



Rank-1 Bimatrix Games: A Homeomorphism and a Polynomial Time Algorithm *

Bharat Adsul Jugal Garg Ruta Mehta Milind Sohoni
 Indian Institute of Technology, Bombay
 Powai, Mumbai, India
 adsul, jugal, ruta, sohoni@cse.iitb.ac.in

ABSTRACT

Given a rank-1 bimatrix game (A, B) , *i.e.*, where $\text{rank}(A + B) = 1$, we construct a suitable linear subspace of the rank-1 game space and show that this subspace is homeomorphic to its Nash equilibrium correspondence. Using this homeomorphism, we give the first polynomial time algorithm for computing an exact Nash equilibrium of a rank-1 bimatrix game. This settles an open question posed in [8, 19]. In addition, we give a novel algorithm to enumerate all the Nash equilibria of a rank-1 game and show that a similar technique may also be applied for finding a Nash equilibrium of any bimatrix game. Our approach also provides new proofs of important classical results such as the existence and oddness of Nash equilibria, and the index theorem for bimatrix games. Further, we extend the rank-1 homeomorphism result to a fixed rank game space, and give a fixed point formulation on $[0, 1]^k$ for solving a rank- k game. The homeomorphism and the fixed point formulation are piece-wise linear and considerably simpler than the classical constructions.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity

General Terms

Algorithms, Theory

Keywords

Rank-1 games, Homeomorphism, Nash equilibrium

1. INTRODUCTION

Non-cooperative game theory is a model to understand strategic interaction of selfish agents in a given organization. In a finite game, there are finitely many agents, each having

*A full version of this paper is available at arxiv.org/abs/1010.3083

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'11, June 6–8, 2011, San Jose, California, USA.

Copyright 2011 ACM 978-1-4503-0691-1/11/06 ...\$10.00.

finitely many strategies. For these games, Nash [13] proved that there exists a steady state where no player benefits by a unilateral deviation. Such a steady state is called a Nash equilibrium of the game.

Finite games with two agents are also called bimatrix games since they may be represented by two payoff matrices (A, B) , one for each agent. The problem of computing a Nash equilibrium of a bimatrix game is said to be one of the most important concrete open questions on the boundary of \mathcal{P} [15]. The classical Lemke-Howson (LH) algorithm [11] finds a Nash equilibrium of a bimatrix game. However, Savani and von Stengel [17] showed that it is not a polynomial time algorithm by constructing an example, for which the LH algorithm takes an exponential number of steps. Chen and Deng [2] showed that this problem is \mathcal{PPAD} -complete, a complexity class introduced by Papadimitriou [14]. They (together with Teng) [3] also showed that the computation of even a $\frac{1}{n^{\Theta(1)}}$ -approximate Nash equilibrium remains \mathcal{PPAD} -complete. These results suggest that a polynomial time algorithm is unlikely.

There are some results for special cases of the bimatrix games. Lipton et al. [12] considered games where both payoff matrices are of fixed rank k and for these games, they gave a polynomial time algorithm for finding a Nash equilibrium. However, the expressive power of this restricted class of games is limited in the sense that most zero-sum games are not contained in this class. Kannan and Theobald [8] defined a hierarchy of bimatrix games using the rank of $(A + B)$ and gave a polynomial time algorithm to compute an approximate Nash equilibrium for games of a fixed rank k . The set of rank- k games consists of all the bimatrix games with rank at most k . Clearly, rank-0 games are the same as zero-sum games and it is known that the set of Nash equilibria of a zero-sum game is a connected polyhedral set and it may be computed in polynomial time by solving a linear program (LP). Moreover, the problem of finding a Nash equilibrium of zero-sum games and solving linear programs are equivalent [4].

The set of rank-1 games is the smallest extension of zero-sum games in the hierarchy, which strictly generalizes zero-sum games. For any given constant c , Kannan and Theobald [8] also construct a rank-1 game, for which the number of connected components of Nash equilibria is larger than c . This shows that the expressive power of rank-1 games is larger than the zero-sum games. Rank-1 games may also arise in practical situations, in particular the *multiplicative games* between firms and workers in [1] are rank-1 games. A polynomial time algorithm to compute an exact Nash equi-

librium for rank-1 games is an important open problem [8, 19]. Kontogiannis and Spirakis [10] defined the notion of mutual (quasi-) concavity of a bimatrix game and provided a polynomial time computation of a Nash equilibrium for mutually concave games (FPTAS for mutually quasi-concave games). However their classification and the games of fixed rank are incomparable.

Shapley's index theory [18] assigns a sign (also called an index) to a Nash equilibrium of a bimatrix game and shows that the indices of the two endpoints of a Lemke-Howson path have opposite signs. The signs of the endpoints of LH paths provide a direction and in turn a "parity argument" that puts the Nash equilibrium problem of a bimatrix game in \mathcal{PPAD} [14, 20].

The set of bimatrix games (Ω), with m and n strategies of the first and the second player respectively, forms a Euclidean space, *i.e.*, $\Omega = \{(A, B) \in \mathbb{R}^{mn} \times \mathbb{R}^{mn}\} = \mathbb{R}^{2mn}$. Kohlberg and Mertens [9] showed that Ω is homeomorphic to its Nash equilibrium correspondence¹ $E_\Omega = \{(A, B, x, y) \in \mathbb{R}^{2mn+m+n} \mid (x, y) \text{ is a Nash equilibrium of } (A, B)\}$. This structural result has been used extensively to understand the index, degree and the stability of a Nash equilibrium of a bimatrix game [5, 9]. Moreover, the homeomorphism result also validates the homotopy methods devised to compute a Nash equilibrium [6, 7]. The structural result has been extended for more general game spaces [16], however, to the best of our knowledge, no such result is known for special subspaces of the bimatrix game space. Such a result may pave a way to device better algorithms for the Nash equilibrium computation or to prove the hardness of computing a Nash equilibrium for the games in the subspace.

Our contributions. For a given rank-1 game $(A, B) \in \mathbb{R}^{mn} \times \mathbb{R}^{mn}$, the matrix $(A + B)$ may be written as $\alpha \cdot \beta^T$, where $\alpha \in \mathbb{R}^m$ and $\beta \in \mathbb{R}^n$. Motivated by this fact, in Section 2.2, we define an m -dimensional subspace $\Gamma = \{(A, C + \alpha \cdot \beta^T) \mid \alpha \in \mathbb{R}^m\}$ of Ω , where $A \in \mathbb{R}^{mn}$, $C \in \mathbb{R}^{mn}$ and $\beta \in \mathbb{R}^n$ are fixed and analyze the structure of its Nash equilibrium correspondence $E_\Gamma = \{(A', B', x, y) \mid (x, y) \text{ is a Nash equilibrium of } (A', B') \in \Gamma\}$. For a given bimatrix game (A', B') , the best response polytopes P and Q may be defined using the payoff matrices A' and B' respectively (described in Section 2.1). There is a notion of fully-labeled points of $P \times Q$, which capture all the Nash equilibria of the game [21]. Note that the polytope P is same for all the games in Γ since the payoff matrix of the first player is fixed to A . However the payoff matrix of the second player varies with α , hence Q is different for every game. We define a new polytope Q' in Section 2.2, which encompasses Q for all the games in Γ . We show that the set of fully-labeled points of $P \times Q'$, say \mathcal{N} , captures all the Nash equilibria of all the games in Γ and in turn captures E_Γ .

Surprisingly, \mathcal{N} turns out to be a set of cycles and a single path on the 1-skeleton of $P \times Q'$ under the non-degeneracy assumption. We refer to the path in \mathcal{N} as the fully-labeled path and show that it captures at least one Nash equilibrium of every game in Γ . The structure of \mathcal{N} also proves the existence and the oddness of the number of Nash equilibria in a non-degenerate bimatrix game. Moreover, an edge of \mathcal{N} may be efficiently oriented, and using this orientation, we determine the index of every Nash equilibria for a bimatrix game.

¹The actual result is for N player game space.

Further, in Section 3 we show that if Γ contains only rank-1 games (*i.e.*, $C = -A$) then \mathcal{N} does not contain cycles and the fully-labeled path exhibits a strict monotonicity. Using this monotonic nature, we establish a homeomorphism map between Γ and E_Γ . This is the first structural result for a subspace of the bimatrix game space. The homeomorphism maps that we derive are very different than the ones given by Kohlberg and Mertens for the bimatrix game space [9], and require a structural understanding of E_Γ .

Using the above facts on the structure of \mathcal{N} , in Section 4 we present two algorithms. For a given non-degenerate rank-1 game $(A, -A + \gamma \cdot \beta^T)$, we consider the subspace $\Gamma = \{(A, -A + \alpha \cdot \beta^T) \mid \alpha \in \mathbb{R}^m\}$. Note that Γ contains the given game and the corresponding set \mathcal{N} is a path which captures all the Nash equilibria of the game. The first algorithm (*BinSearch*) finds a Nash equilibrium of a rank-1 game in polynomial time by applying binary search on the fully-labeled path using the monotonic nature of the path. The algorithm works for degenerate games as well with a minor modification. This is the first polynomial time algorithm to find an exact Nash equilibrium of a rank-1 game.

The second algorithm (*Enumeration*) enumerates all the Nash equilibria of a rank-1 game. Using the fact that \mathcal{N} contains only the fully-labeled path, the *Enumeration* algorithm traces this path and locates all the Nash equilibria of the game. For an arbitrary bimatrix game, we may define a suitable Γ containing the game. Since the fully-labeled path of the corresponding \mathcal{N} covers at least one Nash equilibrium of all the games in Γ , the *Enumeration* algorithm locates at least one Nash equilibrium of the given bimatrix game. Theobald [19] also gave an algorithm to enumerate all the Nash equilibria of a rank-1 game, however it may not be generalized to find a Nash equilibrium of any bimatrix game. Moreover, our algorithm is much simpler and a detailed comparison is given in Section 4.2. There, we also compare our algorithm with the Lemke-Howson algorithm, which follows a path of *almost* fully-labeled points [21].

