

Incorporating User Utility Into Sponsored-Search Auctions

(Short Paper)

Yagil Engel
University of Michigan
2260 Hayward St
Ann Arbor, MI 48109-2121, USA
yagil@umich.edu

David Maxwell Chickering
Microsoft Live Labs
1 Microsoft Way
Redmond, WA 98052, USA
dmax@microsoft.com

ABSTRACT

We study principled methods for incorporating user utility into the selection of sponsored search ads. We describe variations of the GSP allocation/pricing mechanism that accommodate these user utility functions, we provide interesting and useful parallels of some of the theoretical properties from the traditional GSP mechanisms in the new GSP variations, and we present simulation results that exemplify the use of the ranking system.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Economics

Keywords

Sponsored Search, Auctions, Preference Handling

1. INTRODUCTION

The most popular search engines currently rank and price ads using a generalized second price (GSP) auction, which works as follows. Advertisers place per-click bids on search keywords that are relevant to their ads. Each time a search user types in a query, all the advertisers that have matching bids are sorted using a ranking function, and the top advertisers have their ads shown along with the “organic” results returned by the search engine. An advertiser is charged only if the search user clicks on his ad, and the amount he is charged is the minimum bid needed to retain his position in the sort.

In most analyses of the GSP auction, the ranking function is assumed to be either (1) the unaltered per-click bids or (2) the “per impression” bids obtained by multiplying each per-click bid by the click-through rate (CTR) of the corresponding ad. Method (2) is sometimes referred to as ranking “by revenue”. We refer to GSP auctions using these ranking functions as GSP_B and GSP_R , respectively.

Cite as: Incorporating User Utility Into Sponsored-Search Auctions (Short Paper), Yagil Engel and David Maxwell Chickering, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16, 2008, Estoril, Portugal, pp. 1565-1568.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

GSP_B and GSP_R may be appropriate for extracting short-term revenue from the advertisers, but by not representing the value of the ads to the user of the search engine, they are missing a potentially important component needed for long-term revenue: search traffic volume. Several existing search-engine companies claim to be using a ranking function that includes a quality measure for keeping out low-quality ads, but the specifics of what these measures are and how they are used is not disclosed. In this paper, we study principled methods for handling ad quality by defining a ranking function that makes explicit trade-offs between user utility and short-term revenue. In Section 2, we define a user-utility model. In Section 3, we define a publisher-utility model that combines user utility with short-term revenue. In Section 4, we describe mechanisms based on GSP that use the publisher-utility model. Finally, in Section 5, we provide simulation results.

2. MODELING USER EXPERIENCE

In this section, we describe models for search-user utility as a function of the interactions that the user has with the ads that are shown. Due to space constraints, we limit our discussion to models that are *additive* over these interactions; that is, the utility for the interactions on a set of ads Ω is simply the sum of the utility for the interaction on each individual ad $\sigma \in \Omega$.

For a given ad σ , we assume there are four mutually exclusive and collectively exhaustive interactions $I(\sigma)$ that the user can have: i_s is the event that the user looks at (“scans”) σ but does not click, i_{cg} is the event that the user clicks σ , and finds the resulting page relevant (“clicks good”), i_{cb} is the event that the user clicks σ and does not find the resulting page relevant (“clicks bad”), and i_{ig} is the event that the user does not look at σ (“ignored”). For simplicity, we assume that a user must look at an ad before clicking, and that the user has exactly one interaction with each ad.

We use $U(i)$ to denote the user utility function over interaction events. We assume that the maximum value \bar{u} for $U(i)$ occurs when $i = i_{cg}$ and that the minimum value \underline{u} occurs when $i = i_{cb}$. Without loss of generality, we assume $\bar{u} - \underline{u} = 1$. Furthermore, we assume that $U(i_{ig}) = 0$; as a result of these assumptions, the function has two free parameters ($U(i_s)$ and either $U(i_{cb})$ or $U(i_{cg})$).

