

Sponsored Search Auctions with Markovian Users

Gagan Aggarwal

Google, Inc.
1600 Amphitheatre Pkwy
Mountain View, CA, 94043
gagana@google.com

S. Muthukrishnan

Google, Inc.
76 Ninth Avenue, 4th Floor
New York, NY, 10011
muthu@google.com

Jon Feldman

Google, Inc.
76 Ninth Avenue, 4th Floor
New York, NY, 10011
jonfeld@google.com

Martin Pál

Google, Inc.
76 Ninth Avenue, 4th Floor
New York, NY, 10011
mpal@google.com

ABSTRACT

Sponsored search involves running an auction among advertisers who bid in order to have their ad shown next to search results for specific keywords. Currently, the most popular auction for sponsored search is the “Generalized Second Price” (GSP) auction in which advertisers are assigned to slots in the decreasing order of their *score*, which is defined as the product of their bid and click-through rate. In the past few years, there has been significant research on the game-theoretic issues that arise in an advertiser’s interaction with the mechanism as well as possible redesigns of the mechanism, but this ranking order has remained standard.

From a search engine’s perspective, the fundamental question is: what is the best assignment of advertisers to slots? Here “best” could mean “maximizing user satisfaction,” “most efficient,” “revenue-maximizing,” “simplest to interact with,” or a combination of these. To answer this question we need to understand the behavior of a search engine user when she sees the displayed ads, since that defines the *commodity* the advertisers are bidding on, and its value. Most prior work has assumed that the probability of a user clicking on an ad is independent of the other ads shown on the page.

We propose a simple Markovian user model that does not make this assumption. We then present an algorithm to determine the most efficient assignment under this model, which turns out to be different than that of GSP. A truthful auction then follows from an application of the Vickrey-Clarke-Groves (VCG) mechanism. Further, we show that our assignment has many of the desirable properties of GSP that makes bidding intuitive. At the technical core of our result are a number of insights about the structure of the optimal assignment.

1. Introduction

Targeted advertisements on search queries is an increasingly important advertising medium, attracting large numbers of advertisers and users. When a user poses a query, the search engine returns search results together with advertisements that are placed into positions, usually arranged linearly down the page, top to bottom. On most major search engines, the assignment of ads to positions is determined by

an auction among all advertisers who placed a bid on a keyword that matches the query. The user might click on one or more of the ads, in which case (in the pay-per-click model) the advertiser receiving the click pays the search engine a price determined by the auction.

In the past few years, the sponsored search model has been highly successful commercially, and the research community is attempting to understand the underlying dynamics, explain the behavior of the market and to improve the auction algorithms. The most common auction being run today is the *Generalized Second Price* (GSP) auction: Each bidder i submits a bid b_i stating the maximum amount they are willing to pay for a click, and the bidders are placed in descending order of $b_i p_i$, where p_i is what is called the *click-through-rate* of advertiser i ; i.e., the probability that a user will click on the ad, given that the user looks at it. Much of previous research on sponsored search auctions has fixed this sort order, and focused on understanding the implications of different pricing schemes, assuming strategic behavior on the part of the advertisers. We now know something about GSP’s equilibrium properties [12, 29, 4], alternative pricing that will make it truthful [4], and to some extent, impact on the revenue in principle [12] and via simulations [28].

However, by fixing this sort order, prior work exogenizes an important third party in sponsored search; i.e., the *search engine user*. Unfortunately, there is very little guidance on this in the literature, even though the user’s behavior is the essential ingredient that defines *the commodity* the advertisers are bidding on, and its value. We suggest a different framework for principled understanding of sponsored search auctions:

- Define a suitable probabilistic model for search engine user behavior upon being presented the ads.
- Once this model is fixed, ask the traditional mechanism design questions of how do assign the ads to slots, and how to price them.
- Analyze the given mechanism from the perspective of the bidders (e.g., strategies) and the search engine (e.g., user satisfaction, efficiency and revenue).

There are certain well-accepted observations about the user’s interaction with the sponsored search ads that should inform the model: (i) The higher the ad is on the page, the more clicks it gets. (ii) The “better” the ad is, the more clicks it gets, where the “goodness” of an ad is related to the inherent quality of the ad, and how well it matches the user’s query. These properties govern not only how the auction is run but also how advertisers think about their bidding strategy (they prefer to appear higher and get more clicks). Thus it is important for an auction to have what we call *intuitive bidding*: a higher bid translates to a higher position and more clicks.

In this paper, we propose a natural Markov model for user clicks, taking the above observations into account and design an algorithm to determine an optimal assignment of ads to positions in terms of economic efficiency. Together with VCG pricing, this gives a truthful auction. We further show that the optimal assignment under this model has certain monotonicity properties that allow for intuitive bidding. In what follows, we will describe our technical contributions in more detail.

Modeling the Search Engine User. Most previous work on sponsored search has (implicitly) modeled the user using two types of parameters: ad-specific click-through rates p_i and position-specific visibility factors α_j . There are some intuitive user behavior models that express overall click-through probabilities in terms of these parameters. One possibility is “for each position j *independently*, the user looks at the ad i in that position with probability α_j then clicks on the ad with probability p_i .” Alternatively: “The user picks a *single* position according to the distribution implied by the α_j ’s, and then clicks on the ad i in that position with probability p_i .” Under both these models, it follows that the probability of an ad i in position j receiving a click is equal to $p_i\alpha_j$, which is the so-called *separability* assumption [4]. From separability it follows that GSP ordering of ads will be suitable, because GSP ordering maximizes the total advertiser value on the page.