For a given rank- k game (A, B) , the matrix $(A + B)$ may be written as $\sum_{l=1}^k \gamma^l \cdot \beta^l{}^T$, where $\forall l, \gamma^l \in \mathbb{R}^m$ and $\beta^l \in \mathbb{R}^n$. We define a km -dimensional affine subspace $\Gamma^k = \{(A, -A + \sum_{l=1}^k \alpha^l \cdot \beta^l{}^T) \mid \alpha^l \in \mathbb{R}^m, \forall l\}$ of Ω . In Section 5, we establish a homeomorphism between Γ^k and its Nash equilibrium correspondence E_{Γ^k} using techniques similar to the rank-1 homeomorphism. Further, to find a Nash equilibrium of a rank- k game we give a piece-wise linear polynomial-time computable fixed point formulation on $[0, 1]^k$ using the homeomorphism result and discuss the possibility of a polynomial time algorithm.

2. GAMES AND NASH EQUILIBRIUM

2.1 Preliminaries

Notations. For a matrix $A = [a_{ij}] \in \mathbb{R}^{mn}$ of dimension $m \times n$, let A_i be the i^{th} row and A^j be the j^{th} column of the matrix. For a vector $\alpha \in \mathbb{R}^m$, let α_i be its i^{th} coordinate. Vectors are considered as column vectors.

For a finite two-player game, let the strategy sets of the first and the second player be $S_1 = \{1, \dots, m\}$ and $S_2 = \{1, \dots, n\}$ respectively. The payoff function of such a game may be represented by the two payoff matrices $(A, B) \in \mathbb{R}^{mn} \times \mathbb{R}^{mn}$, each of dimension $m \times n$. If the played strategy profile is $(i, j) \in S_1 \times S_2$, then the payoffs of the first and

second players are a_{ij} and b_{ij} respectively. Note that the rows of these matrices correspond to the strategies of the first player and the columns to that of the second player, hence the first player is also referred to as the *row-player* and second player as the *column-player*.

The strategies in S_1 and S_2 are called *pure strategies*. A *mixed strategy* is a probability distribution over the available set of strategies. The set of mixed strategies for the row-player is $\Delta_1 = \{(x_1, \dots, x_m) \mid x_i \geq 0, \forall i \in S_1, \sum_{i \in S_1} x_i = 1\}$ and for the column-player, it is $\Delta_2 = \{(y_1, \dots, y_n) \mid y_j \geq 0, \forall j \in S_2, \sum_{j \in S_2} y_j = 1\}$. If the strategy profile $(x, y) \in \Delta_1 \times \Delta_2$ is played, then the payoffs of the row-player and column-player are $x^T A y$ and $x^T B y$ respectively.

A strategy profile is said to be a Nash equilibrium strategy profile (NESP) if no player achieves a better payoff by a unilateral deviation [13]. Formally, $(x, y) \in \Delta_1 \times \Delta_2$ is a NESP iff $\forall x' \in \Delta_1, x^T A y \geq x'^T A y$ and $\forall y' \in \Delta_2, x^T B y \geq x^T B y'$. These conditions may also be equivalently stated as,

$$\begin{aligned} \forall i \in S_1, x_i > 0 &\Rightarrow A_i y = \max_{k \in S_1} A_k y \\ \forall j \in S_2, y_j > 0 &\Rightarrow x^T B^j = \max_{k \in S_2} x^T B^k \end{aligned} \quad (1)$$

From (1), it is clear that at a Nash equilibrium, a player plays a pure strategy with non-zero probability only if it gives the maximum payoff with respect to (w.r.t.) the opponent's strategy. Such strategies are called the *best response strategies* (w.r.t. the opponent's strategy). The polytope P in (2) is closely related to the best response strategies of the row-player for a given strategy (y) of the column-player [21] and it is called the *best response polytope* of the row-player. Similarly, the polytope Q is called the best response polytope of the column-player. In the following expression, x and y are vector variables, and π_1 and π_2 are scalar variables.

$$\begin{aligned} P &= \{(y, \pi_1) \in \mathbb{R}^{n+1} \mid A_i y - \pi_1 \leq 0, \forall i \in S_1; \\ &\quad y_j \geq 0, \forall j \in S_2; \sum_{j=1}^n y_j = 1\} \\ Q &= \{(x, \pi_2) \in \mathbb{R}^{m+1} \mid x_i \geq 0, \forall i \in S_1; \\ &\quad x^T B^j - \pi_2 \leq 0, \forall j \in S_2; \sum_{i=1}^m x_i = 1\} \end{aligned} \quad (2)$$

Note that for any $y' \in \Delta_2$, a unique (y', π'_1) may be obtained on the boundary of P , where $\pi'_1 = \max_{i \in S_1} A_i y'$. Clearly, the pure strategy $i \in S_1$ is in the best response against y' only if $A_i y' - \pi'_1 = 0$, hence indices in S_1 corresponding to the tight inequalities at (y', π'_1) are in the best response. Note that, in both the polytopes the first set of inequalities correspond to the row-player, and the second set correspond to the column player. Since $|S_1| = m$ and $|S_2| = n$, let the inequalities be numbered from 1 to m , and $m+1$ to $m+n$ in both the polytopes. Let the *label* $L(v)$ of a point v in the polytope be the set of indices of the tight inequalities at v . A pair $(v, w) \in P \times Q$ is called *fully-labeled pair* if $L(v) \cup L(w) = \{1, \dots, m+n\}$.

LEMMA 1. [21] A strategy profile (x, y) is a NESP of the game (A, B) iff $((y, \pi_1), (x, \pi_2)) \in P \times Q$ is a fully-labeled pair, for some π_1 and π_2 .

A game is called non-degenerate if both the polytopes are non-degenerate. Note that for a non-degenerate game, $|L(v)| \leq n$ and $|L(w)| \leq m, \forall (v, w) \in P \times Q$, and the equality holds iff v and w are the vertices of P and Q respectively. Therefore, a fully-labeled pair of a non-degenerate game has to be a vertex-pair.

Let $\Omega = \{(A, B) \in \mathbb{R}^{mn} \times \mathbb{R}^{mn}\} = \mathbb{R}^{2mn}$ be the bimatrix game space and $E_\Omega = \{(A, B, x, y) \in \mathbb{R}^{mn} \times \mathbb{R}^{mn} \times \Delta_1 \times$

$\Delta_2 \mid (x, y) \text{ is a NESP of the game } (A, B)\}$ be it's Nash equilibrium correspondence. Kohlberg and Mertens [9] proved that E_Ω is homeomorphic to the bimatrix game space $\mathbb{R}^{2mn}(\Omega)$. No such structural result is known for a subspace of the bimatrix game space \mathbb{R}^{2mn} . With the hope of establishing such a result for a subspace, we define an m -dimensional affine subspace of \mathbb{R}^{2mn} and analyze the structure of it's Nash equilibrium correspondence in the next section.

2.2 Game Space and the Nash Equilibrium Correspondence

Let $\Gamma = \{(A, C + \alpha \cdot \beta^T) \mid \alpha \in \mathbb{R}^m\}$ be a game space, where $A \in \mathbb{R}^{mn}$ and $C \in \mathbb{R}^{mn}$ are $m \times n$ dimensional non-zero matrices, and $\beta \in \mathbb{R}^n$ is an n -dimensional non-zero vector. Note that for a game $(A, B) \in \Gamma$, there exists a unique $\alpha \in \mathbb{R}^m$, such that $B = C + \alpha \cdot \beta^T$. Therefore, Γ may be parametrized by α , and let $G(\alpha)$ be the game $(A, C + \alpha \cdot \beta^T) \in \Gamma$. Clearly, Γ forms an m -dimensional affine subspace of the bimatrix game space \mathbb{R}^{2mn} . Let $E_\Gamma = \{(\alpha, x, y) \in \mathbb{R}^m \times \Delta_1 \times \Delta_2 \mid (x, y) \text{ is a NESP of the game } G(\alpha) \in \Gamma\}$ be the Nash equilibrium correspondence of Γ . We wish to investigate: *Is E_Γ homeomorphic to the game space Γ ($\equiv \mathbb{R}^m$)?*

For a game $G(\alpha) \in \Gamma$, let the best response polytopes of row-player and column-player be $P(\alpha)$ and $Q(\alpha)$ respectively. Since the row-player's matrix is fixed to A , $P(\alpha)$ is the same for all α and we denote it by P . However, $Q(\alpha)$ varies with α . We define a new polytope Q' in (3), which encompasses $Q(\alpha)$, for all $G(\alpha) \in \Gamma$.

$$Q' = \{(x, \lambda, \pi_2) \in \mathbb{R}^{m+2} \mid x_i \geq 0, \forall i \in S_1; \\ x^T C^j + \beta_j \lambda - \pi_2 \leq 0, \forall j \in S_2; \sum_{i=1}^m x_i = 1\} \quad (3)$$

Note that the inequalities of Q' may also be numbered from 1 to $m+n$ in a similar fashion as in Q . For a game $G(\alpha)$, the polytope $Q(\alpha)$ may be obtained by replacing λ by $\sum_{i=1}^m \alpha_i x_i$ in Q' . In other words, $Q(\alpha)$ is the projection of $Q' \cap \{(x, \lambda, \pi_2) \mid \sum_{i=1}^m \alpha_i x_i - \lambda = 0\}$ on the (x, π_2) -space. Let $\mathcal{N} = \{(v, w) \in P \times Q' \mid L(v) \cup L(w) = \{1, \dots, m+n\}\}$ be the set of fully-labeled pairs in $P \times Q'$. Let $\Psi : E_\Gamma \rightarrow \mathcal{N}$ be such that,

$$\begin{aligned} \Psi(\alpha, x, y) &= ((y, \pi_1), (x, \lambda, \pi_2)), \text{ where } \pi_1 = x^T A y, \\ &\quad \lambda = \sum_{i=1}^m \alpha_i x_i, \pi_2 = x^T (C + \alpha \cdot \beta^T) y \end{aligned}$$

LEMMA 2. The map Ψ is well defined and surjective.

PROOF. For a point $(\alpha, x, y) \in E_\Gamma$, the corresponding $\Psi(\alpha, x, y)$ is fully-labeled (clear from Lemma 1) and hence lies in \mathcal{N} .

Let $(v, w) \in \mathcal{N}$ be a fully-labeled pair with $v = (y, \pi_1)$ and $w = (x, \lambda, \pi_2)$. Let $\alpha \in \mathbb{R}^m$ be such that $\sum_{i=1}^m \alpha_i x_i - \lambda = 0$, then clearly $(v, (x, \pi_2)) \in P(\alpha) \times Q(\alpha)$ is a fully-labeled pair. Therefore, $(\alpha, x, y) \in E_\Gamma$ and $(v, w) = \Psi(\alpha, x, y)$. \square

We further strengthen the connection between E_Γ and \mathcal{N} with the following lemma.

