Lecture 14

Markets, Mechanisms and Machines

David Evans and Denis Nekipelov

Auctions: an overview

- Explains price formation
- Widely used selling game

Explore strategic behavior of:

- Bidders (usually buyers)
  - What bid to submit?
- Sellers
  - Which auction format to use?
  - Which selling game
  - Whether to restrict participation
  - Whether to charge entry fees

Examples

- Auctions used for many transactions in the Ancient world (marriage auctions in Mesopotamia, auctions for debt claims in ancient Greece)
- Art auctions (Christie's, Sotheby's)
- Real estate, treasury bills, electricity, livestock
- Large corporations are sold at auction
- Government procurement (highway construction), spectrum licenses
- Online advertising auctions

Auction formats

- A variety of formats are used to sell items
- Single item auction formats
  - English auction
    - Bidders call out successively higher prices until one bidder remains (Sotheby's and Christy's: Hammer auctions)
    - Japanese auction: seller continuously increases price, bidders drop out gradually and irrevocably by pressing a button
  - Vickrey or 2nd price auction:
    - Bidders submit sealed bids; high bidder wins and pays second highest bid
  - Dutch or descending price auction
    - opposite of English auction, Price falls until one bidder presses button, bidder gets object at the current price (Dutch flower auction)

Auction formats

- First Price sealed bid auction
  - Bidders submit sealed bids; high bidder wins and pays his bid
  - Construction contracts, governmental procurement
- Multi items auction formats
- Discriminatory Auction
  - A seller has an supply of items (possibly increasing in p)
  - Buyers submit downward sloping demand schedules (p; q combinations)
  - Equilibrium supply where aggregate demand equals supply
  - Buyers pay their bid for sold items

Auction formats

- Uniform Price
  - A seller has an supply of items (possibly increasing in *p*)
  - Buyers submit downward sloping demand schedules (*p*; *q* combinations)
  - Equilibrium supply where aggregate demand equals supply
  - Buyers pay the equilibrium price (where aggregate demand equals supply)
- Vickrey Auction
  - Win k units, then pay k highest opponents' losing bids (first highest losing bid for top unit, second highest losing bid for second unit, ...)

Auction formats

- Simultaneous ascending price auction (Milgrom (2000))
  - Each bidder demands one unit,
  - Bids are raised in multiple rounds,
  - In each round bidders specify which object that they are bidding for, and may switch from bidding for one object to bidding for another object
  - Auction closes when no further bids are raised
- Combinatorial Auction
  - Submit bids for stand-alone items and also for combination of items
  - Most expensive bidder/item allocation wins

Strategic equivalence

- When do auctions yield the same outcome? When are the bidding strategies identical?
- Example:
  - First price and Dutch auctions
  - Rational bidders, think about bidder giving instructions to an agent
  - In Dutch auction: a price at which to jump in.
  - Would do the same in a first price auction
  - Intuition: no information is revealed in a Dutch auction.
- Under some conditions there is also a strategic equivalence between the 2nd price and the English auction

Informational environment

- Private values:
  - Each bidder *i* values the item at a (privately) known value v_i
  - Other bidders do not know vi but know that v_i is drawn from some probability distribution
  - Example: construction contract in which firms know their own cost but not other firms' costs
- Common values
  - Same value for all bidders
  - Each bidder has a signal of the true value
  - Example: oil field as the value of oil is the same to everyone

Informational environment

- Affiliated values
  - a mixture between private and common values
- Interdependent values
- Reserve price: R
  - seller announces a minimum price prior to the auction, $b \geq R$
  - Reserve price may be kept secret

Vickrey (2^{nd} price) auction

- Rule: High bidder wins and pays the second highest bid
- N bidders
- Common model: private values, each bidder's value $v_i \in [0,V]$ known to bidder i but not known to other bidders
- Bidder i wins if her bid is the highest
  - Gets payoff $v_i - b_{(2)}$ ($b_{(1)}, b_{(2)}, \ldots, b_{(N)}$ are order statistics of the set of submitted bids)
- Otherwise she gets 0
- **Theorem**: *Every bidder bids their true value is a dominant strategy equilibrium.*

Vickrey (2nd price) auction

- Vickrey auction is efficient (item is allocated to bidder with the highest value)
- Expected Revenues (for the seller) of the Vickrey auction equal the expected second highest valuation
- Other equilibria?
  - yes
- Suppose $b_1 = V$ and all other bidders bid $b_i = 0$
- This is an equilibrium as nobody benets from deviating, but it is not a dominant strategy equilibrium

English auction

Setup (with private values):

- Button auction with continuously increasing prices
- Observe other bidders drop out prices
- No bidding costs

- Strategy: Press button until the price reaches your value v_i
- Is this an equilibrium?
  - Follow the proof for the Vickrey auction
- Other variants of the English auction may feature:
  - Discrete price increases; Open access: bidders may re-enter later-on; Bidding costs.

