

Maintaining Equilibria During Exploration in Sponsored Search Auctions

Jennifer Wortman¹, Yevgeniy Vorobeychik²,
Lihong Li³, and John Langford⁴

¹ Computer and Information Science, University of Pennsylvania
`wortmanj@seas.upenn.edu`*

² Computer Science & Engineering, University of Michigan
`yvorobey@umich.edu`*

³ Computer Science, Rutgers University
`lihong@cs.rutgers.edu`*

⁴ Yahoo! Research
`jl@yahoo-inc.com`

Abstract. We introduce an exploration scheme aimed at learning advertiser click-through rates in sponsored search auctions with minimal effect on advertiser incentives. The scheme preserves both the current ranking and pricing policies of the search engine and only introduces one parameter which controls the rate of exploration. This parameter can be set so as to allow enough exploration to learn advertiser click-through rates over time, but also eliminate incentives for advertisers to alter their currently submitted bids. When advertisers have much more information than the search engine, we show that although this goal is not achievable, incentives to deviate can be made arbitrarily small by appropriately setting the exploration rate. Given that advertisers do not alter their bids, we bound revenue loss due to exploration.

1 Introduction

Recent years have seen an explosion of interest in sponsored search auctions, due in large part to the unique opportunity for targeted advertising and the resulting billions of dollars in revenue. Most sponsored search auctions display a list of advertisements on the sidebar or other sections of a search engine's results page, ranked by some function of advertisers' revealed willingness-to-pay for every click on their ad. The advertisers in turn pay the search engine for every click their ad receives. While several pricing schemes have been circulated in the literature [7], by far the most popular is a generalization of second-price auctions, under which each advertiser pays the lowest bid that is sufficient to ensure that the ad remain in its current slot. Typically the number of available slots for advertisements on the first search page is fixed, and thus only high ranking advertisements are displayed.

* This work was done while J. Wortman, Y. Vorobeychik, and L. Li were at Yahoo! Research, New York.

An essential part of both designing sponsored search auction mechanisms and bidding in them is the knowledge of the probability that a given ad is clicked each time it is displayed in a particular slot for a particular search query or keyword. This probability is known as the *click-through rate* or *CTR* of the ad. Knowledge of these click-through rates helps advertisers determine optimal bidding behavior. CTRs can also be an integral part of the ad ranking policy. For example, it is common for policies to rank bidders by the product of their bid and some function of their *relevance*, a slot-independent measure of CTR. Throughout the paper, we assume that CTRs do not change over time.

Most of the existing literature on sponsored search auctions treats CTRs as known. When advertisers first enter the system, however, their CTRs are not yet known either by the search engine or even by the advertisers themselves, and can only be estimated over time based on the observed clicks. Observations are inherently limited to slots in which ads appear, and estimates are generally poor for advertisers with low rank that do not usually appear at all. Furthermore, without the assumption of factorable CTRs, little can be said about CTRs of an ad in slots in which it has not previously appeared (or has appeared only a small number of times). Thus there is a need for an exploration policy that periodically perturbs the current slate of displayed ads, showing some in alternate slots and occasionally displaying those ads that are ranked below the last slot. Ideally, this exploration policy should not be difficult to incorporate into the current sponsored search mechanisms. Additionally, if the advertisers' bids have reached an equilibrium, the exploration policy should, when possible, eliminate the incentives for bidders to change their bids, thereby destabilizing the auction. Such destabilization can result in negative user and advertiser experience, as well as unnecessary loss in revenue to the search engine, and can make exploration harder to control.

In this paper, we address the problem of learning the click-through rates for each ad in every slot. Our primary goal is to maintain an equilibrium bid configuration if the bidders did indeed play according to an equilibrium prior to exploration. When this is not possible, we provide bounds on the amount that any advertiser could gain by deviating. This incentive to deviate can be minimized by reducing exploration, at the cost of slowing down the process of learning the CTRs. Additionally, we bound the revenue loss that the search engine incurs due to exploration, as compared to maintaining a policy based on current estimates of CTRs.

