31E11100 - Microeconomics: Pricing

Pauli Murto

Aalto University

Part 4: Auctions
Lectures on 2.10., 4.10., 9.10. and 11.10.2023

Plan for Part 4: Auctions

- Lecture 2.10., 4.10.: Introduction to auctions
  - Why auctions?
  - Different auction formats
  - Auction design in real world (reading assignment)
  - Auction design in theory
- Lecture 9.10.: Formal analysis of auctions
  - Auctions as Bayesian games
  - Envelope formula
  - Revenue equivalence theorem
- Lecture 11.10.: Common value auctions
  - Winner's curse
  - How prices aggregate dispersed information

Why auctions?

- Suppose a seller has a single item to sell and a number of potential buyers. How to sell?
  - So far in this course: seller sets a price (or menu)
  - Buyer: take it or leave it
- Why use an auction?
  - What is the right price? If too high, no one buys. If too low, excess demand.
  - Auction is a mechanism for *price discovery*
  - Buyers know what they would pay, but why should they tell?
    - Auction *induces competition* between buyers
  - Auctions can also *aggregate dispersed information* in prices (e.g. markets for financial assets)
- Important applications
  - Telecommunication licences, electricity markets, public procurement, online ad auctions, etc.
  - How to design an auction?

Most common auction formats (1)

- Sealed bid auctions
  - Seller asks for a single bid from each participant
  - Highest bid wins and pays her bid
  - Common in selling real estate and different commodities
  - Also very common in procuring services
    - Governments and public sector procures services through competitive tendering
    - Suppliers make bids for service contracts and lowest bid wins
    - This is a "reverse" auction, since buyer seeks the lowest price from competing suppliers
- An important variant: second price auction.
  - Highest bidder wins but pays the second highest bid.

Most common auction formats (2)

- Ascending price auction
  - Price starts low and increases gradually.
  - Bidders drop out.
  - The bidder who stays longest wins and pays the price where second last bidder drops out
  - Common for art, antique, company take-overs, ...
- A variant: descending price auction
  - Price starts high and falls until someone buys
  - Also called Dutch auction (as in Dutch flower auctions)

Simple example

- A seller with a single object to sell and two possible buyers.
- Valuation of the object is zero for the seller, and v_1 and v_2 to the buyers.
- Valuations v_1 and v_2 are
  - Independently drawn from uniform distribution $[0, 1]$.
  - Private information of the buyers.
- What is the best way for the seller to sell the object?

What is the best way to sell in terms of *revenue?*

- Posted price?
- First-price auction?
- Second-price auction?
- Ascending auction?
- Something else (what?)

What is the best way to sell in terms of *efficiency?*

- Posted price?
- First-price auction?
- Second-price auction?
- Ascending auction?
- Something else (what?)

Posted price

- Seller posts a price and buyers announce whether or not to buy
- If both want to buy, object allocated randomly (rationing)
- If none wants to buy, seller keeps the object
- What is the optimal price?
- What is the expected revenue?
- Is allocation efficient?

Second price auction

- Let us next consider second-price sealed bid auction.
- Both bidders submit simultaenously a sealed bid (e.g. write it on a paper and submit to the seller).
- Bidder who submitted the highest bid wins, but pays the second highest bid.
- This is a game between buyers:
  - The strategy for each bidder is simply the bid.
  - How should you bid?

Second price auction

- Claim: irrespective of the other bidder's strategy, it is optimal to bid one's valuation.
- In the terminology of game theory: bidding own valuation is a *dominant strategy*

- Why?
  - Consider an alternative strategy (bid above/below your valuation).
  - Would such a deviation affect what you pay if you win?
  - Would such a deviation affect whether or not you win? If so, when? Would you be happy about that effect?
- As a result, in equilibrium every bidder bids their true value.
  - Bidder with the highest value wins.
  - Pays an amount equal to the the second highest value.
  - Allocation is efficient

- What is the expected revenue by the seller?
  - Revenue is equal to the second highest valuation (i.e., with two bidders, the lowest valuation).
  - Hence, expected revenue is the expectation of the second highest value.
  - How to compute this? Derive the probability distribution for the second-highest valuation (*second order statistic*), and compute its expectation.
- Let $G(b)$ denote the cumulative distribution function (c.d.f.) of the second order statistic:

$$G(b) = 1 - (1 - b)^2$$

- Can you derive this? How to compute expected revenue from here?

