CS269I: Incentives in Computer Science
Lecture #13: Introduction to Auctions*

Tim Roughgarden†

November 7, 2016

1 Preamble

Why should a computer scientist care about auctions? Until recently (and maybe even still), the first picture that comes to mind is probably an art auction in an auction house. (Or maybe an estate sale, or maybe eBay.) Auctions have been around a long time, and economists have been thinking hard about them for over 50 years.

Computer scientists started to care about auctions in the mid- to late-90s. eBay started in 1995, and initially all transactions were done using auctions.[1] Auctions have been still more important in this century, and all of the major search engines and social networking sites derive most of their revenue from real-time auctions used to sell online advertising. Thus auctions are one of the major drivers of the modern Internet economy. These days, many Internet companies are regularly looking for students trained in both computer science and in auction theory.

2 Single-Item Auctions

Let's start with the canonical auction design problem: *single-item auctions*. Consider a seller with a single good, such as a slightly antiquated smartphone. This is the setup in a typical eBay auction, for example. There is some number n of (strategic!) bidders who are potentially interested in buying the item.

We want to reason about bidder behavior in various auction formats. To do this, we need a model of what a bidder wants. The first key assumption is that each bidder i has a *valuation* v_i — its maximum willingness-to-pay for the item being sold. Thus bidder i wants

†Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

[1]Auctions are now a relatively small percentage of eBay's transactions, but they remain useful and profitable, especially in niche markets [2].

to acquire the item as cheaply as possible, provided the selling price is at most v_i. Another important assumption is that this valuation is *private*, meaning it is unknown to the seller and to the other bidders.

Our bidder utility model, called the *quasilinear utility model*, is as follows. If a bidder loses an auction, her utility is 0. If the bidder wins at a price p, then her utility is $v_i - p$. This is perhaps the simplest natural utility model, and it is the one we will focus on in this course.

3 Sealed-Bid Auctions

For now, we'll focus on a particularly simple class of auction formats: *sealed-bid auctions*. Here's what happens:

(1) Each bidder i privately communicates a bid b_i to the auctioneer — in a sealed envelope, if you like.

(2) The auctioneer decides who gets the good (if anyone).

(3) The auctioneer decides on a selling price.

There is an obvious way to implement step (2) — give the good to the highest bidder. Today, this will be the only selection rule that we study.[2]

There are multiple reasonable ways to implement step (3), and the implementation significantly affects bidder behavior. For example, suppose we try to be altruistic and charge the winning bidder nothing. This idea backfires badly, with the auction devolving into a game of "who can name the highest number?"

4 First-Price Auctions

A much more reasonable choice is to ask the winning bidder to pay her bid. This is called a *first-price auction*, and such auctions are common in practice.

First-price auctions are surprisingly hard to reason about. First, as a participant, it's hard to figure out how to bid. Second, as a seller or auction designer, it's hard to predict what will happen. To drive this point home, imagine participating in the following experiment. There's an item being sold via a first-price auction. Your valuation (in dollars) is the number of your birth month plus the day of your birth. Thus, your valuation is somewhere between 2 (for January 1st) and 43 (for December 31st). Suppose there is exactly one other bidder (drawn at random from the world) whose valuation is determined in the same way. What bid would you submit to maximize your (quasilinear) utility? Would your answer change if you knew there were two other bidders in the auction, rather than one?

[2]We'll see later why other winner selection rules are sometimes preferable.

5 Second-Price Auctions

Let's now focus on a different auction format, which is also common in practice, and is much easier to reason about. To motivate it, think about what happens when you win an eBay auction. If you bid \$100 and win, do you necessarily pay \$100? Not necessarily — eBay uses a "proxy bidder" that increases your bid on your behalf until your maximum bid is reached, or until you are the highest bidder (whichever comes first). For example, if the highest other bid is only \$90, then you will only pay \$90 (plus a small increment), rather than your maximum bid \$100. The upshot is: *if you win an eBay auction, the sale price equals the highest other bid (the second highest overall), plus a small increment.*

A *second-price* or *Vickrey* auction is basically a sealed-bid version of eBay, where the highest bidder wins and pays a price equal to the second-highest bid.[3] The most important property of a second-price auction is the following, as proved in the paper that originally founded auction theory [4].

Claim 5.1 *In a second-price auction, every bidder has a dominant strategy: set her bid b_i equal to her private valuation v_i.*[4] *That is, this strategy maximizes the utility of bidder i, no matter what the other bidders do.*

This claim implies that second-price auctions are particularly easy to participate in — you don't need to reason about the other bidders in any way (how many there are, what their valuations are, whether or not they bid truthfully, etc.) to figure out how you should bid. Note this is completely different from a first-price auction. You should never bid your valuation in a first-price auction (that would guarantee zero utility), and the ideal amount to underbid depends on the bids of the other players.

