A Study of Second Price Internet Ad Auctions

Andrew Nahlik

1 Introduction

Internet search engine advertising is becoming increasingly relevant as consumers shift to a digital world. In 2010 the Internet search giant Google generated over 28 billion dollars in ad revenue while its main competitor Yahoo! generated almost 7 billion dollars of revenue. Ad revenue has grown at a tremendous pace and clearly is a major force behind many of the business models of the new web 2.0 movement. Almost all of the revenue on Yahoo! and Google is generated by an auctioning mechanism that generalizes a second-price auction to sell multiple items.

The general idea of search advertising auctions is that advertisers place a single bid on a keywords or a particular combination of keywords such that when a consumer searches for these terms their ad will appear in a special advertising slot on the search results page. This auction operates as a calculation in real time so all bids are submitted by companies in advance and there is much learning by the bidder. There are several different types of auctions that could be used to allocate these slots. The simplest version would be to effectively have a different second-price auction auction for each slot with a very small bidding cost. Bidders would participate in the auction for a single slot with all the standard rules of a second-price auction and then the seller would move on to the next slot to auction off. In this case, when bidders know all the valuations of their opponents, the seller would generate no revenue because bidders would know the valuations of all other bidders and only the bidder with the top valuation for each slot would bid for that slot because they would not want to incur the bidding cost if they were not going to win. Thus the bidder with the top valuation would win the auction with the only bid submitted. By the fact that it is a second price auction, he would pay only the reservation price for the slot. This would happen for the auction of every slot generating no revenue for the seller.

A much more complicated auction is the auction currently used by Google which the literature has called a generalized second price auction. This auction is similar to a second price auction in that bidders pay not according to their own bid but according to the bids of others but is generalized in that the seller auctions off multiple slots at once. Bidders submit a single bid for the maximum amount they are willing to pay for a click on their ad. They cannot chose a single ad position for their ad or specify different valuations for different slots. Their only signaling mechanism for all slots is the single bid they input. The position of their ad is based on only this solitary bid.[1] The seller makes this a second price auction by stating in the rules that the advertiser will not pay more than than the bid that would put them in next lowest slot. Several papers have explored the generalized second-price auction in depth including the work done characterizing the auction by Edelman, Ostrovsky, and Schwraz (2007) and Varian (2007). However, one major assumption made throughout this literature and by the ad companies themselves is that the position of the ad does not matter, only the rate at which people click through the ad's link is the rate that matters.[2]

In this paper, I will investigate the so called generalized second price auctions in relation to the previous literature that has been done. I will discuss the online ad auction market in general and how click-through versus independent valuation is an important distinction in section II. In section III I will set up the model. In section IV I will do a illustrative simulations between bidders. In section V I will conclude.

2 Internet Search Ads

Even in the short history of the Internet, online advertising has changed quite a bit. At first, starting in about 1994, advertisers were charged on a per-impression basis to show consumers their ads. Starting in 1997 a company called Overture implemented the first auctions where they auctioned off ads. In these auctions, bidders bid on the amount they would pay for a click on their ad and when their ad was clicked on this amount was automatically billed to the advertiser. This generalized first price auction was plagued by the fact that bidders had the incentive to misrepresent their bids and change their bids frequently. Those who could change their bids the fastest had the most to gain from this type of auction. Google introduced the innovation of a generalized second price auction in 2002. As many advertisers

[1] Google also takes into account a quality score which will be described later

[2] http://adwords.blogspot.com/2009/08/conversion-rates-dont-vary-much-with-ad.html

are shifting to the online world, these methods of determining prices for advertising and the method in which advertising is being displayed is becoming increasingly important. The current process to be listed in Google's search results is fairly simple.

When a company wants their ad to be listed next to the so called "organic" search results they go through Google's adwords program.[3] Entry costs are practically zero[4]; all you need to do is take the time to set up an account and input specific information. The advertiser chooses the keyword or keywords that it would like to be listed next to and enters them into an online form.[5] For example, if the advertiser is a wedding shop that operates online and/or nationally, they may choose the keywords "wedding" or "wedding dress" to have their ad, which contains a link to their site, displayed when the search query containing those words is entered into Google. They then submit a bid for the maximum amount that they are willing to pay for each "click-through" on their ad. A click-through on an ad occurs when the customer searches for one of the stated keywords, sees the ad, and clicks the link to the advertisers website.

To distribute the text based ad slots that show up to the right of the organic search, Google uses the bids entered by the advertisers and creates an auction for those slots whenever a search containing the keywords is performed. Google uses these single bids and carries out a generalized second price auction to allocate the ads into slots. The moment an advertiser enters their bid their ad can show up in any keyword search and can be included in every future auction to display next to that keyword until their bid is rescinded.[6] In fact all bids, bids changes, and bid deletions are instantaneous and will be used in the auctions immediately. When a bid is clicked on a page the advertiser is not charged their own bid but instead is charged the highest bid that would allocate the bidder to the next slot down.

[3]Organic search results are the results that the search engine actually finds using their algorithms when you search for a keyword and are not paid for.

[4]In fact you can find many coupons and deals that allow you to get your first $100 in adwords credit for free.

[5]I will specifically reference Google but the process for Yahoo! is similar.

[6]There are also options to stop your bidding based on a daily budget but because this is a static model these budgets will never be applicable.

Thus is it similar to a second price auction but it allocates more than one spot. Edelman (2007) proved that when this method is used to auction off more than one slot at a time it was not equivalent to a Vickery Clarke Groves auction and does not have the desirable properties of this type of auction, namely truthtelling.

