Lecture 3: Incentive Compatibility, Revenue of the Vickrey auction, Sponsored Search

- (Last lecture) Lemma: An auction satisfies incentive compatibility if and only if it is threshold based.
- We proved necessity (see last lecture's notes). Formal proof of easy, useful sufficiency direction, which we hand waved earlier:
    - Recall definition of incentive compatibility: For all bidders i, for all bids b_{-i} of other bidders, setting $b_i = v_i$ maximizes bidder utility.
    - Fix bidder i, bids b_{-i} of other bidders. This sets up a threshold $t(b_{-i})$. Two cases.
        - If v_i less than $t(b_{-i})$, then bidding value yields utility of zero. Bidding anything lower than $t(b_{-i})$ also has utility zero, bidding higher than $t(b_{-i})$ has negative utility. So lying is not profitable.
        - If v_i larger than $t(b_{-i})$, then bidding value yields utility of $v_i - t(b_{-i})$. Bidding anything larger than $t(b_{-i})$ also has the same utility, bidding lower than $t(b_{-i})$ has a utility of zero. So lying is not profitable.
- Optimizing for Economic Efficiency in auctions requires **optimization over 'private' values**
    - Both first price and second price auctions allocate item to highest bidder.
    - For the second price auction, buyers have no incentive to lie about value. So item goes to buyer with highest *value*. Consequently, efficiency is maximized.
    - In contrast, in first price auction, hard to claim that item goes to bidder with highest *value.* Example: Two buyers A, B. Value of A for item is 100 and value of B for item is 70. Buyer A _thinks_ that B's value for the item is no larger than 50, and so considers 55 a bid that guarantees a win, and bids this. Buyer B , an inherently truthful bidder bids 70 and wins the item. The allocation is thus inefficient.
- **Collusion:** Though the second price auction is incentive compatible, it is not collusion resistant. Example: <insert your own: How would two bidders get together to lower the price that the winner pays?>
    - In practice, how would you collude in eBay?
    - Fixed price is collusion-resistant.

Money as a side-effect

- Assume bids/values are drawn IID from Uniform[0,1]
- Density function: f(x) = 1; Cumulative Distribution Function: F(x) = x
- What is the expected value of the second highest number?
- $\int_0^1 n(n-1)x\, F(x)^{n-2} f(x)(1-F(x))\, dx$

- expectation of a random variable = $\int_0^\infty y\, f(y)\, dy$
- n ways to pick identity of highest bid, n-1 ways to pick identity of second highest bid. Highest bid takes on value larger than ‘x’, second highest bid takes on value exactly ‘x’, remaining numbers take on value less than ‘x’

- $= \int_0^1 n\,(n-1) x\,(x)^{n-2}\,(1-x)\,dx \;=\; (n-1)/(n+1)$
- Can we do better?
- Clearly, we cannot extract more than expected value of the highest number: $n/(n+1)$
- Therefore: $(n-1)/(n+1)$ approximation

Sponsored Search Market Basics

- Why advertising?
  - Merchant pushes information to potential consumer with the hope of future transaction
  - First advertising: Wall paintings, papyrus
- Interestingly, advertising did not change much till 20th century, or even the 21st!
  - John Wannamaker: “Half the money I spend on advertising is wasted; the trouble is I don't know which half.” E.g. Billboards
- Advantages of Internet advertising
  - Engagement can be measured: Views-->Clicks-->Conversions
    - Contrast with billboards, Nielsen+TV
  - Relevance can be tuned based on ‘context’ and user
    - context=website content / user query
    - In practice, context far more important than user features
    - In practice, search ad clicks much more valuable than display ads clicks
- Pricing:
  - Queries have a long tail http://www.seomoz.org/img/upload/search-demand-curve%281%29.gif
  - Variation in search cpcs
    - http://www.fetch123.com/SEM/the-most-expensive-keywords-in-google
  - Hard to identify value of each query to each advertiser
  - Need auction to discover value, same premise as eBay
  - Differences from eBay?
    - user happiness is important (next lecture), repeated setting, every bidder in many auctions, automated matching of queries to advertisers (soon)
  - Ok, so let us model each advertiser as having a value per impression for each view

Running keyword auctions

- Many queries convey the same intent
  - spellings
  - re-ordering of words
  - synonyms
- Advertisers bid on keywords
- Keywords are 'matched' to queries
  - Advertisers can control the looseness of the match
    - precision v/s recall
    - can result in humor
- For a given query, first identify all eligible advertisers and their bids per impression
- Run efficient, multi-item auction with as many items as positions on sale.
- Notice that this is a k+1st price auction. (Must have bid independent threshold, must allocate to top k bidders.)