A similar problem has been addressed by Pandey and Olston [9] and Gonen and Pavlov [5]. The former work addresses the learning problem without considering advertiser incentives. The latter addresses both. Our model differs from existing ones in three primary ways:

1. We avoid imposing a particular ranking policy or introducing a new pricing scheme so that changes to existing systems are minimal.
2. The data gathered by our approach can be incorporated into general learning algorithms using sample selection debiasing techniques.[6]
3. We avoid the standard but unrealistic assumption that click-through rates can be factored into advertiser- and slot-specific components.

2 Notation and Definitions

We consider an auction for a particular keyword in which there are N advertisers (alternately called bidders or players) placing bids.¹ We assume that the search engine has K slots with non-negligible CTRs. Throughout the discussion on incentives, we assume that the CTRs depend only on the ad being displayed and the slot in which it is shown. Thus, we use c_i^s to denote the true CTR of player i in slot s . We assume that for each player i , $c_i^s > c_i^t$ whenever $1 \leq s < t \leq K$. For convenience, we define $c_i^s = 0$ for $s > K$ and $s < 1$. In most of our analysis we deal explicitly with estimated click-through rates; the search engine estimates are denoted by \hat{c}_i^s , whereas the advertiser i 's estimates are denoted by \tilde{c}_i^s . Finally, we let v_i denote the value of a click to player i .

For now we assume that throughout the exploration process, advertisers are ranked according to their bid b_i multiplied by a weight w_i which is an increasing function of their estimated relevance scores for the particular keyword. Setting this weight equal to relevance recovers the standard rank-by-revenue model. Without loss of generality, assume that advertisers are indexed in the order in which they are ranked when playing equilibrium, i.e. advertiser i is in slot i in the ranking. Each advertiser pays a price per click equal to the lowest bid that maintains his current position; thus the price paid by bidder i in rank s is $p_i^s = w_{s+1}b_{s+1}/w_i$.

The relevance score of an advertiser, which we denote by e_i , can be thought of as an average CTR over all slots for the given keyword. We might choose to define this relevance as $\sum_{s=1}^K c_i^s$ or alternately as $\sum_{s=1}^K c_i^s/c_s$ where c_s is the ‘‘average’’ CTR that any ad might expect to receive on slot s .² We can fix the weights for each advertiser prior to (each phase of) exploration and reveal the new estimates of CTRs at the end of the exploration period only, allowing greater control of exploration.

We assume that prior to exploration the advertisers converge to a symmetric Nash equilibrium, a variant of Nash equilibrium introduced simultaneously by Varian [10] and Edelman et al.[3]. We slightly alter the standard definition to take into account CTR estimates as follows.

Definition 1. *A symmetric Nash equilibrium (SNE) is an ordering and a set of bids such that for every player i and for every slot s , $\tilde{c}_i^i (v_i - p_i^i) \geq \tilde{c}_i^s (v_i - p_i^s)$, where \tilde{c}_i^s denotes advertiser i 's CTR estimate at slot s .*

Existence of at least one symmetric Nash Equilibrium was proved in a slightly different setting than ours by Börgers et al. [1]. Their proof applies essentially without change to our setting.

¹ Since our analysis can be repeated for each keyword, the restriction to a single keyword is without loss of generality. Indeed, the analysis can even be generalized to incorporate arbitrary context information, as long as the number of contexts is finite and advertisers may submit separate bids for each. [4]

² Observe that when c_i^s is factorable into the product $e_i c_s$, both of these relevance scores are proportional to e_i .

3 An Algorithm for Exploration

We begin by describing a simple algorithm for learning click-through rates. Below (in Section 4) we show that we can set parameters of this algorithm in such a way as to minimize or entirely eliminate incentives for advertisers to deviate from a pre-exploration SNE. Our key condition will be that throughout the entire run of the algorithm the prices which the advertisers pay are fixed to their pre-exploration equilibrium prices.