- With two bidders, expected revenue is

$$\mathbb{E} \min\{v_1, v_2\} = \frac{1}{3}.$$

(can you compute this?)

- Expected value of the winner is

$$\mathbb{E} \max\{v_1, v_2\} = \frac{2}{3}.$$

- Hence, surplus is split equally between seller and winning bidder (on expectation)

- What if there are more bidders?
  - With 3 bidders, it is easy to show that expected revenue is $1/2$
  - Expected value of the winner is $3/4$
  - Hence, total surplus increases, but the share that goes to seller increases too
- This generalizes: as N increases, the seller gets a larger and larger share of the total surplus
  - With 10 bidders, expected price is $9/11$ and expected value of winner is $10/11$

First price auction

- Next, consider the first price sealed bid auction.
- As above, bidders submit bids simultaneously.
- Highest bid wins, but now the winner pays her own bid, i.e. the highest bid.
- Does this imply a higher revenue to the seller?

- Is it now optimal to pay your own bid?
  - Clearly you should bid less.
  - But how much less?
- Submitting a lower bid will
  - Increase the surplus if winning.
  - Decrease chances of winning.
- Optimal bid will depend on what you think the other(s) will do (unlike with second price auction).
- We need to consider a full *equilibrium analysis.*

Bayesian Nash equilibrium

- This is a game of incomplete information: each bidder knows privately her own value.
- Each bidder's equilibrium strategy must maximize her expected payoff accounting for the uncertainty about other bidders' values:

Definition

A set of bidding strategies is a Bayesian Nash equilibrium if each bidder's strategy maximizes her expected payoff given the strategies of the other bidder(s).

- We will analyze this thoroughly in the next lecture, but for now it suffices to note that since each bidder know privately her valuation, a strategy must determine what a bidder bids as a function of her valuation.

Finding the equilibrium bid function

- This example with two players and uniform value distributions can be solved easily by a simple trick (we will analyze the more general model later).
- Suppose bidder 2 uses bidding strategy $b_2(v_2) = \beta v_2$ for some $\beta > 0$.
- What is then the optimal bid for bidder 1? Suppose bidder 1 has value v_1, and consider payoff of bidding b:

$$\begin{aligned}
\pi(b; v_1) &= \Pr(win)(v_1 - b) \\
&= \Pr(\beta v_2 < b)(v_1 - b) \\
&= \Pr\left(v_2 < \frac{b}{\beta}\right)(v_1 - b) \\
&= \frac{b(v_1 - b)}{\beta}.
\end{aligned}$$

- This is maximized by choosing $b = \frac{1}{2} v_1$.

Finding the equilibrium bid function

- So, if bidder 2 uses a linear bidding strategy, the *best response* of bidder 1 is to use a linear bidding strategy $b_1(v_1) = \frac{1}{2}v_1$.
- Hence, if both bidders bid half of their value, they are both best-responding to each other.
- In other words, this is a Bayesian Nash equilibrium. In this equilibrium, both bidders use strategy

$$b_i(v_i) = \frac{1}{2}v_i,\ i = 1, 2.$$

Efficiency and revenue

- How do the properties of the equilibrium contrast with second price auction?

- Bidder with the highest value wins here too: auction is efficient.
- How about expected revenue? Let us compute:
  - Remember, expected highest value is $\mathbb{E}\left(\max\left\{v_1, v_2\right\}\right) = \frac{2}{3}$
  - Therefore, expected price is $\mathbb{E}\left(\max\left\{b_1\left(v_1\right), b_2\left(v_2\right)\right\}\right) = \frac{1}{2}\frac{2}{3} = \frac{1}{3}$.
  - This is the same as with second price auction!
- Is this a coincidence?

