Ad Auction Design and User Experience

Zoë Abrams[1] and Michael Schwarz[2]

[1] Yahoo!, Inc., 2821 Mission College Blvd., Santa Clara, CA, USA, za@yahoo-inc.com
[2] Yahoo! Research, 1950 University Ave., Berkeley, CA, USA, mschwarz@yahoo-inc.com

Abstract. When users click on poor quality advertisements, there is a hidden cost to the search engine due to the user dissatisfaction (for instance, users are less likely to click on ads in the future). We describe how to incorporate hidden costs into the GSP auction for internet ads such that it is in an advertiser's self interest to create a user experience that maximizes efficiency.

1 Introduction

In sponsored search, the behavior of users in the long run is endogenous: users continue to click on advertisements only if on average the value that a user derives from clicking on ads exceeds the cost of time required to click and to evaluate the contents of the offer. Sometimes, the value of a click to a user may be a large negative number (e.g. an unscrupulous advertiser may mislead an unsuspecting user to infect his computer with spyware).

We consider that a user's future propensity to click on ads is influenced by his experience with past clicks. An ad with disappointing quality of landing page imposes a negative externality on the search engine because the future stream of revenue from a user is reduced by some amount (some of the future clicks are lost since a disappointed user may learn to ignore ads). Of course, the externality may also be positive. A good experience with an ad may train users to pay more attention to other ads. We will refer to this externality as a hidden cost. Obviously, if an ad's hidden cost is greater than its bid, a search engine should never show the ad. How should a search engine incorporate hidden costs into an auction mechanism?

Our main contribution is the design of a mechanism that encourages advertisers to create an experience for users that maximizes efficiency.

A classic method for encouraging or discouraging certain behavior towards the social optimum is the Pigovian tax [26]. Pigovian taxes charge (or subsidize) an agent its *externality* from the transaction (i.e. the effect on agents that occurs outside of the transaction). A Pigovian tax causes an agent to internalize the externality that he imposes by his actions on other agents, creating an incentive for individual agents to take actions that maximize social welfare.

The concept of pricing externalities to encourage social welfare is one of the key insights of economic theory [29, 3, 2]. For instance, the celebrated VCG mechanism [31, 8, 16] builds on this idea.

A relevant example of Pigovian taxation is the market for renting space at a shopping mall. That market has much in common with the virtual "real estate" allocated for ads on a search result page. The shopping mall owner, much like a search engine, tends to allocate space to the stores that value space the most. The auctions for sponsored search considered in the academic literature allocate space based on the advertiser's willingness to pay per click, sometimes adjusted by ad clickabilty (the ads are ranked based on the product of the bid per click and the probability the ad is clicked when it is seen). The current mechanism charges higher prices for ads with low clickability but does not charge a higher price for ads with high hidden cost. No model in the literature that we know of incorporates hidden costs: usually advertisement specific information is the expected clickthrough rate [30, 11, 1, 25, 21, 28, 17]. In the case of shopping malls, reputable stores draw in customer traffic to the shopping mall, thus imposing a long term positive externality on other tenants in the mall. These externalities are huge, as apparent from the contracts between stores and mall operators: the rent paid by premium brand stores may be a few times lower than the rent paid by less reputable tenants [15]. This is because the shopping mall uses pricing to attract tenants that will impose a positive externality on other tenants by creating more traffic. Exactly the same logic applies in sponsored search. To foster efficiency, a search engine should encourage ads that give users a positive experience because it makes users more likely to click on other ads, thus enhancing the value of the virtual real estate. However, at this time, the auctions for sponsored search do not explicitly take into account the hidden costs that are created by advertisers and the literature on sponsored search does not address hidden costs. This causes a large inefficiency if the role of hidden costs in sponsored search is comparable to the role of externalities in an off line retail environment. The experience of shopping malls is not fully transferable to the sponsored search environment, because the rental agreements for shopping mall space are negotiated, while the ad space has to be sold via auction due to fluidity and the high volume of small transactions in this market. Also, the nature and volume of user traffic enables search engines to effectively measure the hidden cost of an advertisement. We show how to incorporate hidden costs in the standard pricing mechanism for ad auctions in a manner that fosters efficiency.

