CSC 196 Week 12 (December 6,8)

- Agenda (based on CSC200 slides with Craig Boutilier)
  - Mechanism design
  - Auctions (Ch.9 in Easley and Kleinberg)
- Announcements
  - We will try to have everything graded as soon as possible. For Q2 and A4 please make any regrading requests within 2 days. I will announce when grades are available for Q2 and A4.
  - Once you get your grades for all assignments and quizzes you will know everything but your participation grade which I will determine based on my records of attendance and class activity.

Games with Incomplete Information

- In basic game theory, we assume everybody knows the moves of the game (their own and opponents') as well as payoffs (their own and opponents')
  - We call this the *complete information* assumption
  - Holds in normal form and extensive form games
- This is also wildly unrealistic in most settings
  - Knowing fully even one's own actions/strategies is not always reasonable
  - Knowing payoffs of other players is even less realistic
    - These are often called *incomplete information games.* We'll focus on uncertainty in payoffs.

Mechanism Design

- Suppose policy maker wants to build new highway or a new hospital and ensure it is used to maximal benefit of society
- Can't determine maximal benefit unless you know individual preferences however
- Could you design a *game* that gets people (in equilibrium) to truthfully reveal their preferences for various outcomes?
  - e.g., the most they would be willing to pay to drive for *x* minutes
  - If so, you could assign outcomes that maximize ``social welfare''
- *Mechanism design:* a branch of game theory that does just this. Informally, one designs games where preferences (or payoffs) are unknown (private information), so when players act "rationally" (in equilibrium) a socially desirable outcome emerges.
- Mechanism design is therefore a type of algorithm design where **inputs are coming from self interested agents.**

Auctions

- *Auctions* provide a canonical application of mechanism design.
- Simple example: I want to give or sell my old iPhone6 to the person in class who values it most (this will maximize social welfare).
  - More importantly. consider gov't auctioning wireless spectrum, logging, fishing, oil exploration rights, etc.
- How should I (the government) do this to insure that I am optimizing the social welfare (ignoring whatever revenue I may or may not get?)
- Lets do this now as a sealed bid auction.
- I am asking everyone in a poll to state how much it is really worth to you and what you are willing to bid for it (the bids are your actions or strategies). I will pretend to only see the "bids".
- My goal is to give it to the person who values it the most. Can I use prices to make this happen? We will assume that a ``truthful'' bid would be to bid your value for the iPhone 6; that is you would rather have it than not have if you pay your value.

Auctions (selling my iPhone6)

- Our simple auction example: I want to give (or sell) my old iPhone6 to one person. I want to give it to person who values it the most. How should I do this? Should I set a price and if so how?

- Procedure 1: Outcome is to give it to highest bidder at no cost.
  - Why wouldn't you all exaggerate (assuming you personally want a used iPhone)?

- Procedure 2: Outcome: give to highest bidder, but charge bid price
  - You won't exaggerate now? But won't you understate your true value?

- Procedure 3: Outcome: give to highest bidder, but charge the second-highest bid price
  - Now how will you bid?

Auctions

- Auctions widely used (to both sell, buy things)
  - our focus will be on *one-sided, sell-side (forward)* auctions: that is, we have a single seller, and multiple buyers
  - examples: rights to use public resources (timber, mineral, oil, wireless spectrum), fine art/collectibles, houses/property (Australia, UK, …), Ebay ($60B volume/yr), online ads (Google, Facebook, Microsoft, …), …
- Variations:
  - *multi-item auctions:* one seller, sells multiple items at once
    - e.g., wireless spectrum, online ads
    - interesting due to substitution, complementarities
  - *procurement (reverse) auctions:* one buyer, multiple sellers
    - common in business for dealing with suppliers
    - government contracts tendered this way
    - aim: purchase items from *cheapest* bidder (meeting requirements)
  - *double-sided auctions:* multiple sellers and buyers
    - stock markets a prime example, *matching* is the critical problem

Single-item Auctions (Sell-side/Forward)

- Assume seller with one item for sale
- Several different formats
  - **Ascending-bid (open-cry) auctions** (aka *English auctions*)
    - price rises over time, bidders drop out when price exceeds their "comfort level"; final bidder left wins item at last drop-out price
  - **Descending-bid (open-cry) auctions** (aka *Dutch auctions*)
    - price drops over time, bidders indicate willingness to buy when price drops to their "comfort level"; first bidder to indicate willingness to buy wins at that price
  - **First-price (sealed bid) auctions**
    - bidders submit "private" bids; highest bidder wins, pays price s/he bid
  - **Second-price (sealed bid) auctions**
    - bidders submit "private" bids; highest bidder wins, pays price bid by the *second-highest bidder*

Why would seller use an auction?

