CS269I: Incentives in Computer Science
Lecture #14: More on Auctions*

Tim Roughgarden†

November 9, 2016

1 First-Price Auction

Last lecture we ran an experiment demonstrating that first-price auctions are not so easy to reason about. (In a first-price single-item auction, the highest bidder wins and pays her bid.) Does theory give us any useful advice about how to bid?

1.1 The Two-Bidder Case

To begin, let's consider the case of $n = 2$ bidders, with private valuations v_1, v_2 drawn independently and uniformly from $[0, 1]$.[1]

We claim that if your opponent always bids half her valuation—that is, uses the bidding strategy $b_2(v_2) = v_2/2$—then your best response is to bid half *your* value. Thus the bidding strategies $b_i(v_i) = v_i/2$ for $i = 1, 2$ constitute an equilibrium.[2,3]

To prove the claim, fix an arbitrary valuation $v_1 \in [0, 1]$. Since v_2 is always at most 1, $b_2(v_2)$ is always at most $\frac{1}{2}$, and there is no reason to bid higher than $\frac{1}{2}$ (recall this is a first-price auction, where you have to pay your bid). The expected utility of a bid $b_1 \in [0, \frac{1}{2}]$ is

$$\underbrace{(v_1 - b_1)}_{\text{utility of winning}} \cdot \mathbf{Pr}[\text{win with bid } b_1]\,. \tag{1}$$

Note that the probability is over the randomness in v_2. Raising one's bid decreases the first term (utility of winning) but increases the second term (likelihood of winning). So what's

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

[1]The valuations in last lecture's experiment weren't exactly uniformly distributed, but they were pretty close.

[2]In this context, called a *Bayes-Nash equilibrium*.

[3]It turns out to be the unique equilibrium, but we won't prove uniqueness.

the sweet spot? With the bid $b_1 \in [0, \frac{1}{2}]$ still fixed, we have

$$\mathbf{Pr}[\text{win}] = \mathbf{Pr}[b_1 \geq b_2] = \mathbf{Pr}[b_1 \geq v_2/2] = \mathbf{Pr}[2b_1 \geq v_2] = 2b_1, \tag{2}$$

where the final equation follows from the fact that v_2 is distributed uniformly in $[0, 1]$. Thus the expected utility of the bid b_1 is

$$(v_1 - b_1)2b_1 = 2b_1v_1 - 2b_1^2.$$

This is a strictly concave function of b_1 (the second derivative is negative) and so it has a unique maximizer, which we can identify by taking the derivative, setting it to zero and solving. This results in $b_1 = v_1/2$. So as claimed, the best response is to bid half your valuation.

Does this mean that you should have bid half your valuation in the experiment? Not immediately, since the claim above only specifies your best response when your opponent is bidding half her value. But actually, if your opponent is using *any* scaling strategy of the form $b_2(v_2) = \alpha v_2$, then your best response is to bid half your value (modulo some details, see Exercise Set #7). So yes, you probably should have bid around half your valuation in the experiment. Only a couple students did this; most shaded their bid by only 25%-35%. This is consistent with other first-price auction experiments, where most bidders bid much more aggressively than in an equilibrium.

It's worth noting that, at equilibrium, the bidder with the highest bid is also the bidder with the highest valuation (since everyone scales their valuations by the same factor). This means that, at equilibrium, the first-price auction is welfare maximizing—the item is always awarded to the bidder with the highest valuation (like in a Vickrey auction).

1.2 The Three-Bidder Case

What about with $n = 3$ bidders, again with valuations drawn independently and uniformly from $[0, 1]$? Here, we claim that the bidding strategies $b_i(v_i) = \frac{2}{3}v_i$ constitute an equilibrium (which is again unique). To see this, adopt bidder 1's perspective and assume that the other two bidders are bidding as above. The expected utility of a bid b_1 is again the expression in (1). But the probability of winning (2) needs to be replaced by the derivation

$$\mathbf{Pr}[\text{win}] = \mathbf{Pr}[b_1 \geq b_2 \wedge b_1 \geq b_3] = \mathbf{Pr}[b_1 \geq b_2] \cdot \mathbf{Pr}[b_1 \geq b_3] = \mathbf{Pr}\left[\tfrac{3}{2}b_1 \geq v_2\right] \cdot \mathbf{Pr}\left[\tfrac{3}{2}b_1 \geq v_3\right] = \left(\frac{3}{2}b_1\right)^2.$$

