Economics, AI, and Optimization
Lecture Note 3: Introduction to Auctions

Christian Kroer*

January 25, 2022

1 Introduction

In this lecture note we will study the problem of how to aggregate a set of agent preferences into an outcome, ideally in a way that achieves some desirable outcome. Desiderata we might care about include *social welfare*, which is just the sum of the agent's utilities derived from the outcome, or revenue in the context of auctions.

Suppose that we have have a car, and we wish to give it to one of n people, with the goal of giving it to the person that would get the most utility out of the car. One thing we could do is ask each person to tell us how much utility they would get out of receiving the car, expressed as some positive number. This, obviously, leads to the "who can name the highest number?" game, since no person will want to tell us how much value they actually place on the car, but will instead try to name as large of a number as possible.

The above, rather silly, example shows that in general we need to be careful about how we design the rules that map the stated preferences by the agents of a mechanism into an outcome. The general field concerned with the design of rules, or *mechanisms* for designing rules that ask agents about their preferences and use that to choose an outcome is called *mechanism design*.

2 Auctions

We will mostly focus on the most classical mechanism-design setting: auctions. We will start by considering single-item auctions: there is a single good for sale, and there is a set of n buyers, with each buyer having some value v_i for the good. The goal will be to sell the item via a *sealed-bid* auction, which works as follows:

1. Each bidder i submits a bid $b_i \geq 0$, without seeing the bids of anyone else.
2. The seller decides who gets the good based on the submitted bids.
3. Each buyer i is charged a price p_i which is a function of the bid vector b.

A few things in our setup may seem strange. First, most people would not think of sealed bids when envisioning an auction. Instead, they typically envision what's called the *English auction*. In the English auction, bidders repeatedly call out increasing bids, until the bidding stops, at which

*Department of Industrial Engineering and Operations Research, Columbia University. Email: christian.kroer@columbia.edu.

point the highest bidder wins and pays their last bid. This auction can be conceptualized as having a price that starts at zero, and then rises continuously, with bidders dropping out as they become priced out. Once only one bidder is left, the increasing price stops and the item is sold to the last bidder at that price. This auction format turns out to be equivalent to the *second-price* sealed-bid auction which we will cover below. Another auction format is the *Dutch auction*, which is less prevalent in practice. It starts the price very high such that nobody is interested, and then continuously drops the price until some bidder says they are interested, at which point they win the item at that price. The Dutch auction is likewise equivalent to the *first-price* sealed-bid auction, which we cover below.

Secondly, it would seem natural to always give the item to the highest bid in step 2, but this is not always done (though we will focus on that rule). Thirdly, the pricing step allows us to potentially charge more bidders than only the winner. This is again done in some reasonable auction designs, though we will mostly focus on auction formats where $p_i = 0$ if i does not win.

When thinking about how buyers are going to behave in an auction, we need to first clarify what each buyers knows about the other bidders. The typical setting is one where each buyer i has some distribution F_i from which they draw their bid. Every buyer knows the distribution of every other buyer, but they only get to observe their own value $v_i \sim F_i$ before choosing their bid b_i. For this model, we need a new game-theoretic equilibrium notion called a *Bayes Nash equilibrium* (BNE). A BNE is a set of mappings $\{\sigma_i\}_{i=1}^n$, where $\sigma_i(v_i)$ specifies the bid that buyer i submits when they have value i, such that for all values v_i and alternative bids b_i, $\sigma_i(v_i)$ achieves at least as much utility as b_i in a Bayesian sense:

$$\mathbb{E}_{v_{-i}\sim F_{-i}}[u_i(\sigma_i(v_i), \sigma_{-i}(v_{-i}))|v_i] \geq \mathbb{E}_{v_{-i}\sim F_{-i}}[u_i(b_i, \sigma_{-i}(v_{-i}))|v_i].$$

In the auction context, $u_i(b_i, \sigma_{-i}(v_{-i})$ is utility that buyer i derives given the allocation and payment rule. The idea of a BNE works more generally for a game setup where u_i is some arbitrary utility function.

We will now introduce some useful mechanism-design terminology. We will introduce it in this single-item auction context, but it applies more broadly.

