31E11100 - Microeconomics: Pricing

Pauli Murto

Aalto University

Lecture 11: Formal analysis of auctions
October 15, 2018

Objectives for today

- In the previous lecture, we worked through simple examples.
  - Two bidders, independent private values drawn from uniform distribution.
  - Ascending price auction, second-price auction, first-price auction.
- It turned out that all these formats resulted in the same expected revenue for the seller.
- We also saw that a reserve price can increase seller's revenue.
- Today, the goal is to understand these findings better.
- In particular, we look for an explanation of the revenue equivalence theorem.
- To do that, we start by defining games of incomplete information.

Bayesian Games: Formal Definitions

- Harsanyi: a game of incomplete information is given by
  1. set of players: $i \in \{1, 2, ..., N\}$
  2. actions available to player i: A_i for $i \in \{1, 2, ..., N\}$. Let $a_i \in A_i$ denote a typical action for player i
  3. sets of possible types for all players: Θ_i for $i \in \{1, 2, ..., N\}$. Let $\theta_i \in \Theta_i$ denote a typical type of player i
  4. let $a = (a_1, ..., a_N)$, $\theta = (\theta_1, ..., \theta_N)$, $a_{-i} = (a_1, ..., a_{i-1}, a_{i+1}, ..., a_N)$, $\theta_{-i} = (\theta_1, ..., \theta_{i-1}, \theta_{i+1}, ..., \theta_N)$ etc.
  5. natures move: θ is selected according to a joint probability distribution $p(\theta)$ on $\Theta = \Theta_1 \times \cdots \times \Theta_N$
  6. strategies: $s_i : \Theta_i \to A_i$, for $i \in \{1, 2, ..., N\}$. $s_i(\theta_i) \in A_i$ is then the action that type θ_i of player i takes
  7. payoffs: $u_i(a_1, ..., a_N; \theta_1, ..., \theta_N)$

Bayesian Games: Formal Definitions

- Game proceeds as follows
  - Nature chooses θ according to $p(\theta)$.
  - Each player i observes realized type $\widehat{\theta}_i$ and updates her beliefs.
    - ★ Each player comes up with conditional probability on remaining types conditional on $\theta_i = \widehat{\theta}_i$.
    - ★ Denote distribution on θ_{-i} conditional on $\widehat{\theta}_i$ by $p_i(\theta_{-i}|\widehat{\theta}_i)$.
    - ★ Recall Bayes' rule:

$$p_i(\widehat{\theta}_{-i}|\widehat{\theta}_i) = \frac{p_i\left(\widehat{\theta}_i, \widehat{\theta}_{-i}\right)}{\sum_{\theta_{-i}\in\Theta_{-i}} p_i\left(\widehat{\theta}_i, \theta_{-i}\right)}.$$

  - Players take actions simultaneously.

Bayesian Games: Formal Definitions

- Important special cases:
- Private values: for all a, i, θ_i and all $\theta_{-i}, \theta'_{-i}$ we have:

$$u_i\left(a; \theta_i, \theta_{-i}\right) = u_i\left(a; \theta_i, \theta'_{-i}\right).$$

  - In words, player i's payoff in the game depends on her own information and the actions chosen by all players, but not on the information of the others.
  - In all other cases, we say that we have interdependent values.
  - Come up with examples where private values make sense and where interdependent values make sense.
- Independent values: for all i, θ_i and θ'_i we have:

$$p_i\left(\theta_{-i}|\theta_i\right) = p_i\left(\theta_{-i}|\theta'_i\right).$$

  - In words, your own type contains no information on the types of the others.
  - Hence $p\left(\theta\right) = p_1\left(\theta_1\right) \cdot p_2\left(\theta_2\right) \cdot ... \cdot p_N\left(\theta_N\right)$, where $p_i\left(\theta_i\right)$ is the marginal distribution on θ_i.

Bayesian Games: Formal Definitions

- Solution Concept: Bayesian Nash Equilibrium:

Definition: A strategy profile $(s_1(\theta_1), ..., s_N(\theta_N))$ is a (pure strategy) Bayesian Nash Equilibrium if $s_i(\theta_i)$ is a best response to $s_{-i}(\theta_{-i})$ for all i and all $\theta_i \in \Theta_i$.

