Semidefinite Programming and Nash Equilibria in Bimatrix Games

Amir Ali Ahmadi and Jeffrey Zhang *

Abstract

We explore the power of semidefinite programming (SDP) for finding additive ϵ-approximate Nash equilibria in bimatrix games. We introduce an SDP relaxation for a quadratic programming formulation of the Nash equilibrium (NE) problem and provide a number of valid inequalities to improve the quality of the relaxation. If a rank-1 solution to this SDP is found, then an exact NE can be recovered. We show that for a strictly competitive game, our SDP is guaranteed to return a rank-1 solution. We propose two algorithms based on iterative linearization of smooth nonconvex objective functions whose global minima by design coincide with rank-1 solutions. Empirically, we demonstrate that these algorithms often recover solutions of rank at most two and ϵ close to zero. Furthermore, we prove that if a rank-2 solution to our SDP is found, then a $\frac{5}{11}$-NE can be recovered for any game, or a $\frac{1}{3}$-NE for a symmetric game. We then show how our SDP approach can address two (NP-hard) problems of economic interest: finding the maximum welfare achievable under any NE, and testing whether there exists a NE where a particular set of strategies is not played. Finally, we show the connection between our SDP and the first level of the Lasserre/sum of squares hierarchy.

Keywords: Nash equilibria, semidefinite programming, correlated equilibria.

1 Introduction

A bimatrix game is a game between two players (referred to in this paper as players A and B) defined by a pair of $m \times n$ payoff matrices A and B. Let $\triangle_m$ and $\triangle_n$ denote the m-dimensional and n-dimensional simplices

$$\triangle_m = \{x \in \mathbb{R}^m |\ x_i \geq 0, \forall i, \sum_{i=1}^{m} x_i = 1\}, \triangle_n = \{y \in \mathbb{R}^n |\ y_i \geq 0, \forall i, \sum_{i=1}^{n} y_i = 1\}.$$

These form the strategy spaces of player A and player B respectively. For a strategy pair $(x, y) \in \triangle_m \times \triangle_n$, the payoff received by player A (resp. player B) is $x^T Ay$ (resp. $x^T By$). In particular, if the players pick vertices i and j of their respective simplices (also called pure strategies), their payoffs will be $A_{i,j}$ and $B_{i,j}$. One of the prevailing solution concepts for bimatrix games is the notion of *Nash equilibrium.* At such an equilibrium, the players are playing mutual best responses, i.e., a payoff maximizing strategy against the opposing player's strategy. In our notation, a Nash equilibrium for the game (A, B) is a pair of strategies $(x^*, y^*) \in \triangle_m \times \triangle_n$ such that

$$x^{*T} Ay^* \geq x^T Ay^*, \forall x \in \triangle_m,$$

and

$$x^{*T} By^* \geq x^{*T} By, \forall y \in \triangle_n.$$[1]

*The authors are partially supported by the DARPA Young Faculty Award, the Young Investigator Award of the AFOSR, the CAREER Award of the NSF, the Google Faculty Award, and the Sloan Fellowship.

[1]In this paper we assume that all entries of A and B are between 0 and 1, and argue at the beginning of Section 2 why this is without loss of generality for the purpose of computing Nash equilibria.

Nash [36] proved that for any bimatrix game, such pairs of strategies exist (in fact his result more generally applies to games with a finite number of players and a finite number of pure strategies). While existence of these equilibria is guaranteed, finding them is believed to be a computationally intractable problem. More precisely, a result of Daskalakis, Goldberg, and Papadimitriou [16] implies that computing Nash equilibria is PPAD-complete (see [16] for a definition) even when the number of players is 3. This result was later improved by Chen and Deng [9] who showed the same hardness result for bimatrix games.

These results motivate the notion of an approximate Nash equilibrium, a solution concept in which players receive payoffs "close" to their best response payoffs. More precisely, a pair of strategies $(x^*, y^*) \in \triangle_m \times \triangle_n$ is an *(additive) ϵ-Nash equilibrium* for the game (A, B) if

$$x^{*T} A y^* \geq x^T A y^* - \epsilon, \forall x \in \triangle_m,$$

and

$$x^{*T} B y^* \geq x^{*T} B y - \epsilon, \forall y \in \triangle_n.^2$$

Note that when $\epsilon = 0$, (x^*, y^*) form an exact Nash equilibrium, and hence it is of interest to find ϵ-Nash equilibria with ϵ small. Unfortunately, approximation of Nash equilibria has also proved to be computationally difficult. Cheng, Deng, and Teng have shown in [10] that, unless PPAD $\subseteq$ P, there cannot be a fully polynomial-time approximation scheme for computing Nash equilibria in bimatrix games. There have, however, been a series of constant factor approximation algorithms for this problem [18, 17, 27, 44], with the current best producing a .3393 approximation via an algorithm by Tsaknakis and Spirakis [44].

We remark that there are exponential-time algorithms for computing Nash equilibria, such as the Lemke-Howson algorithm [32, 41]. There are also certain subclasses of the problem which can be solved in polynomial time, the most notable example being the case of zero-sum games (i.e. when $B = -A$). This problem was shown to be solvable via linear programming by Dantzig [14], and later shown to be polynomially equivalent to linear programming by Adler [2]. Aside from computation of Nash equilibria, there are a number of related decision questions which are of economic interest but unfortunately NP-hard. Examples include deciding whether a player's payoff exceeds a certain threshold in some Nash equilibrium, deciding whether a game has a unique Nash equilibrium, or testing whether there exists a Nash equilibrium where a particular set of strategies is not played [21, 12].

Our focus in this paper is on understanding the power of semidefinite programming[3] (SDP) for finding approximate Nash equilibria in bimatrix games or providing certificates for related decision questions. The goal is not to develop a competitive solver, but rather to analyze the algorithmic power of SDP when applied to basic problems around computation of Nash equilibria. Semidefinite programming relaxations have been analyzed in depth in areas such as combinatorial optimization [22], [33] and systems theory [8], but not to such an extent in game theory. To our knowledge, the appearance of SDP in the game theory literature includes the work of Stein for exchangeable equilibria in symmetric games [43], of Parrilo on zero-sum polynomial games [38], of Parrilo and Shah for zero-sum stochastic games [42], and of Laraki and Lasserre for semialgebraic min-max problems in static and dynamic games [28].

1.1 Organization and Contributions of the Paper

In Section 2, we formulate the problem of finding a Nash equilibrium in a bimatrix game as a nonconvex quadratically constrained quadratic program and pose a natural SDP relaxation for it. In

[2]There are also other important notions of approximate Nash equilibria, such as ϵ-well-supported Nash equilibria [20] and relative approximate Nash equilibria [15] which are not considered in this paper.

[3]The unfamiliar reader is referred to [45] for the theory of SDPs and a description of polynomial-time algorithms for them based on interior point methods.

Section 3, we show that our SDP is exact when the game is strictly competitive (see Definition 3.3). In Section 4, we design two continuous but nonconvex objective functions for our SDP whose global minima coincide with rank-1 solutions. We provide a heuristic based on iterative linearization for minimizing both objective functions. We show empirically that these approaches produce ϵ very close to zero (on average in the order of 10^{-3}). In Section 5, we establish a number of bounds on the quality of the approximate Nash equilibria that can be read off of feasible solutions to our SDP. In Theorems 5.5, 5.6, and 5.8, we show that when the SDP returns solutions which are "close" to rank-1, the resulting strategies have have small ϵ. We then present an improved analysis in the rank-2 case which shows how one can recover a $\frac{5}{11}$-Nash equilibrium from the SDP solution (Theorem 5.10). We further prove that for symmetric games (i.e., when $B = A^T$), a $\frac{1}{3}$-Nash equilibrium can be recovered in the rank-2 case (Theorem 5.17). We do not currently know of a polynomial-time algorithm for finding rank-2 solutions to our SDP. If such an algorithm were found, it would, together with our analysis, improve the best known approximation bound for symmetric games. In Section 6, we show how our SDP formulation can be used to provide certificates for certain (NP-hard) questions of economic interest about Nash equilibria in symmetric games. These are the problems of testing whether the maximum welfare achievable under any symmetric Nash equilibrium exceeds some threshold, and whether a set of strategies is played in every symmetric Nash equilibrium. In Section 7, we show that the SDP analyzed in this paper dominates the first level of the Lasserre hierarchy (Proposition 7.1). Some directions for future research are discussed in Section 8. The four appendices of the paper add some numerical and technical details.

2 The Formulation of our SDP Relaxation

In this section we present an SDP relaxation for the problem of finding Nash equilibria in bimatrix games. This is done after a straightforward reformulation of the problem as a nonconvex quadratically constrained quadratic program. Throughout the paper the following notation is used.

- $A_{i,}$ refers to the i-th row of a matrix A.
- $A_{,j}$ refers to the j-th column of a matrix A.
- e_i refers to the elementary vector $(0, \ldots, 0, 1, 0, \ldots, 0)^T$ with the 1 being in position i.
- $\triangle_k$ refers to the k-dimensional simplex.
- 1_m refers to the m-dimensional vector of one's.
- 0_m refers to the m-dimensional vector of zero's.
- $J_{m,n}$ refers to the $m \times n$ matrix of one's.
- $A \succeq 0$ denotes that the matrix A is positive semidefinite (psd), i.e., has nonnegative eigenvalues.
- $A \geq 0$ denotes that the matrix A is nonnegative, i.e., has nonnegative entries.
- $A \succeq B$ denotes that $A - B \succeq 0$.
- $\mathbb{S}^{k \times k}$ denotes the set of symmetric $k \times k$ matrices.
- $\mathrm{Tr}(A)$ denotes the trace of a matrix A, i.e., the sum of its diagonal elements.
- $A \otimes B$ denotes the Kronecker product of matrices A and B.
- $vec(M)$ denotes the vectorized version of a matrix M.
- For a vector v, $diag(v)$ denotes the diagonal matrix with v on its diagonal. For a square matrix M, $diag(M)$ denotes the vector containing its diagonal entries.

We also assume that all entries of the payoff matrices A and B are between 0 and 1. This can be done without loss of generality because Nash equilibria are invariant under certain affine transformations in the payoffs. In particular, the games (A, B) and $(cA + dJ_{m\times n}, eB + fJ_{m\times n})$

have the same Nash equilibria for any scalars c, d, e, and f, with c and e positive. This is because

$$\begin{aligned} x^{*T}Ay &\geq x^TAy \\ \Leftrightarrow c(x^{*T}Ay^*) + d &\geq c(x^TAy^*) + d \\ \Leftrightarrow c(x^{*T}Ay^*) + d(x^{*T}J_{m\times n}y^*) &\geq c(x^TAy^*) + d(x^TJ_{m\times n}y^*) \\ \Leftrightarrow x^{*T}(cA + dJ_{m\times n})y^* &\geq x^T(cA + dJ_{m\times n})y \end{aligned}$$

Identical reasoning applies for player B.

2.1 Nash Equilibria as Solutions to Quadratic Programs

Recall the definition of a Nash equilibrium from Section 1. An equivalent characterizaiton is that a strategy pair $(x^*, y^*) \in \triangle_m \times \triangle_n$ is a Nash equilibrium for the game (A, B) if and only if

$$\begin{aligned} x^{*T}Ay^* &\geq e_i^TAy^*, \forall i \in \{1, \ldots, m\}, \\ x^{*T}By^* &\geq x^{*T}Be_i, \forall i \in \{1, \ldots, n\}. \end{aligned} \tag{1}$$

The equivalence can be seen by noting that because the payoff from playing any mixed strategy is a convex combination of payoffs from playing pure strategies, there is always a pure strategy best response to the other player's strategy.

We now treat the Nash problem as the following quadratic programming (QP) feasibility problem:

$$\begin{aligned} \min_{x\in\mathbb{R}^m, y\in\mathbb{R}^n} \quad & 0 \\ \text{subject to} \quad & x^TAy \geq e_i^TAy, \forall i \in \{1, \ldots, m\}, \\ & x^TBy \geq x^TBe_j, \forall j \in \{1, \ldots, n\}, \\ & x_i \geq 0, \forall i \in \{1, \ldots, m\}, \\ & y_i \geq 0, \forall j \in \{1, \ldots, n\}, \\ & \sum_{i=1}^m x_i = 1, \\ & \sum_{i=1}^n y_i = 1. \end{aligned} \tag{2}$$

Similarly, a pair of strategies $x^* \in \triangle_m$ and $y^* \in \triangle_n$ form an ϵ-Nash equilibrium for the game (A, B) if and only if

$$\begin{aligned} x^{*T}Ay^* &\geq e_i^TAy^* - \epsilon, \forall i \in \{1, \ldots, m\}, \\ x^{*T}By^* &\geq x^{*T}Be_i - \epsilon, \forall i \in \{1, \ldots, n\}. \end{aligned}$$

Observe that any pair of simplex vectors (x, y) is an ϵ-Nash equilibrium for the game (A, B) for any ϵ that satisfies

$$\epsilon \geq \max\{\max_i e_i^TAy - x^TAy, \max_i x^TBe_i - x^TBy\}.$$

We use the following notation throughout the paper:

- $\epsilon_A(x, y) := \max_i e_i^TAy - x^TAy$,
- $\epsilon_B(x, y) := \max_i x^TBe_i - x^TBy$,
- $\epsilon(x, y) := \max\{\epsilon_A(x, y), \epsilon_B(x, y)\}$,

and the function parameters are later omitted if they are clear from the context.

2.2 SDP Relaxation

The QP formulation in (2) lends itself to a natural SDP relaxation. We define a matrix

$$\mathcal{M} := \begin{bmatrix} X & P \\ Z & Y \end{bmatrix},$$

and an augmented matrix

$$\mathcal{M}' := \begin{bmatrix} X & P & x \\ Z & Y & y \\ x & y & 1 \end{bmatrix},$$

with $X \in S^{m\times m}, Z \in \mathbb{R}^{n\times m}, Y \in S^{n\times n}, x \in \mathbb{R}^m, y \in \mathbb{R}^n$ and $P = Z^T$.

The SDP relaxation can then be expressed as

$$\begin{aligned}
&\min_{\mathcal{M}' \in \mathbb{S}^{m+n+1,m+n+1}} && 0 && \text{(SDP1)} \\
&\text{subject to} && \operatorname{Tr}(AZ) \geq e_i^T Ay, \forall i \in \{1,\ldots,m\}, && (3) \\
& && \operatorname{Tr}(BZ) \geq x^T Be_j, \forall j \in \{1,\ldots,n\}, && (4) \\
& && \sum_{i=1}^{m} x_i = 1, && (5) \\
& && \sum_{i=1}^{n} y_i = 1, && (6) \\
& && \mathcal{M}' \geq 0, && (7) \\
& && \mathcal{M}'_{m+n+1,m+n+1} = 1, && (8) \\
& && \mathcal{M}' \succeq 0. && (9)
\end{aligned}$$

We refer to the constraints (3) and (4) as the relaxed Nash constraints and the constraints (5) and (6) as the unity constraints. This SDP is motivated by the following observation.

