Settling Some Open Problems on 2-Player Symmetric Nash Equilibria

Ruta Mehta   Vijay V. Vazirani   Sadra Yazdanbod

College of Computing, Georgia Tech
rmehta,vazirani,syazdanb@cc.gatech.edu

Abstract. Over the years, researchers have studied the complexity of several decision versions of Nash equilibrium in (symmetric) two-player games (bimatrix games). To the best of our knowledge, the last remaining open problem of this sort is the following; it was stated by Papadimitriou in 2007: find a non-symmetric Nash equilibrium (NE) in a symmetric game. We show that this problem is NP-complete and the problem of counting the number of non-symmetric NE in a symmetric game is #P-complete.

In 2005, Kannan and Theobald defined the *rank of a bimatrix game* represented by matrices (A, B) to be rank$(A + B)$ and asked whether a NE can be computed in rank 1 games in polynomial time. Observe that the rank 0 case is precisely the zero sum case, for which a polynomial time algorithm follows from von Neumann's reduction of such games to linear programming. In 2011, Adsul et. al. obtained an algorithm for rank 1 games; however, it does guarantee symmetric NE in symmetric rank 1 game. We resolve this problem.

1 Introduction

One of the major achievements of complexity theory in recent years is obtaining a fairly complete understanding of the complexity of computing a Nash equilibrium (NE) in a two-player game in various situations; such a game can be represented by two payoff matrices (A, B), and therefore is also known as *bimatrix game*. Of the few remaining open questions, we settle two in this paper regarding symmetric bimatrix games. We note that symmetry arises naturally in numerous strategic situations. In fact, while providing game theory with its central solution concept, Nash [12] felt compelled to also define the notion of a symmetric game and prove, in a separate theorem, that such (finite) games always admit a *symmetric equilibrium*, *i.e.,* where all players play the same strategy. Examples of well-known bimatrix games that are symmetric are Prisoners' Dilemma and Rock-Paper-Scissors. With the growth of the Internet, on which typically users are indistinguishable, the relevance of symmetric games has further increased.

In a *symmetric game* all players participate under identical circumstances, i.e., strategy sets and payoffs. Thus the payoff of a player i depends only on the strategy, s, played by her and the multiset of strategies, S, played by the others, without reference to their identities; moreover, if any other player j were to play

s and the remaining players S, the payoff to j would be identical to that of i. In case of a bimatrix game (A, B) such a symmetry translates to $B = A^T$.

We first provide a brief summary of the known results. The seminal works of Daskalakis, Goldberg and Papadimitriou [6], and Chen, Deng and Teng [2] proved that finding a NE in a bimatrix game, or a symmetric NE in a symmetric bimatrix game, is PPAD-complete [13]. Before the resolution of this long-standing question, researchers studied the complexity of computing a NE with desired special properties. For numerous properties, these problems turned out to be NP-hard[1], even for the case of symmetric games [8,5].

In 2005, Kannan and Theobald [9] defined the *rank of a bimatrix game* (A, B) to be $\text{rank}(A + B)$ and asked whether a NE can be computed in rank 1 games in polynomial time. They also gave an example of a bimatrix rank 1 game that has disconnected NE, thereby providing evidence that the problem would be a difficult one[2]. Observe that the rank 0 case is precisely the zero sum case, for which a polynomial time algorithm follows from von Neumann's reduction of such games to linear programming. In 2011, Adsul et. al [1] answered this question in the affirmative; this appears to be the first efficient algorithm for a problem having disconnected solutions. More recently, Mehta [11] showed that for games of rank 3 or more, and for symmetric games of rank 6 or more, the problem is PPAD-complete.

We now list the open problems we are aware of. In 2007, Papadimitriou [14] asked for the complexity of finding a non-symmetric NE in a symmetric game. One motivation for this problem may be the following. In some situations, symmetric equilibrium may imply both the player exhausting the same resources, for instance if they access the same web site, which may be undesirable and non-symmetric equilibrium is called for.

Mehta [11] left open the problem of determining the complexity of the following problems: finding a NE in a rank 2 game, and finding a symmetric NE in a symmetric game of rank $1, 2, 3, 4$, or 5.

In this paper, we show that the problem of finding asymmetric NE in a symmetric game (Papadimitriou's problem) is NP-complete. We further show that the problem of counting the number of non-symmetric NE in a symmetric game is #P-complete. We also give a polynomial time algorithm for finding a symmetric NE in a symmetric game of rank 1. In Section 5, we give some reasons to believe that finding a symmetric NE in a symmetric game of rank 2 or more should not be in P.

Next, we note that given a symmetric bimatrix rank 1 game, the algorithm of Adsul et. al. [1] is not guaranteed to produce a symmetric Nash equilibrium, as required in the definition of a symmetric game. Furthermore, the symmetric NE of a symmetric rank 1 game can also be disconnected, thereby making it a

---

[1] One exception is the problem of checking if an evolutionarily stable strategy exist in symmetric 2-player game. This problem was recently shown to be $\sum_p^2$-hard by Conitzer [4]; containment in $\sum_p^2$ was shown by Etessami and Lochbihler [7].

[2] von Stengel [16] went further to give a symmetric bimatrix rank 1 game that has exponentially many disconnected symmetric Nash equilibria.

difficult problem from the viewpoint of obtaining a polynomial time algorithm. In Section 1.1 we give the new ideas that are needed, in addition to those of [1], to solve this problem. For an example of a well-known game having disconnected symmetric equilibria, consider Battle of Sexes, with appropriate payoffs so that the game is symmetric.

Recently, McLennan and Tourky [10] gave the notion of imitation games, which simplified the existing proofs of NP-completeness of [5] considerably for the case of symmetric games and led to even more such results. In Appendix B we study further properties of imitation games.

1.1   New techniques

First problem: Next we give an overview of the approach to show NP (#P) hardness for computing (counting) non-symmetric NE in symmetric games. A quick look at the set of Nash equilibrium problems proven NP-hard suggests Non-Unique NE [8], i.e., whether the given bimatrix game has two or more NE, as the most suitable problem to reduce from. Furthermore, an obvious approach is to use the standard reduction from a bimatrix game (A, B) (where both $A > 0$ and $B > 0$ are $m \times n$ matrices) to a symmetric game (M, M^T), where

$$M = \begin{bmatrix} 0 & A \\ B^T & 0 \end{bmatrix}$$

For any non-zero vector $\boldsymbol{z} \geq 0$, let $\eta(z)$ denote the *normalized vector*, i.e., its components are non-negative and add to 1. It is easy to see that the Nash equilibria $(\boldsymbol{a}, \boldsymbol{b})$ of (A, B) are in one-to-one correspondence with the symmetric NE, $(\eta(\boldsymbol{a}/v, \boldsymbol{b}/w), \eta(\boldsymbol{a}/v, \boldsymbol{b}/w))$ of M, where $v = \boldsymbol{a}^T B \boldsymbol{b}$ and $w = \boldsymbol{a}^T A \boldsymbol{b}$. Furthermore, from a non-symmetric NE $((\boldsymbol{a}, \boldsymbol{b}), (\boldsymbol{a}', \boldsymbol{b}'))$ of game (M, M^T), one can obtain two potential equilibria for (A, B), namely $(\eta(\boldsymbol{a}'), \eta(\boldsymbol{b}))$ and $(\eta(\boldsymbol{a}), \eta(\boldsymbol{b}'))$; indeed, one can readily confirm that they satisfy all complementarity conditions. However, there is a snag, namely all four of vectors $\{\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{a}', \boldsymbol{b}'\}$ may not be nonzero, or one set of vectors may be a scaled up version of the other, thereby not yielding 2 NE for the game (A, B). In fact every NE $(\boldsymbol{a}, \boldsymbol{b})$ of (A, B) yields a non-symmetric NE $((\boldsymbol{a}, \boldsymbol{0}), (\boldsymbol{0}, \boldsymbol{b}))$ for the symmetric game (M, M^T), and such a non-symmetric NE yields only one NE for (A, B).

Let us say that a NE has *full support* if both players play all their strategies. We first seek a small dimensional symmetric game which has a unique NE and moreover it has full support. Recently, [3] showed that all symmetric 2×2 games always have pure NE, leading us to consider 3×3 games. As shown in Lemma 1, the game (D, D^T), for the matrix D specified below, has the right properties.