In both these models there is no reason *a priori* that the position factors α_j should be decreasing; this is simply imposed because it makes sense, and it is verifiable empirically. Also, both suggested models assume that the probability of an ad getting clicked is independent of *other ads* that appear with it on the page, an assumption made without much justification. It is hard to imagine that seeing an ad, perhaps followed by a click, has no effect on the subsequent behavior of the user.

In designing a user model, we would like to have the monotonicity of the positions arise naturally. Also, each ad should have parameters dictating their effect on the user both in terms of clicking on that ad, as well as looking at other ads. We propose a model based on a user who starts to scan the list of ads from the top, and makes decisions (about whether to click, continue scanning, or give up altogether) based on what he sees. More specifically, we model the user as the following Markov process: “Begin scanning the ads from the top down. When position j is reached, click on the ad i with probability p_i . Continue scanning with probability q_i .” In

this model, if we try to write the click probability of an ad i in position j as $p_i\alpha_j$, we get that $\alpha_j = \prod_{i' \in A} q_{i'}$, where A is the set of ads placed above¹ position j . Thus the “position factor” in the click probability decreases with position, and does so naturally from the model. Also note that we do not have separability anymore, since α_j depends on which ads are above position j . Consequently, it can be shown that GSP assignment of ads is no longer the most efficient.

Auction with Markovian users. Given this new user model, we can now ask what the best assignment is of ads to slots. We will study the most efficient assignment; i.e., the one that maximizes total advertiser value derived from user clicks. It turns out that the structure of this assignment is different than that of GSP, and indeed is more sophisticated than any simple ranking. The presence of the q_i ’s requires a delicate tradeoff between the click probability of an ad and its effect on the slots below it. In this paper, we identify certain structural properties of the optimal assignment and use them to find such an optimal assignment efficiently, not only in polynomial time, but in near-linear time. Given this algorithm, a natural candidate for pricing is VCG [30, 10, 16], which is clearly truthful in this setting.

Intuitive Bidding. One of the reasons why GSP is successful is perhaps because bidding strategy is intuitive: Under GSP ranking, if an advertiser bids more, they get to a higher position, and consequently, if they bid more, their click probability increases. Now that we have defined a more sophisticated assignment function, even though VCG pricing is truthful, the auction still may not have these intuitive properties. Our main technical result is to show that in our model, if a mechanism uses the most efficient assignment, indeed position and click probabilities are monotonic in an ad’s bid (with all other bids fixed), thus preserving this important property. Monotonicity of click probability follows from the general result of Archer and Tardos [5] on single-parameter mechanisms—for completeness we provide a proof from first principles. In contrast, position monotonicity turns out to be rather involved to prove, requiring some delicate combinatorial arguments, and insights into the optimal substructure of bidder assignments.

In summary, we approach sponsored search auctions as a three party process by modeling the behavior of users first and then designing suitable mechanisms to affect the game theory between the advertiser and the search engine. Our work sheds some light on the intricate connection between the user models and the mechanisms; for example, the sort order of GSP that is currently popular (sort by $b_i p_i$) is not optimal under the Markov user model we propose here. More powerful models will be of great interest, such as making the continuation probability q_i a function of position as well, endogenizing the actions of the user as they navigate on the landing page, etc. We leave it open to design truthful auctions under such extended models.

¹Throughout the paper, we will often refer to a position or an ad being “higher” or “above” another position or ad; this means that it is earlier on the list, and is looked at first by the user.

1.1 Related Work. Sponsored search has been an active area of research in the last several years after the early papers explored the foundational models [12, 4, 29, 21]. In general, the motivation for the work that followed is that sponsored search in practice is much more complex than as described by the first models. Some papers have taken on the effect of advertiser budgets [8, 25, 2], as well as analyzing bidder strategy and dynamics [7, 27, 9, 13, 32, 31, 22]. There have also been several papers offering extensions to GSP, or entirely new models and mechanisms [3, 20, 23, 14, 26, 24, 1].

Only very recently are alternate user models that break the separability assumption starting to receive some attention. Ghosh and Mahdian [15] study a very general model and show hardness results for the allocation (winner determination) problem; they also give algorithms for several special cases, but none of those imply the algorithms discussed in this work. Craswell et al. [11] give an empirical study of several user click models. The “cascade” model, which was found to fit the data the best, is a special case of the model we study here (with $p_i = 1 - q_i$ and the events being mutually exclusive). Gunawardana and Meek [17] performed an empirical study of ad aggregators with the goal of detecting the affect of an ad on the other ads on the page. Their findings were consistent with our model; i.e., the presence of an ad can have a significant affect on the ads below it. Athey and Ellison [6] present a model where users have an inherent need, and click until that need is filled (or there is little chance of it getting filled). They analyze user behavior, advertiser bidding strategies and Bayesian equilibria in their model.

Independently of our work, Mahdian and Kempe [18] study the same model we do here. They also provide an $O(n \log n + nk)$ dynamic program for allocation; however at that point they generalize to the case of position-dependent continuation probabilities and provide an approximation algorithm for this case, whereas we go on to study deeper structural and incentive properties in the original model.