- Action specified by strategy of any given player has to be optimal given strategies of all other players and beliefs of player.
- To compute the expected payoff, note:
  - Given strategy $s_i(\cdot)$, type θ_i of player i plays action $s_i(\theta_i)$
  - With vector of types $\theta = (\theta_1, ..., \theta_N)$ and strategies $(s_1, ..., s_N)$, realized action profile is $(s_1(\theta_1), ..., s_N(\theta_N))$
  - Player i of type $\widehat{\theta}_i$ has beliefs about types of other players given by conditional probability distribution $p_i(\theta_{-i}|\widehat{\theta}_i)$

Bayesian Games: Formal Definitions

- The expected payoff from action s_i is

$$\sum_{\boldsymbol{\theta}_{-i}} u_i(s_i, s_{-i}(\theta_{-i}), \theta) p_i(\theta_{-i} \mid \widehat{\theta}_i)$$

- Best Response: action $s_i(\widehat{\theta}_i)$ is a best response to $s_{-i}(\theta_{-i})$ if and only if for all $a_i' \in A_i$

$$\sum_{\boldsymbol{\theta}:_{-i}} u_i(s_i(\widehat{\theta}_i), s_{-i}(\theta_{-i}), \theta) p_i(\theta_{-i} \mid \widehat{\theta}_i)$$

$$\geq \sum_{\boldsymbol{\theta}:_{-i}} u_i(a_i', s_{-i}(\theta_{-i}), \theta) p_i(\theta_{-i} \mid \widehat{\theta}_i)$$

Bayesian Games: Auctions

- An auction is a particular Bayesian game.
- A seller with an indivisible item for sale, zero cost.
- N bidders: $i = 1, ..., N$.
- Each bidder i has private information $\theta_i \in \Theta_i$.
- Given the profile $\theta = (\theta_i, \theta_{-i})$, bidder i's valuation is $u_i(\theta_i, \theta_{-i})$ if he gets the item and zero otherwise.
- The prior distribution over $\Theta \equiv \times_{i=1}^{N} \Theta_i$ is $F(\theta)$. After knowing one's own θ_i, bidder i forms the posterior distribution of others' valuation payoff as $F_i(\theta_{-i}|\theta_i)$.
- All bidders and seller are risk-neutral expected utility maximizers.

Bayesian Games: Auctions

- B_i : (pure) action space for bidder i ($b_i \in B_i$ the amount i can bid in auction, most typically $B_i = \mathbb{R}_+$).
- Pure strategies: $s_i : \Theta_i \to B_i$.
- Let $P_i(b_1, \cdots, b_N)$ be the probability that bidder i wins.
- Let $T_i(b_1, \cdots, b_N)$ be the monetary payment that bidder i transfers to seller (no matter i wins or not) if $(b_1, \cdots, b_N)$ is the vector of bids.
  - $T_i(b_i, b_{-i})$ can even be negative.
- Payoffs to i if θ is the realized type vector and b is the realized bid vector:

$$P_i(b_1, ..., b_N)\, u_i(\theta_i, \theta_{-i}) - T_i(b_i, b_{-i}).$$

Bayesian Games: Auctions

- **Private values**: if for all θ_i, θ'_{-i}, θ_{-i}, $u_i(\theta_i, \theta_{-i}) = u_i(\theta_i, \theta'_{-i})$.
- **Interdependent values**: if the above condition is violated.
- **Common values**: For all i, j and $\theta \in \Theta \equiv \times_{i=1}^{N} \Theta_i$,

$$u_i(\theta) = u_j(\theta).$$

- **Independent value** model: if θ_i, $i = 1, ..., N$, are independentlty drawn.
- **Symmetric** case: if $f_i(\theta) = f_j(\theta)$ and $u_i = u_j$ for any i and j.