Efficiency. [*** write? ***]

Revenue. [*** write? ***]

Truthfulness, strategyproofness, and incentive compatibility. Informally, we say that an auction is *truthful* or *incentive compatible* (IC) if buyers maximize their utility by bidding their true value (i.e. $b_i = v_i$). More formally, an auction is *dominant strategy incentive compatible* (DSIC) if a buyer maximizes their utility by bidding their value, *no matter what everyone else does.* Saying that an auction (or more generally a mechanism) is "truthful" or "strategyproof" typically means that it is DSIC. DSIC auctions are very attractive because it means that buyers do not need to worry about what the other buyers will do: no matter what happens, they should just bid their value. This also means that, as auction designers, we can reasonably expect that buyers will bid their true value (or at least try to, if they are not perfectly capable of estimating it themselves). This makes it much easier to reason about aspects such as efficiency or revenue.

A slightly weaker degree of truthfulness is that of *Bayes-Nash incentive compatibility*: an auction is Bayes-Nash IC if there exists a BNE where every buyer bids their value. It is clear why this notion is less appealing: Now buyers need to worry about whether other buyers are going to bid truthfully. If they think that they will, then we might expect them to bid their value as well. However, if the system starts out in some other state, we might worry that buyers will adapt their bidding over time in a way that pushes them into some other non-truthful equilibrium.

2.1 First-price auctions

First-price auctions are perhaps what most people imagine when we say that we are selling our good via a sealed-bid auctions. In first-price auctions, each buyer submits some bid $b_i \geq 0$, and then we allocate the item to the buyer i^* with the highest bid and charge that buyer b_{i^*}. This is also sometimes referred to as *pay-your-bid*.

Let's briefly try to reason about what might happen in a first-price auction. Clearly, no buyer should bid their true value for the good under this mechanism; in that case they receive no utility even when they win. Instead, buyers should *shade* their bids, so that they sometimes win while also receiving strictly positive utility. The problem is that buyers must strategize about how other buyers will bid, in order to figure out how much to shade by.

This issue of shading and guessing what other buyers will bid happened in early Internet ad auctions, where first-price auctions were initially adopted. *Overture* was an early pioneer in selling Internet sponsored search ads via auction. They initially ran first-price auctions, and provided services to MSN and Yahoo (which were popular search engines at the time). Bidding and pricing turned out to be very inefficient, because buyers were constantly changing their bids in order to best respond to each other. Plots of the price history show a clear "sawtooth pattern," where a pair of bidders will take turns increasing their bid by 1 cent each, in order to beat the other bidder. Finally, one of the bidders reaches their valuation, at which point they drop their bid much lower in order to win something else instead. Then, the winner realizes that they should bid much lower, in order to decrease the price they pay. At that point, the bidder that dropped out starts bidding 1 cent more again, and the pattern repeats. This leads to huge price fluctuations, and inefficient allocations, since about half the time the item goes to the bidder with the lower valuation.

All that said, it turns out that there does exist at least one interesting characterization of how bidding should work in a single-item first-price auction (the Overture example technically consists of many "independent" first-price auctions; though that independence does not truly hold as we shall see later).

For this characterization, we assume the following symmetric model: we have n buyers as before, and buyer i assigns value $v_i \in [0, \omega]$ for the good. Each v_i is sampled IID from an increasing distribution function F. F is assumed to have a continuous density f and full support. Each bidder knows their own value v_i, but only knows that the value of each other buyer is sampled according to F.

Given a bid b_i, buyer i earns utility $v_i - b_i$ if they win, and utility 0 otherwise. If there are multiple bids tied for highest then we assume that a winner is picked uniformly at random among the winning bids, and only the winning bidder pays.

It turns out that there exists a *symmetric equilibrium* in this setting, where each bidder bids according to the function

$$\beta(v_i) = \mathbb{E}[Y_1 | Y_1 < v_i],$$

where Y_1 is the random variable denoting the maximum over $n-1$ independently-drawn values from F.

Theorem 1. *If every bidder in a first-price auction bids according to β then the resulting strategy profile is a Bayes-Nash equilibrium.*

Proof. Let $G(y) = F(y)^{n-1}$ denote the distribution function for Y_1.