Next, let us define a "blown up" version of matrix D. For every $k \in \mathbb{R}$, let $K_{n \times m}(k)$ be a matrix with n rows and m columns with all entries equal to k, and define K to be the following $(1 + m + n) \times (1 + m + n)$ matrix:

$$D = \begin{bmatrix} 0 & 4 & 0 \\ 2 & 0 & 4 \\ 3 & 2 & 0 \end{bmatrix} \qquad K = \begin{bmatrix} 0 & K_{1 \times m}(4) & K_{1 \times n}(0) \\ K_{m \times 1}(2) & K_{m \times m}(0) & K_{m \times n}(4) \\ K_{n \times 1}(3) & K_{n \times m}(2) & K_{n \times n}(0) \end{bmatrix} \qquad G = K + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & A \\ 0 & B^T & 0 \end{bmatrix}$$

Let $\boldsymbol{a}, \boldsymbol{a}' \in \mathbb{R}^m$, $\boldsymbol{b}, \boldsymbol{b}' \in \mathbb{R}^n$ and $c, c' \in \mathbb{R}$. Now let $\boldsymbol{x}$ and $\boldsymbol{y}$ be the following $(1 + m + n)$-dimensional vectors, $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{y} = (c', \boldsymbol{a}', \boldsymbol{b}')$. Define the *collapse*

of a $(1 + m + n)$-dimensional vector say $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ to be the 3-dimensional vector whose first component is c, the second is the sum of components of $\boldsymbol{a}$ and the third is the sum of components of $\boldsymbol{b}$; we will denote this by cl($\boldsymbol{x}$). Now it is easy to see that if $(\boldsymbol{x}, \boldsymbol{y})$ is a NE of (K, K^T) then $(\text{cl}(\boldsymbol{x}), \text{cl}(\boldsymbol{y}))$ must be a NE of (D, D^T). Therefore, the NE of (K, K^T) must inherit the properties of the NE of D, and hence all four vectors $\{\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{a}', \boldsymbol{b}'\}$ must be nonzero!

The next key idea is to insert the given game (A, B) in this setup in such a way that the composite game has not only the property established above but also captures certain essential features of the game (A, B). For this we will first make the assumption that w.l.o.g. the entries of A and B are positive and $<< 1$, and we construct the matrix G given above. We further define certain 3×3 matrices, $D_{\epsilon_1, \epsilon_2}$, by perturbing D appropriately (see Section 3).

We then show that the Nash equilibria of the symmetric game (G, G^T) have an "image" on the Nash equilibria of the perturbed 3×3 matrix, e.g., we show that payoff of the first player, assuming $(\boldsymbol{x}, \boldsymbol{y})$ is played on the symmetric game (G, G^T), is the same as the payoff of first player if $(\text{cl}(\boldsymbol{x}), \text{cl}(\boldsymbol{y}))$ is played on $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon'_1, \epsilon'_2})$, for a suitable choice of $\epsilon_1, \epsilon_2, \epsilon'_1, \epsilon'_2$. Eventually this leads to showing that (G, G^T) has a non-symmetric NE iff (A, B) has at least two NE, and moreover, the non-symmetric NE of (G, G^T) are in a one-to-one correspondence with ordered pairs of NE of (A, B). These give the NP-hardness and #P-completeness results, respectively.

Second problem: An obvious approach to designing an algorithm for finding symmetric NE in rank-1 symmetric games is to impose symmetry in the approach of Adsul et. al. However, this approach fails and a new approach is called for. In order to describe the salient features of the latter, it is important to recall their approach and show where it fails.

Their approach was to start with the standard quadratic program (QP) that captures all Nash equilibria of a given bimatrix game as optimal solutions. Since $\text{rank}(A + B) = 1$, $A + B = \boldsymbol{c}\boldsymbol{d}^T$, for a suitable choice of vectors $\boldsymbol{c}, \boldsymbol{d}$. After making this substitution, the objective function of the QP becomes the product of two linear forms. [1] replaces one of the linear forms by a parameter λ, thereby getting a parameterized linear program LP(λ). They show that the optimal solutions of this linear program, over all choices of $\lambda \in \mathbb{R}$, are precisely all NE of a certain space of rank 1 games, i.e., $(A, \boldsymbol{u}\boldsymbol{d}^T - A)$, for all choices of $\boldsymbol{u} \in \mathbb{R}^m$; we will denote the bimatrix game $(A, \boldsymbol{u}\boldsymbol{d}^T - A)$ by $(A, \boldsymbol{u}, \boldsymbol{d})$. They further show that the union of all the polyhedra defined by the constraints of LP(λ), over all λ, is yet another polyhedron. The one-skeleton of the latter polyhedron contains a path whose points are in one-to-one correspondence with the optimal solutions of LP(λ), $\forall \lambda$. Additionally, λ is monotonic on this path. Therefore, they are able find a NE of game (A, B) via a binary search for the "correct" value of λ.

Adapting this approach to symmetric rank one games will involve the following. We are given game (A, A^T), where $A + A^T = \boldsymbol{c}\boldsymbol{d}^T$, and we seek a symmetric NE $(\boldsymbol{x}, \boldsymbol{x})$. Clearly, we must start with the standard QP that captures symmetric equilibria of symmetric bimatrix games. The optimal solutions of the analogous parameterized linear program are not even NE of games in the corresponding

space of rank 1 games, i.e., $(A, \boldsymbol{u}, \boldsymbol{c})$, for all choices of $\boldsymbol{u} \in \mathbb{R}^m$. The reason is that this is not a space of symmetric games.

We rectify this situation by moving to a space of symmetric bimatrix games, but of rank 2. This is made possible by the observation that matrix A can be written as the sum of a skew-symmetric matrix K and the rank one matrix $\frac{1}{2}\boldsymbol{c}\boldsymbol{d}^T$, using the fact that $\boldsymbol{c}\boldsymbol{d}^T$ is a symmetric matrix. Now, replacing the vector $\boldsymbol{d}$ by $\boldsymbol{u}$, for all choices of $\boldsymbol{u} \in \mathbb{R}^m$, we get the space of rank 2 symmetric games

$$\left(\left(K + \frac{1}{2}\boldsymbol{c}\boldsymbol{u}^T \right), \left(K + \frac{1}{2}\boldsymbol{u}\boldsymbol{c}^T \right)^T \right).$$

The new LP(λ) will capture all symmetric NE of this space of symmetric games.

At this stage we introduce another idea, thereby achieving a substantial simplification. We bypass the polyhedra mentioned above completely and reduce the problem of finding the "correct" λ to a one-dimensional fixed-point computation in which every fixed point is guaranteed to be rational. Such a fixed point can be found efficiently by a binary search and yields the "correct" λ, which in turn yields the desired symmetric NE.

In what follows we discuss details of main results, however due to space constraint proofs of some of the lemmas and theorems have been moved to Appendix A.

2   Preliminaries

A bimatrix game is a two player game, each player having finitely many pure strategies (moves). Let S_i, $i = 1, 2$ be the set of strategies for player i, and let $m \stackrel{\text{def}}{=} |S_1|$ and $n \stackrel{\text{def}}{=} |S_2|$. Then such a game can be represented by two payoff matrices A and B, each of $m \times n$ dimension. If the first player plays strategy i and the second plays j, then the payoff of the first player is A_{ij} and that of the second player is B_{ij}. Note that the rows of these matrices correspond to the strategies of the first player and the columns to the strategies of second player.

Players may randomize among their strategies; a randomized play is called a *mixed strategy*. The set of mixed strategies for the first player is $X = \{\boldsymbol{x} = (x_1, \ldots, x_m) \mid \boldsymbol{x} \geq 0, \sum_{i=1}^{m} x_i = 1\}$, and for the second player is $Y = \{\boldsymbol{y} = (y_1, \ldots, y_n) \mid \boldsymbol{y} \geq 0, \sum_{j=1}^{n} y_j = 1\}$. By playing $(\boldsymbol{x}, \boldsymbol{y}) \in X \times Y$ we mean strategies are picked independently at random as per $\boldsymbol{x}$ by the first-player and as per $\boldsymbol{y}$ by the second-player. Therefore the expected payoffs of the first-player and second-player are, respectively $\sum_{i,j} A_{ij} x_i y_j = \boldsymbol{x}^T A \boldsymbol{y}$ and $\sum_{i,j} B_{ij} x_i y_j = \boldsymbol{x}^T B \boldsymbol{y}$.

