Epsilon Approximate Equilibria with Multiple Payoffs and No Priors

Anshul Sawant[1], John Dickerson[2], Mohammad T. Hajiaghayi[1], V. S. Subrahmanian[1]

[1]UMIACS
University of Maryland
College Park, MD 20740

[2]Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

2012-11-14 Wed

Problem Overview

Analyze a simultaneous normal form game with multiple payoff matrices.

Motivating Example

- Consider a real world situation modeled as a simultaneous normal form game.
- Five players:
  - US
  - India
  - Pakistani Government
  - Pakistani Military
  - Lashkar - e - Taiba
- In addition, we are given a set of actions each player can take.
- We are interested in analyzing this game to find target equilibria in which the terrorist group's terrorist acts could be significantly reduced. For that we need *payoff matrices*.

Motivating Example

- We approached three subject matter experts.
- They gave us significantly diverging payoffs leading to widely differing sets of Nash Equilibria.
- Experts cannot be too sure about payoffs.

Motivating Example

- We approached three subject matter experts.
- They gave us significantly diverging payoffs leading to widely differing sets of Nash Equilibria.
- Experts cannot be too sure about payoffs.
- How to analyze a game when:

Motivating Example

- We approached three subject matter experts.
- They gave us significantly diverging payoffs leading to widely differing sets of Nash Equilibria.
- Experts cannot be too sure about payoffs.
- How to analyze a game when:
  - Payoffs are uncertain

Motivating Example

- We approached three subject matter experts.
- They gave us significantly diverging payoffs leading to widely differing sets of Nash Equilibria.
- Experts cannot be too sure about payoffs.
- How to analyze a game when:
    - Payoffs are uncertain
    - There are many payoff matrices

Motivating Example

- We approached three subject matter experts.
- They gave us significantly diverging payoffs leading to widely differing sets of Nash Equilibria.
- Experts cannot be too sure about payoffs.
- How to analyze a game when:
  - Payoffs are uncertain
  - There are many payoff matrices
  - It is difficult to statistically model this uncertainty

Motivating Example

- We approached three subject matter experts.
- They gave us significantly diverging payoffs leading to widely differing sets of Nash Equilibria.
- Experts cannot be too sure about payoffs.
- How to analyze a game when:
  - Payoffs are uncertain
  - There are many payoff matrices
  - It is difficult to statistically model this uncertainty
  - The output of analysis can lead to policy decisions that can significantly influence allocation of security resources

A Word on Notation

- In the following slides $u_j(\sigma)$ denotes payoff of player j where σ is a *strategy profile* (defined later).
- $u_j(\sigma^{j'}, \sigma_{-j})$ denotes payoff of player j when j's *strategy* (defined later) is σ'^j and rest of players play strategy given by strategy profile σ.
- A vector e_k represent's a vector whose k^{th} component is 1 and all other components are 0.
- Set of players is denoted by $[n] = \{1..n\}$. Set of actions is denoted by $[m] = \{1..m\}$.

Some Basic Concepts

- A *strategy* is a distribution over actions of players.
- A *strategy profile* specifies a strategy for each player.
- A *t-uniform strategy* is one where all probabilities over actions are multiples of $\frac{1}{t}$.

Definition

A strategy profile σ is a Nash equilibrium iff:

$$u_j(\sigma^{j\prime}, \sigma^{-j}) \leq u_j(\sigma), \forall \sigma^{j\prime} \in \Delta_m, j \in [n]$$

- Unilateral deviation does not lead to any gains for a player.
- Finding Nash equilibrium is PPAD-complete. So, we need some notion of approximate Nash equilibrium.

Some Basic Concepts

Definition

A strategy profile σ is an ϵ-approximate Nash equilibrium for some $0 \leq \epsilon \leq 1$ iff:

$$u_j(\sigma^{j\prime}, \sigma^{-j}) \leq u_j(\sigma) + \epsilon, \forall \sigma^{j\prime} \in \Delta_m, j \in [n]$$

- Unilateral deviation leads to at most ϵ gain over expected payoff.
- A stricter notion of approximate Nash equilibrium is the approximate well supported Nash equilibrium.