End of lecture.

Lecture 4: Sponsored Search Auctions: k-item auction,

prominence, pay-per-click

eBay: The rise of fixed price (Drop in fraction of revenue from auctions for eBay)
Courtesy: Sarah Boisseree

Sarah Boisseree
MS&E 233

Year	Auction	Fixed
2006	64%	36%
2007	60%	40%
2008	55%	45%
2009	47%	53%
2010	40%	60%
2011	37%	63%
2012	34%	66%

*All data comes from the annual reports (10-k) for eBay found on the SEC website
**Percentage of Gross Merchandise Volume (GMV) used in comparisons

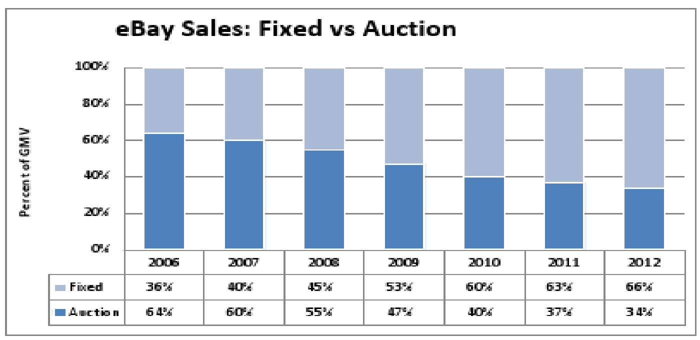

Likely cause: Sellers know buyer value better and switch to auctions.
Other possible causes:

- Buyers are less patient and most auctions don't result in sales.
- Sellers who were merchants on Amazon joined eBay and they were accustomed to fixed price sales.
- ...

Identifying causality is hard :)
See for example: Simpson's Paradox

Back to Sponsored Search...

First-Cut: Efficient auction for k identical items

- For a given query, first identify all eligible advertisers and their bids per impression
- Run efficient, multi-item auction with as many items as positions on sale.
- Notice that this is a k+1st price auction. (Must have bid independent threshold, must allocate to top k bidders.)
    - Notice that the auction of running a second price auction for each item in succession, removing the winner at each step is not incentive compatible
    - E.g. items positions, and 5 bidders with values 10, 9, 8, 7, 6. Highest bidder pays 9, though it would still have won if it dropped its bid to 7.5, in which case it would only pay 7.

Second Cut: Deal with varying prominence

- Problem: different positions have different amounts of prominence to search users
    - see ads for query ‘camera’ on Bing