Ascending auction

- Finally, consider the ascending auction.
- Price starts ascending from 0 and bidders indicate their willingness to buy by staying in the game.
- As soon as one bidder drops out (e.g. say "I give up"), the remaining bidder wins and pays the standing price.
- This is a game, where the strategy of each bidder is to decide when to "stop" (i.e. drop out).
- When should you stop?

Ascending auction

- The optimal strategy is: stay in the game until price hits your valuation.
- This strategy is optimal *irrespective of the strategy of the other player.* (Why?)
- Bidder with the highest valuation wins and pays the second highest value.
  - Outcome is equivalent to the second-price auction.

Revenue equivalence theorem

- The equivalence of expected revenue in first price auction and ascending/second price auction is a manifestation of so called *Revenue equivalence theorem.*
- As we will see formally in the next lecture, it holds to any auction format where highest value bidder always wins.
- For example, the expected revenue would be the same in All-pay auction
  - Bidders submit bids, high bidder wins, and everyone has to pay their own bid.
  - Winner pays on average less than in standard formats, but expected total payment is the same since also losers pay.
  - Not commonly seen as an auction format, but used as a stylized model of contests (e.g. political lobbying or R&D race).

Reserve price

- Is there any way for the seller to increase expected revenue?
- Suppose the seller sets a reserve price r, i.e. minumum accepted price.
- Is it a good idea?
  - Potential benefit: higher price.
  - Drawback: maybe no sale (if all bidders have value below r).
- Consider second-price auction with reserve price $r = \frac{1}{2}$ and compute expected revenue. Note:
  - if $\min \{v_1, v_2\} > r$, then price is $\min \{v_1, v_2\}$.
  - if $\min \{v_1, v_2\} < r < \max \{v_1, v_2\}$, then price is $p = r$.
  - if $\max \{v_1, v_2\} < r$, then there is no trade.

- Can you compute the expected revenue? (it is indeed higher than without reserve price)
- One can show that $r = \frac{1}{2}$ is the optimal reserve price in this case
- The auction is not efficient: sometimes there is no trade at all even when bidders have positive values.
- Standard lesson about monopoly power applies in auctions too:
  - Monopolist distorts allocation (causes inefficiency) in order to transfer consumer surplus into profit.

Auction design

- We saw that the seller can increase profits by using a reserve price
- Are there other instruments that the seller could use?
- Are there other issues that should be taken into account in designing the auction?
- In real world, auction design is often a complicated problem:
  - Think about your reading assignment. What makes things complicated there?
- We consider next three important issues thorough examples:
  - How to treat asymmetric bidders?
  - How to ensure sufficient entry?
  - How to deter collusion?

Bidder subsidies and set-asides

- In real auction it is common that seller treats some bidders preferentially. Why?
- Distributional reasons:
  - Government favoring domestic bidders, municipal favoring local producers in procurement, etc.
  - Favoring of small businesses by subsidies or restricting entry (exclusions, or set-asides)
- Competition, or other post-auction market reasons:
  - Make sure there is sufficient competition in the market after auction
- Is it possible to increase revenue by subsidies?
- Let us look at a specific example with asymmetric bidders

Example of bid subsidies

- Two bidders with private values v_1 and v_2.
- Suppose the bidders are ex-ante asymmetric in the following sense:
  - Valuations are independently drawn from

$$\begin{aligned} v_1 &\sim U[0,1], \\ v_2 &\sim U[0,2]. \end{aligned}$$

- Consider an ascending auction (or equivalently, second price auction)
  - Both bidders bid up to their values and the higher value bidder wins.
  - This is more likely to be bidder 2.
- What is the expected price?