LEMMA 3. E_Γ is connected iff \mathcal{N} is a single connected component.

PROOF. (\Rightarrow) Since Ψ is a continuous surjective function from E_Γ to \mathcal{N} (Lemma 2), if E_Γ is connected then \mathcal{N} is connected as well.

(\Leftarrow) For a $(v, w) \in \mathcal{N}$, where $w = (x, \lambda, \pi_2)$, all the points in $\Psi^{-1}(v, w)$ satisfy $\sum_{i=1}^m \alpha_i x_i = \lambda$, hence $\Psi^{-1}(v, w)$ is homeomorphic to \mathbb{R}^{m-1} . Since \mathcal{N} is connected, Ψ is continuous and the fact that the fibers $\Psi^{-1}(v, w), \forall (v, w) \in \mathcal{N}$ are connected imply that E_Γ is connected. \square

Lemma 2 and 3 imply that E_Γ and \mathcal{N} are closely related. Henceforth, we assume that the polytopes P and Q' are non-degenerate. Recall that when the best response polytopes (P and Q) of a game are non-degenerate, all the fully-labeled pairs are vertex pairs. However Q' has one more variable λ than Q , which gives one extra degree of freedom to form the fully-labeled pairs. We show that the structure of \mathcal{N} is very simple by proving the following proposition.

PROPOSITION 1. *The set of fully-labeled points \mathcal{N} admits the following decomposition into mutually disjoint connected components: $\mathcal{N} = \mathcal{P} \cup \mathcal{C}_1 \cup \dots \cup \mathcal{C}_k$, $k \geq 0$, where \mathcal{P} and \mathcal{C}_k s respectively form a path and cycles on 1-skeleton of $P \times Q'$.*

In order to prove Proposition 1, first we identify the points in P and Q' separately, which participate in the fully-labeled pairs and then relate them. For a $v \in P$, let $\mathcal{E}_v = \{w' \in Q' \mid (v, w') \in \mathcal{N}\}$, and similarly for a $w \in Q'$, let $\mathcal{E}_w = \{v' \in P \mid (v', w) \in \mathcal{N}\}$. Let $\mathcal{N}^P = \{v \in P \mid \mathcal{E}_v \neq \emptyset\}$ and $\mathcal{N}^{Q'} = \{w \in Q' \mid \mathcal{E}_w \neq \emptyset\}$.

For neighboring vertices u and v in either polytopes, let $\overline{u, v}$ be the edge between u and v . Recall that P and Q' are non-degenerate, therefore $\forall v \in P$, $|L(v)| \leq n$ and $\forall w \in Q'$, $|L(w)| \leq m+1$. Using this fact, it is easy to deduce the following observations for points in P . Similar results hold for the points in Q' .

- O_1 . If $(v, w) \in \mathcal{N}$, then both v and w lie on either 0 or 1-dimensional faces of P and Q' respectively, and at least one of them is a 0-dimensional face, i.e., a vertex.
- O_2 . If $v \in P$ is a vertex, then \mathcal{E}_v is either empty or an edge of Q' . If v is on an edge, then \mathcal{E}_v , if non-empty, is exactly one vertex of Q' .
- O_3 . If $v \in P$ is neither a vertex nor on an edge, then $\mathcal{E}_v = \emptyset$.
- O_4 . Let $(v, w) \in \mathcal{N}$ and both v and w be vertices. Since $|L(v)| = n$, $|L(w)| = m+1$ and $|L(v) \cup L(w)| = m+n$, $|L(v) \cap L(w)| = 1$ and the element in the intersection is called the *duplicate label* of the pair (v, w) .
- O_5 . Let $v \in P$ be a vertex and \mathcal{E}_v be an edge of Q' . If $w \in \mathcal{E}_v$ is a vertex, then (v, w) has a duplicate label (see O_4). Let the duplicate label be i , then there exists a unique vertex $v' \in P$ adjacent to v such that $v, v' \in \mathcal{N}^P$, where v' is obtained by relaxing the inequality i at v . This also implies that $\mathcal{E}_w = \overline{v, v'}$ and $\mathcal{E}_v \cap \mathcal{E}_{v'} = w$.

The above observations, bring out the structure of \mathcal{N} significantly. Every point in \mathcal{N} is a pair (v, w) where $v \in P$ and $w \in Q'$. From O_1 , one of them is a vertex (say v), and the other is on the corresponding edge ($w \in \mathcal{E}_v$). Hence \mathcal{N} contains only 0 and 1-dimensional faces of $P \times Q'$. Clearly, an edge of \mathcal{N} is of type (v, \mathcal{E}_v) or (\mathcal{E}_w, w) , where v and w are the vertices of P and Q' respectively.

Note that a vertex (v, w) of \mathcal{N} corresponds to a fully-labeled vertex-pair of $P \times Q'$, and hence it has a duplicate label (by O_4). Relaxing the inequality corresponding to the duplicate label in P and Q' separately, we get the edges (\mathcal{E}_w, w) and (v, \mathcal{E}_v) of \mathcal{N} respectively. Clearly, these are the only adjacent edges of the vertex (v, w) in \mathcal{N} . Hence, in a component of \mathcal{N} , edges alternate between type (v, \mathcal{E}_v) and (\mathcal{E}_w, w) , and the degree of every vertex of \mathcal{N} is exactly two.

Therefore, \mathcal{N} consists of infinite paths and cycles on the 1-skeleton of $P \times Q'$. Note that a path in \mathcal{N} has unbounded edges on both the sides.

Using the above analysis, we only need to show that there is exactly one path in \mathcal{N} to prove Proposition 1. Let the *support-pair* of a vertex $(y, \pi_2) \in P$ be (I, J) where $I = \{i \in S_1 \mid A_i y - \pi_2 = 0\}$ and $J = \{j \in S_2 \mid y_j > 0\}$. Note that $|L(y, \pi_2)| = n$, hence $|I| = |J|$. Let $\beta_{j_s} = \min_{j \in S_2} \beta_j$, $i_s = \arg \max_{i \in S_1} a_{ij_s}$, $\beta_{j_e} = \max_{j \in S_2} \beta_j$, and $i_e = \arg \max_{i \in S_1} a_{ij_e}$. In other words, the indices j_s and j_e correspond to the minimum and maximum entries in β respectively, and the indices i_s and i_e correspond to the maximum entry in A^{j_s} and A^{j_e} respectively. It is easy to see that $j_s \neq j_e$, since Q' is non-degenerate.

LEMMA 4. *There exist two vertices v_s and v_e in P , with support-pairs $(\{i_s\}, \{j_s\})$ and $(\{i_e\}, \{j_e\})$ respectively.*

PROOF. Let $y \in \Delta_2$ be such that $y_{j_s} = 1$ and $y_j = 0$, $\forall j \neq j_s$. Clearly, $v_s = (y, a_{i_s j_s}) \in P$ and $|L(v_s)| = n$. Similarly, the vertex $v_e \in P$ may be obtained by setting $y_{j_e} = 1$ and the remaining y_j s to zero. \square

Next we show that there are exactly two unbounded edges of type (v, \mathcal{E}_v) in \mathcal{N} , all other edges have two bounding vertices.

LEMMA 5. *An edge $(v, \mathcal{E}_v) \in \mathcal{N}$ has exactly one bounding vertex if v is either v_s or v_e , otherwise it has two bounding vertices.*

PROOF. Let $v = v_s$. The points in \mathcal{E}_v satisfy

$$x_{i_s} = 1 \text{ and } \forall i \neq i_s, x_i = 0, \quad \pi_2 = c_{i_s j_s} + \beta_{j_s} \lambda \quad (4)$$

$$\forall j \neq j_s, c_{i_s j} + \beta_j \lambda \leq c_{i_s j_s} + \beta_{j_s} \lambda$$

Since $\beta_j \geq \beta_{j_s}$, we get $\lambda \leq \frac{c_{i_s j_s} - c_{i_s j}}{\beta_j - \beta_{j_s}}$. Let $\lambda_s = \min_{j \neq j_s} \frac{c_{i_s j_s} - c_{i_s j}}{\beta_j - \beta_{j_s}}$, then $\mathcal{E}_v = \{(x, \lambda, \pi_2) \mid \lambda \in (-\infty, \lambda_s], x \text{ and } \pi_2 \text{ satisfy (4)}\}$. Note that on \mathcal{E}_v , x is a constant and λ varies from $-\infty$ to λ_s . Moreover the point corresponding to $\lambda = \lambda_s$ is a vertex, because one more inequality becomes tight there. Similarly for $v = v_e$, λ varies from $\lambda_e = \max_{j \neq j_e} \frac{c_{i_e j} - c_{i_e j_e}}{\beta_{j_e} - \beta_j}$ to ∞ on \mathcal{E}_v , and $\lambda = \lambda_e$ corresponds to a vertex of \mathcal{E}_v .

Let a vertex $v \in P$ be such that $v \neq v_s$, $v \neq v_e$ and $\mathcal{E}_v \neq \emptyset$. We show that \mathcal{E}_v has exactly two bounding vertices. Let (I, J) be the support-pair corresponding to v . There are two cases.

Case 1 - $|I| = |J| = 1$: Let $I = \{i_1\}$ and $J = \{j_1\}$. Then for all the points in \mathcal{E}_v , $x_{i_1} = 1$ and all other x_i s are zero. Let $J_l = \{j \mid \beta_j < \beta_{j_1}\}$ and $J_g = \{j \mid \beta_j > \beta_{j_1}\}$. Clearly $j_s \in J_l$ and $j_e \in J_g$. All the points in \mathcal{E}_v must satisfy the inequalities $c_{i_1 j} + \beta_j \lambda \leq c_{i_1 j_1} + \beta_{j_1} \lambda$, $\forall j \notin J$, and using them, we get the following upper and lower bounds on λ .

$$\max_{j \in J_g} \frac{c_{i_1 j_1} - c_{i_1 j}}{\beta_j - \beta_{j_1}} \leq \lambda \leq \max_{j \in J_l} \frac{c_{i_1 j} - c_{i_1 j_1}}{\beta_{j_1} - \beta_j}$$

Therefore, the values of λ on \mathcal{E}_v , form a closed and bounded interval, and for each extreme point of this interval, there is a vertex in \mathcal{E}_v .

