Mechanism Design & Auctions

Game Theoretic Analysis

Kevin Leyton-Brown

University of British Columbia
Canada CIFAR AI Chair, Amii

James R. Wright

University of Alberta
Canada CIFAR AI Chair, Amii

Lecture Overview

Recap: Bayesian Games

Mechanism Design

Auctions

Second-Price Analysis

First-Price Analysis

Revenue Equivalence

Definition 2: Epistemic Types

- Represent uncertainty over utility function using the notion of **epistemic type**.

Definition

A **Bayesian game** is a tuple (N, A, Θ, p, u) where

- N is a set of agents,
- $A = (A_1, \ldots, A_n)$, where A_i is the set of actions available to player i,
- $\Theta = (\Theta_1, \ldots, \Theta_n)$, where Θ_i is the type space of player i,
- $p : \Theta \to [0, 1]$ is the common prior over types,
- $u = (u_1, \ldots, u_n)$, where $u_i : A \times \Theta \to \mathbb{R}$ is the utility function for player i.

Strategies

- **Pure strategy**: $s_i : \Theta_i \to A_i$
  - a mapping from every type agent i could have to the action he would play if he had that type.
- **Mixed strategy**: $s_i : \Theta_i \to \Pi(A_i)$
  - a mapping from i's type to a probability distribution over his action choices.
- $s_j(a_j|\theta_j)$
  - the probability under mixed strategy s_j that agent j plays action a_j, given that j's type is θ_j.

Nash equilibrium

Definition (Bayes-Nash equilibrium)

A **Bayes-Nash equilibrium** is a mixed strategy profile s that satisfies $\forall i \;\; s_i \in BR_i(s_{-i})$.

- we can also construct an induced normal form for Bayesian games
- the numbers in the cells will correspond to ***ex-ante*** expected utilities
  - however as argued above, as long as the strategy space is unchanged, best responses don't change between the *ex-ante* and *ex-interim* cases.

Lecture Overview

Recap: Bayesian Games

Mechanism Design

Auctions

Second-Price Analysis

First-Price Analysis

Revenue Equivalence

Bayesian Game Settings

A Bayesian game setting is a Bayesian game with no actions, but instead with a set of outcomes over which agents have utilities that depend on their types.

Definition (Bayesian game setting)

A **Bayesian game setting** is a tuple (N, O, Θ, p, u), where

- N is a finite set of n agents;
- O is a set of outcomes;
- $\Theta = \Theta_1 \times \cdots \times \Theta_n$ is a set of possible joint type vectors;
- p is a (common prior) probability distribution on Θ; and
- $u = (u_1, \ldots, u_n)$, where $u_i : O \times \Theta \to \mathbb{R}$ is the utility function for each player i.

Mechanisms

Definition (Mechanism)

A **mechanism** (for a Bayesian game setting (N, O, Θ, p, u)) is a pair (A, M), where

- $A = A_1 \times \cdots \times A_n$, where A_i is the set of actions available to agent $i \in N$; and
- $M : A \to \Pi(O)$ maps each action profile to a distribution over outcomes.

Thus, the designer gets to specify

- the action sets for the agents *(though these may be constrained by the environment)*
- the mapping to outcomes, over which agents have utility
- **can't** change outcomes; agents' preferences or type spaces

Mechanism Design

- The problem is to pick a mechanism that will **cause rational agents to behave in a desired way**, specifically maximizing the mechanism designer's own "utility" or objective function
  - each agent holds private information, in the Bayesian game sense
  - often, we're interested in settings where agents' action space is identical to their type space, and an action can be interpreted as a declaration of the agent's type
- In other words:
  - perform an optimization problem, given that the values of (some of) the inputs are unknown
  - choose the Bayesian game out of a set of possible Bayesian games that maximizes some performance measure
- The **strength of the solution concept** used determines the mechanism's robustness to irrational behavior, miscoordination between agents, divergence among agents' beliefs, etc.

Lecture Overview

Recap: Bayesian Games

Mechanism Design

Auctions

Second-Price Analysis

First-Price Analysis

Revenue Equivalence

Motivation

- Auctions are mechanisms for **allocating resources among self-interested agents**
- Very widely used
  - government sale of resources; privatization
  - stock market
  - request for quote
  - real estate sales; used goods (e.g., eBay, police auctions)
  - advertisements on Google and Facebook
  - computational resources, network bandwidth, ...

