CSC200: Lecture 15

- Today
  - Recap, brief wrap up: First-price auctions: Ch.9.5, 9.7A
    - get a flavor for analysis in 9.7A (not mathematical details)
    - ignore “All Pay Auctions” in 9.5 and 9.7C
  - Additional auction topics:seller revenue, reserve prices (read 9.7B) and sellers curse (read 9.6)
  - Next topic is combinatorial auctions (not in text as a topic but related to material in Chapter 15, sponsored search)

- Announcements
  - Assignment 2 due November 25.
  - Term test : in tutorial Nov.27 (will announce scope later)
  - Quiz 4, final quiz of term on December 4

Revenue Equivalence

- Does seller make more in 1st or 2nd price auction?
  - In 2nd price: bidders *bid higher* (no shading), but seller *charges less than bid price*
  - In 1st price: bidders *bid lower* (shading), but seller *charges the actual bid price*
- How do they compare?
  - *seller gets same expected revenue (price) in both!*

Revenue Equivalence

- Consider example of uniform valuations on [0,m], and *n* bidders
- *Observation:* if you draw *n* numbers uniformly at random from *[0,m],* the expected value of the k^{th} lowest number is $\frac{k}{n+1}m$
- 2nd price auction (bidders bid true valuation):
  - expected value (EV) of highest bid (n^{th} valuation) is $\frac{n}{n+1}m$
  - EV of 2nd highest bid ($n\text{-}1^{st}$ valuation) is $\frac{n-1}{n+1}m$, which is *expected revenue*
- 1st price auction (bidders hedge):
  - expected value (EV) of highest (n^{th}) valuation is $\frac{n}{n+1}m$
  - highest bidder expected to bid $\frac{n-1}{n}$ times valuation: $\frac{n-1}{n} \cdot \frac{n}{n+1}m = \frac{n-1}{n+1}m$, which is *expected revenue*
  - same expected revenue as 2nd price auction!

Revenue Equivalence

E.g., 3 bidders drawn from U[0,100]

25

Expected valuation of $b_{(3)}$ (3^{rd} highest v)

50

Expected valuation of $b_{(2)}$ (2^{nd} highest v)

75

Expected valuation of $b_{(1)}$ (highest v)

2^{nd} price : $b_{(1)}$ wins and pays $v_{(2)}$: in expectation, 50

1^{st} price : $b_{(1)}$ wins with bid $2/3*v_{(1)}$: in expectation, $2/3*75 = 50$

- *Revenue equivalence:* making certain assumptions about bidder valuations, it has been shown that many different forms of auctions satisfying certain reasonable conditions must give same expected revenue to the seller! [Vickrey, Myerson and others]
- Neither of these maximize seller revenue… as we will see in a moment.

Third-Price Auction

- Is a 3rd-price auction (assuming enough bidders) "truthful"? Is it a dominant strategy to bid your true value?
  - *No : why not? Who can gain by not being truthful?*
- Why would a seller use a 3rd price auction instead of a 2nd price auction?
- What if I had two iPhones to sell at once?
  - what if I could charge different prices for each one?
  - what if I had to *charge the same price* for each one?
    - 3rd price auction ensures truthful bidding since 2nd-highest bidder will still be willing to pay 3rd highest price
- Idea can generalize to *kth-price auctions* for *k-1* identical goods that don't admit price discrimination

Reserve Prices

- New Zealand began auctioning wireless spectrum in 1990: 2nd price auction

- *"In one extreme case, a firm that bid NZ$100,000 paid the second-highest bid of NZ$6.*

In another the high bid was NZ$7 million and the second bid NZ$5,000. (NZ$1 equaled US$0.55.)

An Otago University, student bid NZ$1 for a television license for a small city; no one else bid so he won and paid nothing.

The total revenue fell far short of the advance estimate: the cellular licenses fetched NZ$36 million, one-seventh the NZ$240 million that NERA had predicted."

