

Online Ad Auctions

By HAL R. VARIAN*

Draft: February 16, 2009

I describe how search engines sell ad space using an auction. I analyze advertiser behavior in this context using elementary price theory and derive a simple way to estimate the producer surplus generated by online search advertising. It appears that the estimated value of online advertising tends to be between 2 and 2.3 times advertising expenditures. JEL: D44 (Auctions), D21 (Firm Behavior) Keywords: Auctions, Online advertising

The ad auctions used by major search engines all have a similar structure. Advertisers enter ad text, keywords and bids into the system. When a user sends a query to the search engine, the system finds a set of ads with keywords that match the query and determines which ads to show and where to show them.

When the search results and ads are displayed, the user may click on an ad for further information. In this case, the advertiser pays the search engine an amount determined by the bids of the other competing advertisers.

The expected revenue received by the search engine is the price per click times the expected number of clicks. In general, the search engine would like to sell the most prominent positions—those most likely to receive clicks—to those ads that have the highest expected revenue. To accomplish this, the ads are ranked by bid times expected click-through rates and those ads with the highest expected revenue are shown in the most prominent positions.

It is natural to suppose that users who have positive experiences from ad clicks will increase their propensity to click on ads in the future, and those with negative experiences will tend to decrease their propensity to click in the future. Hence search engines may also consider various measures of “ad quality” in their choice of which ads to display.

I. Auction rules

How do search engines determine which ads are shown, where they are shown, and how much they pay per click? We start with some notation. Let $a = 1, \dots, A$ index advertisers and $s = 1, \dots, S$ index slots. Let (v_a, b_a, p_a) be the value, bid, and price per click of advertiser a for a particular keyword. We assume that the expected clickthrough rate of advertiser a in slot s (z_{as}) can be written as the product of an ad-specific effect (e_a) and a position-specific effect (x_s), so we write $z_{as} = e_a x_s$. Other formulas for predicting clicks could be used, but this one leads to particularly simple results.

Here are the rules of the Generalized Second Price Auction used by the major search engines. (1) Each advertiser a chooses a bid b_a . (2) The advertisers are ordered by bid times predicted clickthrough rate ($b_a e_a$). (3) The price that advertiser a pays for a click is the minimum necessary to retain its position. (4) If there are fewer bidders than slots, the last bidder pays a reserve price r .

* UC Berkeley and Google. hal@ischool.berkeley.edu

Consider a specific auction and identify the advertisers with their slot position. Neglecting ties for the sake of exposition, the rules of the auction imply $b_1e_1 > b_2e_2 > \dots > b_me_m$ where m less than or equal to the number of possible slots. The price paid by advertiser in slot s is the minimum necessary to retain its position so $p_s e_s = b_{s+1} e_{s+1}$ which implies so $p_s = b_{s+1} e_{s+1} / e_s$. The price paid per click by the last advertiser is the reserve price if $m < S$ or determined by the bid of the first omitted advertiser if $m = S$.

II. Equilibrium revenue

We assume that advertisers are interested in maximizing surplus: the value of clicks they receive minus the cost of those clicks. In order to simplify the algebra, we will assume that the advertisers have the same ad quality, so $e_a \equiv 1$ for all advertisers.

In equilibrium, the advertiser in slot $s + 1$ doesn't want to move up to slot s , so $(v_{s+1} - p_{s+1})x_{s+1} \geq (v_{s+1} - p_s)x_s$. Rearranging, we have

$$(1) \quad p_s x_s \geq p_{s+1} x_{s+1} + v_{s+1} (x_s - x_{s+1}).$$

This inequality shows that the cost of slot s must be at least as large as the cost of slot $s + 1$ plus the value of the *incremental* clicks attributable to the higher position. It is advertiser $s + 1$'s value for those clicks that is relevant since that is the bid that the advertiser in slot s must beat.

The price of the last ad shown on the page is either the reserve price or the bid of the first omitted ad. Denoting this price by p_m we solve the recursion in inequality (1) repeatedly to get the inequalities

$$\begin{aligned} p_1 x_1 &\geq v_2(x_1 - x_2) + v_3(x_2 - x_3) + v_4(x_3 - x_4) + \dots + p_m x_m \\ p_2 x_2 &\geq \quad \quad \quad + v_3(x_2 - x_3) + v_4(x_3 - x_4) + \dots + p_m x_m \\ p_3 x_3 &\geq \quad \quad \quad \quad \quad \quad + v_4(x_3 - x_4) + \dots + p_m x_m \end{aligned}$$

Summing, we have an upper bound on total revenue:

$$\sum_s p_s x_s \geq v_2(x_1 - x_2) + 2v_3(x_2 - x_3) + \dots + (m - 1)p_m x_m.$$

In the same way, in equilibrium each advertiser prefers its slot to the slot above it, which yields a lower bound on total revenue

$$\sum_s p_s x_s \leq v_1(x_1 - x_2) + 2v_2(x_2 - x_3) + \dots + (m - 1)p_m x_m.$$

One can think of the advertiser values as being drawn from a distribution. If there are S slots, the ads that are shown are those with the S largest values out of the set of ads available. If there are many advertisers competing for a small number of slots, the upper and lower bounds will be close together, so this simple calculation will essentially determine the auction revenue.