1.2 Outline. In Section 2 we define our model formally. In Section 3, we establish several properties of optimal assignments in this model, including our main technical result that position and click probability will be monotone in bid and match our intuition. We give our algorithm for finding an optimal assignment in Section 4 which gives the truthful auction via VCG pricing, and conclude in Section 5.

2. Markov User Click Model

We consider a sponsored search auction with n bidders $\mathcal{B} = \{1, \dots, n\}$ and k positions. We will also refer to “ad i ,” meaning the advertisement submitted by bidder i . Each bidder $i \in \mathcal{B}$ has two parameters, p_i and q_i . The click-through-rate p_i is the probability that a user will click on ad i , given that they *look* at it. The continuation probability q_i is the probability that a user will look at the next ad in a list, given that they look at ad i .

Each bidder submits a bid b_i to the auction, representing the amount that they value a click. The quantity $p_i b_i$ then represents the value of an “impression,” i.e., how much they

value a user looking at their ad. This is commonly referred to as their “ecpm.”² Throughout, we will use the notation $e_i = p_i b_i$ for convenience.

Given an assignment (x_1, \dots, x_k) of bidders to the k positions, the user looks at the first ad x_1 , clicks on it with probability p_{x_1} , and then continues looking with probability q_{x_1} .³ This is repeated with the second bidder, etc., until the last ad is reached, or some continuation test has failed. Thus the overall expected value of the assignment to the bidders is

$$e_{x_1} + q_{x_1}(e_{x_2} + q_{x_2}(e_{x_3} + q_{x_3}(\dots q_{x_{n'-1}}(e_{x_n}))))).$$

The goal of the auctioneer is to compute an assignment of ads to positions that maximizes the overall expected value. Given this assignment, prices can be computed using VCG [30, 10, 16]: for each assigned bidder we compute the change in others’ value if that bidder were to disappear. This assures truthful reporting of bids under a profit-maximizing utility function.

3. Properties of Optimal Assignments

We will start analyzing some basic properties of the optimal assignment. Our insights will allow us to give our main results regarding monotonicity of position and click probability, as well as an efficient algorithm for finding this assignment.

3.1 Adjusted ECPM. It turns out that the quantity $e_i/(1 - q_i)$, which we will refer to as the “adjusted ecpm (a-ecpm),” plays a central role in this model. Intuitively, this quantity is the impression value adjusted by the negative effect this bid has on the ads below it. We use $a_i = e_i/(1 - q_i)$ for convenience. The following theorem⁴ tells us how to assign a set of k selected ads to the k positions:

THEOREM 1. *In the most efficient assignment, the ads that are placed are sorted in decreasing order of adjusted ecpm $a_i = e_i/(1 - q_i)$.*

PROOF. Suppose not. Then in the ranking there are two consecutive ads i and i' in positions j and $j + 1$ where

$$\frac{e_i}{1 - q_i} < \frac{e_{i'}}{1 - q_{i'}}. \quad (1)$$

The contribution of positions $j \dots n$ to the efficiency of the ranking (given that position j is reached) is

$$e_i + q_i(e_{i'} + q_{i'} \hat{e})$$

²The acronym ecpm stands for “expected cost per thousand” impressions, where M is the roman numeral for one thousand. We will drop the factor of one thousand and refer to $p_i b_i$ as the “ecpm.”

³The click event and the continuation event could in principle have some correlation, and all our results will still hold. However since we only consider expected value, we never use this correlation explicitly in our analysis.

⁴Interestingly, this theorem essentially follows from a classical result on optimizing database queries [19].

where \hat{e} is the efficiency of positions $j + 2 \dots k$ given that position $j + 2$ is reached. If i and i' are switched, then the contribution would change to

$$e_{i'} + q_{i'}(e_i + q_i \hat{e}),$$

and nothing else would change. So since the former is the most efficient assignment, we have

$$e_i + q_i(e_{i'} + q_{i'} \hat{e}) > e_{i'} + q_{i'}(e_i + q_i \hat{e})$$

and so

$$e_i(1 - q_{i'}) \geq e_{i'}(1 - q_i).$$

This contradicts (1). \square

While this theorem tells us how to sort the ads selected, it does not tell us *which* k ads to select. One is tempted to say that choosing the top k ads by a-ecpm would do the trick; however the following example proves otherwise:

EXAMPLE 1. *Suppose we have three bidders and two slots, and the bidders have the following parameters:*

Bidder	e_i	q_i	$a_i = e_i/(1 - q_i)$
1	\$1	.75	4
2	\$2	.2	2.5
3	\$0.85	.8	4.25

Let's consider some possible assignments and their efficiency. If we use simple ranking by ecpm e_i , we get the assignment (2, 1), which has efficiency $\$2 + .2(\$1) = \$2.20$. If we use simple ranking by a-ecpm a_i we get the assignment (3, 1) with efficiency $\$0.85 + .8(\$1) = \$1.65$. It turns out that the optimal assignment is (1, 2) with efficiency $\$1 + .75(\$2) = \$2.50$. The assigned bidders are ordered by a-ecpm in the assignment, but are not the top 2 bidders by a-ecpm.

Now suppose we have the same set of bidders, but now we have three slots. The optimal assignment in this case is (3, 1, 2); note how bidder 3 goes from being unassigned to being assigned the first position.

