American Economic Review: Papers & Proceedings 2014, 104(5): 442–445
http://dx.doi.org/10.1257/aer.104.5.442

MARKET DESIGN FOR AUCTION MARKETS[‡]

The VCG Auction in Theory and Practice[†]

By Hal R. Varian and Christopher Harris*

It is now common to sell online ads using an auction. Auctions are used for search ads by Google and Microsoft, for display ads by DoubleClick and other ad exchanges, and for social network ads by Facebook. However, different auction designs are used in each of these cases. Search ads use a Generalized Second Price (GSP) auction, display ad exchanges generally use a Vickrey (second price) auction, and Facebook uses a Vickrey-Clarke-Groves (VCG) auction.

It turns out that these auctions are all closely related. The VCG auction encompasses the traditional Vickrey auction as a special case. It has the attractive property that bidding the true value is a dominant strategy for all players and the equilibrium revenue should, in theory, be about the same as the GSP auction. However, it also has some drawbacks; see Ausubel and Milgrom (2006) and Rothkopf, Teisberg, and Kahn (1990) for a list of potential issues.

In this note we describe two simple theoretical properties of the VCG ad auction and some of the practical lessons learned in implementing a VCG auction for contextual ads.

I. Search Ad Auctions

In a search ad auction advertisers submit keywords and bids. When the advertiser's keyword matches a user's query, the advertiser enters an auction. The advertiser with the highest bid gets the most prominent slot, the advertiser with the second highest bid gets the second most prominent slot, and so on. (In the actual auction, the bids are adjusted by a "quality score," but we ignore this additional complexity in this exposition.)

II. How the GSP Auction Works

Let v_s be the value of a click to an advertiser in slot $s = 1, \ldots, S$, and let x_s be the clicks (or clickthrough rate) associated with that slot. We assume that the slots have been ordered with the most prominent slots first, so that $x_1 > x_2 > \cdots > x_S$.

The GSP auction produces a price for each slot. These prices must satisfy the revealed preference conditions that an advertiser who purchases slot s prefers that slot to other slots it could have purchased:

$$v_s x_s - p_s x_s \geq v_s x_t - p_t x_t. \tag{1}$$

It turns out that, if these inequalities are satisfied for $t = s + 1$, they are satisfied for all slots. After some manipulation we find the following system of inequalities that characterizes equilibrium prices.

$$\begin{aligned} v_s(x_s - x_{s+1}) + p_{s+1} x_{s+1} &\geq p_s x_s \\ &\geq v_{s+1}(x_s - x_{s+1}) + p_{s+1} x_{s+1}. \end{aligned} \tag{2}$$

We note that these inequalities imply that

$$(v_s - v_t)(x_s - x_t) \geq 0, \tag{3}$$

so that advertisers with higher values get more prominent slots, which shows that the GSP equilibria are efficient.

The same manipulations work in reverse. That is, we can start with an efficient assignment of advertisers to slots, which must satisfy

[‡] *Discussant:* Al Roth, Stanford University.

* Varian: Google MTV-1055, 1600 Amphitheatre Parkway, Mountain View, CA 94043 (e-mail: hal@google.com); Harris: Google MTV-900, 1600 Amphitheatre Parkway, Mountain View, CA 94043 (e-mail: ckharris@google.com). We thank Gagan Aggarwal, Josh Dillon, Adam Juda, and Tim Lupus for helpful discussions on these topics.

[†] Go to http://dx.doi.org/10.1257/aer.104.5.442 to visit the article page for additional materials and author disclosure statement(s).

inequality (3) and show that there must exists prices that satisfy the equilibrium inequalities (2). Thus, this simple position auction has mini versions of the First and Second Welfare Theorems.

