Simulation-Based Analysis of Keyword Auctions

YEVGENIY VOROBEYCHIK
Computer and Information Science
University of Pennsylvania
yev@seas.upenn.edu

Keyword auctions account for an enormous proportion of revenue for the major search engines. Consequently, substantial literature analyzing alternative auction designs has sprouted in recent years. We contribute to this growing literature by engaging in a simulation-based analysis of strategic interactions in keyword auctions which are difficult to analyze in closed form. In this letter we provide an overview of two specific efforts in this direction. The first effort presents analysis of the *dynamic* bidding strategies, while the second effort is performed in a *static* context, but involves a close analysis of several classes of Bayesian bidding strategies. Both these efforts attempt to bridge the gap between the mostly theoretical literature to date and real auctions.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics; I.6.6 [**Simulation and Modeling**]: Simulation Output Analysis

General Terms: Economics, Experimentation

Additional Key Words and Phrases: Electronic commerce, keyword auctions, simulation-based game theory

1. MOTIVATION

Keyword auctions have become one of the primary sources of revenue for the major search engines. Since their beginnings these auctions have undergone a series of changes, from a first-price to a next-price format, from rank-by-bid to rank-by-revenue. Roughly, the consensus today stands at the rank-by-revenue schemes, which augment the ranking rule to include advertiser *quality scores*[1] and a next-price (or *generalized second-price (GSP)*) format.[2]

In the academic literature, much progress has been made in modeling sponsored search auctions as one-shot games of complete information, in which the players' values per click and click-through-rates are common knowledge. As the complexity of modeling the full dynamic game between advertisers that is actually taking place is quite intractable, static models provide a good first approximation. However, it ultimately pays to understand how relevant the dynamics really are to strategic choices of players. Furthermore, since actual valuations of bidders are not mutually observed and likely change over time, it is also valuable to understand the auctions from the incomplete information perspective.

One question which has been addressed in the dynamic setting is whether it is reasonable to expect simple dynamic strategies to converge to Nash equilibria. Cary et al. [2007] explored several *greedy bidding strategies*, that is, strategies under which players submit bids with the goal of obtaining the most profitable slot given

[1] Quality scores reflect the effect of the ad quality on the probability it is clicked.
[2] In a *GSP* auction an advertiser pays the minimum amount sufficient to remain in the currently allocated slot.

that other players' bids are fixed. One of these strategies, *balanced bidding*, provably converges to a minimum revenue symmetric Nash equilibrium of the static game of complete information.

Our own approach ([Vorobeychik and Reeves 2008]) focuses on a small set of greedy bidding strategies described in Cary et al. [2007]. In motivating greedy bidding strategies, Cary et al. have argued that advertisers are unlikely to engage in highly fine-grained strategic reasoning and will rather prefer to follow relatively straightforward strategies. This motivation, however, only restricts attention to a set of plausible candidates. To identify which are likely to be selected by advertisers, we need to assess their relative stability to profitable deviations. For example, while we would perhaps like advertisers to follow a convergent strategy like balanced bidding, it is unclear whether players will find it more profitable to follow a non-convergent strategy.

An important shortcoming of greedy bidding strategies is the unrealistic assumption that all advertisers observe the bids previously submitted by their competitors. Eliminating observability entirely, we are back to a one-shot incomplete information setting, which is the subject of Vorobeychik [2009]. The reality of keyword auctions, of course, is somewhere between these two extremes, and our hope is that in studying the extremes carefully, we shed considerable light on the real problem and, perhaps, pave way to more sophisticated future analyses of yet more accurate models of actual strategic interactions in keyword auctions.

2. ANALYSIS OF DYNAMIC BIDDING STRATEGIES

The analysis of dynamic bidding strategies draws on the work of Cary et al. [2007] in identifying both the restricted class of dynamic strategies and the specific candidates to consider. Balanced bidding (the convergent strategy in the analysis of Cary et al.), along with three other dynamic strategies (two of which proposed by Cary et al.) form the strategic landscape which is analyzed by Vorobeychik and Reeves [2008]. Since valuations are not common knowledge, but each advertiser can be reasonably assumed to at least estimate his own value per click, the analysis is performed by transforming the dynamic game into a *static game of incomplete information* in which each advertiser is allowed to choose one of the restricted set of *dynamic* strategies (such as balanced bidding), *as a function of his value*. Based on the simulation-based analysis of the dynamic strategies, we make the following observations: (a) balanced bidding cedes the fewest benefits for deviation (i.e., players rarely prefer to play another of the four dynamic strategies analyzed) and (b) balanced bidding yields the least payoffs to players (highest revenue to the designer). Thus, even though players would prefer to coordinate on an alternative strategy, inability to do so greatly benefits the search engine.[3]