Case 2 - $|I| = |J| > 1$: Note that exactly m inequalities of Q' are tight at \mathcal{E}_v because $|L(v)| = n$ and Q' is non-degenerate. These m tight inequalities with $\sum_{i=1}^m x_i =$

1 form a 1-dimensional line L in the (x, λ, π_2) -space, and clearly $\mathcal{E}_v = L \cap Q'$. Let $w = (x, \lambda, \pi_2) \in L$ and d be a unit vector along the line L . For a $w' \in L$, there exists a unique $\epsilon \in \mathbb{R}$ such that $w' = w + \epsilon d$. Let $d(x_i)$ be the coordinate of d corresponding to x_i . Note that $\sum_{i=1}^m d(x_i) = 0$, because L satisfies $\sum_{i=1}^m x_i = 1$. Further $\exists i \in I$ such that $d(x_i) \neq 0$, otherwise x becomes constant on L , which in turn imply that λ and π_2 are also constants on L . Hence $\exists i_1, i_2 \in I$ s.t. $d(x_{i_1}) > 0$ and $d(x_{i_2}) < 0$. For all the points in \mathcal{E}_v , the inequalities $x_i \geq 0$, $\forall i \in I$ hold. Using these, we get

$$x_{i_1} + \epsilon d(x_{i_1}) \geq 0, x_{i_2} + \epsilon d(x_{i_2}) \geq 0 \Rightarrow \frac{x_{i_1}}{d(x_{i_1})} \leq \epsilon \leq \frac{x_{i_2}}{d(x_{i_2})}$$

From the above observations, we may easily deduce that the set $\{\epsilon \mid w + \epsilon d \in \mathcal{E}_v\}$ is a closed and bounded interval $[b_l, b_u]$. Moreover, at the extreme points $w_u = w + b_u d$ and $w_l = w + b_l d$ of \mathcal{E}_v , one more inequality is tight. Therefore, w_u and w_l are the vertices in \mathcal{E}_v . \square

Now we are in a position to prove Proposition 1.

Proof of Proposition 1:

For a vertex $w = (x, \lambda, \pi_2) \in \mathcal{N}^{Q'}$, $\exists r \leq m$ such that $x_r > 0$ since $\sum_{i=1}^m x_i = 1$. In that case, $A_r y = \pi_1$ holds on the corresponding edge $\mathcal{E}_w \in \mathcal{N}^P$ (O_2). This implies that the edge \mathcal{E}_w is bounded from both the sides, since $\forall j \in S_2$, $0 \leq y_j \leq 1$ and $A_{min} \leq \pi_1 \leq A_{max}$ on the edge \mathcal{E}_w , where $A_{min} = \min_{(i,j) \in S_1 \times S_2} a_{ij}$ and $A_{max} = \max_{(i,j) \in S_1 \times S_2} a_{ij}$. Therefore, there are exactly two unbounded edges in the set \mathcal{N} namely (v_s, \mathcal{E}_{v_s}) and (v_e, \mathcal{E}_{v_e}) (Lemma 5). This proves that \mathcal{N} contains exactly one path \mathcal{P} , with unbounded edges (v_s, \mathcal{E}_{v_s}) and (v_e, \mathcal{E}_{v_e}) at both the ends. All the other components of \mathcal{N} form cycles (\mathcal{C}_i s). \square

From Proposition 1, it is clear that \mathcal{N} contains at least the path \mathcal{P} . We show the importance of \mathcal{P} in the next two lemmas.

LEMMA 6. For every $a \in \mathbb{R}$, there exists a point $((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{P}$ such that $\lambda = a$.

PROOF. Since \mathcal{P} is a continuous path in $P \times Q'$ (Proposition 1), therefore λ changes continuously on \mathcal{P} . Moreover, in the proof of Lemma 5, we saw that on the edge $(v_s, \mathcal{E}_{v_s}) \in \mathcal{P}$, λ varies from $-\infty$ to λ_s and on the edge $(v_e, \mathcal{E}_{v_e}) \in \mathcal{P}$ it varies from λ_e to ∞ . Therefore for any $a \in \mathbb{R}$, there is a point $((y, \pi_1), (x, \lambda, \pi_2))$ in \mathcal{P} such that $\lambda = a$. \square

Consider a game $\alpha \in \Gamma$, and the corresponding hyper-plane $H : \lambda - \sum_{i=1}^m \alpha_i x_i = 0$. Note that, every point in $\mathcal{N} \cap H$ corresponds to a NESP of the game $G(\alpha)$ and vice-versa.

LEMMA 7. The path \mathcal{P} of \mathcal{N} covers at least one NESP of the game $G(\alpha)$.

PROOF. If there are points in \mathcal{P} on opposite sides of H , then the set $\mathcal{P} \cap H$ has to be non-empty. Let $w_1 = (x^1, \lambda_1, \pi_1^1)$ and $w_2 = (x^2, \lambda_2, \pi_1^2)$ in \mathcal{P} be s.t. $\lambda_1 = \min_{i \in S_1} \alpha_i$ and $\lambda_2 = \max_{i \in S_1} \alpha_i$. Note that w_1 and w_2 exist (Lemma 6) and they satisfy $\lambda_1 - \sum_{i=1}^m \alpha_i x_i^1 \leq 0$ and $\lambda_2 - \sum_{i=1}^m \alpha_i x_i^2 \geq 0$. \square

Remark 1. The proof of Lemma 7 in fact shows the existence of a Nash equilibrium for a bimatrix game. It is also easy to deduce that the number of Nash equilibria of a

non-degenerate bimatrix game is odd from the fact that \mathcal{N} contains a set of cycles and a path (Proposition 1), simply because a cycle must intersect the hyper-plane H an even number of times, and the path must intersect H an odd number of times.

From the proof of Proposition 1, it is clear that every vertex of \mathcal{N} has a duplicate label and the two edges incident on a vertex may be easily obtained by relaxing the inequality corresponding to the duplicate label in P and in Q' . Therefore, given a point of some component of \mathcal{N} , it is easy to trace the full component by leaving the duplicate label in P and Q' alternately at every vertex. The next lemma shows that the components of \mathcal{N} may be easily oriented.

LEMMA 8. Let E be the set of edges of \mathcal{N} , and $E' = \{\overrightarrow{u, u'}, \overleftarrow{u, u'} \mid \overline{u, u'} \in \mathcal{N}\}$ be the set of directed edges. There exists a (efficiently computable) function $\rightarrow : E \rightarrow E'$ such that it maps a cycle of \mathcal{N} to a directed cycle and the path \mathcal{P} to a path oriented from (v_s, \mathcal{E}_{v_s}) to (v_e, \mathcal{E}_{v_e}) .

The direction of the edges of \mathcal{N} , defined by function \rightarrow , may be used to determine the index (see [20] for definition) of every Nash equilibrium for a game in Γ . Let H^- and H^+ be the half-spaces corresponding to the hyper-plane H ($\lambda - \sum_{i=1}^m \alpha_i x_i = 0$) of $G(\alpha)$.

PROPOSITION 2. Let an edge $\overline{u, u'} \in \mathcal{N}$ intersect H at a NESP (x, y) of $G(\alpha)$, and let $\rightarrow(\overline{u, u'}) = \overrightarrow{u, u'}$. If $u \in H^-$ and $u' \in H^+$ then index of (x, y) is +1, otherwise it is -1.

From Proposition 2, it is easy to see that in a component, the index of Nash equilibria alternates. Further, both the first and the last Nash equilibria, on the path \mathcal{P} , have index +1. This proves that the number of Nash equilibria with index +1 is one more than the number of Nash equilibria with index -1, which is an important known result [20, 18].

Recall that if \mathcal{N} is disconnected, then E_Γ is also disconnected (Lemma 3). Example 1 shows that E_Γ may be disconnected in general by illustrating a disconnected \mathcal{N} .

Example 1. Consider the following A , C and β .

$$A = \begin{bmatrix} 0 & 9 & 9 \\ 6 & 6 & 5 \\ 9 & 7 & 2 \end{bmatrix}, \quad C = \begin{bmatrix} 6 & 8 & 6 \\ 5 & 8 & 8 \\ 4 & 3 & 0 \end{bmatrix}, \quad \beta = \begin{bmatrix} 9 \\ 7 \\ 8 \end{bmatrix}$$

The set \mathcal{N} of the corresponding game space Γ contains a path \mathcal{P} and a cycle \mathcal{C}_1 . From Proposition 1, it is clear that a component of \mathcal{N} may be obtained from a component of \mathcal{N}^P and the corresponding component of $\mathcal{N}^{Q'}$. Therefore we demonstrate the path \mathcal{P}^P and the cycle \mathcal{C}_1^P of \mathcal{N}^P , and using them \mathcal{P} and \mathcal{C}_1 of \mathcal{N} may be easily obtained. The path \mathcal{P}^P is $\overline{v_s, v_1}, \overline{v_1, v_e}$, where $v_s = ((0, 1, 0), 9)$, $v_1 = ((0.18, 0.82, 0), 7.36)$ and $v_e = ((1, 0, 0), 9)$. The cycle \mathcal{C}_1^P is $\overline{v_2, v_3}, \overline{v_3, v_4}, \overline{v_4, v_2}$, where $v_2 = ((0.5, 0, 0.5), 5.5)$, $v_3 = ((0.38, 0.18, 0.44), 5.56)$ and $v_4 = ((0.4, 0, 0.6), 5.4)$. Note that v_s and v_e correspond to the minimum and maximum β_j respectively (Lemma 5). \square

Since Γ ($\equiv \mathbb{R}^m$) is connected, if E_Γ is disconnected then it is not homeomorphic to Γ .

3. RANK-1 SPACE: HOMEOMORPHISM

From the discussion of the last section, we know that Γ and E_Γ are not homeomorphic in general (illustrated by Example 1). Surprisingly, they turn out to be homeomorphic if Γ consists of only rank-1 games, i.e., $C = -A$. Recall that E_Γ forms a single connected component iff \mathcal{N} has only one component (Lemma 3). First we show that when $C = -A$, the set \mathcal{N} consists of only a path.

For a given matrix $A \in \mathbb{R}^{m \times n}$ and a vector $\beta \in \mathbb{R}^n$, we fix the game space to $\Gamma = \{(A, -A + \alpha \cdot \beta^T) \mid \alpha \in \mathbb{R}^m\}$. We assume that A and β are non-zero and the corresponding polytopes P and Q' are non-degenerate. Lemma 9 shows that the set \mathcal{N} may be easily identified on the polytope $P \times Q'$.

