CS599: Algorithm Design in Strategic Settings
Fall 2012
Lecture 2: Game Theory Preliminaries

Instructor: Shaddin Dughmi

Administrivia

- Website: http://www-bcf.usc.edu/~shaddin/cs599fa12
  - Or go to www.cs.usc.edu/people/shaddin and follow link
- Emails?
- Registration

Outline

1. Games of Complete Information
2. Games of Incomplete Information
   - Prior-free Games
   - Bayesian Games

Outline

1. Games of Complete Information

2. Games of Incomplete Information
   - Prior-free Games
   - Bayesian Games

Example: Rock, Paper, Scissors

Figure: Rock, Paper, Scissors

Rock, Paper, Scissors is an example of the most basic type of game.

Simultaneous move, complete information games

- Players act simultaneously
- Each player incurs a utility, determined only by the players' (joint) actions. Equivalently, player actions determine "state of the world" or "outcome of the game".
- The payoff structure of the game, i.e. the map from action vectors to utility vectors, is common knowledge

Standard mathematical representation of such games:

Normal Form

A game in normal form is a tuple (N, A, u), where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $A = A_1 \times \ldots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \ldots, a_n) \in A$ is called an action profile.
- $u = (u_1, \ldots u_n)$, where $u_i : A \to \mathbb{R}$ is the utility function of player i.

Standard mathematical representation of such games:

Normal Form

A game in normal form is a tuple (N, A, u), where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $A = A_1 \times \ldots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \ldots, a_n) \in A$ is called an action profile.
- $u = (u_1, \ldots u_n)$, where $u_i : A \to \mathbb{R}$ is the utility function of player i.

- Typically thought of as an n-dimensional matrix, indexed by $a \in A$, with entry $(u_1(a), \ldots, u_n(a))$.
- Also useful for representing more general games, like sequential and incomplete information games, but is less natural there.

Figure: Generic Normal Form Matrix

Strategies in Normal Form Games

It will be convenient down the line to distinguish actions from strategies

- Strategies of player i
  - Pure strategy: a choice of action $a_i \in A_i$
    - Example: rock
  - Mixed strategy: a choice of distribution over actions.
    - Example: uniformly randomly choose one of rock, paper, scissors
- Let S_i, $\overline{S}_i$ denote the set of mixed and pure strategies of player i, respectively.
  - $S = S_1 \times \ldots \times S_n$ is the set of mixed strategy profiles (similarly, $\overline{S}$)
  - For strategy $s \in S_i$ and $a \in A_i$, let $s(a)$ denote the probability of action a in strategy s.
- Extending utilities to mixed strategies:
  - $u_i(s_1, \ldots, s_n) = \sum_{a \in A} u_i(a) \prod_{j=1}^{n} s_j(a_j)$

Example: Prisoner's Dilemma

Figure: Prisoner's Dilemma

Example: Battle of the Sexes

Figure: Battle of the Sexes

Example: First Price Auction

Two players, with values $v_1 = 1$ and $v_2 = 2$, both common knowledge.

- $A_1 = A_2 = \mathbb{R}$ (note: infinite!)
- $u_i(a_1, a_2) = v_i - a_i$ if $a_i > a_{-i}$, and 0 otherwise.

Aside: Sequential Games

But . . .

what about “sequential” games like the english auction, chess, etc?

- More naturally modeled using the extensive form tree representation
  - Each non-leaf node is a step in the game, associated with a player
  - Outgoing edges = actions available at that step
  - leaf nodes labelled with utility of each player
  - Pure strategy: choice of action for each contingency (i.e. each non-leaf node)
- Can be represented as a normal form game by collapsing pure strategies to actions of a large normal form game
- In any case, the revelation principle suggests that simultaneous move games often suffice in mechanism design.

Solution Concepts

A solution concept identifies, for every game, some strategy profiles of interest. Solution concepts either serve as a prediction of the outcome of the game, or as a way of identifying desirable outcomes.

Examples

- Welfare maximizing outcome
- Pareto optimal outcome
- 2-approximately welfare maximizing outcome
- Pure Nash equilibrium
- Mixed Nash equilibrium
- Dominant Strategy equilibrium
- Others: undominated strategies, rationalizable equilibrium, iterated removal ...

Figure: Prisoners' Dilemma

Nash Equilibrium

A mixed strategy $s_i \in S_i$ of player i is a best response to a mixed strategy profile s_{-i} of the other players if $u_i(s) \geq u_i(s'_i, s_{-i})$ for every other possible strategy s'_i.

- Note: There is always a pure best response
- The set of mixed best responses is the randomizations over pure best responses.

Nash Equilibrium

A mixed strategy $s_i \in S_i$ of player i is a best response to a mixed strategy profile s_{-i} of the other players if $u_i(s) \geq u_i(s'_i, s_{-i})$ for every other possible strategy s'_i.

- Note: There is always a pure best response
- The set of mixed best responses is the randomizations over pure best responses.

