Position Auctions

One application of auctions which has become increasingly important over the last few years is their application to online advertising. Hallerman[1] estimates that online advertising revenues in the US in 2007 were at least $21 billion. [?] estimates that revenues associated with the auction of search terms by industry leaders Google and Yahoo in 2005 were at least $11 billion. [?] provide a less specific and more conservative estimate of $10 billion. Either way, a very large chunk of total advertising revenue on the internet.

Auctions are used to determine which urls are displayed in the "Sponsored Links" column of the google search page, and the "Sponsored Results" column on Yahoo. These auctions are typically referred to as *Position Auctions.* The technology behind these is remarkable. When a search word is entered, a page is returned with a set of search results (sometimes referred to as 'organic links') along with a series of sponsored links. The primary difference between the two is that the sponsored links almost always offer to sell you something. The organic links may or may not involve sales. The sponsored links are useful provided they contain items the person looking at the page is likely to want to buy.

The way the urls are made relevant is to associate with each search term, a set of *positions.*[2] The first position is the place at the top right of the page where the first sponsored link will be placed, the second position is the second highest link that will be displayed, and so on. Advertisers who are interested in a particular search term can then bid for that position in an auction. When you enter your search term, google checks for bids associated with that search term, then puts the link associated with the advertiser who submitted the highest bid in the first position[3], the second

[1]Hallerman, D (2008) "US Online Advertising: Resilient in a Rough Economy", *eMarketer,* March.

[2]It is even more sophisticated than this since Google has quite a bit of information about the viewer who has entered the search term. When the viewer sends his search request to Google, his browser passes a cookie to the google server. The cookie doesn't identify the viewer directly, but it provides an id that can be used to view a database that contains information about past searches that have been made from the same browser. By viewing these, Google can make a good guess about whether the viewer is male or female, rich or poor, young or old, etc. The ip address of the browser can be used to determine where the viewer is located geographically. Finally, Google knows how likely it is that the viewer using the current browser will actually follow a sponsored link. Bids that advertisers submit can be conditioned on all these things.

[3]This isn't quite correct. In fact what they do is to give each advertiser a score based on how much revenue they expect to earn from that advertiser given the bid they have submitted. The winner of the auction is the advertiser whose bid generates the highest expected revenue given his score. We use the simpler high bid auction to make things a bit simpler below.

highest bidder in the second highest position, etc, before sending the page back to you. In other words, every Google page view results in a new auction being held. In this way, we are all involved in a large number of auctions.

The payment that a bidder makes depends on the number of viewers who actually click on the ad. To give some idea of what is involved, Google estimates that bidders who want to have their link displayed in the third position of the auction associated with the search phrase "luxury hotels in London" would have to pay about \$3.04 per click.

Given the enormous number of times Google and Yahoo are used to conduct internet searches each day, even a tiny probability of clicking on a sponsored ad will result in large revenues for the search sites. Tweaking the design of the auction by modifying reserve prices and limiting search slots could make a lot of difference to total revenue. Squeezing revenue is not the only consideration that is important of course. The search engines also want to attract bidders and viewers who will click on their ads. We deal with some of these issues later in the chapter on competition among auctions. This chapter focuses on the case where the search site faces an exogenous set of advertisers bidding for positions whose value to them is also exogenous.

1. Complete Information

To begin, we will simply assume that there are K positions to be auctioned. Each position has a value x_i describing the expected number of clicks associated with that position. Assume $x_1 > x_2 > \cdots > x_K > 0$. The idea is that the slot in the first position is the one that most consumers will click on. For the formalism, think of x_i as the expected number of consumers who will click on the ad in the i^{th} position.

There are $m > K$ firms bidding for the various positions on the web page. Firm i has characteristic v_i that describes its expected profit it receives for any consumer who clicks on the ad. For simplicity, assume that $v_1 \geq v_2 \geq \ldots \geq v_m$. The expected revenue of a firm who successfully places their ad in position j is then $v_i x_j$, the amount they expect to earn from each consumer who click, times the number who are expected to click on their ad. Every firm wants the first position, but not all firms will be willing to pay the same amount for it. All firms, on the other hand, still earn revenue by placing their ads in lower positions on the web page.