-- J. McMillan, Selling Spectrum Rights, J. Econ. Perspectives, 1994

Reserve Prices

- Sellers can set a *reserve price*: a minimum price below which the item won't be sold
  - e.g., won't sell my iPhone unless highest bid exceeds $150
- Two possible goals:
  - *maximize social welfare:* if seller has true value for item, should keep item if value is higher than those of all bidders
  - *maximize revenue:* announcing a reserve can increase selling price
- In 2^{nd} price auction: reserve price is like seller inserting a bid
  - truthful bidding still dominant for bidders
  - if seller reserve/bid is true valuation, item goes to person who values it the most (high bidder, or seller): *maximizes social welfare*
  - prevents seller from giving item for less than it is worth (to the seller)

Reserve in 2nd-Price Auction

v_{seller} = $150

reserve = $150

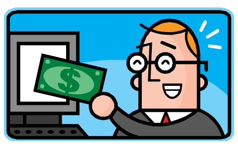

bid = $200?

bidder gets item

bid = $100?

seller keeps item

- Seller *values* item at $150
- Single bidder
- 2nd price-auction
- Seller sets *reserve* at $150
  - equal to seller value
- If highest bidder bids more than reserve (e.g., $200), 2nd highest less, bidder wins item, pays $150
  - since bid should be truthful, person valuing item the most gets it
- If highest bidder bids less (e.g. $100), seller keeps item
  - since bid should be truthful, person valuing item the most gets it

Reserve in 2nd-Price Auction

$v_{seller} = \$150$

reserve = $170!!

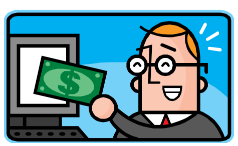

bid = $200?

bidder gets item

bid = $160?

seller keeps item

- Seller *values* item at $150
- Single bidder
- 2nd price-auction
- Seller sets *reserve* at $170
  - higher than seller value
- If (highest) bidder bids more than reserve (e.g., $200), bidder wins item and pays $170
  - maximizes social welfare
  - *improves seller revenue*
- If bidder bids $160, seller keeps
  - lost $10 in social welfare (person who values it most does not get it)
  - seller loses sale (possible surplus of $10)

Reserve Prices

- Should seller set a reserve price above her valuation?
  - depends on the tradeoff between the probability of an increased sale price and the probability that the item doesn't sell
  - in general, *revenue is maximized by setting a reserve above seller's valuation*, and running the risk of botching a potentially profitable sale
- For example, in our single bidder, 2nd price auction
  - suppose bidder values are uniform between *[0,200],* seller value = *150*
  - seller sets reserve price $r \geq 150$
  - probability of sale is: $(200-r)/200$ (e.g., at $r=180$, 10% chance of sale)
  - surplus of sale is $r-150$
  - so r should maximize expected surplus: $(r-150)(200-r)/200$
  - this happens at $r = 175$: which is midpoint between true seller value and highest possible bid
  - this is generally true for multiple bidders, uniform random values

Revenue Maximizing Auctions

- This shows: neither 1st nor 2nd price auction maximize seller revenue
- Maximizing revenue requires setting an optimal reserve price:
  - trade off probability of losing a sale with probability improving revenue if you make the sale
  - in general, depends on distribution of valuations among bidders
- The analysis on previous slides can be applied to more bidders in 2nd-price, to 1st-price as well
- General revenue maximizing auction: the *Myerson auction* **(FYI only):**
  - Assume all bidders i.i.d. with (same) CDF *F* (and PDF *f*)
  - Define *virtual valuation* of bidder k to be $\varphi_k(v_k) = v_k - \left({}^{1-F_k(v_k)}/_{f_k(v_k)}\right)$
  - Optimal revenue auction: run 2nd price with reserve set to $\varphi^{-1}(0)$
  - e.g., uniform valuations on [0,1]: *optimal reserve price* is $\varphi^{-1}(0) = {}^{1}/_{2}$
    - Why? *F(v) = v, f(v) = 1;* so v that has virtual valt'n of 0 is ½: ½ - (½ /1) = 0
    - *exercise:* convince yourself that higher/lower reserve is worse in this case
  - Myerson's analysis much more general

Winner's Curse

- Five companies bidding (1st-price) for oil drilling rights in area *A*
  - ultimate value is pretty much the same for each: a certain amount of oil *(B bbls);* each will sell it at market price (ignore technology differences)
- Why run an auction? It's worth the same to everyone…
  - but seller (gov't), and the companies *are unsure of the value!*
  - each produces its own *(private) estimate* of the reserves (quantity *B*)
    - e.g., each sends in own geologists, has personal interpretation of satellite images, has own test-drill results, etc.
    - value is now *random (probabilistic):* bid based on your *expected value*
- Estimates are related to *B*, but noisy (error-prone):
  - e.g., *U* estimates 50M bbl; *V*: 47M bbl; *W*: 42M bbl; *X*: 40M bbl; *Y*: 38M
  - who will bid most? who will win this auction?
  - Assume *U* wins. Once it does, it learns something about other's estimates:
    - *"everyone else's estimate was lower than mine"*
  - suggests *U*'s estimate was too high: perhaps *U* overpaid!