First price sealed bid auction

Setup:

- One object auctioned off and $R = 0$
- N bidders with private values $v_i \in [0,1]$ (normalization of support of values)
- Beliefs of other bidders about v_i drawn from distribution with density f: $[0, 1] \rightarrow \mathbf{R}_+$ (cdf $F(x) = \int_0^x f(z)dz$)
- Note: $v_i \perp v_j$
- Pay-your-bid payment rule:
  - If b_i is the highest bid, payoff is $v_i - b_i$

First price sealed bid auction

- Expected payoff (interim utility)

  $U_i(b_i\,; v_i) = (v_i - b_i)\Pr(b_i > b_j\ \forall j \neq i)$

- Strategy $\beta_i : [0, 1] \rightarrow \mathbf{R}_+$ (maps values to bids)
- In BNE: each bidder chooses b_i that maximizes expected payoff given v_i and beliefs regarding values of other bidders
- Symmetric equilibrium: bidder i with value b_i picks $b_i = \beta(v_i)$
- Claim: if density $f>0$ on its support then $\beta(v)$ is strictly monotone

First price sealed bid auction

- Determining the probability of winning

$\Pr(b_i > b_j\ \forall j \neq i) = \Pr(b_i > \beta(v_j)\ b_j\ \forall j \neq i)$

$= \Pr(\beta^{-1}(b_i) > \beta^{-1}(\beta(v_j))\ \forall j \neq i)$

$= \Pr(\beta^{-1}(b_i) > v_j\ \forall j \neq i)$ (monotonicity)

$= \Pr(\beta^{-1}(b_i) > v_1, \beta^{-1}(b_i) > v_2, \ldots, \beta^{-1}(b_i) > v_N)$

$= \Pr(\beta^{-1}(b_i) > v_1) \ldots \Pr(\beta^{-1}(b_i) > v_N)$ (independence)

$= F(\beta^{-1}(b_i)) \ldots F(\beta^{-1}(b_i)) = F^{N-1}(\beta^{-1}(b_i))$ (symmetry+beliefs)

- Expected payoff (interim utility)

$U_i(b_i\,; v_i) = (v_i - b_i)\Pr(b_i > b_j\ \forall j \neq i) = (v_i - b_i)\ F^{N-1}(\beta^{-1}(b_i))$

First price sealed bid auction

- Utility is determined by interim allocation rule $x_i(b_i) = F^{N-1}(\beta^{-1}(b_i))$ (allocation probability as a function of bid)
- In a symmetric BNE: Bids maximize utilities; same values should lead to same bids
- Thus
  - $\frac{d}{db_i} U_i(b_i\,; v_i) = 0$, which produces $b_i = \beta(v_i)$
  - $\beta^{-1}(b_i) = v_i$ and $F^{N-1}(\beta^{-1}(b_i)) = F^{N-1}(v_i)$

First price sealed bid auction

$$\frac{d}{dv_i}U_i(\beta(v_i); v_i) = \frac{\partial}{\partial v_i}U_i(\beta(v_i); v_i) + \frac{\partial}{\partial b_i}U_i(\beta(v_i); v_i)\ \beta'(v_i)$$

$$= F^{N\text{-}1}(v_i) = x_i(v_i) \text{ (allocation rule for values)}$$

Because:

- FOC holds and $\frac{\partial}{\partial b_i}U_i(\beta(v_i); v_i) = 0$
- Utility $U_i(b_i; v_i)=(v_i - b_i)\ F^{N\text{-}1}(v_i)$ and $\frac{\partial}{\partial v_i}U_i(\beta(v_i); v_i) = x_i(v_i)$