- Let's assume the following:
  - seller *s* has a single item for sale
  - *s* values item at v_s (would rather keep than sell for less than v_s)
  - *s* is trying to maximize sale price (its own revenue)
  - set of potential buyers *B*, each *b* in *B* has value v_b for item
  - buyer values are *independent and private*
- *Independent, private valuations* means each buyer's value is "personal", does not depend on values of other buyers
  - consider items that will be consumed/used directly by *b*
- *Common values*
  - if *b* could resell item to another *b'*, then values are no longer independent or purely personal (*b* could speculate about value of item to others and buy it just for purpose of reselling)
  - values become correlated too if buyers are uncertain about the value of the item, and others have (different) private information about that value
    - *classic example: bidding for oil-drilling rights*

Why use an auction: what if values are known?

- If s knows values v_b of all buyers
  - s can offer item to b for price v_b (or a bit less), *if greater than* v_s
  - s would select the b^* with *highest value* v_{b^*}
    - alternatively, could just announce a *price* (a shade below) v_{b^*}
  - s extracts entire *surplus*: $v_{b^*} - v_s$
  - if b purchases at price v_b then net benefit or gain to b is zero
  - if b^* could bargain, could potentially reduce price
    - but s won't accept price below second-highest value: $v_{b(2)}$
- Notice that selling item this way maximizes social welfare
  - item goes to buyer (or stays with seller) that values it most
  - any price paid simply redistributes some of the surplus between the buyer and the seller (any price between v_{b^*} and v_s works)

Seller knows values: illustration

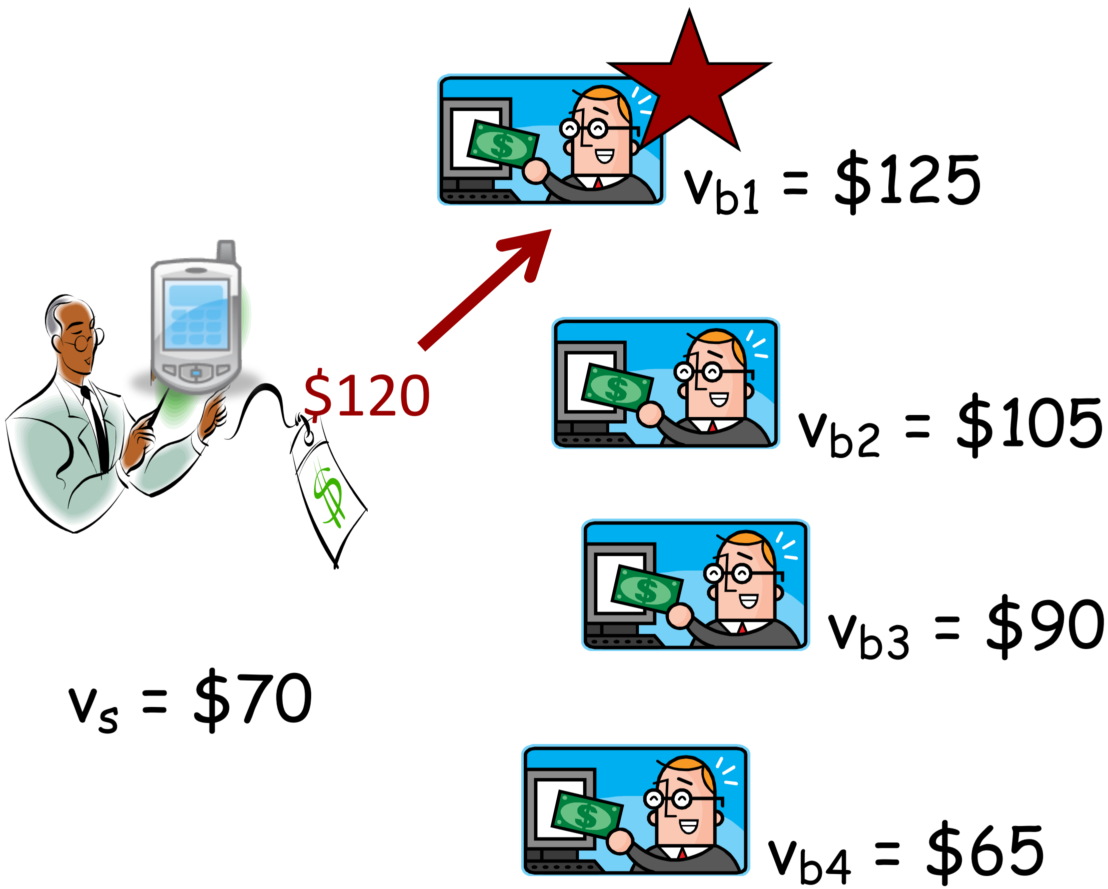

- Seller can set price anywhere between \$105 and \$125 (let's say \$120)
- Buyer b1 with highest value will accept price; gain to b1: \$5 (125-120); gain to s: \$50 (120-70)
- Social welfare increases by \$55 (125-70)
- Different price (or bargaining) splits the surplus between b1 and s (but s doesn't need to accept price below 105)