Thus the winning probability is now a quadratic function of the bid, rather than a linear function. This changes the best response: the expected utility of bid b_1 is $\frac{9}{4}(v_1 - b_1)b_1^2$; taking the derivative, setting it to zero, and solving shows that the expression is maximized when $b_1 = \frac{2}{3}v_1$, as claimed.[4]

[4] In the experiment, there was roughly a 50/50 split between students who used the same bid in both the 2-person and 3-person auctions and those who used different bids. Of those who used different bids, most (but not all) of them bid more aggressively in the 3-bidder case than in the 2-bidder case.

1.3 Generalizations

You can probably guess the pattern for the equilibrium when there are n bidders, all with independent and uniform valuations: $b_i(v_i) = \frac{n-1}{n} v_i$ for each i. Thus, the more competition in the auction, the more aggressively you should bid.

Even more generally, suppose there are n bidders with valuations drawn i.i.d. from some distribution D (not necessarily uniform). Then we still have a clean understanding of the equilibrium: $b_i(v_i)$ equals the expected value of the second-highest valuation, conditioned on the event that v_i is the highest valuation. (You can check that this gives the $\frac{n-1}{n} v_i$ formula for the uniform case.) Even at this level of generality, it remains true the first-price auction is welfare maximizing at equilibrium.

If different bidders have valuations drawn from different distributions, then first-price auctions are a mess to analyze.[5] They are also no longer welfare maximizing—at equilibrium, there can be cases where the highest bidder (i.e., the winner) does not have the highest valuation.

1.4 Revenue Equivalence

We won't really think about the revenue of auctions until Lecture #16. But while we're talking about first- and second-price auctions, it's worth noting an interesting fact that is a special case of a much more general phenomenon.

Let's go back to the case of two bidders with independent and uniformly distributed valuations. Which earns more expected revenue, a first- or second-price auction? (The expectation is over the random choice of the valuations v_1, v_2.) In a second-price auction, where bidders bid truthfully, the selling price is the second-highest (i.e., the smaller) bid, which equals the smaller valuation. So the expected revenue of the second-price auction (assuming truthful bidding) is

$$\mathbf{E}_{v_1,v_2}[\min\{v_1, v_2\}]\,. \tag{3}$$

In a first-price auction, there are two differences: first, the selling price is the higher bid, not the lower one; second, bidders bid half their valuations, not their true valuations. Thus the expected revenue of a first-price auction (assuming equilibrium bidding) is

$$\mathbf{E}_{v_1,v_2}\left[\max\{\tfrac{v_1}{2}, \tfrac{v_2}{2}\}\right] = \tfrac{1}{2} \cdot \mathbf{E}_{v_1,v_2}[\max\{v_1, v_2\}]\,. \tag{4}$$

So which expression is bigger, (3) or (4)? It turns out that the order statistics of i.i.d. samples from a uniform distribution have some interesting properties. Of course, the expected value of a single draw from $U[0,1]$ is $\frac{1}{2}$—that is, this expectation splits the unit interval into 2 equal-length pieces (Figure 1). It turns out that the same thing is true for any number of samples: the expected order statistics of k samples from $U[0,1]$ (the expected highest, the expected second-highest, etc.) split the unit interval into $k+1$ equal-sized pieces (Figure 1).[6]

[5] In Vickrey's original paper [6] he asked the seemingly innocuous question about what the equilibrium looks like with two bidders with valuations drawn from *different* uniform distributions (i.e., uniform on different supports). The answer came only a half-century later [4]!

[6] Nothing deep here, just some calculus.

In particular, $\mathbf{E}[\min\{v_1, v_2\}] = \frac{1}{3}$ while $\mathbf{E}[\max\{v_1, v_2\}] = \frac{2}{3}$. Thus (3) and (4) are both $\frac{1}{3}$, and the first- and second-price auctions have the same expected revenue! You might have expected the first-price auction to earn more revenue—for each fixed bid vector $\mathbf{b}$, it charges more—but at equilibrium the bidders shade their bids just enough to counteract the higher prices.[7]

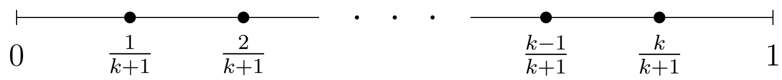

Figure 1: The expected order statistics of k samples split the unit interval into $k+1$ equal-sized pieces.