At the same time

$$U_i(\beta(v_i); v_i)=\int_0^{v_i} \frac{\partial}{\partial v_i}\, U_i(\beta(z); z)dz = \int_0^{v_i} x_i(z)dz$$

First price sealed bid auction

- Collecting information
  - $U_i(\beta(v_i); v_i) = (v_i - b_i)\, F^{N-1}(v_i)$
  - $U_i(\beta(v_i); v_i) = \int_0^{v_i} F^{N-1}(z)\, dz$
- This produces explicit expression for bid function

$$\beta(v_i) = v_i - \frac{\int_0^{v_i} F^{N-1}(z) dz}{F^{N-1}(v_i)}$$

- This demonstrates that bid function is indeed differentiable and monotone

First price sealed bid auction

- **Example:** N bidders with uniformly distributed values on [0,1]
- **Theorem:** $\beta(v)$=E[$v_{(2)} \mid v_{(2)}<v$]
- Note that $\Pr(v_{(2)}<v)$ = (cdf of 2nd order statistic) = $x(v)$
- Density $f_{v(2)}(v)=\frac{d}{dv}\Pr(v_{(2)}<v) = x'(v)$ and

$$f_{v(2)}(z \mid v_{(2)}<v) = x'(z)/\, x(v)$$

- Therefore E[$v_{(2)} \mid v_{(2)}<v$]=$\int z\, f_{v(2)}(z \mid v_{(2)}<v)\, dz$

$$= \int_0^v zx'(z)dz \,/\, x(v) = v - \int_0^v x(z)dz/x(v) = \beta(v)$$

Revenue equivalence

- How do auction formats compare?

Theorem: *The expected revenue to the seller is the same in a first-price and second-price auction.*

- First price auction bid equals the expected second highest valuation
- Second price auction payment equals the second highest valuation
- In expectation they are the same. Thus, seller's expected revenues are identical

Bidder surplus equivalence

- How do auction formats compare?
- Note that both first and second price auctions allocate to the highest value bidder (since the bid function in the first price auction is monotone in values)
- Thus, both auction formats lead to the same welfare
- Auction welfare is equal to the sum of revenue of the auctioneer and the surplus of bidders
- Since the revenue of the auctioneer is the same for both formats (by revenue equivalence theorem), the surplus of bidders has to be the same too

Revenue optimization

- First and second price auctions are efficient: they maximize social welfare by allocating the item to the bidder with the highest value
- Can auctioneer optimize her revenue by, possibly, sacrificing efficiency
- Yes, by setting a reserve price
  - Reserve price excludes some bidders (whose values are below the reserve)
  - BUT it also incentivizes the remaining bidders to raise their bids
  - This results in an increased revenue of the auctioneer

Revenue optimization

$$\text{Revenue}(R) = N\int_R^1 \beta(v)x(v)f(v)dv$$

$$= N\int_R^1 (v\text{-}\int_R^v x(z)dz/x(v))x(v)f(v)dv$$

$$= N\int_R^1 v\, x(v)f(v)dv \text{ - } \int_R^1 \int_R^v x(z)f(v)dvdz$$

Note that $\int_R^1 \int_0^v x(z)f(v)dvdz = \int_R^1 x(z) \int_z^1 f\,(v)dvdz$

$$=\int_R^1 x(z)(1\text{-}F(z))dz = \int_R^1 x(v)(1\text{-}F(v))dv$$

Revenue optimization

Revenue(R) $= N\int_R^1 \beta(v)x(v)f(v)dv$

$= N \int_R^1 v\, x(v)f(v)dv - \int_R^1 x(v)(1\text{-}F(v))dv$

$= N \int_R^1 (v - (1\text{-}F(v))/f(v))\, f(v)\, x(v)\, dv$

Also note that

Welfare(R) $= N\int_R^1 vx(v)f(v)dv$, so the surplus of bidders is $N \int_R^1 (1 - F(v))\, x(v)\, dv$

Revenue optimization

$$\text{Revenue}(R) = N\int_R^1 \left(v - \frac{1-F(v)}{f(v)}\right) F^{N-1}(v) f(v)\,dv$$

- Revenue is maximized when $R = (1 - F(R))/f(R)$
- Function $(1 - F(v))/f(v)$ is called Myerson's "virtual value"
- We maximize the revenue by discarding all bidders whose values are less than their virtual values
- Nobel prize, 2007

Using Myerson’s reserve prices

Reserve Prices in Internet Advertising Auctions:
A Field Experiment*

Michael Ostrovsky† Michael Schwarz‡

November 2016

Abstract

We present the results of a large field experiment on setting reserve prices in auctions for online advertisements, guided by the theory of optimal auction design suitably adapted to the sponsored search setting. Consistent with the theory, revenues increased substantially after the new reserve prices were introduced.