- Consider two equally likely events:
  - Bidder 2 has value $v_2 > 1$
  - Bidder 2 has value $v_2 < 1$
- In the former case, bidder 2 wins and pays on expectation $1/2$
- In the latter case, each bidder as likely to win, and expected price $1/3$
- So, bidder 2 wins with probability $\frac{3}{4}$ and the expected revenue is $\frac{1}{2}\frac{1}{2} + \frac{1}{2}\frac{1}{3} = \frac{5}{12}$.

- Suppose the seller gives 50% discount to the weaker bidder (bidder 1)
- What is the optimal bidding strategy of bidder 1?
  - Bid up to $2v_1$
- Behavior of bidders is as if both bidders have values drawn uniformly from $[0, 2]$
- As a result, both bidders are as likely to win
- Expected "clock price" is now $\frac{2}{3}$
- But taking into account the subsidy payment, the expected revenue of the seller is

$$R = \frac{1}{2}\frac{2}{3} + \frac{1}{2}\frac{1}{3} = \frac{1}{2}.$$

- Effect of subsidies:
- With no subsidy
  - Strong bidder is more likely to win ($\frac{3}{4}$ against $\frac{1}{4}$)
  - Expected revenue is $\frac{5}{12}$
  - Auction is efficient: higher value bidder always wins
- With subsidy:
  - Both bidders equally likely to win
  - Expected revenue is $\frac{1}{2} > \frac{5}{12}$
  - Auction is inefficient
- Again: seller gives up on efficiency to increase revenue

Entry of bidders

- A common problem in organizing auctions: how to ensure there are enough bidders participating?
- More bidders guarantees more competition
- But if bidders expect tough competition, why would they participate if entry is costly?
- This is a typical problem for example in procurement auctions, where it takes some work and effort for the participants to prepare offers
- Asymmetries can also be problematic

- Take the same example as above. Two bidders with independently drawn valuations:

$$\begin{aligned} v_1 &\sim U[0,1], \\ v_2 &\sim U[0,2]. \end{aligned}$$

- Second price auction / ascending auction
- Ex-ante expected payoffs of the two bidders (before they learn their valuations):
  - Bidder 1 expects to get $\frac{1}{12}$ (why?)
  - Bidder 2 expects to get $\frac{1}{2} + \frac{1}{12}$ (why?)

- Suppose now that there is a cost of $\frac{1}{10}$ to enter
  - Think of this as the cost of learning how much you value the good (cost of inspecting the procurement contract, cost of learning the production cost of service, etc.)
- Given this, bidder 1 should not enter at all
- Therefore, bidder 2 is the only one to enter and bids zero
- Not good for the seller...

How to promote entry of bidders in practice?

- Subside weaker bidders
  - Increase their payoff of entering, hence encourage entry
- Subsidize the entry costs directly
  - E.g. reimburse costs of preparing documentation for procurement contract offers
- Restrict the strong bidders from participating: set-asides
  - Excluding a strong incumbent may increase profits by inducing more competitive entry
- How about auction format?
  - In ascending price auction, the strong bidders can always respond in real time to weaker bidders.
  - Not good for entry (see your reading assignment).

Collusion

- Collusion occurs if bidders agree in advance or during the auction to let price settle at some low level.
  - This is illegal, but happens anyway.
- This occurs most naturally in situations, where there are multiple items for sale.
  - All bidders get a fair share, why raise price?
  - In extreme situations, incentives for price competition can be very low, even without formal collusion.
  - E.g. three similar objects, three bidders. Each bidder gets one, why raise prices?
  - Spectrum auctions?
- With a single object, collusion may rely on:
  - Side agreements: you win and share profits with me.
  - Intertemporal arrangement: you win today, I win tomorrow.

How to deter collusion?

- Tougher law enforcement?
- What about the auction format?
- Ascending auction
  - Suppose bidders 1 and 2 agree in advance that 1 should win.
  - What happens if bidder 2 deviates the agreement, and keeps on bidding as price increases?
  - Bidder 1 can bid back - makes deviation unprofitable and helps the collusion.
- Sealed bid auction
  - Again, suppose bidders 1 and 2 agree on bids such that bidder 2 wins.
  - But then bidder 1 can secretly outbid and steal the auction.
  - Deviation from agreement more tempting - makes it harder to sustain collusion.

