

On the interest of introducing randomness in ad-word auctions

Patrick Maillé¹ and Bruno Tuffin²

¹ Institut Telecom; Telecom Bretagne
2 rue de la Châtaigneraie CS 17607
35576 Cesson Sévigné Cedex, France
patrick.maill@telecom-bretagne.eu

² INRIA Rennes - Bretagne Atlantique
Campus universitaire de Beaulieu
35042 Rennes Cedex, France
btuffin@irisa.fr

Abstract. Search engines play and will still play a major role in the use of networks. Sponsored search auctions is the basic tool for a return on investment in this industry, accounting for an increasing part of the business. We introduce here a model for consumer behavior in the context of ad-word auctions. Considering that unsatisfying answers of the ad-word engine will lead some consumers to perform again the same request later on, we show that displaying only the highest bidding or highest revenue-producing advertisers in a deterministic way is not always the best strategy for the ad-word engine. Instead, some randomization among advertisers can provide higher revenues. We also design a Vickrey-Clarke-Groves auction rule for a display probability and compare it with the current generalized-second-price scheme.

Keywords: Auctions, Random processes

1 Introduction

Search engines play a crucial role in the Internet, allowing, just by giving keywords, to reach the most relevant web pages. This role is expected to be at least as important in next generation networks. Search engines make money by proposing advertising slots to potential advertisers, usually displayed at the top and/or at the right of the page of results corresponding to the keyword(s). Actually search engine advertising has become an important business, the combined revenue of the two main actors in the area, Yahoo! and Google, being more than \$11 billion in 2005 for instance [1], and this business is expected to count for about 40% of total advertising revenue [2].

Keyword ads work as follows: an Internet user seeking for a web site types keywords on a search engine. Based on those words, the search engine proposes a list of links ranked by relevance, but also some commercial links chosen according to criteria which depend on the engine. Our goal is to investigate the

allocation of commercial slots to advertisers. For a nice overview of the general and recent modeling issues, an interested reader can for instance look at [3]. We focus in this paper on a single keyword, the best strategy of advertisers in terms of keyword selection being not addressed here as in most of the literature. Slot allocation is based on auctions, and most of existing works deal with a game of complete information [1,4,5,6,7] where advertisers make a bid for the keyword corresponding to the maximum price they would accept to pay if the link is clicked through. Advertisers are ranked according to a prespecified criterion, for instance the bid value (initially for Yahoo!), or the revenue they will generate (more or less corresponding to Google's situation), taking into account the click-through rate (CTR) for each advertiser. The K available slots are then allocated according to the above ranking to the K first advertisers. Advertisers pay the search engine each time their link is clicked through, the amount they are then charged being, in the literature, either exactly the bid (corresponding to the so-called first price auction), the opportunity cost that the considered advertiser's presence introduces to all the other advertisers (the so-called Vickrey-Clarke-Groves (VCG) auction), or the generalized second price (GSP) where each advertiser pays the bid of the advertiser just below itself in the ranking (that latter scheme is applied by Google and Yahoo!, and yields the same incentives to bid truthfully than VCG auction if the game is not one-shot and advertisers anticipate the effects of the strategies of their competitors [8]). Most models differ on the assumptions about the CTR modeling, the budget limit of the advertisers, the game being static or played several times, or the pricing rule applied.

But all the above works, and to our knowledge all the literature, assume that an Internet user will make a search only once and will never come again. We propose here to investigate the consequences of users potentially composing several times the same keyword, the interval between two searches being random. This is typical of users looking for new or additional information, or who do not remember the previous results. In that situation, *always* presenting the same advertisers on the sponsored slots may be a bad option since if a user has not clicked through the links once, he is likely to proceed the same way again. A random strategy for the allocation, still dependent on the bid of the advertisers but allowing advertisers with small bids to be displayed with a positive even if small probability, may produce a higher expected revenue than the current deterministic strategy. This paper aims at illustrating the relevance of that approach. We therefore introduce here a model describing users searching the same keyword a random number of times, and compare according to the parameter values the revenues generated by the random or deterministic allocation of slots. Only $K = 1$ slot and two advertisers are considered to simplify the analysis and for comparison purposes. We show that *for fixed bids and prices*, applying a random allocation can increase revenue. We then build a VCG auction scheme based on this allocation rule, and compare the revenues and social welfare at equilibrium of the bidding game among advertisers with those obtained using a deterministic GSP rule.