LEMMA 9. *For all $(v, w) = ((y, \pi_1), (x, \lambda, \pi_2))$ in $P \times Q'$, we have $\lambda(\beta^T \cdot y) - \pi_1 - \pi_2 \leq 0$, and the equality holds iff $(v, w) \in \mathcal{N}$.*

PROOF. Recall that $C = -A$, hence from (2) and (3), we get $x^T \cdot (A \cdot y - \pi_1) \leq 0$ and $(x^T \cdot (-A) + \beta^T \lambda - \pi_2) \cdot y \leq 0$. By summing up these two inequalities, we get $\lambda(\beta^T \cdot y) - \pi_1 - \pi_2 \leq 0$. If $(v, w) \in \mathcal{N}$, then $\forall i \leq m, x_i > 0 \Rightarrow A_i \cdot y - \pi_1 = 0$ and $\forall j \leq n, y_j > 0 \Rightarrow x^T(-A^j) + \beta_j \lambda - \pi_2 = 0$, hence $\lambda(\beta^T \cdot y) - \pi_1 - \pi_2 = 0$.

If $(v, w) \notin \mathcal{N}$, then at least one label $1 \leq r \leq m + n$ is missing from $L(v) \cup L(w)$. Let $r \leq m$ (wlog), then $x_r > 0$ and $A_r \cdot y - \pi_1 < 0$, which imply that $x^T \cdot (A \cdot y - \pi_1) < 0$. Therefore, $\lambda(\beta^T \cdot y) - \pi_1 - \pi_2 < 0$. \square

Motivated by the above lemma, we define the following parametrized linear program $LP(\delta)$.

$$LP(\delta) : \quad \max \quad \delta(\beta^T \cdot y) - \pi_1 - \pi_2 \\ (y, \pi_1) \in P; (x, \lambda, \pi_2) \in Q'; \lambda = \delta \quad (5)$$

Note that the above linear program may be broken into a parametrized primal linear program and its dual, with δ being the parameter. The primal may be defined on polytope P with the cost function *maximize*: $\delta(\beta^T \cdot y) - \pi_1$ and its dual is on polytope Q' with additional constraint $\lambda = \delta$ and the cost function *minimize*: π_2 .

Remark 2. $LP(\delta)$ may look similar to the parametrized linear program, say $TLP(\xi)$, by Theobald [19]. However the key difference is that $TLP(\xi)$ is defined on the best response polytopes of a given game (i.e., $P(\alpha) \times Q(\alpha)$ for the game $G(\alpha)$), while $LP(\delta)$ is defined on a bigger polytope ($P \times Q'$) encompassing best response polytopes of all the games in Γ . A detailed comparison is given in Section 4.2.

Let $OPT(\delta)$ be the set of optimal points of $LP(\delta)$. In the next lemma, we show that $\forall \delta \in \mathbb{R}$, $OPT(\delta)$ is exactly the set of points in \mathcal{N} , where $\lambda = \delta$.

LEMMA 10. $\forall a \in \mathbb{R}$, $OPT(a) = \{((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{N} \mid \lambda = a\}$ and $OPT(a) \neq \emptyset$.

PROOF. Clearly the feasible set of $LP(a)$ consists of all the points of $P \times Q'$, where $\lambda = a$. Therefore the set $\{((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{N} \mid \lambda = a\}$ is a subset of the feasible set of $LP(a)$. The set $\{((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{N} \mid \lambda = a\}$ is non-empty (Lemma 6). From Lemma 9, it is clear that the maximum possible value, the cost function of $LP(a)$ may achieve is 0, and it is achieved only at the points of \mathcal{N} . Therefore, $OPT(a) = \{((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{N} \mid \lambda = a\}$ and $OPT(a) \neq \emptyset$. \square

Next we show that \mathcal{N} in fact consists of only one component.

PROPOSITION 3. \mathcal{N} does not contain cycles.

PROOF. From Proposition 1, it is clear that there is always a path \mathcal{P} in the set \mathcal{N} . Moreover, for every $a \in \mathbb{R}$, there exists a point $((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{P}$ with $\lambda = a$ (Lemma 6). Further, $OPT(a)$ is connected, since it is the solution set of $LP(a)$. Therefore, $\forall a \in \mathbb{R}$, $OPT(a)$ is contained in the path \mathcal{P} . Therefore \mathcal{N} consists of only the path. \square

From Proposition 3, it is clear that \mathcal{N} consists of only the path \mathcal{P} , henceforth we refer to \mathcal{N} as a path. To construct homeomorphism maps between E_Γ and Γ , we need to encode a point $(\alpha, x, y) \in E_\Gamma$ (of size $2m + n$) into a vector $\alpha' \in \Gamma$ (of size m), such that α' uniquely identifies the point (α, x, y) (i.e., a bijection). First we show that there is a bijection between \mathcal{N} and \mathbb{R} and using this, we derive a bijection between Γ and E_Γ . Consider the function $g : \mathcal{N} \rightarrow \mathbb{R}$ such that

$$g((y, \pi_1), (x, \lambda, \pi_2)) = \beta^T \cdot y + \lambda \quad (6)$$

LEMMA 11. *Each term of g , namely $\beta^T \cdot y$ and λ , monotonically increases on the directed path \mathcal{N} , and the function g strictly increases on it.*

PROOF. From the proof of Proposition 1, we know that the edges of type (v, \mathcal{E}_v) (where $v \in \mathcal{N}^P$ is a vertex) and of type (\mathcal{E}_w, w) (where $w \in \mathcal{N}^{Q'}$ is a vertex) alternate in \mathcal{N} . Clearly $\beta^T \cdot y$ is a constant on an edge of type (v, \mathcal{E}_v) and λ is a constant on an edge of type (\mathcal{E}_w, w) . Now, consider the two consecutive edges (\mathcal{E}_w, w) and (v, \mathcal{E}_v) , where $\mathcal{E}_w = w', v$ and $\mathcal{E}_v = w, w'$. It is enough to show that λ and $\beta^T \cdot y$ are not constants on (v, \mathcal{E}_v) and (\mathcal{E}_w, w) respectively, and $\beta^T \cdot y$ increases from (v', w) to (v, w) (i.e., on (\mathcal{E}_w, w)) iff λ also increases from (v, w) to (v, w') (i.e., on (v, \mathcal{E}_v)).

Let $w = (x, \gamma, \pi_2)$, $w' = (x', \gamma', \pi_2')$, $v = (y, \pi_1)$ and $v' = (y', \pi_1')$. Clearly, $OPT(\gamma) = (\mathcal{E}_w, w)$ and $(v, w') \in OPT(\gamma')$ (Lemma 10). Further $\gamma \neq \gamma'$, since $OPT(\gamma)$ contains only one edge.

Claim. $\beta^T \cdot y' \neq \beta^T \cdot y$, and $\beta^T \cdot y' < \beta^T \cdot y \Leftrightarrow \gamma < \gamma'$.

PROOF. Since the feasible set of $LP(\gamma')$ contains all the points of $P \times Q'$ with $\lambda = \gamma'$, the point (v', w') is a feasible point of $LP(\gamma')$. Note that (v', w') is a suboptimal point of $LP(\gamma')$ otherwise $\mathcal{E}_{w'} = v, v'$ and $\mathcal{E}_{v'} = w', w$, which creates a cycle in \mathcal{N} . Further, $(v, w') \in OPT(\gamma')$, hence $\gamma'(\beta^T \cdot y) - \pi_1 - \pi_2' > \gamma'(\beta^T \cdot y') - \pi_1' - \pi_2'$. Since both (v', w) and (v, w) are in $OPT(\gamma)$, we get $\gamma(\beta^T \cdot y') - \pi_1' - \pi_2 = \gamma(\beta^T \cdot y) - \pi_1 - \pi_2$. Summing up these two, we get $\gamma(\beta^T \cdot y') + \gamma'(\beta^T \cdot y) > \gamma(\beta^T \cdot y) + \gamma'(\beta^T \cdot y') \Rightarrow (\beta^T \cdot y - \beta^T \cdot y')(\gamma' - \gamma) > 0$. \square

The above claim shows that $\beta^T \cdot y$ is strictly monotonic on (\mathcal{E}_w, w) and λ is strictly monotonic on (v, \mathcal{E}_v) . Further, if $\beta^T \cdot y$ increases on (\mathcal{E}_w, w) from (v', w) to (v, w) then λ increases on (v, \mathcal{E}_v) from (v, w) to (v, w') and vice-versa.

Recall that on the directed path \mathcal{N} , (v_s, \mathcal{E}_{v_s}) is the first edge and (v_e, \mathcal{E}_{v_e}) is the last edge (Lemma 8). Further, λ varies from $-\infty$ to λ_s on the first edge (v_s, \mathcal{E}_{v_s}) , and it varies from λ_e to ∞ on the last edge (v_e, \mathcal{E}_{v_e}) (proof of Lemma 5). Therefore, λ and $\beta^T \cdot y$ increase monotonically on the directed path \mathcal{N} , and in turn g strictly increases from $-\infty$ to ∞ on the path. \square

Lemma 11 implies that g is a continuous, bijective function with a continuous inverse $g^{-1} : \mathbb{R} \rightarrow \mathcal{N}$. Now consider the following candidate function $f : E_\Gamma \rightarrow \Gamma$ for the homeomorphism map.

$$f(\alpha, x, y) = (\beta^T \cdot y + \alpha^T \cdot x, \alpha_2 - \alpha_1, \dots, \alpha_m - \alpha_1)^T \quad (7)$$

Using the properties of g , next we show that f indeed establishes a homeomorphism between Γ and E_Γ .

THEOREM 1. E_Γ is homeomorphic to Γ .

PROOF. The function f of (7) is continuous because it is a quadratic function. To show the homeomorphism we need to show that it has a continuous inverse. Define function $f^{-1} : \Gamma \rightarrow E_\Gamma$ using g of (6) as follows. Given $\alpha' \in \Gamma$, let $(v, w) = ((y, \pi_1), (x, \lambda, \pi_2)) = g^{-1}(\alpha')$ be the corresponding point in \mathcal{N} . This gives the values of x, y and λ . Using these values, we solve the following equalities with the variable vector $\mathbf{a} = (a_1, \dots, a_m)$.

$$\forall i > 1, a_i = \alpha'_i + a_1; \quad \sum_{i=1}^m x_i a_i = \alpha'_1 - \beta^T \cdot y$$

It is easy to see that the above equations have a unique solution, which gives a unique value for the vector \mathbf{a} and a unique point $(\mathbf{a}, x, y) \in E_\Gamma$. Let, $f^{-1}(\alpha') = (\mathbf{a}, x, y)$, then clearly $f^{-1} \circ f$ and $f \circ f^{-1}$ are identity maps.