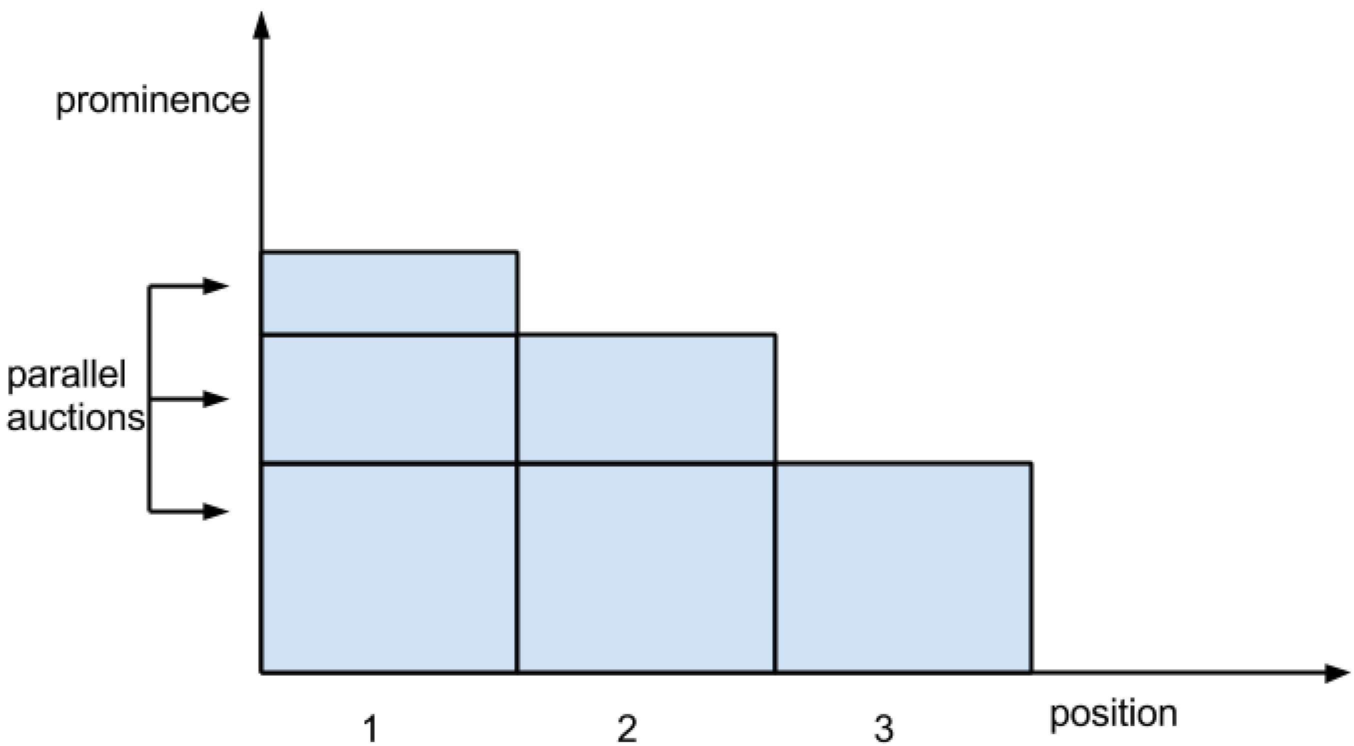

- Thus, from an advertiser’s perspective, these auctions sell ‘views’ and not impressions
    - So the bids are for a view.
- As we will see, this is mathematically equivalent to many simultaneous k+1st price auctions; of course what runs is just one auction
- Define position prominence q_j as the probability of user noticing position j.
    - Higher slots are more prominent: $q_j > q_{j+1}$

- From the picture, there is an auction of j identical items with weight $q_j - q_{j+1}$ for every $j \in 1 \ldots \#positions$
- Putting auctions together: Rank bidders in decreasing sequence of bid per *view* (b_j). Allocate according to rank: i th highest bidder pays: $\sum_{j>i} (q_{j-1} - q_j)\, b_j$; assume that $q_j = 0$ for positions which don't exist

Pay-per click advertising

- Naive sales pitch: "You only have to pay if you get a click"
- Practical truth: How can advertiser estimate value per view?
  - Easier for search engine to predict click-through-rate than for the advertiser to do so.
    - Probability c_i that user viewing the ad will click on it. Notice that we did not say 'probability that a shown ad will get a click'; this difference is necessary because of the prominence issue discussed earlier.
    - This is a function of the query and the advertiser
  - Possible for advertiser to identify value for a click. Impossible for search engine because it does not know what clicks lead to sales.
  - This results in concise communication from the advertiser to the search engine. Imagine communicating a bid per query type!
  - Results in lower risk for the advertiser: Suppose that advertiser bids 0.5$ per impression/view, knowing that it has a value of 5$ for each click, and presuming that it has a click-through-rate of 0.1. Suppose its actual click-through rate is lower, it risks losing money.
- Therefore, inventory is impressions, but we should sell clicks
  - Given a bid per click, we can multiply it by the predicted click-through rate, and derive a bid per view.
  - Then, we run the per-view auction from the previous section.
  - When charging, divide per-view price by *predicted* click-through-rate (c_i)
- Example: Consider a second price auction for a single position: Rank bidders by $b_i c_i$, Give item to highest ranked bidder; charge $b_2 c_2 / c_1$ per click. Notice that this price is less than b_1 for the click. So there is no risk for the advertiser.
- Notice, poor click-through rate implies costlier clicks. Therefore, incentive to improve click-through rate