Definition 1. *(Nash Equilibrium [15]) A strategy profile is said to be a Nash equilibrium strategy profile (NESP) if no player achieves a better payoff by a unilateral deviation [12]. Formally, $(\boldsymbol{x}, \boldsymbol{y}) \in X \times Y$ is a NESP iff $\forall \boldsymbol{x}' \in X$, $\boldsymbol{x}^T A \boldsymbol{y} \geq \boldsymbol{x}'^T A \boldsymbol{y}$ and $\forall \boldsymbol{y}' \in Y$, $\boldsymbol{x}^T B \boldsymbol{y} \geq \boldsymbol{x}^T B \boldsymbol{y}'$.*

Given strategy $\boldsymbol{y}$ for the second-player, the first-player gets $(A\boldsymbol{y})_k$ from her k^{th} strategy. Clearly, her best strategies are $\arg\max_k (A\boldsymbol{y})_k$, and a mixed strategy

fetches the maximum payoff only if she randomizes among her best strategies. Similarly, given $\boldsymbol{x}$ for the first-player, the second-player gets $(\boldsymbol{x}^T B)_k$ from k^{th} strategy, and same conclusion applies. These can be equivalently stated as the following complementarity type conditions,

$$\forall i \in S_1, \ x_i > 0 \ \Rightarrow \ (A\boldsymbol{y})_i = \max_{k \in S_1}(A\boldsymbol{y})_k$$
$$\forall j \in S_2, \ y_j > 0 \ \Rightarrow \ (\boldsymbol{x}^T B)_j = \max_{k \in S_2}(\boldsymbol{x}^T B)_k$$

It is easy to get the following from the above discussion: $(\boldsymbol{x}, \boldsymbol{y}) \in X \times Y$ is a NE if and only if the following holds, where π_1 and π_2 are scalars.

$$\begin{aligned} \forall i \in S_1, (A\boldsymbol{y})_i \leq \pi_1; & \quad x_i((A\boldsymbol{y})_i - \pi_1) = 0 \\ \forall j \in S_2, (\boldsymbol{x}^T B)_j \leq \pi_2; & \quad y_j((\boldsymbol{x}^T B)_j - \pi_2) = 0 \end{aligned} \tag{1}$$

Game (A, B) is said to be symmetric if $B = A^T$. In a symmetric game the strategy sets of both the players are identical, i.e., $m = n$, $S_1 = S_2$ and $X = Y$. Therefore, we will use n, S and X for both. A Nash equilibrium profile $(\boldsymbol{x}, \boldsymbol{y}) \in X \times X$ is called *symmetric* if $\boldsymbol{x} = \boldsymbol{y}$. Note that at a symmetric strategy profile $(\boldsymbol{x}, \boldsymbol{x})$ both the players get payoff $\boldsymbol{x}^T A \boldsymbol{x}$. Using (1) it follows that $\boldsymbol{x} \in X$ is a symmetric NE of game (A, A^T), with payoff π to both players, iff,

$$\forall i \in S, (A\boldsymbol{x})_i \leq \pi; \quad x_i((A\boldsymbol{x})_i - \pi) = 0 \tag{2}$$

3   NP-hardness of Non-Symmetric NE in a Symmetric Game

As discussed in the introduction, existence of symmetric NE in a symmetric game is guaranteed [12], however, a symmetric game may not have a non-symmetric equilibrium. In this section, we show that checking existence of non-symmetric NE in general symmetric game is NP-complete and counting such NE is #P-complete. For the NP-completeness result, we will reduce the problem of checking *non-uniqueness of Nash equilibria in bimatrix games*, which is known to be NP-complete [8], to checking if *symmetric game has a non-symmetric equilibrium*. Refer to the second part of Section 1.1 for an overview of the reduction. The reduction is strong enough to also give #P-hardness result since counting equilibria in bimatrix games is known to be #P-hard [5].

We will use the definitions and notation established in Section 1.1. Given $0 \leq \epsilon_1, \epsilon_2 << 1$, define $D_{\epsilon_1, \epsilon_2}$ as below (similar to D defined in Section 1.1):

$$D_{\epsilon_1, \epsilon_2} = \begin{bmatrix} 0 & 4 & 0 \\ 2 & 0 & 4 + \epsilon_1 \\ 3 & 2 + \epsilon_2 & 0 \end{bmatrix} \tag{3}$$

See Appendix A for the proof of the following lemma.

Lemma 1. *Consider the bimatrix game $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'}{}^T)$ where $0 \leq \epsilon_i, \epsilon_i' << 1, i = 1, 2$. The game has a unique NE which has full support, and if $D_{\epsilon_1, \epsilon_2} = D_{\epsilon_1', \epsilon_2'}$ then it is a symmetric NE.*

As observed in Section 1.1, the well known reduction from a bimatrix game (A, B) to symmetric game $G = \begin{bmatrix} 0 & A \\ B^T & 0 \end{bmatrix}$ is not useful for our purpose. This is because, non-symmetric NE of (G, G^T) can be of the form $((\boldsymbol{a}, \boldsymbol{0}), (\boldsymbol{0}, \boldsymbol{b}))$, and therefore fails to produce more than one NE of game (A, B). Next we show how to circumvent this issue by constructing a suitable matrix G using the game of Lemma 1 such that no component in non-symmetric NE is zero, and it relates to a unique pair of NE of game (A, B). This one-to-one correspondence gives #P-hardness result as well.

Recall the following, where $K_{c \times d}(k)$ is a $c \times d$ dimensional matrix with all entries set to $k \in \mathbb{R}$.

$$K = \begin{bmatrix} 0 & K_{1 \times m}(4) & K_{1 \times n}(0) \\ K_{m \times 1}(2) & K_{m \times m}(0) & K_{m \times n}(4) \\ K_{n \times 1}(3) & K_{n \times m}(2) & K_{n \times n}(0) \end{bmatrix} \qquad G = K + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & A \\ 0 & B^T & 0 \end{bmatrix}$$

Before we go into proving our claims, we define a few terms and functions, to be used in the rest of the section. For any non-zero, non-negative vector $\boldsymbol{z}$, of any dimension, let $\eta(\boldsymbol{z})$ denote the *normalized vector*, i.e., its components are non-negative and add to 1. For a matrix M and two non-zero vectors $\boldsymbol{z}_1, \boldsymbol{z}_2 \geq 0$ of appropriate dimensions, we define,

$$P(M; \boldsymbol{z}_1, \boldsymbol{z}_2) \overset{\text{def}}{=} \eta(\boldsymbol{z}_1)^T M \eta(\boldsymbol{z}_2)$$

i.e., the payoffs obtained by player with payoff matrix M at strategy profile $(\eta(\boldsymbol{z}_1), \eta(\boldsymbol{z}_2))$. Given strategy $(\boldsymbol{x}, \boldsymbol{y})$ of game (G, G^T), where $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{y} = (c', \boldsymbol{a}', \boldsymbol{b}')$, or given NE $(\boldsymbol{a}, \boldsymbol{b}')$ and $(\boldsymbol{a}', \boldsymbol{b})$ for game (A, B), define

$$\epsilon_1 \overset{\text{def}}{=} P(A; \boldsymbol{a}, \boldsymbol{b}'), \quad \epsilon_2 \overset{\text{def}}{=} P(B; \boldsymbol{a}', \boldsymbol{b}), \quad \epsilon_1' \overset{\text{def}}{=} P(A; \boldsymbol{a}', \boldsymbol{b}) \quad \epsilon_2' \overset{\text{def}}{=} P(B; \boldsymbol{a}, \boldsymbol{b}') \quad (4)$$

For $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$, recall that $\text{cl}(\boldsymbol{x}) \overset{\text{def}}{=} (c, \sum_i a_i, \sum_j b_j)$. Next we show a connection between payoffs in game (G, G^T) and in game $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'}^T)$.

Lemma 2. $P(G; \boldsymbol{x}, \boldsymbol{y}) = P(D_{\epsilon_1, \epsilon_2}; \text{cl}(\boldsymbol{x}), \text{cl}(\boldsymbol{y}))$, *and* $P(G^T; \boldsymbol{x}, \boldsymbol{y}) = P(D_{\epsilon_1', \epsilon_2'}^T; \text{cl}(\boldsymbol{x}), \text{cl}(\boldsymbol{y}))$.

Proof. We will prove the first part, and the second part follows similarly. Let $\boldsymbol{v} = (v_1, v_2, v_3) = \text{cl}(\boldsymbol{x})$ and $\boldsymbol{w} = (w_1, w_2, w_3) = \text{cl}(\boldsymbol{y})$. Then,

$$P(G; \boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x}^T G \boldsymbol{y} = \boldsymbol{x}^T K \boldsymbol{y} + \boldsymbol{x}^T \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & A \\ 0 & B^T & 0 \end{bmatrix} \boldsymbol{y} = \boldsymbol{x}^T K \boldsymbol{y} + \boldsymbol{a}^T A \boldsymbol{b}' + \boldsymbol{a}'^T B \boldsymbol{b}$$

where, $\boldsymbol{x}^T K \boldsymbol{y} = 2v_2 w_1 + 3v_3 w_1 + 4v_1 w_2 + 2v_3 w_2 + 4v_2 w_3$. Note that $\boldsymbol{a} = (\sum_{i \leq m} a_i) * \eta(\boldsymbol{a}) = v_2 \eta(\boldsymbol{a})$, and similarly $\boldsymbol{b} = v_3 \eta(\boldsymbol{b})$, $\boldsymbol{a}' = w_2 \eta(\boldsymbol{a}')$ and $\boldsymbol{b}' = w_3 \eta(\boldsymbol{b}')$. Thus, $\boldsymbol{a}^T A \boldsymbol{b}' + \boldsymbol{a}'^T B \boldsymbol{b} = \eta(\boldsymbol{a})^T A \eta(\boldsymbol{b}') v_2 w_3 + \eta(\boldsymbol{a}')^T B \eta(\boldsymbol{b}) v_3 w_2$ and hence

$$P(G; \boldsymbol{x}, \boldsymbol{y}) = 2v_2 w_1 + 3v_3 w_1 + 4v_1 w_2 + 2v_3 w_2 + 4v_2 w_3 + \eta(\boldsymbol{a})^T A \eta(\boldsymbol{b}') v_2 w_3 + \eta(\boldsymbol{a}')^T B \eta(\boldsymbol{b}) v_3 w_2.$$