Google slightly complicates this auction by not ranking the slots purely on the bid per click where the highest bid would win the highest slot. Instead, Google first computes a quality score for each advertiser which takes into account its previous click through rates and other factors to come up with a number that that estimates what the quality of an advertisers ad will be.[7] Google states that this is to provide their search users with a better experience. Google then multiplies each bid by the quality score to get another number which they rank the bids by. It then charges an advertiser the bid amount that would move it down 1 slot in the rankings instead of the more transparent simplified generalized second price auction. This allows Google to increase their revenue over the less complicated method of just ranking by bid.

The key assumption made in both Google and Yahoo!'s auctions is that the only thing that matters to the advertiser is the click-through. Advertisers only pay when a link is clicked. This assumption ignores two important considerations. First impressions could matter even if there was no click. For example, if a consumer was searching for a widget on Google and a smaller widget company had the number one slot it could be making an impression in that consumers mind. The next time the consumer goes to buy widgets he remembers the small company and makes a purchase. The search engine is not paid for a click-through but clearly that impression was valuable to the advertiser because it generated sales. Second as shown in Agarwal (2011) not all clicks are created equal. The value of a click may differ depending on the position of the ad containing that link. For example assume there are two types of shoppers who click a link for a product, determined shoppers and casual shoppers. Casual shoppers will only click on the first link and if they do not find what they want they will give up. Determined shoppers will click on the first link, and if

[7] Google does not release how these scores are determined.

they do not find what they want, they are more patient and also click on the second link. Assume there are 100 shoppers of each type. Casual shoppers are mostly browsing so they only buy 20% of the time. Determined shoppers are more likely to buy and they click on the first link 50% of the time, and given that they do not buy from the first link, they click on the second link 40% of the time. In this simple example, the first position will generate 70 sales off of 200 clicks (all types have clicked the first link. The second position will only generate 20 sales, but they only have to pay for 50 clicks. The first position has generated 0.35 sales per click while the second position has generated 0.4 sales per click. Clearly the value of a click is worth more in the second position.[8] Not only do the value of clicks depend on position but they also will vary by advertiser. Advertisers will certainly have different values for the sales that are generated per click or the value of a sale to their company. This will make the value of a click through depend on both the advertiser and the position.

In the context of this paper we will assume that a company wishing to place an ad will have a different valuation for each slot that they have predetermined. Of course the companies will take into account the estimated click-through rate for that position and how much these click-throughs are worth when they come up with their valuation. However, just as importantly, they may care about getting the first slot for the impression value rather than the direct click value. In fact this structure allows a company to come up with any valuation that applies to their specific outlook for what kind of revenue they can earn. If Google is only contracting on click-through rates they are not capturing the true valuation to the companies. If companies significantly differ in the weighting of click-through's versus impressions, and they differ in the value of these types of exposure the auction could easily be inefficient in that it would not be capturing the true values to the advertiser of each ad slot.

[8] Note the profits of the first position are still probably higher because they have more sales, but I am concerned not with the value of the position but the value of the click because that is the restrictive assumption made in the literature.

3 Model

Assume that there are K advertising spots with I bidders who will bid on these slots. Each bidder i has a valuation V_{ik} for each advertising slot k. Also assume for each bidder i slot 1 has the highest value, slot 2 has the second highest value, et cetera. Thus the valuations are ordered such that $V_{i1} \geq V_{i2} \geq \ldots \geq V_{im}$. We will denote the bids similarly with bid for position k by company i denoted B_{ik}. Because we are describing an internet auction for placement of ads, this is an easy assumption to make because one would naturally assume that ads placed higher on a page would have a higher value to the advertiser. We will assume that there is full information so that all advertisers know all the valuations of their competitors. Also assume that we order the bidders by their valuation on the top slot. Thus bidder 1 will have the highest valuation for slot 1 and bidder I will have the lowest valuation for slot 1. Note that this provides no insight into the ordering of bids for any slot besides the first slot. The seller institutes a rule that a company can win at most one.

The full information condition is not an obvious choice but is suited to this application. First, these auctions are repeated games and for some general keywords the auction could occur many thousands of time in a single day. Thus a bidder over time should be able to infer the valuations that other bidders place on particular spots by noting what bids win those slots and being able to change their bids instantaneously. Second, Google actually offers a service that estimates the bid you will have to submit to be in a certain range of spots. It can be found at https://adwords.google.com/select/TrafficEstimatorSandbox. The tool is fairly sophisticated and will give you estimates based on many different variables you input or will determine "optimum" bids for you if you wish.[9] For a given keyword, it will give you an estimated price per click, the estimated ad position that matches that price per click, the corresponding estimated number of clicks per day you will receive from that slot and finally the estimated cost per day that keeps your bids under a maximum price per click

[9] Google allows you to enter a list of keywords, the maximum amount you will pay for a click, daily budget, and targeting. It will even suggest keywords for you if you input your website or tell what it is that you are trying to promote.

and your daily budget under a set limit. Thus we can model this as a full information game because bidders have resources available to determine valuations for other bidders.

Now that we have the model, we would like to develop an auction to distribute the ad positions and associated payments to the advertisers. Because we are interested in the value for a position not for a click the auction will be for a position on a page regardless of whether or not a click occurs.[10] As Milgrom (2007) points out, one of the simplest auctions that could be held to efficiently distribute the ad slots is an individual second-price auction for each advertising slot with a very small cost of bidding. Google could auction slot 1, take bids for the slot, determine a winner (where the winner would be barred from winning any auction for slots further down), and continue to the next auction. In a full information setting, this type of auction will efficiently distribute the ad slots, however, it will result in companies only paying the reservation price for each slot. Furthermore, the search engine will earn no profits. This is easily seen from the fact that for slot 1 everyone knows who has the highest valuation. This company bids its true valuation and because other companies know they will not win they will not bid because of the small bid cost. In a second price auction this means that the company must only pay the reservation price. The auction continues with slot 2 and the same scenario plays out. In the end the companies with the highest valuations for each slot (given that they have not already won a slot) will get their correct slot and the auction is efficient but it will generate no revenue for the search engine.

