



Contents lists available at ScienceDirect

Electronic Commerce Research and Applications

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

The value of location in keyword auctions

Maurizio Naldi^{a,*}, Giuseppe D'Acquisto^b, Giuseppe Francesco Italiano^a^a Università di Roma "Tor Vergata", Dipartimento di Informatica, Sistemi e Produzione, Via del Politecnico 1, Rome, Italy^b Garante per la Protezione dei Dati Personali, Piazza di Monte Citorio n. 121, Rome, Italy

ARTICLE INFO

Article history:

Received 20 November 2008

Received in revised form 7 September 2009

Accepted 7 September 2009

Available online 15 September 2009

Keywords:

Search engines

Keywords

Auctions

Sponsored search

Advertising

ABSTRACT

Sponsored links on search engines are an emerging advertising tool, whereby a number of slots are put on sale through keyword auctions. This is also known as *contextual advertising*. Slot assignment and pricing in keyword auctions are then essential for the search engine's management since provide the main stream of revenues, and are typically accomplished by the Generalized Second Price (GSP) mechanism. In GSP the price of slots is a monotone function of the slot location, being larger for the highest slots. Though a higher location is associated with larger revenues, the lower costs associated with the lowest slots may make them more attractive for the advertiser. The contribution of this research is to show, by analytical and simulation results based on the theory of order statistics, that advertisers may not get the optimal slot they aim at (the slot maximizing their expected profit) and that the GSP mechanism may be unfair to all the winning bidders but the one who submitted the lowest bid.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Search engines, such as Google or Yahoo!, present users with a set of hyperlinks in response to their queries. In addition to the links deemed relevant to the query by the search engine (often named *organic links*), a number of sponsored links are presented as well (Battelle 2005), associated to the query through the keywords specified in it. The addition of sponsored links represents a form of *contextual advertising* (Kenny and Marshall 2000, Peyrat 2009) whereby advertisements are proposed to customers on the basis of the content that is being displayed to them.

Advertisers are willing to pay for their ads to appear on the search engine's response page. Such sponsored links are then generally assigned through auctions, and the resulting revenues represent a significant source of income for search engines (Edelman et al. 2007). On the other hand, advertisers are interested in making their advertising strategy as effective as possible. Such auctions have been studied for some time now in the context of game theory, where the players in the game are the auctioneer (the search engine) and the advertisers (see e.g., Varian 2007). The related studies have been devoted mainly to examine if, and under which conditions, the game exhibits a Nash equilibrium. Such conditions are typically linked to the assignment and pricing rule on one side and to the advertisers' bidding strategies on the other side

(in the following the terms *advertiser* and *bidder* will be used interchangeably).

As to the former issue, the Generalized Second Price (GSP) rule has reached a wide consensus; in GSP the bidders are ranked in descending order according to their bids, and the *slots* containing the sponsored links are assigned to the bidders submitting the largest bids. The price is set so that each bidder pays for each click on its hyperlink a price equal to the next lower bid. When the advertiser pursues a strategic goal, it has to define its bidding strategy, but the GSP mechanism is not helpful for this purpose. In fact, GSP has been shown not to be an incentive-compatible mechanism (it does not induce bidders into bidding their actual valuation), since it leaves the field open as to the bidding strategy for the advertiser as long as the bid is lower than the valuation of the slot on sale.

Recently the *balanced bidding* (BB) strategy has been proposed by Cary et al. (2008) for such purpose, where advertisers update their bid at each auction round by exploiting the intelligence gathered in the previous rounds. In this process each advertiser identifies at each round its optimal slot as that maximizing its profit. In Cary et al. (2008), the convergence to a Nash equilibrium of the game where users apply the BB strategy has been studied, but the optimal slot determination process has not been explored in detail, though its relationship with the subsequent slot assignment is central to the advertiser's satisfaction.

Actually, in the GSP mechanism the price paid per click by advertisers is linked to the slot position, the highest slots achieving prices that are higher than the lowest positioned ones. The rationale for such correspondence between location and price is that

* Corresponding author.

E-mail addresses: naldi@disp.uniroma2.it (M. Naldi), g.dacquisto@garanteprivacy.it (G. D'Acquisto), italiano@disp.uniroma2.it (G.F. Italiano).

top-positioned slots receive more clicks (hence they are assumed to generate more sales and revenues for the advertiser). However, when the advertiser has to evaluate the return on its advertising investment, it has to consider costs as well as revenues, that is, it has to take into account that (clicks on) top-positioned slots cost more. When balancing costs and revenues (or clicks, playing the role of proxy for the revenues), the advertiser may then find out that top-positioned slots are not the most convenient for it. It has been argued, on the basis of a behavioural model of the search engine user, that this may indeed be the case (Agarwal et al. 2008). Supporting evidence is therefore needed to investigate the fairness of GSP to advertisers.

In this paper we explore the issue of the actual value associated to slots in the context of both *balanced bidding* and *truthful bidding* (i.e., when the bidder submits its true valuation of the item for sale). Our aim is to assess if the slot pricing mechanism in the current GSP approach is aligned with the advertisers' expectations. Rather than assuming a particular navigational behaviour for the user, as is done in Agarwal et al. (2008), we rely on the concept of profit for the advertiser. We also exploit a variety of probability models to describe the preference of users for slots (embodied by the click-through rate), and the private valuations assigned by advertisers to clicks. In addition, though Bakos and Katsamakis (2008) have shown that the presence of the intermediary providing the platform for the exchange (in our case the exchange consists in bids vs. slots) has relevant effects on the value distribution in the resulting two-sided network (the bidders on one side and the seller on the other), we assume that either the intermediary is transparent to the business process or that it coincides with the auctioneer.

Through both analytical derivations and simulation, we provide the following results in this article: (1) we provide closed form expressions for the profit function for a large number of scenarios; (2) we indicate the location of the optimal slot (i.e., that maximizing the advertiser's profit); and (3) we provide evidence of the mismatch between the optimal slot and the location actually assigned to the same advertiser by the GSP mechanism. We show that bidders pay the large prices associated with higher slots while their profit might be maximized by lower-positioned ones, hence the pricing mechanism embedded in GSP may penalize all but the winning bidder with the lowest bid. The consequences of this finding may be far reaching. In fact the GSP pricing model would be unfair to top bidders, that end up paying too much for what they receive in return. A new pricing model, taking into account the actual slot value rather than its location only, is then needed.

2. Fundamentals of keyword auctions

2.1. Sponsored search advertising

Search engines act in response to users' queries for websites containing the information of interest. In such queries the information of interest is synthetically expressed as a string of keywords, possibly connected through Boolean operators. For example, in response to the query "sea AND winds NOT ice" the search engine will return pointers to all the documents related to the first two terms but not to the third one. The hyperlinks returned by the search engine are typically named organic links. The search engine can add to this list (and show, e.g., on the right-hand side of the screen) a number of sponsored links, which typically accomplish an advertising task (hence the name sponsored search advertising) (Jansen and Mullen 2008). The available positions for sponsored links are named *slots*. Such links are provided by advertisers, who are willing to pay to have their ad appear on the screen in relation to a query containing a specific keyword. Hence, for any query, there are a number of potential fillers of the screen space devoted to

sponsored links. It is assumed that the advertisers choose to acquire keywords that are actually related to their product.

The payment rules may be freely defined in the contract relationship between the search engine manager and the advertiser. The pricing schemes for advertising may be classified under three classes (Hoffman and Novak 2000), respectively, based on exposure, click-through, and outcomes (e.g., the actual purchase), corresponding to different stages of the purchase process (and to different scenarios of risk sharing between the auctioneer and the advertisers). While schemes based on outcomes, where the advertiser payment is related to the final result of the advertising process, may appear as the most natural ones to evaluate the effectiveness of advertising, they require the search engine owner to have visibility over the whole process through the completion of the order. They may be difficult to apply when the aim of the advertiser is not selling a good or a service but rather increasing the awareness, for example, of a brand or an event. In the *click-through-based schemes* (also known as *pay-per-click business models*) the advertiser pays a pre-determined amount of money each time a user actually clicks on the ad. Instead, in the *impression-based schemes* the advertiser is charged for the number of times the user is exposed to the advertisement. The main drawback of the pay-per click model is that it may be fooled by automated agents repeatedly clicking on an ad in order to increase the payments due by the advertiser. Learning algorithms have however been proposed to fight against such click frauds (Immorlica et al. 2005). Despite such problems the pay-per-click model appears to be the most widely used, for example, by Google and Yahoo!'s Overture.

