CS269I: Incentives in Computer Science
Lecture #15: The VCG Mechanism*

Tim Roughgarden†

November 14, 2016

1 GSP vs. VCG

First we prove the important result stated last lecture about sponsored search auctions. Recall the setup: there are k slots with click-through-rates (CTRs) $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$, and each bidder $i = 1, 2, \ldots, n$ has a private valuation per-click v_i. (Bids and prices are also per-click.)

Our first attempt to generalize the Vickrey auction to sponsored search resulted in the "Generalized Second Price (GSP)" auction, where each bidder pays the bid of the next-highest bidder (per-click).

The GSP Auction

1. Accept a bid from each bidder. Relabel the bidders so that $b_1 \geq b_2 \geq \cdots \geq b_n$.
2. For $i = 1, 2, \ldots, k$, assign the ith bidder to the ith slot.
3. For $i = 1, 2, \ldots, k$, charge bidder i a price of b_{i+1} per-click.

We saw in Lecture #13 that the GSP auction is not truthful, as bidders can have an incentive to underbid to get slightly fewer clicks at a much cheaper price.

We also saw a sponsored search auction that *is* truthful (as proved in Section 3), the "Vickrey-Clarke-Groves (VCG)" auction.

The VCG Auction (for Sponsored Search)

1. Accept a bid from each bidder. Relabel the bidders so that $b_1 \geq b_2 \geq \cdots \geq b_n$.

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

2. For $i = 1, 2, \ldots, k$, assign the ith bidder to the ith slot.

3. For $i = 1, 2, \ldots, k$, charge bidder i

$$\frac{1}{\alpha_i} \sum_{j=i+1}^{k+1} b_j(\alpha_{j-1} - \alpha_j) \tag{1}$$

per click. (With α_{k+1} interpreted as 0.)

Recall that a bidder's payment corresponds to the "externality" she imposes on the other bidders, meaning the loss in their social welfare attributable to i's presence. For example, by existing, bidder 1 deprives bidder 2 of an additional $\alpha_1 - \alpha_2$ clicks, bidder 3 of an additional $\alpha_2 - \alpha_3$ clicks, ..., and finally bidder $k+1$ of α_k clicks. The CTR differences get multiplied by the valuation of the bidder being thus deprived (or rather the bid, as a proxy for the private valuation). This is the per-impression externality; the leading $\frac{1}{\alpha_i}$ factor translates this externality to per-click payments (so that the expected payment for an impression equals the externality).

Last time we noticed that the VCG price of a slot i is a weighted average of the lower bids $b_{i+1}, \ldots, b_n$. Since the GSP price would simply be b_{i+1} (the highest of $b_{i+1}, \ldots, b_n$), we can conclude that VCG prices are only smaller than GSP prices (for a fixed bid vector). In effect, the overly aggressive prices in the GSP auction incentivize bidders to shade their bids (reminiscent of a first-price single-item auction, but more subtle).

The main result of this section is the following.

Theorem 1.1 ([2, 7]) *For any valuations and click-through rates, the GSP auction has an equilibrium equivalent to the truthful outcome of the VCG auction.*

By "equivalent," we mean that the assignment of the bidders to slots is the same, and the payments charged are the same. (Hence both the social welfare and the revenue are also identical.)

Theorem 1.1 does not assert that the equilibrium is unique, nor does it assert that all equilibria are equivalent to the truthful VCG outcome. Last time we saw that there can be other equilibria of GSP that have strictly more or strictly less revenue than the truthful VCG outcome. There can even be equilibria of GSP where bidders are not assigned to the same slots as in the truthful VCG outcome (Exercise Set #8). Thus no sweeping revenue comparison of the VCG and GSP auctions is possible. But still, Theorem 1.1 suggests that there's no real reason to expect GSP to be worse than VCG from a revenue or social welfare standpoint.

Before proving Theorem 1.1, it's useful to have a better understanding of the prices in the VCG auction. The next lemma identifies some of their properties. For simplicity, we assume from now on that the CTRs are strictly decreasing and that bidders' valuations are distinct. (Otherwise, the strict inequalities in the theorem below become weak inequalities.)

