SDP Relaxations for Nash Equilibria in Bimatrix Games

Jeffrey Zhang
Princeton University
Dept. of Operations Research and Financial Engineering (ORFE)

Joint work with:
Amir Ali Ahmadi
Princeton, ORFE

INFORMS Annual Meeting
October 22, 2017

Introduction to Bimatrix Games

Introduction to Bimatrix Games

1. Two payoff matrices A and B.

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

Introduction to Bimatrix Games

1. Two payoff matrices A and B.
2. The players choose strategies x and y which denote probabilities with which each player plays each row/column.

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

Introduction to Bimatrix Games

1. Two payoff matrices A and B.
2. The players choose strategies x and y which denote probabilities with which each player plays each row/column.
3. Expected payoffs will be $x^T Ay$ and $x^T By$.

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

Introduction to Bimatrix Games

1. Two payoff matrices A and B.
2. The players choose strategies x and y which denote probabilities with which each player plays each row/column.
3. Expected payoffs will be $x^T Ay$ and $x^T By$.
4. A Nash equilibrium is a pair of strategies which are a "mutual best response" to each other.

Example (Rock Paper Scissors)

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0
	Rock	Paper	Scissors
Rock	0	1	-1
Paper	-1	0	1
Scissors	1	-1	0

1. Nash equilibrium: $x = y = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
2. Not a Nash equilibrium: $x = y = (\frac{1}{2}, \frac{1}{2}, 0)$

Complexity of Nash Equilibria

Complexity of Nash Equilibria

1. Existence?

Complexity of Nash Equilibria

1. Existence?
    - Exists for every finite game!
    - Proved by Nash (1951).

Complexity of Nash Equilibria

1. Existence?
   - Exists for every finite game!
   - Proved by Nash (1951).
2. Is computationally hard to find.

Complexity of Nash Equilibria

1. Existence?
    - Exists for every finite game!
    - Proved by Nash (1951).
2. Is computationally hard to find.
    - Lemke-Howson can find, but worst-case exponential time.

Complexity of Nash Equilibria

1. Existence?
   - Exists for every finite game!
   - Proved by Nash (1951).
2. Is computationally hard to find.
   - Lemke-Howson can find, but worst-case exponential time.
3. ϵ-Approximate Nash equilibrium: players are playing strategies which give them within ϵ of their best response.

Complexity of Nash Equilibria

1. Existence?
   - Exists for every finite game!
   - Proved by Nash (1951).
2. Is computationally hard to find.
   - Lemke-Howson can find, but worst-case exponential time.
3. ϵ-Approximate Nash equilibrium: players are playing strategies which give them within ϵ of their best response.
4. Note: Any x and y form an ϵ-Nash Equilibrium where
   $$\epsilon = max(\max_i e_i^T Ay - x^T Ay, \max_j x^T Be_j - x^T By))$$

Complexity of Nash Equilibria

1. Existence?
   - Exists for every finite game!
   - Proved by Nash (1951).
2. Is computationally hard to find.
   - Lemke-Howson can find, but worst-case exponential time.
3. ϵ-Approximate Nash equilibrium: players are playing strategies which give them within ϵ of their best response.
4. Note: Any x and y form an ϵ-Nash Equilibrium where
   $$\epsilon = max(\max_i e_i^T Ay - x^T Ay, \max_j x^T Be_j - x^T By))$$
5. Approximating Nash Equilibria is also computationally hard.

QP Formulation (Nonconvex)

QP Formulation (Nonconvex)

Observation

The solutions to the following nonconvex QCQP are the Nash equilibria of the game defined by A and B:

$$\begin{aligned} \min \quad & 0 \\ \text{subject to} \quad & x^T Ay - e_i^T Ay \geq 0, \ \forall i, \\ & x^T By - x^T Be_i \geq 0, \ \forall i, \\ & x \in \triangle_m, \\ & y \in \triangle_n. \end{aligned}$$

Relaxations

Nonconvex Set ⇒ Convex Relaxation ⇒ Tightened Convex Relaxation with Valid Inequalities

