Games and Economic Behavior 102 (2017) 20–43

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

An experimental study of sponsored-search auctions ☆

Yeon-Koo Che [a], Syngjoo Choi [b], Jinwoo Kim [b,*]

[a] Department of Economics, Columbia University, 420 West 118th St, 10161AB, New York, NY 10027, USA
[b] Department of Economics, Seoul National University, 1 Gwanak-ro Gwanak-gu, Seoul 151-742, South Korea

ARTICLE INFO

Article history:
Received 23 April 2014
Available online 12 November 2016

JEL classification:
C92
D44
M3

Keywords:
Online advertising
Sponsored search auction
Generalized second price auction
Experiment

ABSTRACT

We study the Generalized Second Price auctions—a standard method for allocating online search advertising—experimentally, considering both the static environment assumed by the prevailing theory and a dynamic game capturing the salient aspects of real-world search advertising auctions. Subjects of our experiment bid consistently with the leading equilibrium notions, but exhibit significant overbidding relative to the Vickrey–Clarke–Groves (VCG) outcome favored as an equilibrium selection in the literature. The observed bidding behavior is well explained by a model that explicitly accounts for the strategic uncertainty facing a bidder, which suggests strategic uncertainty as a source of the observed departure from the VCG outcome. Meanwhile, the observed bidding behavior in static environment approximates those of dynamic environments for important cases. Our finding thus provides some empirical support for the use of a static game as a valid modeling proxy, but calls into question the prevailing equilibrium selection.

1. Introduction

Search engines such as Google, Yahoo! and Microsoft sell online search spaces to advertisers. In comparison with conventional advertising, online search advertising is highly targetable, and thus is an effective means for finding buyers. Naturally, the sponsored search auctions have become a major revenue source for search firms. In 2007, search advertising accounted for more than $21 billion of revenue for search firms in US.[1] The auction format used for selling ad spaces has evolved, with a few adjustments along the way, to what is now known as the *generalized second price* (GSP) auctions.

Under the GSP, advertisers bid per-click prices, and these bids are converted into per-impression bids—their per-click bids multiplied by the estimated click through rates—to determine the assignment of ad positions. Specifically, the highest

☆ We are grateful to Jacob and Michelle Goeree, John Kagel, Dan Levin, Michael Ostrovsky, Michael Schwarz, Yan Chen, and the participants at the AMMA 2011, Monash University, Ohio State University, University of East Anglia, University of Michigan, Universitat Pompeu Fabra, University of Zurich, and the WCU Market Design Conference held at Yonsei University in August 2010, for valuable comments, and to Brian Wallace for writing the experimental program, and to Tom Rutter for helping us run the experiment. Che acknowledges the support by NSF (#1023818). Che and Kim acknowledge the support by the National Research Foundation of Korea funded by the Ministry of Education, Science and Technology through its grant of Global Research Network (2013S1A2A2035408). Choi acknowledges the support by the Housing and Commercial Bank Economic Research Fund for Institute of Economic Research of Seoul National University.

* Corresponding author.
E-mail addresses: yc2271@columbia.edu (Y.-K. Che), syngjooc@snu.ac.kr (S. Choi), jikim72@snu.ac.kr (J. Kim).
URLs: http://www.columbia.edu/~yc2271 (Y.-K. Che), https://sites.google.com/site/syngjoochoi/ (S. Choi), https://sites.google.com/site/jikim72/home (J. Kim).

[1] See http://www.iab.net/media/file/IAB_PwC_2007_full_year.pdf. This revenue includes over 90% of Google's revenue and 50% of Yahoo! and MS's revenues (more precisely, the revenue of Online Services Division of MS) according to each firm's report of annual revenue.

http://dx.doi.org/10.1016/j.geb.2016.10.008
0899-8256/

bidder (in per-impression bid) is assigned the top position, the second-highest bidder is assigned the next best, and so on. A winner of each ad position then pays the smallest price per click that would have won that position. If the number of clicks depends only on one's position, as is often assumed, per-impression bids essentially coincide with per-click bids, so each winning bidder simply pays a per-click price that equals the bid submitted by the next-highest bidder.