A Mixed Nash equilibrium is a mixed strategy profile $s \in S$ such that, for each player i, s_i is a best response to s_{-i}. If $s \in \overline{S}$, then it is a pure Nash equilibrium.

Dominant-strategy Equilibrium

Some games admit a very special kind of equilibrium, where one strategy profile "dominates"

A mixed strategy $s_i \in S_i$ of player i is a dominant strategy if it is a best response for every mixed strategy (equivalently, every pure strategy) s_{-i} of the other players.

- Note: If there is a mixed dominant strategy, then there is a pure dominant strategy
- The set of mixed dominant strategies is the set of randomizations over pure dominant strategies

Dominant-strategy Equilibrium

Some games admit a very special kind of equilibrium, where one strategy profile "dominates"

A mixed strategy $s_i \in S_i$ of player i is a dominant strategy if it is a best response for every mixed strategy (equivalently, every pure strategy) s_{-i} of the other players.

- Note: If there is a mixed dominant strategy, then there is a pure dominant strategy
- The set of mixed dominant strategies is the set of randomizations over pure dominant strategies

A (pure/mixed) dominant-strategy equilibrium is a strategy profile where each player plays a dominant strategy.

- Every dominant strategy equilibrium is also a Nash equilibrium

Example: prisoner's dillemma

Existence of Equilibria

- Pure Nash equilibria and Dominant strategy equilibria do not always exist (e.g. rock paper scissors)
- However, mixed Nash equilibrium always exists!

Theorem (Nash 1951)

Every finite game admits a mixed Nash equilibrium.

Existence of Equilibria

- Pure Nash equilibria and Dominant strategy equilibria do not always exist (e.g. rock paper scissors)
- However, mixed Nash equilibrium always exists!

Theorem (Nash 1951)

Every finite game admits a mixed Nash equilibrium.

Note: generalizes to infinite continuous games

Existence of Equilibria

- Pure Nash equilibria and Dominant strategy equilibria do not always exist (e.g. rock paper scissors)
- However, mixed Nash equilibrium always exists!

Theorem (Nash 1951)

Every finite game admits a mixed Nash equilibrium.

Note: generalizes to infinite continuous games
Example: battle of the sexes. (solve in class)

Outline

- In settings of complete information, Nash equilibria are a defensible prediction of the outcome of the game.
- In many settings, as in auctions, the payoff structure of the game itself is private to the players.
- How can a player possibly play his part of the Nash equilibrium if he's not sure what the game is, and therefore where the equilibrium is?
    - i.e. the set of Nash equilibria depends on opponents' private information.

- In settings of complete information, Nash equilibria are a defensible prediction of the outcome of the game.
- In many settings, as in auctions, the payoff structure of the game itself is private to the players.
- How can a player possibly play his part of the Nash equilibrium if he's not sure what the game is, and therefore where the equilibrium is?
  - i.e. the set of Nash equilibria depends on opponents' private information.

Example

Example: First price auction $v_1 = 3$, v_2 is either 1 or 2. In both cases, Nash equilibrium bids are $b_1 = b_2 = v_2$ (unique if with small probability we give the item to each player at his bid). Player 1's equilibrium bid depends on player 2's private information!

- In settings of complete information, Nash equilibria are a defensible prediction of the outcome of the game.
- In many settings, as in auctions, the payoff structure of the game itself is private to the players.
- How can a player possibly play his part of the Nash equilibrium if he's not sure what the game is, and therefore where the equilibrium is?
  - i.e. the set of Nash equilibria depends on opponents' private information.

Example

Example: First price auction $v_1 = 3$, v_2 is either 1 or 2. In both cases, Nash equilibrium bids are $b_1 = b_2 = v_2$ (unique if with small probability we give the item to each player at his bid). Player 1's equilibrium bid depends on player 2's private information!

To explicitly model uncertainty, and devise credible solution concepts that take it into account, games of incomplete information were defined.

Modeling Uncertainty

Two main approaches are used to model uncertainty:

1. Prior-free:
   - A player doesn't have any beliefs about the private data of others (other than possible values it may take), and therefore about their strategies.
   - Only consider a strategy to be a "credible" prediction for a player if it is a best response in every possible situation.
2. Bayesian Common Prior:
   - Players' private data is drawn from a distribution, which is common knowledge
   - Player only knows his private data, but knows the distribution of others'
   - Bayes-Nash equilibrium generalizes Nash to take into account the distribution.

Though there are other approaches...

Prior-free Games

A game of strict incomplete information is a tuple (N, A, T, u), where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $A = A_1 \times \ldots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \ldots, a_n) \in A$ is called an action profile.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times A \to \mathbb{R}$ is the utility function of player i.

Prior-free Games

A game of strict incomplete information is a tuple (N, A, T, u), where

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $A = A_1 \times \ldots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \ldots, a_n) \in A$ is called an action profile.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times A \to \mathbb{R}$ is the utility function of player i.

