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.

ence 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 aformentioned 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

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).

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 clickstream 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.