Proposition 2.1. *Let $\mathcal{M}'$ be any rank-1 feasible solution to SDP1. Then the vectors x and y from its last column constitute a Nash equilibrium for the game (A, B).*

Proof. We know that x and y are in the simplex from the constraints (5), (6), and (7).
If the matrix $\mathcal{M}'$ is rank-1, then it takes the form

$$\begin{bmatrix} xx^T & xy^T & x \\ yx^T & yy^T & y \\ x^T & y^T & 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T. \tag{10}$$

Then, from the relaxed Nash constraints we have that

$$e_i^T Ay \leq \operatorname{Tr}(AZ) = \operatorname{Tr}(Ayx^T) = \operatorname{Tr}(x^T Ay) = x^T Ay,$$
$$x^T Ae_i \leq \operatorname{Tr}(BZ) = \operatorname{Tr}(Byx^T) = \operatorname{Tr}(x^T By) = x^T By.$$

The claim now follows from the characterization given in (1). □

Remark 2.1. Because a Nash equilibrium always exists, there will always be a matrix of the form (10) which is feasible to SDP1. Thus we can disregard any concerns about SDP1 being feasible, even when we add valid inequalities to it in Section 2.3.

Remark 2.2. It is intuitive to note that the submatrix $P = Z^T$ of the matrix $\mathcal{M}'$ corresponds to a probability distribution over the strategies, and that seeking a rank-1 solution to our SDP can be interpreted as making P a product distribution.

The following theorem shows that SDP1 is a weak relaxation and stresses the necessity of additional valid constraints.

Theorem 2.2. *Consider a bimatrix game with payoff matrices bounded in* $[0,1]$. *Then for any two vectors* $x \in \triangle_m$ *and* $y \in \triangle_n$, *there exists a feasible solution* $\mathcal{M}'$ *to SDP1 with* $\begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$ *as its last column.*

Proof. Consider any $x, y, \gamma > 0$, and the matrix

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T + \begin{bmatrix} \gamma J_{m+n,m+n} & 0_{m+n} \\ 0_{m+n}^T & 0 \end{bmatrix}.$$

This matrix is the sum of two nonnegative psd matrices and is hence nonnegative and psd. By assumption x and y are in the simplex, and so constraints (5) – (9) of SDP1 are satisfied. To check that constraints (3) and (4) hold, note that since A and B are nonnegative, as long as the matrices A and B are not the zero matrices, the quantities $\mathrm{Tr}(AZ)$ and $\mathrm{Tr}(BZ)$ will become arbitrarily large as γ increases. Since $e_i^T Ay$ and $x^T Be_i$ are bounded by 1 by assumption, we will have that constraints (3) and (4) hold for γ large enough. In the case where A or B is the zero matrix, the Nash constraints are trivially satisfied for the respective player. □

2.3 Valid Inequalities

In this subsection, we introduce a number of valid inequalities to improve upon the SDP relaxation in SDP1. These inequalities are justified by being valid if the matrix returned by the SDP is rank-1. The terminology we introduce here to refer to these constraints is used throughout the paper. Constraints (11) and (12) will be referred to as the *row inequalities*, and (13) and (14) will be referred to as the *correlated equilibrium inequalities*.

Proposition 2.3. *Any rank-1 solution* $\mathcal{M}'$ *to SDP1 must satisfy the following:*

$$\sum_{j=1}^{m} X_{i,j} = \sum_{j=1}^{n} P_{i,j} = x_i, \forall i \in \{1, \dots, m\}, \tag{11}$$

$$\sum_{j=1}^{n} Y_{i,j} = \sum_{j=1}^{m} Z_{i,j} = y_i, \forall i \in \{1, \dots, n\}. \tag{12}$$

$$\sum_{j=1}^{n} A_{i,j} P_{i,j} \geq \sum_{j=1}^{n} A_{k,j} P_{i,j}, \forall i, k \in \{1, \dots, m\}, \tag{13}$$

$$\sum_{j=1}^{m} B_{j,i} P_{j,i} \geq \sum_{j=1}^{m} B_{j,k} P_{j,i}, \forall i, k \in \{1, \dots, n\}. \tag{14}$$

Proof. Recall from (10) that if $\mathcal{M}'$ is rank-1, it is of the form

$$\begin{bmatrix} xx^T & xy^T & x \\ yx^T & yy^T & y \\ x^T & y^T & 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T.$$

To show (11), observe that

$$\sum_{j=1}^{m} X_{i,j} = \sum_{j=1}^{m} x_i x_j = x_i, \forall i \in \{1, \dots, m\}.$$

An identical argument works for the remaining matrices $P, Z,$ and Y. To show (13) and (14), observe that a pair (x, y) is a Nash equilibrium if and only if

$$\forall i, x_i > 0 \Rightarrow e_i^T Ay = x^T Ay = \max_i e_i^T Ay,$$

$$\forall i, y_i > 0 \Rightarrow x^T Be_i = x^T By = \max_i x^T Be_i.$$

This is because the Nash conditions require that $x^T Ay$, a convex combination of $e_i^T Ay$, be at least $e_i^T Ay$ for all i. Indeed, if $x_i > 0$ but $e_i^T Ay < x^T Ay$, the convex combination must be less than $\max_i x^T Ay$.

For each i such that $x_i = 0$ or $y_i = 0$, inequalities (13) and (14) reduce to $0 \geq 0$, so we only need to consider strategies played with positive probability. Observe that if $\mathcal{M}'$ is rank-1, then

$$\sum_{j=1}^{n} A_{i,j} P_{i,j} = x_i \sum_{j=1}^{n} A_{i,j} y_j = x_i e_i^T Ay \geq x_i e_k^T Ay = \sum_{j=1}^{n} A_{k,j} P_{i,j}, \forall i, k$$

$$\sum_{j=1}^{m} B_{j,i} P_{j,i} = y_i \sum_{j=1}^{m} B_{j,i} x_j = y_i x^T Be_i \geq y_i x^T Be_k = \sum_{j=1}^{m} B_{j,i} P_{j,k}, \forall i, k.$$

□

Remark 2.3. There are two ways to interpret the inequalities in (13) and (14): the first is as a relaxation of the constraint $x_i(e_i^T Ay - e_j^T Ay) \geq 0, \forall i, j$, which must hold since any strategy played with positive probability must give the best response payoff. The other interpretation is to have the distribution over outcomes defined by P be a correlated equilibrium [4]. This can be imposed by a set of linear constraints on the entries of P as explained next.

Suppose the players have access to a public randomization device which prescribes a pure strategy to each of them (unknown to the other player). The distribution over the assignments can be given by a matrix P, where $P_{i,j}$ is the probability that strategy i is assigned to player A and strategy j is assigned to player B. This distribution is a correlated equilibrium if both players have no incentive to deviate from the strategy prescribed, that is, if the prescribed pure strategies a and b satisfy

$$\sum_{j=1}^{n} A_{i,j} Prob(b = j | a = i) \geq \sum_{j=1}^{n} A_{k,j} Prob(b = j | a = i),$$

$$\sum_{i=1}^{m} B_{i,j} Prob(a = i | b = j) \geq \sum_{i=1}^{m} B_{i,k} Prob(a = i | b = j).$$

If we interpret the P submatrix in our SDP as the distribution over the assignments by the public device, then because of our row constraints, $Prob(b = j | a = i) = \frac{P_{i,j}}{x_i}$ whenever $x_i \neq 0$ (otherwise the above inequalities are trivial). Similarly, $P(a = i | b = j) = \frac{P_{i,j}}{y_j}$ for nonzero y_j. Observe now that the above two inequalities imply (13) and (14). Finally, note that every Nash equilibrium generates a correlated equilibrium, since if P is a product distribution given by xy^T, then $Prob(b = j | a = i) = y_j$ and $P(a = i | b = j) = x_i$.

2.3.1 Implied Inequalities

In addition to those explicitly mentioned in the previous section, there are other natural valid inequalities which are omitted because they are implied by the ones we have already proposed. We give two examples of such inequalities in the next proposition. We refer to the constraints in (15) below as the *distribution constraints*. The constraints in (16) are the familiar McCormick inequalities [34] for box-constrained quadratic programming.

Proposition 2.4. *Let* $z := \begin{bmatrix} x \\ y \end{bmatrix}$. *Any rank-1 solution* $\mathcal{M}'$ *to SDP1 must satisfy the following:*

$$\sum_{i=1}^{m}\sum_{j=1}^{m} X_{i,j} = \sum_{i=1}^{n}\sum_{j=1}^{m} Z_{i,j} = \sum_{i=1}^{n}\sum_{j=1}^{n} Y_{i,j} = 1. \tag{15}$$

$$\begin{aligned} \mathcal{M}_{i,j} &\leq z_i, \forall i,j \in \{1,\dots,m+n\}, \\ \mathcal{M}_{i,j} + 1 &\geq z_i + z_j, \forall i,j \in \{1,\dots,m+n\}. \end{aligned} \tag{16}$$

Proof. The distribution constraints follow immediately from the row constraints (11) and (12), along with the unity constraints (5) and (6).

The first McCormick inequality is immediate as a consequence of (11) and (12), as all entries of $\mathcal{M}$ are nonnegative. To see why the second inequality holds, consider whichever submatrix X, Y, P, or Z that contains $\mathcal{M}_{i,j}$. Suppose that this submatrix is, e.g., P. Then, since P is nonnegative,

$$0 \leq \sum_{k=1,k\neq i}^{m} \sum_{l=1,l\neq j}^{n} P_{k,l} \overset{(11)}{=} \sum_{k=1,k\neq i}^{m} (x_k - P_{k,j}) \overset{(12)}{=} (1 - x_i) - (y_j - P_{i,j}) = P_{i,j} + 1 - x_i - y_j.$$

The same argument holds for the other submatrices, and this concludes the proof. □

2.4 Simplifying our SDP

We observe that the row constraints (11) and (12) along with the correlated equilibrium constraints (13) and (14) imply the relaxed Nash constraints (3) and (4). Indeed, if we fix an index $k \in \{1,\dots,m\}$, then

$$\mathrm{Tr}(AZ) = \sum_{i=1}^{m}\sum_{j=1}^{n} A_{i,j}P_{i,j} \overset{(13)}{\geq} \sum_{i=1}^{m}\sum_{j=1}^{n} A_{k,j}P_{i,j} \geq \sum_{j=1}^{n} A_{k,j}(\sum_{i=1}^{m} P_{i,j}) \overset{(12),P=Z^T}{\geq} \sum_{j=1}^{n} A_{k,j}y_j = e_k^T Ay.$$

The proof for player B proceeds identically. Then, after collecting the valid inequalities and removing the relaxed Nash constraints, we arrive at an SDP given by

$$\begin{aligned} \min_{\mathcal{M}' \in S^{(m+n+1)\times(m+n+1)}} \quad & 0 \\ \text{subject to} \quad & (5)-(9), (11)-(14). \end{aligned} \tag{SDP1'}$$

We make the observation that the last row and column of $\mathcal{M}'$ can be removed from this SDP, that is, there is a one-to-one correspondence between solutions to SDP1' and those to the following SDP (where $\mathcal{M} := \begin{bmatrix} X & P \\ Z & Y \end{bmatrix}$, with $P = Z^T$):

$$\min_{\mathcal{M}\in S^{(m+n)\times(m+n)}} \quad 0 \tag{SDP2}$$

subject to

$$\mathcal{M} \succeq 0, \tag{17}$$

$$\mathcal{M} \geq 0, \tag{18}$$

$$\sum_{i=1}^{n}\sum_{j=1}^{n} P_{i,j} = 1, \tag{19}$$

$$\sum_{j=1}^{m} X_{i,j} = \sum_{j=1}^{n} P_{i,j}, \forall i \in \{1,\dots,m\}, \tag{20}$$

$$\sum_{j=1}^{n} Y_{i,j} = \sum_{j=1}^{m} Z_{i,j}, \forall i \in \{1,\dots,n\}, \tag{21}$$

$$\sum_{j=1}^{n} A_{i,j}P_{i,j} \geq \sum_{j=1}^{n} A_{k,j}P_{i,j}, \forall i,k \in \{1,\dots,m\}, \tag{22}$$

$$\sum_{j=1}^{m} B_{j,i}P_{j,i} \geq \sum_{j=1}^{m} B_{j,k}P_{j,i}, \forall i,k \in \{1,\dots,n\}. \tag{23}$$

Indeed, it is readily verified that the submatrix $\mathcal{M}$ from any feasible solution $\mathcal{M}'$ to SDP1' is feasible to SDP2. Conversely, let $\mathcal{M}$ be any feasible matrix to SDP2. Consider an eigendecomposition $\mathcal{M} = \sum_{i=1}^{k} \lambda_i v_i v_i^T$ and let $\begin{bmatrix} x \\ y \end{bmatrix} := \mathcal{M}\frac{1_{m+n}}{2}$. Then the matrix

$$\mathcal{M}' := \begin{bmatrix} \mathcal{M} & \begin{bmatrix} x \\ y \end{bmatrix} \\ \begin{bmatrix} x^T & y^T \end{bmatrix} & 1 \end{bmatrix} = \sum_{i=1}^{k} \lambda_i \begin{bmatrix} v_i \\ 1_{m+n}^T v_i/2 \end{bmatrix} \begin{bmatrix} v_i \\ 1_{m+n}^T v_i/2 \end{bmatrix}^T \tag{24}$$

is easily seen to be feasible to SDP1'.

Given any feasible solution $\mathcal{M}$ to SDP2, observe that the submatrix P is a correlated equilibrium. We take our candidate approximate Nash equilibrium to be the pair $x = P1_n$ and $y = P^T 1_m$. If the correlated equilibrium P is rank-1, then the pair (x, y) so defined constitutes an exact Nash equilibrium. In Section 4, we will add certain objective functions to SDP2 with the interpretation of searching for low-rank correlated equilibria.

3 Exactness for Strictly Competitive Games

In this section, we show that SDP1 recovers a Nash equilibrium for any zero-sum game, and that SDP2 recovers a Nash equilibrium for any strictly competitive game (see Definition 3.3 below). Both these notions represent games where the two players are in direct competition, but strictly competitive games are more general, and for example, allow both players to have nonnegative payoff matrices. These classes of games are solvable in polynomial time via linear programming. Nonetheless, it is reassuring to know that our SDPs recover these important special cases.