3.2 Notation for Assignments. Theorem 1 implies that the optimal assignment can be described by the set of assigned bidders, since this set will always be sorted by a-ecpm. For a set X of bidders, let $e(X)$ be the efficiency of an assignment that places only the set X , regardless of the number of bidders in X ; i.e., if $X = \{1, 2, 3, \dots, n'\}$ and wlog is sorted by a-ecpm, then

$$e(X) = e_1 + q_1(e_2 + q_2(e_3 + q_3(\dots q_{n'-1}(e_{n'}))))).$$

Similarly, let $q(X)$ be the overall continuation probability of the set X of ads: $q(X) = \prod_{i' \in X} q_{i'}$. If $X = \emptyset$ we have $e(X) = 0$ and $q(X) = 1$.

Throughout the paper, we will also make use of the notation (A, x, B, C, y, \dots) for a particular assignment, where uppercase letters denote sets of bidders and lowercase letters denote single bidders. This denotes a solution where the

order of the bidders is as given, and where the order within a set is by a-ecpm (breaking ties using a lexicographic order on bidders). So in this example, the solution would put all the bidders of A first in a-ecpm order, followed by bidder x , followed by the bidders of B in a-ecpm order, etc. Note that this notation could express suboptimal solutions (for the particular set of bidders) if the elements are not in order of ecpm. We use $e(\cdot)$ and $q(\cdot)$ to denote the efficiency and continuation probability of such solutions.

3.3 Bidder Dominance. In classical sponsored search with simple ranking, a bidder j can dominate another bidder i by having higher ecpm; i.e., bidder j will always appear whenever i does, and in a higher position. Example 1 above shows that having a higher ecpm (or a-ecpm) does not allow a bidder to dominate another bidder in our new model. However, we show in this section that if she has higher ecpm *and* a-ecpm, then this does suffice. This is not only interesting in its own right, it is essential for proving deeper structural properties in later sections.

THEOREM 2. *For all bidders i in an optimal assignment, if some bidder j is not in the assignment, and $a_j \geq a_i$ and $e_j \geq e_i$, then we may substitute j for i , and the assignment is no worse.*

PROOF. Consider some bidder i that appears in an optimal assignment, and some j that does not appear in the assignment, such that $e_j \geq e_i$ and $a_j \geq a_i$. Let (X, i, Y) be the optimal assignment, where X is the sequence of bidders above i and Y is the sequence below i . The efficiency of the assignment (X, i, Y) is $e(X) + q(X)(e_i + q_i e(Y))$. The efficiency of the assignment (X, j, Y) is $e(X) + q(X)(e_j + q_j e(Y))$. Suppose $q_j \geq q_i$; then clearly (X, j, Y) is as efficient as (X, i, Y) since $e_j \geq e_i$, and the theorem is proven. Thus we assume that $q_i > q_j$ for the remainder of the proof. Note that $a_j \geq a_i$ is equivalent to:

$$e_j - e_i \geq \frac{e_i(q_i - q_j)}{1 - q_i} \quad (2)$$

Now consider the assignment (X, Y) , with efficiency $e(X) + q(X)e(Y)$. Since (X, i, Y) is optimal, we get

$$e(X) + q(X)e(Y) \leq e(X) + q(X)(e_i + q_i e(Y)),$$

i.e., $e(Y) \leq e_i/(1 - q_i)$. Combining this with (2), and using the fact that $q_i > q_j$, we get $e_j - e_i \geq e(Y)(q_i - q_j)$ which can be rewritten as

$$e_j + q_j e(Y) \geq e_i + q_i e(Y).$$

This implies that the solution (X, j, Y) is as efficient as (X, i, Y) . \square

3.4 Subset Substructure in Optimal Assignments.

In this section we give a theorem that shows some subset structure between optimal assignments to different numbers of slots. This theorem is used to prove position monotonicity, and is an essential ingredient of our algorithm. Let $\text{OPT}(C, j)$ denote the set of all optimal solutions for filling j positions with bidders from the set C .

THEOREM 3. *Let $j \in \{1, \dots, k\}$ be some number of positions, and let C be an arbitrary set of bidders. Then, for all $S \in \text{OPT}(C, j-1)$, there is some $S' \in \text{OPT}(C, j)$ where $S' \supset S$.*

PROOF. We proceed by induction on j , the base case $j = 1$ being simple. Let S be some solution in $\text{OPT}(C, j-1)$, and let S' be the solution in $\text{OPT}(C, j)$ with the most bidders in common with S . We will refer to an ad being ‘‘above’’ another ad if it has higher a-ecpm. Let x be the highest bidder in S' .

If x does not appear in S , then we claim that the solution (x, S) must be in $\text{OPT}(C, j)$: Once x is chosen for S' , taking any set in $\text{OPT}(C \setminus x, j-1)$ for the remaining positions will result in an optimal solution; the set S is such a set, since by assumption it does not include x , and is in $\text{OPT}(C, j-1)$. But if $(x, S) \in \text{OPT}(C, j)$ the theorem is proven, so we may assume $x \in S$.