This turns out to be a special case of a result known as *revenue equivalence*. Revenue equivalence says that whenever two auctions always select the same winner (at equilibrium), then they must have equal expected revenue. For example, we already saw that both the first- and second-price auctions always give the item to the bidder with the highest valuation (at equilibrium), and we verified directly their revenue equivalence.

2 The VCG Auction for Sponsored Search

2.1 Recap

Recall the sponsored search setting introduced last lecture. 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; bidder i has a private valuation v_i for each click its link receives. Each slot j has a *click-through-rate (CTR)* α_j, representing the probability that the end user clicks on this slot. We order the slots so that $\alpha_1 > \alpha_2 > \cdots > \alpha_k$. We also assume that the CTR of a slot is independent of its occupant, but this is easy to relax (see Exercise Set #7). The expected value of an impression (i.e., of being displayed) in slot j for bidder i is then $v_i\alpha_j$.

We saw last lecture that, assuming truthful bids, a simply greedy algorithm maximizes the social welfare $\sum_{i=1}^n v_i\alpha_{s(i)}$, where $s(i)$ denotes the slot assigned to bidder i (if there is no such slot, interpret $\alpha_{s(i)}$ as 0.) Namely, assign the ith highest bidder to the ith slot. This is the analog of the Vickrey auction's decision to always award the item to the highest bidder. But what's the analog of the Vickrey auction's rule that the selling price is the second-highest bid?

Last lecture, we tried the most natural extension, where the occupant of slot i pays (per-click) the next-highest bid. (So the first bidder pays the second bid, the second bidder the third bid, and so on.) This is called the *Generalized Second Price (GSP)* auction, and

[7]Never forget: when you change a system with self-interested users, the response of the users will also change!

we saw that it is not truthful, because bidders sometimes have an incentive to underbid to receive slightly fewer clicks at a deep discount.

2.2 Externalities

There *is* a truthful sponsored search auction. To derive it, we need to think about the Vickrey auction's payment rule in a different way. So here's a thought experiment: consider a single-item auction, and relabel the bidders so that $v_1 > v_2 > \cdots > v_n$.[8] First question: in the outcome of the Vickrey auction, what is the total welfare enjoyed by bidders $2, 3, \ldots, n$? First answer: Zero—all of them lose the item to the highest bidder. Second question: what if we removed bidder #1 and reran the Vickrey auction? Now how much welfare is enjoyed by bidders $2, 3, \ldots, n$? Second answer: v_2, since the item now goes to the highest bidder remaining (bidder 2) and she gets welfare v_2 from it. The difference between the social welfare of the others without and with bidder 1 ($v_2 - 0 = v_2$) is called the *externality* imposed by bidder 1 on the others.[9] Note that the winner's payment in the Vickrey auction is exactly the externality she imposes (v_2). Intuitively, charging a bidder her externality forces her to care about others' welfare and aligns her individual objective (quasilinear utility) with the collective objective (social welfare).

The idea of "charging bidders their externalities" is defined much more generally than just in single-item auctions. Let's go through exactly the same thought experiment with sponsored search. Consider again n bidders, labeled so that $v_1 > v_2 > \cdots > v_n$. Suppose we make the usual greedy assignment, with the ith highest bidder assigned to the ith best slot ($i = 1, 2, \ldots, k$). What's the total welfare of bidders $2, 3, \ldots, n$? Well, each bidder $j = 2, 3, \ldots, k$ get assigned slot j and hence obtains welfare $v_j \alpha_j$, for a total of

$$\sum_{j=2}^{k} v_j \alpha_j. \tag{5}$$

Now suppose we remove bidder 1 and rerun the algorithm that assigns bidders to slot. Now bidder 2 will move up to the first slot, the third bidder to the second slot, and so on, with bidder $k+1$ now getting the last slot. The welfare earned by bidders $2, 3, \ldots, n$ in this case is

$$\sum_{j=2}^{k+1} v_j \alpha_{j-1}. \tag{6}$$

[8]We'll assume no ties for the rest of the lecture, for simplicity.