Bayesian Games: Auctions

- We work today with the independent, symmetric and private value model in which all θ_is are i.i.d. drawn from a common distribution.
- We also assume that all bidders and seller are risk neutral.
- Hence, given the bid profile (b_i, b_{-i}), bidder i's payoff is

$$\theta_i P_i(b_i, b_{-i}) - T_i(b_i, b_{-i}).$$

Standard Auction Formats

- First Price Auction (High-bid Auction)
  - buyers simultaneously submit bids
  - the highest bidder wins (tie broken by flip coin)
  - winner pays bid (losers pay nothing)
- $$P_i(b_i, b_{-i}) = \begin{cases} 1 & \text{if } b_i > b_j, \forall j \neq i \\ \frac{1}{K} & \text{if } b_i \text{ ties for highest with } K-1 \text{ others} \\ 0 & \text{otherwise} \end{cases}.$$
- $$T_i(b_i, b_{-i}) = \begin{cases} b_i & \text{if } i \text{ wins} \\ 0 & \text{otherwise} \end{cases}.$$

Standard Auction Formats

- Dutch Auction (Open Descending Auction)
  - Auctioneer starts with a high price and continuously lowers it until some buyer agrees to buy at current price
  - the highest bidder wins (tie broken by flip coin)
-
$$P_i(b_i, b_{-i}) = \begin{cases} 1 & \text{if } b_i > b_j, \forall j \neq i \\ \frac{1}{K} & \text{if } b_i \text{ ties for highest with } K-1 \text{ others} \\ 0 & \text{otherwise} \end{cases}.$$
-
$$T_i(b_i, b_{-i}) = \begin{cases} b_i & \text{if } i \text{ wins} \\ 0 & \text{otherwise} \end{cases}.$$
- This is the same as the case in FPA.
- Dutch Auction and First Price Auction are *strategically* equivalent.

Standard Auction Formats

- Second Price Auction (Vickrey Auction)
  - same rules as FPA except that winner pays *second* highest bid
  - proposed in 1961 by William Vickrey
- $P_i(b_i, b_{-i}) = \begin{cases} 1 & \text{if } b_i > b_j, \forall j \neq i \\ \frac{1}{K} & \text{if } b_i \text{ ties for highest with } K-1 \text{ others} \\ 0 & \text{otherwise} \end{cases}.$
- $T_i(b_i, b_{-i}) = \begin{cases} \max_{j \neq i} b_i & \text{if } i \text{ wins} \\ 0 & \text{otherwise} \end{cases}.$

Second-price auction (SPA)

- Claim: It is optimal for each player i to bid according to $b_i(\theta_i) = \theta_i$.
- Proof: Let $V_i(\theta_i, b_i, b_{-i})$ be the payoff to i of type θ_i when the others bid vector is b_{-i}.Then

$$V_i(\theta_i, b_i, b_{-i}) = \begin{cases} \theta_i - \max_{j \neq i} b_i \text{ if } b_i \geq \max_{j \neq i} b_i, \\ 0 \text{ otherwise.} \end{cases}$$

  Hence it is optimal to set $b_i \geq \max_{j \neq i} b_i$ if and only if
  $\theta_i - \max_{j \neq i} b_i \geq 0$.
  Clearly setting $b_i(\theta_i) = \theta_i$ accomplishes exactly this.
- We say that $b_i(\theta_i) = \theta_i$ is a dominant strategy since the optimal bid amount does not depend on the strategies of the other players.

Standard Auction Formats

- English Auction (Ascending Price Auction)
  - buyers announce bids, each successive bid higher than previous one
  - the last one to bid the item wins at what he bids
- As long as the current price p is lower than θ_i, bidder i has a chance to get positive surplus. He will not drop out until p hits θ_i .
- Only when anyone else drops out before bidder i, i.e., $p = \max_{j \neq i} \theta_j$ can he win by paying p, the second highest valuation.
- This shows that English Auction and Second Price Auction are equivalent.

First-price auction (FPA)

- Derivining the equilibrium bid function for the first-price auction is more tricky, since there is no dominant strategy
- The equilibrium is derived in a direct way in the appendix of this slide set (at the end)
- Instead, we next derive the Revenue Equivalence Theorem and use that to derive the equilibrium of the first-price auction

Envelope Formula and Revenue Equivalence Theorem

- How to explain the revenue equivalence between first and second price auctions that we observed in the example last week?
- Consider an IPV auction with symmetric type distributions (do not yet specify auction format)
- Suppose that i with type θ_i bids b_i.
- Her probability of winning P_i and her expected payment T_i are then determined by b_i, and not by θ_i.
- We write the expected payoff then as:

$$V_i(\theta_i, b_i) = \theta_i P_i(b_i) - T_i(b_i).$$

- The expected maximized payoff to i of type θ_i is then:

$$U_i(\theta_i) = \max_{b_i} \theta_i P_i(b_i) - T_i(b_i).$$

- The envelope theorem tells us that $U'(\theta_i) = P_i(b_i)$ (Check that you know what envelope theorem says)