- In summary: For a given query, first identify all eligible advertisers and their bids per click, and their predicted click-through rates, **compute** bids per-view, then run auction as above. When charging, divide per-view price by *predicted* click-through-rate (c_i)

User Happiness

- Search Engine Competition
    - Users want good search, will tend to use one search engine
    - Advertisers need eyeballs, can advertise on multiple search engines
    - Consequence: Search engines compete for users, and not advertisers (unless budget constrained)
    - Ads on search need to be non-annoying and useful. Important to not blind users to ads.
        - Contrast with pop-ups, spam email, banner ads

End of lecture

HW1:

- **Problem 1:** Recall in the lecture that we showed that the Vickrey auction is incentive-compatible and efficient. Show that it is the *unique* single-item auction with both these properties.
- **Problem 2:** Your grandmother passed away and left you an antique rifle. You have no value for the rifle. A local collector is willing to pay a 250$ for it. eBay has an *insertion fee* of 2$ per listing, and takes a 10% cut on the eventual sales price. You are guaranteed to get one bidder on eBay (the bonus question makes the problem harder and more meaningful), whose willingness to pay is drawn from uniform distribution with range [0,1000]. What is your optimal strategy, and what is your profit from it? We are seeking an optimal strategy of the form of picking a reserve price on eBay, or selling directly to the collector. **Bonus**: Suppose instead that you have 5 potential buyers with valuations distributed IID in the range uniform[0,1000]. What is your optimal strategy then? Note that we insist on an incentive compatible auction, i.e., you must run a second price auction with a reserve.
- **Problem 3:** Either prove that these single-item auctions are incentive compatible or supply an example that demonstrates otherwise:
    - Given a weight w_i for each buyer, give item to bidder with highest product $w_i b_i$, charge $w_2 b_2 / w_1$, where b_2 is the second highest bid.
    - Give item to highest bidder, charge second highest bid if second highest bid is at least half of the highest bid. Charge half the highest bid otherwise.
    - Fix a sequence on the bidders. Inspect the first three bids in this sequence. Run a second price auction among the other bidders with a reserve price = max of the first three bids.
- **Problem 4:** We discussed a model of eBay auctions in class. Identify a few different ways that this model was not realistic based on evidence from this paper. Mention two insightful statistics about eBay from this paper.
- **Problem 5:** Write a short paragraph comparing eBay and Amazon resolution processes.

End of HW1

- **Problem 6:** What is revenue of the single-item Vickrey auction with n bidders each with value drawn IID from the exponential distribution with mean 1. What is the efficiency? What fraction of efficiency is the revenue? How does this fraction compare with the case where the value distributions are IID Uniform[0,1]? Intuitively, why does the exponential distribution behave differently from the uniform distribution?

Solutions to HW1

Problem 1.

Probem 2. Suppose exactly one user bids for our item. The true value of this bidder is expected to be 500. In an incentive compatible auction she bids her true value, 500. Thus our revenue is 91% $91\% \times 500 = 455$. We have to pay 2$ to post the item. Thus, the total profit is 453. Since 453>250, we should use eBay to sell the item.

Bonus: If there are 5 bidders, then the expected true value of the second highest bid, as dervied in the class, is $\frac{5-1}{5+1} \times 1000 = 666.7\$$. Thus, our revenue is $91\% \times 666.7 = 606.67\$$. Subtracting the initial posting price, the total benefit will be: $606.67 - 2 = 604.67\$$. So again, we should use the eBay. Not surprisingly, as the number of bidders increases, the expected revenue increases.

The above conclusions are based on the assumption that we base our decisions only on the expected value of payoffs and do not care about the uncertainty (risk) in the final profit. In other words, we consider expected value of the profit and not the variance of the profit. In these situations, the decision maker is called *risk neutral*, meaning that it is neutral to risk and does not account for it.

