Auction Theory II

Lecture 19

Lecture Overview

1. Recap
2. First-Price Auctions
3. Revenue Equivalence
4. Optimal Auctions

Motivation

- Auctions are any mechanisms for allocating resources among self-interested agents
- resource allocation is a fundamental problem in CS
- increasing importance of studying distributed systems with heterogeneous agents
- currency needn't be real money, just something scarce

Intuitive comparison of 5 auctions

	English	Dutch	Japanese	1st-Price	2nd-Price
Duration	#bidders, increment	starting price, clock speed	#bidders, increment	fixed	fixed
Info Revealed	2nd-highest val; bounds on others	winner's bid	all val's but winner's	none	none
Jump bids	yes	n/a	no	n/a	n/a
Price Discovery	yes	no	yes	no	no
Regret	no	yes	no	yes	no

Second-Price proof

Theorem

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

English and Japanese auctions

- A much more complicated strategy space
  - extensive form game
  - bidders are able to condition their bids on information revealed by others
  - in the case of English auctions, the ability to place jump bids
- intuitively, though, the revealed information doesn't make any difference in the IPV setting.

English and Japanese auctions

- A much more complicated strategy space
  - extensive form game
  - bidders are able to condition their bids on information revealed by others
  - in the case of English auctions, the ability to place jump bids
- intuitively, though, the revealed information doesn't make any difference in the IPV setting.

Theorem

Under the independent private values model (IPV), it is a dominant strategy for bidders to bid up to (and not beyond) their valuations in both Japanese and English auctions.

Lecture Overview

First-Price and Dutch

Theorem

First-Price and Dutch auctions are strategically equivalent.

- In both first-price and Dutch, a bidder must decide on the amount he's willing to pay, conditional on having placed the highest bid.
  - despite the fact that Dutch auctions are extensive-form games, the only thing a winning bidder knows about the others is that all of them have decided on lower bids
    - e.g., he does not know *what* these bids are
    - this is exactly the thing that a bidder in a first-price auction assumes when placing his bid anyway.
- Note that this is a stronger result than the connection between second-price and English.

Discussion

- So, why are both auction types held in practice?
  - First-price auctions can be held asynchronously
  - Dutch auctions are fast, and require minimal communication: only one bit needs to be transmitted from the bidders to the auctioneer.
- How should bidders bid in these auctions?

Discussion

- So, why are both auction types held in practice?
  - First-price auctions can be held asynchronously
  - Dutch auctions are fast, and require minimal communication: only one bit needs to be transmitted from the bidders to the auctioneer.
- How should bidders bid in these auctions?
  - They should clearly bid less than their valuations.
  - There's a tradeoff between:
    - probability of winning
    - amount paid upon winning
  - Bidders don't have a dominant strategy any more.

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 expected utility is

$$E[u_1] = \int_0^1 u_1 dv_2. \tag{1}$$

Note that the integral in Equation (1) can be broken up into two smaller integrals that differ on whether or not player 1 wins the auction.

$$E[u_1] = \int_0^{2s_1} u_1 dv_2 + \int_{2s_1}^1 u_1 dv_2$$

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. 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. Observe that we can ignore the case where the agents have the same valuation, because this occurs with probability zero.

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

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 find bidder 1's best response to bidder 2's strategy by taking the derivative of Equation (2) and setting it equal to zero:

$$\frac{\partial}{\partial s_1}(2v_1 s_1 - 2s_1^2) = 0$$

$$2v_1 - 4s_1 = 0$$

$$s_1 = \frac{1}{2}v_1$$

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.

More than two bidders

- Very narrow result: two bidders, uniform valuations.
- Still, first-price auctions are not incentive compatible
  - hence, unsurprisingly, not equivalent to second-price auctions

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
- a broader problem: that proof only showed how to *verify* an equilibrium strategy.
  - How do we identify one in the first place?

Lecture Overview

1. Recap
2. First-Price Auctions
3. Revenue Equivalence
4. Optimal Auctions

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.

Revenue Equivalence Proof

Proof.

Consider any mechanism (direct or indirect) for allocating the good. Let $u_i(v_i)$ be i's expected utility given true valuation v_i, assuming that all agents including i follow their equilibrium strategies. Let $P_i(v_i)$ be i's probability of being awarded the good given (a) that his true type is v_i; (b) that he follows the equilibrium strategy for an agent with type v_i; and (c) that all other agents follow their equilibrium strategies.

$$u_i(v_i) = v_i P_i(v_i) - E[\text{payment by type } v_i \text{ of player } i] \tag{1}$$

From the definition of equilibrium, for any other valuation $\hat{v}_i$ that i could have,

$$u_i(v_i) \geq u_i(\hat{v}_i) + (v_i - \hat{v}_i) P_i(\hat{v}_i). \tag{2}$$

To understand Equation (2), observe that if i followed the equilibrium strategy for a player with valuation $\hat{v}_i$ rather than for a player with his (true) valuation v_i, i would make all the same payments and would win the good with the same probability as an agent with valuation $\hat{v}_i$. However, whenever he wins the good, i values it $(v_i - \hat{v}_i)$ more than an agent of type $\hat{v}_i$ does. The inequality must hold because in equilibrium this deviation must be unprofitable.

Revenue Equivalence Proof

Proof (continued).