In Section 5, we show experiments using three specific utility functions. In the first, which we name *UER* for “utility equals relevance”, we set $U(i_s) = U(i_{cb}) = 0$; the result is that the expected utility of an advertisement is simply the product of its click-through rate and the probability of rele-

vance. In our second utility function we use $U(i_s) = -0.05$ and $U(i_{cb}) = -0.4$ (and consequently $U(i_{cg}) = 0.6$); we call this model *LC* for “low cost” due to the fact that the cost of clicking an irrelevant ad is smaller than the gain from clicking a relevant ad. Our third utility function has $U(i_s) = -0.05$ and $U(i_{cb}) = -0.5$; we call this model *EQ* for “equal” because the cost of clicking an irrelevant ad is equal to the gain from clicking a relevant ad.

3. PUBLISHER UTILITY MODEL

Search users currently do not participate in any monetary transfer with the search engine, and therefore it is up to the publisher to incorporate user utility into the search mechanism through the choice of which ads are shown. In this section, we show how to construct a publisher ranking function that makes an appropriate tradeoff between user utility and revenue.

3.1 Publisher Value of User Utility

Let $v(u, r)$ denote the publisher’s value, in dollars, of providing utility u to the user and receiving short-term revenue r . We assume a multi-linear form for $v(u, r)$ that accommodates the following indifference constraint:

$$v(\underline{u}, r) = v(\bar{u}, (1 - t_m)r - t_a). \quad (1)$$

That is, the publisher is willing to spare a fraction t_m of the revenue and an additional t_a dollars from the remainder in order to increase user utility from the minimum \underline{u} to the maximum \bar{u} . By specifying a non-zero t_m , the publisher indicates that he values the user experience in high stake searches more than the experience in searching for a cheap item. On the other hand, if the publisher is always willing to spare a constant small amount t_a in order to improve a user’s utility, then for low-revenue auctions he can improve user utility dramatically with low cost.

A simple multi-linear form that satisfies (1), given that $\bar{u} - \underline{u} = 1$, is

$$v(u, r) = t_a u + t_m(u - \underline{u})r + (1 - t_m)r. \quad (2)$$

This form is consistent with assuming that the publisher’s preferences exhibit mutual independence of u and r [5]. In our experiments, we consider publisher utility functions of the form of (2).

3.2 Probabilistic Model

A utility maximizing publisher will rank sets of ads by expected (publisher) utility given a particular query. In this section, we define the probability distribution (over the interaction events) that allows us to compute this expectation.

We first assume that the interaction events for a set of ads are mutually independent, and thus the joint probability for all events decomposes into a product of probabilities $P_{\sigma_k, k}(\cdot)$ that are specific to each ad σ_k and corresponding position k ¹. Our independence assumption is violated if the probability of the user clicking on one ad depends on the quality and/or quantity of other ads he clicked. More accurate models can be estimated from data; for the purpose of this work, however, we keep the independence assumption for simplicity.

¹Note that $P_{\sigma_k, k}(\cdot)$ will certainly depend on the given query q ; we leave this dependence implicit to simplify notation.

User-interaction events decompose into up to three sub-events: the user looks at the ad or not, the user clicks on the ad or not, and a clicked ad is relevant or not. As common in the sponsored-search-auction literature [3, 7, 2, 6], we assume that the click probability α_σ^j of an ad σ in position j is the product of an ad effect π_σ , and a positional effect β^j , so that $\alpha_\sigma^j = \pi_\sigma \beta^j$. Both the positional effect β^j and the ad-specific factor π_σ are computed based on data². In our model we interpret the positional effect as the probability that the user looks at the ad. To complete our probabilistic model, we need the probability ρ_σ that the landing page of an ad is relevant to the query given that the ad was clicked. In our experiments, we used the following algorithm to compute ρ_σ : we computed a similarity score between the text of the ad-destination page and the text of the organic search-result summaries shown in the results page, and converted this score to a probability using a calibration model constructed from human-labeled data.