Solution I expected:
Suppose you commit to selling on eBay: Then pick a reserve price t that optimizes:
$0.9 \, * \, t\,[1-F(t)] \, + \, 250 * F(t) = \, 0.9 \, * t\,(1-t/1000) \, + 250 \, * t/1000$
The first term is our revenue, the second term is the money we would make from giving the item

to the collector if it does not sell
Differentiate wrt t and set to zero.
Plug optimal t back in to compute revenue. If this revenue minus the listing price (2$) is greater than 250$ (directly give item to collector), then use eBay. Else give item to collector

Problem 3.
Part (a). Incentive compatible. Consider the following equivalent auction.
User i is offered a TILI price equal to $w_{-i}^{max} b_{-i}^{max} / b_i$, where b_{-i}^{max} is the maximum bid offered by all users other than the i^{th} user, and w_{-i}^{max} is the weight of the the user that has this bid. First, we note that this auction obeys the conditions mentioned in the lemma in page 3 of lecture 2, and thus is incentive compatible. Second, we note that only the person with the highest $w_1 b_1$ wins the item (because for all other bidders the price offered to them will be larger than their true bid) and this person pays $w_2 b_2 / w_1$. So, the above item is equivalent to the auction suggested in this part.

Part (b). Not incentive compatible. The condition mentioned in the lemma in page 3 of lecture note 2, is a necessary condition. The auction provided in this part does not satisfy the condition mentioned in the lemma.

Part (c). Incentive compatible. A bidder would think as follows: If I am unlucky enough to be chosen among the first three individuals, then no matter what I bid, I will have gain nothing. So, telling the truth does not hurt me in this case. On the other hand, if I happen to be placed in the set of bidders attending the second-price auction, then given that second-price auctions are truthful auctions, it is to my advantage to tell my true value. That is, it never hurts me to tell the truth and sometime it indeed benefits me. Thus, overall the auction is truthful.

Problem 6:
Suppose, we have n independent samples derived from a probability density function (PDF) $f(t)$ with corresponding cumulative distribution function (CDF) $F(t)$. Then using an argument similar to what was given in the class for the PDF of the second highest number, the PDF of the first highest number is given by:
$g_1(x) = nF(x)^{n-1} f(x)$

Also, as discussed in the class, PDF of the second highest number is given by:

$g_2(x) = n(n-1)F(x)^{n-2}(1 - F(x))f(x)$

Thus, in the case when there are n bids coming from the distribution $f(x)$, expected value of the highest bid, and expected value of the second highest bid are, respectively, given by:

$E\, b_1 = \int x g_1(x) dx$

$E\, b_2 = \int x g_2(x) dx$

In the single-item auction, expected total welfare (efficiency) is equal $E\, b_1$ and expected revenue is equal to $E\, b_2$.

In the case of Exponential distribution with means 1, $f(x) = e^{-x}$ and $F(x) = 1 - e^{-x}$. Thus, we have:

$$E\, b_1 = \int x g_1(x) dx = \int_0^\infty x\, n\, (1 - e^{-x})^{n-1} e^{-x}\, dx =$$

$$E\, b_2 = \int x g_2(x) dx = \int x\, n(n-1)\, (1 - e^{-x})^{n-2} e^{-2x}\, dx =$$

And the fraction of seller revenue over total welfare is:

fraction$= \frac{E\, b_2}{E\, b_1} =$

In the case of uniform distribution between 0 and 1, we have:

$$E\, b_1 = \int x g_1(x) dx = \int_0^1 x\, n\, (x)^{n-1}\, dx = \frac{n}{n+1}$$

$$E\, b_2 = \int x g_2(x) dx = \int_0^1 x\, n(n-1)\, (x)^{n-2} (1-x)\, dx = \frac{n-1}{n+1}$$

And the fraction of seller revenue over total welfare is:

fraction$= \frac{E\, b_2}{E\, b_1} = \frac{n-1}{n+1}$

Lecture 5: The Incorrectly Generalized Generalized Second Price Auction, Reserve Prices

Instructor: Mukund Sundararajan

Reserve Prices

- Reserve prices are useful in a variety of situations:
  - seller values item
  - seller has a 'offline' buyer lined up
  - item costs money to produce
  - there are too few bidders
  - top bid is way higher than second highest one