Auction Types

English Auction

- auctioneer starts the bidding at some “reservation price”
- bidders then shout out **ascending prices**
- once bidders stop shouting, the high bidder gets the good at that price

First-Price Auction

- bidders **write down bids** on pieces of paper
- auctioneer awards the good to the bidder with the highest bid
- that bidder **pays the amount of her bid**

Auction Types II

Second-Price Auction

- bidders write down bids on pieces of paper
- auctioneer awards the good to the bidder with the highest bid
- that bidder pays **the amount bid by the second-highest bidder**

All-Pay Auction (sealed bid)

- bidders write down bids on pieces of paper
- auctioneer awards the good to the bidder with the highest bid
- **everyone pays the amount of their bid** regardless of whether or not they win

Representing Sealed-Bid Auctions as Bayesian Games

- **Set of agents:** bidders
- **Actions** for each agent i: bid amounts b_i
- **Types** for each agent i: valuations v_i
- **Common prior** over types:
  - **Independent Private Value** model: agents' types are drawn independently
  - Distribution: can be anything; uniform often easiest to analyze
- **Risk attitude:** how do we translate money into utility? We'll consider the **risk neutral** case, where the relationship between money and utility is linear.
- **Allocations and Payments:** determined based on the vector of bid amounts b
- **Utility function** (risk-neutral, IPV case): if agent i is allocated the good, $v_i - p_i$; else 0

Lecture Overview

Recap: Bayesian Games

Mechanism Design

Auctions

Second-Price Analysis

First-Price Analysis

Revenue Equivalence

Second-Price Analysis

Theorem

*Bidding one's value ("**truth-telling**") is a dominant strategy in a second-price auction.*

Proof.

Assume that the other bidders bid in some arbitrary way. We must show that i's best response is always to bid truthfully. We'll break the proof into two cases:

1. Bidding honestly, i would win the auction
2. Bidding honestly, i would lose the auction

□

Second-Price Analysis (2)

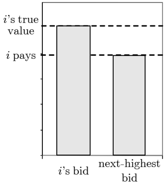

- Bidding honestly, i is the winner

Second-Price Analysis (2)

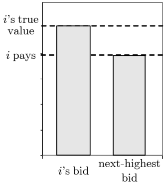

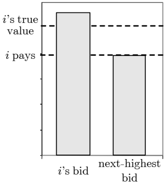

- Bidding honestly, i is the winner
- If i bids higher, he will still win and still pay the same amount

Second-Price Analysis (2)

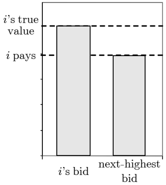

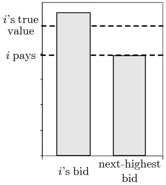

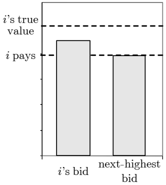

- Bidding honestly, i is the winner
- If i bids higher, he will still win and still pay the same amount
- If i bids lower, he will either still win and still pay the same amount...

Second-Price Analysis (2)

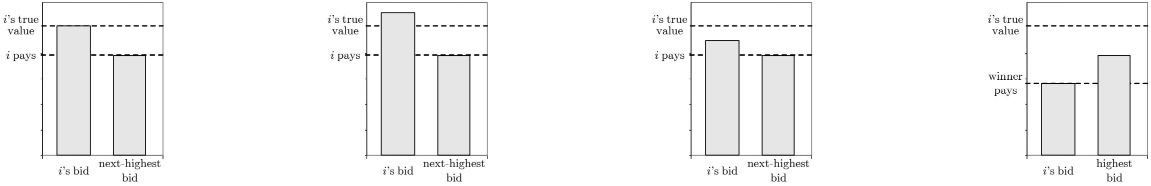

- Bidding honestly, i is the winner
- If i bids higher, he will still win and still pay the same amount
- If i bids lower, he will either still win and still pay the same amount...or lose and get utility of zero.