Why use an auction: values are unknown

- All sounds good, but *s* usually doesn't know values
  - what if *s* sets price too high?
    - No transaction, lose social welfare (and revenue)
  - what if *s* sets price too low?
    - Transaction occurs, but item could go to bidder with lower value than *b**, lose social welfare (and revenue)
- Auction format is a way of discovering preferences/values
  - can be used to maximize social welfare, revenue, other objectives such as ``fairness''

Second-price Auction

- Bidders submit “sealed” bids; highest bidder wins, pays price bid by *second-highest bidder*
  - also known as *Vickery auction*
  - special case of *Groves mechanism*, *Vickery-Clarke-Groves (VCG) mechanism* (which we’ll see in later chapters)
- 2^{nd}-price seems weird but is quite remarkable
  - truthful bidding, i.e., bidding your true value, is a *dominant strategy*
- To see this, let’s formulate it as a game

The Second-Price Auction Game

- *n* players (bidders)
- each player *k* has value v_k for item
  - assume v_k between [0,1] (for concreteness only)
- strategies/actions for player *k*: any bid b_k between [0,1]
- outcomes: player *k* wins, pays price *p* (2nd highest bid)
  - *more than n outcomes: outcome includes price paid by winner*
- payoff for player *k*:
  - if *k* loses: payoff is 0
  - if *k* wins, payoff depends on price *p*: payoff is $v_k - p$
- Notice: game differs in critical way from usual matrix form
  - *no player actually knows the payoffs of the other players*

Equilibrium: Second-Price Auction Game

- Even without knowing payoffs of others, it turns out that *bidding its true valuation is dominant* for every player k
  - strategy depends on valuation: but k selects b_k equal to v_k

- Let's see why deviation from *truthful bid* can't help (and could harm) k, regardless of what others do
- We'll consider two cases:
  a) if k wins with truthful bid $b_k = v_k$
  b) if k loses with truthful bid $b_k = v_k$

Equilibrium: Second-Price Auction Game

- Suppose k wins with truthful bid v_k
  - Notice k's payoff must be zero (if tied for 1st place) or positive
  - Now let's consider if k could have done better…
- Bidding b_k higher than v_k:
  - v_k already highest bid, so k still wins and still pays price p equal to second-highest bid $b_{(2)}$; so k is no better off
- Bidding b_k lower than v_k:
  - If b_k remains higher than second-highest bid $b_{(2)}$ then no change in winning status or price; so k is no better off
  - If b_k falls below second-highest bid $b_{(2)}$ then k now loses and is worse off (assuming v_k is great than $b_{(2)}$

Equilibrium: Second-Price Auction Game

- Suppose k loses with truthful bid v_k
  - Notice k's payoff must be zero and highest bid $b_{(1)} > v_k$
  - Now let's consider if k could have done better…
- Bidding b_k lower than v_k:
  - v_k already a losing bid, so k still loses and still gets payoff zero
- Bidding b_k higher than v_k:
  - If b_k remains lower than highest bid $b_{(1)}$, no change in winning status (k still loses)
  - If b_k is above highest bid $b_{(1)}$, k now wins, but pays price p equal to $b_{(1)} > v_k$ (payoff is negative since price is more than it's value)

- So a truthful bid is *dominant*: optimal no matter what others are bidding

Truthful Bidding in Second-Price Auction

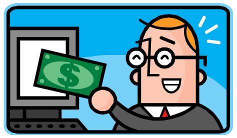

b_1 = \$125

v_2 = \$105
b_2 = ???

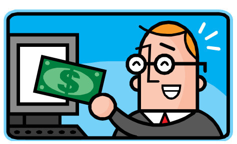

b_3 = \$90

b_4 = \$65

- Consider actions of bidder 2
  - Ignore values of other bidders, consider only their bids. Their values don't impact outcome, only bids do.
- What if bidder 2 bids:
  - truthfully \$105?
    - loses (payoff 0)
  - too high: \$120
    - loses (payoff 0)
  - too high: \$130
    - wins (payoff -20)
  - too low: \$70
    - loses (payoff 0)

Truthful Bidding in Second-Price Auction

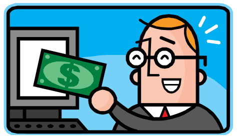

b_1 = \$95

v_2 = \$105

b_2 = ???

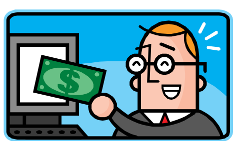

b_3 = \$90

b_4 = \$65

- Consider actions of bidder 2
  - Ignore values of other bidders, consider only their bids. Their values don't impact outcome, only bids do.
- What if bidder 2 bids:
  - truthfully \$105?
    - wins (payoff 10)
  - too high: \$120
    - wins (payoff 10)
  - too low: \$98
    - wins (payoff 10)
  - too low: \$90
    - loses (payoff 0)

Other Properties: Second-Price Auction

- Elicits true values (payoffs) from players in game even though they were unknown a priori. Is this a good or bad property?
- Allocates item to bidder with highest value (maximizes social welfare)
- Surplus is divided between seller and winning buyer
  - splits based on second-highest bid (which is the lowest price that the winner could reasonably expect)
- **Outcome is similar to English auction (ascending auction)**
  - consider process of raising prices (very slowly), bidders dropping out, until one bidder remains
  - until price exceeds *k*'s value, *k* should stay in auction
    - drop out too soon: you lose when you might have won
    - drop out too late: will pay too much if you win
  - last bidder remaining has highest value, pays 2^{nd} highest value!
  - games are in fact "strategically equivalent"; seller get same price (but now we don't learn the true value of item for the winner).
  - What can be different "in practice"?