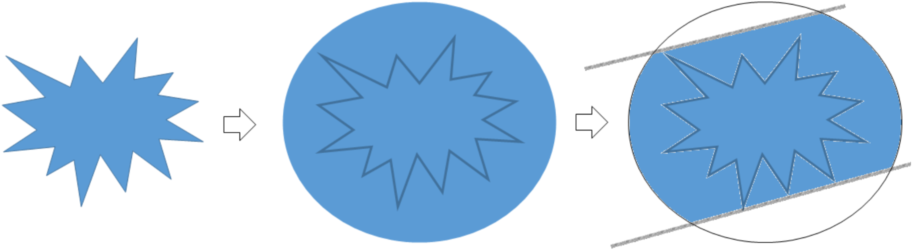

SDP Relaxation

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T = \begin{bmatrix} xx^T & xy^T & x \\ yx^T & yy^T & y \\ x^T & y^T & 1 \end{bmatrix} \qquad M := \begin{bmatrix} X & P & x \\ P^T & Y & y \\ x^T & y^T & 1 \end{bmatrix}$$

$$\begin{aligned} \min_{x,y} \quad & 0 \\ \text{subject to} \quad & x^T A y - e_i^T A y \geq 0, \\ & x^T B y - x^T B e_i \geq 0, \\ & x \in \triangle_m, \\ & y \in \triangle_n. \end{aligned} \quad \Rightarrow \quad \begin{aligned} \min_{x,y,X,Y,P} \quad & 0 \\ \text{subject to} \quad & \mathrm{Tr}(AP^T) - e_i^T A y \geq 0, \\ & \mathrm{Tr}(BP^T) - x^T B e_i \geq 0, \\ & x \in \triangle_m, \\ & y \in \triangle_n, \\ & M \succeq 0, \\ & + \text{ Valid Inequalities.} \end{aligned}$$

SDP Relaxation

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T = \begin{bmatrix} xx^T & xy^T & x \\ yx^T & yy^T & y \\ x^T & y^T & 1 \end{bmatrix} \qquad M := \begin{bmatrix} X & P & x \\ P^T & Y & y \\ x^T & y^T & 1 \end{bmatrix}$$

$$\begin{aligned} \min_{x,y} \quad & 0 \\ \text{subject to} \quad & x^T Ay - e_i^T Ay \geq 0, \\ & x^T By - x^T Be_i \geq 0, \\ & x \in \triangle_m, \\ & y \in \triangle_n. \end{aligned} \quad \Rightarrow \quad \begin{aligned} \min_{x,y,X,Y,P} \quad & 0 \\ \text{subject to} \quad & \mathrm{Tr}(AP^T) - e_i^T Ay \geq 0, \\ & \mathrm{Tr}(BP^T) - x^T Be_i \geq 0, \\ & x \in \triangle_m, \\ & y \in \triangle_n, \\ & M \succeq 0, \\ & + \text{Valid Inequalities}. \end{aligned}$$

SDP Relaxation

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}^T = \begin{bmatrix} xx^T & xy^T & x \\ yx^T & yy^T & y \\ x^T & y^T & 1 \end{bmatrix} \qquad M := \begin{bmatrix} X & P & x \\ P^T & Y & y \\ x^T & y^T & 1 \end{bmatrix}$$

$$\begin{aligned} \min_{x,y} \quad & 0 \\ \text{subject to} \quad & x^T A y - e_i^T A y \geq 0, \\ & x^T B y - x^T B e_i \geq 0, \\ & x \in \triangle_m, \\ & y \in \triangle_n. \end{aligned} \quad \Rightarrow \quad \begin{aligned} \min_{x,y,X,Y,P} \quad & 0 \\ \text{subject to} \quad & \mathrm{Tr}(AP^T) - e_i^T A y \geq 0, \\ & \mathrm{Tr}(BP^T) - x^T B e_i \geq 0, \\ & x \in \triangle_m, \\ & y \in \triangle_n, \\ & M \succeq 0, \\ & + \text{ Valid Inequalities.} \end{aligned}$$

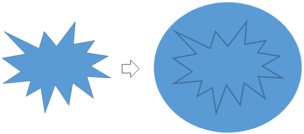

Zero-Sum Game

Definition (Zero-Sum Game)

A zero-sum game is a game in which B = -A.

Theorem (Zero-Sum Game)

This SDP recovers a Nash Equilibrium in Zero Sum games.

Approximation Quality

Theorem

Let $\lambda_1, ...\lambda_k$ be the eigenvalues of the matrix M. Then x and y are an ϵ-Nash Equilibrium with $\epsilon \leq \frac{1}{2}(m+n)\sum_{i=2}^{k} \lambda_i$.

Theorem

If the matrix M is rank-2, then a $\frac{5}{11}$-Nash Equilibrium can be recovered from the solution.