Second-Price Analysis (3)

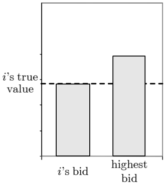

- Bidding honestly, i is not the winner

Second-Price Analysis (3)

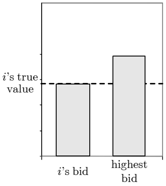

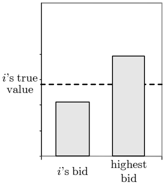

- Bidding honestly, i is not the winner
- If i bids lower, he will still lose and still pay nothing

Second-Price Analysis (3)

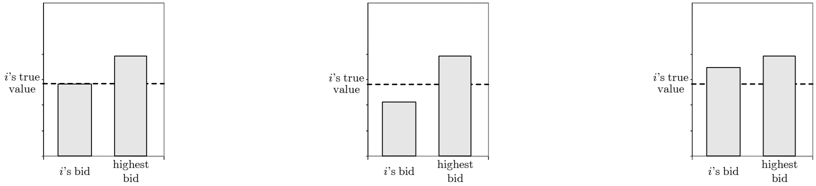

- Bidding honestly, i is not the winner
- If i bids lower, he will still lose and still pay nothing
- If i bids higher, he will either still lose and still pay nothing...

Second-Price Analysis (3)

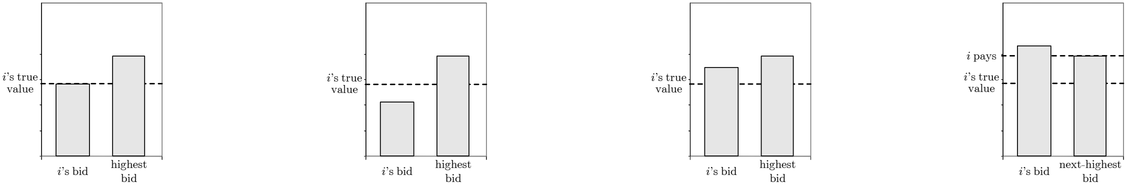

- Bidding honestly, i is not the winner
- If i bids lower, he will still lose and still pay nothing
- If i bids higher, he will either still lose and still pay nothing...or win and pay more than his valuation.

Lecture Overview

Recap: Bayesian Games

Mechanism Design

Auctions

Second-Price Analysis

First-Price Analysis

Revenue Equivalence

First-Price Analysis

Theorem

In a first-price auction with two risk-neutral bidders whose valuations are drawn independently and uniformly at random from $[0, 1]$, $(\frac{1}{2}v_1, \frac{1}{2}v_2)$ is a ***Bayes-Nash equilibrium*** *strategy profile.*

Proof.

Assume that bidder 2 bids $\frac{1}{2}v_2$, and bidder 1 bids s_1. From the fact that v_2 was drawn from a uniform distribution, all values of v_2 between 0 and 1 are equally likely. Bidder 1's ***ex-interim* expected utility** is

$$u_1\left(\left[s_1, \frac{1}{2}v_2\right] \middle| v_1\right) = \int_0^1 u_1\left(\left[s_1, \frac{1}{2}v_2\right] \middle| [v_1, v_2]\right) dv_2.$$

This integral can be broken up into two smaller integrals, splitting at the point $v_2 = 2s_1$.

$$u_1\left(\left[s_1, \frac{1}{2}v_2\right] \middle| v_1\right) = \int_0^{2s_1} u_1\left(\left[s_1, \frac{1}{2}v_2\right] \middle| [v_1, v_2]\right) dv_2 + \int_{2s_1}^1 u_1\left(\left[s_1, \frac{1}{2}v_2\right] \middle| [v_1, v_2]\right) dv_2$$

□

First-Price Analysis

Theorem

*In a first-price auction with two risk-neutral bidders whose valuations are drawn independently and uniformly at random from $[0,1]$, $(\frac{1}{2}v_1, \frac{1}{2}v_2)$ is a **Bayes-Nash equilibrium** strategy profile.*

Proof (continued).

We can now substitute in values for $u_1([s_1, \frac{1}{2}v_2]|[v_1, v_2])$. In the first case, because 2 bids $\frac{1}{2}v_2$, 1 wins when $v_2 < 2s_1$, and gains utility $v_1 - s_1$. In the second case 1 loses and gains utility 0. We can ignore the case where the agents have the same valuation, because this occurs with probability zero.

$$\begin{aligned} u_1\left(\left[s_1, \frac{1}{2}v_2\right] \middle| v_1\right) &= \int_0^{2s_1} (v_1 - s_1) dv_2 + \int_{2s_1}^{1} (0) dv_2 \\ &= (v_1 - s_1) v_2 \Big|_0^{2s_1} \\ &= [(v_1 - s_1)(2s_1)] - [(v_1 - s_1)(0)] \\ &= 2v_1 s_1 - 2(s_1)^2 \end{aligned}$$

First-Price Analysis

Theorem

*In a first-price auction with two risk-neutral bidders whose valuations are drawn independently and uniformly at random from $[0,1]$, $(\frac{1}{2}v_1, \frac{1}{2}v_2)$ is a **Bayes-Nash equilibrium** strategy profile.*

Proof (continued).