First-price Auction

- Bidders submit "private" bids; highest bidder wins, pays her bid
- 1st-price auction seems more intuitive than 2nd-price
  - but it makes bidders reason more *strategically*
  - in fact it is not obvious exactly how to determine your bid
  - question: why don't you bid your true value for the item?
    - *if you win, your payoff is zero (you pay exactly what it's worth, i.e., the price at which you're indifferent between taking the item of leaving it)*
    - *this suggests you want to "shade" your bid lower than true value; but by how much?*
- To better understand a first-price auction, let's formulate it as a game

The First-Price Auction Game

- n players (bidders)
- each player k has value v_k for item
  - assume v_k between [0,1] (for concreteness only)
- strategies/actions for player k: any bid b_k between [0,1]
- outcomes: player k wins, pays price p (equal to her bid)
  - *more than n outcomes: outcome includes price paid by winner*
- payoff for player k:
  - if k loses: payoff is 0
  - if k wins, payoff depends on price p: payoff is $v_k - p$
- Notice: game is incomplete information (like 2^{nd}-price)
  - *no player actually knows the payoffs of the other players*

First-Price Auction: No dominant strategy

- Claim: there is no dominant strategy for any player k
- Other players bid: let highest *"bid from others"* be $b_{(1)}$
  - If value v_k is greater than $b_{(1)}$ then *k*'s *best bid* is b_k that is just a "shade" greater than $b_{(1)}$ (depends on how ties are broken)
  - This gives *k* a payoff of (just shade under) $v_k - b_{(1)} > 0$
  - If *k* bids less than $b_{(1)}$: *k* loses item (payoff 0)
  - If *k* bids "a lot" more than $b_{(1)}$: pays more than necessary (so *k*'s payoff is reduced)
  - Notice *k* should never bid more than v_k
- So *k*'s optimal bid *depends on what others do*
- Thus *k* needs some prediction of how others will bid
  - requires genuine *equilibrium analysis*