Definition 3.1. *A* zero-sum game *is a game in which the payoff matrices satisfy $A = -B$.*

Theorem 3.2. *For a zero-sum game, the vectors x and y from the last column of any feasible solution $\mathcal{M}'$ to SDP1 constitute a Nash equilibrium.*

Proof. Recall that the relaxed Nash constraints (3) and (4) read

$$\text{Tr}(AZ) \geq e_i^T Ay, \forall i \in \{1, \ldots, m\},$$

$$\text{Tr}(BZ) \geq x^T Be_j, \forall j \in \{1, \ldots, n\}.$$

Since $B = -A$, the latter statement is equivalent to

$$\text{Tr}(AZ) \leq x^T Ae_j, \forall j \in \{1, \ldots, n\}.$$

In conjunction these imply

$$e_i^T Ay \leq \text{Tr}(AZ) \leq x^T Ae_j, \forall i \in \{1, \ldots, m\}, j \in \{1, \ldots, n\}. \tag{25}$$

We claim that any pair $x \in \triangle_m$ and $y \in \triangle_n$ which satisfies the above condition is a Nash equilibrium. To see that $x^T Ay \geq e_i^T Ay, \forall i \in \{1, \ldots, m\}$, observe that $x^T Ay$ is a convex combination of $x^T Ae_j$, which are at least $e_i^T Ay$ by (25). To see that $x^T By \geq x^T Be_j \Leftrightarrow x^T Ay \leq x^T Ae_j, \forall j \in \{1, \ldots, n\}$, observe that $x^T Ay$ is a convex combination of $e_i^T Ay$, which are at most $x^T Ae_j$ by (25). □

Definition 3.3. *A game (A, B) is* strictly competitive *if for all $x, x' \in \triangle_m, y, y' \in \triangle_n$, $x^T Ay - x'^T Ay'$ and $x'^T By' - x^T By$ have the same sign.*

The interpretation of this definition is that if one player benefits from changing from one outcome to another, the other player must suffer. Adler, Daskalakis, and Papadimitriou show in [3] that the following much simpler characterization is equivalent.

Theorem 3.4 (Theorem 1 of [3])**.** *A game is strictly competitive if and only if there exist scalars c, d, e, and f, with $c > 0, e > 0$, such that $cA + dJ_{m\times n} = -eB + fJ_{m\times n}$.*

One can easily show that there exist strictly competitive games for which not all feasible solutions to SDP1 have Nash equilibria as their last columns (see Theorem 2.2). However, we show that this is the case for SDP2.

Theorem 3.5. *For a strictly competitive game, the vectors $x := P1_n$ and $y := P^T 1_m$ from any feasible solution $\mathcal{M}$ to SDP2 constitute a Nash equilibrium.*

To prove Theorem 3.5 we need the following lemma, which shows that feasibility of a matrix $\mathcal{M}$ in SDP2 is invariant under certain transformations of A and B.

Lemma 3.6. *Let c, d, e, and f be any set of scalars with $c > 0$ and $e > 0$. If a matrix $\mathcal{M}$ is feasible to SDP2 with input payoff matrices A and B, then it is also feasible to SDP2 with input matrices $cA + dJ_{m\times n}$ and $eB + fJ_{m\times n}$.*

Proof. It suffices to check that constraints (22) and (23) of SDP2 still hold, as only the correlated equilibrium constraints use the matrices A and B. We only show that constraint (22) still holds because the argument for constraint (23) is identical.

Note from the definition of x that for each $i \in \{1, \ldots, m\}, x_i = \sum_{j=1}^n (J_{m\times n})_{i,j} P_{i,j}$. To check that the correlated equilibrium constraints hold, observe that for scalars $c > 0, d$, and for all

$i, k \in \{1, \ldots, m\}$,

$$\sum_{j=1}^{n} A_{i,j} P_{i,j} \geq \sum_{j=1}^{n} A_{k,j} P_{i,j}$$
$$\Leftrightarrow c \sum_{j=1}^{n} A_{i,j} P_{i,j} + d \sum_{j=1}^{n} P_{i,j} \geq c \sum_{j=1}^{n} A_{k,j} P_{i,j} + d \sum_{j=1}^{n} P_{i,j}$$
$$\Leftrightarrow c \sum_{j=1}^{n} A_{i,j} P_{i,j} + d \sum_{j=1}^{n} (J_{m\times n})_{i,j} P_{i,j} \geq c \sum_{j=1}^{n} A_{k,j} P_{i,j} + d \sum_{j=1}^{n} (J_{m\times n})_{k,j} P_{i,j}$$
$$\Leftrightarrow \sum_{j=1}^{n} (cA_{i,j} + dJ_{m\times n})_{k,j} P_{i,j} \geq \sum_{j=1}^{n} (cA_{i,j} + dJ_{m\times n})_{k,j} P_{i,j}.$$

□

Proof (of Theorem 3.5). Let A and B be the payoff matrices of the given strictly competitive game and let $\mathcal{M}$ be a feasible solution to SDP2. Since the game is strictly competitive, we know from Theorem 3.4 that $cA + dJ_{m\times n} = -eB + fJ_{m\times n}$ for some scalars $c > 0, e > 0, d, f$. Consider a new game with input matrices $\tilde{A} = cA + dJ_{m\times n}$ and $\tilde{B} = eB - fJ_{m\times n}$. By Lemma 3.6, $\mathcal{M}$ is still feasible to SDP2 with input matrices $\tilde{A}$ and $\tilde{B}$. By the arguments in Section 2.4, the matrix $\mathcal{M}' := \begin{bmatrix} \mathcal{M} & \begin{bmatrix} x \\ y \end{bmatrix} \\ \begin{bmatrix} x^T & y^T \end{bmatrix} & 1 \end{bmatrix}$ is feasible to SDP1', and hence also to SDP1. Now notice that since $\tilde{A} = -\tilde{B}$, Theorem 3.2 implies that the vectors x and y in the last column form a Nash equilibrium to the game $(\tilde{A}, \tilde{B})$. Finally recall from the arguments at the beginning of Section 2 that Nash equilibria are invariant to scaling and shifting of the payoff matrices, and hence (x, y) is a Nash equilibrium to the game (A, B). □

4 Algorithms for Lowering Rank

In this section, we present heuristics which aim to find low-rank solutions to SDP2 and present some empirical results. Recall that our SDP2 in Section 2.4 did not have an objective function. Hence, we can encourage low-rank solutions by choosing certain objective functions, in particular the trace of the matrix $\mathcal{M}$, which is a general heuristic for minimizing the rank of symmetric matrices [40, 19]. This simple objective function is already guaranteed to produce a rank-1 solution in the case of strictly competitive games (see Proposition 4.1 below). For general games, however, one can design better objective functions in an iterative fashion (see Section 4.1).

Notational Remark: For the remainder of this section, we will use the shorthand $x := P1_n$ and $y := P^T 1_m$, where P is the upper right submatrix of a feasible solution $\mathcal{M}$ to SDP2.

Proposition 4.1. *For a strictly competitive game, any optimal solution to SDP2 with* $\text{Tr}(\mathcal{M})$ *as the objective function must be rank-1.*

Proof. Let

$$\mathcal{M} := \begin{bmatrix} X & P \\ P^T & Y \end{bmatrix}$$

be a feasible solution to SDP2. In the case of strictly competitive games, from Theorem 3.5 we know that that (x, y) is a Nash equilibrium. Then because the matrix $\mathcal{M}$ is psd, from (24) and an application of the Schur complement (see, e.g. [7, Sect. A.5.5]) to $\begin{bmatrix} \mathcal{M} & \begin{bmatrix} x \\ y \end{bmatrix} \\ \begin{bmatrix} x^T & y^T \end{bmatrix} & 1 \end{bmatrix}$, we have that

$\mathcal{M} \succeq \begin{bmatrix} x \\ y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}^T$. Hence, $\mathcal{M} = \begin{bmatrix} xx^T & xy^T \\ yx^T & yy^T \end{bmatrix} + \mathcal{P}$ for some psd matrix $\mathcal{P}$ and the Nash equilibrium (x, y). Given this expression, the objective function $\text{Tr}(\mathcal{M})$ is then $x^T x + y^T y + \text{Tr}(\mathcal{P})$. As (x, y) is a Nash equilibrium, the choice of $\mathcal{P} = 0$ results in a feasible solution. Since the zero matrix has the minimum possible trace among all psd matrices, the solution will be the rank-1 matrix $\begin{bmatrix} x \\ y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}^T$. □

Remark 4.1. If the row constraints and the nonnegativity constraints on X and Y are removed from SDP2, then this SDP with $\text{Tr}(\mathcal{M})$ as the objective function can be interpreted as searching for a minimum-rank correlated equilibrium P via the nuclear norm relaxation; see [40, Section 2].

4.1 Linearization Algorithms

The algorithms we present in this section for minimzing the rank of the matrix $\mathcal{M}$ in SDP2 are based on iterative linearization of certain nonconvex objective functions. Motivated by the next proposition, we design two continuous (nonconvex) objective functions that, if minimized exactly, would guarantee rank-1 solutions. We will then linearize these functions iteratively.

Proposition 4.2. *Let the matrices X and Y and vectors $x := P1_n$ and $y := P^T 1_m$ be taken from a feasible solution to SDP2. Then the matrix $\mathcal{M}$ is rank-1 if and only if $X_{i,i} = x_i^2$ and $Y_{i,i} = y_i^2$ for all i.*

Proof. Note that if $\mathcal{M}$ is rank-1, then it can be written as zz^T for some $z \in \mathbb{R}^{m+n}$. The i-th diagonal entry in the X submatrix will then be equal to

$$z_i^2 \overset{(15)}{=} \frac{1}{4} z_i^2 (1_{m+n}^T z z^T 1_{m+n}) = (\frac{1}{2} \mathcal{M}_{i,} 1_{m+n})^2 \overset{(11)}{=} (P_{i,} 1_n)^2 = x_i^2,$$

where the second equality holds because $\mathcal{M}_{i,}$—the i-th row of $\mathcal{M}$—is $z_i z^T$. An analogous statement holds for the diagonal entries of Y, and hence the condition is necessary.

To show sufficiency, let $z := \begin{bmatrix} x \\ y \end{bmatrix}$. Since $\mathcal{M}$ is psd, we have that $\mathcal{M}_{i,j} \leq \sqrt{\mathcal{M}_{i,i} \mathcal{M}_{j,j}}$, which implies $\mathcal{M}_{i,j} \leq z_i z_j$ by the assumption of the proposition. Recall from the distribution constraint (15) that $\sum_{i=1}^{m+n} \sum_{j=1}^{m+n} \mathcal{M}_{i,j} = 4$. Further, the same constraint along with the definitions of x and y imply that $\sum_{i=1}^{m+n} z_i = 2$, which means that $\sum_{i=1}^{m+n} \sum_{j=1}^{m+n} z_i z_j = 4$. Hence in order to have the equality

$$4 = \sum_{i=1}^{m+n} \sum_{j=1}^{m+n} \mathcal{M}_{i,j} \leq \sum_{i=1}^{m+n} \sum_{j=1}^{m+n} z_i z_j = 4,$$

we must have $\mathcal{M}_{i,j} = z_i z_j$ for each i and j. Consequently $\mathcal{M}$ is rank-1. □

We focus now on two nonconvex objectives that as a consequence of the above proposition would return rank-1 solutions:

Proposition 4.3. *All optimal solutions to SDP2 with the objective function $\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}}$ or $\text{Tr}(\mathcal{M}) - x^T x - y^T y$ are rank-1.*

Proof. We show that each of these objectives has a specific lower bound which is achieved if and only if the matrix is rank-1.

Observe that since $\mathcal{M} \succeq \begin{bmatrix} x \\ y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}^T$, we have $\sqrt{X_{i,i}} \geq x_i$ and $\sqrt{Y_{i,i}} \geq y_i$, and hence

$$\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}} \geq \sum_{i=1}^{m} x_i + \sum_{i=1}^{n} y_i = 2.$$

Further note that

$$\text{Tr}(\mathcal{M}) - \begin{bmatrix} x \\ y \end{bmatrix}^T \begin{bmatrix} x \\ y \end{bmatrix} \geq \begin{bmatrix} x \\ y \end{bmatrix}^T \begin{bmatrix} x \\ y \end{bmatrix} - \begin{bmatrix} x \\ y \end{bmatrix}^T \begin{bmatrix} x \\ y \end{bmatrix} = 0.$$

We can see that the lower bounds are achieved if and only if $X_{i,i} = x_i^2$ and $Y_{i,i} = y_i^2$ for all i, which by Proposition 4.2 happens if and only if $\mathcal{M}$ is rank-1. □

We refer to our two objective functions in Proposition 4.3 as the "*square root objective*" and the "*diagonal gap objective*" respectively. While these are both nonconvex, we will attempt to iteratively minimize them by linearizing them through a first order Taylor expansion. For example, at iteration k of the algorithm,

$$\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k)}} \simeq \sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k-1)}} + \frac{1}{2\sqrt{\mathcal{M}_{i,i}^{(k-1)}}}(\mathcal{M}_{i,i}^{(k)} - \mathcal{M}_{i,i}^{(k-1)}).$$

Note that for the purposes of minimization, this reduces to minimizing $\sum_{i=1}^{m+n} \frac{1}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}} \mathcal{M}_{i,i}^{(k)}$.

In similar fashion, for the second objective function, at iteration k we can make the approximation

$$\text{Tr}(\mathcal{M}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)} \simeq \text{Tr}(\mathcal{M}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} (\begin{bmatrix} x \\ y \end{bmatrix}^{(k)} - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)}).$$

Once again, for the purposes of minimization this reduces to minimizing $\text{Tr}(\mathcal{M}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)}$. This approach then leads to the following two algorithms.[4]

Algorithm 1 Square Root Minimization Algorithm

Let $x^{(0)} = 1_m, y^{(0)} = 1_n, k = 1$.
while !convergence **do**
    Solve SDP2 with $\sum_{i=1}^{m} \frac{1}{\sqrt{x_i^{(k-1)}}} X_{i,i} + \sum_{i=1}^{n} \frac{1}{\sqrt{y_i^{(k-1)}}} Y_{i,i}$ as the objective, and let $\mathcal{M}^*$ be an optimal solution.
    Let $x^{(k)} = diag(X^*), y^{(k)} = diag(Y^*)$.
    Let $k = k + 1$.
end while

---

[4] An algorithm similar to Algorithm 2 is used in [24].

Algorithm 2 Diagonal Gap Minimization Algorithm

```
Let $x^{(0)} = 0_m, y^{(0)} = 0_n, k = 1$.
while !convergence do
    Solve SDP2 with $\mathrm{Tr}(X) + \mathrm{Tr}(Y) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)}$ as the objective, and let $\mathcal{M}^*$ be an optimal solution.
    Let $x^{(k)} = P^* 1_n, y^{(k)} = P^{*T} 1_m$.
    Let $k = k + 1$.
end while
```

Remark 4.2. Note that the first iteration of both algorithms uses the nuclear norm (i.e. trace) of $\mathcal{M}$ as the objective.

The square root algorithm has the following property.

Theorem 4.4. *Let $\mathcal{M}^{(1)}, \mathcal{M}^{(2)}, \dots$ be the sequence of optimal matrices obtained from the square root algorithm. Then the sequence*

$$\{\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k)}}\} \tag{26}$$

is nonincreasing and is lower bounded by two. If it reaches two at some iteration t, then the matrix $\mathcal{M}^{(t)}$ is rank-1.