Clearly, f^{-1} is continuous as well. Hence, f and f^{-1} establishes homeomorphism between Γ and E_Γ . \square

4. ALGORITHMS

In this section, we present two algorithms to find Nash equilibria of a rank-1 game using the structure and monotonicity of \mathcal{N} . First we discuss a polynomial time algorithm to find a Nash equilibrium of a non-degenerate rank-1 game and later extend it for degenerate games. It does a binary search on \mathcal{N} using the monotonicity of λ . Later we give a path-following algorithm which enumerates all Nash equilibria of a rank-1 game, and finds at least one for any bimatrix game (Lemma 7).

Recall that the best response polytopes P and Q (of (2)) of a non-degenerate game are non-degenerate, and hence it's Nash equilibria set is finite. Consider a non-degenerate rank-1 bimatrix game $(A, B) \in \mathbb{R}^{2mn}$ such that $A + B = \gamma \cdot \beta^T$, where $\gamma \in \mathbb{R}^m$ and $\beta \in \mathbb{R}^n$. We assume that β is a non-zero and non-constant² vector, and both A and B are rational matrices. Let c be the LCM of the denominators of the a_{ij} s, β 's and γ 's. Note that multiplying both A and B by c^2 makes A, γ and β integers, and the total bit length of the input gets multiplied by at most $O(m^2 n^2)$, which is a polynomial increase. Since scaling both the matrices of a bimatrix game by a positive integer does not change the set of Nash equilibria, we assume that entries of A, γ and β are integers.

Now consider the game space $\Gamma = \{(A, -A + \alpha \cdot \beta^T) \mid \alpha \in \mathbb{R}^m\}$. Clearly, $G(\gamma) = (A, B) \in \Gamma$ and the corresponding polytopes P and Q' of (3) are non-degenerate. Let \mathcal{N} be the set of fully-labeled points of $P \times Q'$ as defined in Section 2.2. By Lemma 2, we know that for every Nash equilibrium of the game $G(\gamma)$, there is a unique point in \mathcal{N} .

²If β is a constant vector, then the game (A, B) may be converted into a zero-sum game without changing it's Nash equilibrium set, by adding constants in the columns and rows of A and B respectively.

Consider the hyper-plane $H : \lambda - \sum_{i=1}^m \gamma_i x_i = 0$ in $(y, \pi_1, x, \lambda, \pi_2)$ -space and the corresponding half spaces $H^+ : \lambda - \sum_{i=1}^m \gamma_i x_i \geq 0$ and $H^- : \lambda - \sum_{i=1}^m \gamma_i x_i \leq 0$. It is easy to see that the intersection of \mathcal{N} with the hyper-plane H gives all the Nash equilibria of $G(\gamma)$. If the hyper-plane H intersects an edge of \mathcal{N} , then it intersects the edge exactly at one point, because $G(\gamma)$ is a non-degenerate game.

Let $\gamma_{min} = \min_{i \in S_1} \gamma_i$ and $\gamma_{max} = \max_{i \in S_1} \gamma_i$. Since $\forall x \in \Delta_1, \gamma_{min} \leq \sum_{i=1}^m \gamma_i x_i \leq \gamma_{max}$, a point $w \in \mathcal{N}$ corresponds to a Nash equilibrium of $G(\gamma)$, only if the value of λ at w is between γ_{min} and γ_{max} . From Proposition 3, we know that \mathcal{N} contains only a path. If we consider the path \mathcal{N} from the first edge (v_s, \mathcal{E}_{v_s}) to the last edge (v_e, \mathcal{E}_{v_e}) , then λ monotonically increases from $-\infty$ to ∞ on it (Lemmas 5 and 11). Therefore all the points, corresponding to the Nash equilibrium of $G(\gamma)$ on the path \mathcal{N} , lie between $OPT(\gamma_{min})$ and $OPT(\gamma_{max})$ (Lemma 10).

4.1 Rank-1 NE: Polynomial Time Algorithm

Recall that finding a Nash equilibrium of the game $G(\gamma)$ is equivalent to finding a point in the intersection of \mathcal{N} and the hyper-plane H . Since λ increases monotonically on \mathcal{N} , and all the points in the intersection are between $\lambda = \gamma_{min}$ and $\lambda = \gamma_{max}$, the *BinSearch* algorithm of Table 1 applies binary search on \mathcal{N} between $\lambda = \gamma_{min}$ and $\lambda = \gamma_{max}$ to locate a point in the intersection.

```

BinSearch( $\gamma_{min}, \gamma_{max}$ )
   $a_1 \leftarrow \gamma_{min}; a_2 \leftarrow \gamma_{max};$ 
  if  $\text{IsNE}(a_1) = 0$  or  $\text{IsNE}(a_2) = 0$  then return;
  while true
     $a \leftarrow \frac{a_1 + a_2}{2}; \text{flag} \leftarrow \text{IsNE}(a);$ 
    if  $\text{flag} = 0$  then break;
    else if  $\text{flag} < 0$  then  $a_1 \leftarrow a;$ 
    else  $a_2 \leftarrow a;$ 
  endwhile
  return;

IsNE( $\delta$ )
  Find  $OPT(\delta)$  by solving  $LP(\delta)$ ;
   $\bar{u}, \bar{v} \leftarrow$  The edge containing  $OPT(\delta)$ ;
   $\mathcal{H} \leftarrow \{w \in \bar{u}, \bar{v} \mid w \in H\};$ 
  if  $\mathcal{H} \neq \emptyset$  then Output  $\mathcal{H}$ ; return 0;
  else if  $\bar{u}, \bar{v} \in H^+$  then return 1;
  else return -1;

```

Table 1: BinSearch Algorithm

The *IsNE* procedure of Table 1 takes a $\delta \in \mathbb{R}$ as the input, and outputs a NESP if possible, otherwise it indicates the position of $OPT(\delta)$ with respect to the hyper-plane H . First it finds the optimal set $OPT(\delta)$ of $LP(\delta)$ and the corresponding edge \bar{u}, \bar{v} containing $OPT(\delta)$. Next, it finds a set \mathcal{H} , which consists of all the points in the intersection of \bar{u}, \bar{v} and the hyper-plane H if any, i.e., Nash equilibria of $G(\gamma)$. Since the game $G(\gamma)$ is non-degenerate, \mathcal{H} is either a singleton or empty. In the former case, the procedure outputs \mathcal{H} and returns 0 indicating that a Nash equilibrium has been found. However in the latter case, it returns 1 if $\bar{u}, \bar{v} \in H^+$ and returns -1 otherwise, indicating the position of $OPT(\delta)$ w.r.t. the hyper-plane H .

The *BinSearch* algorithm maintains two pivot values a_1

and a_2 of λ such that the corresponding $OPT(a_1) \in H^-$ and $OPT(a_2) \in H^+$, i.e., always on the opposite sides of the hyper-plane H . Clearly \mathcal{N} crosses H at least once between $OPT(a_1)$ and $OPT(a_2)$. Since $OPT(\gamma_{min}) \in H^-$ and $OPT(\gamma_{max}) \in H^+$, the pivots a_1 and a_2 are initialized to γ_{min} and γ_{max} respectively. Initially it calls *IsNE* for both a_1 and a_2 separately and terminates if either returns zero indicating that a NESP has been found. Otherwise the algorithm repeats the following steps until *IsNE* returns zero: It calls *IsNE* for the mid point a of a_1 and a_2 and terminates if it returns zero. If *IsNE* returns a negative value, then $OPT(a) \in H^-$ implying that $OPT(a)$ and $OPT(a_2)$ are on the opposite sides of H , and hence the lower pivot a_1 is reset to a . In the other case $OPT(a) \in H^+$, the upper pivot a_2 is set to a , as $OPT(a_1)$ and $OPT(a)$ are on the opposite sides of H .

Note that, the index of the Nash equilibrium obtained by *BinSearch* algorithm is always +1, since $a_1 < a_2$ is an invariant (Proposition 2). For $X \in \mathbb{R}^{kl}$, let $\tilde{X} = \max_{i,j} |x_{ij}|$.

Since the column-player's payoff matrix is represented by $-A + \gamma \cdot \beta^T$ of the game $G(\gamma)$, let $|B| = \max\{\tilde{A}, \tilde{\beta}, \tilde{\gamma}\}$. Let $\Delta = (m+2)! (|B|)^{(m+2)}$.

THEOREM 2. *Let \mathcal{L} be the bit length of the input. The *BinSearch* terminates in time $\text{poly}(\mathcal{L}, m, n)$.*

PROOF. Clearly, the algorithm terminates when the call *IsNE*(a) outputs a NESP of $G(\gamma)$. Let the range of λ for an edge $(v, \mathcal{E}_v) \in \mathcal{N}$ be $[\lambda_1 \ \lambda_2]$.

Claim. $\lambda_2 - \lambda_1 \geq \frac{1}{\Delta^2}$.

PROOF. Note that λ_1 and λ_2 correspond to the two vertices of $\mathcal{E}_v \in Q'$. Since Q' is in a $(m+2)$ -dimensional space, there are $m+2$ equations tight at every vertex of it. Hence both λ_1 and λ_2 are rational numbers with denominator at most Δ . Therefore $\lambda_2 - \lambda_1$ is at least $\frac{1}{\Delta^2}$. \square

When $a_2 - a_1 \leq \frac{1}{\Delta^2}$, $OPT(a_1)$ and $OPT(a_2)$ are either part of the same edge or adjacent edges. In either case, the algorithm terminates after one more call to *IsNE*.

Clearly $a_2 - a_1 = \frac{\gamma_{max} - \gamma_{min}}{2^l}$ after l iterations of the *while* loop. Let k be such that

$$\frac{\gamma_{max} - \gamma_{min}}{2^k} = \frac{1}{\Delta^2} \Rightarrow k = 2 \log \Delta + \log(\gamma_{max} - \gamma_{min})$$

BinSearch makes at most $\lceil k \rceil$ calls to the procedure *IsNE*, which is polynomial in \mathcal{L}, m , and n . The procedure *IsNE* solves a linear program and computes a set \mathcal{H} , both may be done in $\text{poly}(\mathcal{L}, m, n)$ time. Therefore the total time taken by *BinSearch* is polynomial in \mathcal{L}, m , and n . \square

Degeneracy. For a degenerate rank-1 game (A, B) the corresponding polytope $P \times Q'$ may be degenerate as well. However, the *BinSearch* algorithm, with a small modification, works with the same polynomial time bound. First we make a few observations and then state a minor modification to the algorithm.