Bid Shading in First-Price Auction

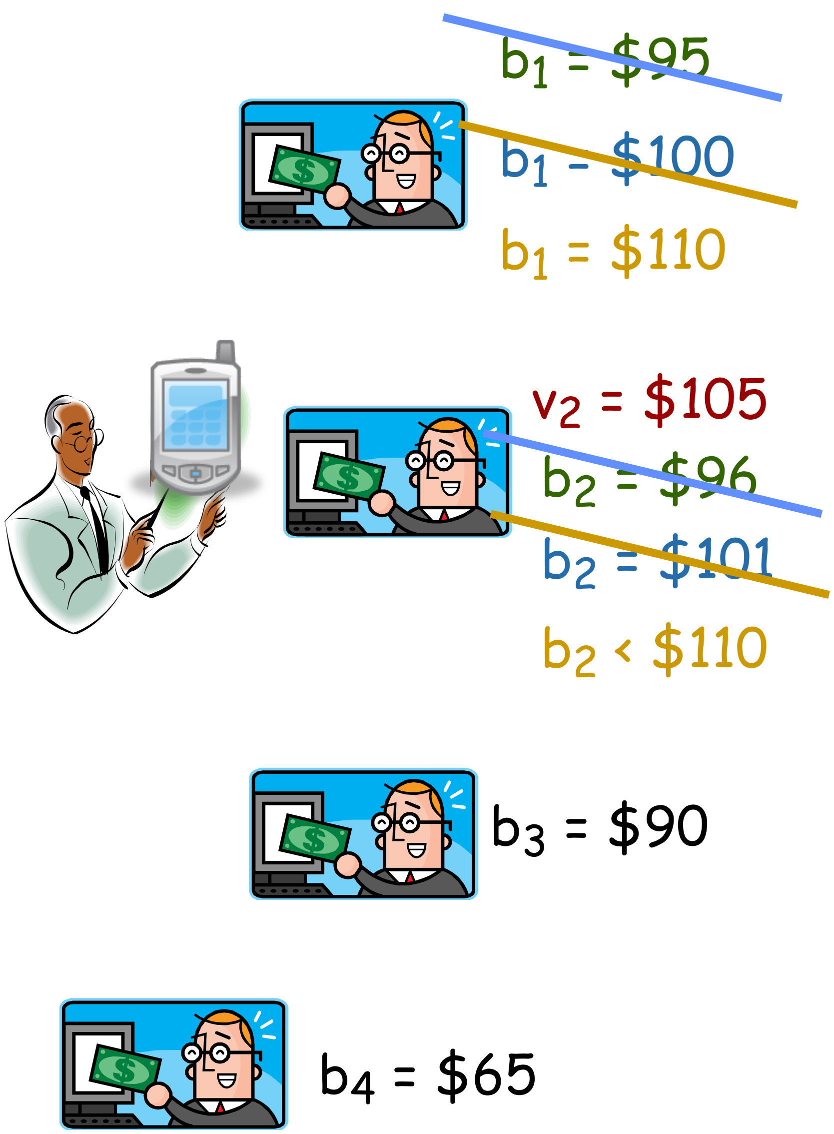

- Consider actions of bidder 2
  - ignore values of other bidders, consider only bids.
  - assume "bid increment" of \$1 and ties are broken against bidder 2
- If bidder 1 bids \$95:
  - bidder 2 *should bid* \$96
    - wins with payoff 9
  - if 2 bids \$94, loses (0)
  - if 2 bids \$97, payoff 8
- If bidder 1 bids \$100
  - bidder 2 *should bid* \$101
    - wins (payoff 4)
- If bidder 1 bids \$110
  - bidder 2 should bid "less"
    - loses (payoff 0)

First-Price Auction: How much to shade?

- Player *k* doesn't know values/bids of other bidders
  - cannot choose a strategy that is *guaranteed* to win or to maximize payoff
  - *k* needs to have some beliefs about how others will bid
  - *but how others will bid depends on their underlying values*
- Unlike earlier games, players need predictions (beliefs) about other players' *payoffs*, not just their strategies
  - This discussion becomes technical and we will skip it. One can work out the analysis without too much difficulty (for example) if we assume that all other bids are random between 0 and 1.
- Suppose bids of others are random between 0 and 1
  - *k* should shade its bid *more* (bid lower) when there are *fewer* competitors (the highest competing bid more likely to be lower)
  - No dominant strategy for any player but can analyze equilbrium.

What is expected value of strategy *s*?

- What is k's expected payoff for playing s?
  - Payoff is zero if k loses
  - Payoff is "value minus bid" if k wins: $v_k - s(v_k)$
  - So if k wins with probability p, expected payoff is $p\ (v_k - s(v_k))$
- What is probability k wins?
  - Since strategies are symmetric, k wins just when $v_k > v_j$
  - This happens with probability v_k
  - So k's expected payoff is $v_k(v_k - s(v_k))$

$Prob(v_j < 0.8) = 0.8$

$Prob(v_j > 0.8) = 0.2$

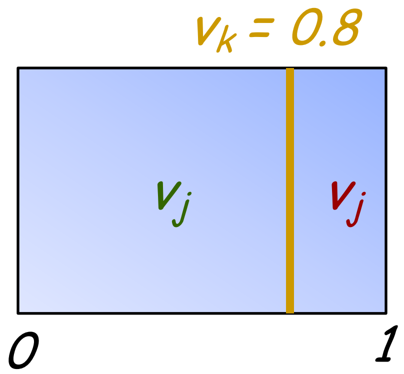

What is optimal bidding strategy?

- Want a strategy s where expected value of *bidding* $s(v_k)$ is better than bidding $s(v)$ for any other value v
  - If true value is v_k and bid is $s(v)$: probability of winning is v, and payoff if bidder wins is $v_k - s(v)$
  - So we want s satisfying: $v_k(v_k - s(v_k)) \geq v(v_k - s(v))$ for all v
- So payoff function $g(v) = v(v_k - s(v))$ should be maximized for the input v_k
  - we do this by finding derivative of g, setting the derivative to zero at point (input) v_k and solving a simple differential equation (see details in text if mathematically inclined)
- Result is: $s(v) = v/2$ (easy to *verify* that $s(v_k)$ is best bid for v_k)
- In other words, *the bidding strategy where both bidders bid half of their value is a Nash equilibrium*

For More Than Two Bidders

- Same analysis can be applied to uniform values: intuitive result
- If we have *n* bidders, a symmetric equilibrium strategy is for any bidder with value v_i to bid $(n-1)/n\ v_i$
  - e.g., if 2 bidders, bid half of your value
  - e.g, if 10 bidders, bid 9/10 of your value
  - e.g, if 100 bidders, bid 99% of your value
- Intuition (again): more competing bidders means that there is a greater chance for higher bids: so you sacrifice some payoff ($v_i - b_i$) to increase probability of winning in a more “competitive” situation
- Analysis is more involved for more general distributions of values
  - each specific form requires its own analysis, but general idea is similar to the uniform distribution case