The algorithm, which we call **k-swap** (Algorithm 1), starts by ranking ads by the product of bid and weight as usual, and repeatedly chooses pairs of ads to swap in order to explore. In particular, each time the given keyword receives an *impression* (i.e. each time a query is made on the keyword), a swapping distance $k \in \{1, \dots, K\}$ is chosen from some distribution (e.g. uniformly at random). The algorithm calculates or looks up a swapping probability for each pair of slots s and $s + k$ that are a distance k apart. (The method for choosing these probabilities will be discussed in Section 4.) Finally, the algorithm uses this set of swapping probabilities to decide which (if any) pair of ads to swap.

We must be careful about how pairs of ads are chosen to be swapped so we can avoid swapping the same ad more than once on a single query. Let S_i denote the event that the ads in slots i and $i + k$ are swapped and let $r_i^k = \Pr(S_i)$ be the probability that this event occurs. We have

$$\Pr(S_i) = \Pr(S_i|S_{i-k}) \Pr(S_{i-k}) + \Pr(S_i|\neg S_{i-k}) \Pr(\neg S_{i-k}).$$

To avoid conflicting swaps, we can set $\Pr(S_i|S_{i-k}) = 0$, which implies that $\Pr(S_i|\neg S_{i-k}) = \Pr(S_i)/\Pr(\neg S_{i-k}) = r_i^k/(1 - r_{i-k}^k)$, which is no greater than one as long as we enforce that $r_{i-1}^k + r_i^k \leq 1$.

For the sake of this algorithm, all ads with rank $K + 1, \dots, N$ can be thought of as sharing slot $K + 1$. Thus whenever an ad in slot $s \leq K$ is chosen to swap with slot $K + 1$, any ad with rank $K + 1, \dots, N$ could be displayed in slot s . Due to lack of space, we do not discuss how the algorithm might decide which losing ad to display, but one could imagine giving preference to ads that have not often been displayed in the past.

4 Maintaining Equilibrium During Pairwise Swapping

In this section, we consider the effect on advertiser incentives of implementing an exploration policy that occasionally chooses pairs of ads that are k slots apart to swap or moves an undisplayed ad into slot $K - k + 1$ for some *fixed* value of k . By ensuring that advertisers do not have incentives to deviate from equilibrium bids for any fixed k , we ensure that the advertisers do not deviate throughout the entire run of **k-swap**.

We assume that the search engine bases the weights w_i on the CTR estimates \hat{c}_i^s , and fix the prices paid by the advertisers through the entire run of **k-swap**. The updated CTR estimates obtained during exploration are only reported to advertisers after the algorithm completes. In practice, the algorithm may need

Algorithm 1. The k -swap algorithm.

```

Calculate all swapping probabilities  $r_i^k$ 
for all queries on the given keyword do
    Randomly select a  $k \in \{1, \dots, K\}$ 
    for  $i = 1$  to  $\min\{k, K - k + 1\}$  do
        Set  $S_i \leftarrow 1$  with probability  $r_i^k$ ,  $S_i \leftarrow 0$  otherwise
    end for
    for  $i = k + 1$  to  $K - k + 1$  do {Note that this statement is null if  $2k > K$ }
        if  $S_{i-k} = 1$  then
            Set  $S_i \leftarrow 0$ 
        else
            Set  $S_i \leftarrow 1$  with probability  $r_i^k / (1 - r_{i-k}^k)$ ,  $S_i \leftarrow 0$  otherwise
        end if
    end for
    for  $i = 1$  to  $K - k$  do
        Swap the ads in slots  $i$  and  $i + k$  if  $S_i = 1$ 
    end for
    if  $S_{K-k+1} = 1$  then
        Choose an  $i \in \{K + 1, \dots, N\}$  to display in slot  $K - k + 1$ 
    end if
end for

```

to be run in multiple phases, interleaving exploration with updates of CTR estimates, and allowing sufficient time for advertisers to reach a new equilibrium after each phase.

Our assumptions raise a conceptual question: if the advertisers care about the *real* CTRs, how can we maintain incentives given only estimates? We posit that often advertisers do not know the CTRs any better than the search engine and formulate their own optimization problem (at least approximately) in terms of the estimates provided by the search engine; that is, we assume that $\tilde{c}_i^s = \hat{c}_i^s \forall i, s$. We consider the case in which advertisers have additional information about their CTRs in Section 6.