Lecture: Formal analysis of auctions

- So far, we have 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.
- The goal now 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}, $u_i(\theta_i, \theta_{-i}) = u_i\left(\theta_i, \theta'_{-i}\right)$.
- **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
- $P_i(b_i, b_{-i}) = \begin{cases} 1 & \text{if } b_i > b_j, \forall j \neq i \\ \frac{1}{K} & \text{if } b_i \text{ ties for highest with } K-1 \text{ others} \\ 0 & \text{otherwise} \end{cases}.$
- $T_i(b_i, b_{-i}) = \begin{cases} \max_{j \neq i} b_j & \text{if } i \text{ wins} \\ 0 & \text{otherwise} \end{cases}.$

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)

- Deriving 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 at the end of this slide set (Additional material)
- 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.

Envelope Formula and Revenue Equivalence Theorem

- 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 $U'(\theta_i) = P_i(b_i(\theta))$, where $b_i(\theta)$ is the optimally chosen bid for type θ (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.

Envelope Formula and Revenue Equivalence Theorem

- 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 additional material at the end of this slide set for a direct derivation of the same formula.

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)

ADDITIONAL MATERIAL (For completeness): 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.

- 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-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,

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

- 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:

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

- 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.

- 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?

- 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 (as it should be revenue equivalence theorem).

Lecture: Common value auctions

- So far we have considered models, where
  - each bidder's value depends on his/her own signal only (private values), and
  - signals are independently drawn
- Recall the example: how much would you bid for a jar of coins?
- Here the value of the object is common to all the bidders, but different bidders have a different estimate about the value
- Do you care about the estimates of the other bidders?

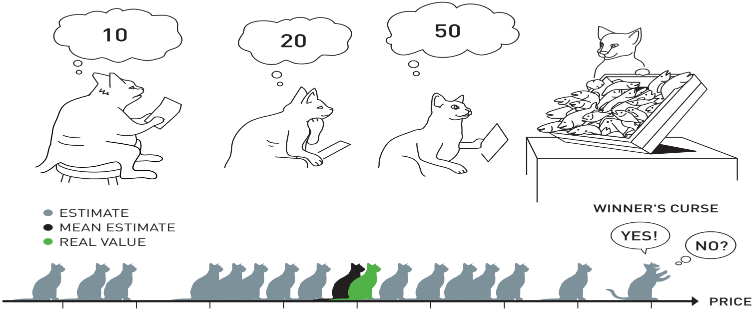
10
20
50
ESTIMATE
MEAN ESTIMATE
REAL VALUE
WINNER'S CURSE
YES!
NO?
PRICE

Winner's curse

- Winning means that all the other bidders were more pessimistic about the value than you.
- Winning is "bad news".
- Equilibrium bidding should take this into account.
- But how exactly?
- Do bidders take it into account in reality?
  - If not, then selling jars of coins is a money printing business
  - Experienced/inexperienced bidders

A simple model of common value auction

- Suppose that there is a common value v for the good, but its value is unknown.
- Formally, v is a random variable with some known probability distribution (e.g. Normally distributed)
- Both bidders observe a private signal that is correlated with the true value v. For example, we might have

$$\begin{aligned} \theta_1 &= v + \varepsilon_1, \\ \theta_2 &= v + \varepsilon_2, \end{aligned}$$

  where ε_1 and ε_2 are some i.i.d. random variables (e.g. Normally distributed noise terms)
- Then a high signal indicates that it is likely that also v is high

A simple model of common value auction

- This model is often called mineral-rights model
  - think of v as the true value of an mineral right, such as oil field
- Note: θ_1 and θ_2 are independently drawn, *conditional on v*
- But because v is unkonwn, θ_1 and θ_2 are correlated with each other through v
- Signals provide information about v (but only imperfect):
  - The expected value for bidder i based on her own signal is $\mathbb{E}(v \mid \theta_1)$
  - The expected value based on both signals is $\mathbb{E}(v \mid \theta_1, \theta_2)$
- It is natural to assume that these are increasing in signal values (a high signal predicts a high value)