Remark that randomness has already been considered in ad-word auctions. In [6,9], a Markovian model is used to represent the user behavior when looking at the ranked ads: she looks at the first ad then clicks and quits or goes to the second with a given probability, which she clicks or not, then goes to the following, up to the number K . In [10], randomness is introduced on the size of the population and the type of advertisers but the allocation rule is kept deterministic. The closest work to ours is maybe [11]. In that paper the authors discuss the ad-slot placement and pricing strategies maximizing the expected revenue of search engine over a given interval of time. They find the optimal allocation rule and show that the $c\mu$ -scheduling rule (i.e., maximizing selling probability times (virtual) valuation) is a good suboptimal policy, but the randomness is rather on the types of advertisers. To the best of our knowledge, no paper is actually considering randomness in the behavior of users, coming several times, and the consequence on slot allocations as we do.

The paper is organized as follows. Section 2 introduces the model and describes users' behavior. Section 3 computes the average revenue generated using a random allocation strategy with fixed price per click for users and shows that this rule can increase revenue. Then Section 4 suggest to use Vickrey-Clarke-Groves auctions to allocate the ad display probabilities among the advertisers, when implementing a randomized instead of deterministic ad display policy. Section 5 illustrates the fact that such an auction scheme can perform better than a deterministic GSP display mechanism in terms of advertiser revenue and overall wealth generated by the ad slot. Our conclusions and some directions for future work are given in Section 6.

2 Model

In this section, we introduce a mathematical model aimed at representing the behavior of a customer when faced with a search engine implementing an ad-word auction. The particularity of that model is to explicitly express the retry rate of customers, when previous answers have not resulted into a sale.

Since the paper is mainly focused on illustrating a phenomenon (i.e., the fact that introducing randomness in the selection of the advertisers to display may be profitable to the search engine), we consider an extreme case to simplify the analysis, both for the auction engine and the customer population. The model can be complicated in order to better fit reality.

2.1 Search engine basic model

We consider a search engine providing only one commercial slot, and two advertisers, say, 1 and 2, competing for that slot on a given ad-word. In an auction, players (here, advertisers) submit bids, from which allocation and pricing rules are applied. Each advertiser $i = 1, 2$, is therefore characterized in the auction by

- b_i the bid of advertiser i for that keyword,

- π_i the probability that advertiser i 's ad is displayed (which should depend on the bid profile (b_1, b_2)),
- p_i the price-per-click that advertiser i is charged, also dependent on (b_1, b_2) .

2.2 User behavior model

The population of users likely to search for that ad-word is assumed to be heterogeneous, made of two types: some users, called type- A users, can only be interested in purchasing the good sold by advertiser 1, but can potentially click on the ad of advertiser 2 without purchasing it eventually. On the other hand, type- B users are potential buyers of advertiser 2 only (likewise, they will never purchase advertiser 1's product even if they click on its ad when displayed).

Formally, customer (new) requests occur randomly over time: we denote by λ_A (resp. λ_B) the average number of *first* requests per time unit of type- A (resp. type- B) customers.

For a type- A user who performs a search with the considered keyword:

- if the search engine displays the ad of advertiser 2, then the probability of the user clicking the ad is denoted by $c_{A,2}$, but recall that even if the user clicks on the ad, he does not purchase the good;
- if the search engine displays the ad of advertiser 1, we denote by $c_{A,1} > 0$ the probability of the user clicking on that ad, and by $h_{A,1} > 0$ the probability that the user purchases the good after clicking on the ad.