Proof. Observe that for any $k > 1$,

$$\sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k)}} \leq \frac{1}{2} \sum_{i=1}^{m+n} (\frac{\mathcal{M}_{i,i}^{(k)}}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}} + \sqrt{\mathcal{M}_{i,i}^{(k-1)}}) \leq \frac{1}{2} \sum_{i=1}^{m+n} (\frac{\mathcal{M}_{i,i}^{(k-1)}}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}} + \sqrt{\mathcal{M}_{i,i}^{(k-1)}}) = \sum_{i=1}^{m+n} \sqrt{\mathcal{M}_{i,i}^{(k-1)}},$$

where the first inequality follows from the arithmetic-mean-geometric-mean inequality, and the second follows from that $\mathcal{M}_{i,i}^{(k)}$ is chosen to minimize $\sum_{i=1}^{m+n} \frac{\mathcal{M}_{i,i}^{(k)}}{\sqrt{\mathcal{M}_{i,i}^{(k-1)}}}$ and hence achieves a no larger value than the feasible solution $\mathcal{M}^{(k-1)}$. This shows that the sequence is nonincreasing.

The proof of Proposition 4.3 already shows that the sequence is lower bounded by two, and Proposition 4.3 itself shows that reaching two is sufficient to have the matrix be rank-1. □

The diagonal gap algorithm has the following property.

Theorem 4.5. *Let $\mathcal{M}^{(1)}, \mathcal{M}^{(2)}, \dots$ be the sequence of optimal matrices obtained from the diagonal gap algorithm. Then the sequence*

$$\{\mathrm{Tr}(\mathcal{M}^{(k)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)}\} \tag{27}$$

is nonincreasing and is lower bounded by zero. If it reaches zero at some iteration t, then the matrix $\mathcal{M}^{(t)}$ is rank-1.

Proof. Observe that

$$\begin{aligned}
\mathrm{Tr}(\mathcal{M}^{(k)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)} &\leq \mathrm{Tr}(\mathcal{M}^{(k)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k)} + (\begin{bmatrix} x \\ y \end{bmatrix}^{(k)} - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)})^T(\begin{bmatrix} x \\ y \end{bmatrix}^{(k)} - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)}) \\
&= \mathrm{Tr}(\mathcal{M}^{(k)}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} + \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} \\
&\leq \mathrm{Tr}(\mathcal{M}^{(k-1)}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} + \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)} \\
&= \mathrm{Tr}(\mathcal{M}^{(k-1)}) - \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)},
\end{aligned}$$

where the second inequality follows from that $\mathcal{M}^{(k)}$ is chosen to minimize

$$\mathrm{Tr}(\mathcal{M}^{(k-1)}) - 2 \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)T} \begin{bmatrix} x \\ y \end{bmatrix}^{(k-1)}$$

and hence achieves a no larger value than the feasible solution $\mathcal{M}^{(k-1)}$. This shows that the sequence is nonincreasing.

The proof of Proposition 4.3 already shows that the sequence is lower bounded by zero, and Proposition 4.3 itself shows that reaching zero is sufficient to have the matrix be rank-1. □

We also invite the reader to also see Theorem 5.6 in the next section which relates the objective value of the diagonal gap minimization algorithm and the quality of approximate Nash equilibria that the algorithm produces.

4.2 Numerical Experiments

We tested Algorithms 1 and 2 on games coming from 100 randomly generated payoff matrices with entries bounded in $[0, 1]$ of varying sizes. Below is a table of statistics for 20×20 matrices; the data for the rest of the sizes can be found in Appendix A.[5] We can see that our algorithms return approximate Nash equilibria with fairly low ϵ (recall the definition from Section 2.1). We ran 20 iterations of each algorithm on each game. Using the SDP solver of MOSEK [1], each iteration takes on average under 4 seconds to solve on a standard personal machine with a 3.4 GHz processor and 16 GB of memory.

Table 1: Statistics on ϵ for 20×20 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0198	0.0046	0.0039	0.0034
Diagonal Gap	0.0159	0.0032	0.0024	0.0032

The histograms below show the effect of increasing the number of iterations on lowering ϵ on 20×20 games. For both algorithms, there was a clear improvement of the ϵ by increasing the number of iterations.

[5] The code that produced these results is publicly available at `aaa.princeton.edu/software`. The function `nash.m` computes an approximate Nash equilibrium using one of our two algorithms as specified by the user.

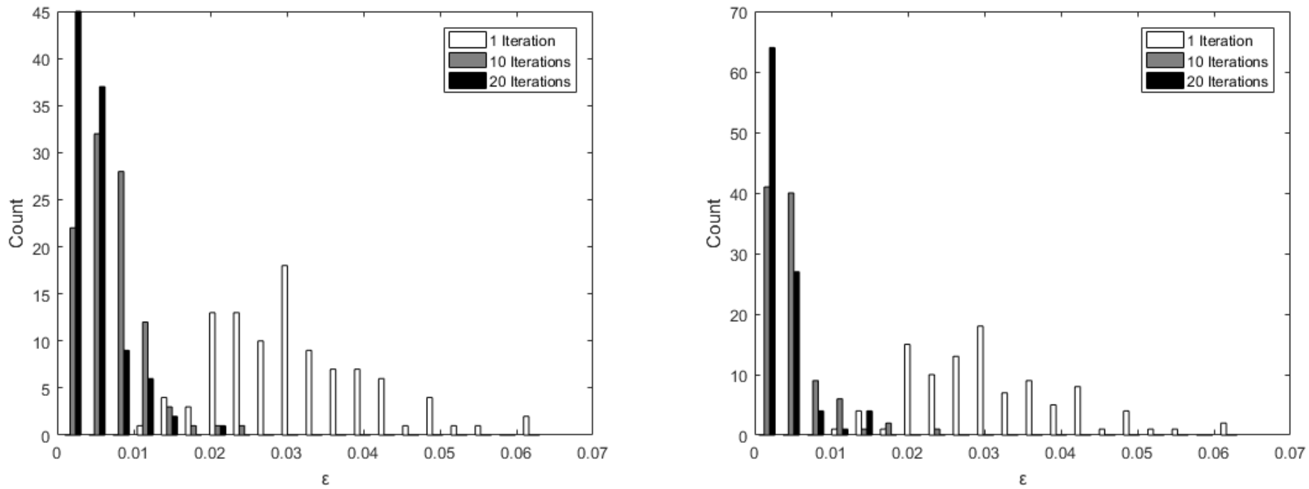

Figure 1: Distribution of ϵ over numbers of iterations for the square root algorithm (left) and the diagonal gap algorithm (right).

5 Bounds on ϵ for General Games

Since the problem of computing a Nash equilibrium to an arbitrary bimatrix game is PPAD-complete, it is unlikely that one can find rank-1 solutions to this SDP in polynomial time. In Section 4, we designed objective functions (such as variations of the nuclear norm) that empirically do very well in finding low-rank solutions to SDP2. Nevertheless, it is of interest to know if the solution returned by SDP2 is not rank-1, whether one can recover an ϵ-Nash equilibrium from it and have a guarantee on ϵ. Our goal in this section is to study this question.

Notational Remark: Recall our notation for the matrix

$$\mathcal{M} := \begin{bmatrix} X & P \\ Z & Y \end{bmatrix}.$$

Throughout this section, any matrices $X, Z, P = Z^T$ and Y are assumed to be taken from a feasible solution to SDP2. Furthermore, x and y will be $P1_n$ and $P^T 1_m$ respectively.

The ultimate results of this section are the theorems in Sections 5.2 and 5.3. To work towards them, we need a number of preliminary lemmas which we present in Section 5.1.

5.1 Lemmas Towards Bounds on ϵ

We first observe the following connection between the approximate payoffs $\text{Tr}(AZ)$ and $\text{Tr}(BZ)$, and $\epsilon(x, y)$, as defined in Section 2.1.

Lemma 5.1. *Consider any feasible solution to SDP2. Then*

$$\epsilon(x, y) \leq \max\{\text{Tr}(AZ) - x^T Ay, \text{Tr}(BZ) - x^T By\}.$$

Proof. Recall from the argument at the beginning of Section 2.4 that constraints (13) and (14) imply $\text{Tr}(AZ) \geq e_i^T Ay$ and $\text{Tr}(BZ) \geq x^T Be_i$ for all i. Hence, we have $\epsilon_A \leq \text{Tr}(AZ) - x^T Ay$ and $\epsilon_B \leq \text{Tr}(BZ) - x^T By$. □

We thus are interested in the difference of the two matrices $P = Z^T$ and xy^T. These two matrices can be interpreted as two different probability distributions over the strategy outcomes.

The matrix P is the probability distribution from the SDP which generates the approximate payoffs $\text{Tr}(AZ)$ and $\text{Tr}(BZ)$, while xy^T is the product distribution that would have resulted if the matrix had been rank-1. We will see that the difference of these distributions is key in studying the ϵ which results from SDP2. Hence, we first take steps to represent this difference.

Lemma 5.2. *Consider any feasible matrix $\mathcal{M}$ to SDP2 with an eigendecomposition*

$$\mathcal{M} = \sum_{i=1}^{k} \lambda_i v_i v_i^T =: \sum_{i=1}^{k} \lambda_i \begin{bmatrix} a_i \\ b_i \end{bmatrix} \begin{bmatrix} a_i \\ b_i \end{bmatrix}^T, \tag{28}$$

so that the eigenvectors $v_i \in \mathbb{R}^{m+n}$ are partitioned into vectors $a_i \in \mathbb{R}^m$ and $b_i \in \mathbb{R}^n$. Then for all i, $\sum_{j=1}^m (a_i)_j = \sum_{j=1}^n (b_i)_j$.

Proof. We know from (19), (20), and (21) that

$$\sum_{i=1}^{k} \lambda_i 1_m^T a_i a_i^T 1_m \overset{(19),(20)}{=} 1, \tag{29}$$

$$\sum_{i=1}^{k} \lambda_i 1_m^T a_i b_i^T 1_n \overset{(19)}{=} 1, \tag{30}$$

$$\sum_{i=1}^{k} \lambda_i 1_n^T b_i a_i^T 1_m \overset{(19)}{=} 1, \tag{31}$$

$$\sum_{i=1}^{k} \lambda_i 1_n^T b_i b_i^T 1_n \overset{(19),(21)}{=} 1. \tag{32}$$

Then by subtracting terms we have

$$(29) - (30) = \sum_{i=1}^{k} \lambda_i 1_m^T a_i (a_i^T 1_m - b_i^T 1_n) = 0, \tag{33}$$

$$(31) - (32) = \sum_{i=1}^{k} \lambda_i 1_n^T b_i (a_i^T 1_m - b_i^T 1_n) = 0. \tag{34}$$

By subtracting again these imply

$$(33) - (34) = \sum_{i=1}^{k} \lambda_i (1_m^T a_i - 1_n^T b_i)^2 = 0. \tag{35}$$

As all λ_i are nonnegative due to positive semidefiniteness of $\mathcal{M}$, the only way for this equality to hold is to have $1_m^T a_i = 1_n^T b_i, \forall i$. This is equivalent to the statement of the claim. $\square$

From Lemma 5.2, we can let $s_i := \sum_{j=1}^m (a_i)_j = \sum_{j=1}^n (b_i)_j$, and furthermore we assume without loss of generality that each s_i is nonnegative. Note that from the definition of x we have

$$x_i = \sum_{j=1}^{m} P_{ij} = \sum_{l=1}^{k} \sum_{j=1}^{m} \lambda_l (a_l)_i (b_l)_j = \sum_{j=1}^{k} \lambda_j s_j (a_l)_i. \tag{36}$$

Hence,

$$x = \sum_{i=1}^{k} \lambda_i s_i a_i. \tag{37}$$

Similarly,

$$y = \sum_{i=1}^{k} \lambda_i s_i b_i. \tag{38}$$

Finally note from the distribution constraint (15) that this implies

$$\sum_{i=1}^{k} \lambda_i s_i^2 = 1. \tag{39}$$

Lemma 5.3. *Let*

$$\mathcal{M} = \sum_{i=1}^{k} \lambda_i \begin{bmatrix} a_i \\ b_i \end{bmatrix} \begin{bmatrix} a_i \\ b_i \end{bmatrix}^T,$$

be a feasible solution to SDP2, such that the eigenvectors of $\mathcal{M}$ *are partitioned into* a_i *and* b_i *with* $\sum_{j=1}^{m}(a_i)_j = \sum_{j=1}^{n}(b_i)_j = s_i, \forall i$. *Then*

$$P - xy^T = \sum_{i=1}^{k} \sum_{j>i}^{k} \lambda_i \lambda_j (s_j a_i - s_i a_j)(s_j b_i - s_i b_j)^T.$$

Proof. Using equations (37) and (38) we can write

$$\begin{aligned}
P - xy^T &= \sum_{i=1}^{k} \lambda_i a_i b_i^T - (\sum_{i=1}^{k} \lambda_i s_i a_i)(\sum_{j=1}^{k} \lambda_j s_j b_j)^T \\
&= \sum_{i=1}^{k} \lambda_i a_i (b_i - s_i \sum_{j=1}^{k} \lambda_j s_j b_j)^T \\
&\overset{(39)}{=} \sum_{i=1}^{k} \lambda_i a_i (\sum_{j=1}^{k} \lambda_j s_j^2 b_i - s_i \sum_{j=1}^{k} \lambda_j s_j b_j)^T \\
&= \sum_{i=1}^{k} \sum_{j=1}^{k} \lambda_i \lambda_j a_i s_j (s_j b_i - s_i b_j)^T \\
&= \sum_{i=1}^{k} \sum_{j>i}^{k} \lambda_i \lambda_j (s_j a_i - s_i a_j)(s_j b_i - s_i b_j)^T,
\end{aligned}$$

where the last line follows from observing that terms where i and j are switched can be combined. □

We can relate ϵ and $P - xy^T$ with the following lemma.

Lemma 5.4. *Let the matrix* P *and the vectors* $x := P1_n$ *and* $y := P^T 1_m$ *come from any feasible solution to SDP2. Then*

$$\epsilon \leq \frac{\|P - xy^T\|_1}{2},$$

where $\|\cdot\|_1$ *here denotes the entrywise L-1 norm, i.e., the sum of the absolute values of the entries of the matrix.*

Proof. Let $D := P - xy^T$. From Lemma 5.1,

$$\epsilon_A \leq \text{Tr}(AZ) - x^T Ay = \text{Tr}(A(Z - yx^T)).$$

If we then hold D fixed and restrict that A has entries bounded in [0,1], the quantity $\text{Tr}(AD^T)$ is maximized when

$$A_{i,j} = \begin{cases} 1 & D_{i,j} \geq 0 \\ 0 & D_{i,j} < 0 \end{cases}.$$

The resulting quantity $\text{Tr}(AD^T)$ will then be the sum of all nonnegative elements of D. Since the sum of all elements in D is zero, this quantity will be equal to $\frac{1}{2}\|D\|_1$.
The proof for ϵ_B is identical, and the result follows from that ϵ is the maximum of ϵ_A and ϵ_B. □

5.2 Bounds on ϵ

We provide a number of bounds on $\epsilon(x, y)$ for $x := P1_n$ and $y := P^T 1_m$ coming from any feasible solution to SDP2. Our first two theorems roughly state that solutions which are "close" to rank-1 provide small ϵ.