The position auction works as follows: each firm places a bid b_i. Lets order the bids so that b_1 is the highest bid, b_2 the second highest, and so on.

The first and most desirable position is awarded to the bidder who submits the highest bid b_1. He or she then pays the second highest bid b_2 for each click. Similarly, the bidder who submits the k^{th} highest bid wins the k^{th} position, and then pays the $k+1^{st}$ highest bid for each click he or she receives. Let the bids be ordered from highest to lowest. If firm i wins the k^{th} best position (with a bid b_k) and pays b_{k+1}, then his or her profit per click is $(v_i - b_{k+1})$. Notice that the firm that wins the K^{th} position (i.e., the lowest or worst position) pays the highest bid of among bidders who did not win a slot. Given some array $b_1, \ldots, b_m$ of bids, bidder i's total profit when he wins the k^{th} slot is

$$(v_i - b_{k+1})\, x_k,$$

while his payoff is zero if he doesn't win a slot.

In the notation above, b_k is the k^{th} highest bid. We want to discuss bids by some bidder i and this bid could have any value b_k - it depends on the bids of the other bidders. It is necessary here to develop a slight bit of new notation.

Lets write the bids submitted by all the players as $b = \{b^1, \ldots, b^m\}$. In the usual way, when we want to focus on bidder i, we can use the notation (b^i, b^{-i}). This is a vector of dimension m, the first component being i's bid, the second component being the bids of the firms other than i. Now turn b_k into a function, and think of $b_k(b)$ to mean the k^{th} highest bid in the array of bids b. So $b_k(b)$ is the k^{th} highest order statistic of whatever vector we use as its argument. For example, $b_1(b)$ is the highest bid submitted by any player, $b_2(b)$ is the second highest bid in the array b, and so on.

However, now we can also write $b_k(b^{-i})$ to be the k^{th} highest bid of the bidders other than i. To deal with ties, lets write $\#_k(b^{-i})$ to mean the number of bids in b^{-i} that are equal to $b_k(b^{-i})$. Also, assume that $x_k = 0$ if $k > K$.

Now we can write the payoff to each player for every array of bids in a position auction where the k^{th} position is awarded to the k^{th} highest bidder who is then charged the $k+1^{st}$ highest bid for each click. I'll take a shortcut here to make the formula somewhat simpler by assuming there are no ties. This will work okay here because we'll only discuss some equilibrium outcomes

$$\pi^i\left(b', b^{-i}\right) = \begin{cases} \left(v_i - b_1\left(b^{-i}\right)\right) x_1 & b' > b_1\left(b^{-i}\right) \\ \left(v_i - b_k\left(b^{-i}\right)\right) x_k & k \leq K; b_k\left(b^{-i}\right) < b' < b_{k-1}\left(b^{-i}\right) \\ 0 & \text{otherwise.} \end{cases} \tag{1.1}$$

This is a complicated expression, so well go over it. When a bidder considers a bid b', she first checks the bids of the other players (i.e., b^{-i}) and puts them in order $\{b_1(b^{-i}), \ldots, b_{m-1}(b^{-i})\}$. She then looks to see if her bid b' is higher than one of the K highest bids that the others have submitted. For example, if there are four bidders (as there will be in the graphical illustration below), and they submit the bids $\{2, 4, 3, 6\}$ (where 2 is submitted by the bidder with identity 1, 4 by bidder with identity 2, etc, then bidder 3 first takes the other bids $\{2, 4, 6\}$ and puts them in descending order. This gives her the vector $b_1(\{2, 4, 6\}) = 6$, $b_2(\{2, 4, 6\}) = 4$ and $b_3(\{2, 4, 6\}) = 2$. Suppose there are $K = 3$ positions to be won in the auction. She sees that there is a bid by bidder 1 , i.e., 2, which is less than the third highest bid (her bid of 3). That means she will win the third highest spot, and pay $b_3(\{2, 4, 6\}) = 2$ for it. Her profit is then the difference between her value and the price 2 multiplied by the average number of clicks associated with slot 3. The math expression above is certainly a more concise way of saying this. It even works when there are more bidders than slots once you think it through.

This profit function isn't complete because it ignores ties. Bidding the same as another player won't normally be a best reply (unless you don't expect to win) since, as we described with auctions, a player who expects to tie with another bidder could raise their payoff by incrementing their bid slightly.