There is strong intuition about why the claim should be true. Suppose you were playing a first-price auction, and you telepathically knew all of the other bids. What would you bid? Either 0 (if the highest other bid is above your valuation), or a tiny amount ϵ more than the highest other bid, to win at the cheapest price possible (otherwise). You're not actually telepathic and don't know the other bids, but the auctioneer does! And in a second-price auction, the auctioneer effectively underbids optimally on your behalf.

Proof of Claim 5.1: Fix an arbitrary player i with valuation v_i, and the bids $\mathbf{b}_{-i}$ of the other players. (Here $\mathbf{b}_{-i}$ means the vector $\mathbf{b}$ of all bids, but with the ith component deleted. It's wonky notation but you should get used to it.) We need to show that bidder i's utility is maximized by setting $b_i = v_i$. (Recall v_i is player i's immutable valuation, while the player can set her bid b_i to whatever she wants.)

Let $B = \max_{j \neq i} b_j$ denote the highest bid by some other bidder. What's special about a second-price auction is that, even though there are an infinite number of bids that i could

[3] eBay is a little more complicated, because an auction takes place over a period of time (e.g., a week), and bidders can bid multiple times. If all bidders "snipe," meaning submit their bids just as the auction is expiring, then an eBay auction is basically the same as a Vickrey auction.

[4] You've seen the concept of dominant strategies before, when we discussed the Prisoner's Dilemma (Lecture #5).

make, only two distinct outcomes can result. If $b_i < B$, then i loses and receives utility 0. If $b_i \geq B$, then i wins at price B and receives utility $v_i - B$.[5]

We now consider two cases. First, if $v_i < B$, the highest utility that bidder i can get is $\max\{0, v_i - B\} = 0$, and she achieves this by bidding truthfully (and losing). Second, if $v_i \geq B$, the highest utility that bidder i can get is $\max\{0, v_i - B\} = v_i - B$, and she achieves this by bidding truthfully (and winning). ■

Economic jargon: an auction with the property above (truthful bidding is a dominant strategy) is called *truthful* or *dominant-strategy incentive compatible (DSIC).*

The second important property is that a truthtelling bidder will never regret participating in a second-price auction.

Claim 5.2 *In a second-price auction, every truthtelling bidder is guaranteed non-negative utility.*

Proof: Losers all get utility 0. If bidder i is the winner, then her utility is $v_i - p$, where p is the second-highest bid. Since i is winner (and hence the highest bidder) and bid her true valuation, $p \leq v_i$ and hence $v_i - p \geq 0$. ■

Economic jargon: an auction with the property above (truthful bidding guarantees non-negative utility) is called *individually rational (IR).*

From a bidder's perspective, the two properties above (truthfulness and IR) make it particularly easy to figure out what to do: just participate in the auction and bid your true valuation. From the seller's or auction designer's perspective, these properties make it much easier to reason about the auction's outcome. Note that *any* prediction of an auction's outcome has to be predicated on assumptions about how bidders behave. In a truthful auction, one only has to assume that a bidder with an obvious dominant strategy will play it — behavioral assumptions don't get much weaker than that.

We also want more. For example, an auction that always gives the item away for free to bidder #7 is truthful and IR, but it makes no effort to identify which bidders actually want the good. This brings us to our third claim, which should be self-evident.

Claim 5.3 *In a second-price auction, if every bidder bids truthfully, then the item is given to the bidder who values it the most.*

Economic jargon: an auction with the property above is called *welfare-maximizing.*

Welfare maximization can be thought of as a strong form of Pareto optimality. It states that no other allocation has more total value to the bidders; in particular, no other allocation could make anyone better off without making someone else worse off.

The welfare-maximization property states something rather amazing: even though the bidder valuations were a priori unknown to the auctioneer, the auction nevertheless successfully identifies the bidder with the highest valuation! (Assuming truthful bids, which is a

[5]We're assuming here that ties are broken in favor of bidder i. The claim holds no matter how ties are broken, as you should check.

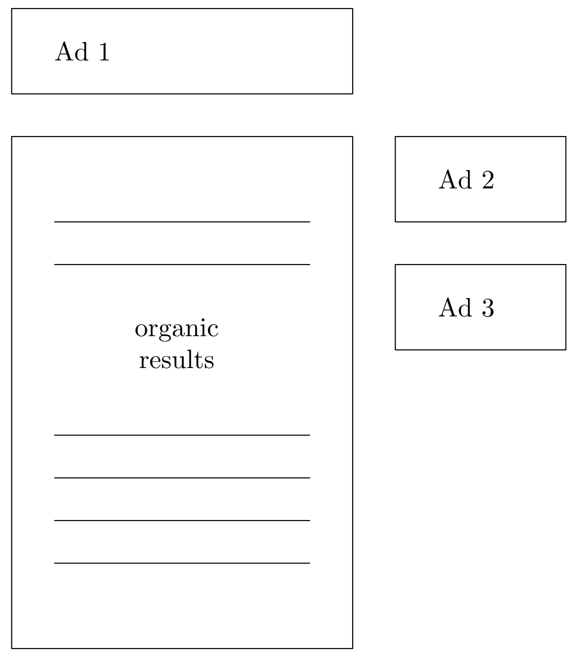

Figure 1: A typical webpage with sponsored search content.