For the analysis that follows, we assume that the search engine knows (or can obtain good estimates of) each advertiser's value per click. If we assume that a SNE is played prior to exploration, we can derive bounds on advertiser values [10] and base our estimates on these bounds. In practice, this assumption will not be necessary; we do not actually advocate setting the swapping probabilities separately for each individual auction, but rather fixing probabilities in such a way that the guarantees will hold for most typical auctions.

Since all analysis in this section is for a fixed value of k , we drop the superscript and use r_i in place of r_i^k to denote the probability that ads i and $i + k$ are swapped. These probabilities can be represented as multiples of r_1 , i.e. $r_i = \alpha_i r_1$. Then, if α_i are set exogenously (for example, $\alpha_i = 1$ for all $1 \leq i \leq K$), k -swap has only one tunable parameter, r_1 , for a fixed value of k . For convenience of notation, we define $\alpha_i = 0$ for all $i < 1$ and $i > K - k + 1$. In order to allow exploration of CTRs of all bidders, we let r_{K-k+1} designate the total

probability that *any* losing bidder is swapped into slot $K - k + 1$. Let q_s denote the probability that a losing bidder with rank $K + 1 \leq s \leq N$ is displayed *conditional* on *some* losing ad being displayed.³ We have that $\sum_{s=K+1}^N q_s = 1$. Finally, define $q_{max} = \max_{K+1 \leq s \leq N} q_s$.

Once we add exploration, the *effective* estimate of CTR for advertiser i in slot s is no longer \hat{c}_i^s . Rather, now with some probability r_{s-k} the ad in slot s is moved to slot $s - k$, and with some probability r_s the ad is moved to slot $s + k$. Then the new effective estimate of CTR of player i for rank s is $\hat{c}_i^{s'} = (1 - r_{s-k} - r_s)\hat{c}_i^s + r_{s-k}\hat{c}_i^{s-k} + r_s\hat{c}_i^{s+k}$.⁴

Let $D_{i,s} = \alpha_s(\hat{c}_i^s - \hat{c}_i^{s+k}) - \alpha_{s-k}(\hat{c}_i^{s-k} - \hat{c}_i^s)$. Observe that $r_1 D_{i,s}$ is the marginal CTR loss of advertiser i in slot s when exploration is allowed. We now define the quantities $J_{i,j}$ and Z_i which are used in Theorem 1:

$$J_{i,j} = (v_i - p_i^i)D_{i,i} - (v_i - p_i^j)D_{i,j} \quad (1)$$

$$Z_i = (v_i - p_i^i)D_{i,i} + \alpha_{K-k+1}q_{max}\hat{c}_i^{K-k+1}v_i. \quad (2)$$

To get some intuition about what these mean, note that $r_1 J_{i,j}$ is the difference between the marginal loss in expected payoff due to exploration that the advertiser i receives in slot j and the marginal loss in expected payoff due to exploration in slot i . Similarly, $r_1 Z_i$ is the difference between the marginal loss in payoff due to exploration that the advertiser i receives by switching to rank above $K + 1$ (and thereby not occupying any slot) and the marginal loss due to exploration in slot i .

The following result gives the conditions under which exploration does not incent advertisers to change their bids and characterizes the settings in which this is not possible. The proof of this theorem and others can be found in the appendix of the extended version of this paper.⁵

Theorem 1. *Assume that each advertiser $i \in \{1, \dots, K\}$ strictly prefers his current slot to all others in equilibrium, i.e. the condition $(v_i - p_i^i)\hat{c}_i^i > (v_i - p_i^j)\hat{c}_i^j$ holds for all $1 \leq i, j \leq K, i \neq j$ whenever $J_{i,j} > 0$ and $v_i - p_i^i > 0 \forall i$ whenever $Z_i > 0$. Then for generic valuations and relevances there exists an $r_1 > 0$ such that no advertiser has incentive to deviate from the pre-exploration SNE bids once exploration is added. In particular, any r_1 satisfying the following set of conditions is sufficient:*

$$r_1 \leq \min \left\{ \begin{array}{l} \min_{2 \leq i \leq K} \frac{1}{\alpha_i + \alpha_{i-k}}, \quad \min_{1 \leq i \leq K; Z_i > 0} \frac{1}{Z_i} (v_i - p_i^i)\hat{c}_i^i, \\ \min_{1 \leq i, j \leq K; i \neq j; J_{i,j} > 0} \frac{1}{J_{i,j}} \left((v_i - p_i^i)\hat{c}_i^i - (v_i - p_i^j)\hat{c}_i^j \right) \end{array} \right\}.$$