Now that we have the profit function (1.1), we basically know the payoffs that each player will get for every array of bids. That is like knowing what payoffs we put in each cell of the matrix in a bi-matrix game. A Nash equilibrium of this game is an array of m bids such that no player can improve his payoff given the bids of the other players.

We can't really do the gambit thing and check every profile of bids for profitable deviations. There is a continuum of profiles to worry about. Solving games like this usually means making a good guess about what the equilibrium looks like, then trying to verify your guess is an equilibrium. The educated guess here comes from the observation that if you win a slot at all, you won't pay what you bid. Instead, you'll pay what one of the other players bid. So your only real choice is to decide which slot, if any, you want. The others' bids completely determine what you pay when you win it.

The second part of the guess comes from the intuition that it is likely to be the case that the bidder with the highest value will want the best slot, the bidder with the second highest value will want the second best slot, and so on. Why? Well, at this stage in your reasoning you can't really articulate why, but is just seems right. Nash equilibrium is going to give us a conceptual tool that will make it crystal clear why that intuition will work.

So lets suppose that the bidder with the highest value per click submits the highest bid, the bidder with the second highest value per click submits the second highest bid, and so on. The auction would then award the top slot to the bidder with the highest value per click. The bidder who wins the k^{th} highest slot would pay the k^{th} highest bid among the other bidders as we described.

As in a second price auction, a bidder can't change the price they pay for any particular slot, but they can change the slot that they win by outbidding someone else. For example, the bidder who wins the second highest slot could get the best slot by raising his or her bid until it is as large as the bid of the highest bidder. If the bidder decides to do this, then he or she could win the top slot, but to get it they would pay the bid of the high bidder for it according to the rules of the position auction.[4]

So from the perspective of an individual bidder who thinks he knows the bids of the others, he really only needs to choose one of the positions whose 'prices' are all defined by the bids of the others. In particular, a set of bids $(b^1, \ldots, b^m)$ will constitute a Nash equilibrium whenever

$$\pi^i \left(b^i, b^{-i}\right) \geq \pi^i \left(b', b^{-i}\right) \tag{1.2}$$

for each bidder i.

This is still complicated, so I'll show how to construct a set of bids that will support a Nash equilibrium using a diagram and a simple algorithm. We illustrate for the case where there are three slots and 4 bidders.

The horizontal axis in the picture measures the possible values that firms can have for slots. The values of the four firms are marked. The dashed line with the lowest slope (the blue line in the picture) is the graph of the function

$$(v - v_4)\, x_3$$

which describes the payoff to different firm types when they buy slot 3 at price v_4. It is flat because the slope of the line is x_3, the click through rate of the worst available slot. The advantage of drawing this is that it makes it possible to see the profits the higher valued firms would make from this slot as well if they decide they want it. For example, if the bidder with value v_2, who is supposed to bid b^2 in this equilibrium decides to cut his bid below b^3, then he or she will win slot 3 instead of slot 2, so we can read their expected profit as the distance from the point v_2 vertically up to the dashed blue line.

[4]Perhaps you can see why exactly matching the bid of the high bidder (a tie) doesn't make sense. If the second highest bidder finds this deviation profitable, then instead of winning the top slot with probability $\frac{1}{2}$, they could do strictly better by adding a bit to their bid so that they are the highest bidder instead of one of the highest.