In addition to being connected with the literature on externalities, this paper is also related to previous work in the field of marketing. A number of papers study the role of the internet on consumer choices and firms' pricing [9, 4, 10, 22]. It is widely appreciated that information obtained by consumers from the internet plays an important role in their purchasing decisions, see for example, [13]. The internet plays a key role in the marketing approach of advertisers, for example, targeting advertisements to consumer segments within a market is investigated in [18], where it is assumed that the cost of reaching out to consumers in different segments is the same (the present paper shows that it is in a search engine's interests to make it cheaper for the advertiser to reach out to the market segment that derives the greatest consumer surplus from the offer). The marketing research also studies brand loyalty [23], the factors that influence a user's propensity to search and click [19, 24], and the use of clickstream data to increase our understanding of user behavior [14, 5, 12]. Unlike the afor-mentioned papers, which focus on the decisions of consumers and advertisers, this paper explores the decisions available to the search engine: namely, how to acheive efficiency by incorporating the user experience into the determination of which advertisements to show to consumers, and at what price. In [7], a similar question is addressed but in the context of display advertising (as opposed to sponsored search).

2 Hidden Cost GSP

We begin by describing the environment and then proceed to describe an auction mechanism that incorporates hidden costs into pricing. We assume that an advertiser can make choices about his offerings that influence the value per click for the advertiser and the user experience. For instance, an advertiser may choose to make a misleading offer that leads to poor user experience but higher profit or a more honest offer that may be somewhat less profitable. An advertiser may also choose the price of the product that he sells. The lower the price, the better the user experience. We denote the choice of the landing page for advertiser i by q_i (q_i may belong to a finite set, for instance a binary choice of being explicit or not being explicit about the shipping charges, or it can be a continuous variable such as the price of a good or the cost of shipping).

The search engine allocates positions on a search result page to advertisers. We assume that the number of clicks that an ad receives depends only on its position. The number of clicks in position j is α_j. Ads in higher positions receive more clicks so that $\alpha_k \geq \alpha_{k+1}$. The value that an advertiser derives per click is denoted by s_i. The value per click may depend on the advertiser's choice of q and can be viewed as a function $s_i(q_i)$. The private values per click for each advertiser are denoted with the set s, indexed by value so that $s_k \geq s_{k+1}$. Let r_i be the rank of advertiser i (not necessarily the same as their value index), and K be the total number of positions on the screen, where ads can be displayed. The payoff of an advertiser is the number of clicks that an advertiser receives multiplied by the value per click minus the total cost that an advertiser pays for clicks. That is, $advertiser\ payoff = s_i(q_i)\alpha_{r_i} - payment$.

We assume that an advertiser's choice of q imposes an externality on the search engine's long term health. The hidden cost per click is denoted by $h_i = h(q_i)$. The hidden cost for advertiser i can be thought of as the change in future revenues due to the change in a user's propensity to click caused by clicking on an ad published by i. An ad that reduces a user's propensity to click will lead to a positive hidden cost (likewise, an increase in future clicks leads to a negative hidden cost). With this interpretation, a search engine can statistically infer a hidden cost of an ad without examining the contents of an advertiser's offer. In light of this, our model assumes that hidden costs are common knowledge for both the advertiser and the search engine. The efficient advertisement quality, that is beneficial to the search engine but gives equal weight to the cost for the advertiser i, is $q_i^* = \arg\max_q s_i(q) - h(q)$. We design a mechanism that (i) for a given vector of offers $\mathbf{q}$ orders ads efficiently (ii) incentivizes advertisers to make the efficient choice of offerings q_i^*. More formally,

Definition 1. ***Efficiency with hidden costs.*** *For ranks and qualities $r_i, q_i, \forall i \in I$, efficiency is $\sum_i (s_i(q_i) - h(q_i))\alpha_{r_i}$.*

The generalized second price auction (GSP) is defined in [11, 30]. We describe it here for completeness. Given a set of bids b_i and clickabilities c_i, $\forall i \in I$, advertisers are ranked according to their bid times clickability. For ease of notation we assume bidder i is ranked i^{th} and bidder i charged a price per click of $\frac{b_{i+1}c_{i+1}}{c_i}$ (if there are not strictly more bidders than slots we add imaginary bidders with bids of zero). In words, their price per click is the smallest bid amount needed to maintain their position. If all advertisers have identical clickability, i.e. $c_i = c$ for every advertiser, then advertisers are ranked according to their bids and charged the bid of the advertiser ranked immediately below. Note that the GSP design does not penalize advertisers with high hidden costs and thus does not lead to an efficient outcome in an environment with hidden costs. Here we describe a modification of GSP that yields an efficient outcome (in the sense of the above definition). To keep notation as simple as possible we will assume the clickability values are all 1.