Putting these pieces together, we have $P_{\sigma, j}(i_s) = \beta^j \cdot (1 - \pi_\sigma)$ (the user looked at the ad but did not click), $P_{\sigma, j}(i_{cg}) = \beta^j \cdot \pi_\sigma \cdot \rho_\sigma$ (the user looked at the ad, clicked, and the landing page was relevant), $P_{\sigma, j}(i_{cb}) = \beta^j \cdot \pi_\sigma \cdot (1 - \rho_\sigma)$ (the user looked, clicked, and the landing page was not relevant), and $P_{\sigma, j}(i_{ig}) = 1 - \beta^j$ (the user did not look at the ad).

3.3 Publisher Ranking Function

We now compute the expected publisher utility for a set of ads by combining the models from the previous two subsections. In place of the short-term revenue r that results from a click, we use the advertiser bid b . Because under GSP the price is the lowest bid needed to retain the given position, it follows that this modification of the publisher utility function will rank ads in the same order as if we had used the true price. Furthermore, we can use a reserve price to guarantee that prices are such that we only show positive-utility ads, and it follows that using the bid b is equivalent to using the price in terms of deciding what ads to show. The advantage we gain is that the expected value for a set of ads decomposes into the sum of per-position expectations.

Given an advertisement σ with bid b_σ in position j , we compute the publisher’s expected utility $v^e(\sigma, b_\sigma, j)$ by summing over the interaction events:

$$v^e(\sigma, b_\sigma, j) = \sum_i P_{\sigma, j}(i) v(U(i), r(i, b_\sigma))$$

where $r(i, b_\sigma) = b_\sigma$ if i is a click event (i_{cg} or i_{cb}) and zero otherwise. We now expand each of the utility terms in the sum above using the publisher-utility model from (2). When $i = i_{cg}$, the publisher utility is

$$v_{cg} = t_a U(i_{cg}) + [t_m(U(i_{cg}) - \underline{u}) + (1 - t_m)] \cdot b_\sigma,$$

and similarly when $i = i_{cb}$, the publisher utility is

$$v_{cb} = t_a U(i_{cb}) + [t_m(U(i_{cb}) - \underline{u}) + (1 - t_m)] \cdot b_\sigma.$$

When $i = i_s$, the publisher utility is $t_a U(i_s)$ (there is no short-term revenue), and when $i = i_{ig}$ the user utility and the revenue are both zero and thus the publisher utility is zero. Combining with our probability model, and using the

²We can estimate the positional effect by observing click rates of the same ad shown in different positions. In the model we used in our experiments, we used a shared positional model for every query.

definitions v_{cg} and v_{cb} above, we have:

$$v^e(\sigma, b_\sigma, j) = \beta^j(1 - \pi_\sigma)t_a U(i_s) + \beta^j \pi_\sigma(\rho_\sigma v_{cg} + (1 - \rho_\sigma)v_{cb}). \quad (3)$$

A utility-maximizing publisher will allocate ads so as to maximize the sum of (3) over all positions. Without loss of generality, assume that the positional effects β^j are monotonically decreasing with j (i.e., the click rate of a particular ad will be higher if it is placed higher in the list). We can achieve the utility-maximizing allocation by using a *position independent* ranking function $\mu(\sigma, b_\sigma)$ to greedily assign ads to positions, starting from $j = 1$ and continuing until all slots are filled or until the best unassigned ad has a negative value. In particular, we define

$$\mu(\sigma, b_\sigma) = \frac{v^e(\sigma, b_\sigma, j)}{\beta^j} = (1 - \pi_\sigma)t_a U(i_s) + \pi_\sigma(\rho_\sigma v_{cg}^e + (1 - \rho_\sigma)v_{cb}^e). \quad (4)$$

The equivalence of ranking by (3) and by (4) follows from the fact that the total expected utility $\sum_j v^e(\sigma, b_\sigma, j)$ for a set of ads is the dot product of the vector of $\mu(\cdot)$ values for each position and the corresponding vector of positional effects. In particular, because the positional effects are monotonically decreasing with j , the total expected utility is maximized by placing ads with highest value of $\mu(\cdot)$ in the positions with the highest positional effects.