Let A be the set of ads in S above x , and so we can write $S = (A, x, Q)$ and $S' = (x, Q')$. We claim $Q' \supset Q$. To see this, consider the set L of all ads that have lower a-ecpm than x . By the optimality of S , we have $Q \in \text{OPT}(L, j')$ for $j' = j - |A| - 1 < j$. By induction there is a $Q'' \in \text{OPT}(L, j-1)$ where $Q'' \supset Q$. Thus we must have $Q' \supset Q$ since S' is the solution with the most bidders in common with S . Decompose Q' as $Q' = (B + X, D, z, E)$ where

- z is the lowest ad in S' that does not appear in S ,
- E is the set of ads below z in S' (this can be empty),
- D is the maximal set of ads immediately above z in S' that also appear in S (this can be empty),
- X are the remaining ads in S' that do not appear in S ,
- B are the remaining ads (besides x) that appear in both S and S' .

Let $B' = B \cup x$. Note that by the definitions above we may write $S = (A, B', D, E)$. We have $e(S) = e(A, B', D, E) \geq e(B' + X, D, E)$ since $S \in \text{OPT}(C, j-1)$ and $|(B' + X, D, E)| = j - 1$. Decomposing this a bit gives

$$\begin{aligned} e(A, B', D) + q(A)q(B')q(D)e(E) \\ \geq e(B' + X) + q(B')q(X)e(D, E). \end{aligned} \quad (3)$$

We also note that

$$\begin{aligned} e(S') &= e(B' + X, D, z, E) \\ &= e(B' + X) + q(B')q(X)e(D, z, E). \end{aligned} \quad (4)$$

Let $S'' = (A, B', D, z, E)$. The remainder of the proof will show that $e(S'') \geq e(S')$, which implies the theorem since $|S''| = j$, $S'' \supset S$ and $S' \in \text{OPT}(C, j)$. We can rewrite

$e(S'')$ as follows:

$$\begin{aligned} e(S'') &= e(A, B', D) + q(A)q(B')q(D)e(z, E) \\ &\geq e(B' + X) + q(B')q(X)e(D, E) \\ &\quad + q(A)q(B')q(D)(e(z, E) - e(E)) \quad (5) \\ &= e(S') - q(B')q(X)(e(D, E) - e(D, z, E)) \\ &\quad + q(A)q(B')q(D)(e(z, E) - e(E)) \quad (6) \end{aligned}$$

In the above, (5) follows from (3), and (6) follows from (4). Rearranging, and using

$$e(D, E) - e(D, z, E) = q(D)(e(E) - e(z, E)),$$

we get

$$e(S'') - e(S') = q(B')q(D)(q(A) - q(X))(e(z, E) - e(E)).$$

We know that $e(z, E) \geq e(E)$ since otherwise $e(B' + X, D, E) > e(B' + X, D, z, E) = e(S')$, and this cannot be since $S' \in \text{OPT}(C, j)$. We claim that $q(A) \geq q(X)$, which would imply $e(S'') \geq e(S')$ and thus complete the proof. This is trivially true if $A = X = \emptyset$. Since $|A| = |X|$ by the definitions above, we can assume both A and X are non-empty. Consider some $y \in A$ and $y' \in X$. We have $y \notin S'$ by the definition of A . Since S' is the solution in $\text{OPT}(C, j)$ with the most bidders in common with S , we must not be able to substitute y for y' in S , and thus by Lemma 2 we must have that $e_{y'} > e_y$ or $a_{y'} > a_y$. But by the definitions of A and X , we have $a_y \geq a_x \geq a_{y'}$. Therefore $e_{y'} > e_y$. The previous two inequalities imply $q_y > q_{y'}$. Since y and y' were arbitrary and $|A| = |X|$, this gives $q(A) > q(X)$. \square

3.5 Monotonicity of Position and Click Probability. In this section we give our main theorem regarding the structure of the optimal assignments in the Markovian click model: that position and click probability are monotonic in a bidder’s bid, with all other bids fixed. This is a fundamental property that makes the bidder’s interaction with the system intuitive, and allows the bidder to adjust her bid intelligently without global knowledge of the other bids.

THEOREM 4. *As a bidder increases her bid (keeping all other bids fixed):*

- (a) *the probability of her receiving a click in the optimal solution does not decrease, and*
- (b) *her position in the optimal solution does not go down.*

PROOF. As bidder x increases her bid b_x with all other bids fixed, the value of a particular solution $S = (A, x, B)$ increases linearly as $q(A)p_x \cdot b_x + [e(A) + q(A)q_x e(B)]$. (Solutions not involving x stay constant.)

Let S_1, \dots, S_k denote the sequence of optimal solutions that occurs as b_x increases from 0. Solution S_1 is the best solution not involving x , and S_k is the best solution that puts x in the first position. By the fact that each solution increases linearly by the term $q(A)p_x \cdot b_x$, which is the probability that x receives a click in that solution, it must be the case that for a new solution to become optimal it gives x a higher click

probability than in the previous solution; i.e., for all $i \geq 0$, S_{i+1} gives x a higher click probability than S_i . This proves the first part of the theorem.⁵

Now suppose the second part of the theorem is false. Then, there must be some consecutive solutions S_i and S_{i+1} where x has a higher position in S_i than in S_{i+1} . Let b be the bid that makes both S_i and S_{i+1} optimal, which must exist since they are consecutive in the list of optimal solutions, and fix b for the remainder of the proof. Decompose the two solutions as $S_{i+1} = (A, x, E)$ and $S_i = (F, x, G)$ where $|A| > |F|$ by assumption, and $q(A) > q(F)$ by the argument that proved the first part of the theorem. Since both solutions are optimal, they are both sorted by a-ecpm, with ties broken lexicographically.