There are many prices that satisfy these inequalities, but a particularly interesting equilibrium is the one with minimal revenue, where the right inequalities hold with equality. Writing these conditions out for the three-slot case gives us this system:

$$p_1 x_1 = v_2(x_1 - x_2) + p_2 x_2 \tag{4}$$

$$p_2 x_2 = v_3(x_2 - x_3) + p_3 x_3 \tag{5}$$

$$p_3 x_3 = v_4 x_3. \tag{6}$$

Adding up the payments gives us a lower bound on revenue to the seller of

$$R_L = v_2(x_1 - x_2) + 2v_3(x_2 - x_3) + 3v_4 x_3. \tag{7}$$

We can perform the same sort of manipulations to get an upper bound on revenue:

$$R_U = v_1(x_1 - x_2) + 2v_2(x_2 - x_3) + 3v_3 x_3. \tag{8}$$

III. How the VCG Auction Works

In the VCG auction, each bidder is required to pay the cost their presence imposes on the other bidders, using their stated bids as the value they place on the slots. We denote the bid by the advertiser who occupies slot s by b_s. If advertiser 1 participates in the auction, the stated value received by the other advertisers is $b_2 x_2 + b_3 x_3$. If advertiser 1 does not participate in the auction, the other advertisers all move up one position and so receive $b_2 x_1 + b_3 x_2 + b_4 x_3$. Thus, the "harm" that advertiser 1 imposes on the other advertisers is the difference between these two expressions, $b_2(x_1 - x_2) + b_3(x_2 - x_3) + b_3 x_4$, so this is the amount advertiser 1 is required to pay. It turns out that in the VCG auction, it is optimal for each advertiser to bid its true value per click. Writing out the VCG payments in the three-slot case, we have:

$$p_1 x_1 = v_2(x_1 - x_2) + v_3(x_2 - x_3) + v_4 x_3, \tag{9}$$

$$p_2 x_2 = \qquad + v_3(x_2 - x_3) + v_4 x_3, \tag{10}$$

$$p_3 x_3 = \qquad + v_4 x_3. \tag{11}$$

It is easy to check that this produces the same outcome as the GSP system (4)–(6). Hence, the minimum-revenue GSP equilibrium has the same revenue as the VCG equilibrium, a result noted by Edelman, Ostrovsky, and Schwarz (2007) and Varian (2007), and is a special case of a result derived by Demange and Gale (1985); Demange, Gale, and Sotomayor (1986) in a different context. See Roth and Sotomayor (1990) for a unified treatment.

IV. Broad Match

We said that the ad is eligible for the auction if the user's query matches the advertiser's keyword. But what counts as a match? It turns out that search engines use several types of matches including "exact match" and "broad match." A keyword [dog food] would be an exact match for the query "dog food" but a broad match for the query "pet food."

A single broad match keyword will generally have different values in auctions associated with different queries. Accordingly, we use v_s^q to denote the value of the keyword to the advertiser in slot s, in the auction for query q. We use $\overline{v}_s$ to denote the expected value of slot s across all the broad-match auctions.

Advertisers who choose broad match have to pick a single bid that applies for a whole range of auctions. In the VCG auction, each advertiser can state its average value for a broad-matched visitor to its website and everything works out neatly. The GSP auction can, in general, be quite messy since advertisers can appear in different positions in different auctions. However, if the impact of broad match on advertiser values is small enough so that the ordering of advertisers in all broad-match auctions is the same, then everything works out neatly in the GSP auction as well. If the same advertiser is in the same slot in each auction, then the equilibrium calculation is the same as before, with $\overline{v}_s$ replacing v_s.

To summarize: the VCG auction handles broad match in general, while the GSP auction does so only under rather special circumstances. This makes the VCG auction attractive by comparison.

V. Unknown Clickthrough Rates

It would seem that, in order to compute payments for in the VCG auction, we would need to know the clicks (or clickthrough rates) associated with each position. However, that is not the case. Varian (2009) provided an (overly) brief sketch of how this can be accomplished, but we spell out the argument in greater detail here.

Consider the following algorithm to compute advertiser 1's net payment:

(i) Each time there is a click on position 1, *charge* advertiser 1 the amount b_2.

(ii) Each time there is a click on position $s > 1$, *pay* advertiser 1 the amount $b_s - b_{s+1}$.

At the end of the day there will be x_1 clicks on position 1, which results in a payment *from* advertiser 1 of $b_2 x_1$. There will be x_2 clicks on position 2, resulting in a payment *to* advertiser 1 of $(b_2 - b_3)x_2$. And finally, there will be x_3 clicks on position 3, yielded a payment *to* advertiser 1 of $(b_3 - b_4)x_3$.

The total payment by advertiser 1 is then

$$b_2 x_1 - (b_2 - b_3)x_2 - (b_3 - b_4)x_3,$$

which is simply a rearrangement of the payment in equation (9).