A more interesting auction would occur if we allowed a second price auction for each individual slot but introduced the concept of "bump-up." Because the valuations are ordered as strictly increasing as the slots get higher on the page, they have the nice property that for any slot the winner of that slot values the positions above at least as high as the current position. We can exploit this fact and create an auction where the bid for a slot k is not only a bid for that exact slot but implicitly it is a bid for each slot above it as well. In other words, the search engine assumes that if an advertiser is bidding for a particular slot they

[10] Clearly the chance that a click does occur and the likelihood that the click leads to revenue would influence the valuation of the position.

can take that bid and apply it to all other higher slots because naturally the advertiser will be indifferent or even prefer to have a higher slot for the same price.

In this auction, if more than one company bids above the reservation price, the seller earns revenue. See this by noting that the winning bidder for any spot will now have to pay the maximum valid bid for any spots below them or any other bid for the slot he bid on because of the implicit bid made by the bump-up condition. In other words the winner for slot k will have to pay the maximum bid from his slot or any of the slots below. Bidders do not bid on a slot if they know they will not win due to small bid costs. This means that the winner of slot k will have to pay the winning bid for the slot below $B_{i(k+1)}$. The winner of the last slot K will still only have to pay the reservation price if there are no other bids for that slot because there are no implicit "bump-up" bids forcing payment from below. When only winners bid on a slot, the search engine receives revenues of the sum of the winning bids from slots 2 through K (because the slot 1 winner pays bid 2, slot 2 pays bid 3 et cetera till the K-1 slot pays bid K and the K slot pays 0).

This auction allows bidders to bid on positions individually however the outcome, under the full information restriction, is the similar to the outcome of a GSP auction. All bidders who will get a slot know which slot they will get and they will bid the amount that gets them that position. Under full information, bidders will know the other bids and therefore it is an easy exercise to bid perfectly to get their known slot. No other bidder will bid on that same slot so the binding payment is the bid of the slot below. So the outcome is exactly one bid per slot where the advertiser pays the bid of the slot beneath him, exactly as it is with a GSP auction.[11] The key difference is the current search engine GSP auctions only take bids and charge on the event of a click. This bump up auction or equivalent new generalized second price auction would take bids and charge on impression.

[11] Clearly imperfect information would make the outcomes of these two auctions very different.

4 Simulation

4.1 Bidders are Truthful

Because values are allowed to change in many different dimensions between positions, there are many ways to simulate the values that advertisers may have. One logical way would be to directly assign values to each advertiser using some probability density function for position value.[12] Another option is modeling the values with an initial value using one probability density function, call it W_{i0}, and a scaling factor d_i, which differs by bidder[13], using a second probability density function:

$$W_{ik} = W_{i0} + d_i \times k$$

Distributions for both the initial values and the scaling factors can be set to any distribution that suits the data. The values are further scaled (as shown in the literature) to follow a Zipf distribution.

$$\text{Click Rate} \propto \frac{1}{k^{\alpha}}$$

Where α is the (positive) Zipf coefficient. Thus the value for bidder i in slot k (with impression value being zero) is

$$V_{ik}(\alpha, k) = \frac{1}{k^{\alpha}} \left[W_{i0} + d_i \times k\right]$$

This simulation determines the outcomes of the auction if all bidders have random valuations for each slot and they truthfully revealed these valuations in the auction. The auction was conducted by taking a fixed number of bidders j and auctioning a fixed number of slots k to them. I first did this for a 1 bidder 1 slot auction (to check for errors) up to a 10 bidder 10 slot auction where the number of bidders equaled the slots. Next, I did a similar simulation but I included one more bidder than the number of slots. So it went from a 2 bidder 1 slot auction up to a 10 bidder 9 slot auction. The simulation gave each bidder k random numbers to serve as their valuation. The random numbers were uniformly distributed from 1 to 1000 (in integer increments). It then sorted these random valuations

[12] Keeping in mind that valuations for a player must be declining as positions increase.

[13] This scaling factor could be any number, positive or negative, and it can be drawn from any distribution.

to conform to the condition that the highest valuation should be for slot 1 and the lowest valuation for slot k.

Next the simulation conducted the actual auction assuming that each bidder truthfully revealed their valuations. Note bidders who know they will not win a slot in this scenario do not bid on that slot because of the small bidding costs. The bidders only bid on the slot that they know they will win.[14] The simulation found the highest bid for slot 1, assigned the slot to that bidder, and removed the bidder from the rest of the auction. It then took the highest bid for slot 2, assigned that slot to that bidder, remove, etc until all the slots were filled. Note that if the maximum bid for a slot was shared by two parties the simulation assigned that slot to a random bidder. It then calculated the profit for the winner of each slot and the bid values for each slot. The profit for each slot was calculated as the difference between the valuation the bidder had for the slot minus the price he had to pay for that slot. In this case the price the bidder paid was the bid for the winner of the slot beneath him because of the bump up rule. I ran this auction 20,000 times for each type of auction and the results follow in Table 1 and Table 2.

The results given in Tables 1 and 3 is the average amount of the bid that won a particular slot on top with the standard deviation in parenthesis below. The results are given where the row heading jXk denotes j bidders and k slots. The results given in Tables 2 and 4 again give the average profit for the winner of each slot and in parenthesis gives the standard deviation. Note these tables depend on all bidders bidding their true valuation for the slot that they win.