[9]In general, an "externality" refers to a cost or benefit of an action that is borne by others, rather than by the decision-maker. Like network effects (Lecture #6), externalities can be positive or negative. For example, in selfish routing (Lecture #7), the externalities are negative—when you travel on a route, you also add to the travel time of everyone else on the route. In a public good, like Wikipedia (Lecture #6), the externalities are positive—when you contribute to the public good (e.g., writing an article), most of the benefit is to others.

Subtracting (5) from (6), we get the externality imposed by bidder 1 on the others:

$$\sum_{j=2}^{k+1} v_j(\alpha_{j-1} - \alpha_j), \tag{7}$$

where we're interpreting α_{k+1} as 0. This expression makes sense: 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. Note that this externality is per-impression (with all v_j's getting multiplied by the appropriate CTR), not per-click.

2.3 The VCG Auction

This brings us to the *VCG auction*, which charges each bidder the externality it imposes.[10]

The VCG Auction (Sponsored Search)

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

$$\frac{1}{\alpha_i} \sum_{j=i+1}^{k+1} b_j(\alpha_{j-1} - \alpha_j) \tag{8}$$

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

There are two differences between the payment (8) in the VCG auction and the externality (7). First, each valuation v_j has been replaced by the corresponding bid b_j. This is necessary for the payment to typecheck—the mechanism doesn't know bidders' valuations, only their bids. (Of course if bidders are truthful, then these will be the same.) Second, recall that the externality (7) is per impression, while in a sponsored search auction we can only charge a bidder for a click. Since only an α_i fraction of impressions lead to a click for bidder i, the per-click payment is $\frac{1}{\alpha_i}$ times the per-impression externality. Thus, with truthful bids, the expected payment per impression is exactly the bidder's externality (7), as desired.

The payment (8) probably looks pretty opaque. But notice that

$$\sum_{j=i+1}^{k+1} \frac{\alpha_{j-1} - \alpha_j}{\alpha_i} = \frac{\alpha_i - \alpha_{k+1}}{\alpha_i} = 1,$$

[10] "VCG" here is for Vickrey [6], Clarke [1], and Groves [3].

and thus *bidder i's payment-per-click is a weighted average of lower bids*, with the weights given by the relevant CTRs. We can immediately conclude that, for a fixed bid vector, payments in the VCG auction are no more than in the GSP auction (the latter uses only the biggest of the lower bids). It also follows that a bidder is never charged more than her bid for each click, and so truthful bidding guarantees nonnegative utility ("individual rationality (IR)").

Let's revisit the example from the end of last lecture, with two slots and three bidders, and with $\alpha_1 = .1$, $\alpha_2 = .05$, $v_1 = 10$, $v_2 = 9$, and $v_3 = 6$. The VCG auction assigns the first and second bidders to the first and second slots, respectively. For the second slot, the payment is $\frac{1}{\alpha_2} \cdot \alpha_2 \cdot v_3$, which is 6 (just as it would be in the GSP auction). The payment for slot 1 is the externality imposed by bidder 1, who deprives the second bidder of $(.1 - .05)$ clicks and the third bidder of $(.05 - 0)$ clicks. This yields a payment of

$$\frac{.05 \cdot 9 + .05 \cdot 6}{.1} = 7.5.$$

This is lower than the payment in the GSP auction (which would be 9), reflecting the opportunities available to the bidder to drop to a lower slot. In the GSP auction, the first bidder had an incentive to drop to the second slot (to get utility $.05(10-6) = .2$ rather than $.1(10-9) = .1$), while in the VCG auction, the bidder receives utility $.1(10-7.5) = .25$ from truthful bidding and hence has no incentive to deviate.

In general, the VCG auction has all of the nice properties of the Vickrey auction.

Theorem 2.1 *The VCG auction is truthful, individually rational, and welfare maximizing.*

We already mentioned individual rationality. The auction is welfare maximizing (assuming truthful bids) because we assign the bidder with the ith highest valuation to the ith best slot (recall the correctness of the greedy algorithm from last lecture). The proof of truthfulness is not too difficult, but it makes sense to do it in a more general setting next lecture. So we'll take the truthfulness of the VCG auction on faith for this lecture.