3. KEYWORD AUCTION DESIGN IN A ONE-SHOT BAYESIAN SETTING

All dynamic bidding strategies analyzed by Vorobeychik and Reeves [2008] make a strong assumption that the actual bids of all advertisers are mutually observable. This may have been true in the past, but is no longer the case for the major search

[3]As a corollary, it may serve the search engines well to implement balanced bidding as a proxy bidding agent.

engines. Furthermore, the analysis was highly restricted, considering only four choices in the corresponding Bayesian game. Finally, another central assumption of this analysis is that valuations of the advertisers do not change over time, an assumption that is likely quite inaccurate in many settings. An alternative analysis undertaken by Vorobeychik [2009] looks at a static setting, but retains the assumption of incomplete information. Instead of making the strong restriction to a small finite number of choices, Vorobeychik considers two infinite low-dimensional strategy classes. The first one-dimensional strategy class would submit bids which are a fraction $\alpha \in [0, 1]$ of the actual valuation; the second extends the first class by adding a quadratic term.

3.1 Incentive Compatibility

While it is well known that generalized second-price auctions (GSPs) are not incentive compatible, to our knowledge this "failure" of incentive compatibility has never been quantified. In practice, quantifying it seems quite useful for two reasons: first, if the failure is negligible, it can be ignored by the designer, since bidding true valuation is an especially convenient and focal choice by the advertisers; and second, the advertisers would certainly want to know how much they can gain by "shading" their valuations. Vorobeychik [2009] demonstrates that the benefits that can accrue to advertisers for bidding well below value are often substantial, and the designer is likely to be misled if he reasons that GSP is approximately truthful.

3.2 Profit Maximization

In past work on revenue analysis in keyword auctions it has been shown that the optimal mechanism (be it rank-by-revenue, rank-by-bid, or something "in the middle" in the restricted design space) is highly dependent on the correlation between bidder values and their click-through-rates [Lahaie and Pennock 2007]. Since the analysis by Lahaie and Pennock has been performed under complete information, a question remains whether it is robust to relaxing the complete information assumption.

Vorobeychik [2009] considers a variety of joint distributions between players' values and click-through-rates (and, in addition, considers a case where values are correlated across players, which has rarely been analyzed but seems more a rule than an exception) and uses simulations to approximate Bayes-Nash equilibria. This paper finds, somewhat in disagreement with Lahaie and Pennock [2007], that rank-by-bid mechanisms are often optimal, and usually better than pure rank-by-revenue schemes. Indeed, Vorobeychik shows that the rank-by-bid mechanism remains optimal even if we assume that bidders do not adjust their bidding strategies.

4. CONCLUSION

We provided a glimpse at several contributions to the analysis of keyword auctions. One contribution considers a small set of dynamic strategies, and finds that the one that is convergent is also highly stable to deviations and yields high revenue to the search engine. The second contribution analyzes a one-shot incomplete information model, and demonstrates that (a) GSP mechanisms are quite far from truthful and (b) the rank-by-bid mechanism tends to yield a higher profit to the designer than rank-by-revenue.

REFERENCES

Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A. R., Mathieu, C., and Schwarz, M. 2007. Greedy bidding strategies for keyword auctions. In *Eighth ACM Conference on Electronic Commerce*. 262–271.

Lahaie, S. and Pennock, D. M. 2007. Revenue analysis of a family of ranking rules for keyword auctions. In *Eighth ACM Conference on Electronic Commerce*. 50–56.

Varian, H. 2007. Position auctions. *International Journal of Industrial Organization 25,* 6, 1163–1178.

Vorobeychik, Y. 2009. Simulation-based game theoretic analysis of keyword auctions with low-dimensional bidding strategies. In *Conference on Uncertainty in Artificial Intelligence*.

Vorobeychik, Y. and Reeves, D. 2008. Equilibrium analysis of dynamic bidding in sponsored search auctions. *International Journal of Electronic Business*.