These simulations reveal some interesting information. If bidders were to reveal their true valuation (and valuations were uniformly distributed), it turns out that on average the person in the first slot (highest on the page) does not make the most profits. On average the bidder making the most profits is the bidder somewhere in the middle of the pack. This is due to the fact that there is so much competition for these top spots. In other words for a bidder y in to win slot one his valuation V_{y1} (which he will make his bid B_{y1}) must be the

[14] They know this because of the full information condition.

Table 1: Average Bids by Slot Number: Equal Number of Bidders and Slots

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8	Slot 9	Slot 10
10X10	990.6	940.6	857.1	753.2	636.0	511.2	382.5	254.7	134.2	36.7
	(9.8)	(30.8)	(48.1)	(60.8)	(69.6)	(74.3)	(74.5)	(69.1)	(57.2)	(32.8)
9X9	988.3	929.3	831.9	710.4	575.2	432.8	289.5	153.2	42.1	
	(12.0)	(36.3)	(56.6)	(69.7)	(78.6)	(80.8)	(77.3)	(64.8)	(37.5)	
8X8	985.1	913.4	797.3	654.7	496.5	333.7	177.5	48.3		
	(15.1)	(44.4)	(66.7)	(81.6)	(88.5)	(86.7)	(74.2)	(43.2)		
7X7	980.4	890.7	749.5	577.9	392.6	210.8	58.1			
	(19.9)	(54.2)	(81.1)	(96.0)	(98.7)	(87.0)	(51.0)			
6X6	973.5	858.4	679.5	482.4	254.0	70.1				
	(26.3)	(69.8)	(99.7)	(115.4)	(101.5)	(60.9)				
5X5	962.2	808.5	576.9	337.9	88.3					
	(37.0)	(92.5)	(124.3)	(130.8)	(76.7)					
4X4	941.3	722.8	414.7	149.0						
	(55.5)	(125.5)	(149.8)	(120.3)						
3X3	899.6	566.5	167.5							
	(90.7)	(173.8)	(136.6)							
2X2	800.2	264.3								
	(163.3)	(198.9)								
1X1	499.4									
	(289.5)									

Table 2: Average Profits by Slot Number: Equal Number of Bidders and Slots

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8	Slot 9	Slot 10
10X10	50.0	83.5	103.9	117.2	124.8	128.6	127.8	120.5	97.5	36.7
	(29.2)	(44.6)	(55.5)	(63.3)	(67.5)	(69.5)	(69.0)	(64.9)	(54.3)	(32.8)
9X9	59.0	97.4	121.6	135.1	142.5	143.2	136.4	111.1	42.1	
	(34.6)	(52.5)	(64.7)	(71.7)	(75.7)	(76.2)	(73.0)	(61.1)	(37.5)	
8X8	71.8	116.0	142.7	158.2	162.7	156.3	129.1	48.3		
	(42.2)	(62.3)	(75.8)	(83.7)	(86.2)	(83.1)	(70.8)	(43.2)		
7X7	89.7	141.2	171.5	185.3	181.8	152.7	58.1			
	(51.1)	(75.0)	(90.4)	(96.7)	(95.1)	(83.5)	(51.0)			
6X6	115.1	178.9	197.1	228.4	183.9	70.1				
	(66.0)	(93.9)	(109.5)	(116.3)	(97.6)	(60.9)				
5X5	153.6	231.7	238.9	249.6	88.3					
	(87.4)	(118.7)	(135.4)	(128.9)	(76.7)					
4X4	218.5	308.1	265.7	149.0						
	(118.6)	(151.3)	(153.9)	(120.3)						
3X3	333.1	399.0	167.5							
	(168.4)	(184.2)	(136.6)							
2X2	535.8	264.3								
	(220.8)	(198.9)								
1X1	499.4									
	(289.5)									

Table 3: Average Bids by Slot Number: 1 More Bidder Than Slots

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8	Slot 9
10X9	989.4	934.2	841.2	724.6	596.0	459.1	320.8	187.7	71.1
	(11.0)	(34.2)	(52.8)	(66.5)	(75.2)	(77.7)	(75.2)	(65.2)	(44.6)
9X8	986.7	920.1	810.3	675.3	524.1	368.2	215.9	81.4	
	(13.5)	(41.1)	(62.4)	(77.1)	(84.5)	(83.8)	(74.9)	(51.1)	
8X7	983.1	900.9	768.1	605.1	430.3	253.9	96.2		
	(17.3)	(49.7)	(74.3)	(90.1)	(94.4)	(86.0)	(60.7)		
7X6	977.2	873.2	709.2	512.3	305.7	117.5			
	(22.8)	(62.4)	(91.6)	(104.4)	(100.1)	(72.1)			
6X5	968.5	832.5	621.2	377.4	146.5				
	(31.2)	(81.6)	(112.4)	(118.2)	(89.1)				
5X4	952.6	761.5	481.0	190.6					
	(45.6)	(109.8)	(139.4)	(112.4)					
4X3	923.2	638.0	264.0						
	(71.0)	(154.5)	(148.8)						
3X2	857.7	401.5							
	(124.1)	(200.0)							
2X1	666.8								
	(234.7)								

Table 4: Average Profits by Slot Number: 1 More Bidder Than Slots

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Slot 6	Slot 7	Slot 8	Slot 9
10X9	55.2	93.0	116.6	128.6	136.8	138.3	133.1	116.6	71.1
	(32.5)	(49.4)	(61.2)	(68.1)	(72.5)	(73.6)	(70.8)	(62.1)	(44.6)
9X8	66.6	109.8	135.0	151.2	155.9	152.2	134.5	81.4	
	(39.2)	(58.0)	(71.1)	(79.5)	(82.3)	(79.9)	(71.3)	(51.1)	
8X7	82.2	132.8	163.0	174.9	176.3	157.7	96.2		
	(47.0)	(68.8)	(84.2)	(91.2)	(90.6)	(82.1)	(60.7)		
7X6	103.9	164.0	196.8	206.6	188.3	117.5			
	(58.8)	(85.3)	(101.0)	(105.3)	(96.6)	(72.1)			
6X5	136.0	211.3	243.8	231.0	146.5				
	(77.1)	(107.7)	(121.6)	(116.2)	(89.1)				
5X4	191.1	280.5	290.4	190.6					
	(104.1)	(138.4)	(141.2)	(112.4)					
4X3	285.1	374.0	264.0						
	(147.0)	(171.1)	(148.8)						
3X2	456.2	401.5							
	(202.1)	(200.0)							
2X1	666.8								
	(234.7)								

