Auction Theory II

CPSC 532A Lecture 22

November 28, 2006

Lecture Overview

Intuitive comparison of 5 auctions

	English	Dutch	Japanese	1^{st}-Price	2^{nd}-Price
Duration	#bidders, increment	starting price, clock speed	#bidders, increment	fixed	fixed
Info Revealed	2^{nd}-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

Theorem

Truth-telling is a dominant strategy in a second-price auction.

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.

First-Price and Dutch

Theorem

First-Price and Dutch auctions are strategically equivalent.

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 $\left(\frac{n-1}{n}v_1, \ldots, \frac{n-1}{n}v_n\right)$.

Lecture Overview

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 $u_i(\hat{v})$ be i's expected utility and let $p_i(\hat{v})$ be i's probability of being awarded the good, in equilibrium of the mechanism if he follows the equilibrium strategy for an agent with type $\hat{v}$ and this were in fact his type.

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

From the definition of equilibrium,

$$u_i(v_i) \geq u_i(\hat{v}) + (v_i - \hat{v})p_i(\hat{v}) \quad (2)$$

By behaving according to the equilibrium strategy for a player of type $\hat{v}$, i makes all the same payments and wins the good with the same probability as an agent of type $\hat{v}$. Because an agent of type v_i values the good $(v_i - \hat{v})$ more than an agent of type $\hat{v}$ does, we must add this term. The inequality holds because this deviation must be unprofitable. Consider $\hat{v} = v_i + dv_i$, by substituting this expression into Equation (2):

$$u_i(v_i) \geq u_i(v_i + dv_i) + dv_i p_i(v_i + d_v i) \quad (3)$$

Revenue Equivalence Proof

Proof.

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 (3) and (4), 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 $dv_i \rightarrow 0$ gives

$$\frac{du_i}{dv_i} = p_i(v_i) \tag{6}$$

Integrating up,

$$u_i(v_i) = u_i(\underline{v}) + \int_{x=\underline{v}}^{v_i} p_i(x)dx \tag{7}$$

Revenue Equivalence Proof

Proof.

Now consider any two mechanisms which satisfy the conditions given in the statement of the theorem. A bidder with valuation $\underline{v}$ will never win (since the distribution is atomless), so his expected utility $u_i(\underline{v}) = 0$. Every agent i has the same $p_i(v_i)$ (his probability of winning given his type v_i) under the two mechanisms, regardless of his type. These mechanisms must then also have the same u_i functions, by Equation (7). 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
  - if he's not the high bidder, he gets no utility anyway, so his strategy should be based on the assumption that he is the high bidder
  - 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

1. Recap
2. Revenue Equivalence
3. Optimal Auctions
4. Beyond IPV and risk-neutrality

Fun game

- Pass around the jar of coins and try to determine how much money is inside.
- Once everyone has seen it, we'll play a game...

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

Bidder i's virtual valuation is

$$\psi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

Definition

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

Designing optimal auctions

Definition

Bidder i's virtual valuation is

$$\psi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}.$$

Definition

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 $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$.

Analyzing optimal auctions

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* $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$.

- Is this VCG?

Analyzing optimal auctions

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* $\max(r_i^*, max_{j \neq i}\psi_j(v_j))$*.*

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

Analyzing optimal auctions

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* $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$.

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

Analyzing optimal auctions

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* $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$.

- 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

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* $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$*.*

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

Analyzing optimal auctions

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* $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$*.*

- 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

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* $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$.

- 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

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* $\max(r_i^*, max_{j \neq i} \psi_j(v_j))$.

- 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

Lecture Overview

Fun game

- Look at the jar of coins
- Bid for it using real money in a sealed-bid second-price auction.

Going beyond IPV

- common value model
  - motivation: oil well
  - winner's curse
  - things can be improved by revealing more information
- general model
  - IPV + common value
  - example motivation: private value plus resale