Lemma 1.2 (Properties of VCG Prices) *Relabel the bidders so that $v_1 > v_2 > \cdots > v_n$, and assume that $\alpha_1 > \alpha_2 > \cdots > \alpha_k$. Let*

$$p_i = \frac{1}{\alpha_i} \sum_{j=i+1}^{k+1} v_j(\alpha_{j-1} - \alpha_j) \tag{2}$$

denote the per-click payment in the ith slot in the truthful outcome of the VCG auction. Then:

(a) $p_i < v_{i+1}$ for every $i = 1, 2, \ldots, k$;

(b) $p_1 > p_2 > \cdots > p_k$;

(c) the p_i's satisfy envy-freeness*, meaning that for all i and ℓ, bidder i prefers her slot i and her per-click price p_i to being assigned slot ℓ at the going per-click price p_ℓ:*

$$\underbrace{\alpha_i(v_i - p_i)}_{i\text{'s current utility}} \geq \underbrace{\alpha_\ell(v_i - p_\ell)}_{i\text{'s utility after changing places with } \ell}. \tag{3}$$

Proof: Part (a) we already know: since the VCG payment for slot i is a weighted average of the lower bids $b_{i+1}, \ldots, b_n$, in the truthful outcome it is at most v_{i+1}.

Part (b) should be intuitive (better slots go for higher prices), and the proof is some simple algebra. Looking at the expression (2), we see that the only differences between p_i and p_{i+1} are that the former has one extra summand and that they have different leading scaling factors. So we can write

$$p_i = \underbrace{\left(\frac{\alpha_i - \alpha_{i+1}}{\alpha_i}\right) v_{i+1}}_{\text{first term of } p_i} + \underbrace{\left(\frac{\alpha_{i+1}}{\alpha_i}\right)}_{\substack{\text{toggles} \\ \text{denominator}}} \cdot \underbrace{p_{i+1}}_{\substack{\leq v_{i+2} \\ < v_{i+1}}},$$

where to bound p_{i+1} by v_{i+2} we are using part (a). Since p_i is a weighted average of p_{i+1} and a number strictly bigger than p_{i+1} (namely v_{i+1}), it follows that $p_i > p_{i+1}$ for every $i = 1, 2 \ldots, k-1$.[1]

For the envy-freeness property in part (c), assume that $\ell > i$. (The case where $\ell < i$ is symmetric.) Rearranging, the inequality (3) holds if and only if

$$\alpha_i p_i - \alpha_\ell p_\ell \leq v_i(\alpha_i - \alpha_\ell).$$

[1]Since we assumed $\alpha_i > \alpha_{i+1}$, the weighted average is nondegenerate.

The following derivation (using (2)) verifies this inequality:

$$
\begin{aligned}
\alpha_i p_i - \alpha_\ell p_\ell &= \left(\sum_{j=i+1}^{k+1} v_j(\alpha_{j-1} - \alpha_j)\right) - \left(\sum_{j=\ell+1}^{k+1} v_j(\alpha_{j-1} - \alpha_j)\right) \\
&= \sum_{j=i+1}^{\ell} \underbrace{v_j}_{\le v_i} (\alpha_{j-1} - \alpha_j) \\
&\le v_i \sum_{j=i+1}^{\ell} (\alpha_{j-1} - \alpha_j) \\
&= v_i(\alpha_i - \alpha_\ell).
\end{aligned}
$$

■

With the lemma in hand, we can return our focus to the GSP auction and prove Theorem 1.1.

Proof of Theorem 1.1: We proceed by the "guess and check" method. Fix valuations $v_1 > v_2 > \cdots > v_n$ and CTRs $\alpha_1 > \alpha_2 > \cdots > \alpha_k$.[2] Let p_i denote the price (2) of slot i in the truthful outcome of the VCG auction. Our guess for the desired equilibrium is that $b_i = p_{i-1}$ for each $i = 2, 3, \ldots, k+1$. Then, if the ith bidder gets assigned to the ith slot (for each i), the GSP price of slot $i-1$ will be the bid of the ith bidder, which by definition is the VCG price p_{i-1}. To complete the description, set $b_i = v_i$ for $i = 1$ and $i > k+1$.[3] Lemma 1.2(b) implies that $b_1 > b_2 > \cdots > b_n$ and hence the bidders are ranked in decreasing order of valuation.[4] With this bid vector $\mathbf{b}$, the outcome of the GSP auction (slot assignments and prices) is exactly the same as in the truthful outcome of the VCG auction. But is $\mathbf{b}$ an equilibrium of the GSP auction?