Bidder 1's **best response** to bidder 2's strategy is the s_1 at which the derivative of his *ex interim* expected utility is 0:

$$\begin{aligned}
\frac{\partial}{\partial s_1}\left(u_1\left(\left[s_1, \frac{1}{2}v_2\right] \middle| v_1\right)\right) &= 0 \\
\frac{\partial}{\partial s_1}(2v_1 s_1 - 2(s_1)^2) &= 0 \\
2v_1 - 4s_1 &= 0 \\
s_1 &= \frac{1}{2}v_1
\end{aligned}$$

Thus when player 2 is bidding half her valuation, player 1's best strategy is to bid half his valuation. The calculation of the optimal bid for player 2 is analogous, given the symmetry of the game and the equilibrium. □

First-Price: More than two bidders

- Very narrow result: two bidders, uniform valuations.

Theorem

In a first-price sealed bid auction with n risk-neutral agents whose valuations are independently drawn from a uniform distribution on the same bounded interval of the real numbers, the ***unique symmetric equilibrium*** *is given by the strategy profile* $(\frac{n-1}{n}v_1, \ldots, \frac{n-1}{n}v_n)$.

- proven using a similar argument, but more involved calculus
- the proof just verifies the equilibrium. How did we know which formula to check?

Lecture Overview

Recap: Bayesian Games

Mechanism Design

Auctions

Second-Price Analysis

First-Price Analysis

Revenue Equivalence

Revenue Equivalence

- Which auction should an auctioneer choose? To some extent, it doesn't matter...

Theorem (Revenue Equivalence Theorem)

Assume that each of n risk-neutral agents has an independent private valuation for a single good at auction, drawn from a common cumulative distribution $F(v)$ that is strictly increasing and atomless on $[\underline{v}, \bar{v}]$. Then ***any auction mechanism*** *in which*

- *the good will be allocated to the agent with the highest valuation; and*
- *any agent with valuation $\underline{v}$ has an expected utility of zero;*

yields the same expected revenue*, and hence results in any bidder with valuation v making the same expected payment.*

First and Second-Price Auctions

- The k^{th} **order statistic of a distribution**: the expected value of the k^{th}-largest of n draws.
- For n IID draws from $[0, v_{\max}]$, the k^{th} order statistic is

$$\frac{n+1-k}{n+1} v_{\max}.$$

First and Second-Price Auctions

- The k^{th} **order statistic of a distribution**: the expected value of the k^{th}-largest of n draws.
- For n IID draws from $[0, v_{\max}]$, the k^{th} order statistic is

$$\frac{n+1-k}{n+1} v_{\max}.$$

- Thus in a second-price auction, the seller's expected revenue is

$$\frac{n-1}{n+1} v_{\max}.$$

First and Second-Price Auctions

- The k^{th} **order statistic of a distribution**: the expected value of the k^{th}-largest of n draws.
- For n IID draws from $[0, v_{\max}]$, the k^{th} order statistic is

$$\frac{n+1-k}{n+1} v_{\max}.$$

- Thus in a second-price auction, the seller's expected revenue is

$$\frac{n-1}{n+1} v_{\max}.$$

- Symmetric equilibria of first and second-price auctions satisfy the requirements of the revenue equivalence theorem
  - every symmetric game has a symmetric equilibrium
  - in a symmetric equilibrium of this auction game, higher bid $\Leftrightarrow$ higher valuation

Applying Revenue Equivalence

- Thus, a bidder in a FPA must bid his expected payment **conditional on being the winner of a second-price auction**
  - this conditioning will be correct if he does win the FPA; otherwise, his bid doesn't matter anyway
  - if v_i is the high value, there are then $n-1$ other values drawn from the uniform distribution on $[0, v_i]$
  - thus, the expected value of the second-highest bid is the first-order statistic of $n-1$ draws from $[0, v_i]$:

$$\frac{n+1-k}{n+1} v_{max} = \frac{(n-1)+1-(1)}{(n-1)+1}(v_i) = \frac{n-1}{n} v_i$$

- This provides a basis for our earlier claim about n-bidder first-price auctions.
  - However, we'd still have to check that this is an equilibrium
  - The revenue equivalence theorem doesn't say that every revenue-equivalent strategy profile is an equilibrium!