(4):778–780.

Mangasarian OL, Stone H (1964) Two-person nonzero-sum games and quadratic programming. *J. Math. Anal. Appl.* 9(3):348–355.

Matsui T (1996) NP-hardness of linear multiplicative programming and related problems. *J. Global Optim.* 9(2):113–119.

Mehrotra S, Ye Y (1993) Finding an interior point in the optimal face of linear programs. *Math. Programming* 62(1-3):497–515.

Mehta R (2018) Constant rank two-player games are PPAD-hard. *SIAM J. Comput*. 47(5):1858–1887.

Murty KG (1980) Computational complexity of parametric linear programming. *Math. Programming* 19(1):213–219.

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

Papadimitriou CH (1994) On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. System Sci.* 48(3):498–532.

Savani R, von Stengel B (2006) Hard-to-solve bimatrix games. *Econometrica* 74(2):397–429.

Schrijver A (1986) *Theory of Linear and Integer Programming* (John Wiley & Sons, Chichester, UK).

Spielman DA, Teng SH (2004) Smoothed analysis of algorithms: why the simplex algorithm usually takes polynomial time. *J. ACM* 51(3):385–463.

Theobald T (2009) Enumerating the Nash equilibria of rank 1-games. Avis D, Bremner D, Deza A, eds. *Polyhedral Computation*, CRM Proceedings and Lecture Notes, vol. 48 (American Mathematical Society, Providence, RI), 115–129.

van Damme E (1991) *Stability and Perfection of Nash Equilibria*, 2nd ed. (Springer, Berlin).

Vazirani VV, Yannakakis M (2011) Market equilibrium under separable, piecewise-linear, concave utilities. *J. ACM* 58(3):Article 10.

von Stengel B (1999) New maximal numbers of equilibria in bimatrix games. *Discrete Comput. Geometry* 21(4):557–568.

von Stengel B (2002) Computing equilibria for two-person games. Aumann RJ, Hart S, eds. *Handbook of Game Theory with Economic Applications*, vol. 3 (North-Holland, Amsterdam), 1723–1759.

von Stengel B (2012) Rank-1 games with exponentially many Nash equilibria. Preprint, submitted November 11, https://arxiv.org/abs/1211.2405.

von Stengel B, van den Elzen A, Talman D (2002) Computing normal form perfect equilibria for extensive two-person games. *Econometrica* 70(2):693–715.

Wardlaw WP (2005) Row rank equals column rank. *Math. Magazine* 78(4):316–318.

Wilson R (1992) Computing simply stable equilibria. *Econometrica* 60: 1039–1070.

Bharat Adsul is a professor of computer science at the Indian Institute of Technology, Bombay. He is interested in formal methods in concurrency, logics and games, and geometric complexity theory.

Jugal Garg is an assistant professor in the Department of Industrial and Enterprise Systems Engineering at the University of Illinois at Urbana–Champaign. He is broadly interested in computational aspects of economics and game theory, design and analysis of algorithms, and mathematical programming.

Ruta Mehta is an assistant professor of computer science at the University of Illinois at Urbana–Champaign. Her research focuses on algorithmic, complexity, strategic, and learning aspects of various game-theoretic and economic problems.

Milind Sohoni is a professor of computer science at the Indian Institute of Technology, Bombay. He is broadly interested in combinatorial optimization, mathematical programming, and algorithms.

Bernhard von Stengel is a professor of mathematics at the London School of Economics and Political Science. He is interested in the geometry and computation of Nash equilibria and other mathematical questions of game theory and operations research.