Theorem

For a symmetric game, if the matrix M is rank-2, then a symmetric ϵ-Nash Equilibrium with $\epsilon \leq \frac{1}{3}$ can be recovered from the solution.

Algorithms

Algorithms

1. In practice we apply many additional improvements to the SDP.

Algorithms

1. In practice we apply many additional improvements to the SDP.
2. Use trace of M as the objective function.

Algorithms

1. In practice we apply many additional improvements to the SDP.
2. Use trace of M as the objective function.
3. Iteratively update the objective function.

Algorithms

1. In practice we apply many additional improvements to the SDP.
2. Use trace of M as the objective function.
3. Iteratively update the objective function.

$$M := \begin{bmatrix} X & P & x \\ P^T & Y & y \\ x^T & y^T & 1 \end{bmatrix}$$

$$\begin{aligned} \text{min} \quad & \mathrm{Tr}(M) \\ \text{subject to} \quad & M \succeq 0, \\ & + \text{Valid Inequalities.} \end{aligned}$$

Linearization Algorithms

Lemma

The following nonconvex objective functions, if minimized, return rank-1 solutions:

- $\mathrm{Tr}(M) - x^T x - y^T y$
- $\sum_{i=1}^{m+n} \sqrt{M_{i,i}}$

Linearization Algorithms

Lemma

The following nonconvex objective functions, if minimized, return rank-1 solutions:

- $\mathrm{Tr}(M) - x^T x - y^T y$
- $\sum_{i=1}^{m+n} \sqrt{M_{i,i}}$

1. We iteratively update the objective functions based on a linearization of those functions "*Diagonal Gap*" and "*Diagonal Square Root*").

Linearization Algorithms

Lemma

The following nonconvex objective functions, if minimized, return rank-1 solutions:

- $\mathrm{Tr}(M) - x^T x - y^T y$
- $\sum_{i=1}^{m+n} \sqrt{M_{i,i}}$

1. We iteratively update the objective functions based on a linearization of those functions "*Diagonal Gap*" and "*Diagonal Square Root*").
2. Solve SDP with $\mathrm{Tr}(M)$ as objective.
   while !convergence **do**
       Solve SDP with updated objective function.
   end while

Properties of Algorithm

Theorem

The diagonal gap linearization algorithm produces a sequence of

$$\mathrm{Tr}(M) - x^T x - y^T y$$

which is nonincreasing and lower bounded by 1. If it reaches 1, then an exact Nash equilibrium can be recovered from the solution.

Theorem

The diagonal square root linearization algorithm produces a sequence of

$$\sum_{i=1}^{m+n} \sqrt{M_{i,i}}$$

which is nonincreasing and lower bounded by 2. If it reaches 2, then an exact Nash equilibrium can be recovered from the solution.

Improvements of ϵ Through Iterations

Histogram of ϵ for 100 20x20 Games (Diagonal Gap)

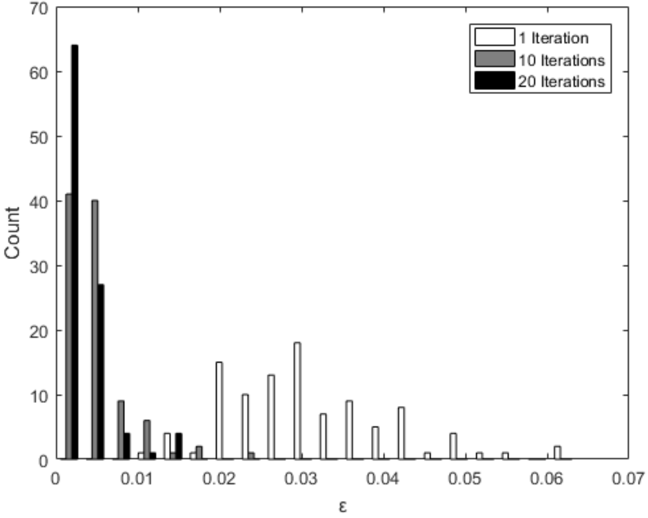

Improvements of ϵ Through Iterations

Histogram of ϵ for 100 20x20 Games (Diagonal Square Root)