Note that Lemma 10 (i.e. $\forall a \in \mathbb{R}, OPT(a) = \mathcal{N}(a)$) holds in case of degeneracy as well. Therefore, the set of fully-labeled points $\mathcal{N} \in P \times Q'$ is a connected set, and λ varies continuously on this set.

LEMMA 12. *Let $f \in \mathcal{N}^{Q'}$ be a maximal face, then λ does not take a unique value on f .*

PROOF. Suppose λ takes a unique value a on f . Let $\mathcal{S} = \cap_{w \in f} \mathcal{E}_w$ (note that \mathcal{S} is a vertex of P). Clearly, there exists a vertex $v \in f$ such that $\mathcal{S} \subset \mathcal{E}_v$. In that case, $\{v \times \mathcal{E}_v\} \cup \{f \times \mathcal{S}\} \in OPT(a)$, which makes $OPT(a)$ non-convex. \square

It is clear from the above lemma that on every alternate maximal face of \mathcal{N} (i.e., similar to (v, \mathcal{E}_v) edges in non-degenerate case), increase in λ is lower bounded by $\frac{1}{\Delta^2}$. It is easy to check that *BinSearch* algorithm with the *IsNE* procedure given in Table 2 works for degenerate case as well.

IsNE (δ)
Find $OPT(\delta)$ by solving $LP(\delta)$;
$f \leftarrow$ The minimal face of $P \times Q'$ containing $OPT(\delta)$;
(Note that f is a maximal face of \mathcal{N})
$\mathcal{H} \leftarrow \{w \in f \mid w \in H\}$;
if $\mathcal{H} \neq \emptyset$ then Output \mathcal{H} ; return 0;
else if $f \in H^+$ then return 1;
else return -1;

Table 2: IsNE Procedure for the Degenerate Case

4.2 Enumeration Algorithm for Rank-1 Games

The *Enumeration* algorithm of Table 3 simply follows the path \mathcal{N} between $OPT(\gamma_{min})$ and $OPT(\gamma_{max})$, and outputs the NESPs whenever it hits the hyper-plane $H : \lambda - \sum_{i=1}^m \gamma_i x_i = 0$.

Enumeration ($\overline{u_1, v_1}, \overline{u_2, v_2}$)
$\overline{u}, \overline{u'} \leftarrow \overline{u_1, v_1}$;
if $\overline{u}, \overline{u'}$ of type (v, \mathcal{E}_v) then $\text{flag} \leftarrow 1$;
else $\text{flag} \leftarrow 0$;
while <i>true</i>
$\mathcal{H} = \{w \in \overline{u}, \overline{u'} \mid w \in H\}$; Output \mathcal{H} ;
if $\overline{u}, \overline{u'} = \overline{u_2, v_2}$ then break ;
if $\text{flag} = 1$ then $\overline{u}, \overline{u'} \leftarrow (\mathcal{E}_{u'}, u')$; $\text{flag} \leftarrow 0$;
else $\overline{u}, \overline{u'} \leftarrow (u', \mathcal{E}_{u'})$; $\text{flag} \leftarrow 1$;
endwhile
return ;

Table 3: Enumeration Algorithm

Let the edges $\overline{u_1, v_1}$ and $\overline{u_2, v_2}$ contain $OPT(\gamma_{min})$ and $OPT(\gamma_{max})$ respectively. The call *Enumeration*($\overline{u_1, v_1}, \overline{u_2, v_2}$) enumerates all the Nash equilibria of the game $G(\gamma)$.

The *Enumeration* algorithm initializes $\overline{u}, \overline{u'}$ to the edge $\overline{u_1, v_1}$. Since the edges alternate between the type (v, \mathcal{E}_v) and (\mathcal{E}_w, w) on \mathcal{N} , the value of *flag* indicates the type of edge to be considered next. It is set to one if the next edge is of type (\mathcal{E}_w, w) , otherwise it is set to zero. In the *while* loop, it first outputs the intersection of the edge $\overline{u}, \overline{u'}$ with the hyper-plane H , if any. Further, $\overline{u}, \overline{u'}$ is set to the next edge and the *flag* is toggled. Recall that the edges incident on a vertex u' in \mathcal{N} may be obtained by relaxing the inequality corresponding to the duplicate label of u' , in P and in Q' (Section 2.2). Let the duplicate label of the vertex u' be i . We may obtain the edge $(\mathcal{E}_{u'}, u')$ by relaxing the inequality i of P and the edge $(u', \mathcal{E}_{u'})$ by relaxing the inequality i of Q' . The algorithm terminates when $\overline{u}, \overline{u'} = \overline{u_2, v_2}$.

Every iteration of the loop takes time polynomial in \mathcal{L}, m and n . Therefore, the time taken by the algorithm is equivalent to the number of edges between $\overline{u_1, v_1}$ and $\overline{u_2, v_2}$.

For a general (non-degenerate) bimatrix game (A, B) , we may obtain C , γ and β such that $B = C + \gamma \cdot \beta^T$, and define the corresponding game space Γ and the polytopes P and Q' accordingly (Section 2.2). There is a one-to-one correspondence between the Nash equilibria of the game (A, B) and the points in the intersection of the fully-labeled set \mathcal{N} and the hyper-plane $\lambda - \sum_{i=1}^m \gamma_i x_i = 0$. Recall that the set \mathcal{N} contains one path (\mathcal{P}) and a set of cycles (Proposition 1). The extreme edges (v_s, \mathcal{E}_{v_s}) and (v_e, \mathcal{E}_{v_e}) of \mathcal{P} may be easily obtained as described in the proof of Lemma 5. Since \mathcal{P} contains at least one Nash equilibrium of every game in Γ (Lemma 7), hence the call $Enumerate((v_s, \mathcal{E}_{v_s}), (v_e, \mathcal{E}_{v_e}))$ outputs at least one Nash equilibrium of the game (A, B) . Note that the time taken by the algorithm again depends on the number of edges on the path \mathcal{P} .

Comparison with Earlier Approaches. The *Enumeration* algorithm may be compared to two previous algorithms. One is the Theobald algorithm [19], which enumerates all Nash equilibria of a rank-1 game, and the other is the Lemke-Howson algorithm [11], which finds a Nash equilibrium of any bimatrix game. The *Enumeration* algorithm enumerates all the Nash equilibria of a rank-1 game and for any general bimatrix game it is guaranteed to find one Nash equilibrium. All three algorithms are path following algorithms. However, the main difference is that both the previous algorithms always trace a path on the best response polytopes of a given game (i.e., $P(\gamma) \times Q(\gamma)$), while the *Enumeration* algorithm follows a path on a bigger polytope $P \times Q'$ which encompasses best response polytopes of all the games of an m -dimensional game space. Therefore, for every game in this m -dimensional game space, the *Enumeration* follows the same path. Further, all the points on the path followed by *Enumeration* algorithm are fully-labeled, and it always hits the best response polytope of the given game at one of its NESP points. However the path followed by previous two algorithms is not fully-labeled and whenever they hit a fully-labeled point, it is a NESP of the game.

In every intermediate step, the Theobald algorithm calculates the range of a variable (ξ) based on the feasibility of primal and dual, and accordingly decides which inequality to relax (in P or Q) to locate the next edge. While *Enumeration* algorithm simply leaves the duplicate label in P or Q' (alternately) at the current vertex to locate the next edge. Further, for a general bimatrix game, the *Enumeration* algorithm locates at least one Nash equilibrium, while Theobald algorithm works only for rank-1 games.

5. RANK- K SPACE: HOMEOMORPHISM

It turns out that the approach used to show the homeomorphism between the subspace of rank-1 games and its Nash equilibrium correspondence may be extended to the subspace with rank- k games. Given a bimatrix game $(A, B) \in \mathbb{R}^{2mn}$ of rank- k , the matrix $A+B$ may be written as $\sum_{l=1}^k \gamma^l \cdot \beta^{lT}$, using the linearly independent vectors $\gamma^l \in \mathbb{R}^m$, $\beta^l \in \mathbb{R}^n$, $1 \leq l \leq k$. Therefore, the column-player's payoff matrix B may be written as $B = -A + \sum_{l=1}^k \gamma^l \cdot \beta^{lT}$, where $\{\beta^l\}_{l=1}^k$ are linearly independent. Consider the corresponding game space $\Gamma^k = \{(A, -A + \sum_{l=1}^k \alpha^l \cdot \beta^{lT}) \in \mathbb{R}^{2mn} \mid \forall l \leq k, \alpha^l \in \mathbb{R}^m\}$. This space is an affine km -dimensional subspace of the bimatrix game space \mathbb{R}^{2mn} , and it contains only rank- k games. Let $\alpha = (\alpha^1, \dots, \alpha^k)$, and $G(\alpha)$ de-

note the game $(A, -A + \sum_{l=1}^k \alpha^l \cdot \beta^{lT})$. The Nash equilibrium correspondence of the space Γ^k is $E_{\Gamma^k} = \{(\alpha, x, y) \in \mathbb{R}^{km} \times \Delta_1 \times \Delta_2 \mid (x, y) \text{ is a NESP of } G(\alpha) \in \Gamma^k\}$.

For all the games in Γ^k , again the row-player's payoff matrix remains constant, hence for all $G(\alpha) \in \Gamma^k$ the best response polytope of the row-player $P(\alpha)$ is P of (2). However, the best response polytope of the column player $Q(\alpha)$ varies, as the payoff matrix of the column-player varies with α . Consider the following polytope (similar to (3)).

$$Q'^k = \{(x, \lambda, \pi_2) \in \mathbb{R}^{m+k+1} \mid x_i \geq 0, \forall i \in S_1; x^T(-A^j) + \sum_{l=1}^k \beta_j^l \lambda_l - \pi_2 \leq 0, \forall j \in S_2; \sum_{i=1}^m x_i = 1\} \quad (8)$$

Note that $\lambda = (\lambda_1, \dots, \lambda_k)$ is a variable vector. The column-player's best response polytope $Q(\alpha)$, for the game $G(\alpha)$, is the projection of the set $\{(x, \lambda, \pi_2) \in Q'^k \mid \forall l \leq k, \sum_{i=1}^m \alpha_i^l x_i - \lambda_l = 0\}$ on (x, π_2) -space. We assume that the polytopes P and Q'^k are non-degenerate. Let the set of fully-labeled pairs of $P \times Q'^k$ be $\mathcal{N}^k = \{(v, w) \in P \times Q'^k \mid L(v) \cup L(w) = \{1, \dots, m+n\}\}$. The following facts regarding the set \mathcal{N}^k may be easily derived.