It turns out that each advertiser is still paying the cost it imposes on the other advertisers, just as in the original VCG argument, but now on a click-by-click basis. Suppose a click arrives on position 1. If advertiser 1 is present, the advertiser in position 2 gets no benefit *from that click*. If advertiser 1 is not present, then the advertiser who was in position 2 would now be in position 1, and would get b_2 from that click. Advertiser 3 would get zero on *that click* whether or not advertiser 1 was present.

Now suppose a click arrives on position 2. If advertiser 1 is present, advertiser 2 gets b_2 from that click. If advertiser 1 is not present, advertiser 2 would be in the first slot and advertiser 3 would receive the click that went to the second slot. So advertiser 1's presence has imposed a net benefit of $(b_2 - b_3)$ on the other advertisers.

Finally, if a click arrives on position 3, then advertiser 1's presence yields a benefit of b_3 to advertiser 3. If advertiser 1 were absent, then advertiser 4 would receive that click, so the net benefit that advertiser 1's presence imposes on the other advertisers is $(b_3 - b_4)$.

VI. Implementing the VCG Auction

Google designed the GSP auction in the Fall of 2001 and implemented it in February of 2002. A few months later, Eric Veach, the computer engineer who was the main architect of the original GSP auction, came up with a way to create a truthful auction for clicks and showed it to Hal, who immediately recognized it as a VCG auction.

We thought very seriously about changing the GSP auction to a VCG auction during the summer of 2002. There were three problems: (i) the existing GSP auction was growing very rapidly and required a lot of engineering attention, making it difficult to develop a new auction; (ii) the VCG auction was harder to explain to advertisers; and (iii) the VCG auction required advertisers to raise their bids above those they had become accustomed to in the GSP auction. The combination of these issues led to shelving the VCG auction in 2002.

In 2012, we reconsidered a version of the VCG auction for use with our contextual ads. These are ads that are displayed based on the textual content on the page; for example, pages about dogs might show dog food ads. Contextual ads can be displayed in a variety of formats, but a common format is an "ad block" of four ads, arranged either horizontally or vertically.

The primary reason for considering the VCG auction for contextual ads was that it is (i) flexible and (ii) truthful.

With respect to flexibility, in 2002, the important decisions were how to rank ads and how to price ads and the GSP handled these decisions well. By 2012, there were other treatments that could be applied to ads. One particularly useful ad treatment is known as "dynamic resizing." It turns out that if you have one highly relevant and three so-so ads, you get more total clicks by enlarging the size of the highly relevant ad and showing it alone. Choosing when to do this and how much to charge was quite difficult with the GSP auction but could be handled easily by VCG.

The fact that the dominant bidding strategy in the VCG is truthful is also important. This is because the contextual ads can participate in other auctions that have different rules. In particular, we mentioned above that display ads run through a (traditional) Vickrey auction. When a publisher doesn't have an ad to show, it can request ads in an ad exchange where contextual ads may compete with pure display ads.

Since ad exchanges are often run using a classic Vickrey auction, the dominant strategy is truthtelling. But equilibrium bids in the GSP auctions are generally not truthful. Changing the GSP auction to a VCG auction resolved this inconsistency and enabled the contextual ads to compete on an equal footing with other ads.

Truthful bidding also helps simplify the advertisers' decisions. We mentioned earlier that ads can be shown in a variety of formats, such as a horizontal list or a vertical list. The clickthrough rates for a horizontal list don't vary much from position to position, but can vary quite a bit in a vertical list.

As we have seen the GSP equilibrium bid depends on the advertisers' estimates of these position effects—but they don't know what configurations will actually occur. The VCG solves this neatly, since the advertiser only has to reveal its value per click which is generally independent of position.[1]

This is not to say that VCG (even in its pure form) does not have some problems. It is incentive compatible for the advertisers but not necessarily for the publishers. In fact, as the celebrated Myerson-Satterthwaite theorem shows, there is generally no mechanism that is incentive compatible for both sellers and buyers at the same time. Ausubel and Milgrom (2006) and Rothkopf, Teisberg, and Kahn (1990) describe some other problematic issues, but most of these are not relevant for the particular situation we face. All auction forms have advantages and disadvantages, so choosing the "best" mechanism will involve tradeoffs of one sort or another.