Theorem 5.5. *Consider any feasible solution $\mathcal{M}$ to SDP2. Suppose $\mathcal{M}$ is rank-k and its eigenvalues are $\lambda_1 \geq \lambda_2 \geq ... \geq \lambda_k > 0$. Then x and y constitute an ϵ-NE to the game (A, B) with $\epsilon \leq \frac{m+n}{2} \sum_{i=2}^{k} \lambda_i$.*

Proof. By the Perron Frobenius theorem (see e.g. [35, Chapter 8.3]), the eigenvector corresponding to λ_1 can be assumed to be nonnegative, and hence

$$s_1 = \|a_1\|_1 = \|b_1\|_1. \tag{40}$$

We further note that for all i, since $\begin{bmatrix} a_i \\ b_i \end{bmatrix}$ is a vector of length $m + n$ with 2-norm equal to 1, we must have

$$\left\| \begin{bmatrix} a_i \\ b_i \end{bmatrix} \right\|_1 \leq \sqrt{m+n}. \tag{41}$$

Since s_i is the sum of the elements of a_i and b_i, we know that

$$s_i \leq \min\{\|a_i\|_1, \|b_i\|_1\} \leq \frac{\sqrt{m+n}}{2}. \tag{42}$$

This then gives us

$$s_i^2 \leq \|a_i\|_1 \|b_i\|_1 \leq \frac{m+n}{4}, \tag{43}$$

with the first inequality following from (42) and the second from (41). Finally note that a consequence of the nonnegativity of $\|\cdot\|_1$ and (41) is that for all i, j,

$$\|a_i\|_1\|b_j\|_1 + \|b_i\|_1\|a_j\|_1 \leq (\|a_i\|_1 + \|b_i\|_1)(\|a_j\|_1 + \|b_j\|_1) = \left\| \begin{bmatrix} a_i \\ b_i \end{bmatrix} \right\|_1 \left\| \begin{bmatrix} a_j \\ b_j \end{bmatrix} \right\|_1 \overset{(41)}{\leq} m + n. \tag{44}$$

Now we let $D := P - xy^T$ and upper bound $\frac{1}{2}\|D\|_1$ using Lemma 5.3.

$$
\begin{aligned}
\frac{1}{2}\|D\|_1 &= \frac{1}{2}\| \sum_{i=1}^{k}\sum_{j>i}^{k} \lambda_i\lambda_j(s_j a_i - s_i a_j)(s_j b_i - s_i b_j)^T\|_1 \\
&\le \frac{1}{2}\sum_{i=1}^{k}\sum_{j>i}^{k} \|\lambda_i\lambda_j(s_j a_i - s_i a_j)(s_j b_i - s_i b_j)^T\|_1 \\
&\le \frac{1}{2}\sum_{i=1}^{k}\sum_{j>i}^{k} \lambda_i\lambda_j\|s_j a_i - s_i a_j\|_1\|s_j b_i - s_i b_j\|_1 \\
&\le \frac{1}{2}\sum_{i=1}^{k}\sum_{j>i}^{k} \lambda_i\lambda_j(s_j\|a_i\|_1 + s_i\|a_j\|_1)(s_j\|b_i\|_1 + s_i\|b_j\|_1) \qquad (45) \\
&\overset{(40),(43)}{\le} \frac{1}{2}\sum_{j=2}^{k}\lambda_1 s_1^2\lambda_j(s_j + \|a_j\|_1)(s_j + \|b_j\|_1) \\
&+ \frac{1}{2}\sum_{i=2}^{k}\sum_{j>i}^{k}\lambda_i\lambda_j(s_j^2\frac{m+n}{4} + s_i^2\frac{m+n}{4} + s_i s_j\|a_i\|_1\|b_j\|_1 + s_i s_j\|a_j\|_1\|b_i\|_1) \\
&\overset{(41),(44),(42)}{\le} \frac{m+n}{2}\lambda_1 s_1^2\sum_{i=2}^{k}\lambda_i \\
&+ \frac{1}{2}\sum_{i=2}^{k}\sum_{j>i}^{k}\lambda_i\lambda_j\frac{m+n}{4}(s_i^2 + s_j^2) + \lambda_i\lambda_j s_i s_j(m+n) \\
&\overset{\text{AMGM}^6}{\le} \frac{m+n}{2}\lambda_1 s_1^2\sum_{i=2}^{k}\lambda_i + \frac{m+n}{2}\sum_{i=2}^{k}\sum_{j>i}^{k}\lambda_i\lambda_j(\frac{s_i^2+s_j^2}{4} + \frac{s_i^2+s_j^2}{2}) \\
&= \frac{m+n}{2}\lambda_1 s_1^2\sum_{i=2}^{k}\lambda_i + \frac{3(m+n)}{8}\sum_{i=2}^{k}\sum_{j>i}^{k}\lambda_i\lambda_j(s_i^2 + s_j^2) \\
&= \frac{m+n}{2}\lambda_1 s_1^2\sum_{i=2}^{k}\lambda_i + \frac{3(m+n)}{8}(\sum_{i=2}^{k}\lambda_i s_i^2\sum_{j>i}^{k}\lambda_j + \sum_{i=2}^{k}\lambda_i\sum_{j>i}^{k}\lambda_j s_j^2) \\
&= \frac{m+n}{2}\lambda_1 s_1^2\sum_{i=2}^{k}\lambda_i + \frac{3(m+n)}{8}(\sum_{j=2}^{k}\lambda_j\sum_{2\le i<j}^{k}\lambda_i s_i^2 + \sum_{i=2}^{k}\lambda_i\sum_{j>i}^{k}\lambda_j s_j^2) \\
&\le \frac{m+n}{2}\lambda_1 s_1^2\sum_{i=2}^{k}\lambda_i + \frac{3(m+n)}{8}(\sum_{j=2}^{k}\lambda_j s_j^2)\sum_{i=2}^{k}\lambda_i \\
&\overset{(39)}{=} \frac{m+n}{2}\lambda_1 s_1^2\sum_{i=2}^{k}\lambda_i + \frac{3(m+n)}{8}(1 - \lambda_1 s_1^2)\sum_{i=2}^{k}\lambda_i \\
&= \frac{m+n}{8}(3 + \lambda_1 s_1^2)\sum_{i=2}^{k}\lambda_i \\
&\overset{(39)}{\le} \frac{m+n}{2}\sum_{i=2}^{k}\lambda_i.
\end{aligned}
$$

□

The following theorem quantifies how making the objective of the diagonal gap algorithm from Section 4 small makes ϵ small. The proof is similar to the proof of Theorem 5.5.

Theorem 5.6. *Let $\mathcal{M}$ be a feasible solution to SDP2. Then, x and y constitute an ϵ-NE to the game (A, B) with $\epsilon \leq \frac{3(m+n)}{8}(\text{Tr}(\mathcal{M}) - x^T x - y^T y)$.*

Proof. Let $\mathcal{M}$ be rank-k with eigenvalues $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_k > 0$ and eigenvectors $v_1, \ldots, v_k$ partitioned as in Lemma 5.2 so that $v_i = \begin{bmatrix} a_i \\ b_i \end{bmatrix}$ with $\sum_{j=1}^m (a_i)_j = \sum_{j=1}^n (b_i)_j$ for $i = 1, \ldots, k$. Let $s_i := \sum_{j=1}^m (a_i)_j$. Then we have $\text{Tr}(\mathcal{M}) = \sum_{i=1}^k \lambda_i$, and

$$x^T x + y^T y \overset{(37),(38)}{=} (\sum_{i=1}^k \lambda_i s_i v_i)^T (\sum_{i=1}^k \lambda_i s_i v_i) = \sum_{i=1}^k \lambda_i^2 s_i^2. \tag{46}$$

We now get the following chain of inequalities (the first one follows from Lemma 5.4 and inequality (45)):

$$\begin{aligned}
\epsilon &\leq \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_j \|a_i\|_1 + s_i \|a_j\|_1)(s_j \|b_i\|_1 + s_i \|b_j\|_1) \\
&\overset{(40),(43)}{\leq} \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_j^2 \frac{m+n}{4} + s_i^2 \frac{m+n}{4} + s_i s_j \|a_i\|_1 \|b_j\|_1 + s_i s_j \|a_j\|_1 \|b_i\|_1) \\
&\overset{(44)}{\leq} \frac{1}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j \frac{m+n}{4} (s_i^2 + s_j^2) + \lambda_i \lambda_j s_i s_j (m+n) \\
&\overset{AMGM}{\leq} \frac{m+n}{2} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (\frac{s_i^2 + s_j^2}{4} + \frac{s_i^2 + s_j^2}{2}) \\
&= \frac{3(m+n)}{8} \sum_{i=1}^k \sum_{j>i}^k \lambda_i \lambda_j (s_i^2 + s_j^2) \\
&= \frac{3(m+n)}{8} (\sum_{i=1}^k \lambda_i s_i^2 \sum_{j>i}^k \lambda_j + \sum_{i=1}^k \lambda_i \sum_{j>i}^k \lambda_j s_j^2) \\
&= \frac{3(m+n)}{8} (\sum_{j=1}^k \lambda_j \sum_{1 \leq i < j}^k \lambda_i s_i^2 + \sum_{i=1}^k \lambda_i \sum_{j>i}^k \lambda_j s_j^2) \\
&= \frac{3(m+n)}{8} (\sum_{i=1}^k \lambda_i \sum_{j \neq i} \lambda_j s_j^2) \\
&\overset{(39)}{=} \frac{3(m+n)}{8} (\sum_{i=1}^k \lambda_i (1 - \lambda_i s_i^2)) \\
&= \frac{3(m+n)}{8} (\sum_{i=1}^k \lambda_i - \sum_{i=1}^k \lambda_i^2 s_i^2) \overset{(46)}{=} \frac{3(m+n)}{8} (\text{Tr}(\mathcal{M}) - x^T x - y^T y).
\end{aligned}$$

□

[6] AMGM is used to denote the arithmetic-mean-geometric-mean inequality.

We now give a bound on ϵ which is dependent on the nonnegative rank of the matrix returned by SDP2. Our analysis will also be useful for the next subsection. To begin, we first recall the definition of the nonnegative rank.

Definition 5.7. *The* nonnegative rank *of a (nonnegative) $m \times n$ matrix M is the smallest k for which there exist a nonnegative $m \times k$ matrix U and a nonnegative $n \times k$ matrix V such that $M = UV^T$. Such a decomposition is called a* nonnegative matrix factorization *of M.*

Theorem 5.8. *Consider the matrix P from any feasible solution to SDP2. Suppose its nonnegative rank is k. Then $x := P1_n$ and $y := P^T 1_m$ constitute an ϵ-NE to the game (A, B) with $\epsilon \leq 1 - \frac{1}{k}$.*

Proof. Since P has nonnegative rank k and its entries sum up to 1, we can write $P = \sum_{i=1}^k \sigma_i a_i b_i^T$, where $a_i \in \triangle_m, b_i \in \triangle_n$, and $\sum_{i=1}^k \sigma_i = 1$. From Lemma 5.4 and inequality (45) (keeping in mind that $s_i = 1, \ \forall\ i$) we have

$$\begin{aligned}
\epsilon &\leq \frac{1}{2} \sum_{i=1}^{k} \sum_{j>i}^{k} \sigma_i \sigma_j (\|a_i\|_1 + \|a_j\|_1)(\|b_i\|_1 + \|b_j\|_1) \\
&\leq 2 \sum_{i=1}^{k} \sum_{j>i}^{k} \sigma_i \sigma_j \\
&= 2(\frac{1}{2}(\sum_{i=1}^{k} \sigma_i \sum_{j=1}^{k} \sigma_j - \sum_{i=1}^{k} \sigma_i^2)) \\
&= 1 - \sum_{i=1}^{k} \sigma_i^2 \\
&\leq 1 - \frac{1}{k},
\end{aligned}$$

where the last line follows from the fact that $\|v\|_2^2 \geq \frac{1}{k}$ for any vector $v \in \triangle_k$. □

5.3 Bounds on ϵ in the Rank-2 Case

We now provide a number of bounds on $\epsilon(x, y)$ with $x := P1_n$ and $y := P^T 1_m$ which hold for rank-2 feasible solutions $\mathcal{M}$ to SDP2 (note that P will have rank at most 2 in this case). This is motivated by our ability to show stronger (constant) bounds in this case, and the fact that we often recover rank-2 (or rank-1) solutions with our algorithms in Section 4. Furthermore, our analysis will use the special property that a rank-2 nonnegative matrix will have nonnegative rank also equal to two, and that a nonnegative factorization of it can be computed in polynomial time (see, e.g., Section 4 of [11]). We begin with the following observation, which follows from Theorem 5.8 when $k = 2$.

Corollary 5.9. *If the matrix P from a feasible solution to SDP2 is rank-2, then x and y constitute a $\frac{1}{2}$–NE.*

We now show how this pair of strategies can be refined.

Theorem 5.10. *If the matrix P from a feasible solution to SDP2 is rank-2, then either x and y constitute a $\frac{5}{11}$-NE, or a $\frac{5}{11}$-NE can be recovered from P in polynomial time.*

Proof. We consider 3 cases, depending on whether $\epsilon_A(x, y)$ and $\epsilon_B(x, y)$ are greater than or less than .4. If $\epsilon_A \leq .4, \epsilon_B \leq .4$, then (x, y) is already a .4-Nash equilibrium. Now consider the case when $\epsilon_A \geq .4, \epsilon_B \geq .4$. Since $\epsilon_A \leq \text{Tr}(A(P - xy^T)^T)$ and $\epsilon_B \leq \text{Tr}(B(P - xy^T)^T)$ as seen in the proof of Lemma 5.1, we have, reusing the notation in the proof of Theorem 5.8,

$$\sigma_1\sigma_2(a_1 - a_2)^T A(b_1 - b_2) \geq .4, \sigma_1\sigma_2(a_1 - a_2)^T B(b_1 - b_2) \geq .4.$$

Since A, a_1, a_2, b_1, and b_2 are all nonnegative and $\sigma_1\sigma_2 \leq \frac{1}{4}$,

$$a_1^T A b_1 + a_2^T A b_2 \geq (a_1 - a_2)^T A(b_1 - b_2) \geq 1.6,$$

and the same inequalities hold for for player B. In particular, since A and B have entries bounded in [0,1] and a_1, a_2, b_1, and b_2 are simplex vectors, all the quantities $a_1^T A b_1, a_2^T A b_2, a_1^T B b_1,$ and $a_2^T B b_2$ are at most 1, and consequently at least .6. Hence (a_1, a_2) and (a_2, b_2) are both .4-Nash equilibria.