Suppose all bidders except i bids according to β. The function β is continuous and monotonically increasing: a higher value for v_i simply adds additional values to the highest end of the distribution. As a consequence, the highest bid other than that of bidder i is $\beta(Y_1)$. It follows that bidder i should never bid more than $\beta(\omega)$, since that is the highest possible other bid. Now consider bidding

$b_i \leq \beta(\omega)$. Letting z be such that $\beta(z) = b_i$, the expected value that bidder i obtains from bidding b_i is:

$$\begin{aligned}
\Pi(b_i, v_i) =& G(z)[v_i - \beta(z)] \\
=& G(z)v_i - G(z)\mathbb{E}[Y_1 | Y_1 < z] && \text{by definition of } \beta(z) \\
=& G(z)v_i - \int_0^z y g(y) dy && \text{by definition of expectation} \\
=& G(z)v_i - G(z)z + \int_0^z G(y) dy && \text{integration by parts} \\
=& G(z)(v_i - z) + \int_0^z G(y) dy
\end{aligned}$$

Now we can compare the values from bidding $\beta(v_i)$ and b_i:

$$\begin{aligned}
\Pi(\beta(v_i), v_i) - \Pi(b_i, v_i) &= G(v_i)(v_i - v_i) + \int_0^{v_i} G(y) dy - G(z)(v_i - z) - \int_0^z G(y) dy \\
&= G(z)(z - v_i) - \int_{v_i}^z G(y) dy
\end{aligned}$$

If $z \geq v_i$ then this is clearly positive since $G(z) \geq G(y)$ for all $y \in [v_i, z]$. If $z \leq v_i$, then $G(z) \leq G(y)$, and so we have a negative number and subtract a more negative number. □

A nice property that follows from the monotonicity of β is that the item is always allocated to the bidder with the highest valuation, and thus the symmetric equilibrium is efficient.

2.2 Second-price auctions

Now we look at another pricing rule: the *second-price auction*. The *second-price auction* turns out to simply allow buyers to submit their true value as their bid. In a second-price auction, the winning bidder i^* is charged the *second-highest bid*. It's easy to see that a bidder should simply bid their valuation in this auction format. There are four cases to consider for a non-truthful bid $b_i \neq v_i$:

1. $b_i > v_i \geq b_2$ where b_2 is the second-highest bid. In that case buyer i would have gotten the same utility from bidding v_i.

2. $b_i > b_2 > v_i$ where b_2 is the second-highest bid. In that case buyer i wins, but gets utility $v_i - b_2 < 0$, and they would have been better off bidding their valuation.

3. $b_i < b_2 < v_i$ where b_2 is the second-highest bid. In that case buyer i does not win, but they could have won and gotten strictly positive utility if they had bid their valuation.

4. $b_2 < b_i < v_i$ where b_2 is the second-highest bid. In that case buyer i wins, but they would have won, and paid the same, if they had bid their true value.

It follows that the second-price auction is DSIC, because an agent should report their true valuation no matter what everybody else does. The second-price auction is also efficient, in the sense that it maximizes social welfare (since the item goes to the buyer with the highest value). Finally, it is *computable*, in the sense that it is easy to find the allocation and payments.

Note that the first-price auction is also computable, and under the symmetric equilibrium given in theorem 1 it is also efficient. But it is not truthful, and it is not hard to come up with a simple discrete setting where there is not even an equilibrium.

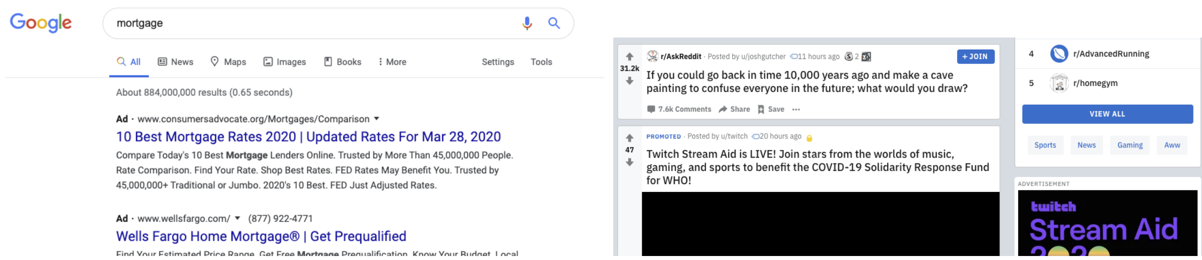

Figure 1: Left: A Google query for "mortgage" shows 2 ads. Organic search results follow further down. Right: The front page of Reddit. The second feed story is an ad.