³ Thus, the probability that a particular losing bidder s gets selected is $q_s r_{K-k+1}$.

⁴ Recall that $r_s = 0$ and $\hat{c}_i^s = 0$ for $s < 1$ and $s > K - k + 1$. We can replace CTR with effective CTR because the prices paid by all advertisers remain fixed for the duration of exploration.

⁵ The extended version is available on the authors' websites.

To get some intuition about how the theorem can be applied and about the magnitude of r_1 , consider the following example.

Example 1. Suppose that there are 3 advertisers bidding on 2 slots. Let $\hat{c}_i^j = \hat{c}_j$ for all players $i \in \{1, 2, 3\}$ and slots $j \in \{1, 2\}$ where $\hat{c}_1 = 1$ and $\hat{c}_2 = 0.5$. Let $v_1 = v_2 = 3$, and $v_3 = 1$. Suppose that prior to exploration each advertiser bids his value per click and pays the next highest bid. One can easily verify that this configuration constitutes a SNE in which player 1 gets slot 1, player 2 gets slot 2, and player 3 gets no slot, and that in this equilibrium, player 1 is indifferent between slots 1 and 2.

Let us fix $\alpha_2 = 3/2$. Now we can determine the setting of r_1 that allows us to swap neighboring ads ($k = 1$) without introducing incentives to deviate during exploration. Applying the first constraint, we find the condition that $r_1 \leq 1/(1 + 3/2) = 2/5$ must hold. By the second constraint, since $Z_1 = 11/4$, we must have $r_1 \leq 4/11$, and since $Z_2 = 7/4$, we must have $r_1 \leq 2/7$. With our setting of α_2 , $J_{1,2} = 0$ and $J_{2,1} = -1/4 < 0$. Consequently, the third constraint on r_1 has no effect. Combining the effects of these constraints, we see that we can set the swapping probabilities as high as $r_1 = 2/7$ and $r_2 = 3/7$ without giving any of the advertisers incentive to deviate during exploration.

Suppose we want to increase r_1 to $2/7 + \epsilon$ and thereby learn a little bit faster. Consider the incentives of the second bidder to switch to rank 3 (i.e., receive no slot). The utility from being ranked third is $3/7 + 3\epsilon/2 > 3/7$, while the utility from remaining in slot two is $3/7 - \epsilon/4 < 3/7$. Consequently, for any $\epsilon > 0$ (and, thus, for any $r_1 > 2/7$) the second bidder wants to deviate from his equilibrium bid.

A similar analysis of constraints and incentives shows that we cannot increase α_2 without decreasing r_1 or altering advertiser incentives. Similarly, any attempt to decrease α_2 can destabilize the equilibrium.

As the example suggests, the bounds in Theorem 1 are close to tight. In fact, the bounds can be made tight simply by replacing q_{max} with the conditional probability with which ad i would be selected if it were not in one of the top K ranks.

Note that we would not expect a search engine to calculate a distinct set of swapping probabilities using Theorem 1 for each individual auction in practice. Indeed it may not be possible for the search engine to estimate advertiser values accurately in all cases. We instead advocate using the theorem to find a single fixed set of swapping probabilities such that advertisers will not wish to deviate when **k-swap** is run for *most* or *all typical* auctions.

5 Learning Bounds

In this section, we bound the error of our estimated click-through rates for each advertiser in each slot after Q queries have been made on the given keyword. Let $n_{i,s}$ denote the number of times we have observed advertiser i in slot s , and let $z_{i,s,j}$ be the indicator random variable which is 1 if ad i is clicked the j th

time it appears in slot s , and 0 otherwise. Finally, let $\pi_{i,s}^k$ be the probability that ad i is displayed at slot s when we are swapping ads that are k slots apart, as discussed in Section 4.