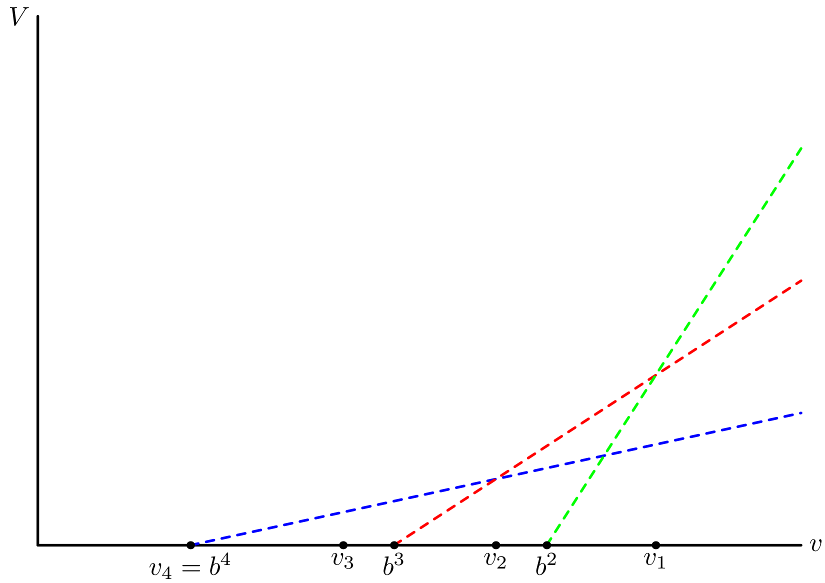

Figure 1.1. Recursive Construction of Slot Prices

I chose b_3 very carefully. It is the bid that the third highest bidder is supposed to make, so it represents payment per click of the second high value bidder. The dashed red line is the graph of the function

$$\left(v - b^3\right) x_2.$$

Since the dashed red and blue lines meet directly above v_2, we have

$$(v_2 - v_4)\, x_3 = \left(v_2 - b^3\right) x_2 \tag{1.3}$$

In other words, bidder 2 will be just indifferent between winning slot 2 at price per click b^3, and cutting her bid below b^3 to win slot 3 at price per click $b^4 = v_4$. Actually, that is exactly how I chose b^3, it is the solution to (1.3).

This argument just shows that bidder 2 doesn't want to lower his or her bid is a way that will give up slot 2 in exchange for slot 3. Bidder 2 could also try to to *raise* her bid to win slot 1. To do that, she will have to match or exceed the bid of the bidder whose value is v_1. This bid is something higher than b^2, which means she will earn no more than

$$\left(v_2 - b^2\right) x_1$$

from this deviation. To understand this one, I have drawn a green dashed line corresponding to the graph of the equation

$$\left(v - b^2\right) x_1.$$

Since v_2 is less than b^2 in the diagram, the expected profit associated with this must be smaller than the vertical distance from v_2 *down* to the extension of the dashed green line - in other words, the deviation involves a loss.

We could now continue and notice that b^2 (which bidder 1 is paying for slot 1) has been chosen so that bidder 1 is just indifferent between winning the top slot and price b^2 and lowering her bid to win the second best slot at price b^3. Setting $b^1$to be anything above b^2 then completes our argument.

One thing to notice about this example, is that each of the bidders (assuming we have bidder 1 bid something higher than v_1) actually bids strictly more than their valuation. In one sense this is completely reasonable since no bidder ever expects to pay their valuation. The reason that is true is that we have modeled this auction as if bidders have complete information - they know each other's value. As we will show later on in this reading, this kind of behavior can't persist when information is incomplete.

Second, observe that the bids we constructed above constitute a Nash equilibrium, but the do not constitute the only Nash equilibrium. There are many more. All we need to do is to make sure that each of the top three firms makes a profit on the slot they win, and that each prefers their slot to any other slot. In the picture, just slide the red line to the left to any position in which firm 3 prefers slot 3 to slot 2 (which just means the blue line has to be higher than the red line at v_3). Similarly, slide the green line to the left lowering p_1 to any position at which the blue line is higher than the green line at v_2. These changes will generate new Nash equilibrium with new prices for the different slots. These prices will be lower for slots 1 and 2.

A second price auction of the kind we are studying here often has the feature that bidders will bid their 'valuations'. This won't necessarily support on equilibrium here as Figure 1 illustrates.

In this Figure, each bidder bids his own value, which would give prices v_4 for slot 3, v_3 for slot 2 and v_2 for slot 1. The it is easy to see from the figure, that at these prices, all three of the top bidders would prefer to win the third slot. The reason is that the higher prices that are supported by the bidders values, aren't warranted given the marginally higher click through rates associated with the better slots. We have no problem constructing an equilibrium for this problem, however, it won't support an outcome where bidders bid their true values.

Exercises.

(1) Unlike a standard auction, an equilibrium in which bidders bid their values may not exist at all. In the following diagram, the colored dash lines are interpreted as the are above and bids are equal to values for all players.

Write down the profitable deviations that are available to each player in this (non-equilibrium) outcome.