We claim that $A \cap G = \emptyset$. If this were not the case then some bidder y would appear in both A and G , but since both S_i and S_{i+1} are sorted by ecpm, and y appears on different sides of x , this must mean that y has the same a-ecpm as x . But, this violates our assumption on how the algorithm breaks ties among different orderings of the same set. Using similar logic, we get $F \cap E = \emptyset$.

By the optimality of S_{i+1} and the fact that $F \cap E = \emptyset$, we get $E \in \text{OPT}(\mathcal{B} \setminus (F \cup A \cup \{x\}), n - |A| - 1)$. Since $|F| < |A|$, Theorem 3 then implies that there is some $G' \in \text{OPT}(\mathcal{B} \setminus (F \cup A \cup \{x\}), n - |F| - 1)$ where $G' \supset E$. Even if $G' \neq G$, the set G' could replace G in S_i (since $A \cap G = \emptyset$) and still be optimal, and so we define $S'_i = (F, x, G')$ and have $e(S'_i) = e(S_i)$. Note that by the definition of G' , we have $A \cap G' = \emptyset$.

Since $A \cap G' = \emptyset$ and $|A| > |F|$, there must be some bidder in A that does not appear in S'_i . Let a be the first such bidder (by a-ecpm). Decompose A into (C, a, D) where C and D are those bidders with higher and lower a-ecpm than a , respectively. Note that $C \subseteq F$, by the definition of a . Let F_1 be the smallest prefix of F (by a-ecpm) that contains all of C , and let $F_2 = F - F_1$.

Similarly, since $A \cap G' = \emptyset$ and $|G| > |E|$, we let $t \in G'$ be the first bidder (by a-ecpm) that does not appear in S_{i+1} . Let G'_1 be the bidders in G' with higher a-ecpm than t , and $G'_2 = G' - G'_1 - t$. Note G'_1 is also a prefix of E , and let $E_2 = E - G'_1$.

Given these definitions, we define nine different solutions that we will use in our proof (renaming S'_i and S_{i+1} for clarity):

$$\begin{aligned} \alpha &= S_{i+1} = (A, x, E) = (C, a, D, x, E) \\ \alpha' &= (C, D, x, E) = (C, D, x, G'_1, E_2) \\ \beta &= S'_i = (F, x, G') = (F, x, G'_1, t, G'_2) \\ \beta' &= (F, x, G'_1, G'_2) \\ \gamma &= (C, D, x, E + t) = (C, D, x, G'_1, t, E_2) \\ \delta &= (F_1, a, F_2, x, G'_1, G'_2) \\ \zeta &= (C, a, F_1 - C, F_2, x, G'_1, G'_2) \\ \zeta' &= (C, F_1 - C, F_2, x, G'_1, G'_2) \\ \eta &= (C, F_1 - C, a, F_2, x, G'_1, G'_2) \end{aligned}$$

In the following claims we will often use the generic identity

$$e(X, y, Y) - e(X, Y) = q(X)(e_y - (1 - q_y)e(Y)) \quad (7)$$

CLAIM 1. $e(\gamma) - e(\alpha') > e(\beta) - e(\beta')$

PROOF. Using (7), we can rewrite the claim as

$$\begin{aligned} q(C, D, x, G'_1)(e_t - (1 - q_t)e(E_2)) \\ > q(F, x, G'_1)(e_t - (1 - q_t)e(G'_2)). \end{aligned} \quad (8)$$

Since $q(C, D, x) \geq q(C, a, D, x) = q(A, x)$, and $q(F) < q(A)$, we get $q(C, D, x, G'_1) > q(F, x, G'_1)$, and so it remains to prove $e(E_2) \leq e(G'_2)$. But this follows from the optimality of β , since E_2 could replace G'_2 in solution β (indeed, $E_2 \subset G'_2$). \square

CLAIM 2. $e(\zeta) - e(\zeta') \geq e(\alpha) - e(\alpha')$

PROOF. Using (7), we can rewrite the claim as

$$\begin{aligned} q(C)(e_a - (1 - q_a))e(F_1 - C, F_2, x, G'_1, G'_2) \\ \geq q(C)(e_a - (1 - q_a))e(D, x, E). \end{aligned} \quad (9)$$

Since $(F_1 - C, F_2, x, G'_1, G'_2)$ has the same length as (D, x, E) , and does not contain a or any bidders in C , it could replace (D, x, E) in α ; but since α is optimal, we may conclude that $e(D, x, E) \geq e(F_1 - C, F_2, x, G'_1, G'_2)$, which proves the claim. \square

Note that $e(\eta) \geq e(\zeta)$, since all bidders in F_1 have a higher a-ecpm than a . Finally, note that a simple application of (7) gives

$$e(\eta) - e(\zeta') = e(\delta) - e(\beta') \quad (10)$$

We now conclude the proof with the following contradiction:

$$\begin{aligned} e(\alpha) - e(\alpha') &\geq e(\gamma) - e(\alpha') && \text{(by opt. of } \alpha) \\ &> e(\beta) - e(\beta') && \text{(Claim 1)} \\ &\geq e(\delta) - e(\beta') && \text{(by opt. of } \beta) \\ &= e(\eta) - e(\zeta') && \text{(by (10))} \\ &\geq e(\zeta) - e(\zeta') \\ &\geq e(\alpha) - e(\alpha') && \text{(by Claim 2)} \end{aligned}$$

⁵This result can also be obtained by applying the general result of Archer and Tardos[5]: since this is a single-parameter setting, and we are using the VCG mechanism (which is truthful), the valuation function (in this case click probability) must be monotone in the bid.