2.3 Sponsored Search Auctions

First and second-price auctions are natural to think of due to traditional ideas of what auctions are. However, in the modern Internet era new types of auction settings have become prevalent that go beyond single-item auctions. This is largely due to Internet advertising, which funds essentially the entirety of Google as well as Facebook and other free major Internet services such as Twitter and Reddit. In these auctions there are multiple reasons why we cannot simply analyze single-item auctions as above. Two major reasons are: 1) advertisers participate in millions of auctions and have budgets that span across these auctions, and 2) each individual auction typically has multiple ad slots for sale. We will now investigate the second reason, while the first reason will be investigated in the next lecture note.

The classical example of a sponsored search auction is a Google query, where a few ads (typically 2) are shown at the top of the search. Figure 1 on the left shows an example search for the keyword "mortgage." The sponsored search auction model can also be used to approximate other settings such as the insertion of ads in a feed. For example, Reddit typically inserts 1 ad in the set of visible results before scrolling (see Figure 1 on the right), with another ad appearing in the next 10-15 results (tested March 28th 2020). Similarly, Facebook and Twitter insert 1-2 sponsored posts near the top of the feed. Truly capturing feed auctions does require some care, however. The assumption of there being a fixed number of items is incorrect for that setting. Instead, the number of ads shown depends on how far the user scrolls, the size of the ads, and what else is being shown in terms of organic content.

In the sponsored search auction model, a set of k *slots* are for sale. The slots are shown in ranked order, and the value that an advertiser derives from showing their ad in a particular slot j decomposes into two terms $v_{ij} = c_i q_j$ where c_i is the value that the advertiser places on a user clicking on their ad, and q_j is the advertiser-independent click probability of slot j. It is assumed that $q_1 \geq q_2 \geq \cdot \geq q_k$, i.e. the top slot is better than the second slot, and so on. It is assumed that advertisers are not inherently interested in getting their ad shown. Instead, their goal is to get the user to click on the ad. Hence, our auction design will only charge an advertiser if their ad is shown.

The *generalized second-price* (GSP) auction sells the k slots as follows: we collect a set of bids $b \in \mathbb{R}^n$ (assume $n \geq k$). Then we sort b (say in the order $b_1 \geq b_2 \geq \cdots \geq b_n$), and allocate the slots in order of bids (so b_1 gets slot 1, up to bid b_k getting slot k). If the user clicks on ad $i \leq k$, then advertiser i is charged the next-highest bid b_{i+1}. GSP generalizes second-price auctions in the sense that if $k = 1$ then this auction format is equivalent to the standard second-price auction (if we take expected values in lieu of the pay-per-click model).

3 Mechanism Design

More generally, in mechanism design:

- There's a set of outcomes O, and the job of the mechanism is to choose some outcome $o \in O$
- Each agent i has a private *type* $\theta_i \in \Theta_i$, that they draw from some publicly-known distribution F_i
- Each agent i has some publicly-known valuation function $v_i(o|\theta_i)$ that specifies a utility for each outcome, given their type
- The goal of the center is to design a mechanism that maximizes some objective, e.g. social welfare $\sum_i u_i(o|\theta_i)$

A mechanism takes as input a vector of reported types θ from the players, and outputs an outcome, formally it is a function $f : \times_i \Theta_i \to O$ that specifies the outcome that results from every possible set of reported types. In mechanism design with money, we also have a *payment function* $g : \times_i \Theta_i \to \mathbb{R}^n$ that specifies how much each agent pays under the outcome. In less formal terms, a mechanism merely specifies what happens, given the reported types from the agents. In first and second-price auctions the outcome function was the same (allocate to the highest bidder), but the payment function was different. We could potentially also allow randomized mechanism $f : \times_i \Theta_i \to \Delta(O)$ that map to a probability distribution over the outcome space.