To show that it is, suppose a bidder i deviates from $\mathbf{b}$ to $(b_i', \mathbf{b}_{-i})$, resulting in the new slot ℓ. To evaluate whether or not this is a beneficial deviation, we only care about the resulting slot ℓ, and not the deviating bid b_i' per se (since the price of i's slot and hence her utility is independent of her bid).

We now consider two cases. First suppose that $\ell > i$, meaning that the deviation drops i down to a lower slot. (Cf., the proof that the GSP auction is not truthful at the end of Lecture #13.) In this case, $b_i' \in (b_{\ell+1}, b_\ell)$, and all of the bidders previously in slots $i+1, \ldots, \ell$ get bumped up one slot (Figure 1(a)). The bidder in slot $\ell+1$, and hence the GSP price of slot ℓ, remains the same: $b_{\ell+1} = p_\ell$. Thus i's deviation results in her getting the ℓth slot at the VCG price p_ℓ of that slot. By the envy-freeness property of the VCG prices (Lemma 1.2(c)), this is not a beneficial deviation.

[2]The theorem remains true even when there are ties, as you should check.

[3]All that matters is that $b_1 > b_2$, and that $b_i \le v_i$ for every $i > k+1$.

[4]We're also using that $b_2 = p_1 < v_1$, which implies that the first bidder gets the first slot, and that $p_k = b_{k+1} = v_{k+1}$, which implies that bidder $k+1$ is indeed the one setting the price for slot k (rather than any of $k+2, k+3, \ldots, n$).

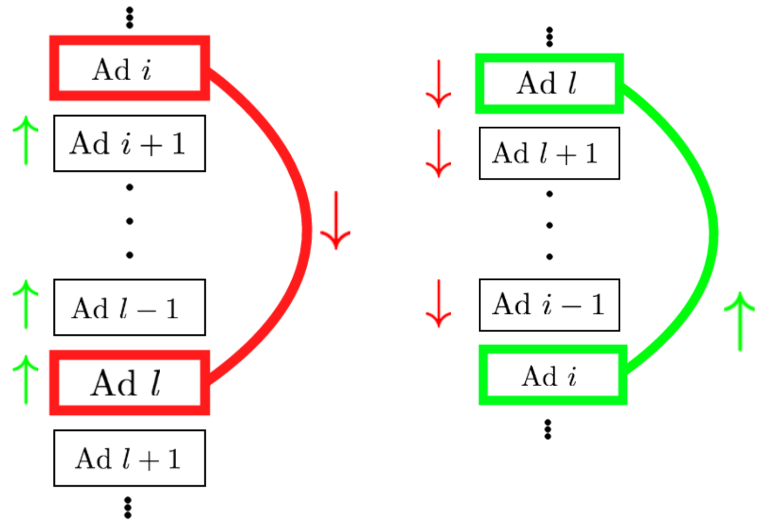

(a) Case $\ell > i$

(b) Case $\ell < i$

Figure 1: When $\ell > i$, bidders previously in slots $i+1, \ldots, \ell$ get bumped up one slot and the price of slot ℓ remains the same: $b_{\ell+1} = p_\ell$. When $\ell < i$, bidders previously assigned to slots $\ell, \ell+1, \ldots, i-1$ drop down a slot and the price for slot ℓ is the bid of its previous occupant: $b_\ell = p_{\ell-1}$.

