Lecture Note 9: Internet Advertising Auctions: Position Auctions

Christian Kroer*

February 28, 2022

1 Introduction

In this lecture note we study more advanced auction concepts, motivated by internet advertising auctions. Internet advertising auctions provide the funding for almost every free internet service such as google search, facebook, twitter, and so on. At the heart of these monetization schemes is a market design based around independently running auctions every time a user shows up. This happens many times per second, advertisers participate in thousands or even millions of auctions, have budget constraints that span the auctions, and each user query generates multiple potential slots for showing ads. For all these reasons, these markets turn out to require a lot of new theory for understanding them. Similarly, the scale of the problem necessitates the design of algorithmic agents for bidding on behalf of advertisers.

First we will introduce the *position auction*, which is a highly structured multi-item auction. There, we will look at the two most practically-important auction formats: the generalized second-price auction (GSP), and the Vickrey-Clarke-Groves (VCG) auction. Then, we will study auctions with budgets and repeated auctions. In internet ad auctions, advertiser budgets must be satisfied across all the auctions they participate in. This causes traditional guarantees such as the strategyproofness of the second-price (SP) auction to break.

1.1 Considerations for internet advertising

First and second-price auctions are natural to think of due to traditional ideas of what auctions are. However, even if we wish to run an independent auction every time a user shows up, we still have multiple potential slots for showing ads, and thus we must either sell those independently (which could lead to the same ad winning both slots, pricing issues, and so on), or design some sort of multi-item auction. In practice the latter approach is taken.

A classical example of an internet advertising 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." Of note here is the fact that two ads are shown: the auction that was run when I searched for "mortgage" was a multi-item auction (and in particular there were two items), and thus two winners of the auction had to be selected. Naturally, the advertisers are going to care about being shown first versus second, since the user is typically more likely to click on the first ad. This setting is referred to as the "sponsored search setting." Sponsored search will motivate the *position auction* setting below, where a ranked set of positions are sold.

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

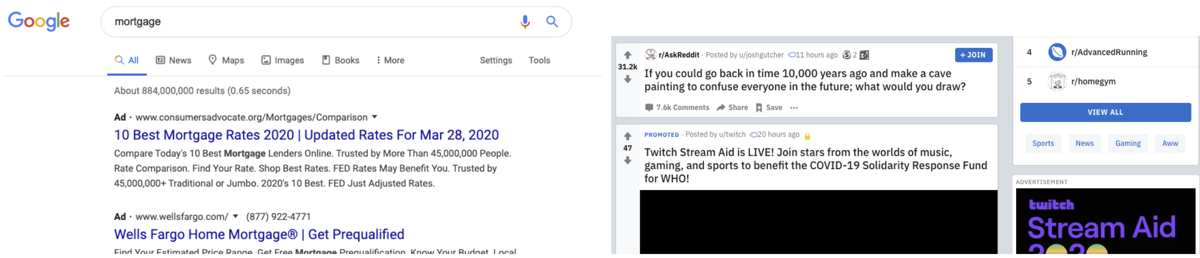

Figure 1: Left: A Google query for "mortgage" shows 2 ads. Organic search results follow further down. Right: The front page of Reddit. The second feed story is an ad.

The position auction model can also be used to approximate other settings such as the insertion of ads in a *news feed*; a news feed is the familiar infinitely-scrolling list of e.g. Facebook posts, Reddit posts, Instagram posts, or Twitter posts. 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. We will focus on the simpler setting with a fixed number of slots, but properly handling feed auctions is an interesting problem.

Beyond the multi-item and budget aspects, internet advertising has a few other interesting quirks. Below these are discussed briefly, though we will mostly abstract away considerations around these issues.

1.1.1 Targeted advertising.

In a classical advertising setting such as TV or newspaper advertising, the same ad is shown to every viewer of a given TV channel, or every reader of a newspaper. This means that it is largely not feasible for smaller, and especially niche, retailers to advertise, since their return on investment is very low due to the small fraction of viewers or readers that fit their niche. All this changed with the advent of internet advertising, where niche retailers can perform much more fine-grained targeting of their ads. This has enabled many niche retailers to scale up their audience reach significantly.