On the other hand we have $P(D_{\epsilon_1,\epsilon_2}; \mathrm{cl}(\boldsymbol{x}), \mathrm{cl}(\boldsymbol{y})) = 2v_2w_1 + 3v_3w_1 + 4v_1w_2 + 2v_3w_2 + 4v_2w_3 + \epsilon_1 v_2 w_3 + \epsilon_2 v_3 w_2$. Since, $\epsilon_1 = P(A; \boldsymbol{a}, \boldsymbol{b}') = \eta(\boldsymbol{a})^T A \eta(\boldsymbol{b}')$ and $\epsilon_2 = P(B; \boldsymbol{a}', \boldsymbol{b}) = \eta(\boldsymbol{a}')^T B \eta(\boldsymbol{b})$ the lemma follows.   □

The above lemma implies equivalence between payoffs in games (G, G^T) and $(D_{\epsilon_1,\epsilon_2}, D_{\epsilon_1',\epsilon_2'}{}^T)$, when the strategies are mapped appropriately. Using this, next we establish relation between their NE.

Lemma 3. *If $(\boldsymbol{x}, \boldsymbol{y})$ is a NE for the game (G, G^T) then $(\mathrm{cl}(\boldsymbol{x}), \mathrm{cl}(\boldsymbol{y}))$ is a NE of game $(D_{\epsilon_1,\epsilon_2}, D_{\epsilon_1',\epsilon_2'}{}^T)$, where ϵs are defined as per (4).*

See Appendix A for the proof of Lemma 3. The next corollary follows using Lemmas 1 and 3.

Corollary 1. *If $(\boldsymbol{x}, \boldsymbol{y})$ is NE for the game (G, G^T), where $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{y} = (c', \boldsymbol{a}', \boldsymbol{b}')$, then vectors $\boldsymbol{a}, \boldsymbol{a}', \boldsymbol{b}, \boldsymbol{b}'$ are non-zero.*

As was our goal, the above corollary establishes non-zeroness of sub-components of a NE $(\boldsymbol{x}, \boldsymbol{y})$ of the symmetric game (G, G^T) that we constructed from (A, B). Using this property we will show how non-symmetric NE of game (G, G^T) give two distinct NE of game (A, B) and vice-versa.

Lemma 4. *If $(\boldsymbol{x}, \boldsymbol{y})$ is a NE for the game (G, G^T), where $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{y} = (c', \boldsymbol{a}', \boldsymbol{b}')$, then $(\eta(\boldsymbol{a}), \eta(\boldsymbol{b}'))$ and $(\eta(\boldsymbol{a}'), \eta(\boldsymbol{b}))$ both are NE for the game (A, B).*

Proof. We will show that $(\eta(\boldsymbol{a}), \eta(\boldsymbol{b}'))$ is NE for the game (A, B), and the proof for $(\eta(\boldsymbol{a}'), \eta(\boldsymbol{b}))$ is analogous. By contradiction, wlog suppose the first player can change $\eta(\boldsymbol{a})$ to $\eta(\boldsymbol{a}'')$ and get a better payoff, where $\sum_{1 \le i \le m} a_i = \sum_{1 \le i \le m} a_i''$. Then, we have $\mathrm{cl}(\boldsymbol{x}) = \mathrm{cl}(\boldsymbol{x}')$ and $\eta(\boldsymbol{a})^T A \eta(\boldsymbol{b}') < \eta(\boldsymbol{a}'')^T A \eta(\boldsymbol{b}')$.

Let $\boldsymbol{x}' = (c, \boldsymbol{a}'', \boldsymbol{b})$. In the proof of Lemma 2 we showed that for $\boldsymbol{v} = \mathrm{cl}(\boldsymbol{x})$ and $\boldsymbol{w} = \mathrm{cl}(\boldsymbol{y})$, $P(G; \boldsymbol{x}, \boldsymbol{y}) = 2v_2w_1 + 3v_3w_1 + 4v_1w_2 + 2v_3w_2 + 4v_2w_3 + \eta(\boldsymbol{a})^T A \eta(\boldsymbol{b}') v_2 w_3 + \eta(\boldsymbol{a}')^T B \eta(\boldsymbol{b}) v_3 w_2$. Then,

$$\begin{aligned}
P(G; \boldsymbol{x}, \boldsymbol{y}) - P(G, \boldsymbol{x}', \boldsymbol{y}) &= \eta(\boldsymbol{a})^T A \eta(\boldsymbol{b}') v_2 w_3 + \eta(\boldsymbol{a}')^T B \eta(\boldsymbol{b}) v_3 w_2 - \\
&\quad \eta(\boldsymbol{a}'')^T A \eta(\boldsymbol{b}') v_2 w_3 - \eta(\boldsymbol{a}')^T B \eta(\boldsymbol{b}) v_3 w_2 \\
&= \eta(\boldsymbol{a})^T A \eta(\boldsymbol{b}') v_2 w_3 - \eta(\boldsymbol{a}'')^T A \eta(\boldsymbol{b}') v_2 w_3 < 0
\end{aligned}$$

A contradiction to $(\boldsymbol{x}, \boldsymbol{y})$ being a NE of (G, G^T).   □

Next we prove the reverse of Lemmas 3 and 4; see appendix A for the proof.

Lemma 5. *If $(\boldsymbol{a}, \boldsymbol{b}')$ and $(\boldsymbol{a}', \boldsymbol{b})$ are NE of game (A, B), and if for ϵs defined in (4), $(\boldsymbol{v}, \boldsymbol{w})$ is a NE of game $(D_{\epsilon_1,\epsilon_2}, D_{\epsilon_1',\epsilon_2'})$, then $((v_1, v_2 * \boldsymbol{a}, v_3 * \boldsymbol{b}), (w_1, w_2 * \boldsymbol{a}', w_3 * \boldsymbol{b}'))$ is a NE of game (G, G^T).*

The next theorem follows directly using Lemmas 3, 4 and 5.

Theorem 1. *For $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{y} = (c', \boldsymbol{a}', \boldsymbol{b}')$, $(\boldsymbol{x}, \boldsymbol{y})$ is a NE of game (G, G^T) iff $(\mathrm{cl}(\boldsymbol{x}), \mathrm{cl}(\boldsymbol{y}))$ is a NE of $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'})$, where ϵs are defined as in (4), and $(\eta(\boldsymbol{a}), \eta(\boldsymbol{b}'))$ and $(\eta(\boldsymbol{a}'), \eta(\boldsymbol{b}))$ are both NE of game (A, B).*

To show NP-hardness of computing non-symmetric NE in symmetric games, we need to establish connection between non-symmetric NE of game (G, G^T) and a pair of distinct NE of game (A, B). Theorem 1 almost does the job except that no such conditions are imposed on the NE of (G, G^T) and of (A, B). Next theorem achieves exactly this (see Appendix A for proof).

Theorem 2. *The symmetric game (G, G^T) has a non-symmetric NE iff the game (A, B) has more than one NE.*

Note that, size of (G, G^T) is $O(size(A, B))$, and hence Theorem 2 implies polynomial-time reduction from the problem of checking if a bimatrix game has more than one NE to checking if a symmetric two-player game has a non-symmetric NE. Since former is NP-complete [8], this shows NP-hardness for the latter. Containment in NP follows since all NE of games (G, G^T) are rational numbers of size polynomial in the size (bit-length) of G [14]. Thus we get the next theorem.

Theorem 3. *Checking existence of a non-symmetric Nash equilibrium in a symmetric game is NP-complete.*

The proof of Theorem 2 does not indicate how the number of equilibria in the games relate to each other. We explore this in the next theorem to show the #P-completeness result (see Appendix A for proof).

Theorem 4. *There is a 1-to-1 correspondence between ordered pairs of (distinct) NE of game (A, B) and non-symmetric NE of the symmetric game (G, G^T).*

The next theorem follows using the fact that counting the number of NE in a bimatrix game is #P-hard [5], and Theorem 4. Here, containment in #P follows from the rationality of NE in bimatrix games.

Theorem 5. *Counting the number of non-symmetric equilibria in a symmetric game is #P-complete.*

4   Efficient Algorithm for Symmetric Rank-1 Games

In this section we consider computing symmetric NE in symmetric constant rank games. Recall that rank of a two player game (A, B) is defined as $rank(A + B)$. Adsul et al. [1] gave a polynomial time algorithm to compute a Nash equilibrium of a rank-1 bimatrix game. In case of a symmetric game, the Nash equilibrium found by the algorithm need not be symmetric. In what follows, we design a polynomial time algorithm to compute a symmetric Nash equilibrium in symmetric rank-1 games. Our algorithm is an extension of the Adsul et al. approach.