Consider $\hat{v}_i = v_i + dv_i$, by substituting this expression into Equation (2):

$$u_i(v_i) \geq u_i(v_i + dv_i) + dv_i P_i(v_i + dv_i). \tag{3}$$

Likewise, considering the possibility that i's true type could be $v_i + dv_i$,

$$u_i(v_i + dv_i) \geq u_i(v_i) + dv_i P_i(v_i). \tag{4}$$

Combining Equations (4) and (5), we have

$$P_i(v_i + dv_i) \geq \frac{u_i(v_i + dv_i) - u_i(v_i)}{dv_i} \geq P_i(v_i). \tag{5}$$

Taking the limit as $dv_i \to 0$ gives $\frac{du_i}{dv_i} = P_i(v_i)$. Integrating up,

$$u_i(v_i) = u_i(\underline{v}) + \int_{x=\underline{v}}^{v_i} P_i(x)dx. \tag{6}$$

Revenue Equivalence Proof

Proof (continued).

Now consider any two efficient auction mechanisms in which the expected payment of an agent with valuation $\underline{v}$ is zero. A bidder with valuation $\underline{v}$ will never win (since the distribution is atomless), so his expected utility $u_i(\underline{v}) = 0$. Because both mechanisms are efficient, every agent i always has the same $P_i(v_i)$ (his probability of winning given his type v_i) under the two mechanisms. Since the right-hand side of Equation (6) involves only $P_i(v_i)$ and $u_i(\underline{v})$, each agent i must therefore have the same expected utility u_i in both mechanisms. From Equation (1), this means that a player of any given type v_i must make the same expected payment in both mechanisms. Thus, i's *ex ante* expected payment is also the same in both mechanisms. Since this is true for all i, the auctioneer's expected revenue is also the same in both mechanisms.

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

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

- 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!

Lecture Overview

Optimal Auctions

- So far we have only considered efficient auctions.
- What about maximizing the seller's revenue?
  - she may be willing to risk failing to sell the good even when there is an interested buyer
  - she may be willing sometimes to sell to a buyer who didn't make the highest bid
- Mechanisms which are designed to maximize the seller's expected revenue are known as optimal auctions.

Optimal auctions setting

- independent private valuations
- risk-neutral bidders
- each bidder i's valuation drawn from some strictly increasing cumulative density function $F_i(v)$ (PDF $f_i(v)$)
  - we allow $F_i \neq F_j$: asymmetric auctions
- the seller knows each F_i

Designing optimal auctions

Definition (virtual valuation)

Bidder i's virtual valuation is $\psi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$.

Definition (bidder-specific reserve price)

Bidder i's bidder-specific reserve price r_i^* is the value for which $\psi_i(r_i^*) = 0$.

Designing optimal auctions

Definition (virtual valuation)

Bidder i's virtual valuation is $\psi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$.

Definition (bidder-specific reserve price)

Bidder i's bidder-specific reserve price r_i^* is the value for which $\psi_i(r_i^*) = 0$.

Theorem

The optimal (single-good) auction is a sealed-bid auction in which every agent is asked to declare his valuation. The good is sold to the agent $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$. If the good is sold, the winning agent i is charged the smallest valuation that he could have declared while still remaining the winner: $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ *and* $\forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ and $\forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- Is this VCG?

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0 \text{ and } \forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- Is this VCG?
  - No, it's not efficient.

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ and $\forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- Is this VCG?
  - No, it's not efficient.
- How should bidders bid?

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ and $\forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- Is this VCG?
  - No, it's not efficient.
- How should bidders bid?
  - it's a second-price auction with a reserve price, held in virtual valuation space.
  - neither the reserve prices nor the virtual valuation transformation depends on the agent's declaration
  - thus the proof that a second-price auction is dominant-strategy truthful applies here as well.

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner,
  $\inf\{v_i^* : \psi_i(v_i^*) \geq 0 \text{ and } \forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- What happens in the special case where all agents' valuations are drawn from the same distribution?

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ and $\forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- What happens in the special case where all agents' valuations are drawn from the same distribution?
  - a second-price auction with reserve price r^* satisfying $r^* - \frac{1-F_i(r^*)}{f_i(r^*)} = 0$.

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner,
  $\inf\{v_i^* : \psi_i(v_i^*) \geq 0 \text{ and } \forall j \neq i, \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- What happens in the special case where all agents' valuations are drawn from the same distribution?
  - a second-price auction with reserve price r^* satisfying $r^* - \frac{1-F_i(r^*)}{f_i(r^*)} = 0$.
- What happens in the general case?

Analyzing optimal auctions

Optimal Auction:

- winning agent: $i = \arg\max_i \psi_i(\hat{v}_i)$, as long as $v_i > r_i^*$.
- i is charged the smallest valuation that he could have declared while still remaining the winner, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0 \text{ and } \forall j \neq i,\ \psi_i(v_i^*) \geq \psi_j(\hat{v}_j)\}$.

- What happens in the special case where all agents' valuations are drawn from the same distribution?
  - a second-price auction with reserve price r^* satisfying $r^* - \frac{1-F_i(r^*)}{f_i(r^*)} = 0$.
- What happens in the general case?
  - the virtual valuations also increase weak bidders' bids, making them more competitive.
  - low bidders can win, paying less
  - however, bidders with higher expected valuations must bid more aggressively