4. SPONSORED SEARCH MECHANISMS

In the previous section, we derived a position-independent ranking function that allows publishers to incorporate user utility into the selection of ads. Importantly, the ranking function from (4) can be expressed as:

$$\mu(\sigma, b_\sigma) = f_\sigma + g_\sigma \cdot b_\sigma. \quad (5)$$

That is, by expanding the definitions of v_{cg}^e and v_{cb}^e and grouping terms as appropriate, the ranking function can be expressed as a linear function of the bid value. In this section, we study properties of the GSP mechanism when using any ranking function of the form of (5). Note that this class of ranking functions includes as special cases the two common ranking functions described in Section 1.

We now provide a formal definition of GSP. Let Φ denote the candidate ads for a particular query. We denote the corresponding set of bids by $B_\Phi = \{b_\sigma \mid \sigma \in \Phi\}$. Let $\mu(\sigma, b_\sigma)$ be an arbitrary position-independent publisher ranking function. We denote the sorted series resulting from applying $\mu(\cdot)$ on B_Φ by Ω_μ , and its j 'th element by $\Omega_\mu(j)$.

DEFINITION 1. *GSP $_\mu$ is an allocation and pricing scheme that allocates slot j to $\sigma = \Omega_\mu(j)$, for all $j \leq M$ for which $\mu(\sigma, b_\sigma) \geq m$. The price paid by σ is p_σ^j such that*

$$\mu(\sigma, p_\sigma^j) = \max\{\mu(\phi, b_\phi), m\}, \quad (6)$$

where $\phi = \Omega_\mu(j + 1)$ and m is the reserve price.

That is, GSP_μ ranks by $\mu(\cdot)$ and charges a bidder the minimum price required to remain in his position. We denote the set of slots allocated by $\mu(\cdot)$ by $\tilde{\Omega}_\mu$; note that $\tilde{\Omega}_\mu$ is Ω_μ truncated to the number of slots actually allocated. When μ represents expected publisher utility, a natural choice for m here is zero, ensuring that an ad is displayed only if it is rational to do so.

When $\mu(\cdot)$ is of the linear form of (5), we use GSP_L to denote the corresponding mechanism. Based on (6), the price assigned by GSP_L to the slot j won by σ is

$$p_\sigma = \frac{f_\psi - f_\sigma + g_\psi b_\psi}{g_\sigma}, \quad (7)$$

where $\psi = \Omega_\mu(j + 1)$.³

Properties of GSP have been analyzed in previous literature [6, 3, 7, 2]. We mainly follow the settings and definitions of Varian [7], which analyzes GSP_B but also considers the Google variation, in which $\mu(\sigma, b_\sigma) = q_\sigma b_\sigma$ for some unknown quality factor q_σ .⁴ GSP_L is different from this variation of GSP due to the additive element that is included in the ranking formula.

We denote the *valuation*—the amount by which a bidder truly values a click—for the advertiser of σ by h_σ . In all of the following discussion, we ignore the case in which the price is determined by m . Further, s and t each refer either to a position in $\tilde{\Omega}_\mu$ or to a losing position (i.e., a position $j > |\tilde{\Omega}_\mu|$ for which $\beta^j = 0$). We use Ω_B in place of Ω_μ to denote the result of ranking by GSP_B .

DEFINITION 2 ([7]). *Let $\sigma = \Omega_B(s)$. A set of bids is a Nash Equilibrium (NE) in GSP_B if it satisfies*

$$\alpha_\sigma^s(h_\sigma - p_\sigma) \geq \alpha_\sigma^t(h_\sigma - p_\psi) \forall t \neq s, \quad (8)$$

where for $t < s$, $\psi = \Omega_B(t)$, and for $t > s$, $\psi = \Omega_B(t + 1)$.