Since the number of slots is generally smaller than the number of interested advertisers (i.e., advertisers who have opted to run for a keyword appearing in the query), slots represent a scarce resource and a natural way to assign them to the advertisers is through auctions, namely keyword auctions (Muthukrishnan 2009). Hence, advertisers declare how much they are willing to pay for a click, and an auction is run for the slots among the advertisers whose keywords match the query. Though, as mentioned before, queries may be conducted for quite complex combinations of keywords, here we consider auctions conducted for a single keyword, though the same advertiser may take part to different auctions independently run for different keywords. We make no attempt to compare auctions for different keywords (that may have a different popularity), or for logical combinations of keywords (e.g., AND or OR), though the number of bidders (which we consider as a parameter in the investigation reported here) may be taken as a proxy for the popularity of the keyword. We have therefore a number of slots $S \in \mathbb{Z}^+$ and a larger number of advertisers $A \in \mathbb{Z}^+$, with $A > S$. Actually, a new auction is run every time a query is submitted, among the advertisers submitting bids for keywords matching the query. For any given keyword we have then a sequence of auctions. As will be seen in Section 3, the repetition of auctions allows advertisers to update their bids by taking into account their past observations of the collective behaviour of the other bidders and of the output of previous auction runs.

In order to make the assignment process as effective as possible, the auctioneer has to carefully design the auctioning rules, which consists of choosing: (a) the assignment rule (i.e., the way advertisers are assigned the slots); and (b) the price setting rule (i.e., the price an advertiser has to pay when a user clicks on its ad).

As to the first issue, this is solved by using a straightforward ordering of slots and advertisers. Slots are indexed progressively by their vertical position on the screen: the slot appearing on top of the screen is assigned index 1 by convention; the slot appearing on the bottom of the ad-devoted space has index S .

If we now denote by b_i the bid submitted by the i th advertiser, and then by $b_{(j)}$ the j th highest bid, the assignment rule states that

the k th slot is assigned to the advertiser submitting the bid $b_{(k)}$. Other ranking mechanisms are possible: a family of ranking mechanisms is examined by Lahaie and Pennock (2007), where the rank depends on the product of the bid and some weight function, while in Vorobeychik and Reeves (2008), the two basic mechanisms of ranking by bid and ranking by revenue are considered. Search engines, such as Google or Yahoo!, are known to deviate from this standard assignment scheme by using additional information (Jansen and Mullen 2008). For example, Google computes a quality score that, in addition to the bid amount, also considers the advertisement's click-through rate, the keyword relevancy along with landing-page and site. However, the details about the quality score computation are not known in the clear. Hence, in this paper we restrict ourselves to the analysis of bid-based ranking, as has been done in the relevant literature (e.g., Varian 2007), and in the seminal paper on the BB strategy that we examine here (Cary et al. 2008).

Setting the price is a less straightforward matter. A well-known mechanism is the truthful Vickrey–Clarke–Groves rule, whereby each advertiser is charged for the harm it causes to the other advertiser; namely, the cost for the i th advertiser when assigned an item is the difference between the social value should that advertiser be removed from the auction and the social value should both that advertiser and that item be removed from the auction. The VCG scheme is known to be truthful, that is, it will lead each participating advertiser to bid its true valuation (Clarke 1971, Groves 1973). Its properties have been shown to hold under a number of scenarios. For the case of private value, when the value that a given buyer attaches to the good being sold is independent of the information owned by other buyers Myerson 1981, Engelbrecht (1988) extended the revenue-equivalence results (i.e., that different auction formats provide the seller with the same expected revenues), for single-item and multiple-item auctions. In addition, Krishna and Perry (1998) have shown that in a multiple-item auction setting, the Vickrey–Clarke–Groves mechanism maximizes the seller's expected revenue among all efficient auctions (i.e., those maximizing social welfare). Dasgupta and Maskin have shown that the Vickrey auction is still efficient for the case of common values, that is, when one buyer's valuation can depend on the private information of another buyer (Dasgupta and Maskin 2000). However, it exhibits a number of weaknesses which limit its practical application (Rothkopf 2007), among which its vulnerability to various kinds of cheating and the possibility of poor revenues. Hence, search engines do not adopt the VCG mechanism in practice, but rather the Generalized Second Price (GSP) rule.

2.2. The click-through rate

In Section 2.1 we have seen that the slots on sale rank differently in the customers' preferences list. Such preferences can be characterized by the click-through rate (CTR): the CTR θ_i of slot i is the probability that the user clicks on that slot. It can be estimated by dividing the number of users who clicked on an ad in a web page by the number of times the ad was delivered (impressions) (Sherman and Deighton 2001). Larger click-through rates are expected to generate more revenues, since they give rise to larger numbers of customers landing on the advertiser's website. It is generally accepted that the click-through rate depends on the slot's position, namely that it declines as the slot gets lower on the screen, that is, $\theta_i > \theta_{i+1}$, where $i = 1, \dots, S - 1$. This assumption is supported by the statistical data reported in Brooks (2004) pertaining to Google AdWords and Overture Precision Match, where the CTR drops monotonically as we go from the first ranked ad down to the tenth ranked ad. Hence, top-positioned slots are more valuable than bottom-positioned ones. As to the precise shape of the click-through rate decaying function, here we consider a Zipf distribution. This power law distribution has been assumed as

the most intuitive for the distribution of clicks on ads in Balakrishna and Kambhampati (2008). In addition, in Xie and O'Hallaron (2002), Regelson and Fain (2006) the popularity of search terms has been shown, by fitting measurement data, to follow a Zipf distribution. In the Zipf model the probability that the user clicks on the slot j is

$$\theta_j \propto \frac{1}{j^\alpha}, \quad (1)$$

where $\alpha \in \mathbb{R}^+$ is the Zipf parameter. For convenience (with no consequence on the following results) we adopt the normalizing condition $\sum_{j=1}^S \theta_j = 1$, so that we are actually considering the probability of clicking on a specific slot conditioned to the user clicking on a slot (or, alternatively, the user clicks on a slot with probability 1). Though in this paper we implicitly consider the click-through rate being a function of the slot's position only, other authors have considered the more general case of click-through rates being a function of the specific advertiser as well (Feldman and Muthukrishnan 2008, Aggarwal et al. 2006).

2.3. The generalized second price mechanism

In this section we review the basic characteristics of GSP as a price setting mechanism. In GSP the natural assignment rule is maintained whereby the advertiser submitting the k th highest bid $b_{(k)}$ is assigned the k th slot. However the price p_k it pays is equal to the next lower bid, that is, $p_k = b_{(k+1)}$. Advertisers who are not assigned a slot pay nothing. The most important decision advertisers have to take is then to choose their bids. As a reference they have their own private valuation of clicks: in the simplest scenario the i th advertiser values a click worth v_i (i.e., the click value does not depend on the slot position itself and does not vary as the auction is repeated). In general any bid of the generic i th advertiser will satisfy the inequality

$$b_i \leq v_i. \quad (2)$$

The most important property of the VCG mechanism is that it induces the advertiser to declare its private valuation, so that $b_i = v_i$ (truthfulness property). On the contrary, GSP is not truthful, hence advertisers' bids are limited by the above inequality only. If we consider the static game associated to GSP-driven auctions, a Nash equilibrium has been shown to exist (Varian 2007). However, in the dynamic version resulting from the repetition of the auction, bidders can update their bid at each new issue of the auction by taking advantage of the knowledge they have gained from the past auction occurrences.

3. The balanced bidding strategy

In Section 2.3 we have seen that the GSP induces an untruthful behaviour in the bidders, whereby the advertiser is led to submit a bid strictly lower than its valuation. If that's the case, the bidder is free to choose its bid, with the only loose constraint represented by inequality (2). It may then subject its bidding behaviour to a strategic intention, for example, reaching a business goal such as maximizing its profit. As stated in the Introduction, we are interested in the specific strategic bidding behaviour represented by the balanced bidding strategy proposed in Cary et al. (2008). In this section we review its main characteristics.

Since keyword auctions are run continuously, they represent a form of repeated auctions. In a repeated auction the bidder may exploit the information it has gathered in the previous runs of the auction, for example, the ranking of bidders, the assignments and the prices of slots (which in the GSP assignment mechanism are just a shifted replica of the set of submitted bids). At each new run each

bidder can then formulate a best response bid, that is, a bid embodying a best response strategy to all the bids submitted previously by the other bidders. In this context that bidder may assume that all the other bidders will maintain their behaviour, so that they will submit the same bids as in the previous run, and may use that information to pursue its strategic goal. In the original formulation of the BB strategy proposed in Cary et al. (2008), the bidding behaviour of an advertiser, depending on the previous outcomes of the auction, is conditional on the collective behaviour of the other advertisers. Actually, at each round each advertiser gets more information about the value of the slots as perceived by the other advertisers and could therefore update its valuation; in this paper we take instead the same limiting assumption of Cary et al. (2008) that the value of a slot for an individual advertiser stays unchanged over the auction's duration. In particular, the bidder may want to focus its bidding by targeting a specific slot. In that case the range of useful bid values is more restricted than that implied by inequality (2). If the advertiser is targeting slot k , whose price has been set at p_k , in order to win that slot it will deem useful to submit any bid in the (p_k, p_{k-1}) range (for the advertiser targeting the top slot the useful range is that defined by just the lower bound p_1). In fact, under the hypothesis that all the other bidders maintain their bids, if the bidder submits a bid lower than p_k it will be ranked not higher than in the $(k + 1)$ th position and will not be awarded the k th slot. On the other hand, if the bidder submits a bid larger than p_{k-1} it will be ranked at least in the $(k - 1)$ th position and will be awarded a higher-positioned slot than the k th it had targeted. Hence the general best response strategy does not suggest a specific value for the bid to be submitted, but rather a range of values.