- We focus for the moment on the last two bullets
- Simple exercise: One bidder, value drawn from distribution F. We would like to maximize expected revenue.
  - Recall incentive compatible auction must set bid-independent threshold
  - So all we can do here is to set a price!
  - To optimize revenue, maximize $p\,(1 - F(p))$
  - In general, higher prices will yield low probability of sale, but better revenue on sale. We need to find the optimal trade-off between revenue on sale and probability of sale.
  - Example: For uniform distribution U[0,1], optimize $p\,(1-p))$. Differentiate, equate to zero: $2p = 1$, so optimal price is a half, probability of sale is a half, revenue is 1/4.
- How should we set a reserve within a single item auction?
  - Consider a single item auction. Suppose we employ a second price auction with a reserve. That is sale occurs if and only if the highest bidder's bid is higher than the reserve, in which case it pays the maximum of the second highest bid and the reserve.
  - Notice that reserves are not always binding, second price can raise revenue too
  - Consider the model where bidder's bids are drawn independently and identically from a specific distribution. How do you expect the optimal reserve to vary as a function of the number of bidders?
  - **First Surprise:** The optimal reserve price does not vary with the number of bidders!
  - Recall that the for the uniform distribution U[0,1], with one bidder, the optimal reserve price was 1/4. Let us see what the optimal reserve price is with 2 bidders.
    - $2 * r * F(r)\,[1 - F(r)] + 2 * \int_r^1 t * f(t) * [1 - F(t)]\,dt$
    - $= 2 * r * r\,[1 - r] + 2 * \int_r^1 t[1 - t]\,dt$
    - $= 2 * r^2 * (1 - r) + 2\,[1/2 - 1/3 - r^2/2 + r^3/3]$

        - Now differentiate and set to zero: $-4r^2 + 2*r = 0$
        - Optimal reserve price is still 1/2!
    - In general, it is not even clear that a second price auction with a reserve is the best way to maximize revenue.
    - **Second Surprise:** The revenue maximizing single-item auction with values distributed IID is a second price auction with a reserve!
    - The form of the revenue optimal auction was identified in a seminal paper by Roger Myerson, who won a Nobel Prize in economics in 2007.

    - How much did we gain by setting a reserve with bids drawn IID from U[0,1]?
        - When there was one bidder, revenue goes from 0 to 1/4.
        - When there are two bidders, revenue goes from 1/3 to 1/3 + 1/12.
    - In general, as intuition would suggest, the gain from reserve prices decreases as the number of bidders goes up.
    - In practice setting good reserves only possible if general range of bids is known. Probably better to err on the lower side, because we make revenue of second pricing anyway. Also, it is very important to encourage bidding. This is useful for efficiency and revenue.

- **On eBay, there are two forms or reserves:**
    - http://ebay.com/help/sell/starting_price.html
    - **Start price**
        - Used to shorten bidding process or as a reserve
    - **Secret reserve**
        - Empirically observed that low starting price encourages bidding. So, secret reserve prevents sale below certain price.

The Generalized Second Price Search Auction

- Initially, Yahoo used a first-price auction, which resulted in volatile bidder behavior
    - see Page 3 of this paper
- Google switched to a second-price auction
    - second price auction for each position, remove winner repeat with next position
    - As we discussed in last lecture, this is not incentive compatible
- Surprisingly, Bing, Yahoo, and Google all still use this auction
- One explanation for this is that the second price auction is stable enough
    - Recall payments of the incentive compatible auction: ith highest bidder pays:

      $$\sum_{j>i} (q_{j-1} - q_j)\, b_j$$

    - In contrast, in a second price auction, the bidder pays: $\sum_{j>i} (q_{j-1} - q_j)\, b_{i+1}$

- Suppose that the position weights and bids both fall off geometrically. E.g., the bids of the bidders are 16, 8, 4, 2, 1. And the position weights are 1, 0.5, 0.25, 0.125, 0.0625. A first price auction would charge the highest bidder 16, a second price auction would charge the bidder 8, the incentive compabtible auction will charge the bidder $\approx$ 5.25. Notice that the reduction in price from the first-price auction to the second price auction goes most of the way towards incentive compatible price.

End of lecture

- Myerson's Lemma (Challenging)
  - Define virtual-value function of a bidder $v - (1-F(v))/f(v)$.
  - *To maximize revenue, maximize virtual value served. Charge threshold price*.
  - Derivation:
    - Incentive compatibility implies bid independent thresholds

        - Expected revenue: $\sum_i E_{b_{-i}}\left[E_{b_i}[t(b_{-i})\,(1 - F(t(b_{-i})))|\, b_{-i}]\right]$
            - Notice use of distributional independence
        - Cannot optimize $t(b_{-i})$ separately for each bidder due to supply constraint.
        - Differentiating and rewriting:

$$\sum_i \int_{b_{-i}} \int_{t(b_{-i})}^{\infty} [v - (1 - F(v))/f(v)\,]\, f(v)\, dv]\, f(v_{-i})\, dv_{-i}$$