As a result, for one search, the probability that a type- A user buys its wanted product is $\pi_1 c_{A,1} h_{A,1}$: indeed, that probability is $c_{A,1} h_{A,1}$ if advertiser 1 is displayed in the ad-word slot, and 0 otherwise.

Inevitably, some users do not end up buying their wanted good, because of the bad ad displayed or for other reasons (e.g., lack of time to finish the purchase, or hesitation). In this paper, we assume that some of those users will try again later and perform a search with the same keyword, after some time, independently of the number of previous search attempts. For those retries, we assume that the user has the same behavior, i.e., the same probabilities of clicking on the ad, of purchasing the good, and of trying again later in the case of no purchase.

As a result, the average number of retries per time unit only depends on the total average number of searches λ_A^{tot} (including new attempts and retries from previous attempts): Let us denote by R_A the probability of a not-buying type- A user to retry the keyword later. Define by θ_A^{retry} the overall probability of retry, given by

$$\theta_A^{\text{retry}} := R_A(1 - \pi_1 c_{A,1} h_{A,1}).$$

The overall behavior of type- A users is illustrated in Figure 1, representing new searches, retries and users definitely leaving the search.

In steady-state -which can be easily shown to exist, the system behaving like an infinite server queue with instantaneous service time-, the mean number of retries per time unit should equal the total mean number of search tries λ_A^{tot} multiplied by the overall probability of retry θ_A^{retry} .




Fig. 1. Type- A users search behavior. Arcs are labeled with the average number of occurrences per time unit.

We clearly see in Figure 1 that the total mean number of requests by type- A users per time unit verifies

$$\lambda_A^{\text{tot}} = \lambda_A + \theta_A^{\text{retry}} \lambda_A^{\text{tot}},$$

which gives, since $\theta_A^{\text{retry}} < 1$,

$$\lambda_A^{\text{tot}} = \frac{\lambda_A}{1 - \theta_A^{\text{retry}}} = \frac{\lambda_A}{1 - R_A(1 - \pi_1 c_{A,1} h_{A,1})}. \quad (1)$$

For type- B users, we define the probabilities $c_{B,1}, c_{B,2}, h_{B,2} > 0$, and R_B in the same way as for type- A users. With the same reasoning, the mean number of requests by type- B users per time unit is $\lambda_B^{\text{tot}} = \frac{\lambda_B}{1 - R_B(1 - \pi_2 c_{B,2} h_{B,2})}$.

3 Search engine revenue

Let us consider fixed values of the advertiser bids b_1 and b_2 , and assume that the corresponding display probabilities are π_1 and $\pi_2 = 1 - \pi_1$, with prices per click p_1 and p_2 , respectively.

The mean revenue (utility) per time unit for the search engine (for the considered keyword auction) is

$$\begin{aligned} U &= p_1 \underbrace{\pi_1 (\lambda_A^{\text{tot}} c_{A,1} + \lambda_B^{\text{tot}} c_{B,1})}_{\text{nb of clicks on ad 1}} + p_2 \underbrace{\pi_2 (\lambda_B^{\text{tot}} c_{B,2} + \lambda_A^{\text{tot}} c_{A,2})}_{\text{nb of clicks on ad 2}} \\ &= \lambda_A \frac{\pi_1 (p_1 c_{A,1} - p_2 c_{A,2}) + p_2 c_{A,2}}{1 - R_A(1 - \pi_1 c_{A,1} h_{A,1})} + \lambda_B \frac{\pi_1 (p_1 c_{B,1} - p_2 c_{B,2}) + p_2 c_{B,2}}{1 - R_B(1 - c_{B,2} h_{B,2} + \pi_1 c_{B,2} h_{B,2})} \quad (2) \end{aligned}$$

where we used the fact that $\pi_1 + \pi_2 = 1$.

To simplify the analysis, we will first assume that users types are symmetric, and that users never click on an ad that does not correspond to their type.