A specific instance of a best response strategy is the balanced bidding (BB) strategy proposed in Cary et al. (2008) and independently analyzed by Vorobeychik and Reeves (2008). In BB the advertiser chooses its next bid b so as to be indifferent between successfully winning the targeted slot k at the price p_k , or winning the slightly more desirable slot $k - 1$ at price b . In a scenario where a set of A advertisers, who compete for S slots and have their private valuations $\{v_1, v_2, \dots, v_A\}$ for a click, are assigned the S slots according to the GSP rule, the BB strategy leads the generic winning i th advertiser to:

1. target the slot k_i^* that maximizes its profit (optimal slot):

$$k_i^* = \operatorname{argmax}_k \{\theta_k [v_i - p_k]\}, \quad (3)$$

where the function $z = \operatorname{argmax}_x \{y(x)\}$ returns the value of the argument x that maximizes the function $y(x)$,

2. set its next bid b according to the expression

$$b_i = v_i - \frac{\theta_{k_i^*}}{\theta_{k_i^*-1}} [v_i - p_{k_i^*}]. \quad (4)$$

For losing advertisers the BB strategy instead prescribes a truthful behaviour, so that they keep submitting bids equal to their valuations. In the first run of the auction, where the advertisers have no information about each other's bids, we assume that each bidder submits a bid equal to a fraction of its valuation, $b_i = \zeta v_i$, with $\zeta < 1$. Since the auction runs continuously, at each new run the bidders have the chance to update their bid. We can therefore envisage two alternative updating models for BB: the synchronous version and the asynchronous one. In the synchronous version, at each auction repetition all the bidders update their bid. In the asynchronous version, at each auction repetition just one of the bidders is allowed to update its bid. If the asynchronous version is played in a deterministic fashion (as in a basic polling mechanism), it therefore takes A auction rounds to have all the bidders update their bids (at each run the set of bid is made of 1 new bid and $A - 1$ stale bids).

The convergence of the BB strategy has been studied by Cary et al. (2008). It has been shown that the dynamic system where all bidders play this strategy in the asynchronous way converge to a unique fixed point, which is also the Nash equilibrium of the static game (Cary et al. 2008), while the convergence for the synchronous updating model is guaranteed just for a limited number of slots. However, the convergence time depends on the number of bidders and may take some hundreds of auction runs, over 300 in the simulation run in Cary et al. (2008). The convergence speed is of course quite faster in the synchronous version: in Naldi and D'Acquisto (2008b), it has been shown, by simulation, that in the synchronous version the steady state is reached within some tens of runs. The convergence to an equilibrium identical to that of the static game is accompanied by the property that advertisers' payments to the search engine in the steady state are identical to payments under the VCG mechanism.

4. Valuations and profit

4.1. Distribution of valuations

A fundamental role in the advertiser's decisions about its bidding behaviour is played by the value it attributes to a click. This value represents the upper bound on the bid that the advertiser submits and hence impacts on the profit the advertiser gains from each click. It is generally assumed that each advertiser has its own private valuation, that is, the value of a click for an advertiser is different from those of the other advertisers and unknown to them. We assume that those values follow a probability distribution, such that the values attributed by any advertiser are drawn from a common probability distribution. In this section we present the probability models we adopt in the following to evaluate the profit gained by advertisers. It is to be remarked that the keyword auction we examine in this paper is not a common value auction. In common value auctions the value of the unit for sale is the same for all the bidders, though none knows it exactly and tries to estimate it (Kagel and Levin 2002). Here (in our keyword auction) each bidder has its own valuation of a click, though all the valuations are drawn from the same probability distribution. In common value auctions the well known phenomenon of the winner's curse is present, whereby winners' estimates are overly optimistic and lead them to overbid. Here bidders are not competing for the same goal, since we have a multi-unit auction (the objects for sale are slots, that are different from one another by their position on the screen), and even non top bidders may end up with the slot they aimed for.

To the best of our knowledge, no statistical results are available to endorse a specific probability model for the value of a click in keyword auctions. Instead, some statistical analyses have been carried out for bids (rather than valuations) in the context of procurement auctions, suggesting mostly the normal probability model (McCaffer and Pettitt 1976, Pin and Scott 1994). In order to reach more reliable results, here we consider instead a variety of models, following the selection done in Naldi and D'Acquisto (2008a). Namely, we use the following models:

- Uniform;
- Normal;
- Exponential;
- Pareto.

These four models allow us to represent a large number of different advertisers' beliefs and can in turn be grouped into two classes. In the first one, including the uniform and the normal model, the valuations concentrate around a central value and may be considered to represent the case where the advertisers share approx-

imately a common opinion about the value of a click. Instead, in the exponential and Pareto models, making up the second class, lower valuations predominate over larger ones, and wide differences may appear in the value attributed to a click.

We then consider A random variables $\{V_1, \dots, V_A\}$, where V_i describes the value of a click for the i th advertiser. We assume that these valuations are independent and identically distributed, following the common probability distribution function (cdf) $\mathbb{P}[V_i < x] = F_V(x)$, $i = 1, \dots, A$. In the following four subsections we describe the cdfs pertaining to the four models named above.

4.2. Uniform model

The probability distribution function is

$$F_V(x) = \begin{cases} 0 & \text{if } x < a \\ \frac{x-a}{b-a} & \text{if } a \leq x < b \\ 1 & \text{if } x \geq b. \end{cases} \quad (5)$$

In the uniform model the valuations are always bounded by the lower bound a and the upper bound b .

4.3. Normal model

The probability distribution function is

$$F_V(x) = \int_{-\infty}^x \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(y-\mu)^2}{2\sigma^2}\right] dy, \quad (6)$$

and depends on the mean value μ and on the standard deviation σ , which describes the degree of scattering around the central value. This model would allow for negative values (absent in real auctions) but not for a minimum bid (present in the other models). In order to reduce the effect of these two shortcomings we introduce two constraints:

1. We choose $(\mu; \sigma)$ pairs such that the probability of having negative values is extremely low;
2. We introduce a conventional definition of the range of values (theoretically infinite in the normal model) as six-times the standard deviation so that the probability of having values outside this range is again very low.

The second condition can be expressed as the following relationship between the standard deviation of the normal model and the conventional values for the highest and the lowest bid (introduced for the uniform model) (Naldi and D'Acquisto 2008a)

$$6\sigma = b - a. \quad (7)$$

4.4. Exponential model

The probability distribution function is

$$F_V(x) = 1 - e^{-\lambda(x-a)}, \quad x \geq a, \quad (8)$$

which depends again on two parameters, the scale factor λ and the lower bound a .

4.5. Pareto model

The probability distribution function is

$$F_V(x) = 1 - \left(\frac{a}{x}\right)^\beta, \quad x > a, \quad (9)$$

where a acts both as the minimum valuation and as a scale factor, and β is the shape parameter. Though low values are commoner than large ones, with the value of the shape parameter dictating

the relative weight of lower bids, very large valuations are possible. Lower values of β correspond to a fatter tail of the probability density function, but if $\beta < 1$ not even the first moment of the valuation is finite. In the following we will therefore consider $\beta > 1$.

4.6. Advertisers' profit

As stated in the Introduction, any figure of merit for the advertising investment of the bidder has to take into account the costs as well as the revenues incurred. In this paper we incorporate both quantities in the profit function. Here we follow the same definition of profit provided by Cary et al. (2008); the authors used the term 'utility' instead, but utility is usually subjective and may not be comparable across different advertisers. Every time a user clicks on the hyperlink pertaining to an advertiser, that advertiser may be deemed to receive a profit gain (profit per click) equal to the difference between the value it attributes to a click and the price it is paying for that click. As previously recalled, here we employ clicks as a proxy for revenues, though not all clicks convert into sales. The imperfect role of the proxy is however a minor problem, since the lower-than-unitary conversion rate may be accounted for in the value attributed to a click: the lower the conversion rate the lower the click value. If we multiply the profit per click by the number of clicks we then get the overall profit for the advertiser. By employing the click-through rate rather than the number of clicks, we obtain the following expression for the expected profit of the i th advertiser receiving the k th slot:

$$U_i(k) = \theta_k [v_i - b_{(k+1)}]. \quad (10)$$

If we introduce the Zipf model defined in Section 2.2 for the click-through rate, the above expression becomes a function of the Zipf parameter as well:

$$U_i(\alpha, k) \propto \frac{1}{k^\alpha} [v_i - b_{(k+1)}]. \quad (11)$$

Here the expected profit is expressed by the product of two terms that move in opposite directions as the target slot lowers. Depending on the balance between the two terms, the profit may then either be a monotone function of the target slot or exhibit a maximum, highlighting an optimal slot different from the top-positioned one.