The second case is similar but not exactly the same. Suppose that $\ell < i$, meaning that i bids higher to get a better slot. In this case, $b'_i \in (b_\ell, b_{\ell-1})$, and the bidders previously assigned to slots $\ell, \ell+1, \ldots, i-1$ drop down a slot (Figure 1(b)). This means that the new GSP price for slot ℓ will be the bid of its previous occupant (who is now in slot $\ell+1$), which is $b_\ell = p_{\ell-1}$. This is only greater than the VCG price of slot ℓ (by Lemma 1.2(b)), and hence envy-freeness (Lemma 1.2(c)) again implies that this is not a beneficial deviation for i. ■

2 Beyond Sponsored Search

Sponsored search auctions continue to be an important part of Internet advertising, but over the past 5+ years several companies have developed real-time auctions for more complex scenarios. The VCG auction is more easily generalized to complex settings than the GSP auction, and for this reason suitable generalizations of the VCG auction are becoming increasingly prevalent in practice.

For example, how do ads on Facebook differ from the basic sponsored search auctions that we've focused on so far? There are many answers to this question. For example, determining which ads are "relevant" in a sponsored search auction is largely determined by the user's search query, whereas Facebook must use other information for this purpose (like a user's friends, recent activity, etc.). Also, ads compete more directly with organic results in Facebook (via its news feed) than in a sponsored search auction. Similarly, real estate devoted to ads is real estate taken away from other things Facebook might want to show you, like friend recommendations. Another difference is that Facebook ads have different sizes and formats, rather than just being sponsored links that each takes up a single slot. Finally, Facebook allows advertisers to bid not only on clicks, but on many other events as well (e.g. likes or downloads of an app).

2.1 Dynamic Resizing

Let's drill down on two practical issues that led to a migration from GSP to VCG auctions (one at Google, one at Facebook). The first issue is *dynamic resizing* of ads. The motivating scenario here is when there is one relevant ad that is much better than all the rest (which is not unusual). In the standard sponsored search auction, the good ad would get the top slot, with the scrubs getting the rest. But you might want to do something more extreme, like show a much bigger version of the good ad and not show the other ads at all. This resizing can result in a higher click-through-rate than in the standard auction. With dynamic resizing, however, the CTRs are not fixed, and depend on the layout chosen by the platform. Does this matter?

Looking back at the sponsored search discussion, who needed to keep track of slots' CTRs? In the GSP auction, the seller doesn't need to know the CTRs—the payment for the jth slot is the $(j+1)$th highest bid, no matter what the CTRs are. It is the *bidders* that need to know the CTRs. For example, recall the basic reason why GSP is not truthful: a bidder might want to drop down to a lower slot to get slightly fewer clicks at a much cheaper price. The cost-benefit analysis of doing this depends on how CTRs of lower slots compare to the current one. Similarly, in the equilibrium constructed in the proof of Theorem 1.1, the bids depend on the CTRs (because the p_i's depend on the CTRs (2)).

In the VCG auction, the situation is reversed. Bidders do not need to know the CTRs—the optimal bid is a truthful one, no matter what the CTRs are. The seller does need to know the CTRs, in order to compute the correct payments (1).

Who's in the best position to know the CTRs? Generally the seller. For example, a search engine has all of the information about CTRs that an advertiser has, plus much more. This suggests a sense in which the VCG auction is better than the GSP auction already in sponsored search. Once the CTRs are not even fixed (as with dynamic resizing), the argument for using VCG becomes even stronger. When the GSP auction is used with variable CTRs, it is not clear how a bidder should bid, nor how a seller should interpret the bids. For this reason, Google now uses (a suitable generalization of) the VCG auction for selling "contextual ads" [8]. These are for ads that are shown on a third-party Web site like `nytimes.com` or `cnn.com`, as opposed to on a Google search results page.[5]

2.2 Different Event Types

An advertiser on Facebook can submit a bid for clicks, likes, app downloads, and a number of other events. An advertiser then pays per event, for whichever event it chose to bid on. The

[5]For sponsored search auctions, Google continues to use the GSP auction.

GSP auction can handle this extension when there is a simple relationship between different events—for example, if the probability of a click is always twice the probability of a like.[6]

It has been empirically determined that there is no simple relationship between all of the different events. For example, the probabilities of some types of events decay qualitatively faster (as you go down the news feed) than those of others. This means that greedily assigning advertisers to slots is no longer guaranteed to maximize social welfare, even assuming truthful bids (see Exercise 34). For this reason and others, Facebook's online advertising system is based on (a suitable generalization) of the VCG auction.