        - Give item to bidder with highest virtual value, unless all virtual values are negative. Notice this implies threshold for each bidder in virtual value space.
        - Virtual values are monotone in value, implying thresholds in value space.
        - Sort by value, with reserve

    - Exercise: Check that this applies to bidders with values distributed independently, but non-identically

- **Start price versus secret reserve:** Low starting price encourages bidding, secret reserve prevents sale below certain price: http://ebay.com/help/sell/starting_price.html

What are we missing from our model?

- **Fixed price v/s Auction**: Features that trade-off value discovery v/s patience.
- **Best offer**: Feature of fixed prices sales that allows buyers to negotiate
    - Fraction of fixed price sales that use this feature
- **Buy it now price**: Brings fixed price feature to auctions. Allows impatient buyers to transact quickly.
    - Fraction of auctions that use this feature
- **Shipping fees**: Seller has responsibility to ship, can vary shipping fee

- Trust issues
    - Buyers trusting seller
        - Seller reputation (will be covered later by Ashish)
        - Payments (PayPal accounts for 40% of eBay revenue!)
            - easier for sellers to create account
            - sellers don't see confidential information
- Item search
- eBay's cut
    - http://pages.ebay.com/help/sell/fees.html

- Points to ponder:
    - How do smart phones affect eBay?

- eBay versus Amazon
- In the long term, auctions versus pricing on eBay?

- Strawman: Run second price auction.
  - Second price?
  - Run next price auction
- Generalized Second Price auction is not truthful
  - Two identical positions, three bidders
  - Bidders will tend to underbid
  - 3rd price is truthful
- What if positions are not identical?
- TODO(mukunds):
  Normalizers?, diagram....
- What if you can show one big ad or two small ones?
- VCG is an efficient, incentive compatible auction for all occasions
  - Informally
    - accept bids
    - compute efficient allocation
      - charge each bidder its externality
      - externality = total value to other bidders *without* this bidder in system - total value to other bidders *with* this bidder in system
      - Sounds sanctimonious, but remarkably objective
      - Revisit 2nd price auction
      - Proof: (Truthful bidding cannot hurt a bidder.)
        - Space of outcomes O.
        - Every bidder i has a valuation function over outcomes $o \in O$: $v_i(o)$
        - Utility from truthtelling: $v_i(o^*) + \sum_{j \neq i} b_j(o^*) - V_{-i}$
        - Utility from lying: $v_i(o') + \sum_{j \neq i} b_j(o') - V_{-i}$
        - By efficiency of VCG:

          $v_i(o') + \sum_{j \neq i} b_j(o') \leq v_i(o^*) + \sum_{j \neq i} b_j(o^*)$

- This is status-quo: Bing, Google, Yahoo
- Why not change to VCG?

Nash Equilibrium of GSP?

Budgets

- Budget Constraints
  - Direct response v/s Brand Advertising
  - Quarterly budgets versus real budgets

Simple algorithms work well

- Advertisers: $a \in A$, Queries/Keywords $k \in K$, bid for every pair $b_{a,k}$
- Keywords arrive 'online'
- First price auction
- Greedy is 2-approximate for revenue
- Define OPT, greedy etc.
- Notice that optimal algorithm is complicated, but proof is simple
- Proof:
  - Consider all keywords K_{low} allocated by OPT that were allocated by Greedy to a lower bid. Let revenue from these keywords be R_{low}. Let all other revenue of OPT be R_{high} Notice, $R_{high} \leq R_{greedy}$
  - Suppose that OPT served a total revenue of R' due to these keywords.
  - Greedy did this because higher bid had exhausted budget. Therefore, $R_{low} \leq R_{greedy}$

Nash Equilibrium

- s

Trust Issues

- Consider a transaction in person with cash. Buyer can inspect item, buy it at the same quality. Seller gets cash. Only trust assumption is by seller on institution that issues cash.
- Consider online transaction with a credit card.

    - Was item delivered on time?
    - Was it as listed?
    - Can buyer trust seller with credit card information?
- First two solved by: http://pages.ebay.com/help/policies/buyer-protection.html
-