Definition

A strategy profile σ is an ϵ-well supported approximate Nash equilibrium for $0 \leq \epsilon \leq 1$ iff:

$$u_j(e_k, \sigma^{-j}) \leq u_j(e_l, \sigma^{-j}) + \epsilon, \forall \sigma^{j\prime} \in \Delta_m, k \in [m], \\ l \in S(\sigma^j), j \in [n]$$

Some Basic Concepts - Rank of Multiplayer Game

- A rank-1 tensor of n-dimensions can be expressed as tensor product of n vectors.
  For example, let $v^i, i \in \{1..n\}$ be n vectors, each of length l. A rank-1 tensor T is:

$$T[i_1, i_2, \ldots, i_n] = \prod_{k \in \{1..l\}} v^k_{i_k}$$

- Rank of a tensor T is minimum number of rank 1 tensors that sum to T.
- A multiplayer game of rank-k is a game where payoff function for any player is a tensor of rank at most k.

Our Contribution

- We extend the concept of approximate Nash equilibrium to the case of games with multiple payoffs. Informally, this is the set of approximate Nash equilibria that are common to all payoff matrices.

  I will call it Multiple payoff ϵ-Approximate Nash Equilibrium (MEAE) for this presentation.

Our Contribution

- We extend the concept of approximate Nash equilibrium to the case of games with multiple payoffs. Informally, this is the set of approximate Nash equilibria that are common to all payoff matrices.

  I will call it Multiple payoff ϵ-Approximate Nash Equilibrium (MEAE) for this presentation.
- For two-player zero-sum games, we show a linear program whose solutions capture the set of MEAE.

Our Contribution

- We extend the concept of approximate Nash equilibrium to the case of games with multiple payoffs. Informally, this is the set of approximate Nash equilibria that are common to all payoff matrices.

  I will call it Multiple payoff ϵ-Approximate Nash Equilibrium (MEAE) for this presentation.
- For two-player zero-sum games, we show a linear program whose solutions capture the set of MEAE.
- For multi-player games of rank-1
  - We prove a necessary and sufficient condition for a strategy profile to be an approximate Nash Equilibrium.
  - We then prove a necessary and sufficient condition for a strategy profile to be an MEAE.

Our Contribution

- We extend the concept of approximate Nash equilibrium to the case of games with multiple payoffs. Informally, this is the set of approximate Nash equilibria that are common to all payoff matrices.

  I will call it Multiple payoff ϵ-Approximate Nash Equilibrium (MEAE) for this presentation.
- For two-player zero-sum games, we show a linear program whose solutions capture the set of MEAE.
- For multi-player games of rank-1
  - We prove a necessary and sufficient condition for a strategy profile to be an approximate Nash Equilibrium.
  - We then prove a necessary and sufficient condition for a strategy profile to be an MEAE.
- For multi-player games of rank-2 we prove conditions under which an MEAE with t-uniform strategy profiles exists and give an algorithm to compute the same.

Past Work

- Our work requires enumeration of Nash Equilibria. It is expected to be a hard problem.
  - Dsakalakis et al. proved that computation of Nash equilibrium for 3-player games is a hard problem.
  - Chen et al. proved that computation of Nash equilibrium for 2-player games is a hard problem.
  - Chen et al. proved that it is unlikely that an FPTAS exists for computing Nash equilibrium.
  - Constant factor multiplicative approximation is PPAD-complete [Daskalakis].
- There is a QPTAS due to Althofer and Lipton et al. based on brute force search over some uniform strategies.
- Kalyanaraman and Umans define constant rank multiplayer games and give a PTAS to find some approximate Nash equilibria of the game.

Summary of Techniques

- Zero-sum games
  - We modify the zero-sum game LP to compute approximate Nash equilibria, given a payoff.
  - Then we search for approximate Nash equilibria by varying payoffs.

Summary of Techniques

- Zero-sum games
  - We modify the zero-sum game LP to compute approximate Nash equilibria, given a payoff.
  - Then we search for approximate Nash equilibria by varying payoffs.
- For rank-1 games
  - We prove that all strategies in support are approximately dominating strategies
  - Any approximately dominating strategy can be in the support of Nash equilibrium strategy.
  - Checking for dominating strategies is easy.

Summary of Techniques

- Zero-sum games
  - We modify the zero-sum game LP to compute approximate Nash equilibria, given a payoff.
  - Then we search for approximate Nash equilibria by varying payoffs.
- For rank-1 games
  - We prove that all strategies in support are approximately dominating strategies
  - Any approximately dominating strategy can be in the support of Nash equilibrium strategy.
  - Checking for dominating strategies is easy.
- For rank-k games with a small number of actions and players
  - We prove that if a common approximate equilibrium exists for all payoff matrices, then an approximate t-uniform equilibrium will also exist with some loss of approximation factor.
  - Searching over t-uniform strategies leads to a PTAS.