3 The General VCG Mechanism

The VCG auction for sponsored search is just a special case of the very general *VCG mechanism* [9, 1, 3]. The abstract setup for the general mechanism is:

- n strategic participants (e.g., Facebook advertisers);
- a finite set Ω of outcomes (e.g., all possible ways of laying out a Facebook page with organic and sponsored content);
- each agent i has a private valuation $v_i(\omega)$ for each outcome $\omega \in \Omega$.

The outcome set Ω is abstract and could be very large. In a single-item auction, Ω has only $n+1$ elements, corresponding to the winner of the item (if any). In sponsored search, the outcome set Ω would be all (exponentially many) ways of assigning bidders to slots. The outcome space for a Facebook advertising auction is still more complex.

While the general setup allows each bidder to have a different valuation for each outcome, in applications one makes some simplifying assumptions. For example, in a single-item auction, we assumed that $v_i(\omega) = 0$ whenever i was not the winner. In sponsored search auctions, we assumed that $v_i(\omega)$ depended (linearly) only on the click-through rate i obtains in ω. In a Facebook auction, where outcomes correspond to page layouts of organic and sponsored content, it's sensible to assume that

$$v_i(\omega) = (\text{value of an event}) \cdot (\mathbf{Pr}[\text{event occurs in } \omega]), \tag{4}$$

for whatever event i chooses to bid on. Facebook devotes enormous effort to learning from data accurate estimates of the probabilities in (4).

Here is the VCG mechanism in its general form.

The VCG Mechanism (in General)

1. Accept a bid $b_i(\omega)$ from each bidder i for each outcome $\omega \in \Omega$. (In practice this information is usually provided implicitly, for example via a bid-per-event.)

[6] In this case, bids can just be scaled according to the event type, and then ranked as usual in GSP. (Cf., Exercise 32.)

2. Implement the outcome ω^* that maximizes the reported social welfare

$$\sum_{i=1}^{n} b_i(\omega) \tag{5}$$

over all $\omega \in \Omega$.

3. Charge each bidder i her externality, meaning the welfare loss caused to the other bidders by i's presence:

$$p_i = \underbrace{\left(\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) \right)}_{\text{without } i} - \underbrace{\sum_{j \neq i} b_j(\omega^*)}_{\text{with } i}, \tag{6}$$

where ω^* is the outcome chosen in the second step.

The payments in (6) specialize to those used in the VCG sponsored search auction (per-impression), as you are invited to check. Recall the intuition for these payments: charging a bidder her externality forces her to care about others' welfare and aligns her own objective (quasilinear utility) with the collective objective (social welfare).

One quick sanity check: VCG prices are always nonnegative. This is because one option for the first term in (6) is to set $\omega = \omega^*$, in which case the expression in (6) is 0; maximizing over $\omega \in \Omega$ can only result in a larger number.

Theorem 3.1 *The VCG mechanism is truthful, meaning that bidding truthfully is a dominant strategy for every bidder.*

Proof: Fix i and bids $\mathbf{b}_{-i}$ for the other bidders. When the chosen outcome is ω^*, i's utility is

$$v_i(\omega^*) - p_i = \underbrace{\left[v_i(\omega^*) + \sum_{j \neq i} b_j(\omega^*) \right]}_{\text{(A)}} - \underbrace{\left[\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega) \right]}_{\text{(B)}}. \tag{7}$$

Observe that the term (B) is always the same number, independent of what b_i is. Thus, the problem of maximizing agent i's payoff reduces to the problem of maximizing the only term under her control, the first term (A). As a thought experiment, let's suppose agent i has the power to choose the outcome ω^* directly, rather than merely influencing the chosen outcome indirectly via her choice of bid b_i. Agent i would, of course, use this extra power to choose an outcome that maximizes the term (A). If agent i bids truthfully and sets $b_i(\omega) = v_i(\omega)$ for every $\omega \in \Omega$, then the term (5) that the mechanism maximizes becomes identical to the term (A) that the agent wants maximized. Thus, bidding truthfully results in the mechanism choosing the best-case scenario for bidder i; no other bid could be better. ■