To simplify the presentation of results, we assume that the swapping distance k is drawn uniformly at random from $\{1, \dots, K\}$ for each query, but the extension to arbitrary distributions is straight-forward.

Theorem 2. *Suppose the k -swap algorithm has been run for Q queries with a fixed set of broadcasted CTR estimates. Let \hat{c}_i^s be our new estimate of CTR, defined as $\hat{c}_i^s = (1/n_{i,s}) \sum_{j=1}^{n_{i,s}} z_{i,s,j}$ for all advertisers i and slots s such that $n_{i,s} \geq 1$. Then for any $\delta \in (0, 1)$, with probability $1 - \delta$, the following holds for all i and s for which we have made at least one observation:*

$$|\hat{c}_i^s - c_i^s| \leq \sqrt{\frac{\ln(2KN/\delta)}{2n_{i,s}}}.$$

Furthermore, with probability $1 - \delta$, for all i and s , we have that $n_{i,s} \geq \max\{(Q/K) \sum_{k=1}^K \pi_{i,s}^k - \sqrt{Q \ln(2KN/\delta)}/2, 0\}$.

Thus as the number of queries Q grows, our estimates of the CTR vectors for each advertiser grow arbitrarily close to the true CTR vectors.

6 Bounds on the Incentives of “Omniscient” Advertisers

If players have much more information about the actual click-through rates than the search engine, it is unlikely that we can entirely eliminate incentives of advertisers to change their bids during exploration. However, if we can bound the error in our estimates of the click-through rates, we can also bound how much advertisers can gain by deviating. When incentives to deviate are small, we may reasonably expect advertisers to maintain their equilibrium bids, since computing the new optimal bids may be costly. The search engine may further dull benefits from deviation by charging a small fee to advertisers when they change their bids.

From this point on, we assume that the error in search engine estimates of the CTRs is uniformly bounded by ϵ ; that is, $|c_i^s - \hat{c}_i^s| \leq \epsilon$ for every i and s .

Assume that r_1^k were set such that the bidders have no incentive to change their bids if they use \hat{c}_i^s as their CTR estimates. We now establish how much incentive they have to deviate if they know their *actual* CTR c_i^s , that is, $\tilde{c}_i^s = c_i^s$; we call such advertisers “omniscient”.

Theorem 3. *The most that any omniscient advertiser can gain by deviating in expectation per impression is $\max_{1 \leq i \leq K} 2\epsilon(v_i - p_i^K)$.*

This bound has the intuitive property that as our CTR estimates improve, the bound on incentives to deviate from equilibrium bids improves as well.⁶ It is also

⁶ Note that given r_1^k the actual payoffs to deviation are not affected as we learn unless we also publicize the learned information.

intuitive, however, that incentives diminish if the exploration probabilities fall. This motivates the following alternate bound which shows that we can make the incentives to deviate arbitrarily small even for omniscient advertisers by appropriately setting r_1^k .

Theorem 4. *The most that any omniscient advertiser can gain by deviating in expectation per impression is*

$$\max_{1 \leq i, j, k \leq K} \left\{ r_1^k \left(\alpha_i (\hat{c}_i^i - \hat{c}_i^{i+k}) + \alpha_{j-k} (\hat{c}_i^{j-k} - \hat{c}_i^j) + 2\epsilon(\alpha_i + \alpha_{j-k}) \right) (v_i - p_i^K) \right\}.$$

7 Bounds on Revenue Loss Due to Exploration

We now assume that the advertisers play according to the symmetric Nash equilibrium that was played prior to exploration and, as in the previous section, assume that the errors of the search engine's estimates of CTRs are uniformly bounded by ϵ with high probability. Given these assumptions, the theorem that follows bounds the loss in revenue due entirely to exploration.