- For every point in E_{Γ^k} there is a unique point in \mathcal{N}^k , and for every point in \mathcal{N}^k there is a point in E_{Γ^k} (Lemma 2). Further the set of points of E_{Γ^k} mapping to a point $(v, w) \in \mathcal{N}^k$, is equivalent to $k(m-1)$ -dimensional space.
- Since there are k more variables in Q'^k , namely $\lambda_1, \dots, \lambda_k$ compared to Q of (2), \mathcal{N}^k is a subset of the k -skeleton of $P \times Q'^k$. If a point $v \in P$ is on a d -dimensional face ($d \leq k$), then the set \mathcal{E}_v is either empty or it is a $(k-d)$ -dimensional face, where $\mathcal{E}_v = \{w \in Q'^k \mid (v, w) \in \mathcal{N}^k\}$ (Observations of Section 2.2).
- $\forall (v, w) = ((y, \pi_1), (x, \lambda, \pi_2))$ in $P \times Q'^k$, $\sum_{l=1}^k \lambda_l (\beta^{lT} \cdot y) - \pi_1 - \pi_2 \leq 0$, and equality holds iff $(v, w) \in \mathcal{N}^k$.

For a vector $\delta \in \mathbb{R}^k$, consider the following parametrized linear program $LP^k(\delta)$.

$$LP^k(\delta) : \max \sum_{l=1}^k \delta_l (\beta^{lT} \cdot y) - \pi_1 - \pi_2 \quad (y, \pi_1) \in P; (x, \lambda, \pi_2) \in Q'^k; \lambda_l = \delta_l, \forall l \leq k \quad (9)$$

Let $OPT^k(\delta)$ be the set of optimal points of $LP^k(\delta)$. Note that for any $a \in \mathbb{R}^k$, all the points on \mathcal{N}^k with $\lambda = a$ may be obtained by solving $LP^k(a)$. In other words, $\{((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{N}^k \mid \lambda = a\} = OPT^k(a)$ (Lemma 10). Using this fact we show that the tuple $(\lambda_1 + \beta^{1T} \cdot y, \dots, \lambda_k + \beta^{kT} \cdot y)$ uniquely identifies a point of \mathcal{N}^k . For a vector $a \in \mathbb{R}^k$, let $S(a) = \{((y, \pi_1), (x, \lambda, \pi_2)) \in \mathcal{N}^k \mid \forall l \leq k, \lambda_l + \beta^{lT} \cdot y = a_l\}$.

LEMMA 13. For a vector $a \in \mathbb{R}^k$, the set $S(a)$ contains exactly one element, i.e., $|S(a)| = 1$.

Consider the function $g^k : \mathcal{N}^k \rightarrow \mathbb{R}^k$ such that,

$$g^k((y, \pi_1), (x, \lambda, \pi_2)) = (\lambda_1 + (\beta^{1T} \cdot y), \dots, \lambda_k + (\beta^{kT} \cdot y))$$

The function g^k is continuous and bijective (Lemma 13), and the inverse $g^{k^{-1}} : \mathbb{R}^k \rightarrow \mathcal{N}^k$ is also continuous, since \mathcal{N}^k is a closed and connected set. Now consider a function $f^k : E_{\Gamma^k} \rightarrow \Gamma^k$ similar to (7).

$$f^k(\alpha, x, y) = (\alpha^1, \dots, \alpha^k), \text{ where } \alpha^l = ((\alpha^{lT} \cdot x) + (\beta^{lT} \cdot y), \alpha_2^l - \alpha_1^l, \dots, \alpha_m^l - \alpha_1^l)^T, \forall l \leq k$$

It is easy to construct the continuous inverse of f^k using g^{k-1} and establish the homeomorphism between Γ^k and E_{Γ^k} .

THEOREM 3. *Function f^k establishes a homeomorphism between NE correspondence E_{Γ^k} and the game space Γ^k .*

Next we give a fixed point formulation to find a Nash equilibrium of a rank- k game $G(\gamma) \in \Gamma^k$. Let $\gamma_{\min} = (\gamma_{\min}^1, \dots, \gamma_{\min}^k)$ and $\gamma_{\max} = (\gamma_{\max}^1, \dots, \gamma_{\max}^k)$, where $\forall l \leq k$, $\gamma_{\min}^l = \min_{i \in S_1} \gamma_i^l$ and $\gamma_{\max}^l = \max_{i \in S_1} \gamma_i^l$. Consider the box $\mathcal{B} \in \mathbb{R}^k$ s.t. $\mathcal{B} = \{a \in \mathbb{R}^k \mid \gamma_{\min} \leq a \leq \gamma_{\max}\}^3$. For the rank-1 case, \mathcal{B} is an interval.

LEMMA 14. *Finding a Nash equilibrium of $G(\gamma)$ reduces to finding a fixed point of a polynomially computable piece-wise linear function $f : \mathcal{B} \rightarrow \mathcal{B}$.*

PROOF. Clearly, $\forall x \in \Delta_1$, $(\sum_{i=1}^m \gamma_i^1 x_i, \dots, \sum_{i=1}^m \gamma_i^k x_i) \in \mathcal{B}$. Now, consider the function $f : \mathcal{B} \rightarrow \mathcal{B}$ such that,

$$f(a) = (\sum_{i=1}^m \gamma_i^1 x_i, \dots, \sum_{i=1}^m \gamma_i^k x_i), \text{ where } (x, \lambda, \pi_2) = \{w \in Q^k \mid (v, w) \in OPT^k(a), v \in P\}$$

The function f is a piece-wise linear function. For every $a \in \mathcal{B}$, the corresponding x is well defined in the above expression (Lemma 13), and may be obtained in polynomial time by solving $LP^k(a)$. It is easy to see that fixed points of f correspond to the Nash equilibria of game $G(\gamma)$ and vice-versa. \square

It seems that for a given $a \in \mathbb{R}^k$, there is a way to trace the points in the intersection of \mathcal{N}^k and $\lambda_l = a_l, l \neq i$, such that λ_i increases monotonically (analysis similar to Lemma 11). Is there a way to locate a fixed point of f in polynomial time using this observation and the simple structure of \mathcal{N}^k , even though finding a fixed point in general is PPAD-complete [14]?

6. CONCLUSION

In this paper, we establish a homeomorphism between an m -dimensional affine subspace Γ of the bimatrix game space and its Nash equilibrium correspondence E_{Γ} , where Γ contains only rank-1 games. To the best of our knowledge, this is the first structural result for a subspace of the bimatrix game space. The homeomorphism maps that we derive are very different than the ones given by Kohlberg and Mertens for the bimatrix game space [9] and it builds on the structure of E_{Γ} . Further, using this structural result we design two algorithms. The first algorithm finds a Nash equilibrium of a rank-1 game in polynomial time. This settles an open question posed by Kannan and Theobald [8] and Theobald [19]. The second algorithm enumerates all the Nash equilibria of a rank-1 game and finds at least one Nash equilibrium of a general bimatrix game.

Further, we extend the above structural result by establishing a homeomorphism between km -dimensional affine subspace Γ^k and its Nash equilibrium correspondence E_{Γ^k} , where Γ^k contains only rank- k games. We hope that this homeomorphism result will help in designing a polynomial time algorithm to find a Nash equilibrium of a fixed rank game.

³For any two vectors $a, b \in \mathbb{R}^n$, by $a \leq b$ we mean $a_i \leq b_i, \forall i \leq n$.

7. REFERENCES

- [1] J. Bulow and J. Levin. Matching and price competition. *American Economic Review*, 96:652–668.
- [2] X. Chen and X. Deng. Settling the complexity of two-player nash equilibrium. In *FOCS*, pages 261–272, 2006.
- [3] X. Chen, X. Deng, and S.-H. Teng. Computing nash equilibria: Approximation and smoothed complexity. In *FOCS*, pages 603–612, 2006.
- [4] G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, 1963.
- [5] S. Govindan and R. Wilson. Equivalence and invariance of the index and degree of nash equilibria. *Games and Economic Behavior*, 21(1):56–61, 1997.
- [6] S. Govindan and R. Wilson. A global newton method to compute nash equilibria. *Journal of Economic Theory*, 110(1):65–86, 2003.
- [7] P. J.-J. Herings and R. Peeters. Homotopy methods to compute equilibria in game theory. *Econ Theory*, 42:119–156, 2010.
- [8] R. Kannan and T. Theobald. Games of fixed rank: A hierarchy of bimatrix games. In *SODA*, pages 1124–1132, 2007.
- [9] E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54(5):1003–1037, 1986.
- [10] S. Kontogiannis and P. Spirakis. Exploiting concavity in bimatrix games: new polynomially tractable subclasses. In *APPROX*, pages 312–325, 2010.
- [11] C. E. Lemke and J. J. T. Howson. Equilibrium points of bimatrix games. *SIAM J. on Applied Mathematics*, 12(2):413–423, 1964.
- [12] R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *EC*, pages 36–41, 2003.
- [13] J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):289–295, September 1951.
- [14] C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *JCSS*, 48(3):498–532, 1992.
- [15] C. H. Papadimitriou. Algorithms, games and the internet. In *STOC*, pages 749–753, 2001.
- [16] A. Predtetchinski. A general structure theorem for the nash equilibrium correspondence. *Games and Economic Behavior*, 66(2):950–958, 2009.
- [17] R. Savani and B. von Stengel. Hard-to-solve bimatrix games. *Econometrica*, 74(2):397–429, 2006.
- [18] L. S. Shapley. A note on the lemke-howson algorithm. In *Mathematical Programming Studies*, pages 175–189, 1974.
- [19] T. Theobald. Enumerating the nash equilibria of rank-1 games. In *Polyhedral Computation, CRM Proceedings, American Mathematical Society*, 2007.
- [20] A. von Schemde and B. von Stengel. Strategic characterization of the index of an equilibrium. In *SAGT*, pages 242–254, 2008.
- [21] B. von Stengel. Equilibrium computation for two-player games in strategic and extensive form. *Chapter 3, Algorithmic Game Theory*, eds. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, pages 53–78, 2007.