reasonable assumption in light of the auction being truthful.) That is, the Vickrey auction maximizes welfare just as effectively as an altruistic dictator who knows the valuations of all of the bidders.[6]

In Exercise Set #7 you will explore further properties of and variations on the Vickrey auction. For example, truthful bidding is the *unique* dominant strategy for a bidder in a Vickrey auction.

Next we try to "scale up" the Vickrey auction to the more complex settings that are relevant to online advertising.

6 Case Study: Sponsored Search Auctions

6.1 Background

A Web search results page comprises a list of organic search results — deemed by some underlying algorithm, such as PageRank, to be relevant to your query — and a list of sponsored links, which have been paid for by advertisers (Figure 1). (Go do a Web search now to remind yourself, preferably on a valuable keyword like "mortgage" or "asbestos") Every time you type a search query into a search engine, an auction is run in real time to decide

[6]Note that the selling price does not participate in the social welfare. The reason is that we treat the auctioneer as a player whose utility is the revenue she earns; her utility then cancels out the utility lost by the auction winner from paying for the item. We will discuss auctions for maximizing seller revenue next week.

which advertisers' links are shown, in what order, and how they are charged. (That's a lot of auctions!) It is impossible to overstate how important such *sponsored search auctions* have been to the Internet economy.[7] Here's one jaw-dropping statistic: around 2006, sponsored search auctions generated roughly 98% of Google's revenue [1]. While online advertising is now sold in many different ways, sponsored search auctions continue to generate tens of billions of dollars of revenue every year.

6.2 The Basic Model of Sponsored Search Auctions

We discuss next a simplistic but useful and influential model of sponsored search auctions, due independently to Edelman et al. [1] and Varian [3]. The goods for sale are the k "slots" for sponsored links on a search results page. The bidders are the advertisers who have a standing bid on the keyword that was searched on. (Each bidder gets at most one slot, and each slot at most one bidder.) For example, Volvo and Subaru might be bidders on the keyword "station wagon," while Nikon and Canon might be bidders on the keyword "camera." Such auctions are more complex than single-item auctions in two ways. First, there are generally multiple goods for sale (i.e., $k > 1$). Second, these goods are *not* identical — slots higher on the search page are more valuable than lower ones, since people generally scan the page from top to bottom.

We quantify the difference between different slots using *click-through-rates (CTRs)*. The CTR α_j of a slot j represents the probability that the end user clicks on this slot. Ordering the slots from top to bottom, we make the reasonable assumption that $\alpha_1 > \alpha_2 > \cdots > \alpha_k$. For simplicity, we also make the unreasonable assumption that the CTR of a slot is independent of its occupant. Exercise Set #7 shows that it's easy to extend the theory we develop to the more general and realistic model in which each advertiser i has a "quality score" β_i (the higher the better) and the CTR of advertiser i in slot j is the product $\beta_i \alpha_j$.

We assume that an advertiser is not interested in an impression (i.e., being displayed on a page) per se, but rather has a private valuation v_i for each *click* on its link. Hence, the expected value derived by advertiser i for an impression in slot j is $v_i \alpha_j$.

6.3 Digression

It's worth noting that pay-per-click advertising is a really good idea. The commercial Internet had about 5 years of online advertising prior to pay-per-click advertising, and most of it used old-school methods, like signing a contract to show x impressions of advertiser y over a time window z for a price of p. Negotiating all of these contracts was a pain for everybody, and everybody was happy to move to a pay-per-click model. The advertisers are happy because they only have to pay for a click (not an impression), reducing their risk, and because the ad can be targeted to the users who are most likely to click on it (e.g., they searched for "camera" and the advertiser is Nikon), rather than the shotgun blast used in display advertising. The search engines, meanwhile, were happy to avoid getting locked into long-term contracts with

[7] Also called *keyword* or *position* auctions.

fixed prices, preferring the competitive prices produced by sponsored search auctions. Even users are happier, because they are likely to see ads that are highly relevant to their goals.

6.4 What We Want

Is there a sponsored search auction as cool as the Vickrey auction for the single-item case? Our desiderata are:

(1) Truthful and IR. That is, truthful bidding should be a dominant strategy, and never leads to negative utility.