Zero-sum Game LP

- Primal

$$\text{maximize: } r$$

$$\text{s.t.:}$$

$$\sum_{i \in [m]} \sigma_i^1 = 1$$

$$\sigma_i^1 \geq 0, \forall i \in [m]$$

$$\sum_{i \in [m]} \sigma_i^1 u(e_i, e_j) \geq r, \forall j \in [m]$$

Zero-sum Game LP

- Primal

$$\text{maximize: } r$$
$$\text{s.t.:}$$
$$\sum_{i \in [m]} \sigma_i^1 = 1$$
$$\sigma_i^1 \geq 0, \forall i \in [m]$$
$$\sum_{i \in [m]} \sigma_i^1 u(e_i, e_j) \geq r, \forall j \in [m]$$

- Dual

$$\text{minimize: } s$$
$$\text{s.t.:}$$
$$\sum_{i \in [m]} \sigma_i^2 = 1$$
$$\sigma_i^2 \geq 0, \forall i \in [m]$$
$$\sum_{j \in [m]} \sigma_j^2 u(e_i, e_j) \leq s, \forall i \in [m]$$

Zero-sum Game - approximate equilibria LP

- It is not straight-forward to modify the previous LP for computing approximate equilbria

Zero-sum Game - approximate equilibria LP

- It is not straight-forward to modify the previous LP for computing approximate equilbria
  - The previous LP works because of minimax theorem.

Zero-sum Game - approximate equilibria LP

- It is not straight-forward to modify the previous LP for computing approximate equilbria
  - The previous LP works because of minimax theorem.
  - When we allow for approximate equilibria, we no longer have a unique minimax payoff.

Zero-sum Game - approximate equilibria LP

- It is not straight-forward to modify the previous LP for computing approximate equilbria
  - The previous LP works because of minimax theorem.
  - When we allow for approximate equilibria, we no longer have a unique minimax payoff.
  - We cannot have bidirectional constraints on a single player's payoffs because that requires us to guess the support.

Zero-sum Game - approximate equilibria LP

- It is not straight-forward to modify the previous LP for computing approximate equilbria
  - The previous LP works because of minimax theorem.
  - When we allow for approximate equilibria, we no longer have a unique minimax payoff.
  - We cannot have bidirectional constraints on a single player's payoffs because that requires us to guess the support.
- We need to combine primal and dual into a single LP to avoid bidirectional constraints

Zero-sum Game - approximate equilibria LP

- The following LP gives ϵ-approximate Nash equilibrium for the game

$$\sum_{i\in[m]} \sigma_i^1 = 1$$

$$\sigma_i^1 \geq 0, \forall i \in [m]$$

$$\sum_{i\in[m]} \sigma_i^2 = 1$$

$$\sigma_i^2 \geq 0, \forall i \in [m]$$

$$\sum_{i\in[m]} \sigma_i^1 u(e_i, e_j) \geq r - \epsilon, \forall j \in [m]$$

$$\sum_{j\in[m]} \sigma_j^2 u(e_i, e_j) \leq r + \epsilon, \forall i \in [m]$$

Nash Equilibrium in Rank-1 Multiplayer Games

Let $\alpha_k^{i,j}$ be value of player i's action k to player j. Then, for a pure strategy profile where player i plays action a_i then utility to player j is equal to $\prod_{i\in[n]} \alpha_{a_i}^{i,j}$.

Nash Equilibrium in Rank-1 Multiplayer Games

Let $\alpha_k^{i,j}$ be value of player i's action k to player j. Then, for a pure strategy profile where player i plays action a_i then utility to player j is equal to $\prod_{i \in [n]} \alpha_{a_i}^{i,j}$.

- For a pure strategy profile, any action a such that $a \in argmax_{k \in [m]} \alpha_k^{j,j}$ can be in support of a Nash Equilibrium.
- This extends to mixed strategy profiles also.

Nash Equilibrium in Rank-1 Multiplayer Games

Let $\alpha_k^{i,j}$ be value of player i's action k to player j. Then, for a pure strategy profile where player i plays action a_i then utility to player j is equal to $\prod_{i \in [n]} \alpha_{a_i}^{i,j}$.

- For a pure strategy profile, any action a such that $a \in argmax_{k \in [m]} \alpha_k^{j,j}$ can be in support of a Nash Equilibrium.
- This extends to mixed strategy profiles also.