Table 5: Optimal Slot Position By Bidder

	Winner of 1	2	3	4	5	6	7	8	9
Mean	5.88	5.84	6.02	6.38	6.78	7.24	7.71	8.23	9.00
	(2.21)	(2.00)	(1.73)	(1.46)	(1.19)	(0.92)	(0.66)	(0.42)	(0)
Median	6	6	6	7	7	7	8	8	9
Mode	8	8	8	8	8	8	8	8	9

maximum of the valuations V_{i1} for all the other bidders. That is to say

$$V_{y1} \geq V_{i1} \qquad \forall \qquad i.$$

However, to win slot two your valuation must only be better than those bidders still left in the auction. The bid by the bidder who was assigned slot one is no longer a valid bid because he has been removed from the auction by virtue of the fact that you may only win a single slot. This idea can be extended and we see that as it gets later in the auction there is less competition remaining and bidder profits increase because of the decreased competition. This idea breaks down towards the end of the auction because if a bidder is still left near the end of the auction their valuations tend to be lower and thus their profits are lower. Because of these two factors, profits for the winners of the middle slots are on average higher than the profits for the slots on the extremes. This demonstrates that bidders do not want to be truthful.

To look at truth telling more concretely, I also looked at the specific 10x9 auction to see if all other bidders bid truthfully what slot the advertiser would want to win, or to say in another way which slot would maximize the profit of the advertiser. Note that they could only win their slot or lower by the assumptions of the truth telling auction. The results in Table 5 again illustrate that most often bidders do not want to tell the truth, so truth-telling is not an equilibrium.

4.2 Bidders are Strategic

The simulation made the unrealistic assumption that bidders would choose to reveal their true valuation as their bids no matter what profit could be had by misrepresenting their bid values. The results of the simulation show that this may not be the dominant strategy for bidders. Because of the random valuation a bidder is likely to win a particular slot in the auction at the rate of $\frac{1}{k}$ where k is again the number of slots being auctioned. Because of the full information condition, if a bidder is able to see others valuations and knows he is in one of the top slots, he can avoid the competition of the upper slots and misrepresent his top bids so as to increase his profits by dropping down in the order. I will show this in a numerical example using a 3 bidder 3 slot auction. The valuations are given in Table 6.

Table 6: Numerical 3X3 Auction

	Slot 1	Slot 2	Slot 3
Bidder A	95	90	85
Bidder B	90	80	70
Bidder C	70	60	50

If bidders bid their true valuation the results of the auction would be bidder A would win slot 1 with a bid of 95, bidder B would win slot 2 with a bid of 80, and bidder C would win slot 3 with a bid of 50. Note in this case A would pay 80, because B and C would not bid on slot 1 but the bump up rule would make an implicit bid for B of 80. Bidder B would pay 50. Bidder C would not pay anything. In this case the profits for bidder A are 15, bidder B 30 and bidder C 50. Now suppose that A was not truthful and instead bid 49 on the last spot (thus implicitly bidding 49 all three slots). In this case B, would win slot 1 with a bid of 90 (or 80 if he believed A would be truthful and he would not bid on slot 1), C would win slot 2 with a bid of 60 (or 50 respectively) , and A would win slot 3. For the sake of the argument we will say that B and C thought A would bid his true valuation and they would do the same. In this case the profit to B would be 40 because he only pays 50 (C's bid that is being bumped up) for a slot he values at 90. The profit to C would be 11 because he gets the slot he values at 60 for a price of 49. The profit to A is now 85 because he gets the slot he values at 85 without having to pay for it (there is no lower bid). So by misrepresenting his bid, bidder A has greatly increased his profits. The argument is similar for B and C to misrepresent their bids.

Because of the large number of possible strategies for each player and because all strategies depend on the both the valuations and the bids of the other players, as the numbers of players and slots increase it becomes nearly impossible to find generic conditions that result in a stable equilibrium where no one has the incentive to move. It is easier when players have specific valuations but it is still a challenge for large numbers of slots or advertisers. In order to compare directly how bids, revenues, and profits were affected by strategic bidding as opposed to truthful bidding, I developed conditions for a stable equilibrium using 3 bidders and 2 slots, and they are shown in Appendix A.

4.3 Strategic Simulation

I ran a 3 bidder 2 bid auction two ways to compare the results if bidders played strategically to if they only bid their true valuation. First, I ran the auction like our previous auctions where bidders bid their true valuation. Next, I did the same simulation with bidders bidding their strategic bids. Note that when a range of bids was possible by assumption the bidder bid the lowest amount possible. Note also that I made epsilon so small that it did not affect the bids. The results are very informative when compared against each other. First, we will examine the average bids by slot in Table 7. It makes sense that bids would go down because bidders are strategically bidding lower than their true valuations to steal a spot at times and it also makes sense because of the assumption that bidders will bid the low end of a range of possible bids.The corresponding average revenue to the seller is in Table 8. The revenue to the seller drops dramatically because of bidder's incentive to misrepresent their valuations and again the assumption that bidders bid low. Note the revenue for the second slot is always 0 because the third highest bidder will never enter any bids. The winner of

Table 7: Average Bids By Slot

.