5. Optimal slot location

When bidding for a strategic goal, such as maximizing their profit, advertisers implicitly look for the slot that best suits them, by expressing the best trade-off between the price they pay (determined by their bid per click and the click-through rate) and the revenue they can expect. These two aspects are summarized in the profit function defined in Section 4.6. In this section we determine the optimal (profit-maximizing) slot both for the case of truthful bidding and for the BB strategy, and examine the mismatch between that optimal slot and the slot the advertiser actually gets.

5.1. Optimal slot location under truthful bidding

An auction mechanism is said to be truthful if the dominant strategy for each bidder is to submit a bid equal to its valuation. The most famous example of a truthful mechanism is the Vickrey-Clarke-Groves (VCG) auction, whose characteristics have been briefly reviewed in Section 2.1. Though this mechanism is not used in the context of keyword auctions, a goal of many research efforts is to design truthful mechanisms (see e.g., Aggarwal et al. 2006). Actually, truthful auctions are of limited interest in auctions that

run repeatedly (such as keyword auctions). In fact, if participants were the same in each repetition, the bids would be the same and the auction results would not change. However, we have to consider that participants do change in the subsequent runs of the auction, so that there is a chance for participants with low valuations to get a slot if new participants have even lower valuations. In addition, inducing truthful bidding has a cost, as shown by Naldi and D’Acquisto (2008a), since, in order to induce bidders into revealing their true valuation, the auctioneer has to relax the payment rule and let bidders pay less than their true valuations, with an impact on its overall revenues. We can however consider truthfulness as a significant property of auction mechanisms, and a useful basis for comparison with other mechanisms. Since, in this paper, we are interested in the class of GSP auctions, we consider the benchmark case where bidders submit their true valuations (truthful bidding), though the assignment is accomplished through the GSP rule. This provides an absolute upper bound for the auctioneer’s profit under the GSP rule, though in a GSP auction bidders would always bid less than their true valuation (hence this upper bound would never be attained), and is therefore a natural benchmark, which allows us to consider in isolation the effect of deviations from truthful bidding that balanced bidding involves.

Since in truthful auctions, bids are equal to valuations, the expected profit $U_i(\alpha, k)$ for the bidder submitting the i th highest bid when obtaining the k th slot is the product of the click-through rate and the expected difference in the valuations of a click. By retaining just the factors depending on the bidder and on the slot location we can write

$$U_i(\alpha, k) \propto \frac{1}{k^\alpha} \mathbb{E}[V_{(i)} - V_{(k+1)}] = \frac{1}{k^\alpha} (\mathbb{E}[V_{(i)}] - \mathbb{E}[V_{(k+1)}]). \quad (12)$$

A general result is that $U_i(\alpha, k) < 0$ when $k < i - 1$ and $U_i(\alpha, k) = 0$ when $k = i - 1$, so that, for the bidder with the i th valuation, only the slots lower than i would exhibit a positive expected profit. Such a condition thus limits the set of choices useful for the advertisers: advertisers with lower valuations can target fewer slots. In this section we provide results for the distribution of preferences for slot locations in truthful auctions under the probability models introduced in Section 4.1.

5.2. Uniform model

For the case of the uniform distribution of valuations, it is straightforward to derive the expected value of the ordered valuation of a click in a set of A advertisers. For the r th valuation we have (Naldi and D’Acquisto 2008a)

$$\mathbb{E}[V_{(r)}] = a + (b - a) \frac{A + 1 - r}{A + 1}. \quad (13)$$

The general expression of the expected profit (12) becomes

$$U_i(\alpha, k) \propto \frac{1}{k^\alpha} \left[\frac{A + 1 - i}{A + 1} - \frac{A + 1 - (k + 1)}{A + 1} \right] \propto \frac{k + 1 - i}{k^\alpha}. \quad (14)$$

It has to be noted that the both the lower and upper bound for valuations, as well as the number of advertisers, do not influence the shape of the profit curve.

We first consider the location of the optimal slot, that is,

$$k_{\max} = \operatorname{argmax}_k U_i(\alpha, k). \quad (15)$$

The optimal location can be easily obtained from Eq. (14). In Fig. 1 we plot the optimal slot for the bidders for the case of 10 slots on sale (quite a reasonable upper bound for the number of slots on sale, considering the available space on the screen) and three different values of the Zipf parameter, namely $\alpha = 0.5, 1, 2$, which span over the single value $\alpha = 1.5$ considered by Balakrishnana and Kambhampati (2008). If the assigned slot (equal to the bid ranking) were optimal,

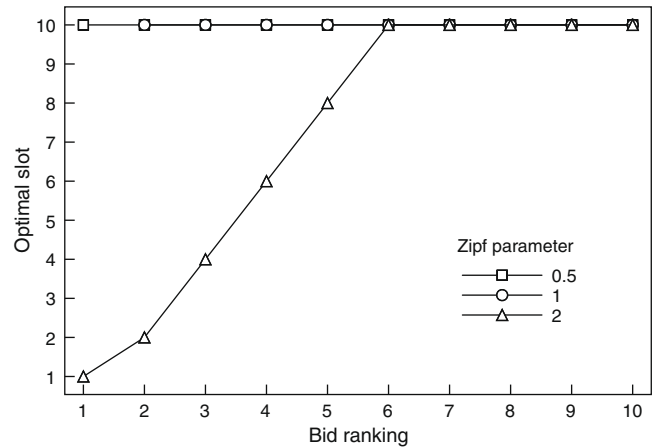


Fig. 1. Optimal slot under the uniform distribution for click valuations.

the curve should be the straight line $k_{\max} = i$. Instead, we can see that, as long as the Zipf parameter is lower than 1, the lowest-positioned slot results to be the profit-maximizing slot for all the bidders. When the Zipf parameter is larger than 1 the preferences distribute more evenly, but the lowest positioned slot still remains the preferred one for most of the bidders. In the case of $\alpha = 2$ we see that the optimal slot is anyway lower positioned than that assigned by the ranking ($k_{\max} > i$), excepting the two bidders submitting the largest bids. We can conclude that the preferences for the lowest slots are actually more marked the lower the Zipf parameter.

We now focus on the advertiser exhibiting the largest valuation and analyse its profit $U_1(\alpha, k)$ as a function of the Zipf parameter in the CTR power law and of the target slot. Eq. (14) representing profit simplifies to

$$U_1(\alpha, k) \propto k^{1-\alpha}. \quad (16)$$

If $\alpha > 1$ this is a decreasing function of the target slot, so that the optimal slot is the top positioned one. If instead $\alpha < 1$ this is an increasing function of the target slot, so that the optimal slot is the lowest-positioned one. Then for the advertiser exhibiting the largest valuation the solution to the optimal slot location problem bounces between the two extremely located slots, depending on the range of the Zipf parameter.

5.3. Normal model

For the Gaussian model no exact general expressions exist for the order statistics. The approximations proposed by David (1981), Cramer (1946) fail when applied to a wide range of orders as in our case. In the Gaussian case we therefore resort to simulation for the evaluation the metrics of interest. For the purpose of simulation we set $\mu = 1$ and $\sigma = 1/6$ in the model of Section 4.3. We begin with plotting the optimal slot locations in Fig. 2. We note a remarkably similar behaviour as in the case of the uniform distribution (the optimal locations for $\alpha = 0.5$ and $\alpha = 2$ are identical to the uniform case). The profit profiles for the bidder submitting the largest bid, shown in Fig. 3, reveal that the top-positioned slot is the preferred one for most values of the Zipf parameter.