Assumption A *Users of both types have a similar behavior with respect to their preferred/less preferred advertisers, in terms of click probability, conversion rate, and retry probability. Formally,*

$$\begin{aligned}c_{A,1} &= c_{B,2} := c \\h_{A,1} &= h_{A,2} := h \\R_A &= R_B := R.\end{aligned}$$

Moreover, users are only likely to click on the ad of their preferred advertiser, i.e., $c_{A,2} = c_{B,1} = 0$.

Under that simplifying assumption, the auctioneer revenue is

$$U = \lambda_{AP_1} c \frac{\pi_1}{1 - R(1 - \pi_1 ch)} + \lambda_{BP_2} c \frac{1 - \pi_1}{1 - R(1 - ch + \pi_1 ch)}. \quad (3)$$

For a real ad-word auction mechanism implementing randomization, the price per click p_i charged to each advertiser $i = 1, 2$ should increase with the probability π_i of displaying each ad. However, if for a moment we fix those prices as given and look at (3) as a function of π_1 , we remark that the revenue-maximizing strategy does not necessarily consist in always displaying the ad with the highest price, or the ad with the highest product $\lambda_{X_i} \times p_i$, with X_i the type of users potentially interested in buying advertiser i 's product. Indeed, if the retry rate R is sufficiently close to 1, the revenue in (3) attains its maximum for a π_1^* in the interior of $(0, 1)$. The rationale is that potential consumers of the not-displayed advertiser will perform again the search, giving a large overall request rate for that type of consumers, which makes it beneficial for the auctioneer to display the ad.

That result is formalized by the following proposition.

Proposition 1. *The value π_1^* of π_1 maximizing the search engine revenue is*

- $\pi_1^* = 0$ if $\sqrt{\frac{\lambda_{AP_1}}{\lambda_{BP_2}}} \leq \frac{1-R}{1-R(1-ch)}$,
- $\pi_1^* = 1$ if $\sqrt{\frac{\lambda_{BP_2}}{\lambda_{AP_1}}} \leq \frac{1-R}{1-R(1-ch)}$,
- $\pi_1^* = \frac{1}{R \cdot ch} \frac{1-R+R \cdot ch - (1-R) \sqrt{\frac{\lambda_{BP_2}}{\lambda_{AP_1}}}}{\sqrt{\frac{\lambda_{BP_2}}{\lambda_{AP_1}}} + 1}$ if $\frac{1-R}{1-R(1-ch)} < \sqrt{\frac{\lambda_{AP_1}}{\lambda_{BP_2}}} < \frac{1-R(1-ch)}{1-R}$.

In other words, when R is sufficiently close to 1, we are in the third case, and the search engine gets a higher revenue than when always showing one of the advertisers.

Proof. It can be easily seen that for $R < 1$ the revenue (3) is a continuous, derivable and strictly concave function of π_1 on $[0, 1]$, with derivative

$$U'(\pi_1) = (1 - R)c \left(\frac{\lambda_{AP_1}}{[1 - R(1 - \pi_1 ch)]^2} - \frac{\lambda_{BP_2}}{[1 - R(1 - (1 - \pi_1)ch)]^2} \right).$$

Due to the strict concavity of U , we have three possibilities:

- either $U'(0) \leq 0$, i.e., $\sqrt{\frac{\lambda_A p_1}{\lambda_B p_2}} \leq \frac{1-R}{1-R(1-ch)}$, then $\pi_1^* = 0$;
- or $U'(1) \geq 0$, i.e., $\sqrt{\frac{\lambda_B p_2}{\lambda_A p_1}} \leq \frac{1-R}{1-R(1-ch)}$, then $\pi_1^* = 1$;
- or $\frac{1-R}{1-R(1-ch)} < \sqrt{\frac{\lambda_A p_1}{\lambda_B p_2}} < \frac{1-R(1-ch)}{1-R}$, which is always the case when R is sufficiently close to 1, then $\pi_1^* \in (0, 1)$, i.e., it is in the interest of the auctioneer to randomize when choosing which ad to display.