One way that targeting can occur is directly through association with the search term in sponsored search. For example, by bidding on the search term "mortgage," a lender is effectively performing a type of targeting. However, a second type of targeting occurs by matching on query and user features (such targeting is used across many types of internet advertising including search, feed ads, and others). For example, a company selling surf boards might wish to target users at the intersection of the categories {age 16-30, lives in California}. Because each individual auction corresponds to a single user query, the idea of targeted advertising can be captured in the valuations that we will use for the buyers in our auction setup: each buyer corresponds to an advertiser, each auction corresponds to a query, and the buyer will have value zero for all items in a given auction if the associated query features do not match their targeting criteria.

Targeted advertising has the potential for some adverse effects. Of particular note are demographic biases in the types of ads being shown (a well-documented example is that in some settings, ads for new luxury housing developments were disproportionately shown to white people). In a later lecture note we will study such questions around demographic fairness. A second potential issue is that of user privacy. This is an interesting topic that we will unfortunately not have too much to say on, as it is outside the scope of the course.

1.1.2 Pay per click.

Another revolution compared to pre-internet advertising is the *pay per click* nature of most internet advertising auctions. Many advertisers are not actually interested in the user simply viewing their ad. Instead, their goal is to get the user to click on the ad, or even something downstream of clicking on the ad, such as selling the advertised product via the linked website. Because the platform, such as google, is in a much better position to predict whether a given user will click on a given ad, these auctions operate on a *cost per click* basis, rather than a cost per impression. What this means is that any given advertiser does not actually pay just because they won the auction and got their ad shown, instead they pay if the user actually clicks on their ad.

From an auction perspective, this means that the valuations used in the auctions must take into account the probability that the user will actually click on the ad. Valuations are typically constructed by breaking down the value that a buyer i (in this case an advertiser) has for an item (which is a particular slot in the search query or user feed) into several components. The *value per click* of advertiser i is the value $v_i > 0$ they place on any user clicking on their ad (modern platforms generalize this concept to a value per *conversion*, where a conversion can be a click, an actual sale of a product, the user viewing a video, etc.) The *click-through-rate* is the likelihood that the user for query j will click on the ad of advertiser i, independently of where on the page the ad is shown. We denote this by CTR_{ij}; we will assume that $CTR_{ij} = 0$ if query j does not fall under the targeting criteria of buyer i. Finally, the *slot qualities* $q_1, \dots, q_S$ are scalar values denoting the quality of each slot that an ad could end up in. These are monotonically decreasing values, indicating the fact that it's generally preferable to be shown higher up on the page. Now, finally, the value that buyer i has for being shown in slot s of query j is modeled as $v_{ijs} = v_i \cdot CTR_{ij} \cdot q_s$.

For the rest of the lecture note, we will assume that v_{ij} is the value that buyer i has for auction j; this value encodes the value per click, the CTR, and the targeting criteria (but can allow for more general valuations that do not decompose). Note that this assumes correct CTR predictions, which is obviously not true in practice. In practice the CTRs are estimated using machine learning, and it is of interest to understand which discrepancies this introduces into the market. Secondly, we are assuming that buyers are ok maximizing their expected utility, rather than observed utility. This is largely a non-problem, since they will participate in thousands or even millions of auctions, and thus their realized value can reasonably be expected to match the expectation (at least if the CTRs are correct). The slot quality q_s will be handled separately in the next section. Once we start discussing budgets, we will keep the presentation simple by assuming a single item per auction, thus avoiding the need for slot qualities.

2 Position Auctions

In the position auction model, a set of S *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 s decomposes into two terms $v_{is} = v_i q_s$ where v_i is the value that the advertiser places on a user clicking on their ad, and q_s is the advertiser-independent click probability of slot s. Here we assume that v_i already incorporates the click-through rate (so in particular it could be that $v_i = v_i' \cdot CTR_i$ where v_i' is their actual value per click, and CTR_i is the click-through rate in the current auction). It is assumed that $q_1 \geq q_2 \geq \cdot \geq q_S$, i.e. the top slot is better than the second slot, and so on.

Now suppose that the n advertisers submit bids $b \in \mathbb{R}^n_+$. Both auction formats we will use then proceed to perform allocation via welfare maximization, assuming that the bids are truthful. We will also refer to this as bid maximization. In particular, we sort b (suppose without loss of generality that the bids are conveniently already ordered by buyer index: $b_1 \geq b_2 \geq \cdots \geq b_n$), and allocate the slots in order of bids (so buyer 1 with bid b_1 gets slot 1, buyer 2 gets slots 2, and so on up to bid b_S getting slot S).