Using Myerson’s reserve prices

- Implemented on Yahoo! advertising auctions (each auction for keyword)
- Took 461,648 keywords
- For each keyword in the sample computed average number of advertisers bidding on this keyword, average bid, average standard deviation of bids
- Assumed that bidders' values were drawn from a lognormal distribution with a mean and a standard deviation to be estimated
- Used function $\beta^{-1}(b_i)$ to recover values from observed bids (remember, we are working with non-truthful bidding)
- Estimated Myerson reserves

Using Myerson’s reserve prices

Figure 1: The distribution of estimated keyword-specific optimal reserve prices (histogram)

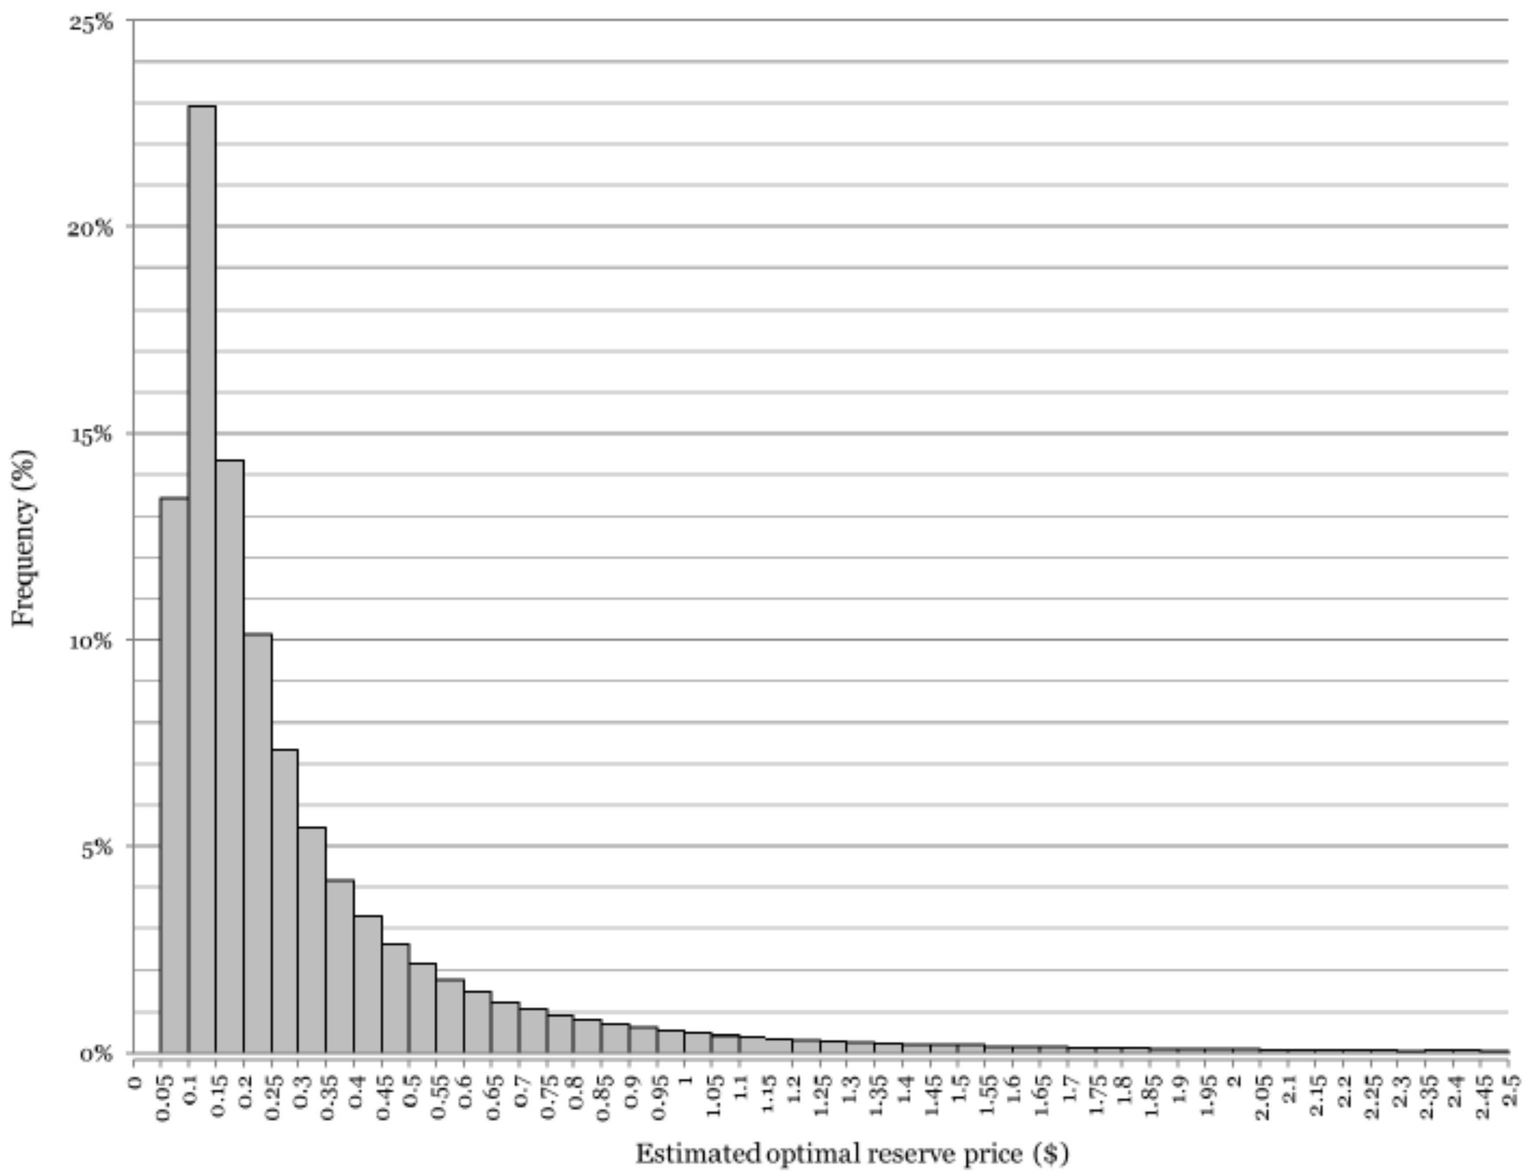

Using Myerson’s reserve prices

- Selected keywords into “treatment” and “control” groups
- 438,198 observations in the treatment group and 22,989 in control group
- Experiment run in May and June of 2008
- Use the diff-in-diffs approach to measure outcome of the experiment
- Outcome measures:
  - “Depth:” average number of advertisers whose bids exceed reserve price (ads are shown to users)