Theorem 5. *The maximum expected loss to the search engine revenue per impression due to exploration is bounded by*

$$\max_{1 \leq k \leq K} \left\{ r_1^k \sum_{i=2}^K p_i^i \left(\alpha_i (\hat{c}_i^i - \hat{c}_i^{i+k}) - \alpha_{i-k} (\hat{c}_i^{i-k} - \hat{c}_i^i) + 2\epsilon \right) \right\}.$$

8 Special Cases

In this section we study the problem of exploration while maintaining a pre-exploration symmetric Nash equilibrium in two special cases. In both cases, it is only necessary to swap adjacent pairs of ads in order to learn reasonable estimates of advertiser CTRs.

8.1 Factorable Click-Through Rates

The first special case we consider is the commonly studied setting where $c_i^s = e_i c_s$; that is, CTRs are factored into a product of advertiser relevance and slot-specific factors. Since there are far more data for estimating c_s than e_i , we assume c_s is known and e_i is to be learned for all advertisers. Under these assumptions, using **k-swap** may seem strange; after all, we can learn e_i for all advertisers $i \leq K$ just as well by leaving them in their current slots! The only problem to be addressed then is to learn CTRs of losing bidders. Consequently, if we truly believe that CTRs are factorable, we need only do adjacent-ad swapping ($k = 1$) and can set $r_1 = \dots = r_{K-1} = 0$ and only allow $r_K > 0$. In this case, we need not worry about deviations by advertisers in slots $1, \dots, K-1$ to alternative slots $1, \dots, K-1$, since the effective CTRs for these deviations are unchanged.

Additionally, no advertiser wants to deviate to slot K , since the CTR in this slot is strictly lower than it was before exploration, and no advertiser ranked $K + 1, \dots, N$ wants a higher slot, since their effective CTRs increase. Thus we need only consider the incentives of the advertiser in slot K . It is not difficult to verify that the condition under which exploration does not affect advertiser K 's incentives is

$$r_K \leq \min \left\{ \min_{1 \leq j \leq K-1} \frac{c_K (v_K - p_K^K) - c_j (v_K - p_K^j)}{c_K (v_K - p_K^K)}, \frac{v_K - p_K^K}{v_K (q_{max} + 1) - p_K^K} \right\},$$

and we can find an $r_K > 0$ when $c_K (v_K - p_K^K) > c_j (v_K - p_K^j)$ for $j < K$.

There is, however, another possible scenario in which exploration might be useful under the factorable CTR assumption. Suppose that we initially posit the factorable CTR model, but want to verify whether this is really the case. To do so, we can use adjacent-ad swapping to form multiple estimates of e_i using data from multiple adjacent slots. By comparing these estimates, we can vet our current model while also improving our CTR estimates for losing bidders.

Since CTR is factorable, our analysis need only consider the effective slot-specific CTRs, which we assume are known, $c'_s = (1 - r_{s-1} - r_s)c_s + r_{s-1}c_{s-1} + r_s c_{s+1}$. Set $\alpha_i = \prod_{j=2}^i [(c_{j-1} - c_j)/(c_j - c_{j+1})]$. By setting the swapping probabilities in this manner, the effective CTRs in slots $2, \dots, K-1$ are unchanged when exploration is added. We can now simplify the bounds and characterization of Theorem 1. In particular, the precondition of the theorem and the second bound on r_1 need only to hold for $i = 1$. Furthermore, it can be shown that in the factorable setting, the necessary precondition $(v_1 - p_1^1)c_1 > (v_1 - p_1^j)c_j$ always holds in the minimum revenue SNE [10,8,2] for generic valuations and relevances. Formal statements and proofs of these results are in the appendix of the extended version.

As in the general setting, it is possible to derive learning bounds that show that as the number of observed queries grow, our estimates of the advertiser CTR vectors grow arbitrarily close to the true CTRs with high probability. Here our estimates of CTR are simply $\hat{c}_i^s = (c_s/c_{s_i} n_{i,s_i}) \sum_{j=1}^{n_{i,s_i}} z_{i,s_i,j}$ for all i and s , where $s_i = \arg \max_s c_s \sqrt{n_{i,s}}$. We once again defer the theorem statement and proof to the appendix of the extended version due to lack of space.