Now suppose that (x, y) is a .4-NE for one player (without loss of generality player A) but not for the other (without loss of generality player B). Then $\epsilon_A \leq .4$, and $\epsilon_B \geq .4$. Let y^* be a best response for player B to x, and let $p = \frac{1}{1+\epsilon_B - \epsilon_A}$. Consider the strategy profile $(\tilde{x}, \tilde{y}) := (x, py + (1-p)y^*)$. This can be interpreted as the outcome (x, y) occurring with probability p, and the outcome (x, y^*) happening with probability $1 - p$. In the first case, player A will have $\epsilon_A(x, y) = \epsilon_A$ and player B will have $\epsilon_B(x, y) = \epsilon_B$. In the second outcome, player A will have $\epsilon_A(x, y^*)$ at most 1, while player B will have $\epsilon_B(x, y^*) = 0$. Then under this strategy profile, both players have the same upper bound for ϵ, which equals $\epsilon_B p = \frac{\epsilon_B}{1+\epsilon_B - \epsilon_A}$. To find the worst case for this value, let $\epsilon_B = .5$ (note from Theorem 5.9 that $\epsilon_B \leq \frac{1}{2}$) and $\epsilon_A = .4$, and this will return $\epsilon = \frac{5}{11}$.

□

We now show a stronger result in the case of symmetric games.

Definition 5.11. *A* symmetric game *is a game in which the payoff matrices A and B satisfy* $B = A^T$.

Definition 5.12. *A Nash equilibrium strategy* (x, y) *is said to be* symmetric *if* $x = y$.

Theorem 5.13 (see Theorem 2 in [36])**.** *Every symmetric bimatrix game has a symmetric Nash equilibrium.*

For the proof of Theorem 5.17 below we modify SDP2 so that we are seeking a symmetric solution. We also need a more specialized notion of the nonnegative rank.

Definition 5.14. *A matrix M is* completely positive *(CP) if it admits a decomposition* $M = UU^T$ *for some nonnegative matrix U.*

Definition 5.15. *The* CP-rank *of an* $n \times n$ *CP matrix M is the smallest k for which there exists a nonnegative* $n \times k$ *matrix U such that* $M = UU^T$.

Theorem 5.16 (see e.g. [26] or Theorem 2.1 in [6])**.** *A rank-2, nonnegative, and positive semidefinite matrix is CP and has CP-rank 2.*

It is also known (see e.g., Section 4 in [26]) that the CP factorization of a rank-2 CP matrix can be found to arbitrary accuracy in polynomial time.

Theorem 5.17. *Suppose the constraint* $P \succeq 0$ *is added to SDP2. Then if in a feasible solution to this new SDP the matrix P is rank-2, either x and y constitute a symmetric* $\frac{1}{3}$*-NE, or a symmetric* $\frac{1}{3}$*-NE can be recovered from P in polynomial time.*

Proof. If (x, y) is already a symmetric $\frac{1}{3}$-NE, then the claim is established. Now suppose that (x, y) does not constitute a $\frac{1}{3}$-Nash equilibrium. Similarly as in the proof of Theorem 5.8, we can decompose P into $\sum_{i=1}^{2} \sigma_i a_i a_i^T$, where $\sum_{i=1}^{2} \sigma_i = 1$ and each a_i is a vector on the unit simplex. Then we have

$$\sigma_1\sigma_2(a_1 - a_2)^T A(a_1 - a_2) \geq \frac{1}{3}.$$

Since A, a_1, and a_2 are all nonnegative, and $\sigma_1\sigma_2 \leq \frac{1}{4}$, we get

$$a_1^T A a_1 + a_2^T A a_2 \geq (a_1 - a_2)^T A(a_1 - a_2) \geq \frac{4}{3}.$$

In particular, at least one of $a_1^T A a_1$ and $a_2^T A a_2$ is at least $\frac{2}{3}$. Since the maximum possible payoff is 1, at least one of (a_1, a_1) and (a_2, a_2) is a (symmetric) $\frac{1}{3}$-Nash equilibrium. □

Remark 5.1. For symmetric games, instead of the construction stated in Theorem 5.17, one can simply optimize over a smaller $m \times m$ matrix (note $m = n$). This is the relaxed version of exchangeable equilibria [43], with the completely positive constraint relaxed to a psd constraint.

Remark 5.2. The statements of Corollary 5.9, and Theorem 5.10, and Theorem 5.17 hold for any rank-2 correlated equilibrium. Indeed, given any rank-2 (equivalently, nonnegative-rank-2) correlated equilibrium P, one can complete it to a (rank-2) feasible solution to SDP2 as follows. Let $P = \sum_{i=1}^{2} \sigma_i a_i b_i^T$, where $a_i \in \triangle_m, b_i \in \triangle_n$, and $\sigma_1 + \sigma_2 = 1$. It is easy to check that

$$\mathcal{M} \ := \ \sum_{i=1}^{2} \sigma_i \begin{bmatrix} a_i \\ b_i \end{bmatrix} \begin{bmatrix} a_i \\ b_i \end{bmatrix}^T$$

is feasible to SDP2.

6 Bounding Payoffs and Strategy Exclusion in Symmetric Games

In addition to finding ϵ-additive Nash equilibria, our SDP approach can be used to answer certain questions of economic interest about Nash equilibria without actually computing them. For instance, economists often would like to know the maximum welfare (sum of the two players' payoffs) achievable under any Nash equilibrium, or whether there exists a Nash equilibrium in which a given subset of strategies (corresponding, e.g., to undesirable behavior) is not played. Both these questions are NP-hard for bimatrix games [21], even when the game is symmetric and only symmetric equilibria are considered [13]. In this section, we consider these two problems in the symmetric setting and compare the performance of our SDP approach to an LP approach which searches over symmetric correlated equilibria. For general equilibria, it turns out that for these two specific questions, our SDP approach is equivalent to an LP that searches over correlated equilibria.

6.1 Bounding Payoffs

When designing policies that are subject to game theoretic behavior by agents, economists would often like to find one with a good socially optimal outcome, which usually corresponds to an equilibrium giving the maximum welfare. Hence, given a game, it is of interest to know the highest achievable welfare under any Nash equilibrium. For symmetric games, symmetric equilibria are of particular interest as they reflect the notion that identical agents should behave similarly given identical options.

Note that the maximum welfare of a symmetric game under any symmetric Nash equilibrium is equal to the optimal value of the following quadratic program:

$$\max_{x \in \triangle_m} \quad 2x^T A x$$
$$\text{subject to} \quad x^T A x \geq e_i^T A x, \forall i \in \{1, \dots, m\}. \tag{47}$$

One can find an upper bound on this number by solving an LP which searches over symmetric correlated equilibria:

$$\max_{P \in \mathbb{S}^{m,m}} \quad \mathrm{Tr}(AP^T) \tag{LP1}$$

$$\text{subject to} \quad \sum_{i=1}^{m} \sum_{j=1}^{m} P_{i,j} = 1 \tag{48}$$

$$\sum_{j=1}^{m} A_{i,j} P_{i,j} \geq \sum_{j=1}^{m} A_{k,j} P_{i,j}, \forall i, k \in \{1, \dots, m\}, \tag{49}$$

$$P \geq 0. \tag{50}$$

A potentially better upper bound on the maximum welfare can be obtained from a version of SDP2 adapted to this specific problem:

$$\max_{P \in \mathbb{S}^{m,m}} \quad \mathrm{Tr}(AP^T) \tag{SDP3}$$
$$\text{subject to} \quad (48), (49), (50)$$
$$P \succeq 0.$$

To test the quality of these upper bounds, we tested this LP and SDP on a random sample of one hundred 5×5 and 10×10 games[7]. The resulting upper bounds are in Figure 2, which shows that the bound returned by SDP3 was exact in a large number of the experiments.[8]

[7]The matrix A in each game was randomly generated with diagonal entries uniform and independent in [0,.5] and off-diagonal entries uniform and independent in [0,1].

[8]The computation of the exact maximum payoffs was done with the `lrsnash` software [5], which computes all extreme Nash equilibria. For a definition of extreme Nash equilibria and for understanding why it is sufficient for us to compare against extreme Nash equilibria (both in Section 6.1 and in Section 6.2), see Appendix C. The computation of the SDP upper bound has been implemented in the file nashbound.m, which is publicly available at `aaa.princeton.edu/software`. This file more generally computes an SDP-based lower bound on the minimum of an input quadratic function over the set of Nash equilibria of a bimatrix game. The file also takes as an argument whether one wishes to only consider symmetric equilibria when the game is symmetric.

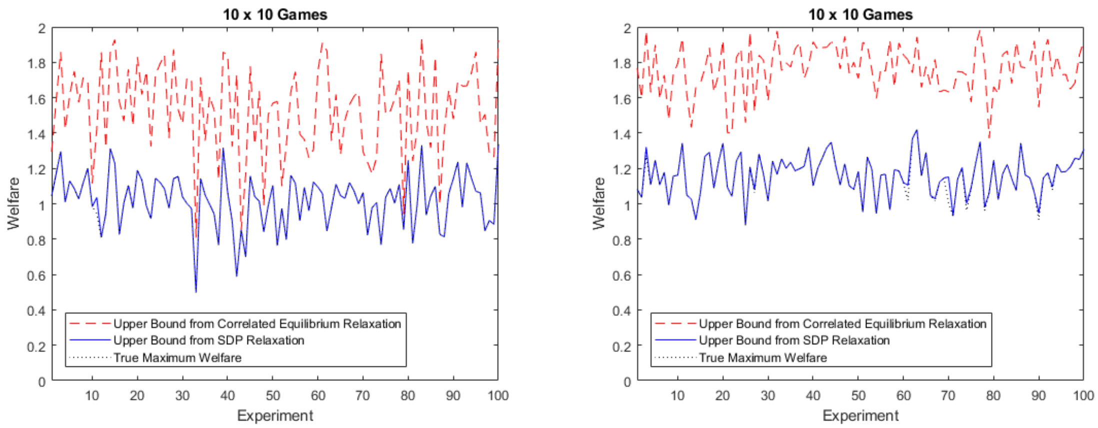

Figure 2: The quality of the upper bound on the maximum welfare obtained by LP1 and SDP3 on 100 5×5 games (left) and 100 10×10 games (right).

6.2 Strategy Exclusion

The strategy exclusion problem asks, given a subset of strategies $\mathcal{S} = (\mathcal{S}_x, \mathcal{S}_y)$, with $\mathcal{S}_x \subseteq \{1, \ldots, m\}$ and $\mathcal{S}_y \subseteq \{1, \ldots, n\}$, is there a Nash equilibrium in which no strategy in $\mathcal{S}$ is played with positive probability. We will call a set $\mathcal{S}$ "persistent" if the answer to this question is negative, i.e. at least one strategy in $\mathcal{S}$ is played with positive probability in every Nash equilibrium. One application of the strategy exclusion problem is to understand whether certain strategies can be discouraged in the design of a game, such as reckless behavior in a game of chicken or defecting in a game of prisoner's dilemma. In these particular examples these strategy sets are persistent and cannot be discouraged.

As in the previous subsection, we consider the strategy exclusion problem for symmetric strategies in symmetric games (such as the aforementioned games of chicken and prisoner's dilemma). A quadratic program which addresses this problem is as follows:

$$\begin{aligned} \min_{x \in \triangle_m} \quad & \sum_{i \in \mathcal{S}_x} x_i \\ \text{subject to} \quad & x^T A x \geq e_i^T A x, \forall i \in \{1, \ldots, m\}. \end{aligned} \tag{51}$$

Observe that by design, $\mathcal{S}$ is persistent if and only if this quadratic program has a positive optimal value. As in the previous subsection, an LP relaxation of this problem which searches over symmetric correlated equilibria is given by

$$\begin{aligned} \min_{P \in \mathbb{S}^{m,m}} \quad & \sum_{i \in \mathcal{S}_x} \sum_{j=1}^{m} P_{ij} \\ \text{subject to} \quad & (48), (49), (50). \end{aligned} \tag{LP2}$$

The SDP relaxation that we propose for the strategy exclusion problem is the following:

$$
\begin{aligned}
\min_{P\in\mathbb{S}^{m,m}} \quad & \sum_{i\in\mathcal{S}_x}\sum_{j=1}^{m} P_{ij} && \text{(SDP4)}\\
\text{subject to} \quad & (48),(49),(50) \\
& P \succeq 0.
\end{aligned}
$$

Our approach would be to declare that the strategy set $\mathcal{S}_x$ is persistent if and only if SDP4 has a positive optimal value.

Note that since the optimal value of SDP4 is a lower bound for that of (51), SDP4 carries over the property that if a set $\mathcal{S}$ is not persistent, then the SDP for sure returns zero. Thus, when using SDP4 on a set which is not persistent, our algorithm will always be correct. However, this is not necessarily the case for a persistent set. While we can be certain that a set is persistent if SDP4 returns a positive optimal value (again, because the optimal value of SDP4 is a lower bound for that of (51)), there is still the possibility that for a persistent set SDP4 will have optimal value zero. The same arguments hold for the optimal value of LP2.

To test the performance of LP2 and SDP4, we generated 100 random games of size 5×5 and 10×10 and computed all their symmetric extreme Nash equilibria[9]. We then, for every strategy set $\mathcal{S}$ of cardinality one and two, checked whether that set of strategies was persistent, first by checking among the extreme Nash equilibria, then through LP2 and SDP4. The results are presented in Tables 2 and 3. As can be seen, SDP4 was quite effective for the strategy exclusion problem.

Table 2: Performance of LP2 and SDP4 on 5×5 games

$\lvert\mathcal{S}\rvert$	1	2
Number of total sets	500	1000
Number of persistent sets	245	748
Persistent sets certified (LP2)	177 (72.2%)	661 (88.7%)
Persistent sets certified (SDP4)	245 (100%)	748 (100%)

Table 3: Performance of LP2 and SDP4 on 10×10 games

$\lvert\mathcal{S}\rvert$	1	2
Number of total sets	1000	4500
Number of persistent sets	326	2383
Persistent sets certified (LP2)	39 (12.0%)	630 (26.4%)
Persistent sets certified (SDP4)	318 (97.5%)	2368 (99.4%)

7 Connection to the Sum of Squares/Lasserre Hierarchy

In this section, we clarify the connection of the SDPs we have proposed in this paper to those arising in the sum of squares/Lasserre hierarchy. We start by briefly reviewing this hierarchy.

[9]The exact computation of the exact Nash equilibria was done again with the `lrsnash` software [5], which computes extreme Nash equilibria. To understand why this suffices for our purposes see Appendix C.