A simple model of common value auction

- Recall from the previous lecture, we can specify an auction environment by defining the utility for a bidder if he wins as $u_i(\theta_i, \theta_{-i})$.
- In this case, we have:

$$u_i(\theta_i, \theta_{-i}) = \mathbb{E}(v | \theta_1, \theta_2).$$

- Hence, the utility of winning depends on both signals
- Moreover, the signals are correlated
- Hence, this is an auction with *interdependent values* and *correlated signals*

How to bid in a common value auction?

- Assume second price auction format
- Suppose bidder 2 uses strategy $b_2(\theta_2)$
- Bidder 1 has signal θ_1. How to bid?
- Consider bidding some p, or slightly more or less:
  - Makes no difference if $b_2(\theta_2) << p$, or if $b_2(\theta_2) >> p$
  - Only matters if $b_2(\theta_2) \approx p$
- The only situation where p is "pivotal" is when $b_2(\theta_2) = p$, i.e. $\theta_2 = b_2^{-1}(p)$.

How to bid?

- If bidder 1 wins being pivotal, her expected value for the object is

$$\mathbb{E}\left(v \left| \theta_1, \theta_2 = b_2^{-1}(p)\right.\right)$$

- To be indifferent between winning and not means

$$p = \mathbb{E}\left(v \left| \theta_1, \theta_2 = b_2^{-1}(p)\right.\right).$$

- Bidding more or less than p would lead to expected loss, so a best response strategy $b_1(\theta_1)$ for bidder one is to bid $b_1(\theta_1)$ that satisfies:

$$b_1(\theta_1) = \mathbb{E}\left(v \left| \theta_1, \theta_2 = b_2^{-1}(b_1(\theta_1))\right.\right).$$

How to bid?

- Hence, a symmetric Bayesian equilibrim is given by $b(\theta)$ that satisfies:

$$b(\theta_i) = \mathbb{E}(v \,|\theta_i, \theta_{-i} = \theta_i).$$

- It is optimal to bid as if the other bidder has exactly the same signal as you
- This generalizes to a symmetric model with N bidders:

$$b(\theta_i) = \mathbb{E}\left(v \,\middle|\theta_i, \max_{-i}\{\theta_{-i}\} = \theta_i\right).$$

- In other words, you should bid as if you have the highest signal, and the second highest signal within all the bidders is the same as your signal
- How would you now bid for the jar of coins?

No regret property

- The strategy that we derived shields against the winner's curse
- Suppose that bidder 1 wins:

$$b(\theta_1) > b(\theta_2) \iff \theta_1 > \theta_2$$

- Bidder 1 expected value post auction is $\mathbb{E}(v \mid \theta_1, \theta_2)$
- But her payment is $\mathbb{E}(v \mid \theta_2, \theta_2) < \mathbb{E}(v \mid \theta_1, \theta_2)$ (note: second price auction)
    - Bidder 1 is happy she won
- Bidder 2 expected value post auction is also $\mathbb{E}(v \mid \theta_1, \theta_2)$
- But to win, she should have outbid bidder 1, in which case she would have paid $\mathbb{E}(v \mid \theta_1, \theta_1) > \mathbb{E}(v \mid \theta_1, \theta_2)$
    - Bidder 2 is happy she lost!

Bidding in common value auctions

- The general idea in bidding in common value auctions: winning or losing conveys information about the information of the other bidders, so take this into account
- There is also a "loser's curse".
- Suppose that there are multiple identical objects for sale, say 10 bidders and 9 objects
- Suppose you lose. What does that tell about the value of the objects?