In words, the advertiser for each ad σ is making more total return in his current position than he would by adjusting his bid to move to another position. The asymmetric conditions under (8) reflect the fact that in order for σ to be moved to a *higher* ranked position t , the bid needs to exceed that of the current holder of position t , whereas to move to a *lower* ranked position t , the bid need only beat the bid of the current holder of $t + 1$.

We now adapt Definition 2 to the more general linear ranking function of (5). To emphasize the ranking function in our notation, we use $L(\sigma, b_\sigma)$ in place of $\mu(\sigma, b_\sigma)$ and we use Ω_L in place of Ω . The main change to the definition results from the fact that the per-click prices depend on the advertiser.

DEFINITION 3. *Let $\sigma = \Omega_L(s)$. A set of bids is a Nash Equilibrium (NE) in GSP_L if it satisfies*

$$\alpha_\sigma^s(h_\sigma - \frac{f_\phi - f_\sigma + g_\phi b_\phi}{g_\sigma}) \geq \alpha_\sigma^t(h_\sigma - \frac{f_\psi - f_\sigma + g_\psi b_\psi}{g_\sigma}) \forall t,$$

where $\phi = \Omega_L(s + 1)$ and for $t < s$, $\psi = \Omega_L(t + 1)$, and for $t > s$, $\psi = \Omega_L(t)$.

Varian [7] also introduces *symmetric Nash equilibrium* in which ψ in Definition 2 refers to $\Omega_B(t + 1)$ regardless of the relative values of s and t . This more restrictive equilibrium has the advantage that it leads to a tractable computation of equilibrium bids. Generalizing the definition to GSP_L yields:

DEFINITION 4. *Let $\sigma = \Omega_L(s)$. A set of bids is a Symmetric Nash Equilibrium (SNE) in GSP_L if it satisfies*

$$\alpha_\sigma^s(h_\sigma - \frac{f_\phi - f_\sigma + g_\phi b_\phi}{g_\sigma}) \geq \alpha_\sigma^t(h_\sigma - \frac{f_\psi - f_\sigma + g_\psi b_\psi}{g_\sigma}) \forall t,$$

³More accurately, the max over the term above and m .

⁴Recently, Yahoo! has started using an unknown quality factor as well.

where $\phi = \Omega_L(s + 1)$ and $\psi = \Omega_L(t + 1)$.

The following observation, which is similar to the one used by Varian in discussing the Google variation of GSP, shows that through the ranking function $L(\cdot)$ there is a simple one-to-one mapping between equilibria in GSP_L and GSP_B .

PROPOSITION 1. *Let $B_\Phi = \{b_\sigma \mid \sigma \in \Phi\}$ be a bids profile, and let $\hat{B}_\Phi = \{L(\sigma, b_\sigma) \mid \sigma \in \Phi\}$. Further, let $H_\Phi = \{h_\sigma \mid \sigma \in \Phi\}$ be the valuations of the advertisers, and $\hat{H}_\Phi = \{L(\sigma, h_\sigma) \mid \sigma \in \Phi\}$. Then: (1) B_Φ is NE in GSP_L with valuations H_Φ if and only if \hat{B}_Φ is NE in GSP_B with valuations \hat{H}_Φ . (2) B_Φ is SNE in GSP_L with valuations H_Φ if and only if \hat{B}_Φ is SNE in GSP_B with valuations \hat{H}_Φ .*

This mapping allows us to easily adapt results about equilibrium points in GSP_B to results about equilibrium points in GSP_L . Using this technique, and given known results on GSP_B in SNE [7], we establish that GSP_L not only optimizes publishers utility with respect to bids, but also with respect to true valuations:

PROPOSITION 2. *Let B_Φ be SNE bids in GSP_L with valuations H_Φ . Then, out of all rankings over Φ , Ω_L maximizes the following term:*

$$\sum_{j=1}^{|\Omega_L|} (f_{\Omega_L(j)} + g_{\Omega_L(j)} h_{\Omega_L(j)}) \beta^j \quad (9)$$

When $L(\cdot)$ is defined using publisher utility as in (4), we see that the optimized term is precisely the total expected publisher’s utility, with the true value h_σ in place of the bid b_σ . Intuitively, Proposition 2 shows that in equilibrium GSP_L “does the right thing” with respect to publisher utility. More precisely, following Lahaie [6], we call a ranking that would have been selected if we knew the true bidders’ valuations a *standard allocation*, and this result shows that in SNE, GSP_L selects a standard allocation.

5. SIMULATION RESULTS

We performed simulation experiments, using real advertisers’ bids submitted to Microsoft Live Search, for an arbitrary set of 470 queries, and a total of 2557 bids. We retrieved up to 8 bids per query, and limited the number of ads shown to 6. We measured the ratio of total expected revenue of the ads (as sum of prices multiplied by CTR) to relevance. The publisher makes an explicit choice of the utility model and the tradeoff between user utility and bids; the results, however, are measured directly in terms of relevance and actual prices. This ratio cannot capture the subtleties of the user experience as we model it, but it helps the publisher find the right tradeoff in terms of the values of t_a and t_m .

We defined the relevance score to be the relevance probability minus 0.5, so that ads more likely to be irrelevant have a negative effect on the score. We aggregated ad relevance across positions using the information retrieval notion of discounted cumulative gain [4]. In all of our simulations, we tested the ratio obtained by varying t_a from 0 to 99 cents.

Figure 1 shows the results of using the utility LC with three choices of t_m (0, 0.4, 0.8), in comparison with EQ and UER (with $t_m = 0.4$). We omit the series of EQ which is slightly above UER. The choices for t_a and t_m should not be made independently: for small values of t_a , a smaller t_m



Figure 1: Revenue vs. relevance

achieves higher relevance for the same amount of revenue loss, but for larger values of t_a , a larger t_m is better.

The simple model UER is clearly dominated by LC, and also by EQ. This is due to the fact that UER never assigns negative utility, and therefore always shows all available ads. LC does not show some of the less relevant ads, resulting in a better score and ratio despite the revenue loss from the dropped ads.

Further, we note the steep slope on the right end of each line, which corresponds to t_a varying from 0 to 30–40 cents. This shows how using the additive factor allows significant gains in relevance for small amounts of revenue loss, verifying our expectations based on the discussion in Section 3.

6. CONCLUSIONS

The revenue model of sponsored search relies on user clicks, which in turn depend on whether or not the ads provide positive utility to the user. Thus to maintain search traffic and ad-revenue for the long-term, it is important to address the user’s experience in the ad selection mechanism as we have done in this paper.

In a related work, Athey and Ellison [1] compute the equilibrium of a game that includes the consumers as well as the advertisers, and model the rational choice made by the consumers whether to click or not based on the probability that the ad will meet the consumer’s need and the cost of clicking on a useless ad. Our work is different in many ways, most notably because we model relevance as an external, measurable term, employed to improve the ranking.

7. REFERENCES

- [1] S. Athey and G. Ellison. Position auctions with consumer search. Unpublished, 2007.
- [2] T. Borgers, I. Cox, M. Pesendorfer, and V. Petricek. Equilibrium bids in sponsored search auctions: Theory and evidence. Unpublished, 2007.
- [3] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97:242–259, 2007.
- [4] K. Järvelin and J. Kekäläinen. IR evaluation methods for retrieving highly relevant documents. In *Proceedings of SIGIR*, pages 41–48, 2000.
- [5] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, 1976.
- [6] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *7th ACM Conference on Electronic Commerce*, pages 218–227, 2006.
- [7] H. Varian. Position auctions. *International Journal of Industrial Organization*, forthcoming.