7.1 Sum of Squares/Lasserre Hierarchy

The sum of squares/Lasserre hierarchy[10] gives a recipe for constructing a sequence of SDPs whose optimal values converge to the optimal value of a given polynomial optimization problem. Recall that a *polynomial optimization problem* (pop) is a problem of minimizing a polynomial over a basic semialgebraic set, i.e., a problem of the form

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f(x) \\ \text{subject to} \quad & g_i(x) \geq 0, \forall i \in \{1, \ldots, m\}, \end{aligned} \tag{52}$$

where f, g_i are polynomial functions. In this section, when we refer to the k-th level of the Lasserre hierarchy, we mean the optimization problem

$$\begin{aligned} \gamma_{sos}^k := \max_{\gamma, \sigma_i} \quad & \gamma \\ \text{subject to} \quad & f(x) - \gamma = \sigma_0(x) + \sum_{i=1}^m \sigma_i(x) g_i(x), \\ & \sigma_i \text{ is sos}, \forall i \in \{0, \ldots, m\}, \\ & \sigma_0, g_i\sigma_i \text{ have degree at most } 2k, \ \forall i \in \{1, \ldots, m\}. \end{aligned} \tag{53}$$

Here, the notation "sos" stands for *sum of squares.* We say that a polynomial p is a sum of squares if there exist polynomials $q_1, \ldots, q_r$ such that $p = \sum_{i=1}^r q_i^2$. There are two primary properties of the Lasserre hierarchy which are of interest. The first is that any fixed level of this hierarchy gives an SDP of size polynomial in n. The second is that, if the set $\{x \in \mathbb{R}^n | g_i(x) \geq 0\}$ is Archimedean (see, e.g. [31] for definition), then $\lim_{k \to \infty} \gamma_{sos}^k = p^*$, where p^* is the optimal value of the pop in (52). The latter statement is a consequence of Putinar's positivstellensatz [39], [29].

7.2 The Lasserre Hierarchy and SDP1

One can show, e.g. via the arguments in [30], that the feasible sets of the SDPs dual to the SDPs underlying the hierarchy we summarized above produce an arbitrarily tight outer approximation to the convex hull of the set of Nash equilibria of any game. The downside of this approach, however, is that the higher levels of the hierarchy can get expensive very quickly. This is why the approach we took in this paper was instead to improve the first level of the hierarchy. The next proposition formalizes this connection.

Proposition 7.1. *Consider the problem of minimizing any quadratic objective function over the set of Nash equilibria of a bimatrix game. Then, SDP1 (and hence SDP2) gives a lower bound on this problem which is no worse than that produced by the first level of the Lasserre hierarchy.*

Proof. To prove this proposition we show that the first level of the Lasserre hierarchy is dual to a weakened version of SDP1.

Explicit parametrization of first level of the Lasserre hierarchy. Consider the formulation of the Lasserre hierarchy in (53) with $k = 1$. Suppose we are minimizing a quadratic function

$$f(x, y) = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T \mathcal{C} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

[10]The unfamiliar reader is referred to [29, 37, 31] for an introduction to this hierarchy and the related theory of moment relaxations.

over the set of Nash equilibria as described by the linear and quadratic constraints in (2). If we apply the first level of the Lasserre hierarchy to this particular pop, we get

$$
\begin{aligned}
\max_{Q,\alpha,\chi,\beta,\psi,\eta} \quad & \gamma \\
\text{subject to} \quad & \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T \mathcal{C} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} - \gamma = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T Q \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} + \sum_{i=1}^{m} \alpha_i (x^T A y - e_i^T A y) \\
& \quad + \sum_{i=1}^{n} \beta_i (x^T B y - x^T B e_i) \\
& \quad + \sum_{i=1}^{m} \chi_i x_i + \sum_{i=1}^{n} \psi_i y_i \\
& \quad + \eta_1 (\sum_{i=1}^{m} x_i - 1) + \eta_2 (\sum_{i=1}^{n} y_i - 1), \\
& Q \succeq 0, \\
& \alpha, \chi, \beta, \psi \geq 0,
\end{aligned}
\tag{54}
$$

where $Q \in \mathbb{S}^{m+n+1 \times m+n+1}, \alpha, \chi \in \mathbb{R}^m, \beta, \psi \in \mathbb{R}^n, \eta \in \mathbb{R}^2$.

By matching coefficients of the two quadratic functions on the left and right hand sides of (54), this SDP can be written as

$$
\begin{aligned}
\max_{\gamma,\alpha,\beta,\chi,\psi,\eta} \quad & \gamma \\
\text{subject to} \quad & \mathcal{H} \succeq 0, \\
& \alpha, \beta, \chi, \psi \geq 0,
\end{aligned}
\tag{55}
$$

where

$$
\mathcal{H} := \frac{1}{2} \begin{bmatrix} 0 & (-\sum_{i=1}^m \alpha_i) A + (-\sum_{i=1}^m \beta_i) B & \sum_{i=1}^n \beta_i B_{,i} - \chi - \eta_1 1_m \\ (-\sum_{i=1}^m \alpha_i) A + (-\sum_{i=1}^n \beta_i) B & 0 & \sum_{i=1}^m \alpha_i A_{i,}^T - \psi - \eta_2 1_n \\ \sum_{i=1}^n \beta_i B_{,i}^T - \chi^T - \eta_1 1_m^T & \sum_{i=1}^m \alpha_i A_{i,} - \psi^T - \eta_2 1_n^T & 2\eta_1 + 2\eta_2 - 2\gamma \end{bmatrix} + \mathcal{C}.
\tag{56}
$$

Dual of a weakened version of SDP1. With this formulation in mind, let us consider a weakened version of SDP1 with only the relaxed Nash constraints, unity constraints, and nonnegativity constraints on x and y in the last column (i.e., the nonegativity constraint is not applied to the entire matrix). Let the objective be $\mathrm{Tr}(C\mathcal{M}')$. To write this new SDP in standard form, let

$$
\mathcal{A}_i := \frac{1}{2} \begin{bmatrix} 0 & A & 0 \\ A^T & 0 & -A_{i,}^T \\ 0 & -A_{i,} & 0 \end{bmatrix}, \mathcal{B}_i := \frac{1}{2} \begin{bmatrix} 0 & B & -B_{,i} \\ B^T & 0 & 0 \\ -B_{,i}^T & 0 & 0 \end{bmatrix},
$$

$$
\mathcal{S}_1 := \frac{1}{2} \begin{bmatrix} 0 & 0 & 1_m \\ 0 & 0 & 0 \\ 1_m^T & 0 & -2 \end{bmatrix}, \mathcal{S}_2 := \frac{1}{2} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1_n \\ 0 & 1_n^T & -2 \end{bmatrix}.
$$

Let $\mathcal{N}_i$ be the matrix with all zeros except a $\frac{1}{2}$ at entry $(i, m+n+1)$ and $(m+n+1, i)$ (or a 1 if $i = m+n+1$).

Then this SDP can be written as

$$
\begin{aligned}
\min_{\mathcal{M}'} \quad & \operatorname{Tr}(\mathcal{C}\mathcal{M}') && \text{(SDP0)}\\
\text{subject to} \quad & \mathcal{M}' \succeq 0, && (57)\\
& \operatorname{Tr}(\mathcal{N}_i\mathcal{M}') \geq 0, \forall i \in \{1,\dots,m+n\}, && (58)\\
& \operatorname{Tr}(\mathcal{A}_i\mathcal{M}') \geq 0, \forall i \in \{1,\dots,m\}, && (59)\\
& \operatorname{Tr}(\mathcal{B}_i\mathcal{M}') \geq 0, \forall i \in \{1,\dots,n\}, && (60)\\
& \operatorname{Tr}(\mathcal{S}_1\mathcal{M}') = 0, && (61)\\
& \operatorname{Tr}(\mathcal{S}_2\mathcal{M}') = 0, && (62)\\
& \operatorname{Tr}(\mathcal{N}_{m+n+1}) = 1. && (63)
\end{aligned}
$$

We now create dual variables for each constraint; we choose α_i and β_i for the relaxed Nash constraints (59) and (60), η_1 and η_2 for the unity constraints (61) and (62), χ for the nonnegativity of x (58), ψ for the nonnegativity of y (58), and γ for the final constraint on the corner (63). These variables are chosen to coincide with those used in the parametrization of the first level of the Lasserre hierarchy, as can be seen more clearly below.

We then write the dual of the above SDP as

$$
\begin{aligned}
\max_{\alpha,\beta,\lambda,\gamma} \quad & \gamma\\
\text{subject to} \quad & \sum_{i=1}^{m}\alpha_i\mathcal{A}_i + \sum_{i=1}^{n}\beta_i\mathcal{B}_i + \sum_{i=1}^{2}\eta_i\mathcal{S}_i + \sum_{i=1}^{m}\mathcal{N}_{i+n}\chi_i + \sum_{i=1}^{n}\mathcal{N}_i\psi_i + \gamma\mathcal{N}_{m+n+1} \preceq \mathcal{C},\\
& \alpha,\beta,\chi,\psi \geq 0.
\end{aligned}
$$

which can be rewritten as

$$
\begin{aligned}
\max_{\alpha,\beta,\chi,\psi,\gamma} \quad & \gamma\\
\text{subject to} \quad & \mathcal{G} \succeq 0, \qquad (64)\\
& \alpha,\beta,\chi,\psi \geq 0,
\end{aligned}
$$

where

$$
\mathcal{G} := \frac{1}{2}\begin{bmatrix} 0 & (-\sum_{i=1}^{m}\alpha_i)A + (-\sum_{i=1}^{m}\beta_i)B & \sum_{i=1}^{n}\beta_i B_{,i} - \chi - \eta_1 1_m\\ (-\sum_{i=1}^{m}\alpha_i)A + (-\sum_{i=1}^{n}\beta_i)B & 0 & \sum_{i=1}^{m}\alpha_i A_{i,}^T - \psi - \eta_2 1_n\\ \sum_{i=1}^{n}\beta_i B_{,i}^T - \chi^T - \eta_1 1_m^T & \sum_{i=1}^{m}\alpha_i A_{i,} - \psi^T - \eta_2 1_n^T & 2\eta_1 + 2\eta_2 - 2\gamma \end{bmatrix} + \mathcal{C}.
$$

We can now see that the matrix $\mathcal{G}$ coincides with the matrix $\mathcal{H}$ in the SDP (55). Then we have

$$
(54)^{opt} = (55)^{opt} = (64)^{opt} \leq \ SDP0^{opt} \leq \ SDP1^{opt},
$$

where the first inequality follows from weak duality, and the second follows from that the constraints of SDP0 are a subset of the constraints of SDP1. □

Remark 7.1. The Lasserre hierarchy can be viewed in each step as a pair of primal-dual SDPs: the sum of squares formulation which we have just presented, and a moment formulation which is dual to the sos formulation [29]. All our SDPs in this paper can be viewed more directly as an improvement upon the moment formulation.

Remark 7.2. One can see, either by inspection or as an implication of the proof of Theorem 2.2, that in the case where the objective function corresponds to maximizing player A's and/or B's payoffs[11], SDPs (55) and (64) are infeasible. This means that for such problems the first level of the Lasserre hierarchy gives an upper bound of $+\infty$ on the maximum payoff. On the other hand, the additional valid inequalities in SDP2 guarantee that the resulting bound is always finite.

8 Future Work

Our work leaves many avenues of further research. Are there other interesting subclasses of games (besides strictly competitive games) for which our SDP is guaranteed to recover an exact Nash equilibrium? Can the guarantees on ϵ in Section 5 be improved in the rank-2 case (or the general case) by improving our analysis? Is there a polynomial time algorithm that is guaranteed to find a rank-2 solution to SDP2? Such an algorithm, together with our analysis, would improve the best known approximation bound for symmetric games (see Theorem 5.17). Can this bound be extended to general games? We show in Appendix D that some natural approaches based on symmetrization of games do not immediately lead to a positive answer to this question. Can SDPs in a higher level of the Lasserre hierarchy be used to achieve better ϵ guarantees? What are systematic ways of adding valid inequalities to these higher-order SDPs by exploiting the structure of the Nash equilibrium problem? For example, since any strategy played with positive probability must give the same payoff, one can add a relaxed version of the cubic constraints

$$x_i x_j (e_i^T Ay - e_j^T Ay) = 0, \forall i, j \in \{1, \dots, m\}$$

to the SDP underlying the second level of the Lasserre hierarchy. What are other valid inequalities for the second level? Finally, our algorithms were specifically designed for two-player one-shot games. This leaves open the design and analysis of semidefinite relaxations for repeated games or games with more than two players.

Acknowledgments.

We would like to thank Ilan Adler, Costis Daskalakis, Georgina Hall, Ramon van Handel, and Robert Vanderbei for insightful exchanges. We are also extremely grateful to an anonymous referee for various insightful questions and comments which have led to a significantly improved version of this manuscript.

References

[1] *MOSEK reference manual*, 2013. Version 7. Latest version available at `http://www.mosek.com/`.

[2] I. Adler. The equivalence of linear programs and zero-sum games. *International Journal of Game Theory*, 42(1):165–177, 2013.

[3] I. Adler, C. Daskalakis, and C. H. Papadimitriou. A note on strictly competitive games. In *International Workshop on Internet and Network Economics*, pages 471–474. Springer, 2009.

[11] This would be the case, for example, in the maximum social welfare problem of Section 6.1, where the matrix of the quadratic form in the objective function is given by

$$\mathcal{C} = \begin{bmatrix} 0 & -A-B & 0 \\ -A-B & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

[4] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[5] D. Avis, G. D. Rosenberg, R. Savani, and B. Von Stengel. Enumeration of Nash equilibria for two-player games. *Economic Theory*, 42(1):9–37, 2010.

[6] A. Berman and N. Shaked-Monderer. *Completely positive matrices.* World Scientific, 2003.

[7] S. Boyd and L. Vandenberghe. *Convex optimization.* Cambridge university press, 2004.

[8] S. P. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix inequalities in system and control theory*, volume 15. SIAM, 1994.

[9] X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *FOCS*, volume 6, page 47th, 2006.

[10] X. Chen, X. Deng, and S.-H. Teng. Computing Nash equilibria: Approximation and smoothed complexity. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 603–612. IEEE, 2006.

[11] J. E. Cohen and U. G. Rothblum. Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Algebra and its Applications*, 190:149–168, 1993.

[12] V. Conitzer and T. Sandholm. Complexity results about nash equilibria. *arXiv preprint cs/0205074*, 2002.

[13] V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621–641, 2008.

[14] G. B. Dantzig. A proof of the equivalence of the programming problem and the game problem. *Activity analysis of production and allocation*, 13:330–338, 1951.

[15] C. Daskalakis. On the complexity of approximating a Nash equilibrium. *ACM Transactions on Algorithms (TALG)*, 9(3):23, 2013.

[16] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[17] C. Daskalakis, A. Mehta, and C. Papadimitriou. A note on approximate Nash equilibria. In *International Workshop on Internet and Network Economics*, pages 297–306. Springer, 2006.

[18] C. Daskalakis, A. Mehta, and C. Papadimitriou. Progress in approximate Nash equilibria. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 355–358. ACM, 2007.

[19] M. Fazel. *Matrix rank minimization with applications.* PhD thesis, PhD thesis, Stanford University, 2002.