Example 1. Consider a simple and arbitrarily chosen example illustrating the kind of gain that can be obtained through randomization. Use the probability values $c = 1/2$, $h = 1/2$, the rates $\lambda_A = 1$, $\lambda_B = 0.8$ and prices per click $p_1 = 1$, $p_2 = 0.8$. The retry probability is taken as $R = 0.8$ so that we are in the third situation presented in the proposition. The revenue (3) is maximized at $\pi_1^* = 2/3$ and given by 1.4. If we compare it with the optimal revenue when only one ad is displayed, $\max(\lambda_A p_1, \lambda_B p_2) = 1$, a gain of 40% is observed. This gain is even increased for a larger value of R .

A similar result applies in the general case, as illustrated by the following proposition.

Proposition 2. *In the non-symmetric case, there also exists a unique π_1^* maximizing the revenue $U(\pi_1)$ of the search engine given by (2). The solution is in the interior of the interval $[0, 1]$ if $U'(0) > 0$ and $U'(1) < 0$.*

Proof. The proof is similar to the one of Proposition 1, given that the revenue defined by (2) is still a strictly concave function of π_1 with $R_A, R_B < 1$.

Remark 1. Notice that the randomization proposed here transfers some of the complexity from the advertiser to the auctioneer, since in most current auctions the advertisers develop quite complex bidding strategies to adapt over time to the opponents' bids and to the flow of requests, so that it appears effectively that not always the same ads are displayed when users perform the same request several times. With the randomization suggested here, we might imagine that advertisers fix their bid for a time period, and the auctioneer computes the ad displaying probabilities based on the bids and some knowledge of the demand (which it is more likely to be able to estimate than the advertiser is).

4 A VCG auction mechanism for display probability

In this section, we highlight a utility function for an advertiser i , that depends on the display probability π_i of its ad. Such a utility function being continuous and increasing in π_i , we propose to define an auction scheme among providers to allocate and price the “probability range”, i.e., the interval $[0, 1]$, as an infinitely divisible resource. We then study a game where advertisers submit their bid representing how much they are accepting to pay per click, and the display probability is allocated and charged according to VCG auction rules.

4.1 Advertiser willingness-to-pay in terms of the ad display probability

We first quantify the average revenue in terms of display probability for each advertiser. With the same notations as in the previous sections and under Assumption A (symmetry among both types of customers, customers only interested in clicking for one advertiser), the mean sales income per time unit due to the keyword searches for, say, provider 1, is

$$\begin{aligned} V_1(\pi_1) &= \lambda_A^{\text{tot}} \pi_1 ch v_1 \\ &= \frac{\lambda_A \pi_1 ch}{1 - R(1 - \pi_1 ch)} v_1, \end{aligned} \quad (4)$$

where v_1 is the benefit that the advertiser makes on each sale (i.e., the selling cost of the product minus its production cost). We call that function V_1 the *valuation* function of provider 1, since it represent the monetary benefit due to the ad display probability. Remark that exactly the same form would be obtained when computing the valuation $V_2(\pi_2)$ of provider 2.

As a result, the overall utility (revenue) for an advertiser i who has to pay a price p_i to have its ad displayed with probability π_i is *quasi-linear*, that is,

$$U_i = V_i(\pi_i) - p_i. \quad (5)$$

Interestingly, we remark that for each $i = 1, 2$, V_i is a continuously derivable, nondecreasing and concave function on $[0, 1]$, and that $V_i(0) = 0$. Such properties are often needed when designing some pricing and allocation mechanisms with nice properties -both from economic and computational perspectives-.

In the following subsection, we therefore propose an auction scheme to efficiently allocate the overall resource (display probabilities) among advertisers.

4.2 VCG auctions for allocating an infinitely divisible good

In this subsection, we use the possibility of randomization to define an auction on ad display probabilities. The auction is interpreted as an auction for an infinitely divisible good, with total quantity 1. Indeed, from the formulation of the previous subsection, advertisers are only sensitive to their display probability and the price they pay, while the total sum of advertiser (probability) allocations cannot exceed 1.