This gives us the following result:

Lemma

Let σ be a mixed strategy profile. Let $u'_{-j}(\sigma) = \prod_{i \in [n]-\{j\}} (\sum_{l \in [m]} \sigma_l^i \alpha_l^{i,j})$. Let support of player j's strategy be $S_j = \{l | \alpha_l^{j,j} = \max(\alpha^{j,j})\}$. σ is a Nash equilibrium iff:

$$u'_{-j}(\sigma) > 0 \implies support(\sigma^j) \subseteq S_j$$

By a very similar line of resoning, we have the following result for approximate Nash equilibrium:

Lemma

A strategy profile $\sigma =< \sigma^1, \sigma^2, \ldots, \sigma^n >$ is an ϵ-well supported approximate (multipicative) Nash equilibrium (with a non-zero payoff for all players) for a rank-1 multiplayer game with payoffs as specified previously iff:

$$\alpha_i^{j,j} \geq (1-\epsilon)(\max \alpha^{j,j}), \forall j \in [n], i \in S(\sigma^j)$$

Since, this is complete characterization of approximate Nash equilibrium for such games, extension to case of multiple payoffs is straightforward.

MEAE in Rank-1 Multiplayer Games

Theorem

Consider an simultaneous game with multiple payoffs of rank 1. For an game, we have f different payoff functions for each player. For $t \in [f]$, let payoff function t for player j be specified by tuple $(\alpha^{1,j,t}, \alpha^{2,j,t}, \ldots, \alpha^{n,j,t})$. Then, a strategy profile $\sigma = (\sigma^1, \sigma^2, \ldots, \sigma^n)$ is an MEAE, iff:

$$\alpha_i^{j,j,t} \geq (\max \alpha^{j,j,t})(1 - \epsilon), \forall j \in [n], i \in S(\sigma^j), t \in [f] \quad (1)$$

Approximate N.E. and Uniform Strategy Profiles

The main result for these games follows from the following result:

Lemma

Let the strategy profile $\sigma = (\sigma^1, \sigma^2, .., \sigma^n)$ *be a well-supported* ϵ*-approximate Nash equilibrium for the given game of rank* k*. Then there exists a* t*-uniform strategy profile* σ' *that is a well-supported* $\epsilon + \frac{2(n-1)mk}{t}$*-approximate Nash equilibrium.*

Approximate N.E. and Uniform Strategy Profiles

The main result for these games follows from the following result:

Lemma

Let the strategy profile $\sigma = (\sigma^1, \sigma^2, .., \sigma^n)$ *be a well-supported* ϵ*-approximate Nash equilibrium for the given game of rank* k*. Then there exists a t-uniform strategy profile* σ' *that is a well-supported* $\epsilon + \frac{2(n-1)mk}{t}$*-approximate Nash equilibrium.*

The intuition behind the above result is:

Approximate N.E. and Uniform Strategy Profiles

The main result for these games follows from the following result:

Lemma

Let the strategy profile $\sigma = (\sigma^1, \sigma^2, .., \sigma^n)$ be a well-supported ϵ-approximate Nash equilibrium for the given game of rank k. Then there exists a t-uniform strategy profile σ' that is a well-supported $\epsilon + \frac{2(n-1)mk}{t}$-approximate Nash equilibrium.

The intuition behind the above result is:

- Due to low rank, payoff function is succinctly representable as a sum of products of a "few" terms.

Approximate N.E. and Uniform Strategy Profiles

The main result for these games follows from the following result:

Lemma

Let the strategy profile $\sigma = (\sigma^1, \sigma^2, .., \sigma^n)$ be a well-supported ϵ-approximate Nash equilibrium for the given game of rank k. Then there exists a t-uniform strategy profile σ' that is a well-supported $\epsilon + \frac{2(n-1)mk}{t}$-approximate Nash equilibrium.

The intuition behind the above result is:

- Due to low rank, payoff function is succinctly representable as a sum of products of a "few" terms.
- Therefore, the effect of perturbing the strategy profiles to make them uniform is much less than it would be for a full rank game.

MEAE for Rank-k Games

Theorem

Let the strategy profile σ be an MEAE with $\epsilon = \tau$ for the given game with multiple payoffs, all of whose constituent games are rank k games. Then, there exists a t-uniform strategy profile σ' that is an MEAE, with $\epsilon = \tau + \frac{2(n-1)mk}{t}$.