Definition 2. ***Hidden Cost GSP: Mechanism $\mathcal{M}'$.*** *Auction $\mathcal{M}'$ on bids and qualities $b_i, q_i, \forall i \in I$ is as follows:*

1. *$\forall i \in I$, $b'_i = b_i - h_i$.*
2. *Run the generalized second price (GSP) auction on b'_i.*
3. *$\forall i \in I$, add h_i to the price per click of bidder i. Leave the ranking from GSP in Step 2 unchanged.*

Mechanism $\mathcal{M}'$ differs from GSP. It subtracts and adds hidden costs before and after (respectively), running GSP. Note that the environment described here is different from the standard models for sponsored search environments [28, 21, 30, 11] because an advertiser controls two input values to the mechanism, b_i and q_i. Despite these differences, $\mathcal{M}'$ maintains many of the attractive properties of the GSP auction, namely there exists an equilibrium solution that has maximum efficiency. Also like in GSP, the bid is a tight upper bound on an advertiser's payment per click in the mechanism $\mathcal{M}'$.

Theorem 1. *Mechanism $\mathcal{M}'$ implements an efficient outcome.*

To prove Theorem 1, we start out by defining VCG and proving some properties of the VCG solution. We then draw conclusions about the GSP auction based on previous work relating it to VCG. Finally, the properties of GSP are used to prove properties of the mechanism $\mathcal{M}'$.

2.1 VCG for Sponsored Search

The VCG auction collects all bids and assigns position k to the bidder with the k^{th} highest bid. Each bidder is charged the loss in other's efficiency created by their presence. Again assuming bidders are indexed by their positions, the VCG price charged to bidder i is $\sum_{j=i+1}^{K} b_j(\alpha_{j-1} - \alpha_j)$, where K is the number of total positions available.

We denote the total VCG payment for the advertiser ranked k_{th} highest in s with $t_k^{(s,\alpha)}$. Let r be a ranking function such that $r(u, s)$ is the rank of value u if it were to be inserted into the set s. Similarly, $r^{-1}(r, s)$ is the value at rank r in vector s. We prove the following property of the VCG outcome:

Lemma 1. *For any set of values V, $argmax_{v \in V}(v \cdot \alpha_{r(v,s)} - t_{r(v,s)}^{(s+v,\alpha)}) = max_{v \in V} v$. In words, given a choice of values per click, an advertiser's utility in the VCG outcome is highest if they choose the highest value per click.*

Proof First, we show that for the VCG outcome, the higher an advertiser's value, the larger their utility. By definition, $s_k - s_{k+1} \geq 0$. Multiplying by α_k and adding $s_{k+1}\alpha_{k+1} - t_{k+1}^{s,\alpha}$ to both sides, we see that the utility of the advertiser with the k^{th} largest value is larger than the utility of the advertiser with the $k+1^{th}$ largest value.

Take two potential values, $u > w$ for an advertiser i. Let x be the value of the advertiser at rank $r(w, s)$ in the set of values $s + u$ (this is analogous to taking the value of the advertiser ranked immediately above w in $s + w$). For conciseness, $\Delta_{\alpha_k} = \alpha_k - \alpha_{k-1}$. Using value u gives utility

$$u\alpha_{r(u,s)} - \sum_{k=P+1}^{k=r(u,s)} \Delta_{\alpha_k} r^{-1}(k-1, s+u) \geq x\alpha_{r(w,s)} - \sum_{k=P+1}^{k=r(w,s)} \Delta_{\alpha_k} r^{-1}(k-1, s+u).$$

Since both w and u are ranked higher than any value of k in the summation,

$$\sum_{k=P+1}^{k=r(w,s)} \Delta_{\alpha_k} r^{-1}(k-1, s+u) = \sum_{k=P+1}^{k=r(w,s)} \Delta_{\alpha_k} r^{-1}(k-1, s+w).$$

Because w is ranked lower than x in the set of values $s + w$,

$$x\alpha_{r(w,s)} - \sum_{k=P+1}^{k=r(w,s)} \Delta_{\alpha_k} r^{-1}(k-1, s+u) \geq w\alpha_{r(w,s)} - \sum_{k=P+1}^{k=r(w,s)} \Delta_{\alpha_k} r^{-1}(k-1, s+w).$$

Combining the inequalities, we see that the larger value u brings higher utility in VCG than the lower value w. $\square$

We define $b(s, \alpha, u)$ to be $\frac{t_{r(u,s-u)}^{s,\alpha}}{\alpha_{r(u,s-u)}}$. The function b relies on the rankings and pricing that VCG outputs (based on s and α).

