Decision Support Systems 56 (2013) 371–378

Contents lists available at ScienceDirect

Decision Support Systems

journal homepage: www.elsevier.com/locate/dss

Bidding behaviors for a keyword auction in a sealed-bid environment☆

Y. Kamijo

Kochi University of Technology, Department of Management, 185 Miyanokuchi, Tosayamada-Cho, Kami-Shi, Kochi 782-8502, Japan

ARTICLE INFO

Article history:
Received 23 March 2012
Received in revised form 27 May 2013
Accepted 8 July 2013
Available online 17 July 2013

JEL classification:
C72
C91
D44

Keywords:
Internet advertisements
Keyword auction
Sealed-bid environment
Dynamic game of incomplete information

ABSTRACT

A keyword auction is conducted by Internet search engines to sell advertising slots listed on the search results page. Although much of the literature assumes the dynamic bidding strategy that utilizes the current bids of other advertisers, such information is, in practice, not available for participants in the auction. This paper explores the bidding behavior of advertisers in a sealed-bid environment, where each bidder does not know the current bids of others. This study considers secure bidding with a trial bid (SBT) as the bid adjustment process used by the advertisers, which is functional in a sealed-bid environment. It is shown that the SBT bid adjustment process converges to some equilibrium point in a one-shot game irrespective of the initial bid profile. Simulation results verify that a sealed-bid environment would be beneficial to search engines.

1. Introduction

Internet advertisements called sponsored links, which are shown along with search results for a keyword or combination of keywords, are sold through keyword auctions. Each time a user enters a search term into a search engine such as Google, Yahoo! or Bing, an auction is run, and advertisement positions and advertisement fees are determined based on the auction result. Over a million keyword auctions are conducted each day all over the world, and Internet advertisements from keyword auctions are a principal source of revenue for search engines.

The generalized second-price (GSP) auction and the auction mechanisms based on it are most widely used for selling advertisements on Internet search engines. In the GSP, based on the bids submitted by advertisers, ad slots are allocated according to the descending order of the bids, that is, the top position is allocated to the bidder with the highest bid, the second-ranked position is allocated to the bidder with the second-highest bid, and so on. Every time a search engine user clicks the advertisement, the advertiser pays the bidding price of the advertiser one position lower. Thus, this is a second-price auction for selling multiple objects with a one-dimensional strategy space.

Since the payment of each advertiser does not depend on his bid, but on the bid submitted by the advertiser one position lower than his, the GSP auction is similar to the Vickrey auction selling one object [25]. In fact, when there is only one ad slot, the GSP auction is equivalent to the Vickrey auction and thus, it has the following property: submitting the true expected revenue from the sponsored link is a dominant strategy for each advertiser.

However, when there are multiple ad slots, the GSP auction does not retain the truth-telling property [9]. This indicates that advertisers participating in the GSP auction have no option but to undertake the complicated task of choosing their bids.

Edelman and Ostrovsky [8] reported that bids observed in GSP auctions fluctuate widely, and proposed that this could be caused by the bidders' strategic behavior.

In this paper, I explore bidding behavior for a hypothetical keyword auction. As explained in the previous paragraph, the bids submitted by advertisers vary over a given period. This suggests that we should pay attention to the dynamic aspect of bidding behavior. After describing the bidding behavior of the advertisers in a keyword auction, I examine whether a stable bid profile exists for the bidding behavior. In the event that it is stable, I investigate the property that the stable bid profile possesses. I also explore how long it takes to realize the stable bid profile.

My analysis considers a simplified model of keyword auctions.[1] I assume that the click-through rates (CTRs) of ad slots are common knowledge. In each period, an advertiser can change his bid according to the result of the keyword auction played in a previous period. The information available to the advertiser is limited to his revenue, his payment to a search engine, and the manner in which ad slots were assigned to advertisers in a previous period. The advertiser does not

E-mail address: yoshio.kamijo@gmail.com.

[1] More complicated, realistic formulations of a keyword auction have been done by, for example [11,15,2,18].

0167-9236/$ – see front matter © 2013 The Author. Published by Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.dss.2013.07.003

know the actual bids of the other advertisers. This means that advertisers cannot follow the greedy bidding strategy, where in each period, they update their bids to provide the best response to others' bids. Since a keyword auction, in practice, is a *sealed-bid* second-price auction, advertisers update their bids according to the limited information.

Since 2002, both Google and Yahoo!, the two leading search engines, have used the GSP auction mechanism.[2]

An important difference between the auctions conducted by Yahoo! and Google was that Google employed a sealed-bid auction, while Yahoo!'s auction was an open-bid auction. In Yahoo!'s keyword auction, the current bids of advertisers were publicly provided through the software (the View Bid Tool). However, this service was discontinued in 2007 when Yahoo! switched its allocation rule to mirror Google's quality-based bidding. Thus, currently, keyword auctions managed by the two leading search engines are sealed-bid auctions.[3] Moreover, search engines generally restrict the bidding information available to the automated bidding software and require a review of any automated bidding code [14].

Even though the current keyword auctions are sealed-bid auctions, most studies concerning the bidding strategy for a dynamic auction assume an open-bid environment. Cary et al. [5,6] considered a type of greedy bidding strategy. Since the payment is calculated by a second-pricing rule, there can be multiple best-response bids even though the best ad slot is uniquely determined. In their analysis, among the best-response bids, the bidder was assumed to choose one bid so as to balance two objectives: to push the prices paid by the other advertisers higher and to limit the risk that a change in other advertisers' bids could result in the bidder paying a higher price than expected. Thus, this bidding strategy is called balanced bidding. Bu et al. [4] analyzed the same bidding behavior.[4] In addition to the greedy bidding strategy, other bidding strategies such as antisocial bidding have also been analyzed in the literature [3,20,28].