8.2 Click-Through Rates with Constant Slot Ratios

In this section, we consider adjacent-ad swapping ($k = 1$) for the case in which for each player i , the click-through rates have constant ratios for adjacent slots. That is, for all i and all $1 \leq s \leq K-1$, we assume that $c_i^{s+1}/c_i^s = \gamma_i \leq 1$ where γ_i is advertiser-dependent and unknown. Let $\hat{\gamma}_i$ denote the search engine estimate of γ_i and suppose as before that advertisers use these as their own estimates. Let $\alpha_j = 1$ for every $j \in \{2, \dots, K-1\}$, so $r_1 = r_2 = \dots = r_{K-1}$. Additionally, let $\alpha_K = \min\{(\hat{\gamma}_i - 1)^2/q_{max}, 1\}$.

As in the previous section, we can considerably simplify the bounds and characterization of Theorem 1 in this special case. In particular, the first and second

bounds on r_1 must hold, but the third bound on r_1 and the precondition need only to hold for $i = 1$ and $i = K$.

We can also prove analogous learning bounds in this setting that show that it is only necessary to explore via adjacent-ad swapping in order to obtain CTR estimates for all advertisers at all slots. This can be accomplished by estimating γ_i for each i as

$$\hat{\gamma}_i = \frac{(1/n_{i,s_i+1}) \sum_{j=1}^{n_{i,s_i+1}} z_{i,s_i+1,j}}{(1/n_{i,s_i}) \sum_{j=1}^{n_{i,s_i}} z_{i,s_i,j}}$$

for a chosen slot s_i at which there is a sufficient amount of data available. The CTR at each slot is then estimated using $\hat{\gamma}_i$ and the estimate of the CTR at the designated slot s_i .

Formal theorems describing the conditions on r_1 necessary to maintain equilibrium in this setting and the corresponding learning bounds can be found in the appendix of the extended version, along with their proofs.

9 Conclusion

We have introduced an exploration scheme which allows search engines to learn click-through rates for advertisements. We showed how, when possible, to set the exploration parameters in order to eliminate the incentives for advertisers to deviate from a pre-exploration symmetric Nash equilibrium. In situations in which we cannot entirely eliminate incentives to change bids, we can make returns to changing bids arbitrarily small. Particularly, we can make these small enough to ensure that bid manipulation is hardly worth advertisers' time. Finally, we derived a bound on worst-case expected per-impression revenue loss due to exploration. Since this loss is zero in the limit of no exploration, we can set exploration parameters in order to make it arbitrarily small, while still ensuring that we eventually learn click-through rates.

Acknowledgements

The authors thank Alexander Strehl and Yiling Chen for insightful comments on early drafts of this paper.

References

1. Börgers, T., Cox, I., Pesendorfer, M., Petricek, V.: Equilibrium bids in auctions of sponsored links: theory and evidence. Technical report, University of Michigan (2007)
2. Cary, M., Das, A., Edelman, B., Giotis, I., Heimerl, K., Karlin, A.R., Mathieu, C., Schwarz, M.: Greedy bidding strategies for keyword auctions. In: Eighth ACM Conference on Electronic Commerce (2007)
3. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review* 9(1), 242–259 (2007)

4. Even-Dar, E., Kearns, M., Wortman, J.: Sponsored search with contexts. In: 3rd International Workshop on Internet and Network Economics (2007)
5. Gonen, R., Pavlov, E.: An incentive-compatible multi-armed bandit mechanism. In: 26th Annual ACM Symp. on Principles of Distributed Computing (2007)
6. Heckman, J.: Sample selection bias as a specification error. *Econometrica* 47, 153–161 (1979)
7. Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: ACM Conference on Electronic Commerce, pp. 218–227 (2006)
8. Lahaie, S., Pennock, D.M.: Revenue analysis of a family of ranking rules for keyword auctions. In: ACM Conference on Electronic Commerce (2007)
9. Pandey, S., Olston, C.: Handling advertisements of unknown quality in search advertising. In: Neural Information and Processing Systems (2006)
10. Varian, H.: Position auctions. *Int'l Journal of Industrial Organization* (to appear)