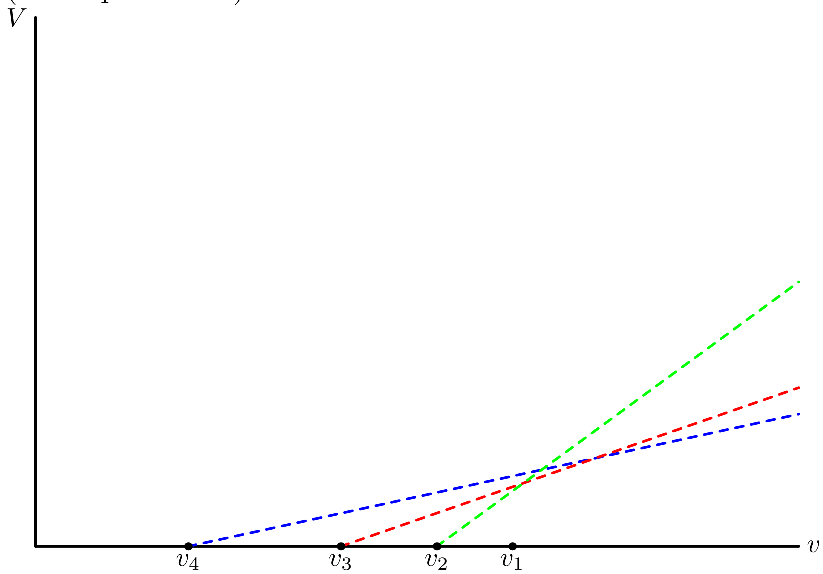

(2) Can you give an algebraic condition (i.e., and inequality) that will ensure that there is a Nash equilibrium for this game in which each bidder bids their value?
(3) Find the two Nash equilibrium in which the sellers expected revenue is highest and lowest
(4) Using the diagrammatic method, show how to compute prices for the case where the bidder with the fourth highest value bids v_3 instead of v_4.
(5) Write down a constrained maximization problem that shows how google would maximize its ad revenues from this action by setting slot prices satisfying (1.2). Can you suggest an algorithmic method to solve this problem? Could you code it?

2. Consumer Search

Apart from the fact that firms bid more than a slot is truly worth to them in the position auctions, there are a number of characteristics of this model that seem unrealistic. First, each firm knows exactly what the values of the other firms are. This is what allows them to confidently bid way more for a slot than it is worth to them. Perhaps more important, the click through rates that the firms are bidding on are unrelated to the characteristics of the firms that are actually making the bids. Most consumers will only click on sponsored ads if they feel that these ads will provide them useful information. Whether they do or not must surely depend on the characteristics they believe the bidding firms have. Secondly, there is a strong connection between click through rates and the position of the ad on the web page. Presumably this reflects the fact that consumers think that ads in a higher slot are more informative. It would be nice to know exactly why this might be.

In a manner that differs a bit from the first section, we'll assume that there is a continuum of consumers. Consumers just want to buy a product. If they succeed in finding the product they want, they get a payoff of 1. When buyers see an ad, they don't know whether or not they will find something they want to buy. Each firm has a *quality* represented by the probability v with which a consumer who visits their site will find something they want to buy (we assume they always buy if they find something they want).

The complication for consumers is two-fold. First, it is costly for them to click on an ad and decide whether they want to buy. The variable s will represent the cost of doing so for some consumer. Second, consumers can't tell the quality of a firm by looking at it's link alone, they actually have to pay the cost s to visit the ad and find out.

One the other side, a firm just wants to sell its product. Assume that if it sells it gets a payoff 1. The probability with which it sells is given by the firm's quality v. The expected payoff when a consumer clicks through an ad link to visit the web page is then v. Firms don't know consumers' search costs.

We'll assume that consumers believe that each firm's quality v is independently drawn from some distribution F which is continuously differentiable with support on $[0, 1]$. Firms believe that the proportion of consumers who have search costs less than s is given by a continuously differentiable function G with domain $[0, 1]$.

Lets examine an example where there are 3 firms bidding on two slots, the highest or top spot, and the lower spot. Slots are auctioned exactly as they were in section 1 - each firm submits a bid for one slot. The highest bid wins the top slot and pays the second highest bid. The second highest bid wins the lower slot and pays the third highest bid. The low bidder doesn't get an advertising slot (or any clicks).

