Eliciting Preferences of Sponsored Search Advertisers: Implications for Mechanism Design

DENIS NEKIPELOV
University of Virginia

Sponsored search advertising attracts hundreds of thousands of advertisers, many with dozens or even thousands of campaigns, leading to tens of millions of distinct keyword bids. Advertiser objectives are heterogeneous. Some advertisers primarily focus on making immediate sales that are referred by clicks, while others want to promote their brand with a top-position placement. In this letter we demonstrate how one can use the empirical bidding data to recover the values of bidders in a sponsored search marketplace when the type of bidder preferences is known (i.e. whether a given bidder values clicks). We also show how one can use the history of bid changes for a given bidder to recover both the type of preferences for this bidder and the value at once. This methodology has direct implications for mechanism design making the case for combining the empirical work and auction design to avoid the optimization of the auction mechanism for the wrong preference type of the bidders.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Economics, Market Design, Econometrics

Additional Key Words and Phrases: Sponsored search advertising, auctions, identification, Baysian inference

1. INTRODUCTION

Sponsored links that appear beside Internet search results on the major search engines are sold using real-time auctions. Advertisers place standing bids that are stored in a database, where bids are associated with search phrases that form part or all of a user's search query. Each time a user enters a search query, applicable bids from the database are entered in an auction. The ranking of advertisements and the prices paid depend on advertiser bids as well as "quality scores" that are assigned for each advertisement and user query. These quality scores vary over time, as the statistical algorithms incorporate the most recent data about user clicking behavior on this and related advertisements and search queries.

[Edelman and Schwarz 2007] and [Varian 2009] assume that biders value user clicks and bids are customized for a single user query and the associated quality scores. However, in practice quality scores do vary from query to query, queries arrive more quickly than advertisers can change their bids,[1] and advertisers cannot perfectly predict changes in quality scores. In this letter we consider the framework (first studied in [Athey and Nekipelov 2010]) where bids apply to many user queries, while the quality scores and the set of competing advertisements may vary

[1]Although bids can be changed in real time, the system that runs the real-time auction is updated only periodically based on the state at the time of the update, so that if bids are adjusted in rapid succession, some values of the bids might never be applied.

Author's address: denis@virginia.edu

from query to query. We further notice a significant heterogenity in the bidder behavior as indicated by the data coming from the actual sponsored search marketplace. Bidders differ by their presence in the markeplace, their approaches to bid adjustment and bid change frequency. We argue that we can account for such a heterogeneity by introducing different types of bidder preferences and demonstrate how one can use the bidding history for each bidder to recover both the preference type and the value for this bidder.

2. INFERENCE WITH VALUES PER CLICK

Consider the model where I bidders are competing for J advertising positions, where bidder scores s_i for $i = 1, \ldots, I$ and position effects α_j for $j = 1, \ldots, J$ determine the probability of click $c_{ij} = \alpha_j \, s_i$ on the advertisement of bidder i placed in position j. We assume that the score of a particular bidder i for a user query is a random variable, denoted s_i, which is equal to $s_i = \bar{s}_i \varepsilon_i$, where ε_i is a shock to the score induced by random variation in the estimates that come from an algorithm the platform uses to infer the click-through rate of a particular bidder.

The ad platform conducts a click-weighted GSP auction. Each advertisement i is assigned score s_i, and bids are ranked in order of the product $b_i \, s_i$. The per-click price p_{i_j} that bidder i_j in position j pays is determined as the minimum price such that the bidder remains in her position

$$p_{i_j} = \min\{b_{i_j} : s_{i_j} b_{i_j} \geq s_{i_{j+1}} b_{i_{j+1}}\} = \frac{s_{i_{j+1}} b_{i_{j+1}}}{s_{k_j}}. \tag{1}$$

Note that advertiser k_j does not directly influence the price that she pays, except when it causes her to change positions, so in effect an advertiser's choice of bid determines which position she attains, where the price per click for each position is exogenous to the bidder and rises with position.

Per click values of bidders v_i $i = 1, \ldots, I$ are commonly observable, and fixed across user queries.The payoff of the bidder in a query where this bidder receives a click is the surplus $v_i - p_i$, where p_i is bidder i's price per click defined by (1).

Suppose that the bidder does not observe the set of competitors in a given query and does not observe neither her own nor her competitors' score shocks. We assume that the bidder forms beliefs regarding the distribution of all the score shocks of all bidders who are eligible for a given query and beliefs regarding the distribution of realizations of the set of her competitors in a use query. We also assume that the bidder can observe the actual bids of her competitors. We assume that each bidder i maximizes the expected payoff (with per click value) with the expectation taken with respect to the bidder's beliefs regarding the distribution of uncertainty of scores and sets of competitors.

Consider observing a large number of queries for a given set of potential bidders, and consider the question of whether the valuations of the bidders can be identified. For each query, we assume that we observe bids, the set of entrants, and the scores.

In this case we can define function $Q_i(\cdot)$ equal to the probability of click on the ad of bidder i in a search query as a function of the profile of bids. We also define function $TE_i(\cdot)$ equal to the expected expenditure of bidder i in a search query as a function of the profile of bids. [Athey and Nekipelov 2010] demonstrate that with a sufficient smoothness and support size of the distribution of scores, functions $Q_i(\cdot)$

and $TE_i(\cdot)$ are strictly increasing and differentiable in b_i, and we can recover the valuation of each bid using the necessary condition for the optimality of a bid

$$v_i = \frac{\partial TE_i(b_i, b_{-i})/\partial b_i}{\partial Q_i(b_i, b_{-i})/\partial b_i}, \tag{2}$$

given that all of the distributions required to evaluate functions $Q_i(\cdot)$ and $TE_i(\cdot)$ are assumed to be observable. Equation (2) provides a simple practical method for estimating values per click: for each bidder we change her bid by a small amout and then compute the change in the outcome of the auctions where the original bid of the bidder was applied. The evaluated change in the expenditure will serve as an estimator for the derivative of $TE_i(\cdot)$ function and the evaluated change in the number of clicks will serve as an estimator for the derivative of $Q_i(\cdot)$ function.