Let (A, A^T) be a symmetric game, where A is an $n \times n$ square matrix. As discussed in Section 2 a mixed-strategy $\boldsymbol{x} \in X$ is a symmetric Nash equilibrium of game (A, A^T) if and only if it satisfies (2). Using this, the next lemma follows.

Lemma 6. *If $\boldsymbol{x} \in X$ satisfies first part of (2) then $\boldsymbol{x}^T A \boldsymbol{x} - \pi \leq 0$. Equality holds iff $\boldsymbol{x}$ is a symmetric NE of (A, A^T).*

Using Lemma 6 we get the following quadratic program which exactly captures the symmetric Nash equilibria of game (A, A^T).

$$\max : \boldsymbol{x}^T A \boldsymbol{x} - \pi$$
$$s.t. \quad A\boldsymbol{x} \leq \pi; \quad \boldsymbol{x} \geq \boldsymbol{0}; \quad \textstyle\sum_{i \in S} x_i = 1$$

Note that, the objective value of the above program is at most zero, and exactly zero at the optimal (Lemma 6). If the rank of game (A, A^T) is one then $A + A^T = \boldsymbol{c} \cdot \boldsymbol{d}^T$, where $\boldsymbol{c}, \boldsymbol{d} \in \mathbb{R}^n$. Note that $\boldsymbol{c} \cdot \boldsymbol{d}^T$ is a symmetric matrix, and therefore, we have $\boldsymbol{c} \cdot \boldsymbol{d}^T = \boldsymbol{d} \cdot \boldsymbol{c}^T$

We will represent matrix A as sum of a skew-symmetric matrix and a rank-1 symmetric matrix. Let K be a matrix such that $k_{ij} = a_{ij} - \frac{c_i d_j}{2}$. This implies $A = K + \frac{1}{2}\boldsymbol{c} \cdot \boldsymbol{d}^T$. Since, $\boldsymbol{c} \cdot \boldsymbol{d}^T = \boldsymbol{d} \cdot \boldsymbol{c}^T$ and $A + A^T = \boldsymbol{c} \cdot \boldsymbol{d}^T$, we get $K + K^T = 0$. Thus, K is skew-symmetric, and therefore $\boldsymbol{z}^T K \boldsymbol{z} = 0$ for any vector $\boldsymbol{z} \in \mathbb{R}^n$. Replacing $A = K + \frac{1}{2}\boldsymbol{c} \cdot \boldsymbol{d}^T$ in the above quadratic program we get,

$$\max : \tfrac{1}{2}(\boldsymbol{x}^T \boldsymbol{c})(\boldsymbol{d}^T \boldsymbol{x}) - \pi$$
$$s.t. \quad K\boldsymbol{x} + \tfrac{\boldsymbol{c}}{2}(\boldsymbol{d}^T \boldsymbol{x}) \leq \pi; \quad \boldsymbol{x} \geq \boldsymbol{0}; \quad \textstyle\sum_i x_i = 1$$

The above formulation is a rank-1 quadratic program, which is NP-hard in general. However, we will show that it can be solved in polynomial time using the Nash equilibrium properties. The feasible region of the above program is linear, while the cost function is quadratic which introduces the difficulty. The idea is to construct an LP-type formulation, using the fact that the quadratic term is a product of two linear terms, while maintaining the fact that optimal value of the new formulation is also zero and it is achieved only when complementarity is satisfied. Towards this, we first replace $\boldsymbol{d}^T \boldsymbol{x}$ by λ in the objective function as well as in the inequality. This gives the following optimization problem where $\boldsymbol{x}$ is a variable vector, and π and λ are scalar variables.

$$\max : \tfrac{1}{2}\lambda(\boldsymbol{x}^T \boldsymbol{c}) - \pi \tag{5}$$
$$s.t. \quad K\boldsymbol{x} + \tfrac{\boldsymbol{c}}{2}\lambda \leq \pi; \quad \boldsymbol{x} \geq \boldsymbol{0}; \quad \textstyle\sum_i x_i = 1$$

Lemma 7. *Let $(\boldsymbol{x}, \lambda, \pi)$ be a feasible point of (5), then $\frac{1}{2}\lambda(\boldsymbol{x}^T \boldsymbol{c}) - \pi \leq 0$. Equality holds iff $x_i(K\boldsymbol{x} + \frac{\boldsymbol{c}}{2}\lambda - \pi)_i = 0, \ \forall i \in [n]$.*

Proof. Since $(\boldsymbol{x}, \lambda, \pi)$ satisfies $(K\boldsymbol{x} + \frac{\boldsymbol{c}}{2}\lambda)_i \leq \pi$ and $x_i \geq 0, \ \forall i \in [n]$, we have $x_i(K\boldsymbol{x} + \frac{\boldsymbol{c}}{2}\lambda - \pi)_i \leq 0$. Summing up over all i, and using $\sum_i x_i = 1$ and $\boldsymbol{x}^T K \boldsymbol{x} = 0$, we get $\frac{1}{2}\lambda(\boldsymbol{x}^T \boldsymbol{c}) - \pi \leq 0$. Further, if $x_i(K\boldsymbol{x} + \frac{\boldsymbol{c}}{2}\lambda - \pi)_i = 0, \forall i \in [n]$ then $\frac{1}{2}\lambda(\boldsymbol{x}^T \boldsymbol{c}) - \pi = 0$, and vice-versa. □

Note that, formulation (5) is independent of vector $\boldsymbol{d}$, an essential for our original game. This seems very counter intuitive at first. However, this very property allows (5) to capture NE of a space of games, as established next. Finally, we will use this rich structure to formulate one-dimensional fixed point to solve our game. The next lemma shows that the solution set (5) is rich enough to contain a point for every value of λ; see Appendix A for its proof.

Lemma 8. *Given $\lambda \in \mathbb{R}$, $\exists (\boldsymbol{x}, \pi) \in \mathbb{R}^{n+1}$ such that $(\boldsymbol{x}, \lambda, \pi)$ is a solution of (5), and the objective value is zero at $(\boldsymbol{x}, \lambda, \pi)$.*

Lemmas 7 and 8 imply that the optimal value of (5) is zero, and for every $a \in \mathbb{R}$ there is an optimal solution with $\lambda = a$. If we substitute some value for λ in (5), then it becomes an LP. Therefore, consider it as a parameterized linear program $LP(\lambda)$. The optimal value of $LP(\lambda)$ for any $\lambda \in \mathbb{R}$ is zero (due to Lemma 8). Therefore, solutions of (5) are exactly the solutions of $LP(\lambda)$, $\forall \lambda \in \mathbb{R}$; see remark 2 in Appendix A for a brief discussion on its structure.

Result of the next lemma is central to the construction of one-dimensional fixed point formulation for solving our original game (A, A^T).

Lemma 9. *Given a $\lambda \in \mathbb{R}$, if $(\boldsymbol{x}, \pi)$ is a solution of $LP(\lambda)$ then for any $\boldsymbol{v} \in \mathbb{R}^n$ satisfying $\boldsymbol{v}^T \boldsymbol{x} = \lambda$, $\boldsymbol{x}$ is a symmetric NE of game (Z, Z^T), where $Z = K + \frac{1}{2} \boldsymbol{c} \boldsymbol{v}^T$.*

Proof. Let $(\boldsymbol{x}, \pi)$ be a solution of $LP(\lambda)$, then since the feasible region of $LP(\lambda)$ is a subset of the feasible region of (5), vector $(\boldsymbol{x}, \lambda, \pi)$ satisfies $(K\boldsymbol{x} - \frac{\boldsymbol{c}}{2}\lambda)_i \leq \pi$; $\boldsymbol{x} \geq \boldsymbol{0}$; $\sum_i x_i = 1$. This ensures that $\boldsymbol{x}$ is a probability distribution vector. Due to Lemma 8, it also satisfies $x_i(K\boldsymbol{x} + \frac{\boldsymbol{c}}{2}\lambda - \pi)_i = 0$, $\forall i \in [n]$. Setting, $\lambda = \boldsymbol{v}^T \boldsymbol{x}$, these conditions are exactly that of (2) for strategy $\boldsymbol{x}$ and game (Z, Z^T) where $Z = K + \frac{1}{2} \boldsymbol{c} \boldsymbol{v}^T$. $\qquad\square$

Remark 1. Note that both the matrices of the games constructed in Lemma 9 change with $\boldsymbol{v}$, and $\boldsymbol{c} \boldsymbol{v}^T$ need not be a symmetric matrix. Therefore, rank$(Z + Z^T) = 2$. In the Adsul et. al. approach, the first matrix is same in all the games, and the solutions of $LP(\lambda)$ are NE of a family of rank-1 games, which crucially uses the fact that $\boldsymbol{y}$ need not be same as $\boldsymbol{x}$ (non-symmetric). For this reason, their approach is not immediately applicable for finding symmetric NE.