III. VCG Auctions

An alternative auction model that has been suggested is the Vickrey-Clarke-Groves (VCG) mechanism. In this mechanism, each advertiser reports a value r_a and each

pays the cost that it imposes on the other advertisers, using the values reported by the other agents.

Suppose, for example that there are 3 slots and 4 advertisers. When advertiser 1 is present, the other three advertisers receive reported values $r_2x_2 + r_3x_3$. (Since advertiser 4 is not shown, it receives zero.) If advertiser 1 is absent, the other three advertisers would each move up a position, so their reported value would be $r_2x_1 + r_3x_2 + r_4x_3$. The difference between these two amounts is $r_2(x_1 - x_2) + r_3(x_2 - x_3) + r_4x_4$, so this is the required payment by advertiser 1.

It can be shown that the dominant strategy equilibrium in the VCG auction is for each advertiser to report its true value, so advertiser 1's payment becomes $v_2(x_1 - x_2) + v_3(x_2 - x_3) + v_4x_4$. Note that this is the same as the lower bound of the equilibrium payments described above. The same calculations go through for the other bidders, so we see that the revenue for the VCG auction is the same as the lower bound of the price equilibrium described above. This is, in fact, a special case of the more general two-sided matching market studied by Gabrielle Demange and David Gale (1985) and Gabrielle Demange, David Gale and Marilda Sotomayor (1986). See also Alvin Roth and Marilda Sotomayor (1990) for a unified description of work in this area.

It might appear that the VCG auction requires knowledge of the expected number of clicks in each position. However, this is not the case. Consider the following algorithm: each time there is a click on position 1, charge advertiser 1 r_2 , and each time there is a click on position $s > 1$ pay advertiser 1 $r_s - r_{s+1}$. In the 3-advertiser example we are considering, the net payment made by advertiser 1 will be $r_2x_1 - (r_2 - r_3)x_2 - (r_3 - r_4)x_3$. This is simply a regrouping of the terms in the VCG payment expression. The argument readily extends to the other advertisers.¹

A more detailed analysis of the General Second Price and VCG auctions is available in Benjamin Edelman, Michael Ostrovsky and Michael Schwartz (2007) and Hal R. Varian (2007) which also fills in some details in the above arguments that were omitted for expositional purposes.

IV. Bidding behavior

The model examined above assumes that an advertiser can choose its bid on an auction-by-auction basis. In reality, advertisers choose a single bid that will apply to many auctions. Let us suppose that there is some reasonably stable relationship between an advertiser's bid and the number of clicks that it receives during some time period, which we will summarize by $b_a = B_a(x_a)$. (Note the change in notation: in the previous section x_s denoted the number of clicks in position s in a given auction; here x_a denotes the number of clicks received by advertiser a in a given time period.)

We also define the cost function $c_a(x_a)$, which is the cost that advertiser a must pay to receive x_a clicks during a given time period. Of course both the bid function and the cost function depend on the interaction with the other advertisers in the ad auction described above, but we take that behavior as fixed.

In this framework, the advertiser's surplus takes the form $v_ax_a - c_a(x_a)$. We call this expression "surplus" rather than "profit" since profit would generally include fixed costs. Of course, surplus maximization is equivalent to profit-maximization, at least as long as profit is non-negative.

Given a cost curve, the advertiser finds its profit-maximizing number of clicks, which is the point where the value equals marginal cost, just as in conventional price theory. Given the optimal number of clicks, we can determine the average cost per click. The advertiser

¹One problem with this algorithm is that it is vulnerable to clickfraud since an advertiser benefits from clicks on lower positions.

can then use the bid function to determine the bid that yields this desired number of clicks. So once an advertiser knows the cost-per-click and bid per-click function, it can determine its optimal behavior. At the present time, an advertiser can only determine these relationships by experimentation.

V. Advertiser surplus

We can use the relationships defined above to construct a bound on the ratio of aggregate value to aggregate cost, which we will refer to as the “surplus ratio.” Let us suppose that advertiser a is choosing some number of clicks x_a at a cost of c_a and that it contemplates changing its bid so that it receives some smaller number of clicks, \hat{x}_a , for which it would pay a smaller cost \hat{c}_a .