Envelope Formula and Revenue Equivalence Theorem

- If we look for equilibria in symmetric increasing strategies, we must have:

$$P_i\left(b_i\left(\theta_i\right)\right) = F\left(\theta_i\right)^{N-1}.$$

- Using envelope theorem, we have:

$$U_i\left(\theta_i\right) = \int_0^{\theta_i} F\left(s\right)^{N-1} ds.$$

- This is really remarkable since we have not said anything about the auction format at this stage.
- The expected payoff to each bidder is the same in all auctions that result in the same probability of winning.
- Hence expected payoff is the same in FPA and SPA.
- But this means that the expected payments that the bidders make must be equal in SPA and FPA.
- But then the expected revenue to the seller must be the same: Revenue Equivalence Theorem

Auctions and Envelope Theorem

- Now we can also use this result to derive equilibria in different auctions!
- For FPA,

$$U_i(\theta_i) = (\theta_i - b(\theta_i)) F(\theta_i)^{N-1}.$$

- But the envelope formula says:

$$U_i(\theta_i) = \int_0^{\theta_i} F(s)^{N-1} ds.$$

Combining these, we get:

$$b(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F(s)^{N-1} ds}{F(\theta_i)^{N-1}}.$$

- See the appendix of this slide set to derive this same formula directly.

Auctions and Envelope Theorem

- We can also compute equilibria for other auctions using this.
- In an all pay auction, all bidders pay their bid and the highest bidder wins the object.
- In a symmetric equilibrium then,

$$U_i(\theta_i) = \theta_i F(\theta_i)^{N-1} - b(\theta_i).$$

- Using the envelope formula, we get:

$$b(\theta_i) = \theta_i F(\theta_i)^{N-1} - \int_0^{\theta_i} F(s)^{N-1} ds.$$

- So in the case with $F(\theta_i) = \theta_i$, we get

$$b(\theta_i) = \frac{N-1}{N}\theta^N.$$

Discussion

- The Revenue Equivalence Theorem shows that whenever two auction formats lead to the same allocation, the expected revenue of the seller is the same
- In particular, this holds for standard first-price and second price auctions, where allocation is efficient (highest valuation bidders gets the object)
- Recall the example in the last lecture with a reserve price:
  - A positive reserve price leads to inefficient allocation
  - But improves expected revenue of the seller
  - Revenue Equivalence also implies that two different auctions with the same distortion lead to the same revenue
- How to design auctions optimally from the seller's perspective?
  - In a significant paper, Myerson (1981): "Optimal Auction Design" (Mathematics of Operations Research) gives the full answer
  - In our environment, an optimally chosen reserve price is indeed the best the seller can do

Further readings

- For a very elegant presentation of the theory of auctions (at advanced MSc/PhD level), see the book Krishna: Auction Theory (Academic Press)
- Another excellent, but a bit advanced book, is Milgrom: Putting Auction Theory to Work (Cambridge University Press)

APPENDIX: direct derivation of equilibrium bids for the first-price auction

- Let all bidders' valuations are independent and have the same cumulative distribution $F(\theta_i)$ an $[0,1]$.
- Let $f(\theta_i)$ be the associated density function.
- Consider symmetric equilibria where all bidders use the same bidding strategy $b(\theta_i)$.
- Assume furthermore that $b(\theta_i)$ is a strictly increasing function so that

$$\theta_i < \theta_i' \Rightarrow b(\theta_i) < b(\theta_i').$$

- Since $F(\cdot)$ has a density ties happen with probability zero and they can be ignored in the analysis.

First-price auction (FPA)

- To find equilibrium, consider optimal bid of bidder i if others use $b(\theta_j)$
- Bidder i wins with bid β_i if and only if

$$b_j = b(\theta_j) < \beta_i \text{ for all } j \neq i.$$

- Hence i wins with bid β_i if and only if

$$\theta_j < b^{-1}(\beta_i), \text{ for all } j \neq i,$$

where $b^{-1}(\cdot)$ is the inverse function of the bid function.