As before, the firms are interested in the click through rates x_1 and x_2 associated with each of the slots. The difference is that these click through rates now represent the proportion of all the consumers who will click through their link to visit their web page.

We are going to find the click through rates as part of the Bayesian Nash equilibrium. However, conditional on the click through rates and bids, the firms' profits are given by (1.2) just as they were in the first section. The complication we'll tackle in the next subsection is that when a firm submits a bid, it can't be sure which slot it will win, or even if it will win a slot at all.

This model also resolves one of the issues associated with the position auction we discussed in the first part of this chapter, since firms won't know exactly the values of their competitors. This helps make bidding strategies of the firms more reasonable.

2.1. **Equilibrium.** We'll use two methods that we've used before. First we'll assume that all three firms use the same bidding rule $\beta(v)$ and that this rule is monotonically increasing - the higher is the firm's quality v, the higher the bid it will submit. Second, we'll use the approach in which a firm *deviates* by bidding *as if* it has a different quality instead of thinking of a different bid. As you'll see, this will get us all the values we want.

If the bidding rule is increasing, then the probability that a firm with value v wins the top slot is

$$F^2(v)$$

which is exactly the same as in the first and second price auctions we looked at previously. The probability that a firm with value v wins the lower slot is

$$2F(v)(1-F(v))$$

Notice that because we are using the assumption that the bidding rule is increasing, we can say this without actually knowing what the equilibrium bidding rule β is. With that said, consumers see the identities of the seller in each slot. We need to condition on this information when we calculate expectations.

The expected quality of the firms in the two slots. In this calculation you can see the name of the firm in the top slot - call it S_1 (for seller 1). The firm in the 2nd slot could be called S_2, while the firm that didn't make it could be called S_3. This happens when S_1 has value $\tilde{v}$, Firm S_2 has value $\tilde{v}' < \tilde{v}$,while S_3 has value $\tilde{v}'' < \tilde{v}'$. The probability that the three sellers have exactly these three values is $f(\tilde{v}) f(\tilde{v}') f(\tilde{v}'')$. The probability that we see S_1 in slot 1, S_2 in slot 2 and $\dot{S}_3$ with no slot is given by summing these probabilities over all the triples of values that satisfy the inequalities.

This is just

$$\int_0^1 \int_0^{\tilde{v}} \int_0^{\tilde{v}'} f(\tilde{v}) f(\tilde{v}') f(\tilde{v}'') \, d\tilde{v}'' d\tilde{v}' d\tilde{v} =$$

$$\int_0^1 \frac{1}{2} F^2(\tilde{v}) f(\tilde{v}) \, d\tilde{v} =$$

$$\frac{1}{6} \int_0^1 dF^3(\tilde{v}) = \frac{1}{6}$$

This must be equal to $\frac{1}{6}$ since there are six possible ways the values of the sellers could have lined up.

What that means is that the density of S_1's quality conditional on the event that S_1 is in the top slot and S_2 is in the second slot and S_3 is in the third slot is

$$\frac{\frac{1}{2}F^2(\tilde{v})\, f(\tilde{v})}{\frac{1}{6}}. \tag{2.1}$$

Using this density, we can find the expected quality of S_1 conditional on the slot allocation to be

$$V_1 = 6\int_0^1 \tilde{v}\frac{1}{2}F^2(\tilde{v})\, f(\tilde{v})\, d\tilde{v}.$$

Similarly, the expected quality of S_2 is

$$V_2 = 6\int_0^1 \tilde{v}F(\tilde{v})\,(1 - F(\tilde{v}))\, f(\tilde{v})\, d\tilde{v}. \tag{2.2}$$

It is completely intuitive that $V_2 < V_1$ as these are expectations of different order statistics.