Here is an alternative interpretation of the payments in the VCG mechanism. Rewrite the expression in (6) as

$$p_i = \underbrace{b_i(\omega^*)}_{\text{bid}} - \underbrace{\left[\sum_{j=1}^{n} b_j(\omega^*) - \left(\max_{\omega \in \Omega} \sum_{j \neq i} b_j(\omega)\right)\right]}_{\text{rebate}}. \tag{8}$$

We can thus think of agent i's payment as her bid minus a "rebate," equal to the increase in social welfare attributable to i's presence. For example, in the Vickrey auction, the highest bidder pays her bid b_1 minus a rebate of $b_1 - b_2$ (where b_2 is the second-highest bid), the increase in welfare that the bidder brings to the table.

The discount in (8) is always nonnegative, implying that $p_i(\omega^*) \leq b_i(\omega^*)$ and hence the VCG mechanism is individually rational (i.e., truthtelling agents are guaranteed nonnegative utility). See Exercise Set #8 for details.[7]

4 Some Practical Issues

There are a number of obstacles to implementing the VCG mechanism is practice.[8] We next mention two of these.

A major challenge in any implementation of the VCG mechanism is the design of a user interface for bidding. There are usually a huge (e.g., exponential) number of possible outcomes ω, in which case it is impossible to elicit a bid for each of them. For example, in sponsored search auctions, a bidder only submits a bid per-click.[9] This induces a bid for every possible outcome, namely the bid-per-click times the expected number of clicks the bidders receives in a given outcome. This user interface effectively assumes that bidders care only about clicks, with everything else being of second-order importance.

In general, bidding interface design aims for a tricky balance between expressivity and simplicity. The ideal sweet spot gives bidders enough flexibility to express their preferences, while being simple enough to use easily and implement efficiently. For example, to better accommodate a variety of advertisers, Facebook wanted to allow a richer set of bids than just a bid-per-click. (But an outcome-by-outcome bid is still impossible, of course.) Their choice of a sweet spot was an interface where a bidder could pick which events they care about and her value for each event (and a daily budget).

A second challenge in implementing more complex versions of the VCG mechanism stems from its computational requirements. For example, Facebook has to run an auction every time a user accesses her news feed, which might be over a billion auctions in a single day. Thus these auctions had better be fast! So how much computation is required to implement the VCG mechanism?

[7]We're assuming here that bids are restricted to be nonnegative.

[8]Indeed, until the rise of the online advertising industry, the VCG mechanism was generally thought to be of academic interest only [6].

[9]And also a daily budget, since this bid will be automatically entered into a possibly large number of auctions.

The second step of the VCG mechanism has to solve the social welfare maximization problem (to pick $\omega^* \in \text{argmax}_\omega \sum_{i=1}^n b_i(\omega)$). The third step of the VCG mechanism has to solve n more social welfare maximization problems (to compute the second term in (6) for each i). In sponsored search, where social welfare maximization reduces to sorting, solving these $n+1$ problems is not a big deal. But for Facebook's problem, where there are many possible page layouts and different probabilities of different events, the social welfare maximization problem is non-trivial. Rumor has it that Facebook does not have enough time to always solve VCG's social maximization problems exactly, and for this reason resorts to heuristics.

Things can go wrong with the VCG mechanism when you substitute an inexact heuristic for exact welfare maximization [4, 5]. For example, the mechanism is no longer truthful (Exercise Set #8). This might not be too big of a deal in practice, since it's usually difficult for a bidder to figure out a profitable non-truthful deviation. A more serious problem is that a heuristic might lead to negative payments (i.e., payments from the mechanism to the bidders)! See Exercise Set #8 for details. Presumably Facebook's implementation of the VCG mechanism replaces negative payments with zero (or small positive) payments.

References

[1] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.

[2] 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.

[3] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.

[4] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[5] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35(1–2):166–196, 2001.

[6] M. Rothkopf, T. Teisberg, and E. Kahn. Why are Vickrey auctions rare? *Journal of Political Economy*, 98(1):94–109, 1990.

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

[8] H. R. Varian and C. Harris. The VCG auction in theory and practice. *American Economic Review*, 104(5):442–445, 2014.

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