The problem of the auctioneer is therefore to share that total amount 1 of resource among bidders. In this work, we look at the well-known Vickrey-Clarke-Groves (VCG) [12,13,14] mechanism, which is the only auction mechanism with the properties below for bidders with quasi-linear utility functions [15]:

- *Incentive compatibility*: truthfully declaring one's value for the good is a dominant strategy, i.e., there is no gain to expect from lying about one's willingness-to-pay.
- *Efficiency*: when bidders are truthful, allocations maximize social welfare, that is the sum of all bidders willingness-to-pay.

- *Individual rationality*: every bidder pays less than his declared willingness-to-pay, so there is always an interest in participating in the auction.
- *Non-negativeness of prices*.

The VCG mechanism works as follows: Each bidder i is asked to declare its entire valuation (willingness-to-pay) function $V_i(\cdot)$. Then,

- **(allocation rule)** the auctioneer computes an allocation vector $(\pi_i)_{i \in \mathcal{I}}$ that maximizes social welfare, i.e., the sum of bidders declared valuations;
- **(pricing rule)** the price that each bidder has to pay equals the loss of social welfare that it imposes the others through its presence.

4.3 Applying VCG

Based on (4), declaring one's willingness-to-pay function consists in revealing the four parameters of V_i , namely for provider 1: λ_A , R , v_1 , and the product ch . However, as pointed out in Remark 1, some parameters like the retry rate R are more likely to be measured by the search engine itself: We can therefore consider that provider 1 only has to declare the parameters λ_A , v_1 , and ch . Since being truthful is a dominant strategy for advertisers, we consider that they reveal their real valuation parameters.

Allocation rule. The auctioneer computes the display probabilities π_i solving the strictly convex optimization problem

$$\max_{\pi_1, \pi_2 \text{ s.t. } \pi_1 + \pi_2 \leq 1} \bar{V}_1(\pi_1) + \bar{V}_2(\pi_2), \quad (6)$$

where \bar{V}_i is the declared willingness-to-pay function of advertiser i ($\bar{V}_i = V_i$ if i bids truthfully). Remark that since the functions \bar{V}_i are strictly increasing, the solution of (6) lies in the set where $\pi_1 + \pi_2 = 1$. Therefore, the declared-welfare maximizing allocation $(\bar{\pi}_1, \bar{\pi}_2)$ is such that

$$\begin{aligned} \bar{\pi}_1 &= \arg \max_{\pi_1 \in [0,1]} \frac{\pi_1 \lambda_A ch}{1 - R(1 - \pi_1 ch)} v_1 + \frac{(1 - \pi_1) \lambda_B ch}{1 - R(1 - (1 - \pi_1) ch)} v_2 \\ \bar{\pi}_2 &= 1 - \bar{\pi}_1. \end{aligned} \quad (7)$$

The optimization problem (7) is exactly the same as in (3), with v_i playing the role of p_i for each advertiser i . Therefore, from Proposition 1, $\bar{\pi}_1$ is given by

- $\bar{\pi}_1 = 0$ if $\sqrt{\frac{\lambda_A v_1}{\lambda_B v_2}} \leq \frac{1-R}{1-R(1-ch)}$,
- $\bar{\pi}_1 = 1$ if $\sqrt{\frac{\lambda_B v_2}{\lambda_A v_1}} \leq \frac{1-R}{1-R(1-ch)}$,
- $\bar{\pi}_1 = \frac{1}{R \cdot ch} \frac{1-R+R \cdot ch - (1-R) \sqrt{\frac{\lambda_B v_2}{\lambda_A v_1}}}{\sqrt{\frac{\lambda_B v_2}{\lambda_A v_1}} + 1}$ if $\frac{1-R}{1-R(1-ch)} < \sqrt{\frac{\lambda_A v_1}{\lambda_B v_2}} < \frac{1-R(1-ch)}{1-R}$.