We assume that the surplus at (x_a, c_a) exceeds that at (\hat{x}_a, \hat{c}_a) . Of course, this would be implied by global profit maximization, but that is overly strong for our results. This assumption gives us

$$v_a x_a - c_a \geq v_a \hat{x}_a - \hat{c}_a.$$

Rearranging the inequality, we have

$$v_a \geq \frac{c_a - \hat{c}_a}{x_a - \hat{x}_a}.$$

Multiply by x_a and sum to find

$$\sum_{a=1}^A v_a x_a \geq \sum_{a=1}^A \frac{c_a - \hat{c}_a}{x_a - \hat{x}_a} x_a.$$

This expression gives us a bound on the total advertiser value. It is convenient to normalize both sides by the total cost, $\sum_b c_b$, which gives us our final expression:

$$\frac{\text{value}}{\text{cost}} \geq \sum_{a=1}^A \frac{c_a - \hat{c}_a}{x_a - \hat{x}_a} \frac{x_a}{\sum_b c_b}.$$

The geometry is shown in Figure 1. The isosurplus line through x_a is described by $\pi_a = v_a x_a - c_a$, so $c_a = v_a x_a - \pi_a$. The point (\hat{x}_a, \hat{c}_a) must lie above this line since it has lower surplus by assumption. Hence the chord connecting (x_a, c_a) and (\hat{x}_a, \hat{c}_a) must intersect the vertical axis at a point above $-\pi_a$. This shows that expression (V) gives us a lower bound on the surplus ratio.

The same logic can be used to construct an upper bound on the surplus ratio by identifying a lower-surplus point with a larger number of clicks and cost. It is clear from inspecting the figure that the tightness of these bounds will depend on the curvature of the cost function. For an affine cost function, the bounds will be equal to each other and be equal to the true surplus. For a highly convex function, the bounds will be wider. Note that the argument for the surplus bounds is perfectly general and applies to any competitive industry.

VI. Estimation method and results

We can observe how many clicks advertiser a is getting currently, and how much these clicks cost. The challenge is to determine how many clicks it *would* receive in some




FIGURE 1: BOUND ON SURPLUS.

different position. Here is the algorithm I used. (1) Cut advertiser a 's bid in half. (2) Determine what position and how much it would pay in the ad auction with this lower bid. (3) Estimate how many clicks it would receive at this lower position.

With respect to step (1), it might seem that cutting bids in half is somewhat extreme. However a 50 percent reduction in bids only results in a 12 percent reduction in costs, on average. Furthermore, a small cut in bids leads to small reduction in clicks and costs, so that $\Delta c/\Delta x$ can be quite noisy.

The outcome of step (2) is determined by the auction rules. Step (3) is the only problematic calculation since it requires a predictive model of how clicks on a given ad vary with ad position. However the assumption that the actual number of clicks is an advertiser-specific effect times a position-specific effect makes this easy. Suppose, for example, that moving from position 4 to position 3 tends on average to increase the number of clicks by 20 percent. Then if we observe that a particular advertiser initially receives 100 clicks in position 4, we would estimate that it would receive 120 clicks if it bid enough to move to position 3.

I have carried out these calculations on a proprietary sample of ad auctions and found that total value enjoyed by advertisers is between 2 and 2.3 times their total expenditure. Note that this figure only reflects the value of the paid clicks; advertisers may well receive many more valuable clicks from search results.

One can perform similar calculations grouping advertisers by vertical, country, ad configuration, or a variety of other groupings. One interesting finding is that ad configurations that are “fully sold”—have all the slots occupied—tend to have lower surplus than those that are “undersold.” This is because advertisers must compete for slots in the fully sold case, which tends to push their surplus down.

Note that the bid of an advertiser also provides a lower bound on the value of a click. However, bids are only about 25-30 percent larger than prices, on average, so using bids dramatically understates advertiser value.

VII. Summary

There is a growing literature concerning the design of online ad auctions; see Sébastien Lahaie, David M. Pennock, Amin Saberi and Rakesh V. Vohra (2007) for a recent survey. These auctions have an interesting, yet simple, theoretical structure as well as significant practical importance. As more economic transactions take place online we will likely see other novel pricing mechanisms arise.

REFERENCES

- Demange, Gabrielle and Gale, David.** “The Strategy Structure of Two-sided Matching Markets.” *Econometrica*, July 1985, 53(4), pp. 873–888.
- Demange, Gabrielle; Gale, David and Sotomayor, Marilda.** “Multi-Item Auctions.” *Journal of Political Economy*. August 1986, 94(4), pp. 863–72.
- Edelman, Benjamin, Ostrovsky, Michael and Schwartz, Michael.** “Internet Advertising and the Generalized Second Price Auction.” *American Economic Review*. March 2007, 97(1), pp. 242–259.
- Lahaie, Sébastien, Pennock, David M., Saberi, Amin and Vohra, Rakesh V.** “Sponsored Search Auctions,” in Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay V. Vazirani, eds., *Algorithmic Game Theory*. Cambridge: Cambridge University Press, 2007, pp. 699–716.
- Roth, Alvin and Sotomayor, Marilda.** *Two-Sided Matching*, Cambridge: Cambridge University Press, 1990.
- Varian, Hal R.** “Position Auctions.” *International Journal of Industrial Organization*, December 2007, 25(7), pp. 1163–1178.