4. Computing the Optimal Assignment

In this section we give algorithms for computing the optimal assignment of bidders to positions using the structural properties we proved in the previous section. We begin with a simple dynamic program that gives an $O(n \log n + nk)$ time algorithm. We then show how our insights from the previous sections give a faster $O(n \log n + k^2 \log^2 n)$ time algorithm.

4.1 Optimal Assignment using Dynamic Programming. The algorithm proceeds as follows. First, sort the ads in decreasing order of a-ecpm in time $O(n \log n)$. Then, let $F(i, j)$ be the efficiency obtained (given that you reach slot j) by filling slots (j, \dots, k) with bidders from the set $\{i, \dots, n\}$. We get the following recurrence:

$$F(i, j) = \max(F(i+1, j+1)q_i + e_i, F(i+1, j)).$$

Solving this recurrence for $F(1, 1)$ yields the optimal assignment, and can be done in $O(nk)$ time.

4.2 Near-linear Time Algorithm. Let $\mathcal{B} = \{1, \dots, n\}$ be the set of bidders, sorted by a-ecpm. Suppose we had an oracle that told us, for any $j, j' \in \mathcal{B}$, the bidder y with $j \leq y \leq j'$ that maximizes $f(q_y, e_y)$ for an arbitrary linear function f . We will later show how to construct this oracle, but first we describe our algorithm that uses this oracle.

Our algorithm will construct a solution $S_i \in \text{OPT}(\mathcal{B}, i)$ for all $i = 1, \dots, k$, the final one S_k being the overall optimum. By Theorem 3, we may assume that $S_{i+1} \supset S_i$. Using this fact, our algorithm builds S_{i+1} from S_i by simply finding $\arg \max_{x \notin S_i} e(S_i \cup \{x\})$. To perform this max, the algorithm first guesses (i.e., searches exhaustively for) the a-ecpm rank of the new bidder x among the bidders in S_i ; this a number ℓ from 1 to $i+1$. Let $\{s_1, \dots, s_i\}$ be the elements of S_i sorted by decreasing a-ecpm. The new bidder x has a-ecpm between $s_{\ell-1}$ and s_ℓ , and so $e(S_i \cup \{x\}) = e(s_1, \dots, s_{\ell-1}) + q(s_1, \dots, s_{\ell-1})(e_x + q_x e(s_\ell, \dots, s_i))$. Since $e(S_i \cup \{x\})$ is linear in (q_x, e_x) , we may appeal to the oracle to find the bidder x that maximizes $e(S_i \cup \{x\})$ among all bidders with a-ecpm between that of $s_{\ell-1}$ and s_ℓ . We make $i+1$ calls to this oracle for each i , and thus $O(k^2)$ calls overall. To get the coefficients of $e(S_i \cup \{x\})$ to pass to the oracle, we precompute the quantities $q(s_1, \dots, s_p)$ and $e(s_p, \dots, s_i)$ for all p . (This can be done in $O(k)$ time per i , for $O(k^2)$ time overall.)

It remains to show how to implement the oracle. We first preprocess the sequence $[1, \dots, n]$ of bidders as follows. We consider the *dyadic* intervals $[\alpha 2^\beta + 1, \dots, (\alpha+1)2^\beta]$ for each possible α, β , for a total of $O(n)$ intervals. Note that any subsequence $[j, \dots, j']$ is made up of at most $O(\log n)$ such intervals. For each such interval, we will make a data structure that can find $\max f(q_x, e_x)$ over bidders in that interval in $O(\log n)$ time. So overall, given $[j, \dots, j']$, the oracle takes the max of $O(\log n)$ calls to the data structure, and completes in $O(\log^2 n)$ time.

The data structure we compute for a particular interval $[\alpha 2^\beta + 1, \dots, (\alpha+1)2^\beta]$ is simply the convex hull of the points (q_x, e_x) in two-dimensional space defined by bidders x in the interval. We can compute all these convex hulls in $O(n \log n)$ time by successively merging convex hulls for

increasing β . Given the convex hull (with the points sorted in order of q_x for example), a simple binary search can find the point maximizing $f(q_x, e_x)$ in $O(\log n)$ time.

This gives,

THEOREM 5. *Consider the auction with n Markovian bidders and k slots. There is an optimal assignment which can be determined in $O(n \log n + k^2 \log^2 n)$ time.*

It follows that using VCG pricing with this optimal assignment, we obtain a truthful mechanism for sponsored search with Markovian users.

5. Concluding Remarks

We approached sponsored search auctions as a three party process by modeling the behavior of users first and then designing suitable mechanisms to affect the game theory between the advertiser and the search engine. This formal approach shows an intricate connection between the user models and the mechanisms.