Pricing rule. Finally, to determine the charge to each advertiser, the auctioneer computes the loss of value (in terms of declared willingness-to-pay) that each advertiser imposes on the other through its presence. Since we consider only two advertisers, then, had one of them been absent from the auction, the whole resource (i.e., display probability 1) would have been given to the other. As a result, the total price t_i per time unit that each advertiser i is charged under the VCG rule is given by $\begin{cases} t_1 = \bar{V}_2(1) - \bar{V}_2(\bar{\pi}_2) \\ t_2 = \bar{V}_1(1) - \bar{V}_1(\bar{\pi}_1). \end{cases}$

Such an average charge per time unit can easily be converted to a price per click, to fit current auctions charging base. The mean number of clicks per time unit for advertiser 1 is $\lambda_A^{\text{tot}}\bar{\pi}_1c$, where λ_A^{tot} is given in (1); as a result, under Assumption A the auctioneer may charge advertiser 1 a price per click

$$p_1 = (\bar{V}_2(1) - \bar{V}_2(\bar{\pi}_2)) \frac{1 - R(1 - \bar{\pi}_1ch)}{\lambda_A\bar{\pi}_1c}.$$

Likewise, the price per click for advertiser 2, whose ad is displayed with probability $\bar{\pi}_2$, should be

$$p_2 = (\bar{V}_1(1) - \bar{V}_1(\bar{\pi}_1)) \frac{1 - R(1 - \bar{\pi}_2ch)}{\lambda_A\bar{\pi}_2c}.$$

4.4 Asymmetric case

The asymmetric case can easily be handled equivalently. The valuation functions $V_1(\pi_1)$ and $V_2(\pi_2)$ are then respectively the first and second terms of the revenue (2), with again p_1 and p_2 replaced by v_1 and v_2 . The rest of the discussion follows as well, invoking Proposition 2 instead of Proposition 1.

5 VCG for randomizing displayed ads versus deterministic GSP

In this section, we consider a simple illustrative example comparing the consequences of a VCG auction on display probability to those obtained with a deterministic GSP pricing scheme. We consider a symmetric case, with $v_1 = 1$, $v_2 = 0.8$, $\lambda_A = \lambda_B = 1$, $c = h = 0.5$. Since the main phenomenon we want to model in this paper is the fact that consumers try out the same research several times, we make the retry probability R vary.

The currently used GSP auction consists in asking each advertiser the price it is willing to pay to be displayed (with probability 1). In the numerical computations we make here, we assume that advertisers bid truthfully, i.e. they declare their real value of a customer click if their ad were to be displayed with probability 1. Under Assumption A, $b_1 = v_1h = 0.5$ and $b_2 = v_2h = 0.4$. As a result, GSP will always result in advertiser 1's ad being displayed. At each click on its ad, advertiser 1 will be charged $b_2 = 0.4$. The corresponding revenue per time unit for the auctioneer is thus $\lambda_A^{\text{tot}}cb_2 = \frac{0.8}{4-3R}$, while the per-unit net benefit

of provider 1 is $\lambda_A^{\text{tot}} c(v_1 h - b_2) = \frac{0.2}{4-3R}$. Provider 2 gets no benefit from the auction, since its ad is never displayed and it does not pay any charge to the search engine.

In Figure 2, we compare the advertiser as well as auctioneer net benefits under VCG and GSP auction schemes. Social welfare, i.e., the sum of those utilities (or in other words, the total wealth generated by the ad slot), is shown in Figure 3.

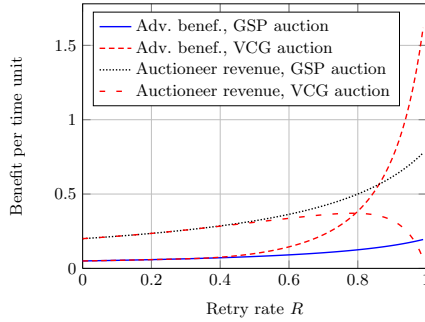


Fig. 2. Benefits of participants.