5.4. Exponential model

For the exponential model we can use the expression for the expected value of the generic ordered valuation provided in Epstein and Sobel (1953):

$$\mathbb{E}[V_{(r)}] = a + \frac{1}{\lambda} \sum_{j=r}^A \frac{1}{j}. \quad (17)$$

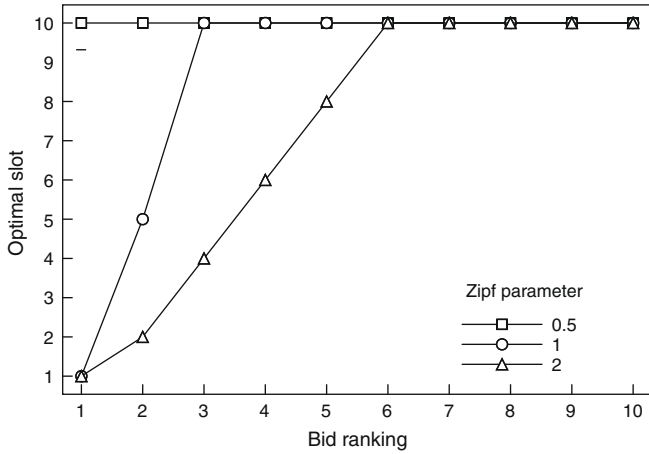


Fig. 2. Optimal slot under the normal distribution for click valuations.

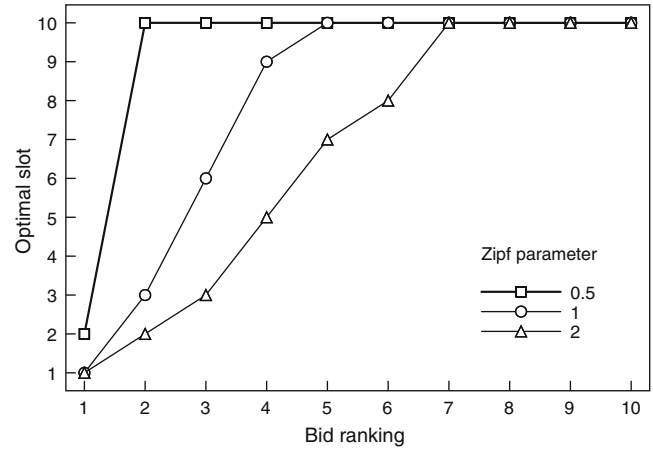


Fig. 4. Optimal slot under the exponential distribution for click valuations.

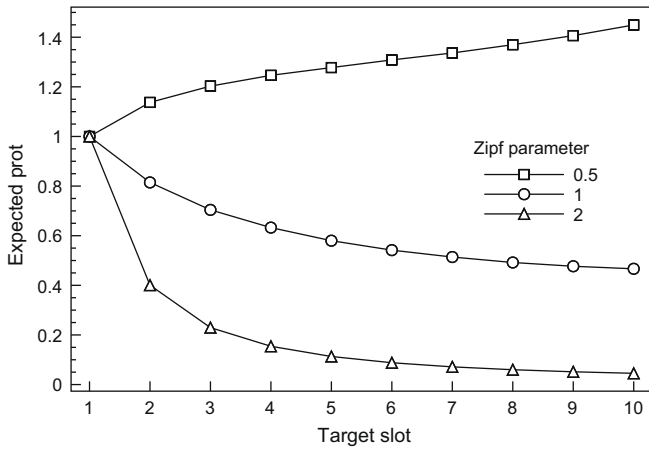


Fig. 3. Expected profit for the largest bid under the normal distribution for click valuations.

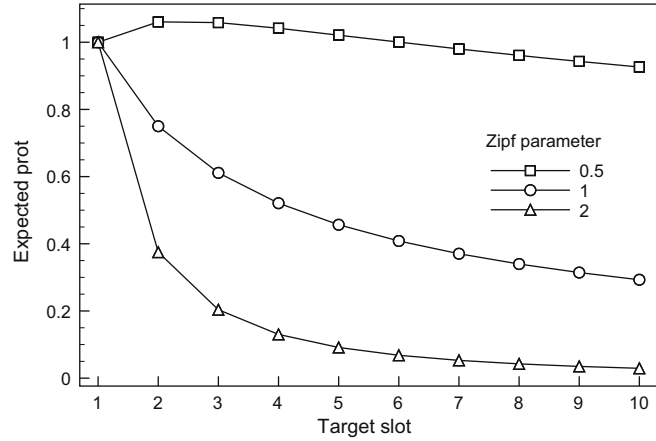


Fig. 5. Expected profit for the largest bid under the exponential distribution for click valuations.

Introducing this expression in the general expression of the expected profit (12) we obtain

$$U_i(\alpha, k) \propto \frac{1}{k^\alpha} \sum_{j=i}^k \frac{1}{j}. \tag{18}$$

Again the profit curve shape results to be independent of the model parameters (in this case the minimum bid a and the scale factor λ) as well as of the number of advertisers. We plot the optimal slot as a function of the bid ranking in Fig. 4 for the case of 10 slots on sale. Though the flight to the lowest positioned slot is less marked than in the case of the uniform distribution, for all but the advertisers with the largest valuations the optimal slot is not that assigned by the bid ranking. And again the phenomenon is more marked as the Zipf parameter gets lower. If we focus on the advertiser with the largest valuation, in the three cases examined its preference goes to a lower slot than that assigned by the ranking (namely the second highest position rather than the top one) just when the Zipf parameter is 0.5, as shown in Fig. 5.

5.5. Pareto model

For the Pareto distribution we employ the expression for the expected value of the generic ordered valuation provided in Malik (1970), Kulldorf and Vannman (1973). For the r th order statistics we have

$$\mathbb{E}[V_{(r)}] = a \frac{A!}{(r-1)!} \frac{\Gamma(r - \frac{1}{\beta})}{\Gamma(A + 1 - \frac{1}{\beta})}. \tag{19}$$

The expected profit is then

$$U_i(\alpha, k) \propto \frac{1}{k^\alpha} \frac{\Gamma(i - 1/\beta)}{(i-1)! \Gamma(A + 1 - 1/\beta)} - \frac{1}{k^\alpha} \frac{\Gamma(k + 1 - 1/\beta)}{k!}. \tag{20}$$

By exploiting some properties of the factorial and of the Gamma function, namely that

$$\Gamma(z + 1) = z\Gamma(z) \quad \forall z \in \mathbb{R}^+, \tag{21}$$

$$k! = (i-1)! \prod_{j=i}^k j$$

and recalling that the useful range of slots is $k \geq i$, we can rewrite the expected profit as

$$U_i(\alpha, k) \propto \frac{1}{k^\alpha} \cdot \frac{\Gamma(i - 1/\beta)}{(i-1)! \Gamma(N + 1 - 1/\beta)} \cdot \left[1 - \prod_{j=i}^k \left(1 - \frac{1}{\beta \cdot j} \right) \right]. \tag{22}$$

Contrary to the other models, now the shape of the profit curve depends on one parameter of the probability model, namely the shape parameter β . However, it is still independent of the minimum valuation. As for the other models we plot in Fig. 6 the optimal slot location when there are 10 slots on sale and the Pareto shape factor

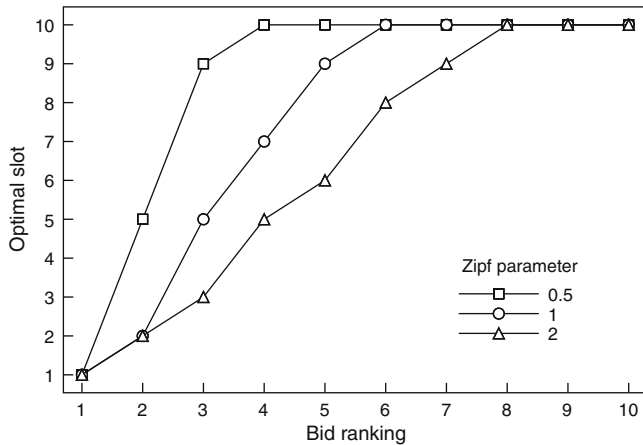


Fig. 6. Optimal slot under the Pareto distribution for click valuations.

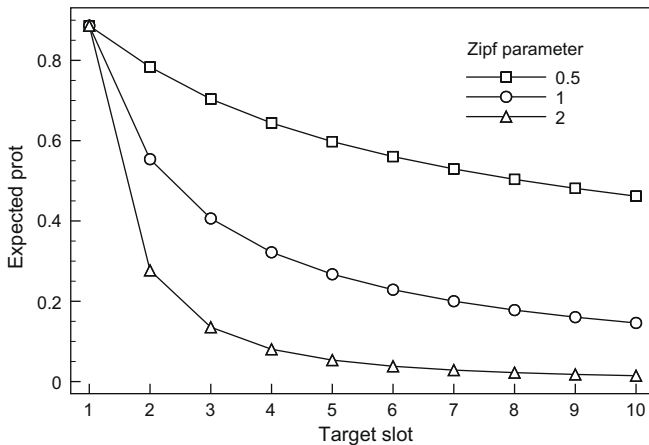


Fig. 7. Expected profit for the largest bid under the Pareto distribution for click valuations.

is $\beta = 2$. We see here a more graceful movement towards the lowest-positioned slot, but the conclusions drawn for other models are confirmed. For the advertiser exhibiting the largest valuation we see instead that the expected profit decreases monotonically with the target slot, hence for this advertiser the top slot appears to be the preferred one for any value of the Zipf parameter, as shown in Fig. 7.