Winner's Curse

- Phenomenon is known as *winner's curse*
  - winning auction: implies value is less than you estimated
  - may still profit (attain a surplus), but could have negative (expected) surplus!
  - occurs in *any* common value auction (e.g., when buying items for resale)
- Can it be avoided?
  - sure: in equilibrium, you should bid based on your expected value "conditional on your bid winning"
  - in other words, anticipate how much you would value "winning", including any informational effects gained from winning (easier said than done!)
- Exercise: why do you think a loser (e.g., V) might regret not bidding higher? Define and explain why a *loser's curse* might occur.

Combinatorial Auctions

- There are much more general types of auctions when selling multiple items under various constraints. Common example: sellers offer many distinct items and buyers want different possible sets of items.
  - buyer's value may depend on the *collection* of items obtained and not just additively; and there may be different collections of interest.
- *Complements:* items whose value increase when combined
  - e.g., a cheap flight to Siena less valuable if you don't have a hotel room
- *Substitutes:* items whose value decrease when combined
  - e.g., you'd like the 10AM flight or the 7AM flight; but not both
- If items are sold separately, knowing how to bid is difficult
  - bidders run an *"exposure" risk:* might win item whose value is unpredictable because unsure of what other items they might win

Flight1

Simultaneous Auctions: Substitutes

Flight1 (7AM, no airmiles, 1 stopover)
Value: $750

Flight2 (10AM, get airmiles, direct)
Value: $950

Bidder can only use ***one*** of the flights:
Value of receiving both flights is $950

- If both flights auctioned *simultaneously*, how should he bid?
- Bid for both? runs the risk of winning both (and would need to hedge against that risk by underbidding, reducing odds of winning either)
- Bid for one? runs the risk of losing the flight he bids for, and he might have won the other had he bid
- If items auctioned *in sequence*, it can mitigate risk a bit; but still difficult to determine how much to bid first time

Simultaneous Auctions: Complements

Flight1

Hotel Room

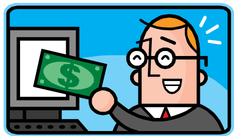

Bidder doesn't want flight without hotel room,
or hotel without flight; but together, value is $1250

- If flight, hotel auctioned *simultaneously*, how should he bid?
- Useless to bid for only one; but if he bids for both, he runs the risk of winning only one (which is worthless in isolation). Requires severe hedging/underbidding to account for this risk.
- If items auctioned *in sequence*, it can mitigate risk only a little bit. If he loses first item, fine. If he wins, will need to bid very aggressively in second (first item a "sunk cost") and could overpay for the pair.

Combinatorial Auction

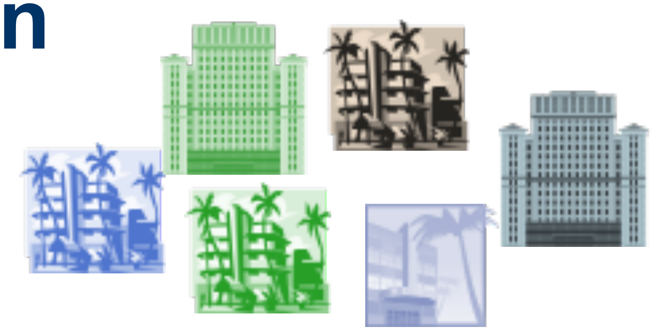

Bidder expresses value for ***combinations*** of items, e.g.:

- Value(flight2, hotel1) = $1250
- Value(flight1, hotel1) = $1050
- Don't want any other package

- *Combinatorial auctions* allow bidders to express *package bids*
  - for any combination of items, bidder can say what he is willing to pay for that combination or package
  - do not pay unless you get exactly that package
  - outcome of auction: assign (at most) one package to each bidder
  - can use 1st-price (pay what you bid) or Vickery-Clarke-Groves (VCG, a generalization of 2nd-price, which we will see later)

Combinatorial Auctions

- Formally:
  - a collection of *goods* *G* for sale
  - *bids* have the form $\{(S_1, v_1), (S_2, v_2), \ldots, (S_k, v_k)\}$, where:
    - each S_i is a subset of *G*, v_i is the price bidder will pay for S_i
    - can assign to any bidder *at most one* subset S_i from his bid
- Goal find an *assignment of goods* to bidders that maximizes the sum of the corresponding prices/valuations
  - i.e., if bidder gets the items that correspond to S_{17} in his bid, he will pay v_{17}; if items correspond to no subset in his bid, he pays nothing
  - sometimes "free disposal" assumed…
- But each item can be assigned to at most one bidder, so some hard choices need to be made by the seller