Lemma 9 implies that if we can find a λ such that the solution $(\boldsymbol{x}, \pi)$ of $LP(\lambda)$ satisfies $\boldsymbol{d}^T \boldsymbol{x} = \lambda$, then $\boldsymbol{x}$ is a symmetric Nash equilibrium of our original rank-1 game (A, A^T). Using this observation, consider a 1-dimensional correspondence $F : [d_{min}, \ d_{max}] \rightarrow 2^{[d_{min}, \ d_{max}]}$, where $d_{min} = \min_{i \in [n]} d_i$ and $d_{max} = \max_{i \in [n]} d_i$.

$$F(\lambda) = \{\boldsymbol{d}^T \boldsymbol{x} \mid \boldsymbol{x} \text{ is a solution of } LP(\lambda)\}$$

By definition we have that $\forall \lambda \in [d_{min}, \ d_{max}]$, $F(\lambda)$ is non-empty (Lemma 8) and convex. Now using the Kakutani fixed-point theorem, F has fixed-points. Clearly, every fixed-point of F gives a Nash equilibrium by Lemma 9, and the next theorem follows.

Theorem 6. *The fixed points of F exactly capture the Nash equilibria of the game (A, A^T).*

Since the Nash equilibrium profiles of game (A, A^T) are rational vectors of size polynomial in the size of A [15], the fixed-points of F are also rational numbers of polynomial sized (using Theorem 6). Thus, one can compute an exact fixed point of F in polynomial time using a simple binary search starting with the pivots d_{min} and d_{max}, and the next theorem follows.

Theorem 7. *The problem of computing a symmetric Nash equilibrium in a symmetric rank-1 game is in P.*

5   Discussion

As stated in the Introduction, the complexity of finding a symmetric NE in a symmetric game of rank 2, 3, 4, or 5 remain open. It is easy to show that for any $k \geq 1$, the rank k problem reduces to the k-dimensional fixed point problem, which is PPAD-hard for $k \geq 2$. Clearly, one way of resolving these questions is to find a reverse reduction. The structure of these problems seem to indicate that such reductions should exist and all these problems should be hard.

References

1. B. Adsul, J. Garg, R. Mehta, and M. Sohoni. Rank-1 bimatrix games: A homeomorphism and a polynomial time algorithm. In *ACM Symposium on the Theory of Computing*, pages 195–204, 2011.
2. X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.
3. S.-F. Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman. Notes on equilibria in symmetric games. In *In Proceedings of International Workshop On Game Theoretic And Decision Theoretic Agents (GTDT)*, pages 71–78, 2004.
4. V. Conitzer. The exact computational complexity of evolutionarily stable strategies. 2013.
5. V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.
6. C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing, Special issue for STOC 2006*, 39(1):195–259, 2009.
7. K. Etessami and A. Lochbihler. The computational complexity of evolutionarily stable strategies. *International Journal of Game Theory*, 37(1):93–113, 2008.
8. I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games Econ. Behav.*, 1:80–93, 1989.
9. R. Kannan and T. Theobald. Games of fixed rank: A hierarchy of bimatrix games. *Economic Theory*, 42(1):157–174, 2010. Preliminary version appeared in SODA'07, and available at arXiv:cs/0511021 since 2005.
10. A. McLennan and R. Tourky. Simple complexity from imitation games. *Games and Economic Behavior*, 68(2):683–688, 2010.
11. R. Mehta. Contant rank bimatrix games are PPAD-hard. In *ACM Symposium on the Theory of Computing*, pages 545–554, 2014.
12. J. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):289–295, September 1951.
13. C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *JCSS*, 48(3):498–532, 1992.
14. C. H. Papadimitriou. The complexity of finding Nash equilibria. *Chapter 2, Algorithmic Game Theory, eds. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani*, pages 29–50, 2007.

15. 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.
16. B. von Stengel. Rank-1 games with exponentially many Nash equilibria. arXiv preprint arXiv:1211.2405, 2012.

Appendix

A   Missing Proofs and Remarks

Proof of Lemma 1.

Let S_1 and S_2 be the support sets of first and second player at a Nash equilibrium $(\boldsymbol{x}, \boldsymbol{y})$. Recall that, we have $S_1 \subseteq \arg\max_{1 \leq i \leq 3}(D_{\epsilon_1, \epsilon_2}\boldsymbol{y})_i$ and $S_2 \subseteq \arg\max_{1 \leq j \leq 3}(\boldsymbol{x}^T D_{\epsilon'_1, \epsilon'_2}{}^T)_j$.

We will first show that $S_1 = S_2 = \{1, 2, 3\}$ for all NE of the game by discarding each of the case where $|S_i| = 1$ and $|S_i| = 2$ for $i = 1, 2$.

Case 1. $|S_1| = 1$ or $|S_2| = 1$.

We will derive a contradiction for $|S_1| = 1$, and the other case follows similarly. Suppose $S_1 = \{1\}$, then $\boldsymbol{x} = (1, 0, 0)$, and $\boldsymbol{x}^T D_{\epsilon'_1, \epsilon'_2}{}^T = \begin{bmatrix} 0 & 2 & 3 \end{bmatrix} \Rightarrow S_2 \subset \{3\}$. Thus only best response to $\boldsymbol{x}$ is $\boldsymbol{y} = (0, 0, 1)$, and it is easy to see $((1, 0, 0), (0, 0, 1))$ is not a NE. Similarly if $S_1 = \{2\}$ then to $\boldsymbol{x} = (0, 1, 0)$ the only best response is $\boldsymbol{y} = (1, 0, 0)$ which is not a NE, and if $S_1 = \{3\}$ then best response to $\boldsymbol{x} = (0, 0, 1)$ is $\boldsymbol{y} = (0, 1, 0)$, again not a NE.

Case 2. $|S_1| = 2$ or $|S_2| = 2$.

Again, we will derive a contradiction for $|S_1| = 2$, and the other case follows similarly. Suppose $S_1 = \{1, 2\}$, then we have $\boldsymbol{x} = (p_1, p_2, 0)$ such that $p_1 + p_2 = 1$, and

$$\boldsymbol{x}^T D_{\epsilon'_1, \epsilon'_2}^T = \begin{bmatrix} 4p_2 & 2p_1 & 3p_1 + (2 + \epsilon'_2)p_2 \end{bmatrix}.$$

Therefore, $S_2 = \{1, 3\}$ in all NE which has $S_1 = \{1, 2\}$ (as $|S_2| > 1$ due to Case 1), so let $\boldsymbol{y} = (q_1, 0, q_3)$. In that case $D_{\epsilon_1, \epsilon_2}\boldsymbol{y} = \begin{bmatrix} 0 & 2q_1 + (4 + \epsilon_1)q_3 & 3q_1 \end{bmatrix}$, and therefore $S_1 \subset \{2, 3\}$, a contradiction. Similarly, $S_1 = \{1, 3\} \Rightarrow S_2 = \{2, 3\} \Rightarrow S_1 \subset \{1, 2\}$, and $S_1 = \{2, 3\} \Rightarrow S_2 = \{1, 2\} \Rightarrow S_1 \subset \{1, 3\}$, contradictions.

Thus the only possibility we are left with is $S_1 = S_2 = \{1, 2, 3\}$ for all NE of the game. Next we show that there is a unique NE with this support. Let $((p_1, p_2, p_3), (q_1, q_2, q_3))$ be a NE of the game. Since $S_2 = \{1, 2, 3\}$, we have:

$$4p_2 = 2p_1 + (4 + \epsilon'_1)p_3 = 3p_1 + (2 + \epsilon'_2)p_2, \quad p_1 + p_2 + p_3 = 1$$

Clearly, the only solution of the above equalities is $p_1 = \frac{2}{7} + \epsilon''_1$, $p_2 = \frac{3}{7} + \epsilon''_2$, $p_3 = \frac{2}{7} + \epsilon''_3$, where $\epsilon''_i > 0$, $i = 1, 2, 3$ depends on ϵ_1 and ϵ_2. We can write similar equalities for (q_1, q_2, q_3) using NE conditions for the first player, which has a unique solution $q_1 = \frac{2}{7} + \epsilon'''_1$, $q_2 = \frac{3}{7} + \epsilon'''_2$, $q_3 = \frac{2}{7} + \epsilon'''_3$.

Since existence of symmetric NE in a symmetric game is guaranteed, if $D_{\epsilon_1, \epsilon_2} = D_{\epsilon'_1, \epsilon'_2}$, then the only NE of such a game is symmetric. $\square$

Proof of Lemma 3.