5.6. Optimal slot location under balanced bidding

In Section 5.1 we have examined the position of the optimal slot in the case of truthful bidding. In this section we turn to the balanced bidding to perform the same analysis, namely the estimation of the expected profit and the extraction of statistics on the optimal slot. It has been proven that the BB strategy converges to the same Nash equilibrium as the truthful VCG, hence we expect the results of the analysis in the BB case to confirm the flight to the lowest slots observed before. Since the bid values are continuously changed in subsequent repetitions of the auction under the BB strategy, obtaining analytical results is not feasible. We have therefore resorted to simulation. In this section we provide the set of simulation results for the same quantities and under the same hypotheses considered in the previous section for truthful bidding.

All the results in this section have been obtained by Monte Carlo simulation, consisting in a number N_{sim} of simulation instances. For each simulation instance we have generated the private valuations

of a click for the bidders (a value for each bidder), to be held constant during the whole simulation instance (but to be changed at the following instance). In turn, during each simulation instance we have repeated the auction for a number of runs and updated the bids of all the bidders at each run in the synchronous way. Actually, the set of auction runs making up each simulation instance has been divided into two phases: the first one, made up of N_{tr} repetitions, represents the transient needed to have the BB strategy converge towards a steady state. The auction outcomes in that stage have not been considered for the purpose of estimating the profit, but just for the purpose of updating the bids and advancing the simulation instance into the steady state. After entering the steady state we have performed an additional number N_{reg} of auction runs. In this stage we have evaluated the profit for each bidder and each slot. The estimate of the expected profit for the ordered bidders (i.e., for the bidder submitting the largest bid, for the bidder submitting the next-highest bid, and so on) has been obtained by averaging the profit values over the $N_{sim} \cdot N_{reg}$ useful auction runs (i.e., those pertaining to the steady state). After obtaining one such estimate for each slot we have identified the slot maximizing the expected profit.

All the results presented in the following have been obtained under the same simulation conditions, namely:

- Number of simulation runs $N_{sim} = 10,000$;
- Duration of transient $N_{tr} = 100$;
- Steady state repetitions $N_{reg} = 200$;
- Number of slots on sale $S = 10$;
- Number of bidders $A = 15$;
- Zipf parameter $\alpha = 0.5, 1, 2$;
- First run factor $\zeta = 0.8$.

The number of simulation runs and the number of steady state repetitions have been chosen so to be large enough to warrant a good statistical accuracy. Since Naldi and D'Acquisto (2008b) found that some tens of repetitions are enough to reach the steady state, the number of transient repetitions we have adopted is a safe choice. As to the relationship between the number of slots and the number of bidders, again from Naldi and D'Acquisto (2008b), we know that, as long as the bidders outnumber the slots on sale, the impact of their exact number is negligible. Finally, the range of values considered for the Zipf parameter (the same as in the analysis of truthful bidding) is quite large and includes the single value considered in Balakrishnana and Kambhampati (2008).

Also in the case of balanced bidding strategy, we have considered the four probability models for the bidders' valuations as described in Section 5.1. Though we are working with a scenario of repeated auctions, we assume that the valuation of a click for a given bidder is maintained during the whole set of auction repetitions (while changing with each simulation instance). We analyse the same characteristics as in Section 5.1, including the location of the optimal slot for all the ranked bidders and the profit profile for the bidder submitting the largest bid. For the profit profile we adopt, for reading convenience, a normalization to the profit value obtained for the top-positioned slot. In the following subsections we report the results separately for each probability model.

5.7. Uniform model

In Section 5.2 we have seen that the results obtained analytically for the uniform model in the case of truthful bidding are independent of the minimum and maximum possible bids, that is, the quantities a and b in Eq. (13). We expect the same to be true in the case of BB. However, for the simulation purposes we have to set values for those two quantities. Here we report the results

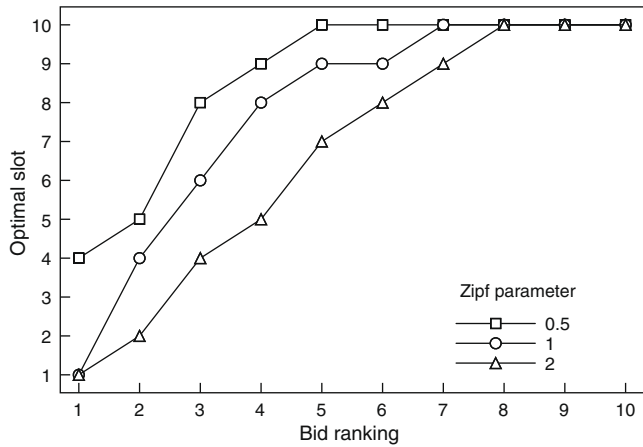


Fig. 8. Optimal slot under the uniform distribution for click valuations (BB).

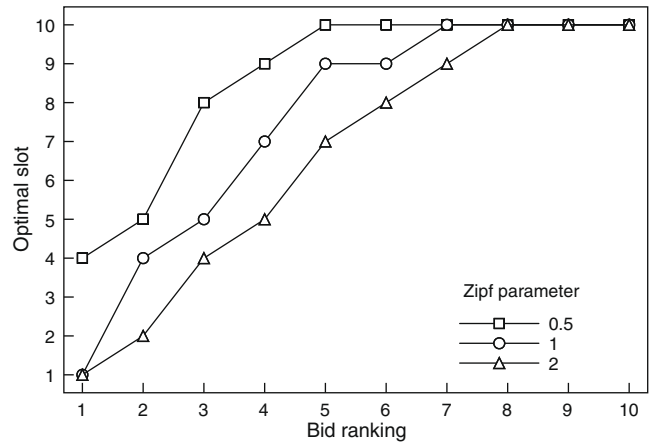


Fig. 10. Optimal slot under the normal distribution for click valuations (BB).

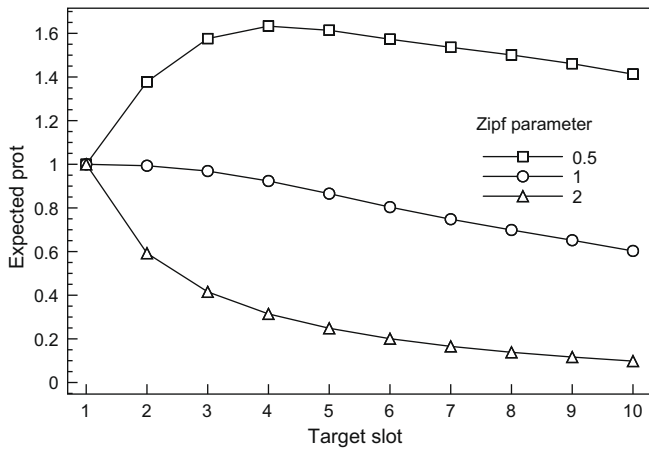


Fig. 9. Expected profit for the largest bid under the uniform distribution for click valuations (BB).

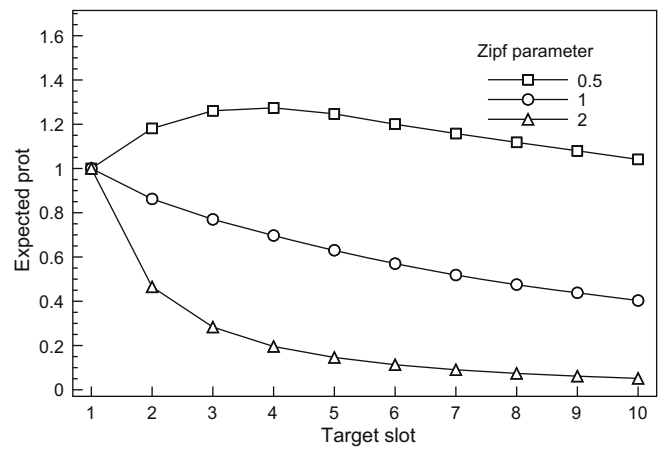


Fig. 11. Expected profit for the largest bid under the normal distribution for click valuations (BB).

obtained for $a = 0$ and $b = 1$. The results obtained for other values of a and b are quite similar. In Fig. 8 we show the location of the optimal slot. Here the flight to the lowest-positioned slots is quite more gradual than in the corresponding Fig. 1, but anyway remarkable: for the lowest value of the Zipf parameter six out of the ten winning bidders have the lowest-positioned slot as their favourite. If we focus on the bidders submitting the largest bid (see Fig. 9) we see again that the unitary value of the Zipf parameter plays a watershed role: when $\alpha < 1$ the optimal slot location becomes progressively lower.