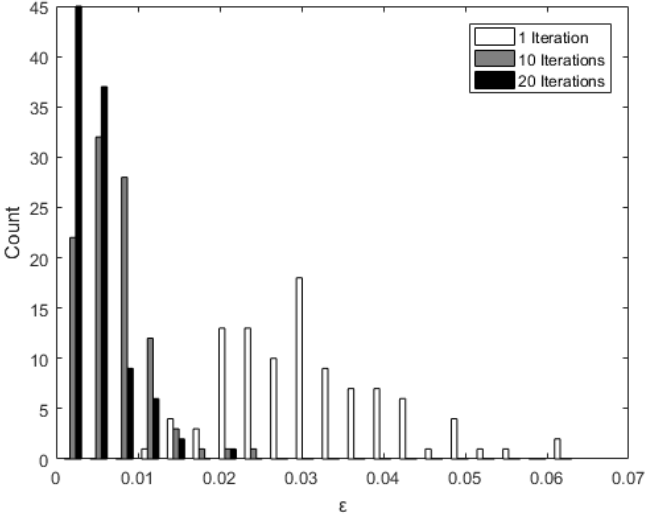

Other Applications of SDP: Strategy Exclusion

Other Applications of SDP: Strategy Exclusion

1. Other interesting questions from an economic perspective:

Other Applications of SDP: Strategy Exclusion

1. Other interesting questions from an economic perspective:
2. Often we seek Nash equilibria with certain properties, or find out whether they exist.

Other Applications of SDP: Strategy Exclusion

1. Other interesting questions from an economic perspective:
2. Often we seek Nash equilibria with certain properties, or find out whether they exist.
3. Given a strategy, is there any Nash Equilibrium in which that strategy is played?

Other Applications of SDP: Strategy Exclusion

1. Other interesting questions from an economic perspective:
2. Often we seek Nash equilibria with certain properties, or find out whether they exist.
3. Given a strategy, is there any Nash Equilibrium in which that strategy is played?
4. This is NP-hard to decide.

Other Applications of SDP: Strategy Exclusion

1. Other interesting questions from an economic perspective:
2. Often we seek Nash equilibria with certain properties, or find out whether they exist.
3. Given a strategy, is there any Nash Equilibrium in which that strategy is played?
4. This is NP-hard to decide.
5. Can be solved through nonconvex QP, which we can also relax with SDP.

Other Applications of SDP: Strategy Exclusion

1. Other interesting questions from an economic perspective:
2. Often we seek Nash equilibria with certain properties, or find out whether they exist.
3. Given a strategy, is there any Nash Equilibrium in which that strategy is played?
4. This is NP-hard to decide.
5. Can be solved through nonconvex QP, which we can also relax with SDP.

Game Size	5×5	10×10
Number of Strategies	1000	2000
Number Correct	996	2000

Other Applications of SDP: Maximum Welfare

Other Applications of SDP: Maximum Welfare

1. We might also seek a Nash equilibrium with high social welfare.
    - Welfare in any Nash Equilibrium - the sum of the payoffs.

Other Applications of SDP: Maximum Welfare

1. We might also seek a Nash equilibrium with high social welfare.
    - Welfare in any Nash Equilibrium - the sum of the payoffs.
2. This quantity is NP-hard to find.

Other Applications of SDP: Maximum Welfare

1. We might also seek a Nash equilibrium with high social welfare.
   - Welfare in any Nash Equilibrium - the sum of the payoffs.
2. This quantity is NP-hard to find.
3. Can be solved through nonconvex QP, which we can also relax with SDP.

True Maximum vs SDP Approximation (10×10 games)

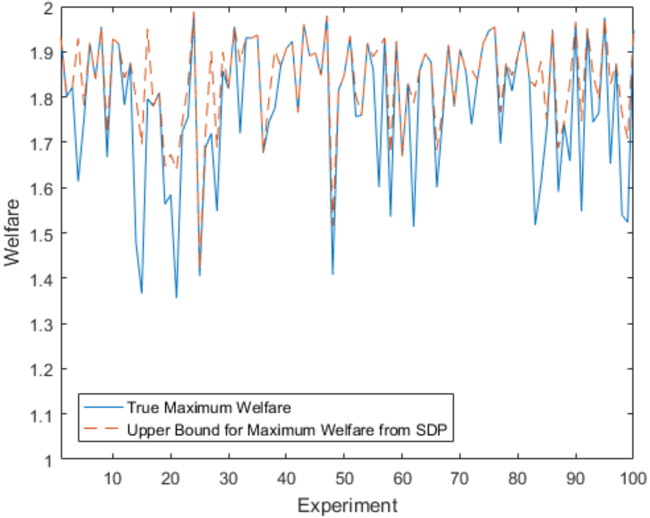

Thank You!

For details see
https://arxiv.org/abs/1706.08550