	Slot 1	Slot 2
Non-Strategic	859.1	401.9
	(123.6)	(199.7)
Strategic	536.3	178.9
	(178.9)	(153.4)

Table 8: Average Revenue to the Seller By Slot

	Slot 1	Slot 2
Non-Strategic	401.9	0
	(199.7)	(0)
Strategic	182.0	0
	(156.1)	(0)

slot 1 will always have to pay the bid of the winner of slot 2's bid. Next look at the average profits to the winner of a particular slot in Table 9. It is not surprising that when bidders play strategically the profits will go up. To see this even more clearly examine the profits by bidder type in Table 10. It is clear that profits have spiked at the expense of revenue to the seller as bidders now bid strategically. Bidder A is able to achieve the highest gains and bidder B and C only achieve modest gains. This is due to the fact that bidder A is doing most of the strategic positioning to maximize his own profit because he is in the strongest position with the highest valuation. One last interesting comparison is the average slot that each bidder won in Table 11.[15] Note that because now bidder A strategically drops down a slot periodically he goes from always winning the first slot which is efficient to sometimes winning the second slot. Also note that bidder A will always win either slot 1 or slot 2.

[15]Note that a bidder who did not gain a slot was counted for averaging purposes as slot 3.

Table 9: Average Profits By Slot

	Slot 1	Slot 2
Non-Strategic	457.0	225.3
	(201.7)	(175.0)
Strategic	656.9	443.7
	(199.3)	(229.8)

Table 10: Average Profits By Bidder

	Bidder A	Bidder B	Bidder C
Non-Strategic	170.3	172.0	56.5
	(143.7)	(192.7)	(114.8)
Strategic	701.4	310.7	88.5
	(172.0)	(255.1)	(172.9)

bidder B benefits (in terms of slot average) by bidder A dropping down a slot because he takes over the top slot position and thus lowers his average. Bidder C can only win slot 2 and because bidder A sometimes drops down a slot when C could have stolen that slot from B his average slot is raised.

5 Conclusion

Online ad sales are a big business. In all previous work and in fact in the search engine auctions themselves, positions have been assumed to be unimportant except for the fact that higher positions gained more clicks. The only thing that matters in the previous analysis of GSP auctions is the click through rate on an ad. If instead advertisers have different values for the different slots, the method of only allowing a single bid dependant on clicks to signal preferences cannot show the bidders true preferences. There are many reasons to assume

Table 11: Average Slot Number By Bidder

	Bidder A	Bidder B	Bidder C
Non-Strategic	1.00	2.33	2.67
	(0)	(0.467)	(0.470)
Strategic	1.14	2.13	2.73
	(0.344)	(0.624)	(0.444)

that valuations do not solely depend on the click through rate especially as ads are moving towards video and multimedia. Impressions if they have value to an advertiser should be included in their valuation and impressions depend greatly on ad position. Also ad position may contribute to the value of the click through itself. Perhaps, bidders who click through on a sixth position ad are actively searching and are more likely to buy a product. Thus, these clicks are actually worth more than a click on the first slot which could be consumers just casually interested and not looking to buy.

Goggle's second price auction and all analysis looking at online auctions have made this critical assumption. This paper makes it clear that if this is not the case the analysis must change. A simple auction would auction each spot separately but this results in zero revenue to the seller under our conditions. An auction that I describe remedies the zero-revenue problem by adding in the "bump-up" rule but achieves the same outcome as the so called generalized second price auction that is run with bids on impressions not clicks. This rule makes the much less harsh assumption that bidders value a higher slot at least as much as a lower slot. Unlike the other analysis, it allows these valuations to take on any form as long as it satisfies this condition. It then states that if you bid an amount k for slot i the auction will implicitly assume that you are willing to bid this amount on every slot higher (of course you are free to make a higher bid on these slots if you wish). This seemingly sensible rule introduces the same incentives as the GSP auctions to not be truth telling and instead to shave bids in many cases. The auction almost always remains inefficient because of these incentives. More analysis is needed to determine how previous analysis will change because of the loosening of the restrictions and changing of how bids are structured.

References

[1] Agarwal, Ashish, Kartik Hosanagar, and Michael Smith. 2011. "Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets," *Journal of Marketing Research*, Forthcoming.

[2] Ashlagi, Itai, Mark Braverman, Avinatan Hassidim, et al. 2010. "Position Auctions with Budgets: Existence and Uniqueness," *The B.E. Journal of Theoretical Economics*, 10: Article 20.

[3] Athey, Susan and Glenn Ellison. 2011. "Position Auctions with Consumer Search," *The Quarterly Journal of Economics*, 126: 1213-1270.

[4] Cary, Matthew, Aparna Das, Benjamin Edelman, et al. 2008. "On Best-Response Bidding in GSP Auctions," *NBER Working Paper* 13788, Draft: January 25, 2008.

[5] Clarke, Edward H. 1971. "Multipart Pricing of Public Goods," *Public Choice*, 11: 1733.

[6] Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *The American Economic Review*, 97: 242-259.

[7] Edelman, Benjamin and Michael Schwarz. 2010. "Optimal Auction Design and Equilibrium Selection in Sponsored Search Auctions," *The American Economic Review*, 100: 597-602.

[8] Feng, Juan. 2008. "Optimal Mechanism for Selling A Set of Commonly Ranked Objects," *Marketing Science*, 27: 501-512.

[9] Ghose, Anindya and Sha Yang. 2009. "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Market," *Management Science*, 55: 1605-1622.

[10] Groves, Theodore. 1973. "Incentives in Teams," *Econometrica*, 41: 61731.