2.2 Relating VCG to GSP

The following Lemma restates the results from [11], which we will use to prove Theorem 1.

Lemma 2. *All bidders i bidding $b(s, \alpha, s_i)$ is an envy-free bidder optimal equilibrium. Furthermore, the ranking and the price charged to each bidder is equivalent to the ranking and price charged to each bidder in the VCG mechanism.*

Corollary 1. *Since in the bidder optimal equilibrium in GSP advertisers receive the same utility as in VCG, an advertiser with a choice of values per click receives highest utility in the bidder optimal equilibrium by bidding their highest choice of value.*

Although there exist other equilibria, it is reasonable to expect the auction to converge to the bidder optimal equilibrium. The result in [11] describes a generalized English auction that leads to the bidder optimal equilibrium. Also, in [6], it is shown that there exists a simple greedy bidding strategy that leads to the bidder optimal equilibrium. For simplicity, we assume there is only one bidder optimal equilibrium solution, since ties can be broken according to a predetermined ordering.

To show the efficiency of $\mathcal{M}'$, we first prove the efficiency of $\mathcal{M}$.

Definition 3. ***Mechanism*** $\mathcal{M}$***.*** *Auction $\mathcal{M}$ on bids $b_i, \forall i \in I$ is as follows:*

1. *Run the generalized second price (GSP) auction on b_i.*
2. *$\forall i \in I$, add h_i to the ppc of bidder i. Leave the ranking from GSP unchanged.*

Theorem 2. *All bidders $i \in I$ implementing websites with quality q_i^* and submitting bids $b(s(q_i^*) - h(q_i^*), \alpha)$ is an equilibrium point for mechanism $\mathcal{M}$. Furthermore, this equilibrium point has maximum efficiency.*

Proof For a fixed quality q_i, the utility of advertiser i in auction $\mathcal{M}$ is $(s(q_i) - h(q_i) - p_i)\alpha_i$, where p_i and α_i are the price and clicks allocated by GSP. By Lemma 2 and the definition of function b, the bid $b(s(q_i) - h(q_i), \alpha)$ must maximize $(s(q_i) - h(q_i) - p_i)\alpha_i$ in the bidder optimal GSP outcome. By Lemmas 1 and 2, the value of q_i that maximizes $(s(q_i) - h(q_i) - p_i)\alpha_i$ is q_i^*. The bidder optimal GSP outcome, since it ranks by $s(q_i^*) - h_(q_i^*)$, maximizes $\sum_i \alpha_i(s(q_i^*) - h_(q_i^*))$. $\square$

Corollary 2. *All bidders $i \in I$ implementing websites with quality q_i^* and submitting bids $b(s(q_i^*) - h(q_i^*), \alpha) + h(q_i^*)$ is an equilibrium point for mechanism $\mathcal{M}'$. Furthermore, this equilibrium point has maximum efficiency.*

Since the addition of $h(q_i^*)$ is immediately subtracted in the first step of $\mathcal{M}'$, and the utility of the bidder is unchanged from $\mathcal{M}$, the maximizing behavior of the bidder is unchanged from Theorem 2.

3 Conclusion

The experience of users is a crucial component for the success of any website. Many leading search engines emphasize the importance of creating a positive user experience, yet fall short of explicitly incorporating the user experience into the auction mechanism (we point out that creating a positive user experience is more general than displaying relevant results). We have presented an approach for incorporating this vital aspect into the current auction mechanism.

An area of future work is to design algorithms that accurately measure hidden costs. We envision that initially, an ad's hidden cost will be determined based on the contents of the ad using machine learning and other techniques. As an advertisement develops a history, the hidden cost estimates will be updated and become more accurate. There are several interesting challenges in creating and maintaining accurate estimates of the hidden costs. For example, there is the issue of how to prevent advertisers from entering as a new advertiser with a new advertisement if they have developed high hidden costs, while still maintaining accurate hidden costs for entering advertisers. We anticipate that accurately measuring hidden costs will be a rich and complex area of future research.

Finally, hidden costs address a specific type of externality for an ad, namely, the impact on the search engine's future revenue. In addition, there are other types of externalities created by an ad such as it's influence on the branding effectiveness of other advertisements [27, 20].