(2) Welfare maximization (assuming truthful bids). In a sponsored search auction, the *social welfare* of an assignment of bidders to slots is the total value of the bidders: $\sum_{i=1}^{n} v_i \alpha_{s(i)}$, where $s(i)$ denotes the slot assigned to i. (If i is not assigned any slot, then interpret $\alpha_{s(i)}$ as 0.)

In the Vickrey auction, we accomplished the second goal by choosing the highest bidder, and the first goal by charging the second-highest bid.

6.5 Assigning Bidders to Slots

Let's start with the second goal: assuming truthful bids, how should we assign bidders to slots to maximize the social welfare? There is an obvious greedy algorithm, and it works. Specifically: assign the highest bidder to the first (best) slot, the second-highest bidder to the second slot, and so on, with the kth-highest bidder getting the last slot.

We can prove the correctness of this greedy algorithm by a canonical exchange argument. Consider an assignment of bidders to slots different from the greedy assignment. Since the greedy assignment is sorted, any other assignment is not sorted, in the sense that there must be bidders i, h assigned to slots j, ℓ with $v_i > v_h$ and yet $\alpha_j < \alpha_\ell$ (so i has the higher valuation but the worse slot). Swapping the positions of bidders i and h increases the social welfare by

$$\underbrace{v_i\left(\alpha_\ell - \alpha_j\right)}_{\text{increase in } i\text{'s value}} - \underbrace{v_h\left(\alpha_\ell - \alpha_j\right)}_{\text{decrease in } h\text{'s value}} = \underbrace{\left(v_i - v_h\right)}_{>0}\underbrace{\left(\alpha_\ell - \alpha_j\right)}_{>0}.$$

Hence, no assignment other than the greedy one maximizes the social welfare.

6.6 Choosing the Payments

We have thus extended the second goal from single-item auctions to sponsored search auctions. What about the first goal? Have we already designed a truthful sponsored search auction? No, since we haven't defined the payments yet. (You can't speak of truthfulness until committing to the payments—recall the very different-behaving single-item auctions we considered.) So the right question is: what is the analog of the second-price payment rule for sponsored search auctions?

Perhaps the most natural thing to try is: for each slot $j = 1, 2, \ldots, k$, charge (per-click) the bidder in slot i the value of $(i+1)$th highest bid. When $k = 1$, this recovers the Vickrey auction. This auction is called the *Generalized Second Price (GSP)* auction, and it has been the dominant paradigm in sponsored search. For example, when Google first rolled out Adwords in 2002, they used GSP auctions (inspired by the per-click online auctions developed earlier by Overture). They were soon adopted by the other major search engines (Yahoo!, Bing, etc.). Sponsored search auctions have been getting fancier lately, though; see the next lecture.

When GSP was developed at Google in 2001, Google did not yet employ any economists, and at that time computer scientists hadn't really studied up on auctions yet. So why GSP? The decision was made by the engineers in charge of building the system, and for engineering reasons. They thought about pay-your-bid auctions (generalizing a first-price auction), but speculated that this would lead to incessant bid updating by automated bidding agents, who would be constantly searching for the minimum bid required to maintain the current position (and this actually happened earlier in Overture's pay-your-bid auctions). Such updating was considered a harmful and pointless load on Google's servers. The solution? Just promise to shade a bidder's bid down automatically to the minimum amount needed to maintain her position.

6.7 Are GSP Auctions Truthful?

The intuition behind the GSP auction resembles that for the Vickrey auction. But is the GSP auction truthful? For example, we know that there is no incentive to submit a bid b_i higher than your value v_i. The only possible effect (compared to a truthful bid) is that you displace some bidder h with $b_h > v_i$. But then b_h becomes your new price-per-click, resulting in negative utility per click.

We know that GSP is truthful when $k = 1$ (since it's then a Vickrey auction). What about $k = 2$? To pick some random (but reasonable numbers), suppose that $\alpha_1 = .1$ and $\alpha_2 = .05$, and that there are three bidders with valuations $v_1 = 10$, $v_2 = 9$, and $v_3 = 6$. Suppose the second and third bidders bid truthfully. Should the first bidder also bid truthfully? If she does ($b_1 = 10$), then her expected utility (per-impression) is $.1(10 - 9) = .1$, since the second-highest bid is 9. If she underbids to claim the second slot (say $b_1 = 8$), then she receives fewer clicks (.05 not .1), but at a much cheaper price (6 not 9), and her expected utility is $.05(10-6) = .2$, twice as much as before. So we can conclude that the GSP auction is not truthful—there can be an incentive to underbid to acquire somewhat fewer clicks at a much reduced price.

Is there a different analog of the second-price payment rule that yields a truthful sponsored search auction? Tune in next time…

References

[1] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the Generalized Second-Price Auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[2] L. Einav, C. Farronato, J. Levin, and N. Sundaresan. Auctions versus posted prices in online markets. *Journal of Political Economy*, 2016. Forthcoming.

[3] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.

[4] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.