Example: Vickrey Auction

- $A_i = \mathbb{R}$ is the set of possible bids of player i.
- $T_i = \mathbb{R}$ is the set of possible values for the item.
- For $v_i \in T_i$ and $b \in A$, we have $u_i(v_i, b) = v_i - b_{-i}$ if $b_i > b_{-i}$, otherwise 0.

Strategies in Incomplete Information Games

- Strategies of player i
  - Pure strategy $s_i : T_i \to A_i$: a choice of action $a_i \in A_i$ for every type $t_i \in T_i$.
    - Example: Truthtelling is a strategy in the Vickrey Auction
    - Example: Bidding half your value is also a strategy
  - Mixed strategy: a choice of distribution over actions A_i for each type $t_i \in T_i$
    - Won't really use... all our applications will involve pure strategies

Strategies in Incomplete Information Games

- Strategies of player i
  - Pure strategy $s_i : T_i \to A_i$: a choice of action $a_i \in A_i$ for every type $t_i \in T_i$.
    - Example: Truthtelling is a strategy in the Vickrey Auction
    - Example: Bidding half your value is also a strategy
  - Mixed strategy: a choice of distribution over actions A_i for each type $t_i \in T_i$
    - Won't really use... all our applications will involve pure strategies

Note

In a strategy, player decides how to act based only on his private info (his type), and NOT on others' private info nor their actions.

Dominant Strategy Equilibrium

$s_i : T_i \to A_i$ is a dominant strategy for player i if, for all $t_i \in T_i$ and $a_{-i} \in A_{-i}$ and $a'_i \in A_i$,

$$u_i(t_i, (s_i(t_i), a_{-i})) \geq u_i(t_i, (a'_i, a_{-i}))$$

Equivalently: $s_i(t_i)$ is a best response to $s_{-i}(t_{-i})$ for all t_i, t_{-i} and s_{-i}.

Illustration: Vickrey Auction

Vickrey Auction

Consider a Vickrey Auction with incomplete information.

Illustration: Vickrey Auction

Vickrey Auction

Consider a Vickrey Auction with incomplete information.

Claim

The truth-telling strategy is dominant for each player.

Prove in class

Bayesian Games

A Bayesian game of Incomplete information is a tuple (N, A, T, u, p),

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $A = A_1 \times \ldots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \ldots, a_n) \in A$ is called an action profile.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times A \to \mathbb{R}$ is the utility function of player i.
- $\mathcal{D}$ is a distribution over T.

Bayesian Games

A Bayesian game of Incomplete information is a tuple (N, A, T, u, p),

- N is a finite set of players. Denote $n = |N|$ and $N = \{1, \ldots, n\}$.
- $A = A_1 \times \ldots A_n$, where A_i is the set of actions of player i. Each $\vec{a} = (a_1, \ldots, a_n) \in A$ is called an action profile.
- $T = T_1 \times \ldots T_n$, where T_i is the set of types of player i. Each $\vec{t} = (t_1, \ldots, t_n) \in T$ is called an type profile.
- $u = (u_1, \ldots u_n)$, where $u_i : T_i \times A \to \mathbb{R}$ is the utility function of player i.
- $\mathcal{D}$ is a distribution over T.

Example: First Price Auction

- $A_i = T_i = [0, 1]$
- $\mathcal{D}$ draws each $v_i \in T_i$ uniformly and independently from $[0, 1]$.
- $u_i(v_i, b) = v_i - b_i$ if $b_i \geq b_{-i}$, otherwise 0.

Bayes-Nash Equilibrium

As before, a strategy s_i for player i is a map from T_i to A_i. Now, we define the extension of Nash equilibrium to this setting.

A pure Bayes-Nash Equilibrium of a Bayesian Game of incomplete information is a set of strategies $s_1, \ldots, s_n$, where $s_i : T_i \to A_i$, such that for all i, $t_i \in T_i$, $a'_i \in A_i$ we have

$$\mathop{\mathbf{E}}_{t_{-i} \sim \mathcal{D}|t_i} u_i(t_i, s(t)) \geq \mathop{\mathbf{E}}_{t_{-i} \sim \mathcal{D}|t_i} u_i(t_i, (a'_i, s_{-i}(t_{-i})))$$

where the expectation is over t_{-i} drawn from p after conditioning on t_i.

- Note: Every dominant strategy equilibrium is also a Bayes-Nash Equilibrium
- But, unlike DSE, BNE is guaranteed to exist.

Example: First Price Auction

Example: First Price Auction

- $A_i = T_i = [0, 1]$
- $u_i(v_i, b) = v_i - b_{(1)}$ if $v_i = b_{(1)}$, otherwise 0.
- $\mathcal{D}$ draws each $v_i \in T_i$ independently from $[0, 1]$.

Show that the strategies $b_i(v_i) = v_i/2$ form a Bayes-Nash equilibrium.

Existence of Bayes-Nash Equilibrium

Theorem

Every finite Bayesian game of incomplete information admits a mixed Bayes-Nash equilibrium.