3. INFERENCE WITH HETEROGENEOUS OBJECTIVES

The model in the previous section is based on the assumption that bidders value clicks. In reality, sponsored search advertisers can observe several parameters in addition to clicks that include their average rank, the number of user impressions, and the average cost per click. These facts resonate with empirical observations in [Athey and Nekipelov 2014] where it was noted that that advertisers significantly differ in the patterns and the speed of their bid changes in response to the changes in the auction platform. In addition, bidders in the sponsored search auctions have access to the automated bidding tools which explicitly allow the bidders to specify objectives which are different from clicks (such as user impressions). Moreover, the bidders can set explicit budget constraints and specify such budget constraint for time intervals of different length. Provided that the auction platform is constantly changing (e.g., due to the fluctuations in the user traffic or the changes in the platform settings) optimizing bidders will be be adjusting their bids to respond to the changes. [Athey and Nekipelov 2014] provide an econometric framework that allows them to estimate both the objective function that each bidder is optimizing *and* the parameters of this objective function using the dynamic bidding data. To explain the idea of [Athey and Nekipelov 2014] consider a very simplified version of their setup and assume that the bidders can have preferences over two mutually exclusive characteristics: clicks and user impressions. We use the previous notation v_i for the value per click of bidder i and the notation $\tilde{v}_i$ for the value per impression of bidder i. Let $I_i(\cdot)$ denote the probability that bidder i appears in the user query (and thus gets a user impression). We assume that bidders have limited attention and adjust their bids in response to large changes in the performance of their ads. Suppose that bidder i changes the bid at time t_1 to the value $b_i^{t_1}$. This bid is optimal for the objective function that this bidder is maximizing (corresponding to either clicks or impressions). Since the type of objective function of the bidder is not known *a priori*, we can recover implied value per click and the value per impression from the respective first-order conditions:

$$v_i^{t_1} = \frac{\partial TE_i^{t_1}(b_i^{t_1}, b_{-i}^{t_1})/\partial b_i}{\partial Q_i^{t_1}(b_i^{t_1}, b_{-i}^{t_1})/\partial b_i}, \quad \tilde{v}_i^{t_1} = \frac{\partial TE_i^{t_1}(b_i^{t_1}, b_{-i}^{t_1})/\partial b_i}{\partial I_i^{t_1}(b_i^{t_1}, b_{-i}^{t_1})/\partial b_i}.$$

Suppose that at time t_2 bidder i adjusts the bid to $b_i^{t_2}$. Using the same principle, we can again recover the implied value per click $v_i^{t_2}$ and the implied value per impression $\tilde{v}_i^{t_2}$. Then we can recover the type of the bidder in the following way: if $v_i^{t_1} = v_i^{t_2}$, we say that bidder i values clicks and estimate this bidder's value per click as $v_i = v_i^{t_1} = v_i^{t_2}$. Otherwise, if $\tilde{v}_i^{t_1} = \tilde{v}_i^{t_2}$ we say that bidder i values impressions.[2] [Athey and Nekipelov 2014] extend this framework to the case where the values are estimated with an error and use the Bayesian framework to select the type for each bidder. They apply their methodology to a subsample of high-revenue search phrases on Bing.com covering a year-long history of bid changes. They consider three types of objective functions including the objective function where bidders value clicks, the objective function where the bidders value clicks *and* the impressions in the top positions, and the objective function where the bidders value impressions with a possible premium on the top position impressions. Their findings show the dominance of the first and the third type of preferences. Moreover, they demonstrate that bidders which are "exact matched" to the search phrase (i.e. the bidder bids on a single search phrase) are more likely to value impressions, while the bidders which are "broad matched" to the search phrase (i.e. their bid can be applied to multiple search prases) are more likely to value clicks.

4. CONCLUSION

We have considered a framework that allows the researcher to use bidding data to infer the values of bidders in sponsored search auctions. Moreover, when a bidding history is available for a given bidder, it becomes possible to also identify the type of preferences for that bidder. Unlike most theoretical literature on sponsored search advertising auctions, this allows the preferences of bidders to be defined over objects like impressions and rank in the auction rather than clicks. Our empirical analysis shows that such new preference types are consistent with observed bidding in some segments of the sponsored search marketplace. The econometric inference may be the important missing link in the automation of the mechanism design, e.g. as in [Conitzer and Sandholm 2002]. First of all, given that the values of the bidders can be directly estimated from the data, these values can be used as inputs into the auction settings thus allowing the auction mechanism to adjust to the data. Second, they demonstrate that empirical analysis is crucial for the auction design. For instance, if the auction mechanism is optimized assuming a specific type of bidder preferences, one needs to validate the consistency of this assumption with the actual behavior of the bidders to avoid significant revenue and welfare loss.

REFERENCES

Athey, S. and Nekipelov, D. 2010. A structural model of sponsored search advertising auctions. In *Sixth Ad Auctions Workshop*.

Athey, S. and Nekipelov, D. 2014. Designing large advertising markets where agents have heterogeneous objectives: A structural empirical approach. *Stanford University Working Paper*.

Conitzer, V. and Sandholm, T. 2002. Complexity of mechanism design. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 103–110.

[2]Note that if neither value remains stable within this framework, we need to reject both models.

Edelman, B., M. O. and Schwarz, M. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 242—259.

Varian, H. R. 2009. Online ad auctions. *The American Economic Review*, 430–434.