3 GSP vs. VCG

Theorem 2.1 crowns the VCG auction as the rightful heir to the Vickrey auction's throne. But last lecture we said that the Generalized Second Price (GSP) auction has been the dominant paradigm in practice. How can we explain the use of the GSP auction over the VCG auction?

When GSP was originally designed by engineers in 2001, Google did not yet employ any economists, and computer scientists hadn't really studied up on auctions yet. So GSP was developed independently of traditional auction theory. Not long after (in 2002), Google engineers rediscovered the VCG auction and also Google hired its first economist (Hal Varian), who knew all about the VCG auction. So why not switch at that point from GSP to VCG?

The first reason is inertia. GSP was getting pretty great results right out of the gate, so it was hard to justify a major change to it ("don't fix what ain't broke"). Second, the

GSP auction is arguably easier to explain to advertisers than the VCG auction.[11] Third, switching from the GSP auction to the VCG auction would likely result in some revenue loss in the short term. Remember that VCG prices are only lower than GSP prices (for fixed bids), and that bidders potentially shade their bids in the GSP auction (and such shading is indeed observed in practice). Most bidders would learn over time to bid truthfully in the VCG auction eventually, but in the meantime the VCG auction would generate less revenue. It is conceivable that this effect would be large enough to affect a quarterly earnings report, for example. Finally, even the long-term gain of switching to the VCG auction is unclear. It would be simpler for bidders, sure, but it's less clear that there would be any gain in social welfare or revenue. The following result (proved next lecture) makes this last point precise.[12]

Theorem 3.1 ([2, 5]) *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.

Theorem 3.1 does not assert that the equilibrium is unique, nor does it assert that all equilibria are equivalent to the truthful VCG outcome. Our final two examples show that the revenue of the GSP auction at an arbitrary equilibrium can be strictly higher or strictly smaller than that of the truthful VCG outcome.[13]

We continue with our running example with three bidders and two slots, with $\alpha_1 = .1$, $\alpha_2 = .05$, $v_1 = 10$, $v_2 = 9$, and $v_3 = 6$. Recall that the VCG prices for the two slots are 7.5 and 6, respectively. Suppose in the GSP auction, the bidders bid $b_1 = 10$, $b_2 = 7.5$, and $b_3 = 6$. Then, the assignment of bidders to slots and the payments made by the bidders match those in the VCG auction. Moreover, these bids constitute an equilibrium. It's easy to see that the second and third bidders have no profitable unilateral deviations; the only thing to check is that the first bidder doesn't want to drop down to the second slot. Since her current utility is $.1(10 - 7.5) = .25$ and deviating to the second slot would give her utility $.05(10 - 6) = .2$, there is no incentive to switch.[14] The revenue of this equilibrium ($.1 \cdot 7.5 + .05 \cdot 6 = 1.05$) is the same as in the truthful VCG outcome.

On the other hand, suppose the third bidder only bids 5 instead of 6. The result is still an equilibrium; the only change is that when the first bidder drops to the second slot her utility would be $.05(10 - 6) = .2$, the same as her utility in the equilibrium. But the price of the second slot drops from 6 to 5, and the revenue from 1.05 to 1.

Yet another equilibrium has $b_1 = 10$, $b_2 = 8$, and $b_3 = 6$. The first bidder is again indifferent between the first and second slots (both would give utility .2). Because of the

[11] This reason is no longer very compelling, since GSP auctions evolved into a deliberately opaque auction design, for example with "ad quality" factors computed without much transparency.

[12] Despite all of these obstacles to the VCG auction, over the past few years there has been movement from GSP to VCG, primarily due to the latter's flexibility. For example, Facebook's advertising system is based on VCG auctions (more on this next lecture).

[13] The GSP auction can even have equilibria with suboptimal social welfare, where the bidders are not assigned to slots in order of valuation (exercise).

[14] The proof of Theorem 3.1 in Lecture #15 shows how to construct such GSP equilibria in general.

increased price of the first slot, the revenue goes up from 1.05 to 1.1.

The point of these examples is to show that no sweeping revenue comparison of the VCG and GSP auctions is possible. But both theory (in the form of Theorem 3.1) and practice suggest that the GSP auction does pretty well from both the social welfare and revenue standpoints.

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] T. Kaplan and S. Zamir. Asymmetric first-price auctions with uniform distributions: Analytic solutions to the general case. *Economic Theory*, 50(2):269–302, 2012.

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

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