There are some interesting open issues to understand about our model and mechanism. For example, in order to implement our mechanism, the search engine needs to devise methods to estimate the parameters of our model, in particular, q_i 's. This is a challenging statistical and machine learning problem. Also, we could ask how much improvement in efficiency and/or revenue is gained by using our model as opposed to VCG without using our model.

More powerful models will also be of great interest. One small extension of our model is to make the continuation probability q_i a function of location as well, which makes the optimization problem more difficult. We can also generalize the Markov model to handle arbitrary configurations of ads on a web page (not necessarily a search results page), or to allow various other user states (such as navigating a landing page). Finally, since page layout can be performed dynamically, we could ask what would happen if the layout of a web page were a part of the mechanism; i.e., a function of the bids.

Acknowledgments

We thank the anonymous reviews for their comments and for pointing us to several references.

6. References

- [1] Z. Abrams and A. Ghosh. Auctions with revenue guarantees for sponsored search. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.
- [2] Z. Abrams, O. Mendelevitch, and J. Tomlin. Optimal delivery of sponsored search advertisements subject to budget constraints. In *ACM Conference on Electronic Commerce*, 2007.
- [3] G. Aggarwal, J. Feldman, and S. Muthukrishnan. Bidding to the top: VCG and equilibria of position-based auctions. In *Proc. Workshop on Approximation and Online Algorithms (WAOA)*, 2006.

- [4] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *ACM Conference on Electronic Commerce (EC)*, 2006.
- [5] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
- [6] Susan Athey and Glenn Ellison. Position auctions with consumer search. Levine's Bibliography 12224700000001633, UCLA Department of Economics, October 2007.
- [7] C. Borgs, J. Chayes, O. Etesami, N. Immorlica, K. Jain, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proc. WWW*, 2007.
- [8] C. Borgs, J. T. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *ACM Conference on Electronic Commerce (EC)*, 2005.
- [9] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. R. Karlin, C. Mathieu, and M. Schwarz. Greedy bidding strategies for keyword auctions. In *ACM Conference on Electronic Commerce*, 2007.
- [10] E. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [11] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. An experimental comparison of click position-bias models. In *WSDM '08: Proceedings of the international conference on Web search and web data mining*, pages 87–94, New York, NY, USA, 2008. ACM.
- [12] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In *Second workshop on sponsored search auctions*, 2006.
- [13] J. Feldman, S. Muthukrishnan, M. Pál, and C. Stein. Budget optimization in search-based advertising auctions. In *ACM Conference on Electronic Commerce*, 2007.
- [14] J. Feldman, E. Nikolova, S. Muthukrishnan, and M. Pál. A truthful mechanism for offline ad slot scheduling. In *Proc. Symposium on Algorithmic Game Theory (SAGT)*, 2008. To appear.
- [15] Arpita Ghosh and Mohammad Mahdian. Externalities in online advertising. In *Proceedings of the 17th International World Wide Web Conference (WWW)*, 2008.
- [16] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- [17] A. Gunawardana and C. Meek. Aggregators and contextual effects in search ad markets. In *WWW Workshop on Targeting and Ranking for Online Advertising*, 2008.
- [18] David Kempe and Mohammad Mahdian. Cascade model for externalities in sponsored search. In *Workshop on Internet Ad Auctions*, 2008.
- [19] Ravi Krishnamurthy, Haran Boral, and Carlo Zaniolo. Optimization of nonrecursive queries. In *VLDB '86: Proceedings of the 12th International Conference on Very Large Data Bases*, pages 128–137, San Francisco, CA, USA, 1986. Morgan Kaufmann Publishers Inc.
- [20] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *ACM Conference on Electronic Commerce (EC)*, 2006.
- [21] S. Lahie, D. Pennock, A. Saberi, and R. Vohra. *Sponsored Search Auctions*, in: *Algorithmic Game Theory*, pages 699–716. Cambridge University Press, 2007.
- [22] Li Liang and Qi Qi. Cooperative or vindictive: Bidding strategies in sponsored search auction. In *WINE*, pages 167–178, 2007.
- [23] M. Mahdian, H. Nazerzadeh, and A. Saberi. Allocating online advertisement space with unreliable estimates. In *ACM Conference on Electronic Commerce*, 2007.
- [24] M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. In *ACM conference on Electronic commerce (EC)*, 2006.
- [25] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. AdWords and generalized online matching. In *FOCS*, 2005.
- [26] S. Muthukrishnan, M. Pál, and Z. Svitkina. Stochastic models for budget optimization in search-based advertising. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.
- [27] P. Rusmevichientong and D. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *Proc. 7th ACM conference on Electronic commerce*, 2006.
- [28] B. K. Szymanski and J.-S. Lee. Impact of roi on bidding and revenue in sponsored search advertisement auctions. In *2nd Workshop on Sponsored Search Auctions in ACM Conference on Electronic Commerce (EC)*, 2006.
- [29] H. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, December 2007.
- [30] W. Vickrey. Counterspeculation, auctions and competitive-sealed tenders. *Finance*, 16(1):8–37, 1961.
- [31] Y. Vorobeychik and D. M. Reeves. Equilibrium analysis of dynamic bidding in sponsored search auctions. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.
- [32] J. Wortman, Y. Vorobeychik, L. Li, and J. Langford. Maintaining equilibria during exploration in sponsored search auctions. In *Proc. Workshop on Internet and Network Economics (WINE)*, 2007.