Example 1. *Suppose we have two slots with quality scores $q_1 = 1, q_2 = 0.5$, and three buyers with values $v_1 = 10,\ v_2 = 8,\ v_3 = 2$, and suppose they all bid their values. Then buyer 1 is allocated slot 1, and they generate a value of $v_1 \cdot q_1 = 10$, buyer 2 is allocated slot 2 and they generate a value $v_2 \cdot q_2 = 4$, and buyer 3 gets nothing.*

2.1 Generalized Second-Price Auctions

The *generalized second-price* (GSP) auction sells the S slots as follows: First, we allocate via bid maximization as described above. If the user clicks on ad $i \leq S$, then advertiser i is charged the next-highest bid b_{i+1}. GSP generalizes second-price auctions in the sense that if $S = 1$ 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). However, this is a fairly superficial generalization, since GSP turns out to lose the core property of the second-price auction: truthfulness!

In particular, consider example 1 again. With GSP prices, buyer 1 gets utility $q_1(v_1 - v_2) = 2$ when everyone bids truthfully. If buyer 1 instead bids some value between 2 and 8, then they get utility $q_2(v_1 - v_3) = 4$. Thus, buyer 1 is better off misreporting. More generally, it turns out that the GSP auction can have several pure-Nash equilibria, and some of these lead to allocations that are not welfare-maximizing. Consider the following bid vector for example 1, $b = (4, 8, 2)$. Buyer 1 gets utility $0.5(10 - 2) = 4$ (whereas they'd get utility 2 for bidding above 8). Buyer 2 gets utility $1(8 - 4) = 4$ (whereas they'd get utility $0.5(8 - 2) = 3$ for bidding below 4). Buyer 3 is priced out.

2.2 VCG for Position Auctions

The second pricing rule we will consider is the VCG rule. Recall that VCG computes the welfare-maximizing allocation (assuming truthful bids), and then charges buyer i their externality (i.e. how much the presence of buyer i decreases the social welfare across the remaining agents).

Let W^S_{-i} be the social welfare achieved by buyers $[n] \setminus i$ if we maximize welfare across only those buyers, and let W^{S-i}_{-i} be the social welfare of $[n] \setminus i$ if we maximize welfare using only the slots $\{1, \dots, i-1, i+1, \dots S\}$.

$$\begin{aligned} W^S_{-i} - W^{S-i}_{-i} &= \sum_{k \in \{i+1, \dots, S+1\}} b_k \cdot s_{k-1} - \sum_{k \in \{i+1, \dots, S\}} b_k \cdot s_k && (1) \\ &= \sum_{k \in \{i+1, \dots, S+1\}} b_k \cdot (s_{k-1} - s_k) && (2) \end{aligned}$$

Theorem 1. *The VCG auction for position auctions is truthful.*

Proof. Suppose again that buyer bids are sorted, with buyer i winning slot i when bidding truthfully. Now suppose buyer i misreports and gets slot k instead. Now we want to show that bidding truthfully maximizes utility, which means:

$$s_i \cdot v_i - [W^S_{-i} - W^{S-i}_{-i}] \geq s_k \cdot v_i - [W^S_{-i} - W^{S-k}_{-i}].$$

Simplifying this expression gives

$$s_i \cdot v_i + W_{-i}^{S-i} \geq s_k \cdot v_i + W_{-i}^{S-k}.$$

Now we see that both the right-hand and left-hand sides correspond to the (reported) social welfare for particular potential allocations. Given that VCG picked the left-hand side, and VCG allocates via welfare maximization, the left-hand side must be larger. □

3 Historical Notes

An early version of the GSP auction was introduced in the early internet search days at Overture, which was an innovator in sponsored search advertising, and they were later acquired by Yahoo, which used this rule as well. Google then started using the more modern version of GSP. From an academic perspective, the GSP rule and position auctions in general started to be studied by Varian [2], Edelman et al. [1], motivated by its use in practice. An interesting historical perspective on why VCG was chosen is discussed by Varian and Harris [3] who worked at Google at the time. The primary reasons are essentially inertial: a lot of engineering work was already going into GSP, and advertisers had gotten used to bidding in GSP. A major concern would be that they would need to raise their bids in VCG due to its truthfulness, which might be hard to explain to them given their existing experience with GSP. Facebook notably uses VCG rather than GSP [3], unlike the prior internet companies.

References

[1] 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.

[2] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.

[3] Hal R Varian and Christopher Harris. The vcg auction in theory and practice. *American Economic Review*, 104(5):442–45, 2014.