How do we analyze what happens in a given mechanism? The ideal answer is that every agent is best off reporting their true type, no matter what everybody else does, i.e. the mechanism should be DSIC. Formally, that would mean that for every agent i, type $\theta_i \in \Theta_i$, any type vector θ_{-i} of the remaining agents, and misreported type $\theta'_i \in \Theta_i$:

$$\mathbb{E}\left[v_i(f(\theta_i, \theta_{-i}))\right] \geq \mathbb{E}\left[v_i(f(\theta'_i, \theta_{-i}))\right],$$

where the expectation is over the the potential randomness of the mechanism. If there is also a payment function g and agents have *quasi-linear utilities* then the inequality is

$$\mathbb{E}\left[v_i(f(\theta_i, \theta_{-i})) - g(\theta_i, \theta_{-i})\right] \geq \mathbb{E}\left[v_i(f(\theta'_i, \theta_{-i})) - g(\theta'_i, \theta_{-i})\right],$$

A less satisfying answer is that there exists a Bayes-Nash equilibrium of the game induced by the mechanism, in which every agent reports their true type. Formally, that would mean that for every agent i, type $\theta_i \in \Theta_i$, and misreported type $\theta'_i \in \Theta_i$:

$$\mathbb{E}\left[v_i(f(\theta_i, \theta_{-i}))\right] \geq \mathbb{E}\left[v_i(f(\theta'_i, \theta_{-i}))\right],$$

where the expectation is over the types θ_{-i} of the other agents, and the potential randomness of the mechanism. This constraint just says that reporting their true type should maximize their expected utility, given that everybody else is truthfully reporting. This can likewise be generalized for a payment function g.

In the setting where we can charge money, the *Vickrey-Clarke-Groves* (VCG) mechanism is DSIC and maximizes social welfare. In VCG, after receiving the type vector θ, we pick the outcome o that maximizes the report welfare. The key to then making this choice incentive compatible is that we charge each agent their *externality*:

$$\max_{o' \in O} \sum_{i' \neq i} v_{i'}(o'|\theta_{i'}) - \sum_{i' \neq i} v_{i'}(o|\theta_{i'}).$$

The externality measures how much better off all the other agents would have been if i were not there. When we add together the value received by player i minus their payment, we get that their utility function is:

$$\sum_{i'} v_{i'}(o|\theta_{i'}) - \max_{o' \in O} \sum_{i' \neq i} v_{i'}(o'|\theta_{i'})$$

Intuitively, we see that i cannot affect the negative term here, and the positive term is exactly the social welfare. Thus we get that each agent i is incentivized to maximize social welfare, which is achieved by reporting their true type θ_i.

4 Historical Notes

The issues with first-price in the context of Overture's sponsored search auctions is described in Edelman and Ostrovsky [2], which also shows plots from real data exhibiting the sawtooth pattern. The derivation of the symmetric equilibrium of the first-price auction follows the proof from Krishna [5]. Interestingly, first-price auctions have experiences a resurgence in the context of display advertising, where many independent ad exchanges moved to first price in 2018, and Google followed suit and moved their Ad Manager to first price in 2019[1].

The second-price auction is sometimes referred to as the *Vickrey auction* after its inventor [7]. The generalized second-price auction was described by Edelman et al. [3], though it had been in use in the Internet ad industry for a while at that point. The VCG mechanism was described in a series of papers by Vickrey [7], Clarke [1], and Groves [4]. A full description of a slightly more general VCG mechanism, and proof of correctness, can be found in Nisan et al. [6, Chapter 9]

References

[1] Edward H Clarke. Multipart pricing of public goods. *Public choice*, pages 17–33, 1971.

[2] Benjamin Edelman and Michael Ostrovsky. Strategic bidder behavior in sponsored search auctions. *Decision support systems*, 43(1):192–198, 2007.

[3] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American economic review*, 97(1):242–259, 2007.

[4] Theodore Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

[5] Vijay Krishna. *Auction theory.* Academic press, 2009.

[6] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory.* Cambridge University Press, 2007.

[7] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[1] see `https://www.blog.google/products/admanager/update-first-price-auctions-google-ad-manager/`