[11] Kuminov, Danny and Moshe Tennenholtz. 2009. "User Modeling in Position Auctions: Re-Considering the GSP and VCG Mechanisms," *AAMAS*, 1: 273-280.

[12] Levin, Jonathan and Paul Milgrom. 2010. "Online Advertising: Heterogeneity and Conflation in Market Design," *The American Economic Review*, 100: 603-607.

[13] Milgrom, Paul. 2009. "Assignment Messages and Exchanges," *American Economic Journal: Microeconomics*, 1, 95-113.

[14] Milgrom, Paul. 2010. "Simplified mechanisms with an application to sponsored-search auctions," *Games and Economic Behavior*, 70, 62-70.

[15] Naldi, Maurizio, Giuseppe D'Acquisto, and Giuseppe Francesco Italiano. 2010. "The value of location in keyword auctions," *Electronic Commerce Research and Applications*, 9, 160-170

[16] Varian, Hal. 2007. "Position Auctions," *International Journal of Industrial Organization*, 25: 1163-1178.

[17] Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16: 8-37.

A General 3X2 Case

Table 12: General 3X2 Auction

	Slot 1	Slot 2
Bidder A	V_{11}	V_{12}
Bidder B	V_{21}	V_{22}
Bidder C	V_{31}	V_{32}

There are several cases to consider. Note for all cases the bidders are labeled such that bidder A has the highest value for slot 1 bidder B has the second highest value and bidder C has the lowest value for slot one that is $V_{11} \geq V_{21} \geq V_{31}$. Note also I will assume that a bidder will bid only the minimum amount to win the bid

A.1 Case 1: $V_{12} < \min(V_{22}, V_{32})$ and $V_{32} < V_{22}$

Note in this case there is no way that bidder A can outbid either bidder for the second slot so the auction will be efficient. Bids and profits will depend on who wins the second slot first we will examine when $V_{32} < V_{22}$

$$B_{11} \in (\max(V_{31}, V_{21} - V_{22}), \infty) \quad B_{12} \in [0, B_{11}] \implies B_{11} = V_{31} + \varepsilon \quad B_{12} = 0$$

$$B_{21} \in [B_{22}, B_{11}) \quad B_{22} \in (V_{32}, B_{11}) \implies B_{21} = V_{32} + \varepsilon \quad B_{22} = V_{32} + \varepsilon$$

$$\pi_A = V_{11} - B_{21} \quad \pi_B = V_{22} \quad \pi_C = 0 \implies \pi_A = V_{11} - V_{32} - \varepsilon \quad \pi_B = V_{22} \quad \pi_C = 0$$

A.2 Case 2: $V_{12} < \min(V_{22}, V_{32})$ and $V_{22} < V_{32}$

Now it follows similarly as

$$B_{11} \in (V_{21}, \infty) \quad B_{12} \in [0, B_{11}] \implies B_{11} = V_{21} + \varepsilon \quad B_{12} = 0$$

$$B_{31} \in [B_{32}, B_{11}) \quad B_{32} \in (V_{22}, B_{11}) \implies B_{31} = V_{22} + \varepsilon \quad B_{32} = V_{22} + \varepsilon$$

$$\pi_A = V_{11} - B_{31} \quad \pi_B = 0 \quad \pi_C = V_{32} \implies \pi_A = V_{11} - V_{22} - \varepsilon \quad \pi_B = 0 \quad \pi_C = V_{32}$$

Note in all further cases I will assume that $V_{12} < \min(V_{22}, V_{32})$

A.3 Case 3: $V_{11} - V_{32} > V_{12}$ and $V_{32} < V_{22}$

In this case we will get an efficient outcome. Because $V_{11} - V_{32} > V_{12}$ we know that if bidder B bids the minimum amount $V_{32} + \varepsilon$ bidder A does prefer slot 1 over the pure profit of slot 2. The bids will be as follows

$$B_{11} \in (\max(V_{31}, V_{21} - V_{22}), \infty) \quad B_{12} \in [0, B_{11}] \implies B_{11} = V_{31} + \varepsilon \quad B_{12} = 0$$

$$B_{21} \in [B_{22}, V_{11} - V_{12}) \qquad B_{22} \in (V_{32}, V_{11} - V_{12}) \qquad \Longrightarrow \qquad B_{21} = V_{32} + \varepsilon \qquad B_{22} = V_{32} + \varepsilon$$

$$\pi_A = V_{11} - B_{21} \quad \pi_B = V_{22} \quad \pi_C = 0 \qquad \Longrightarrow \qquad \pi_A = V_{11} - V_{32} - \varepsilon \quad \pi_B = V_{22} \quad \pi_C = 0$$

A.4 Case 4: $V_{11} - V_{32} < V_{12}$, $V_{21} - V_{32} > V_{22}$, and $V_{32} < V_{22}$

In this case the auction is not efficient. Just as in our first example bidder A will be able to undercut bidder B to get slot 2 and hence obtain a profit of V_{12}. Note that bidder A does not want a fight so his bid B_{11} must be chosen so that $V_{21} - B_{11} > V_{22} \implies B_{11} < V_{21} - V_{22}$ Therefore, B will be satisfied with slot 1 and obtain a profit of $V_{21} - B_{11}$ which will be greater than V_{22} (so he will not want to undercut) but less than V_{21}.

$$B_{11} \in (V_{32}, V_{21} - V_{22}) \qquad B_{12} \in (V_{32}, B_{11}] \qquad \Longrightarrow \qquad B_{11} = V_{32} + \varepsilon \qquad B_{12} = V_{32} + \varepsilon$$

$$B_{21} \in (\max[B_{11}, V_{31}], \infty) \qquad B_{22} \in [0, B_{21}] \implies B_{21} = V_{31} + \varepsilon \qquad B_{22} = 0$$

$$\pi_A = V_{12} \quad \pi_B = V_{21} - B_{11} \quad \pi_C = 0 \qquad \Longrightarrow \qquad \pi_A = V_{12} \quad \pi_B = V_{21} - V_{32} - \varepsilon \quad \pi_C = 0$$