The prevailing theory considers the GSP in a static model in which advertisers bid *simultaneously with complete information about others' preferences* (Edelman et al., 2007, henceforth EOS; and Varian, 2007). EOS and Varian then focus on a class of Nash equilibria, called *locally envy-free* or *symmetric*, in which no bidder wishes to exchange his winning position and the associated price with others' positions and the prices they are paying for them. The Symmetric Nash Equilibrium (SNE) concept predicts efficient allocation of ad positions but admits a plethora of equilibrium prices, including those that would obtain if the Vickrey-Clarke-Groves (VCG) mechanism were employed. This VCG equilibrium is the most preferred by bidders among all SNE's, and is suggested as the most plausible.[2]

While the theory provides useful insights on GSP auctions, it raises two issues. First, unlike the theory, sponsored-search auctions in practice take place continuously in real time—in principle, whenever a user types in a search query—and also bidders are unlikely to have complete information about one another's preferences. This means that advertisers face complex dynamic interactions which may provide them with opportunities to learn and adjust their behavior over time. The actual practice is therefore best described by a dynamic game in which bidders with incomplete information play repeatedly over time. It is unclear whether the static complete information model can adequately represent this rich dynamic environment. Second, the theory lacks a sharp prediction due to the multiplicity of equilibria. Although the literature suggests the VCG outcome as the most salient, there is no compelling theoretical argument or empirical evidence supporting the selection of this outcome.

The current paper investigates these issues via a laboratory experiment. At its core, our experiment induces two environments. The basic control environment is the **static complete-information** game (henceforth **SC**) used in theory, wherein *subjects play one-shot GSP game with complete information about one another's preferences*. The main treatment environment is a **dynamic incomplete-information** game (henceforth **DI**) that captures the salient features of the GSP game in practice, wherein *subjects play the GSP games repeatedly, with possible feedback and learning, but without complete information about their opponents' preferences.* Since the dynamic game differs from the static game with respect to *both* timing and information, we also consider a **static incomplete-information** game (henceforth **SI**) in which *bidders play one-shot GSP game under incomplete information about their opponents' preferences*. This additional game serves as a bridge between the static complete information game and dynamic incomplete information game.

Specifically, our experiment considers three bidders competing to obtain one of the two bundles, A and B, each containing c_A and c_B units of a (fictitious) homogeneous commodity, respectively. The two bundles represent two advertising positions, and the units of the commodity in a bundle represent the number of clicks an ad position receives for a given period.[3] Bundle A contains more units of the commodity than does bundle B, i.e., $c_A > c_B > 0$, and the ratio of the units in bundle B and A, $\frac{c_B}{c_A}$, captures "clicks decays" across ad positions. In fact, the nature of strategic environment depends crucially on the magnitude of click decays. When the ratio $\frac{c_B}{c_A}$ is close to zero (i.e., the unit difference is large), A becomes so much more attractive than B, so the competition becomes essentially about winning bundle A. The game thus becomes close in nature to standard second-price auctions. By contrast, if the ratio $\frac{c_B}{c_A}$ is close to one (i.e., the unit difference is small), both bundles are almost equally good, so bidders compete to win *either* bundle at a cheaper price per unit. The strategic environment thus resembles that of Bertrand game. In the design, we use two different values of the ratio $\frac{c_B}{c_A}$ in each of the aforementioned three games. The difference in the strategic interaction provides an opportunity to test whether players respond strategically to the environments, and this difference provides testable restrictions on the subjects' strategic responses, which is exploited in our experimental design.

Our experiment yields several results. First, the GSP auctions turned out reasonably good efficiency performances in our experiment, with the average surplus across the treatments being 76–93% of the maximum possible surplus improvement over random assignment of the positions. Second, the GSP in the lab yields revenue that is within the upper bound of the symmetric Nash equilibrium, but exceeds consistently the revenue corresponding to the particular selection, namely the VCG outcome, specifically at the median by 3% and 4% in the SC treatments, by 30% and 40% in the SI treatments, and by 9% and 18% in the DI treatments. The findings in efficiency and revenue, while suggestive, may not provide detailed account of the behavior of the subjects, so our main analysis focuses on the subjects' bidding behavior.

A closer inspection of the subjects' bidding behavior shows a broad consistency with Nash equilibrium and symmetric (locally envy free) Nash equilibrium, but a systematic departure from its particular selection, the VCG outcome. Specifically, the majority of bidders with the lowest per-unit values bid close to their true values, as predicted by the theory, while a significant fraction of subjects also bid above their values. This latter pattern of behavior is in common with the overbidding patterns documented in the experimental literature of standard, single-unit, second-price auctions (see Kagel et al., 1987; Kagel and Levin, 1993; Andreoni et al., 2007; and Cooper and Fang, 2008). On the other hand, the bidders with the interme-

[2] See Section 2 for several arguments that have been made in the literature in support of selecting the VCG outcome.

[3] We are thus assuming that the number of clicks an ad receives depends only on the position it is placed in. This assumption may not be realistic, but it serves to simplify the strategic environment for the experimental subjects, and more importantly, to facilitate the testing of the theory, which makes the same assumption for the most part. EOS and Athey and Nekipelov (2010) extend the theory to introduce the advertiser-specific factor in click generation.