If you want a proof, notice that the distribution S_2's quality is given by

$$\int_0^v F(\tilde{v})\,(1 - F(\tilde{v}))\, f(\tilde{v})\, d\tilde{v} =$$

$$\int_0^v F(\tilde{v})\, f(\tilde{v})\, d\tilde{v} - \int_0^{\tilde{v}} F^2(\tilde{v})\, f(\tilde{v})\, d\tilde{v} =$$

$$\frac{1}{2}\int_0^v dF^2(\tilde{v}) - \frac{1}{3}\int_0^v dF^3(\tilde{v}) =$$

$$\frac{3}{6}F^2(v) - \frac{2}{6}F^3(v) > \frac{1}{6}F^3(v).$$

What that says is that for every v, the probability that S_2's value is less than or equal to v is strictly larger than the probability that $S_1's$ value is less than or equal to v. That is the definition of first order stochastic dominance - i.e., the distribution of values for S_1 first order stochastically dominates the distribution of values for S_2. One of the most basic properties of first order stochastic dominance is that if one distribution first order stochastically dominates another, then the expectation of any non-decreasing function with respect to the first distribution must be at least as large as the corresponding expectation with respect to the second distribution.

Since our consumers prefer firms with higher values, it now becomes clear why they will always click on the top link first. If consumers expect firms to bid for the top ad using a monotonic (in their value) bidding rule, which is our running assumption, then it is in their own interest to click on the top ad because that is where they are most likely to find the good they want.

Determine the click through rate. Now we can immediately address an important question. Each consumer will click on the ad in the top position if the expected value V_1 is larger than their search cost s. Since s is distributed G, the proportion of consumers who click on the top ad is $G(V_1)$, in other words, if we are right about the monotonic bidding strategy, the click through rate x_1 on the top ad must be $G(V_1)$.

Now the problem turns to another issue that we have already addressed. If a consumer doesn't click on the top ad because their search costs are too high (for example if they just don't have time), then these same consumers certainly won't click on the second ad. The only consumers who click through the link to the second ad, are those who click on the top ad but do not succeed in finding what they want.

Computing the payoff at the second slot when you fail to trade at the first. You might think that at this point we should repeat what we have already done to find the proportion of consumers whose search cost is below V_2 who fail to trade with S_1. This would be incorrect because consumers who fail to find what they wanted with S_1 should be more pessimistic about $\dot{S}_2$, they'll think the expected quality of S_2 is strictly less than V_2.[5]

The reason is that consumers believe that the firm behind the first slot has the highest expected quality. If they can't find what they want at the best firm, it seems less likely that they'll find what they want at the second best firm.

To do this calculation, we need to figure out the expected quality of the firm behind the second ad *conditional* on having failed to find something good at the first firm, and, as before, on the order of the ads we see. We'll first do it using conditional probability as we did above. Then we'll repeat the exercise using Bayes rule so you can see the difference.

We want to find the expected quality of S_2 conditional on clicking on the ad from S_1 but failing to trade (using the distribution conditional on the stores' order). Recall the rule for conditional probability

$$\Pr(A|B) = \frac{\Pr(A \cap B)}{\Pr(B)}.$$

In this expression A and B are *events* - that is, the are just a list of things that happen. Here *event* A is just the event in which a consumer buys if he or she goes to S_2. The event B is the probability with which the consumer failed to trade when he/she visited S_1.

We can build these things up from scratch as before by trying to construct $\Pr(A \cap B)$ from first principles. We need to find the collection of all the profiles of values for the three firms which are consistent with the ranking we see. For each such profile,

[5]Recall that the consumer doesn't directly observe the quality of S_1, he/she can only infer what it was from the subsequent outcome.

we can calculate the probability that the consumer will fail to trade with S_1, but then trade with S_2. Then we can take the expectation of this probability given the conditional distributions we derived above.

If S_1 has value v while S_2 has value $\tilde{v}$, while S_3 has a value below $\tilde{v}$ the joint density of the pair $(v, \tilde{v})$ conditional on what the consumer observes - i.e. S_1 has a higher value than S_2 who has a higher value than S_3 is given by

$$6 f(v) f(\tilde{v}) F(\tilde{v})$$

which follows from our description of the conditional probability distribution as described above.

The probability the consumer fails to trade with v then trades with $\tilde{v}$ is

$$(1 - v) \tilde{v}$$

So we just have to take the expectation of this over all the pairs $(v, \tilde{v})$ for which $v > \tilde{v}$:

$$\Pr(A \cap B) = 6 \int_0^1 \int_0^v (1 - v) \tilde{v} F(\tilde{v}) f(\tilde{v}) d\tilde{v} f(v) dv$$

from the calculations we did above.