A.5 Case 5: $V_{11} - V_{32} < V_{12}$, $V_{21} - V_{32} < V_{22}$, and $V_{32} < V_{22}$

In this case as in our second numerical example both bidders will want slot 2. If we again impose a minimum bid precision of 10 cents again both will bid $V_{31} + .1$ and slot 1 will be randomly assigned. Slot 2 bids will not matter as long as they are in the range $[V_{32} + .1, V_{31} + .1]$ the profits can now be written as

$$\pi_A = \frac{1}{2}(V_{11} - V_{31} - .1) + \frac{1}{2}V_{12} \qquad \pi_B = \frac{1}{2}(V_{21} - V_{31} - .1) + \frac{1}{2}V_{22}$$

Because we have the condition that $V_{31} \geq V_{32}$ we can show

$$V_{31} \geq V_{32} \implies V_{31} + .1 > V_{32} \implies V_{11} - V_{31} - .1 < V_{11} - V_{32} \implies V_{12} > V_{11} - V_{31} - .1$$

Similar for bidder two so we know that unambiguously the expected value of the profits is less than the pure profit from winning slot 2 so a bidder would always prefer to outright win slot 2.

However, if the profit from slot 1 is greater than the expected value of profits with a fight that is

$$V_{11} - V_{31} - .1 > \frac{1}{2}V_{12} + \frac{1}{2}V_{11} - \frac{1}{2}V_{31} - .05$$

then A or similarly B would would just prefer winning slot 1 to fighting over slot 2. However if we rearrange this condition we see that

$$V_{12} < V_{11} - V_{31} - .1$$

We proved above that this can never be the case so the profit from conceding will never exceed the profit from fighting so there will always be a fight.

$$B_{11} = V_{31} + \varepsilon \qquad B_{12} \in (V_{32}, B_{11}] \qquad \Longrightarrow \qquad B_{11} = V_{31} + .1 \qquad B_{12} = V_{32} + .1$$

$$B_{21} = V_{31} + \varepsilon \qquad B_{22} \in (V_{32}, B_{21}] \Longrightarrow B_{21} = V_{31} + .1 \qquad B_{22} = V_{32} + .1$$

$$\pi_A = \frac{1}{2}(V_{11} - V_{31} - \varepsilon) + \frac{1}{2}V_{12} \quad \pi_B = \frac{1}{2}(V_{12} - V_{31} - \varepsilon) + \frac{1}{2}V_{22} \quad \pi_C = 0 \Longrightarrow$$

$$\pi_A = \frac{1}{2}(V_{11} - V_{31} - .1) + \frac{1}{2}V_{12} \quad \pi_B = \frac{1}{2}(V_{12} - V_{31} - .1) + \frac{1}{2}V_{22} \quad \pi_C = 0$$

A.6 Case 6: $V_{11} - V_{22} > V_{12}$ and $V_{22} < V_{32}$

In this case we also will get an efficient outcome. Because $V_{11} - V_{22} > V_{12}$ we know that if bidder C bids the minimum amount $V_{22} + \varepsilon$ Bidder A does prefer slot 1 over the pure profit of slot 2. The bids will be as follows

$$B_{11} \in (V_{21}, \infty) \qquad B_{12} \in [0, B_{11}] \qquad \Longrightarrow \qquad B_{11} = V_{21} + \varepsilon \qquad B_{12} = 0$$

$$B_{31} \in [B_{32}, V_{11} - V_{12}) \qquad B_{32} \in (V_{22}, V_{11} - V_{12}) \qquad \Longrightarrow \qquad B_{31} = V_{22} + \varepsilon \qquad B_{32} = V_{22} + \varepsilon$$

$$\pi_A = V_{11} - B_{31} \quad \pi_B = 0 \quad \pi_C = V_{32} \qquad \Longrightarrow \qquad \pi_A = V_{11} - V_{22} - \varepsilon \quad \pi_B = 0 \quad \pi_C = V_{32}$$

A.7 Case 7: $V_{11} - V_{22} < V_{12}$ and $V_{22} < V_{32}$

In this case the outcome is inefficient because bidder A once again takes the pure profit of slot 2 over the reduced profits from slot 1. In this case bidder B will not ever want to undercut because he will never win slot 2 he can only win slot 1. Note that A must ensure that his bid is above bidder C's valuation of the third slot to keep him out.

$$B_{11} \in (V_{32}, V_{21}) \qquad B_{12} \in (V_{32}, B_{11}] \qquad \Longrightarrow \qquad B_{11} = V_{32} + \varepsilon \qquad B_{12} = V_{32} + \varepsilon$$

$$B_{21} \in (\max[B_{11}, V_{31}], \infty) \qquad B_{22} \in [0, B_{21}] \Longrightarrow B_{21} = V_{31} + \varepsilon \qquad B_{22} = 0$$

$$\pi_A = V_{12} \quad \pi_B = V_{21} - B_{11} \quad \pi_C = 0 \qquad \Longrightarrow \qquad \pi_A = V_{12} \quad \pi_B = V_{21} - V_{32} - \varepsilon \quad \pi_C = 0$$

This satisfies all possible cases. Note that all cases when $V_{22} > V_{32}$ are explored in the cases number 3-5 and the cases when $V_{22} < V_{32}$ are explored in cases 6-7. Thus we have exhausted the space. The value of V_{12} in relation to V_{22} and V_{32} is explored in cases 1-2.