Other Properties: First-Price Auction

- Bidders generally shade bids (as we've seen)
  - Does seller lose revenue compared to second-price auction? What do you think?
- *<u>If</u> bidders all use same (increasing) strategy,* item goes to bidder with highest value (this maximizes social welfare, like second-price)
  - but note that our symmetric equilibrium need not be only equilibrium!
- Outcome is similar to Dutch auction (descending auction)
  - lower prices until one bidder accepts the announced price
  - until price drops below *k*'s value, *k* should not accept it
    - jump in too soon: will pay more than necessary (equivalent to bid shading)
    - jump in too late: you lose when you might have won
  - first bidder jumping in pays the price she jumped in at (1^{st} price)
  - games are in fact "strategically equivalent"; seller gets same price
    - with some "slop" due to bid decrement in Dutch auction

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 2nd 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)

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

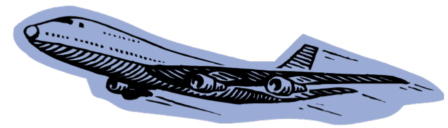

Flight2 (10AM, get airmiles, direct)
Value: $950

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

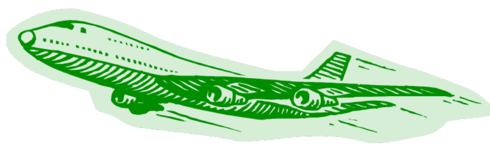

Flight1

Hotel Room

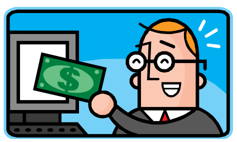

Bidder doesn't want flight without hotel room,
or hotel without flight; but together, value is $1250

- 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

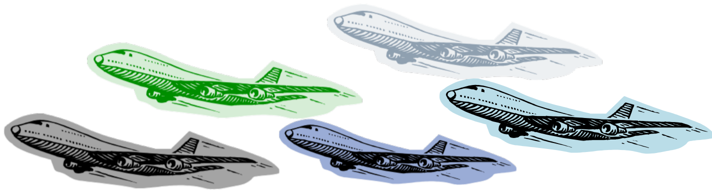

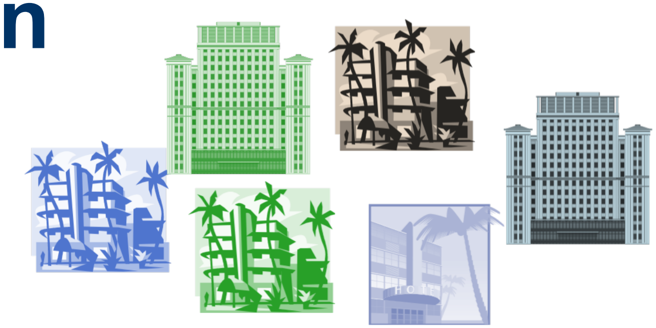

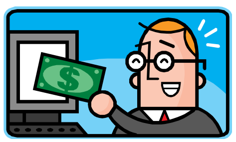

Bidder expresses value for ***combinations*** of items, e.g.:

- 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 1^{st}-price (pay what you bid) or Vickery-Clarke-Groves (VCG, a generalization of 2^{nd}-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)