5.8. Normal model

In the case of the normal model we set the same model parameters as in Section 5.3. In Fig. 10 we see a less steep movement of the optimal slot location than in the truthful bidding case depicted in Fig. 10, but the main trend is preserved. For the bidder with the largest bid we see instead that the profit's growth is more marked now that in the truthful bidding case (see the curve plotted for $\alpha = 0.5$ in Fig. 11).

5.9. Exponential model

Again we note that the profit profiles obtained in Section 5.4 depend neither on the expected value of the bids nor on the minimum possible bid value. In the simulation conducted for the balanced bidding strategy, we have set $a = 0$ and $\lambda = 1$ in the model

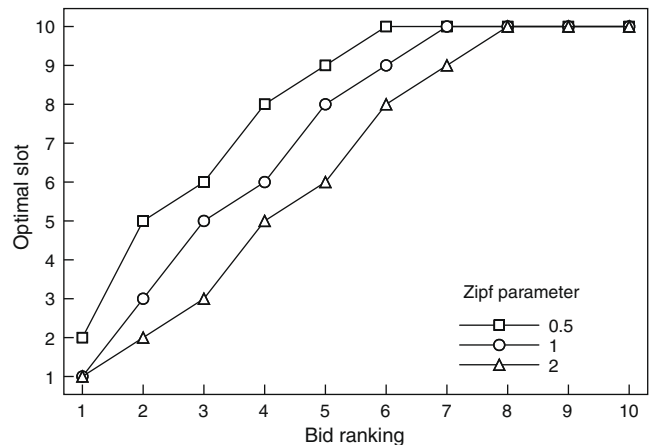


Fig. 12. Optimal slot under the exponential distribution for click valuations (BB).

(17). The optimal slot location is shown in Fig. 12, where the preference for the lowest-positioned slots is confirmed once more. The curves confirm that the trend is more marked as the Zipf parameter gets lower. A quick comparison with Fig. 4 reveals that the trend observed for BB is more gradual than that taking place in truthful auctions: for the three values of the Zipf parameter we note that in the case of BB the landing on slot 10 (the lowest-positioned slot)

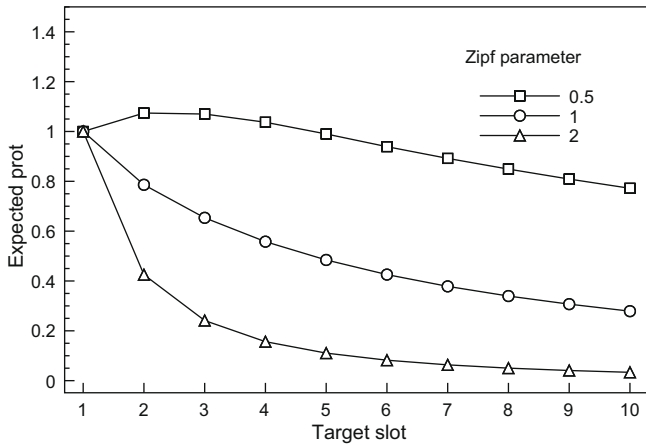


Fig. 13. Expected profit for the largest bid under the exponential distribution for click valuations (BB).

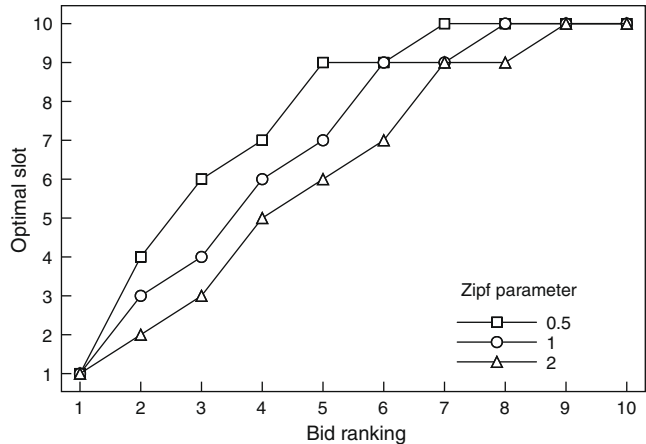


Fig. 14. Optimal slot under the Pareto distribution for click valuations (BB).

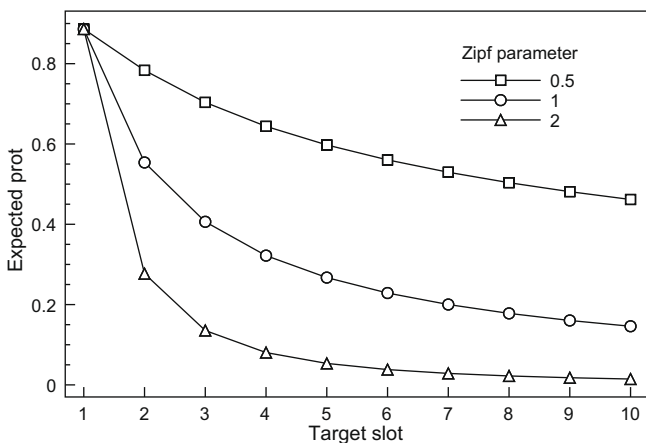


Fig. 15. Expected profit for the largest bid under the Pareto distribution for click valuations (BB).

takes place first, respectively, for the bidders ranking 6, 7, and 8, rather than for the bidders ranking 2, 5, and 7 of the truthful case. In Fig. 13 we show some profit profiles for the bidders submitting the largest bid, which resemble very closely those sketched for the truthful case in Fig. 5.

5.10. Pareto model

In the case of the Pareto model we have to set both the scale and the shape factor. As the reader will recall from Section 4.5, both parameters are subject to constraints. The scale factor (which also plays the role of minimum possible valuation) a has to satisfy the inequality $a > 0$, and we must set the shape factor $\beta > 1$ to have a finite expected value. In our simulation study we have set $a = 1$ and $\beta = 2$, in accordance with the value adopted in Section 5.5. We recall that in that same section it has been shown that the profit profile does not depend on the scale factor, but just on the shape factor. The analysis of Fig. 14 shows that the influence of the Zipf parameter acts in exactly the same way as seen in all the previous cases: lower values of α lead to a faster movement of preferences towards the lowest-positioned slots. And the trend is confirmed to be more gradual in the BB case than in truthful bidding (see Fig. 6 for sake of comparison). If we turn to the profit profile for the bidder submitting the largest bid, we see that, unlike in the uniform, normal, or exponential model, in the Pareto model the top-positioned slot is always the most preferred one (the profit function is a monotonically decreasing function of the target slot) (see Fig. 15).

6. Implications

In Section 5 we have found out that a mismatch exists between what would be optimal for the advertisers and what they actually get. In most cases, purely profit-maximizing advertisers would prefer lower positioned slots than those they actually get. Hence, the current assignment and pricing mechanisms may be considered unfair to all advertisers excepting the one submitting the lowest bid among the winners. In this section we discuss the major implications of these findings.

In the current configuration of displaying mechanisms the advertisements are textual and they differ from one another just by their position, hence their value is a function of their position only. This is reflected in the price paid by advertisers, which depends on the CTR, which in turn depends on the position only. Our findings show that the degree of unfairness is linked to the relevance of the position on the popularity of the slot (i.e., the value of the slope parameter governing the Zipf law that models the decay of the CTR with the slot position). Hence, the tight dependence of the auction mechanisms on the position is a major factor in the fairness of the auction.

The auctioneers who manage the sponsored space associated with the search engine have then to revisit their assignment and pricing mechanisms to make them fairer to bidders; at the same time they wish to safeguard their revenues. A possible solution for the first goal is to relax the position-based mechanisms and introduce a pricing rule based on other factors in addition to position. These factors may even add value to the sponsored link. A recently proposed solution is to drive away from pure textual advertising, for example, by introducing video and images. This feature is not new to web-based advertising since banners, which have a predominant graphic characteristic, have been present since long. Well known search engines appear to move in that direction, as testified by the “Rich Ads In Search” pilot service of Yahoo! and “PlusBox” of Google. The use of videos and images also increases the appeal of the ad and hence its click-through rate, thus satisfying also the second goal of auctioneers.