- We can then calculate the expected payoff to bidder i with valuation θ_i from bid β_i :

$$(\theta_i - \beta_i)\left(F\left(b^{-1}(\beta_i)\right)\right)^{N-1}.$$

- Optimal bid for θ_i is then found by

$$\max_{\beta_i}(\theta_i - \beta_i)\left(F\left(b^{-1}(\beta_i)\right)\right)^{N-1}.$$

First-price auction (FPA)

- First-order condition for optimal β_i :

$$(\theta_i - \beta_i)(N-1)\left(F\left(b^{-1}(\beta_i)\right)\right)^{N-2}\frac{dF\left(b^{-1}(\beta_i)\right)}{d\beta_i} = \left(F\left(b^{-1}(\beta_i)\right)\right)^{N-1}$$

- By chain rule,

$$\frac{dF\left(b^{-1}(\beta_i)\right)}{d\beta_i} = f\left(b^{-1}(\beta_i)\right) d\frac{b^{-1}(\beta_i)}{d\beta_i},$$

and by inverse function rule,

$$\frac{dF\left(b^{-1}(\beta_i)\right)}{d\beta_i} = \frac{f\left(b^{-1}(\beta_i)\right)}{b'\left(b^{-1}(\beta_i)\right)}$$

First-price auction (FPA)

- Since in equilibrium, $\beta_i = b(\theta_i)$ must be optimal, we have:

$$(\theta_i - b(\theta_i))(N-1)(F(\theta_i))^{N-2}\frac{f(\theta_i)}{b'(\theta_i)} - F(\theta_i)^{N-1} = 0.$$

- Multiplying both sides by $b'(\theta_i)$, we get

$$(\theta_i - b(\theta_i))(N-1)(F(\theta_i))^{N-2} f(\theta_i) - b'(\theta_i) F(\theta_i)^{N-1} = 0,$$

or

$$\frac{d}{d\theta_i}(\theta_i - b(\theta_i)) F(\theta_i)^{N-1} - F(\theta_i)^{N-1} = 0,$$

or by integrating:

$$(\theta_i - b(\theta_i)) F(\theta_i)^{N-1} = \int_0^{\theta_i} F(\theta)^{N-1} d\theta.$$

- Hence the symmetric equilibrium bid function is:

$$b(\theta_i) = \theta_i - \frac{\int_0^{\theta_i} F(\theta)^{N-1} d\theta}{F(\theta_i)^{N-1}}.$$

First-price auction (FPA)

- Properties of the bid function:
  - $b(\theta_i) < \theta_i$ for all $\theta_i > 0$
  - $b(\theta_i) > 0$ for all $\theta_i > 0$
  - Increasing in θ_i (i.e. $b'(\theta_i) > 0$, can you see this?)
  - How does $b(\theta_i)$ depend on N?
    - ★ Look at special case $F(\theta_i) = \theta_i$.
    - ★ Then $b(\theta_i) = \theta_i - \frac{1}{N}\theta_i$.
    - ★ Hence the equilibrium bid is increasing in the number of competing bidders.

Comparing FPA and SPA

- We know by revenue equivalence theorem that FPA and SPA lead to the same allocation and the same expected revenue to the seller
- This can of course be checked also directly
- For simplicity, assume uniform distribution here: $F(\theta_i) = \theta_i$.
- The revenue in SPA is simply the second highest θ_i.
- In FPA, revenue is $\left(\frac{N-1}{N}\right)$ times highest θ_i.
- Which one is greater?

Comparing FPA and SPA

- Let $\theta^{(2)}$ be the second highest valuation.
  - It has density function $N(N-1)\,\theta^{N-2}(1-\theta)$ for $\theta \in [0,1]$.
  - Hence it has expected value

$$\mathbb{E}\left(\theta^{(2)}\right) = \int_0^1 N(N-1)\left(\theta^{N-1} - \theta^N\right) d\theta = \frac{N-1}{N+1}$$

- The highest valuation $\theta^{(1)}$ has density $N\theta^{N-1}$ for $\theta \in [0,1]$.
  - Hence

$$\mathbb{E}\left(\theta^{(1)}\right) = \int_0^1 N\theta^N d\theta = \frac{N}{N+1}.$$

  - Expected revenue is then

$$\mathbb{E}\left(b(\theta^{(1)})\right) = \frac{N-1}{N+1}.$$

- We observe that the expected revenue is the same in the two auctions.