Fig. 3. Social welfare

This example illustrates that when a VCG auction is applied on the display probability, the revenue of advertisers is larger than when applying a deterministic GSP pricing, and the difference increases with R , while it is the opposite for the search engine revenue. In the case of a search engine in the situation of a monopoly, a deterministic GSP would be preferable than a random VCG in terms of revenue, but in the case of competitive engines with comparable quality (for instance Yahoo! against Google), advertisers would prefer the one applying VCG due to their larger revenue, and the advantage may change. Also, as seen on Figure 3, social welfare is larger with our new VCG scheme, which points out a better use of the resource (here, the ad slot).

6 Conclusion

This paper proposes a new randomized allocation rule for ad-word auctions, based on the idea that users often retry the same search. We have shown that *for fixed bids and prices*, applying a random allocation can increase revenue. We have also designed a VCG auction scheme based on that allocation rule, and compared the revenues and social welfare at equilibrium of the bidding game among advertisers with those obtained using a deterministic GSP rule. The search engine revenue is lower with our scheme, but not social welfare nor advertisers' revenue. Due to that last point, we claim that in a competition

between engines, our new scheme will be preferred, and therefore could also generate more revenue.

Next works will go into several directions. First, we would like to study theoretically the competition game between two engines, to support our claim. Studying the case of multiple slots, as well as the game between advertisers under the randomized allocation rule but with a GSP pricing rule and a reserve price are also important issues we intend to address.

References

1. Varian, H.: Position auctions. *International Journal of Industrial Organization* **25** (2005) 1163–1178
2. Liu, D., Chen, J., Whinston, A.: Competing keyword auctions. In: *Proc. of 4th Workshop on Ad Auctions*, Chicago, IL, USA (2008)
3. Jordan, P., Wellman, M.: Designing an ad auctions game for the trading agent competition. In: *IJCAI-09 Workshop on Trading Agent Design and Analysis (TADA)*, Pasadena, California, USA (2009)
4. Edelman, B., Ostrovsky, M., 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
5. Lahaie, S., Pennock, D.: Revenue analysis of a family of ranking rules for keyword auctions. In: *Proc. of the EC'07*, San Diego, California, USA (2007)
6. Aggarwal, G., Feldman, J., Muthukrishnan, S., Pál, M.: Sponsored search auctions with markovian users. In: *Proc. of 4th Workshop on Ad Auctions*, Chicago, IL, USA (2008)
7. Chakrabarty, D., Zhou, Y., Lukose, R.: Budget constrained bidding in keyword auctions and online knapsack problems. In: *Proc. of the 3rd Workshop on Ad Auctions*, Banff, Canada (2007)
8. Bu, T., Deng, X., Qi, Q.: Dynamics of strategic manipulation in ad-words auction. In: *Proc. of the 3rd Workshop on Ad Auctions*, Banff, Canada (2007)
9. Kempe, D., Mahdian, M.: A cascade model for externalities in sponsored search. In: *Proc. of 4th Workshop on Ad Auctions*, Chicago, IL, USA (2008)
10. Dellarocas, C., Viswanathan, S.: The holy grail of advertising? Allocative efficiency and revenue implications of “pay-per-action” advertising in environments with quality uncertainty. In: *Proc. of 4th Workshop on Ad Auctions*, Chicago, IL, USA (2008)
11. Menache, I., Ozdaglar, A., Srikant, R., Acemoglu, D.: Dynamic online-advertising auctions as stochastic scheduling. In: *Proc. of NetEcon.* (2009)
12. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* **16**(1) (1961) 8–37
13. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* **11** (1971) 17–33
14. Groves, T.: Incentives in teams. *Econometrica* **41**(3) (1973) 617–631
15. Green, J., Laffont, J.J.: Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* **45**(2) (1977) 427–438