In this paper, a bid adjustment process in a sealed-bid environment is analyzed. While the existing literature provides a good perspective on how bidders change and adjust their bids in a dynamic auction, the analysis of bidding behavior based on a more realistic setting is also encouraged. Even though a sealed-bid environment may be temporary because of the bidders' actual experience in a dynamic auction, the question of how bidders adjust their bids and how bidders learn of other bids is answered only by considering the sealed-bid environment.

First, I consider a conservative bidding strategy called secure bidding. The idea of secure bidding was derived partly from balanced bidding, which was proposed by Cary et al. [5] for the open-bid environment. The bidder who follows secure bidding adjusts his bid, given his revenue, his payment, and his ad slot, and never searches for information about the bids of other advertisers. I show that there exist multiple stable bid profiles against secure bidding (or the fixed point of bidding behavior according to secure bidding) and that some of them are not an envy-free equilibrium [9,23], a Nash equilibrium, or efficient. This implies that to achieve equilibrium, the searching behavior for other bids should be incorporated.

Next, I consider the bidding behavior based on secure bidding that entails a trial bid in a short period as a partial exploration of the competitor's bid in one higher ad slot. I show that the fixed point of secure bidding with a trial bid (SBT) exists uniquely. Moreover, at the fixed point, the ad slots are efficiently assigned to advertisers, the bid profile is an envy-free equilibrium, and the revenue of a search engine is the same as that in the truth-telling equilibrium in the Vickrey-Clarke-Groves (VCG) mechanism [25,7,17].

I also examine whether advertisers' bids converge to the stable bid profile if they update their bids repeatedly according to SBT. I consider an asynchronous model of bid adjustment, where in each period, one bidder is randomly selected and this bidder changes his bid according to SBT. I show that in the resulting Markov process, convergence occurs with probability one in the sealed-bid repeated keyword auction. This is similar to the observation in an open-bid environment reported by Bu et al. [4] and Cary et al. [5].

I also consider greedy bidding in a sealed-bid environment. Since the bids of others are not revealed to the bidder, he has to search for their bids on his own or through an automated bidding agent. A bidder finds others' bids randomly, and from among these, he calculates the best ad slots to acquire and submits a secure bid for the ad slot. I show that imperfect greedy bidding converges with probability one to the same fixed point as that in SBT.

Finally, I compare the bidding behavior in a sealed-bid environment with that in an open-bid environment using a computer simulation. I compare the convergence time, search engines' revenues, and advertisers' utilities in the SBT bid adjustment process and the bid adjustment process in the literature. The simulation results suggest that in the sealed-bid environment, the convergence time becomes longer and the average revenue of search engines becomes higher compared to the open bid environments. Thus, the sealed-bid environment can be beneficial to a search engine. However, advertisers can improve payoffs by switching their bidding behavior from SBT to the greedy bidding strategy even though the search for other advertisers' bids is imperfect.

2. A keyword auction

2.1. The environment

There are N, $N \geqq 2$, advertisers (bidders) participating in a keyword auction. Each advertiser i has an expected revenue v_i per ad click, called a value, and it is assumed that $v_1 > v_2 > \ldots > v_N$. There are K ad slots with CTR $\alpha_1 \geqq \alpha_2 \geqq \ldots \geqq \alpha_K$, where α_k is the estimated probability of being clicked, or the estimated number of clicks in a given period for an advertiser in the k-th ad slot. We also set $\alpha_k = 0$ for all $k > K$ and assume $N \geqq K$.[5]

2.2. The generalized second-price auction

Each advertiser submits a bid in the auction. Let b_i be an advertiser i's bid. I denote the bid profile of N advertisers by $\mathbf{b} = (b_1, \ldots, b_N)$.

In the GSP auction, advertisers are allocated ad slots in descending order of their bids $b_1, b_2, \ldots, b_N$. Let $d(k)$ denote the bidder who submits the k-th highest bid among $\mathbf{b}$. In the GSP auction, bidder $d(k)$ acquires ad slot k.

The advertiser obtaining the k-th ad slot pays the bid of the advertiser obtaining the next ad slot lower down (i.e., the k + 1-th ad slot) for each click.

Hence, the payment is $\alpha_k b_{d(k+1)}$.

To complete the definition of the payments, I assume that $b_{d(k)} = 0$ if $k > N$. Accordingly, when $K = N$, the payment of $d(K)$ is assumed to be zero, and for $k > K$, bidder $d(k)$ pays $\alpha_k b_{d(k+1)} = 0$ (as per the definition of α_k).

[2] Although they used the same payment rule, they follow slightly different rules for the allocation of ad slots to advertisers. Yahoo! ranks advertisers by their bids alone, while Google computes a quality score for each ad and ranks the advertisers by their bids and quality scores. For the analysis of the weighted scheme for bids, see Refs. [21,22].

[3] Even though a keyword auction is in practice a sealed-bid auction, services provided by search engines such as "Bid Simulator", might make the advertisers' environment similar to an open-bid environment. However, the information supplied by the Bid Simulator is the average price of each ad slot in the past and not the current bids of competitors.

[4] They call this forward-looking behavior.

[5] In actual, this is not a restriction, because when $N < K$, it suffices to redefine K by $K = N$.