To the contrary suppose $(\boldsymbol{v}, \boldsymbol{w}) = (\mathrm{cl}(\boldsymbol{x}), \mathrm{cl}(\boldsymbol{y}))$ is not a NE of game $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'}^T)$. Without loss of generality (wlog) suppose first player can deviate to $\boldsymbol{v}'$ and get a better payoff. Let $\boldsymbol{x}' = (v_1', v_2'\eta(\boldsymbol{a}), v_3'\eta(\boldsymbol{b}))$. We claim that in game (G, G^T) first player can deviate to $\boldsymbol{x}'$ and gain, i.e., $P(G; \boldsymbol{x}', \boldsymbol{y}) > P(G; \boldsymbol{x}, \boldsymbol{y})$, contradicting $(\boldsymbol{x}, \boldsymbol{y})$ being NE.

Note that the values of $\epsilon_1, \epsilon_2, \epsilon_1', \epsilon_2'$ for $(\boldsymbol{x}, \boldsymbol{y})$ and $(\boldsymbol{x}', \boldsymbol{y})$ are the same (see (4)). Then, using Lemma 2 we have: $P(G; \boldsymbol{x}, \boldsymbol{y}) = P(D_{\epsilon_1, \epsilon_2}; \boldsymbol{v}, \boldsymbol{w}) < P(D_{\epsilon_1, \epsilon_2}; \boldsymbol{v}', \boldsymbol{w}) = P(G; \boldsymbol{x}', \boldsymbol{y})$. $\qquad\square$

Proof of Lemma 5.

Let $\boldsymbol{x} = (v_1, v_2 * \boldsymbol{a}, v_3 * \boldsymbol{b})$ and $\boldsymbol{y} = (w_1, w_2 * \boldsymbol{a}', w_3 * \boldsymbol{b}')$. We will check the complementary conditions (1) for $(\boldsymbol{x}, \boldsymbol{y})$. Let S_1 be the support set of first player and S_2 be the support set of second player w.r.t. profile $(\boldsymbol{x}, \boldsymbol{y})$. In particular, we need to show:

$$\forall i \in S_1, (G\boldsymbol{y})_i = max_{1 \leq i \leq m}(G\boldsymbol{y})_i, \qquad \forall i \in S_2, (\boldsymbol{x}^T G^T)_i = max_{1 \leq i \leq n}(\boldsymbol{x}^T G^T)_i.$$

We next show the first condition, and the proof for the second condition is similar. Let $\boldsymbol{z} = G\boldsymbol{y}$, then, $z_1 = 4w_2$; $2 \leq i \leq m+1$, $z_i = 2w_1 + 4w_3 + w_3(A\boldsymbol{b}')_{i-1}$; $m+2 \leq j \leq n+m+1$, $z_j = 3w_1 + 2w_2 + w_2(B^T\boldsymbol{a}')_{j-m-1}$. The max among these three sets of strategies are $max_1 = 4w_1$, $max_2 = 2w_1 + 4w_3 + w_3 max_{i \leq m}(A\boldsymbol{b}')_i$ and $max_3 = 3w_1 + 2w_2 + w_2 max_{j \leq n}(\boldsymbol{a}'^T B)_j$. If we show that $max_1 = max_2 = max_3$ then lemma follows using the fact that $\boldsymbol{a}$ is a best response of the first-player against $\boldsymbol{b}'$, $\boldsymbol{b}$ is a best response of the second-player against $\boldsymbol{a}'$ in game (A, B) (as $(\boldsymbol{a}, \boldsymbol{b}')$ and $(\boldsymbol{a}', \boldsymbol{b})$ are its NE).

Since $(\boldsymbol{v}, \boldsymbol{w})$ is NE for the game $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'}^T)$ where $\epsilon_1 = \boldsymbol{a}^T A\boldsymbol{b}'$ and $\epsilon_2 = \boldsymbol{a}'^T B\boldsymbol{b}$, and it is of full support (Lemma 1), we have $4w_2 = 2w_1 + 4w_3 + w_3\epsilon_1 = 3w_1 + 2w_2 + w_2\epsilon_2$. On the other hand, $(\boldsymbol{a}, \boldsymbol{b}')$ and $(\boldsymbol{a}', \boldsymbol{b})$ being NE for the game (A, B) implies $\epsilon_1 = \boldsymbol{a}^T A\boldsymbol{b}' = max_{1 \leq i \leq m}(A\boldsymbol{b}')_i$ and $\epsilon_2 = \boldsymbol{a}'^T B\boldsymbol{b} = max_{1 \leq j \leq n}(\boldsymbol{a}'^T B)_j$. Thus we get $max_1 = max_2 = max_3$ as desired. $\qquad\square$

Proof of Theorem 2.

The symmetric game (G, G^T) has a non-symmetric NE iff the game (A, B) has more than one NE. $(\Rightarrow)$ Let $(\boldsymbol{x}, \boldsymbol{y})$ be a non-symmetric NE of (G, G^T), where $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{y} = (c', \boldsymbol{a}', \boldsymbol{b}')$. Then, using Theorem 1, $(\eta(\boldsymbol{a}), \eta(\boldsymbol{b}'))$ and $(\eta(\boldsymbol{a}'), \eta(\boldsymbol{b}))$ are NE for the game (A, B). We need to show that these are distinct. To the contrary suppose $\eta(\boldsymbol{a}) = \eta(\boldsymbol{a}')$ and $\eta(\boldsymbol{b}) = \eta(\boldsymbol{b}')$. Then, $\epsilon_1 = \epsilon_1'$ and $\epsilon_2 = \epsilon_2'$. Further, $(\mathrm{cl}(\boldsymbol{x}), \mathrm{cl}(\boldsymbol{y}))$ being NE of $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'}^T)$ (Theorem 1), we have $\mathrm{cl}(\boldsymbol{x}) = \mathrm{cl}(\boldsymbol{y})$ (Lemma 1). This together with $\eta(\boldsymbol{a}) = \eta(\boldsymbol{a}')$ and $\eta(\boldsymbol{b}) = \eta(\boldsymbol{b}')$ implies $c = c'$, $\boldsymbol{a} = \boldsymbol{a}'$ and $\boldsymbol{b} = \boldsymbol{b}'$, a contradiction.

$(\Leftarrow)$Let $(\boldsymbol{a}, \boldsymbol{b}')$ and $(\boldsymbol{a}', \boldsymbol{b})$ be two different NE for the game (A, B), and let $(\boldsymbol{v}, \boldsymbol{w})$ be NE for the game $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'}^T)$. Let $\boldsymbol{x} = (v_1, v_2\boldsymbol{a}, v_3\boldsymbol{b})$ and $\boldsymbol{y} = (w_1, w_2\boldsymbol{a}', w_3\boldsymbol{b}')$. then using Theorem 1 $(\boldsymbol{x}, \boldsymbol{y})$ is a NE of game (G, G^T). Further, $(\boldsymbol{a}, \boldsymbol{b}') \neq (\boldsymbol{a}', \boldsymbol{b})$ implies $\boldsymbol{x} \neq \boldsymbol{y}$. $\qquad\square$

Proof of Theorem 4.

First we show that for every ordered pair of NE of the game (A, B), $((\boldsymbol{a}, \boldsymbol{b}'), (\boldsymbol{a}', \boldsymbol{b}))$, there is a unique non-symmetric NE for the game (G, G^T). Let $(\boldsymbol{v}, \boldsymbol{w})$ be NE for the game $(D_{\epsilon_1, \epsilon_2}, D_{\epsilon_1', \epsilon_2'}{}^T)$. Let $\boldsymbol{x} = (v_1, v_2\boldsymbol{a}, v_3\boldsymbol{b})$ and $\boldsymbol{y} = (w_1, w_2\boldsymbol{a}', w_3\boldsymbol{b}')$, then $(\boldsymbol{x}, \boldsymbol{y})$ is a non-symmetric NE of (G, G^T) (Theorem 2). Similarly if another NE pair $((\boldsymbol{\alpha}, \boldsymbol{\beta}'), (\boldsymbol{\alpha}', \boldsymbol{\beta}))$ also gives $(\boldsymbol{x}, \boldsymbol{y})$, then using Lemma 4 $\boldsymbol{\alpha} = \boldsymbol{a}$, $\boldsymbol{\beta}' = \boldsymbol{b}'$, $\boldsymbol{\alpha}' = \boldsymbol{a}'$ and $\boldsymbol{\beta} = \boldsymbol{b}$.