Winner's curse and IPO:s

- Winner's curse may have implications in other environments too
- Consider an initial public offering (IPO) of a company at price p:
  - All buyers have essentially the same value v for shares (unknown future trading price)
  - You should buy if you think $v > p$
  - If there is a lot of demand, then there is rationing (not every buyer gets shares)
  - What does it tell about other's information if you get shares?
  - Winner's curse?
- IPO:s are often underpriced. Why?

Revenue comparison between auction formats

- When signals are not independent, the Revenue equivalence theorem does not hold
- There is another principle called *Linkage Principle,* which allows for revenue comparison between different formats
- This important result is due Milgrom and Weber (1982): "A Theory of Auctions and Competitive Bidding", Econometrica.
- It turns out that second price auction is better for revenue than first price auction.
- The linkage principle also suggests that it is typically beneficial for the seller to release additional information about the object for sale (if she has any)

Information aggregation in common value auctions

- Where do asset prices come from?
- One view: prices reflect all the information that the traders have about asset values
- But how does price get to reflect that information?
- To investigate this question, we can model a financial market using an auction model
- The question is: can equilibrium price in an auction *aggregate* the bidders' information?

Information aggregation in common value auctions

- What is information aggregation?
- Suppose the value of an asset is v
- N bidders have an independent signal $\theta_i = v + \varepsilon_i$
- If N is large, then the median signal gives a very precise estimate of v:

$$Median(\theta_i) = v + median(\varepsilon_i) \approx v$$

  if for example $\varepsilon_i \sim N(0, \sigma^2)$
- "Wisdom of the crowds"
- But can the price in an auction aggregate information?
- If there is only one object, then not likely.

Information aggregation in common value auctions

- Assume a common value auction, with N bidders and K identical objects (think of N as a very large number)
- For simplicity, assume $K = N/2$
- Think of this as a market for an asset (K units, e.g. shares, and N bidders)
- The value of the asset is v and each bidder has a signal $\theta_i = v + \varepsilon_i$
- Auction format is a generalization of second price auction: $K + 1^{st}$ price auction
- Equilibrium bidding function can be shown to be

$$b(\theta_i) = \mathbb{E}\left(v \,\middle|\, \theta_i \text{ ties with the } K + 1^{st} \text{ highest signal}\right).$$

- Intuitively: bid as if you were just pivotal

Information aggregation in common value auctions

- But then

$$
\begin{aligned}
b(\theta_i) &\approx \mathbb{E}(v \,|\theta_i \text{ is median signal}) \\
&= \mathbb{E}(v \,|v + median(\varepsilon_i) = \theta_i) \\
&= \theta_i
\end{aligned}
$$

- Price will be $b\left(\theta^{(K+1)}\right)$, where $\theta^{(K+1)}$ is the $K + 1^{st}$ highest signal
- So the auction price will be approximately the median signal, and hence aggregates information!
- This model is a very simplified version of Pesendorfer and Swinkels (1997): "The loser's curse and information aggregation in common value auctions", Econometrica.

Conclusions

- Winning (or losing) reveals information about others' estimates
- Taking into account winner's curse requires caution in bidding
- Auction price can aggregate information

Some further readings on auctions

- These books were already mentioned:
  - A great book at advanced MSc/PhD level: Krishna: Auction Theory (Academic Press)
  - Another one is: Milgrom: Putting Auction Theory to Work (Cambridge University Press)
- A broad (but a bit old by now) survey on auctions is Klemperer (2002): "Auction Theory: A Guide to the Literature", Journal of Economic Surveys.
- An empirical analysis of collusion in auctions: Asker (2010): "A Study of the Internal Organization of a Bidding Cartel", American Economic Review.
- For discussion on practical issues on auction design, see Klemperer (2002): "What Really Matters in Auction Design", Journal of Economic Perspectives.

- For on-line auction applications, see e.g.
  - Edelman, Ostrovsky, Schwarz (2007): "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords", American Economic Review.
  - Varian (2009): "Online Ad Auctions", American Economic Review (Papers and Proceedings)
  - Varian and Harris (2014): "The VCG Auction in Theory and Practice", American Economic Review (Papers and Proceedings).