Advanced Single-Good; Multiunit Auctions

Lecture 20

Lecture Overview

1. Recap
2. Optimal Auctions
3. Beyond IPV and risk-neutrality
4. Multiunit auctions

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.

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.

Applying Revenue Equivalence

- A bidder in a FPA must bid his expected payment conditional on being the winner of a second-price auction
  - 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. Optimal Auctions
3. Beyond IPV and risk-neutrality
4. Multiunit auctions

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 (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, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ 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,
  $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ 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, $\inf\{v_i^* : \psi_i(v_i^*) \geq 0$ 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

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

Affiliated Values

- Definition: a high value of one bidder's signal makes high values of other bidders' signals more likely
  - common value model is a special case
- generally, ascending auctions lead to higher expected prices than second price, which in turn leads to higher expected prices than first price
  - intuition: winner's gain depends on the privacy of his information.
  - The more the price paid depends on others' information (rather than expectations of others' information), the more closely this price is related to the winner's information, since valuations are affiliated
  - thus the winner loses the privacy of his information, and can extract a smaller "information rent"

Affiliated Values

- Definition: a high value of one bidder's signal makes high values of other bidders' signals more likely
  - common value model is a special case
- generally, ascending auctions lead to higher expected prices than second price, which in turn leads to higher expected prices than first price
  - intuition: winner's gain depends on the privacy of his information.
  - The more the price paid depends on others' information (rather than expectations of others' information), the more closely this price is related to the winner's information, since valuations are affiliated
  - thus the winner loses the privacy of his information, and can extract a smaller "information rent"
- Linkage principle: if the seller has access to any private source of information which will be affiliated with the bidders' valuations, she should precommit to reveal it honestly.

Risk Attitudes

What kind of auction would the auctioneer prefer?

- Buyer is not risk neutral:
  - no change under various risk attitudes for second price
  - in first-price, increasing bid amount increases probability of winning, decreases profit. This is good for risk-averse bidder, bad for risk-seeking bidder.
  - Risk averse, IPV: First $\succ$ [Japanese $=$ English $=$ Second]
  - Risk seeking, IPV: Second $\succ$ First

Risk Attitudes

What kind of auction would the auctioneer prefer?

- Buyer is not risk neutral:
  - no change under various risk attitudes for second price
  - in first-price, increasing bid amount increases probability of winning, decreases profit. This is good for risk-averse bidder, bad for risk-seeking bidder.
  - Risk averse, IPV: First $\succ$ [Japanese $=$ English $=$ Second]
  - Risk seeking, IPV: Second $\succ$ First
- Auctioneer is not risk neutral:
  - revenue is fixed in first-price auction (the expected amount of the second-highest bid)
  - revenue varies in second-price auction, with the same expected value
  - thus, a risk-averse seller prefers first-price to second-price.

Lecture Overview

Multiunit Auctions

- now let's consider a setting in which
  - there are k identical goods for sale in a single auction
  - every bidder only wants one unit
- what is VCG in this setting?

Multiunit Auctions

- now let's consider a setting in which
  - there are k identical goods for sale in a single auction
  - every bidder only wants one unit
- what is VCG in this setting?
  - every unit is sold for the amount of the $k+1$st highest bid

Multiunit Auctions

- now let's consider a setting in which
  - there are k identical goods for sale in a single auction
  - every bidder only wants one unit
- what is VCG in this setting?
  - every unit is sold for the amount of the $k + 1$st highest bid
- how else can we sell the goods?

Multiunit Auctions

- now let's consider a setting in which
  - there are k identical goods for sale in a single auction
  - every bidder only wants one unit
- what is VCG in this setting?
  - every unit is sold for the amount of the $k+1$st highest bid
- how else can we sell the goods?
  - pay-your-bid: "discriminatory" pricing, because bidders will pay different amounts for the same thing
  - lowest winning bid: very similar to VCG, but ensures that bidders don't pay zero if there are fewer bids than units for sale
  - sequential single-good auctions

Revenue Equivalence

Theorem (Revenue equivalence theorem, multiunit version)

Assume that each of n risk-neutral agents has an independent private valuation for a single unit of k identical goods at auction, drawn from a common cumulative distribution $F(v)$ that is strictly increasing and atomless on $[\underline{v}, \bar{v}]$. Then any efficient auction mechanism in which 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_i making the same expected payment.

Sequential Auctions

Although we can apply the revelation principle, for greater intuition we can also use backward induction to derive the equilibrium strategies in finitely-repeated second-price auctions.

- everyone should bid honestly in the final auction
- we can also compute a bidder's expected utility (conditioned on type) in that auction
- in the second-last auction, bid the difference between valuation and the expected utility for losing
  - i.e., bid valuation minus the expected utility for playing the second auction
- combining these last two auctions together, there's some expected utility to playing both of them
- now this is the "expected utility of losing"
- apply backward induction