CS364A: Algorithmic Game Theory
Lecture #2: Mechanism Design Basics*

Tim Roughgarden†

September 25, 2013

1 Single-Item Auctions

The most sensible place to start our discussion of mechanism design — the science of rule-making — is *single-item auctions*. Recall our overarching goal in this part of the course.

Course Goal 1 Understand how to design systems with strategic participants that have good performance guarantees.

Consider a seller that had 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 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, its utility is 0. If the bidder wins at a price p, then its 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.[1]

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

[1]More complex utility models are well motivated and have been studied — to model risk attitudes, for example.

2 Sealed-Bid Auctions

For now, we'll focus on a particular 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"?

3 First-Price Auctions

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

First-price auctions are 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. We'll elaborate on the theory of first-price auctions in Problem Set #1 and in later advanced material. For now, to drive the 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?

4 Second-Price Auctions

Let's now focus on a different auction format, which is also common in practice, that 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

[2] When we study revenue maximization a few lectures hence, we'll see why other winner selection rules are important.

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 a sealed-bid auction in which the highest bidder wins and pays a price equal to the second-highest bid.

Claim 4.1 *In a second-price auction, every bidder has a dominant strategy: set its bid b_i equal to its private valuation v_i. 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, 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.

Proof of Claim 4.1: Fix an arbitrary player i, its 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 need to 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 i's immutable valuation, while it can set its bid b_i to whatever it 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 make, only 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 $v_i - B$.[3]

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 it 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 it achieves this by bidding truthfully (and winning). ■

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

Claim 4.2 *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 its utility is $v_i - p$, where p is the second-highest bid. Since i is winner (and hence the highest bidder) and bid its true valuation, $p \leq v_i$ and hence $v_i - p \geq 0$. ■

[3] 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.

The exercises ask you to 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.

5 Awesome Auctions

Taking a step back, we can claim the following.

Theorem 5.1 (Vickrey [3]) *The Vickrey auction is awesome. Meaning, it enjoys three quite different and desirable properties:*

(1) [**strong incentive guarantees**] *It is* dominant-strategy incentive-compatible (DSIC), *i.e., Claims 4.1 and 4.2 hold.*

(2) [**strong performance guarantees**] *If bidders report truthfully, then the auction maximizes the* social surplus

$$\sum_{i=1}^{n} v_i x_i, \tag{1}$$

where x_i is 1 if i wins and 0 if i loses, subject to the obvious feasibility constraint that $\sum_{i=1}^{n} x_i \leq 1$ (i.e., there is only one item).[4]

(3) [**computational efficiency**] *The auction can be implemented in polynomial (indeed, linear) time.*

All of these properties are important. From a bidder's perspective, the DSIC property, which guarantees that truthful reporting is a dominant strategy and never leads to negative utility, makes it particularly easy to choose a bid. From the seller's or auction designer's perspective, the DSIC property makes 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 DSIC 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.

The DSIC property is great when you can get it, but we also want more. For example, an auction that gives the item away for free to a random bidder is DSIC, but it makes no effort to identify which bidders actually want the good. The surplus-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 reasonable assumption in light of the DSIC property.) That is, the Vickrey auction solves the surplus-maximization optimization problem as well as if the data (the valuations) were known in advance.

[4]Note that the sale price is not part of the surplus. The reason is that we treat the auctioneer as a player whose utility is the revenue it earns; its utility then cancels out the utility lost by the auction winner from paying for the item. We will discuss auctions for maximizing seller revenue in a few lectures.

The importance of the third property is self-evident to computer scientists. To have potential practical utility, an auction should run in a reasonable amount of time — or even in real time, for some applications. Auctions with super-polynomial running time are useful only for fairly small instances.

The next several lectures strive for awesome auctions, in the sense of Theorem 5.1, for applications beyond single-item auctions. The two directions we focus on are well motivated: more complex allocation problems, like those that inevitably arise in sponsored search and combinatorial auctions; and maximizing seller revenue in lieu of social surplus.

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. (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 which advertisers' links are shown, in what order, and how they are charged. It is impossible to overstate how important such *sponsored search auctions* have been to the Internet economy. Here's one jaw-dropping statistic: around 2006, sponsored auctions generate 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 [2]. 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. 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 $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$. For simplicity, we also make the unreasonable assumption that the CTR of a slot is independent of its occupant. The exercises show how to extend the following results 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 value derived by advertiser i from slot j is $v_i\alpha_j$.

6.3 What We Want

Is there an awesome sponsored search auction? Our desiderata are:

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

(2) Social surplus maximization. That is, the assignment of bidders to slots should maximize $\sum_{i=1}^{n} v_i x_i$, where x_i now denotes the CTR of the slot to which i is assigned (or 0 if i is not assigned to a slot). Each slot can only be assigned to one bidder, and each bidder gets only one slot.

(3) Polynomial running time. Remember zillions of these auctions need to be run every day!

6.4 Our Design Approach

What's hard about mechanism design problems is that we have to jointly design two things: the choice of who wins what, and the choice of who pays what. Even in single-item auctions, it is not enough to make the "correct" choice to first design decision (i.e., giving the good to the highest bidder) — if the payments are not just right, then strategic participants will game the system.

Happily, in many applications including sponsored search auctions, we can tackle this two-prong design problem one step at a time.

Step 1: Assume, without justification, that bidders bid truthfully. Then, how should we assign bidders to slots so that the above properties (2) and (3) hold?

Step 2: Given our answer to Step 1, how should we set selling prices so that the above property (1) holds?

If we successfully answer both these questions, then we have constructed an awesome auction. Step 2 ensures the DSIC property, which means that bidders will bid truthfully (assuming as usual that a bidder with an obvious dominant strategy does indeed play that strategy). This means that the hypothesis in Step 1 is satisfied, and so the final outcome of the auction is indeed surplus-maximizing (and is computed in polynomial time).

We conclude this lecture by executing Step 1 of sponsored search auctions. Given truthful bids, how should we assign bidders to slots to maximize the surplus? As an exercise, you should show that the natural greedy algorithm is optimal (and runs in near-linear time): assign the jth highest bidder to the jth highest slot for $j = 1, 2, \ldots, k$.

Can we implement Step 2? Is there an analog of the second-price rule — sale prices that render truthful bidding a dominant strategy for every bidder? Next lecture we'll derive an affirmative answer via Myerson's Lemma, a powerful tool in mechanism design.

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

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