References

1. Z. Abrams and A. Ghosh. Auctions with revenue guarantees for sponsored search. *Workshop on Sponsored Search Auctions*, 2007.
2. M. Backmann, C. B. McGuire, and C. B. Winsten. Pigouvian pricing and stochastic evolutionary implementation. *Studies in the Economics of Transportation*, 1956.
3. E. Basker. Job creation or destruction? labor market effects of wal-mart expansion. *Review of Economics and Statistics*, 87:1:327–353, 1982.
4. E. Biyalogorsky and P. Naik. Clicks and mortar: The effect of on-line activities on off-line sales. *The Journal of Marketing Letters*, 14(1):21–32, 2004.
5. Bucklin, J. M. Lattin, A. Ansari, S. Gupta, D. Bell, E. Coupey, J. D. C. Little, C. Mehla, A. Montgomery, and J. Steckel. Choice and the internet: From click-stream to research stream. *Marketing Letters*, pages 245–258, 2002.
6. M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. *Conference on Electronic Commerc (EC)*, 2007.
7. P. Chatterjee, D. L. Hoffman, and T. P. Novak. Modeling the clickstream: Implications for web-based advertising efforts. *Marketing Science*, 22(4):520–541, 2003.
8. E. Clarke. Multipart pricing of public goods. *Public Choice, 11:17-33*, 1971.
9. P. J. Danaher. Modeling page views across multiple websites with an application to internet reach and frequency prediction. *Marketing Science*, 26(3):422–437, 2007.
10. J. Deighton. Commentary on 'exploring the implications of the internet for consumer marketing'. *The Journal of the Academy of Marketing Science*, 25(4):347–351, 1997.

11. B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
12. A. Eyton and B. Huberman. The economics of surfing. *Working Paper No. 42, Center for eBusiness at MIT*, 1999.
13. F. S. M. Florian Zettelmeyer and J. Silva-Risso. How the internet lowers prices: Evidence from matched survey and auto transaction data. *The Journal of Marketing Research*, 43(2):168–181, 2006.
14. A. Goldfarb. Analyzing website choice using clickstream data. *Working Paper, Northwestern University, Department of Economics*, 2001.
15. E. D. Gould, B. P. Pashigian, and C. J. Prendergast. Contracts, externalities, and incentives in shopping malls. *The Review of Economics and Statistics*, 2005.
16. T. Groves. Incentives in teams. *Econometrica, 41:617-631*, 1973.
17. N. Immorlica, K. Jain, M. Mahdian, and K. Talwar. Click fraud resistant methods for learning click-through rates. *Workshop on Internet and Network Economics*, 2005.
18. G. Iyer, D. Soberman, and J. M. Villas-Boas. Targeting of advertising. *Marketing Science*, 24(3):461–476, 2005.
19. E. Johnson, W. Moe, P. Fader, S. Bellman, and G. Lohse. On the depth and dynamics of online search behavior. *Working Paper, Graduate School of Business, Columbia University*, 2000.
20. K. Kim and J. Meyers-Levy. Context effects in diverse-category brand environments: The influence of target product positioning and consumers' processing mind-set. *The Journal of Consumer Research*, 0, 2007.
21. S. Lahaie and D. M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. *ACM Conference on Electronic Commerce*, 2007.
22. R. Lal and M. Sarvary. When and how is the internet likely to decrease price competition. *Marketing Science*, 18(4):485–503, 1999.
23. L. M. Lodish and C. F. Mela. If brands are built over years, why are they managed over quarters? *Harvard Business Review*, pages 104–112, 2007.
24. S. Moorthy, B. T. Ratchford, and D. Talukdar. Consumer information search revisited: Theory and empirical analysis. *The Journal of Consumer Research*, 23:263–277, 1997.
25. S. Pandey and C. Olston. Handling advertisements of uknown quality in search advertising. *Neural Information Processing Systems Conference*, 2006.
26. A. C. Pigou. *Wealth and Welfare.* 1912.
27. I. Poncina, R. Pietersb, and M. Ambaye. Cross-advertisement affectivity: The influence of similarity between commercials and processing modes of consumers on advertising processing. *Journal of Business Research*, 59:6:745–754, 2006.
28. T. Roughgarden and M. Sundararajan. Is efficiency expensive? *Third Workshop on Sponsored Search Auctions*, 2007.
29. B. Sandholm. Negative externalities and evolutionary implementation. *Review of Economic Studies*, 72:885–915, 2005.
30. H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 2005.
31. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance, 16:8-37*, 1961.