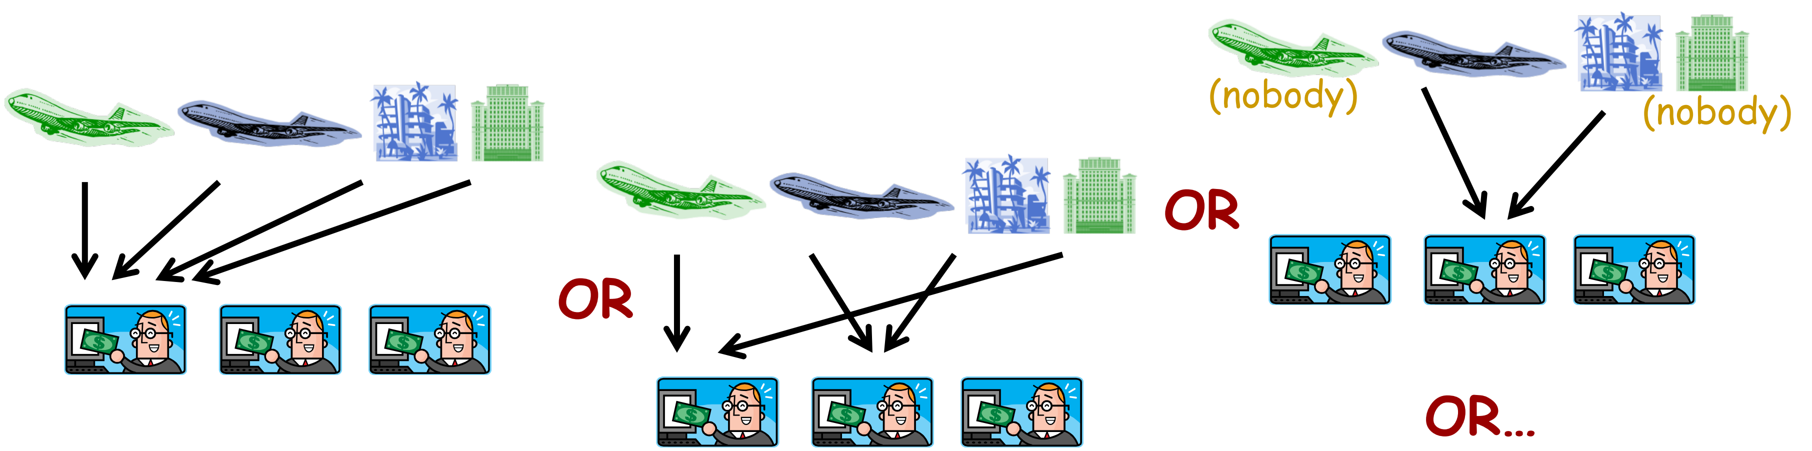

How Many Assignments?

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	$1.2\text{x } 10^{40}$
100	1000	$1.1 \text{ x } 10^{300}$
200	10,000	$1.0 \text{ x } 10^{800}$

- 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

Combinatorial Auctions in Practice

- Despite this, combinatorial auctions are now routinely solved involving hundreds of bidders and hundreds of thousands of items
- Algorithms exploit the fact that most bidders are interested in very few combinations or that these have "structure" (see next slides)
- The valuation functions of the bidders may be very restricted

- Some of these structures do not make the problem easier from a formal perspective: many interesting algorithmic issues remain
- But some of them do make things easier theoretically… and often in practice.

Structure in Combinatorial Auctions

- For example:
- *Single-minded bidders*: every bidder bids on only one subset
  - this is still an NP-hard problem: still 2^m subset assignments
  - even this can't be approximated efficiently according to complexity theory assumptions even ignoring game theory aspects
  - but you can get within a factor of $\sqrt{m}$ of optimizing social welfare using a truthful "greedy allocation mechanism" for single minded bidders whereas there is evidence that there may not be such a deterministic truthful mechanism for multi-minded bidders.
- *Ordered sets* of items: single-minded bidders only bid on "consecutive" subsets of items
  - e.g., only interested in adjacent facilities, plots of land (on a line), consecutive bands of wireless spectrum, adjacent time slots in schedule
  - easier computational problem: solvable in polynomial time
  - Ordered sets can be generalized to items that are arranged in a graph (e.g., an interval or a tree graph) with a special structure.

Bidding Languages

- Bidders may need to specify value for large number of combinations
  - If n items, there are 2^n packages that bidder must consider
- But there are usually a lot of packages bidder doesn't care about, and a lot of structure in values
  - e.g., suppose items are strict substitutes: here are the 10 flights I care about, here's how much each is worth, give me only one
  - e.g., independent/additive values: here are the 10 items are I care about, here's how much each is worth, regardless of how many others I get
  - e.g., k-of complements: I will pay *\$d* for any 10 items from this set of 100
- *Bidding languages* let bidders express values concisely, naturally
  - algorithms can often exploit these languages in practice
  - still not in general theoretically tractable, but often practically so…

Other Issues

- A variety of interesting strategic issues:
  - *envy-free*: find an allocation (and prices to charge) so that no bidder would prefer the bundle of goods allocated to a different bidder
  - *stability*: find payments so no group of bidders could offer to pay more for a set of goods allocated to others, divide it up, and be better off
- *What about pricing:* if people pay what they bid (1st-price), they will obviously hedge their bids; is there an analog of a 2nd-price auction in the combinatorial setting? What would the 2nd highest bid be?

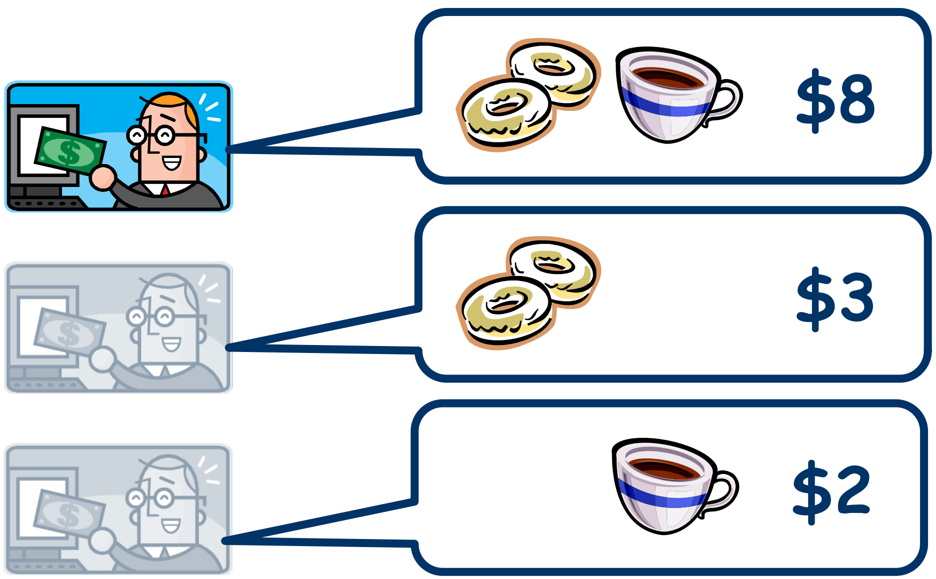

Winner gets bundle of both items. But no other bidder offered a bid on the same bundle: so what price should we charge using "2nd price"?