The attractive feature of the VCG auction is that the bids are true structural parameters that do not change as other features of the auction change. This is a consequence of our assumption that the value of a visitor to the advertiser's webpage is constant. In a more general model where the probability of purchase varies depending on auction design, this may not be true. However, it appears to be a good approximation in practice.

[1] Note that if all S slots have the same clickthrough rates, the VCG auction and the GSP auction both reduce to an auction that charges the bid of advertiser $S + 1$.

A. *Implementation*

The design of the Vickrey auction is so elegant, one might hope that it would be relatively easy to implement. Alas, it is not so. There were many edge cases that needed to be dealt with, adding to design complexity. On the other hand, once the system was built, other aspects of the ad auction, such as dynamic resizing, and bid optimization by advertisers became much simpler.

The final system, which rolled out in late 2012, cannot be considered a "pure" Vickrey auction, but it is reasonably close to one, given the design challenges involved. From what we can tell, it seems to be working well.

REFERENCES

Ausubel, Lawrence M., and Paul Milgrom. 2006. "The Lovely but Lonely Vickrey Auction." In *Combinatorial Auctions,* edited by Peter Cramton, Yoav Shoham, and Richard Steinberg, 17–40. Cambridge, MA: MIT Press.

Demange, Gabrielle, and David Gale. 1985. "The Strategy Structure of Two-sided Matching Markets." *Econometrica* 53 (4): 873–88.

Demange, Gabrielle, David Gale, and Marilda Sotomayor. 1986. "Multi-item Auctions." *Journal of Political Economy* 94 (4): 863–72.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second-price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review* 97 (1): 242–59.

Roth, Alvin E., and Marilda Sotomayor. 1990. *Two-sided Matching: A Study in Game-theoretic Modeling and Analysis.* Cambridge, UK: Cambridge University Press.

Rothkopf, Michael H., Thomas J. Teisberg, and Edward P. Kahn. 1990. "Why Are Vickrey Auctions Rare?" *Journal of Political Economy* 98 (1): 94–109.

Varian, Hal R. 2007. "Position Auctions." *International Journal of Industrial Organization* 25 (6): 1163–78.

Varian, Hal R. 2009. "Online Ad Auctions." *American Economic Review* 99 (2): 430–34.

This article has been cited by:

1. Bowei Chen, Jingmin Huang, Yufei Huang, Stefanos Kollias, Shigang Yue. 2020. Combining guaranteed and spot markets in display advertising: Selling guaranteed page views with stochastic demand. *European Journal of Operational Research* **280**:3, 1144-1159. [Crossref]
2. Ayesha Afzaal, Faria Kanwal, Ayesha Haider Ali, Komal Bashir, Fatima Anjum. 2020. Agent-Based Energy Consumption Scheduling for Smart Grids: An Auction-Theoretic Approach. *IEEE Access* **8**, 73780-73790. [Crossref]
3. Argyrios G. Tasiopoulos, Onur Ascigil, Sergi Rene, Michal Krol, Ioannis Psaras, George Pavlou. DEEM: Enabling Microservices via DEvice Edge Markets 1-9. [Crossref]
4. Yang Zhang, Arnob Ghosh, Vaneet Aggarwal, Tian Lan. 2019. Tiered Cloud Storage via Two-Stage, Latency-Aware Bidding. *IEEE Transactions on Network and Service Management* **16**:1, 176-191. [Crossref]
5. Shunit Agmon, Orna Agmon Ben-Yehuda, Assaf Schuster. Preventing Collusion in Cloud Computing Auctions 24-38. [Crossref]
6. Murat Ali Bayir, Mingsen Xu, Yaojia Zhu, Yifan Shi. Genie 465-473. [Crossref]
7. Seohyeon Jeong, Woongsoo Na, Joongheon Kim, Sungrae Cho. 2018. Internet of Things for Smart Manufacturing System: Trust Issues in Resource Allocation. *IEEE Internet of Things Journal* **5**:6, 4418-4427. [Crossref]
8. Nicola Dimitri. 2018. Combinatorial advertising internet auctions. *Electronic Commerce Research and Applications* **32**, 49-56. [Crossref]
9. Amin Sayedi, Kinshuk Jerath, Marjan Baghaie. 2018. Exclusive Placement in Online Advertising. *Marketing Science* **37**:6, 970-986. [Crossref]
10. Alison Watts. 2018. Generalized Second Price Auctions over a Network. *Games* **9**:3, 67. [Crossref]
11. Amin Sayedi. 2018. Real-Time Bidding in Online Display Advertising. *Marketing Science* **37**:4, 553-568. [Crossref]
12. Patrick Hummel. 2018. Hybrid mechanisms for Vickrey–Clarke–Groves and generalized second-price bids. *International Journal of Game Theory* **47**:1, 331-350. [Crossref]
13. Martin Spann, Robert Zeithammer, Marco Bertini, Ernan Haruvy, Sandy D. Jap, Oded Koenigsberg, Vincent Mak, Peter Popkowski Leszczyc, Bernd Skiera, Manoj Thomas. 2018. Beyond Posted Prices: the Past, Present, and Future of Participative Pricing Mechanisms. *Customer Needs and Solutions* **5**:1-2, 121-136. [Crossref]
14. Tobias Grubenmann, Abraham Bernstein, Dmitry Moor, Sven Seuken. Financing the Web of Data with Delayed-Answer Auctions 1033-1042. [Crossref]
15. Ruggiero Cavallo, Maxim Sviridenko, Christopher A. Wilkens. Matching Auctions for Search and Native Ads 663-680. [Crossref]
16. Jean Pouget-Abadie, Vahab Mirrokni, David C. Parkes, Edoardo M. Airoldi. Optimizing Cluster-based Randomized Experiments under Monotonicity 2090-2099. [Crossref]
17. Orcun Karaca, Maryam Kamgarpour. Game theoretic analysis of electricity market auction mechanisms 6211-6216. [Crossref]
18. Rong Ma, Jinbo Xiong, Mingwei Lin, Zhiqiang Yao, Hui Lin, Ayong Ye. Privacy Protection-Oriented Mobile Crowdsensing Analysis Based on Game Theory 990-995. [Crossref]
19. Pier Giuseppe Sessa, Neil Walton, Maryam Kamgarpour. 2017. Exploring the Vickrey-Clarke-Groves Mechanism for Electricity Markets * *This work is partially funded under M. Kamgarpour's European Union ERC Starting Grant CONENE. *IFAC-PapersOnLine* **50**:1, 189-194. [Crossref]

20. Christopher A. Wilkens, Ruggiero Cavallo, Rad Niazadeh. GSP 25-32. [Crossref]
21. Alexey Drutsa. Horizon-Independent Optimal Pricing in Repeated Auctions with Truthful and Strategic Buyers 33-42. [Crossref]
22. Ruggiero Cavallo, Prabhakar Krishnamurthy, Maxim Sviridenko, Christopher A. Wilkens. Sponsored Search Auctions with Rich Ads 43-51. [Crossref]
23. Amin Sayedi. 2017. Real-Time Bidding in Online Display Advertising. *SSRN Electronic Journal* . [Crossref]
24. Tyler Cowen, Alex Tabarrok. 2016. A Skeptical View of the National Science Foundation's Role in Economic Research. *Journal of Economic Perspectives* **30**:3, 235-248. [Abstract] [View PDF article] [PDF with links]
25. Frank Kelly, Peter Key, Neil Walton. 2016. Efficient Advert Assignment. *Operations Research* **64**:4, 822-837. [Crossref]
26. Mingfan Ma, Jun Gao, Ning Lu, Wenbo Shi. A Privacy-Preserving Online Reverse Multi-attributes Auction Scheme Based on Degree-Matching 432-442. [Crossref]
27. Yoram Bachrach, Sofia Ceppi, Ian A. Kash, Peter Key, Mohammad Reza Khani. Mechanism Design for Mixed Bidders 215-225. [Crossref]
28. Nicola Gatti, Alessandro Lazaric, Marco Rocco, Francesco Trovò. 2015. Truthful learning mechanisms for multi-slot sponsored search auctions with externalities. *Artificial Intelligence* **227**, 93-139. [Crossref]
29. D. C. Parkes, M. P. Wellman. 2015. Economic reasoning and artificial intelligence. *Science* **349**:6245, 267-272. [Crossref]