Next, for every non-symmetric NE $(\boldsymbol{x}, \boldsymbol{y})$ of game (G, G^T), we will construct a unique pair of distinct NE for game (A, B). Let $\boldsymbol{x} = (c, \boldsymbol{a}, \boldsymbol{b})$ and $\boldsymbol{y} = (c', \boldsymbol{a}', \boldsymbol{b}')$, then $(\eta(\boldsymbol{a}), \eta(\boldsymbol{b}'))$ and $(\eta(\boldsymbol{a}'), \eta(\boldsymbol{b}))$ are two different NE for the game (A, B) (Theorem 2). Suppose another symmetric NE $(\boldsymbol{x}', \boldsymbol{y}')$ of (G, G^T), where $\boldsymbol{x}' = (o, \boldsymbol{p}, \boldsymbol{q})$ and $\boldsymbol{y}' = (o', \boldsymbol{p}', \boldsymbol{q}')$, gives the same pair, i.e., $\eta(\boldsymbol{p}) = \eta(\boldsymbol{a})$, $\eta(\boldsymbol{q}) = \eta(\boldsymbol{b})$, $\eta(\boldsymbol{p}') = \eta(\boldsymbol{a}')$ and $\eta(\boldsymbol{q}') = \eta(\boldsymbol{b}')$. Then, ϵ_i and ϵ_i', $i = 1, 2$ defined in (4) are the same for both $(\boldsymbol{x}, \boldsymbol{y})$ and $(\boldsymbol{x}', \boldsymbol{y}')$. Therefore, using Lemma 1 and Theorem 1, we have $(\mathrm{cl}(\boldsymbol{x}), \mathrm{cl}(\boldsymbol{y})) = (\mathrm{cl}(\boldsymbol{x}'), \mathrm{cl}(\boldsymbol{y}'))$, implying $(\boldsymbol{x}, \boldsymbol{y}) = (\boldsymbol{x}', \boldsymbol{y}')$.   □

Proof of Lemma 8.

Let $\boldsymbol{v} = \lambda * \mathbf{1}$, where $\mathbf{1}$ is an n-dimensional vector with all 1s. Now, consider a symmetric game with matrix $Z = K + \frac{1}{2}\boldsymbol{c}\boldsymbol{v}^T$. Let $\boldsymbol{z}$ be a symmetric Nash equilibrium of this game and π be the corresponding payoff $\boldsymbol{z}^T Z \boldsymbol{z}$, then $\boldsymbol{v}^T \boldsymbol{z} = \lambda$. Therefore, using (2) we have

$$\forall i \in [n]: \quad (K\boldsymbol{z} + \frac{\boldsymbol{c}}{2}\lambda)_i \leq \pi; \quad \boldsymbol{z}_i((K\boldsymbol{z} + \frac{\boldsymbol{c}}{2}\lambda)_i - \pi) = 0$$

The first inequality implies that $(\boldsymbol{z}, \lambda, \pi)$ is feasible in (5), and the second implies that objective value is zero at it. Therefore, it is optimal due to Lemma 7.   □

Remark 2. Let S be the set of optimal solutions of (5). Since the optimal value of (5) is zero, it is easy to see that $(\boldsymbol{x}, \lambda, \pi) \in S$ if and only if $\forall i \in [n]$, $x_i(K\boldsymbol{x} - \frac{\boldsymbol{c}}{2}\lambda - \pi)_i = 0$, which requires n equalities. Let us assume that the polyhedron of (5) in $(\boldsymbol{x}, \pi, \lambda)$-space is non-degenerate. Due to the equality $\sum_i x_i = 1$, the polyhedron is in $(n + 1)$-dimension. Now, since at least n equalities are tight on every point of S, we get $S \subset$ 1-skeleton of the polyhedron. Further, we can show that every vertex of S has degree two in S, and thus S forms cycles and paths (with unbounded edges on both ends). In addition, for any $a \in \mathbb{R}$, the set of points of S with $\lambda = a$ are exactly the solutions of $LP(a)$ which must be a closed and convex set. Using all these properties it is easy to see that S forms a single path, on which λ changes monotonically. In this paper, we bypass the polyhedron completely and give a much simpler approach in the next lemma.

B Imitation Games

Given an $n \times n$ matrix A, consider a bimatrix game (A, I), where I is the $n \times n$ identity matrix is called an *imitation matrix*. The name arises from the fact that

when restricted to pure strategies, the second player gets a payoff of 1 if she plays the same strategy as the first player and 0 otherwise. Imitation games lead to simple reductions between games. We illustrate this in Lemma 10 which leads to a reduction from NE in a bimatrix game to a symmetric NE in a symmetric bimatrix game and is well known.

Lemma 10. *[10] Let $(\boldsymbol{x}, \boldsymbol{y})$ be a Nash equilibrium for bimatrix game (A, I) where A is an $n \times n$ matrix of positive entries and I is the $n \times n$ identity matrix. Then $(\boldsymbol{y}, \boldsymbol{y})$ is a symmetric Nash equilibrium for (A, A^T).*

Proof. Let $\alpha = \max_i (A\boldsymbol{y})_i$ and $\beta = \max_i (I\boldsymbol{x})_i$. Clearly, $\alpha > 0$ and $\beta > 0$. Since $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium for (A, I),

- $\forall i: \ \ x_i > 0 \ \Rightarrow \ (A\boldsymbol{y})_i = \alpha \ \ldots \ (1)$.
- $\forall i: \ \ y_i > 0 \ \Rightarrow \ (I\boldsymbol{x})_i = \beta \ \ldots \ (2)$.

It is enough show that $\forall i: \ \ y_i > 0 \ \Rightarrow \ (A\boldsymbol{y})_i = \alpha$. Clearly, $y_i > 0 \ \Rightarrow$ $(I\boldsymbol{x})_i = \beta > 0$, (by (2)) $\Rightarrow x_i > 0 \ \Rightarrow \ (A\boldsymbol{y})_i = \alpha$ (by (1)). $\qquad \square$

An $n \times n$ matrix will be said to be a *positive diagonal matrix* if each of its diagonal entries is a positive number and each of the non-diagonal entries is zero. Observe that Lemma 10 holds even if I is replaced by an arbitrary positive diagonal matrix. This immediately raises the questions, "By changing the positive diagonal matrix, are we guaranteed to get a different symmetric Nash equilibrium for (A, A^T)? If so, can we get all symmetric Nash equilibria in this manner?" Theorem 8 provides a negative answer to the first question and Theorem 9 provides additional insights to the second question.

Theorem 8. *Let $(\boldsymbol{x}, \boldsymbol{y})$ be a Nash equilibrium for bimatrix game (A, I) and D be a positive diagonal matrix with $d_i > 0$ in the i^{th} diagonal entry. Let $s = \sum_i x_i/d_i$ and let $\boldsymbol{x}'$ be the vector whose i^{th} entry is $x_i' = x_i/(sd_i)$. Then $(\boldsymbol{x}', \boldsymbol{y})$ is a Nash equilibrium for bimatrix game (A, D).*

Proof. Clearly $x_i' \geq 0$ for $1 \leq i \leq n$ and $\sum_i x_i' = s/s = 1$, hence $\boldsymbol{x}'$ is a probability vector. Clearly, $x_i > 0$ iff $x_i' > 0$. Therefore, since $\boldsymbol{x}$ is a best response to $\boldsymbol{y}$ in the game (A, I), $\boldsymbol{x}'$ is a best response to $\boldsymbol{y}$ in the game (A, D).

Let $\max_i (I\boldsymbol{x})_i = \beta$. Clearly, $\max_i (D\boldsymbol{x}')_i = \beta/s$; furthermore, the set of indices which achieve maximum are the same in both equalities. Therefore, $y_i > 0 \ \Rightarrow \ (I\boldsymbol{x})_i = \beta \ \Rightarrow (D\boldsymbol{x}')_i = \beta/s$. Therefore, $\boldsymbol{y}$ is a best response to $\boldsymbol{x}'$ in the game (A, D), hence proving the theorem. $\qquad \square$

Lemma 10 and Theorem 8 raise the question, "Consider the bimatrix game (A, D) for a fixed positive diagonal matrix D. Can one characterize the set of symmetric Nash equilibria $\boldsymbol{y}$ of (A, A^T) such that $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium of the game (A, D) for some probability vector $\boldsymbol{x}$?" Theorem 9 implies that this set consists of all symmetric Nash equilibria of (A, A^T).

Theorem 9. *Let $\boldsymbol{y}$ be any symmetric Nash equilibrium for bimatrix game (A, A^T) and D be a positive diagonal matrix. Then there is a probability vector $\boldsymbol{x}$ such that $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium for bimatrix game (A, D).*

Proof. We will construct a probability vector $\boldsymbol{x}$ such that $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium for bimatrix game (A, I). Then, the needed result will follow from Theorem 8.

Let $S = \{i \mid (A\boldsymbol{y})_i > 0\}$. Let $\boldsymbol{x}$ be the probability vector whose i^{th} coordinate is $1/|S|$ if $i \in S$ and 0 otherwise. Let $\max_i (A\boldsymbol{y})_i = \alpha$. Since $\boldsymbol{y}$ is a symmetric Nash equilibrium for bimatrix game (A, A^T), $\forall i : \quad y_i > 0 \quad \Rightarrow \quad (A\boldsymbol{y})_i = \alpha$. Therefore, the set of indices which achieve maximum in $\boldsymbol{x}$ are the same as those in $\boldsymbol{y}$, hence showing that $(\boldsymbol{x}, \boldsymbol{y})$ is a Nash equilibrium for bimatrix game (A, I).
$\square$