  - Monthly revenue per keyword
  - “Per search revenue:” average revenue generated by search engine every time a user searches for this keyword

Using Myerson's reserve prices

Table 6: Results (full sample, split by search volume)

	Full sample	< 10 searches per day	≥ 10 searches per day
Δ-in-Δ Depth	-0.91^{***}	-0.90^{***}	-0.97^{***}
• t-statistic	[−80.4]	[−75.5]	[−27.8]
• p-value	(< 0.0001)	(< 0.0001)	(< 0.0001)
Δ-in-Δ Revenue	3.80%	10.34%	2.06%
• t-statistic	[0.94]	[1.19]	[0.45]
• p-value	(0.347)	(0.235)	(0.653)
Δ-in-Δ Revenue per search	−1.45%	$-2.53\%^{**}$	$3.90\%^{**}$
• t-statistic	[−1.55]	[−2.36]	[2.31]
• p-value	(0.121)	(0.018)	(0.021)
N. obs. in treatment group	438,198	382,860	55,338
N. obs. in control group	22,989	20,133	2,856
Fraction of total revenue	100%	24.9%	75.1%

*, **, ***—*significant at 10%, 5%, and 1% levels. Changes in Revenue and Revenue per search are reported relative to the average revenue and average revenue per search in the corresponding subsample before the experiment.*