Combinatorial Auctions: Complexity

- Deciding how to allocate goods to bidders to maximize revenue (or social welfare) is computationally difficult (i.e. set packing problem)
  - formally, it is an *NP-complete problem*, which means that it is widely believed to require *exponential time to solve in the worst-case*
  - informally, it is not known whether you can do much better (in the worst case) than exhaustively searching all ways of assigning items to bidders
    - even *approximately* maximizing the social welfare is NP-hard
    - Recall our early discussion of computational difficulty
  - if I have *n* items and *m* bidders, there are $(m+1)^n$ such assignments (allow for the possibility of an item going to no bidder)

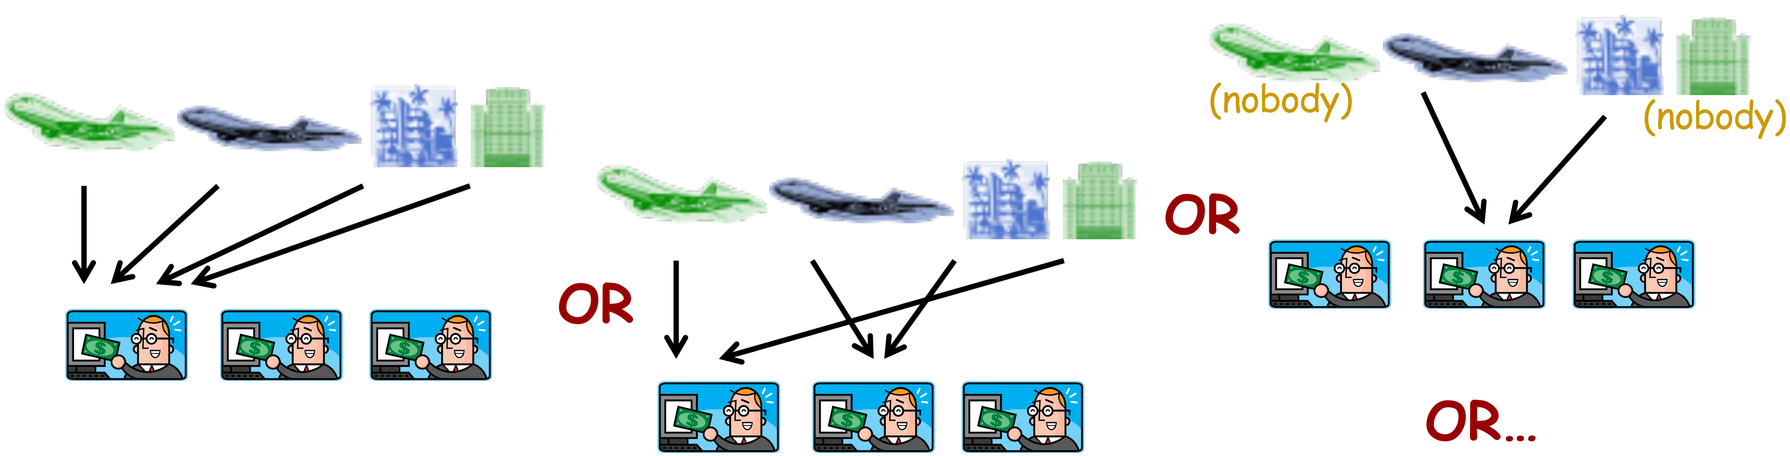

How Many Assignments?

Bidders	Items	Number of Assignments
2	10	59049
4	20	95,367,431,640,625
10	30	1.7×10^{31}
20	100	1.7×10^{132}
100	1000	2.1×10^{2004}

- Of course, you needn't look at assigning arbitrary subsets of items to a bidder: bidder only cares about certain subsets
- Suppose each bidder specifies *k* different subsets in his bid
  - Most bidders may be interested in "relatively few" combinations of things…
- Need only consider possible selections of a single subset from each bidder: test for feasibility (no overlap) and social welfare
- $(k+1)^m$ such subset assignments ignoring potential overlap (allowing "no subset" for a bidder)

How Many Subset Assignments?

Bidders	Number of Subsets in Bid	Number of Subset Assignments
2	10	121
4	20	194,481
10	30	819,628,286,980,801
20	100	$1.2x\ 10^{40}$
100	1000	1.1×10^{300}
200	10,000	1.0×10^{800}

- This is more manageable, but still not possible to try all.
- Still, CAs used and solved more and more widely used in practice
  - wireless spectrum, airport landing gates, industrial sourcing and procurement, etc…
  - how are they solved?
  - sometimes a good approximation (to optimal) is sufficient