Another combinatorial auction

- Agents can have very different valuations:
  - *What about pricing:* if people pay what they bid (1st-price), they will obviously hedge their bids; is there an analog of a 2nd-price auction in the combinatorial setting?

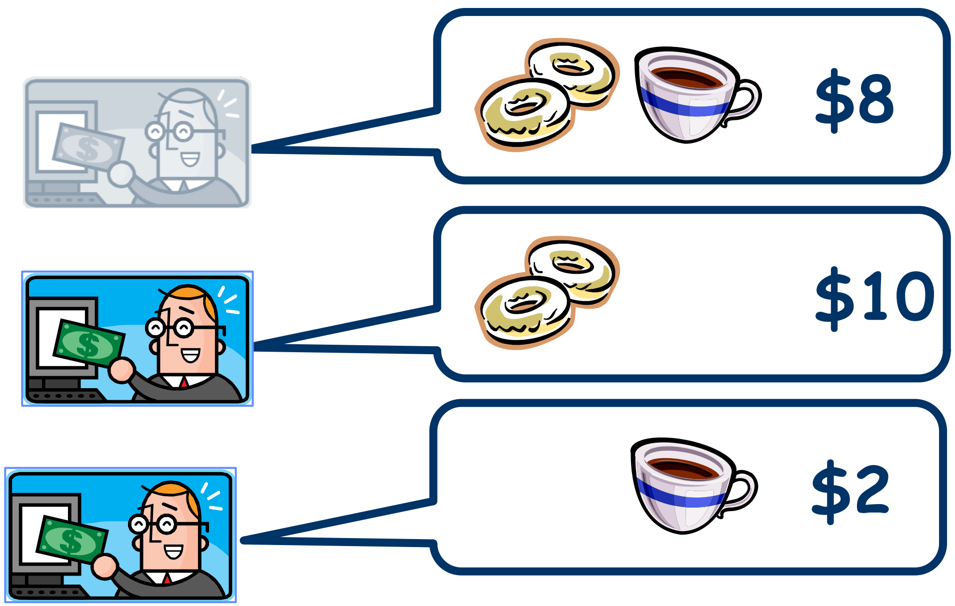

To maximize social welfare we have two winners if agents bid truthfully: second agent gets donuts while third agent gets coffee. So what price should they pay?

The VCG Mechanism

- There is a generalization of the 2nd-price (or Vickrey) auction to CAs
  - The *Vickrey-Clarke-Groves (VCG)* mechanism
- Lets assume (and this is a big assumption) that we know how to allocate items optimally. Now how to charge?
- Roughly speaking, you charge someone based on the “externality” they inflicted on other players by their presence
  - e.g., bidder *X* gets bundle *B*, this (potentially) prevented other bidders *Y, Z, …* from getting (some parts of) bundle *B*
  - figure out social welfare that *Y, Z, …* got in the actual auction
  - then pretend *X* didn’t exist and figure *SW* that *Y, Z, … would have attained* if X hadn’t bid/didn’t exist (can’t be any worse, could be higher)
  - charge X the difference of the two: what he cost *Y, Z,…* by his presence
- In slide 13 winner pays $5 for coffee and donuts. Slide 14 outcome??
- Notice that 2nd-price auction is a special case of VCG with one item
- We’ll see VCG in more detail in Ch.15 (advertising auctions)

Simple auctions

- “Simplicity is the ultimate sophistication” (Leonardo da Vinci

- In an auction, if possible we would like
  - Computational simplicity (eg computationally fast)
  - Strategic simplicity : agents should be able to easily and quickly know what to do
  - Conceptual simplicity : the rules of the auction should be easy to understand and perceived as being fair

Do combinatorial auctions using VCG satisfy these?

Conceptually simple auctions

- If you were a buyer in a combinatorial auction, would you understand what to do?
- Even knowing about (say) the VCG mechanism (but not understanding exactly how the mechanism chooses an optimal allocation) would you feel confidant in bidding?
- Sandholm and Gilpin [2004] argue that mechanisms such as VCG often fail in the real world for various reasons:
  - Buyers may be unwilling to reveal their true values
  - A buyer may be unwilling to participate in an auction where the rules are complex, not fully understood or unintuitive
  - The computational difficulty makes an optimal allocation impossible and VCG in general requires an optimal allocation

Pricing to achieve objectives

- There can be many objectives for a mechanism as we have already seen; social welfare, seller revenue, market clearing, fairness. A number of mechanisms just use item pricing to achieve desired objectives.
- These prices may be offered all at once or dynamically in some manner and then agents choose an allocation based on their valuation function and the item prices.