Event B is the probability that the consumer fails to trade with S_1. The density of the value at the seller who wins the top position conditional on the outcomes for each of the firms was given above by (2.1). Then we probability with which the buyer fails to trade at the seller in the first position is

$$6 \int_0^1 (1 - v) \frac{1}{2} F^2(v) f(v) dv.$$

Now using conditional probability, the probability of trading with S_2 conditional on failing with S_1 is

$$V_2^* = \frac{\int_0^1 \int_0^v (1 - v) \tilde{v} F(\tilde{v}) f(\tilde{v}) d\tilde{v} f(v) dv}{\int_0^1 (1 - v) \frac{1}{2} F^2(v) f(v) dv.}$$

The click through rate for the second slot is equal to the proportion of people of all buyers who fail to trade at S_1 who have costs low enough to warrant the second click. This is

$$G(V_2^*) = G\left(\frac{\int_0^1 \int_0^v (1 - v) \tilde{v} F(\tilde{v}) f(\tilde{v}) d\tilde{v} f(v) dv}{\int_0^1 (1 - \tilde{v}) \frac{1}{2} F^2(\tilde{v}) f(\tilde{v}) d\tilde{v}.}\right).$$

This gives us the click through rate on the second slot as $(1 - V_1) G(V_2^*)$.

Again, notice that we have been able to compute a pair of click through rates

$$x_1 = G(V_1)$$

and

$$x_2 = (1 - V_1)\, G(V_2^*)$$

without having to know the actual bidding rule β that firms are using. All we need to know so far was that this rule is increasing.

2.2. **Bidding Rules.** Now that we have click through rates for each of the slots, we can use exactly the same approach as we used for the standard auctions to try to figure out what the bidding rule is.

The payoff to the firm whose value is v and who bids as if its value were v' is

$$2\int_0^{v'}\int_0^{\tilde{v}} x_1 * (v - \beta(\tilde{v}))\, f(\tilde{v}'')\, d\tilde{v}'' f(\tilde{v})\, d\tilde{v} +$$

$$2\int_{v'}^{1}\int_0^{v'} x_2 * (v - \beta(\tilde{v}))\, f(\tilde{v})\, d\tilde{v} f(\tilde{v}'')\, d\tilde{v}'' =$$

$$2\int_0^{v'} x_1 (v - \beta(\tilde{v}))\, F(\tilde{v})\, f(\tilde{v})\, d\tilde{v} + 2(1 - F(v')) \int_0^{v'} (x_2 (v - \beta(\tilde{v})))\, f(\tilde{v})\, d\tilde{v}. \tag{2.3}$$

Now we can find the bidding rule by using the logic that if the bidding rule β is part of an equilibrium, then it should be optimal for each firm to bid as if their value were equal to their true value v. The first order condition for maximization of (2.3) with respect to v' is:

$$x_1 (v - \beta(v'))\, F(v')\, f(v') + (1 - F(v'))\, (x_2 (v - \beta(v')))\, f(v') =$$

$$f(v') \int_0^{v'} (x_2 (v - \beta(\tilde{v})))\, f(\tilde{v})\, d\tilde{v}.$$

This condition should hold uniformly in v so that in equilibrium

$$v - \beta(v) = \frac{\int_0^v (x_2 (v - \beta(\tilde{v})))\, f(\tilde{v})\, d\tilde{v}}{x_1 F(v) + x_2 (1 - F(v))} \tag{2.4}$$

Notice something interesting about this formula. Unlike the auction with complete information we studied first in this reading, no bidder will ever bid strictly more that v in equilibrium. In the auction with complete information, bidders knew enough about the values of the other bidders that they could submit bids above their valuation without worrying that they would every have to pay more that their value. Here, bidders have much less information, so they can't rule out the possibility that bids above their value will result in them paying too much.

This suggests that the term in the numerator of the expression on the right hand side of (2.4) must be positive. Since the denominator is clearly positive, (2.4) says that bidders will always bid strictly less than their value.

It might seem strange to you that a second price auction leads to bids below value. Though this is indeed a second price auction, it isn't a *Vickery mechanism* - the bidder can manipulate the price he or she pays by trying to manipulate the probability with which it wins the second rather than the first slot.