[20] J. Fearnley, P. W. Goldberg, R. Savani, and T. B. Sørensen. Approximate well-supported Nash equilibria below two-thirds. *Algorithmica*, 76(2):297–319, 2016.

[21] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.

[22] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.

[23] J. Griesmer, A. Hoffman, and A. Robinson. On symmetric bimatrix games. *IBM Research Paper RC-959. IBM Corp, Thomas J Watson Research Center, Yorktown Heights, New York*, 1963.

[24] S. Ibaraki and M. Tomizuka. Rank minimization approach for solving BMI problems with random search. In *American Control Conference, 2001. Proceedings of the 2001*, volume 3, pages 1870–1875. IEEE, 2001.

[25] A. Jurg, M. Jansen, J. A. Potters, and S. Tijs. A symmetrization for finite two-person games. *Zeitschrift für Operations Research*, 36(2):111–123, 1992.

[26] V. Kalofolias and E. Gallopoulos. Computing symmetric nonnegative rank factorizations. *Linear Algebra and its Applications*, 436(2):421–435, 2012.

[27] S. C. Kontogiannis, P. N. Panagopoulou, and P. G. Spirakis. Polynomial algorithms for approximating Nash equilibria of bimatrix games. In *International Workshop on Internet and Network Economics*, pages 286–296. Springer, 2006.

[28] R. Laraki and J. B. Lasserre. Semidefinite programming for min–max problems and games. *Mathematical Programming*, 131(1-2):305–332, 2012.

[29] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.

[30] J. B. Lasserre. Convex sets with semidefinite representation. *Mathematical programming*, 120(2):457–477, 2009.

[31] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, pages 157–270. Springer, 2009.

[32] C. E. Lemke and J. T. Howson, Jr. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2):413–423, 1964.

[33] L. Lovász. On the shannon capacity of a graph. *IEEE Transactions on Information theory*, 25(1):1–7, 1979.

[34] G. P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I : Convex underestimating problems. *Mathematical Programming*, 10(1):147–175, 1976.

[35] C. D. Meyer. *Matrix analysis and applied linear algebra*, volume 2. SIAM, 2000.

[36] J. Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

[37] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2):293–320, 2003.

[38] P. A. Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 2855–2860. IEEE, 2006.

[39] M. Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3):969–984, 1993.

[40] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

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

[42] P. Shah and P. A. Parrilo. Polynomial stochastic games via sum of squares optimization. In *Decision and Control, 2007 46th IEEE Conference on*, pages 745–750. IEEE, 2007.

[43] N. D. Stein. *Exchangeable equilibria*. PhD thesis, Massachusetts Institute of Technology, 2011.

[44] H. Tsaknakis and P. G. Spirakis. An optimization approach for approximate Nash equilibria. In *International Workshop on Web and Internet Economics*, pages 42–56. Springer, 2007.

[45] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.

A Statistics on ϵ from Algorithms in Section 4

Below are statistics for the ϵ recovered in 100 random games of varying sizes using the algorithms of Section 4.

Table 4: Statistics on ϵ for 5×5 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0702	0.0040	0.0004	0.0099
Diagonal Gap	0.0448	0.0027	0	0.0061

Table 5: Statistics on ϵ for 10×5 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0327	0.0044	0.0021	0.0064
Diagonal Gap	0.0267	0.0033	0.0006	0.0053

Table 6: Statistics on ϵ for 10×10 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0373	0.0058	0.0039	0.0065
Diagonal Gap	0.0266	0.0043	0.0026	0.0051

Table 7: Statistics on ϵ for 15×10 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0206	0.0050	0.0034	0.0045
Diagonal Gap	0.0212	0.0038	0.0025	0.0039

Table 8: Statistics on ϵ for 15×15 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0169	0.0051	0.0042	0.0039
Diagonal Gap	0.0159	0.0038	0.0029	0.0034

Table 9: Statistics on ϵ for 20×15 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0152	0.0046	0.0035	0.0036
Diagonal Gap	0.0119	0.0032	0.0022	0.0027

Table 10: Statistics on ϵ for 20×20 games after 20 iterations.

Algorithm	Max	Mean	Median	StDev
Square Root	0.0198	0.0046	0.0039	0.0034
Diagonal Gap	0.0159	0.0032	0.0024	0.0032

B Comparison with an SDP Approach from [28]

In this section, at the request of a referee, we compare the first level of the SDP hierarchy given in [28, Section 4] to SDP2 using $\text{Tr}(M)$ as the objective function on 100 randomly generated games for each size given in the tables below. The first level of the hierarchy in [28] optimizes over a matrix which is slightly bigger than the one in SDP2, though it has a number of constraints linear in the size of the game considered, as opposed to the quadratic number in SDP2. We remark that the approach in [28] is applicable more generally to many other problems, including several in game theory.

The scalar ϵ reported in Table 11 is computed using the strategies (x, y) extracted from the first row of the optimal matrix M_1 as described in Section 4.1 of [28]. The scalar ϵ reported in Table 12 is computed using $x = P1_n$ and $y = P^T 1_m$ from the optimal solution to SDP2 with $\text{Tr}(\mathcal{M})$ as the objective function.

Table 11: Statistics on ϵ for first level of the hierarchy in [28].

	5×5	10×5	10×10	15×10	15×15	20×15	20×20
Max	0.3357	0.3304	0.2557	0.2189	0.1987	0.1837	0.1828
Mean	0.1883	0.1889	0.1513	0.1446	0.1262	0.1217	0.1087
Median	0.1803	0.1865	0.1452	0.1418	0.1271	0.1208	0.1070

Table 12: Statistics on ϵ for SDP2 with $\text{Tr}(M)$ as the objective function.

	5×5	10×5	10×10	15×10	15×15	20×15	20×20
Max	0.1581	0.1589	0.115	0.1335	0.0878	0.082	0.0619
Mean	0.0219	0.0332	0.0405	0.04	0.0366	0.0356	0.0298
Median	0.0046	0.0233	0.036	0.0346	0.0345	0.0325	0.0293

We also ran the second level of the hierarchy in [28] on the same 100 5×5 games. The maximum ϵ observed was .3362, while the mean was .1880 and the median was .1800. The size of the variable matrix that needs to be positive semidefinite for this level is 78×78.

C Lemmas for Extreme Nash Equilibria

The results reported in Section 6 were found using the `lrsnash` [5] software which computes extreme Nash equilibria (see definition below). In particular the true maximum welfare and the persistent strategy sets were found in relation to extreme symmetric Nash equilibria only. We show in this appendix why this is sufficient for the claims we made about *all* symmetric Nash equilibria. We prove a more general statement below about general games and general Nash equilibria since this could be of potential independent interest. The proof for symmetric games is identical once the strategies considered are restricted to be symmetric.

Definition C.1. *An* extreme Nash equilibrium *is a Nash equilibrium which cannot be expressed as a convex combination of other Nash equilibria.*

Lemma C.2. *All Nash equilibria are convex combinations of extreme Nash equilibria.*

Proof. It suffices to show that any extreme point of the convex hull of the set of Nash equilibria must be an extreme Nash equilibrium, as any point in a compact convex set can be written as a convex combination of its extreme points. Note that this convex hull contains three types of points: extreme Nash equilibria, Nash equilibria which are not extreme, and convex combinations of Nash equilibria which are not Nash equilibria. The claim then follows because any extreme point of the convex hull cannot be of the second or third type, as they can be written as convex combinations of other points in the hull. □

The next lemma shows that checking extreme Nash equilibria are sufficient for the maximum welfare problem.

Lemma C.3. *For any bimatrix game, there exists an extreme Nash equilibrium giving the maximum welfare among all Nash equilibria.*

Proof. Consider any Nash equilibrium $(\tilde{x}, \tilde{y})$, and let it be written as $\begin{bmatrix} \tilde{x} \\ \tilde{y} \end{bmatrix} = \sum_{i=1}^{r} \lambda_i \begin{bmatrix} x^i \\ y^i \end{bmatrix}$ for some set of extreme Nash equilibria $\begin{bmatrix} x^1 \\ y^1 \end{bmatrix}, \ldots, \begin{bmatrix} x^r \\ y^r \end{bmatrix}$ and $\lambda \in \triangle_r$. Observe that for any i, j,

$$x^{iT} A y^j \leq x^{jT} A y^j, x^{iT} B y^j \leq x^{iT} B y^i, \tag{65}$$

from the definition of a Nash equilibrium. Now note that

$$\begin{aligned}
\tilde{x}^T (A+B) \tilde{y} &= (\sum_{i=1}^{r} \lambda_i x^i)^T (A+B) (\sum_{i=1}^{r} \lambda_i y^i) \\
&= \sum_{i=1}^{r} \sum_{j=1}^{r} \lambda_i \lambda_j x^{iT} (A+B) y^j \\
&= \sum_{i=1}^{r} \sum_{j=1}^{r} \lambda_i \lambda_j x^{iT} A y^j + \sum_{i=1}^{r} \sum_{j=1}^{r} \lambda_i \lambda_j x^{iT} B y^j \\
&\overset{(65)}{\leq} \sum_{i=1}^{r} \sum_{j=1}^{r} \lambda_i \lambda_j x^{jT} A y^j + \sum_{i=1}^{r} \sum_{j=1}^{r} \lambda_i \lambda_j x^{iT} B y^i \\
&= \sum_{i=1}^{r} \lambda_i x^{iT} A y^i + \sum_{i=1}^{r} \lambda_i x^{iT} B y^i \\
&= \sum_{i=1}^{r} \lambda_i x^{iT} (A+B) y^i.
\end{aligned}$$

In particular, since each (x^i, y^i) is an extreme Nash equilibrium, this tells us for any Nash equilibrium $(\tilde{x}, \tilde{y})$ there must be an extreme Nash equilibrium which has at least as much welfare. □

Similarly for the results for persistent sets in Section 6.2, there is no loss in restricting attention to extreme Nash equilibria.

Lemma C.4. *For a given strategy set $\mathcal{S}$, if every extreme Nash equilibrium plays at least one strategy in $\mathcal{S}$ with positive probability, then every Nash equilibrium plays at least one strategy in $\mathcal{S}$ with positive probability.*

Proof. Let $\mathcal{S}$ be a persistent set of strategies. Since all Nash equilibria are composed of nonnegative entries, and every extreme Nash equilibrium has positive probability on some entry in $\mathcal{S}$, any convex combination of extreme Nash equilibria must have positive probability on some entry in $\mathcal{S}$. □

D A Note on Reductions from General Games to Symmetric Games

An anonymous referee asked us if our guarantees for symmetric games transfer over to general games by symmetrization. Indeed, there are reductions in the literature that take a general game, construct a symmetric game from it, and relate the Nash equilibria of the original game to symmetric Nash equilibria of its symmetrized version. In this Appendix, we review two well-known reductions of this type [23, 25] and show that the quality of approximate Nash equilibria can differ greatly between the two games. We hope that our examples can be of independent interest.

D.1 The Reduction of Griesmer, Hoffman, and Robinson [23]

Consider a game (A, B) with $A, B > 0$ and a Nash equilibrium (x^*, y^*) of it with payoffs $p_A := x^{*T} A y^*$ and $p_B := x^{*T} B y^*$. Then the symmetric game (S_{AB}, S_{AB}^T) with

$$S_{AB} := \begin{bmatrix} 0 & A \\ B^T & 0 \end{bmatrix}$$

admits a symmetric Nash equilibrium in which both players play $\begin{bmatrix} \frac{p_A}{p_A+p_B} x^* \\ \frac{p_B}{p_A+p_B} y^* \end{bmatrix}$. In the reverse direction, any symmetric equilibrium $\left(\begin{bmatrix} x \\ y \end{bmatrix}, \begin{bmatrix} x \\ y \end{bmatrix} \right)$ of (S_{AB}, S_{AB}^T) yields a Nash equilibrium $(\frac{x}{1_m^T x}, \frac{y}{1_n^T y})$ to the original game (A, B).

To demonstrate that high-quality approximate Nash equilibria in the symmetrized game can map to low-quality approximate Nash equilibria in the original game, consider the game given by $(A, B) = \left(\begin{bmatrix} \epsilon & 0 \\ 1 & 1 \end{bmatrix}, \begin{bmatrix} \epsilon^2 & 0 \\ 0 & 1 \end{bmatrix} \right)$ for some $\epsilon > 0$. The symmetric strategy

$$\left(\begin{bmatrix} \frac{1}{1+\epsilon} \\ 0 \\ \frac{\epsilon}{1+\epsilon} \\ 0 \end{bmatrix}, \begin{bmatrix} \frac{1}{1+\epsilon} \\ 0 \\ \frac{\epsilon}{1+\epsilon} \\ 0 \end{bmatrix} \right)$$

is an $\epsilon \frac{1-\epsilon}{1+\epsilon}$-NE for (S_{AB}, S_{AB}^T), but the strategy pair $\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix} \right)$ is a $(1-\epsilon)$-NE for (A, B).

D.2 The Reduction of Jurg, Jansen, Potters, and Tijs [25]

Consider a game (A, B) with $A > 0, B < 0$ and a Nash equilibrium (x^*, y^*) of it with payoffs $p_A := x^{*T} A y^*$ and $p_B := x^{*T} B y^*$. Then the symmetric game (S_{AB}, S_{AB}^T) with

$$S_{AB} := \begin{bmatrix} 0_{m\times m} & A & -1_m \\ B^T & 0_{n\times n} & 1_n \\ 1_m^T & -1_n^T & 0 \end{bmatrix}$$

admits a symmetric Nash equilibrium in which both players play

$$\begin{bmatrix} \frac{x^*}{2-p_B} \\ \frac{y^*}{2+p_A} \\ 1 - \frac{1}{2-p_B} - \frac{1}{2+p_A} \end{bmatrix}.$$

In the reverse direction, any symmetric equilibrium $\left(\begin{bmatrix} x \\ y \\ z \end{bmatrix}, \begin{bmatrix} x \\ y \\ z \end{bmatrix}\right)$ of (S_{AB}, S_{AB}^T) yields a Nash equilibrium $(\frac{x}{1_m^T x}, \frac{y}{1_n^T y})$ to the original game (A, B). This reduction has some advantages over the previous one (see [25, Section 1]).

To demonstrate that high-quality approximate Nash equilibria in the new symmetrized game can again map to low-quality approximate Nash equilibria in the original game, consider the game given by $(A, B) = \left(\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} -1 & -1 \\ 0 & 0 \end{bmatrix}\right)$. Let $\epsilon \in (0, \frac{1}{2})$. The symmetric strategy

$$\left(\begin{bmatrix} \epsilon \\ 0 \\ 1-\epsilon \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \epsilon \\ 0 \\ 1-\epsilon \\ 0 \\ 0 \end{bmatrix}\right)$$

is an $\frac{\epsilon}{2}(1-\epsilon)$-NE[12] for (S_{AB}, S_{AB}^T), but the strategy pair $\left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}\right)$ is a 1-NE for (A, B).

[12] Note that approximation factor is halved since the range of the entries of the payoff matrix in the symmetrized game is $[-1, 1]$.