7. Conclusions

The value of slot locations in keyword auctions has been examined using the profit function under a variety of probability models

for the users' preferences for slots and for the valuation of clicks as assessed by the advertisers. Analytical and simulation results have been obtained for the profit profile and the optimal slot location, with consistent conclusions for the cases of truthful bidding and strategic bidding employing the BB strategy. In the overwhelming majority of cases a slot lower than the assigned one results to be the optimal one, in other words, that maximizing the expected profit of the advertiser. The bottom-positioned slot appears to be the most preferred one, thanks to its low price that compensates for the lower proportion of clicks. The entity of the effect depends on the decay rate of the click-through and is different for the two types of bidding behaviour under examination. As expected, the flight to lower-positioned slots is more marked for the lower values of the Zipf parameter (i.e., when the click-through rate gets flatter). A less obvious finding is that the sliding of the optimal location towards the bottom slot is more graceful in the balanced bidding case than in truthful bidding. The overall consequence is that in most cases advertisers pay more for something that is not the optimal solution for them. The lack of alignment between optimal and actual assignment appears to be significant especially for the bidders ranking in the mid-position. These results corroborate the view that the GSP pricing rule is generally unfair to advertisers and that a revision of pricing rules could help align the interests of the search engine owners with those of the advertisers. Possible moves for auctioneers towards a larger degree of fairness, and possibly larger revenues, envisage the abandonment of tightly position-based mechanisms, for example, through the use of non-textual (employing images and video) advertising.

Some lines of research can be envisaged to extend the domain covered by this paper, among which the most prominent concern the dynamics of the population of bidders and the initial bidding behaviour. As to the former issue, we have considered a fixed number of bidders, while in reality we will have to consider that the number of bidders varies during the repetitions of the auction, with losing bidders reneging to submit new bids and prospective bidders entering the auction. As to the latter issue, we have considered so far that the initial bid (for which strategic bidding provides no prescriptions and bidders may not exploit the information derived from the observation of past auctions) is a fixed percentage of the true valuation, while other strategies may be explored (e.g., submitting one or more random bids with a prescribed upper bound).

Acknowledgements

The authors wish to thank the Editor-in-Chief, the three Guest Editors and the anonymous reviewers.

References

- Agarwal, A., Hosanagar, K., Smith, M. D. Location, location, location: an analysis of profitability of position in online advertising markets. Working paper 1151537, Social Science Research Network, November 26, 2008. Available at <papers.ssrn.com/sol3/papers.cfm?abstract_id=1151537>. Last accessed August 2, 2009.
- Aggarwal, G., Goel, A., and Motwani, R. Truthful auctions for pricing search keywords. In Proceedings of the Seventh ACM Conference on Electronic Commerce, New York, 2006, NY, 1–7.
- Bakos, Y., and Katsamakos, E. Design, ownership of two-sided networks: implications for Internet platforms. *Journal of Management Information Systems*, 25, 2, 2008, 171–202.
- Balakrishnan, R., and Kambhampati, S. Optimal ad-ranking for profit maximization. In Proceedings of the 11th International Workshop on the Web and Databases, Vancouver, BC, Canada, June 13, 2008, ACM Press, New York, NY, 2008.
- Battelle, J. *The Search: How Google and its Rivals Rewrote the Rules of Business and Transformed Our Culture*. Penguin Group, New York, NY, 2005.
- Brooks, N. The Atlas rank report: how search engine rank impacts traffic. Technical report, Atlas Institute, Microsoft, Seattle, WA, July 2004. Available at <www.atlasonsolutions.com>. Last accessed August 2, 2009.
- Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A. R., Mathieu, C., Schwarz, M. On best-response bidding in GSP auctions. Working paper no. W13788, National Bureau of Economic Research, Cambridge, MA, February 2008.
- Clarke, E. H. Multipart pricing of public goods. *Public Choice*, 11, 1, 1971, 17–33.
- Cramer, H. *Mathematical Methods of Statistics*. Princeton University Press, Princeton, NJ, 1946.
- Dasgupta, P., and Maskin, E. Efficient auctions. *Quarterly Journal of Economics*, 115, 2, 2000, 341–388.
- David, H. A. *Order Statistics*. John Wiley and Sons, New York, NY, 1981.
- Edelman, B., Ostrovsky, M., and Schwarz, M. Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *American Economic Review*, 97, 1, 2007, 242–259.
- Engelbrecht-Wiggans, R. Revenue equivalence in multi-object auctions. *Economics Letters*, 26, 1, 1988, 15–19.
- Epstein, B., and Sobel, M. Life testing. *Journal of the American Statistical Association*, 48, 263, 1953, 486–502.
- Feldman, J., Muthukrishnan, S. Algorithmic methods for sponsored search advertising. Paper 0805.1759, Computer Science and Game Theory Section, arXiv, May 12, 2008. Available at <arxiv.org/abs/0805.1759>. Last accessed August 2, 2009.
- Groves, T. Incentives in teams. *Econometrica*, 41, 4, 1973, 617–631.
- Hoffman, D. L., and Novak, T. P. Advertising and pricing models for the Web. In D. Hurley, B. Kahin, and H. Varian (eds.), *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*, MIT Press, Cambridge, MA, 2000.
- Immorlica, N., Jain, K., Mahdian, M., Talwar, K. Click fraud resistant methods for learning click-through rates. In C. Papadimitriou, S. Zhang (eds.), Proceedings of the First Workshop on the Internet and Economics, December 15–17, 2005, 34–45. Also published in C. Papadimitriou, S. Zhang (eds.), *Lecture Notes in Computer Science 5385*, Springer, New York, NY.
- Jansen, B. J., and Mullen, T. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business*, 6, 2, 2008, 114–131.
- Kagel, J. H., and Levin, D. *Common Value Auctions and the Winner's Curse*. Princeton University Press, Princeton, NJ, 2002.
- Kenny, D., and Marshall, J. Contextual marketing: the real business in the Internet. *Harvard Business Review*, 78, 6, 2000, 119–125.
- Krishna, V., Perry, M. Efficient mechanism design. Working paper, Center for the Study of Rationality, Hebrew University of Jerusalem, Israel May 1998. Also SSRN No. 64934. Available at <papers.ssrn.com/sol3/papers.cfm?abstract_id=64934>. Last accessed August 2, 2009.
- Kulldorf, G., and Vannman, K. Estimation of the location and scale parameters of a Pareto distribution by linear functions of order statistics. *Journal of the American Statistical Association*, 68, 341, 1973, 218–227.
- Lahaie, S., and Pennock, D. M. Revenue analysis of a family of ranking rules for keyword auctions. In Proceedings of the Eighth ACM conference on Electronic Commerce, ACM Press, New York, NY, USA, 2007, 50–56.
- Malik, H. J. Exact moments of order statistics from the Pareto distribution. *Skandinavisk Aktuarietidskrift*, 53, 1–2, 1970, 6–9.
- McCaffer, R., and Pettitt, A. N. Distribution of bids for buildings and roads contracts. *Operations Research Quarterly*, 27, 4, 1976, 835–843.
- Muthukrishnan, S. Internet ad auctions: insights and directions. In Proceedings of the 35th International Colloquium on Automata, Language and Programming, Reykjavik, Iceland, July 7–11, 2008, 14–23. Also available in Lecture Notes on Computer Science 5125 and 5126, Springer, New York, NY, 2009.
- Myerson, R. Optimal auction design. *Mathematics of Operations Research*, 6, 1, 1981, 58–73.
- Naldi, M., and D'Acquisto, G. Performance of the Vickrey auction for digital goods under various bid distributions. *Performance Evaluation*, 65, 2008, 10–31.
- Naldi, M., D'Acquisto, G. Transient simulation of balanced bidding in keyword auctions. In: The Proceedings of the 20th European Modeling and Simulation Symposium, Amantea, Italy, September 17–19, 2008, 487–492.
- Peyrat, M. La publicité ciblée en ligne (Online advertising). Technical report, Commission Nationale de l'Informatique et des Libertés, Paris, France, February 5, 2009. Available at <www.cnil.fr>. Last accessed August 2, 2009.
- Pin, T. H., and Scott, W. F. Bidding model for refurbishment work. *Journal of Construction and Engineering Management*, 120, 2, 1994, 257–273.
- Regelson, M., Fain, D. C. Predicting click-through rate using keyword clusters. Presented at the Second Workshop on Sponsored Search Auctions, Ann Arbor, Michigan, June 11, 2006. Available at <www.business.ualberta.ca/kasdemir/ssa2/schedule.htm>. Last accessed August 2, 2009.
- Rothkopf, M. H. Thirteen reasons why the Vickrey–Clarke–Groves process is not practical. *Operations Research*, 55, 2, 2007, 191–197.
- Sherman, L., and Deighton, J. Banner advertising: measuring effectiveness and optimizing placement. *Journal of Interactive Marketing*, 15, 22, 2001, 60–64.
- Varian, H. Position auctions. *International Journal of Industrial Organization*, 25, 2007, 1163–1178.
- Vorobeychik, Y., and Reeves, D. M. Equilibrium analysis of dynamic bidding in sponsored search auctions. *International Journal of Electronic Business*, 6, 2, 2008, 172–193.
- Xie, Y., O'Hallaron, D. R. Locality in search engine queries and its implications for caching. In Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002), vol. 3, New York, NY, June 23–27, 2002, 1238–1247.