Maintaining Equilibria During Exploration in Sponsored Search Auctions

Jennifer Wortman[1], Yevgeniy Vorobeychik[2],
Lihong Li[3], and John Langford[4]

[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.[5] 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.[6] 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,* $\tilde{c}_i^i \left(v_i - p_i^i\right) \geq \tilde{c}_i^s \left(v_i - p_i^s\right)$, *where* $\tilde{c}_i^s$ *denotes advertiser i's CTR estimate at slot s.*

[5] 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]

[6] 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.

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.

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 $k \in \{1, \cdots, K\}$ 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 $K+1, \cdots, N$ 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 $K+1, \cdots, N$ 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

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 ∈ {1, ··· , K}
  for i = 1 to min{k, K − k + 1} do
    Set S_i ← 1 with probability r_i^k, S_i ← 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 ← 0
    else
      Set S_i ← 1 with probability r_i^k/(1 − r_{i−k}^k), S_i ← 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 ∈ {K + 1, ··· , N} to display in slot K − k + 1
  end if
end for
```

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 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 be 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 $1 \le i \le K$), `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 $K+1 \le s \le N$ is displayed *conditional* on *some* losing ad being displayed.[7] We have that $\sum_{s=K+1}^{N} q_s = 1$. Finally, define $q_{max} = \max_{K+1 \le s \le N} q_s$.

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 $\hat{c}_i^{\prime s} = (1 - r_{s-k} - r_s)\hat{c}_i^s + r_{s-k}\hat{c}_i^{s-k} + r_s\hat{c}_i^{s+k}$.[8]

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:

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

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

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.

Theorem 1. *Assume that each advertiser $i \in \{1, \cdots, K\}$ 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 $1 \le i, j \le K, i \ne j$ 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*

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

[8] 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.

conditions is sufficient:

$$r_1 \le \min \left\{ \min_{2\le i\le K} \frac{1}{\alpha_i + \alpha_{i-k}}, \qquad \min_{1\le i\le K; Z_i>0} \frac{1}{Z_i}(v_i - p_i^i)\hat{c}_i^i, \right.$$
$$\left. \min_{1\le i,j\le K; i\ne j; J_{i,j}>0} \frac{1}{J_{i,j}} \left((v_i - p_i^i)\hat{c}_i^i - (v_i - p_i^j)\hat{c}_i^j\right) \right\}.$$

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 $r_1 \le 1/(1 + 3/2) = 2/5$ must hold. By the second constraint, since $Z_1 = 11/4$, we must have $r_1 \le 4/11$, and since $Z_2 = 7/4$, we must have $r_1 \le 2/7$. 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 jth 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 $\{1, \cdots, K\}$ 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* $\max_{1\le i\le K} 2\epsilon(v_i - p_i^K)$.

This bound has the intuitive property that as our CTR estimates improve, the bound on incentives to deviate from equilibrium bids improves as well.[9] It is also 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*

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

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*

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

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.

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

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 $r_1 = \cdots = r_{K-1} = 0$ and only allow $r_K > 0$. In this case, we need not worry about deviations by advertisers in slots $1, \ldots, K-1$ to alternative slots $1, \ldots, K-1$, 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 $K+1, \ldots, N$ 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

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

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 $2, \cdots, K-1$ 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.

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.

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 $j \in \{2, \ldots, K-1\}$, so $r_1 = r_2 = \cdots = r_{K-1}$. 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 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. Tilman Börgers, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. Equilibrium bids in auctions of sponsored links: theory and evidence. Technical report, University of Michigan, 2007.
2. Matthew Cary, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. Greedy bidding strategies for keyword auctions. In *Eighth ACM Conference on Electronic Commerce*, 2007.
3. Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 9(1):242–259, March 2007.
4. Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Sponsored search with contexts. In *3rd International Workshop on Internet and Network Economics*, 2007.
5. Rica Gonen and Elan Pavlov. An incentive-compatible multi-armed bandit mechanism. In *26th Annual ACM Symp. on Principles of Distributed Computing*, 2007.
6. James Heckman. Sample selection bias as a specification error. *Econometrica*, 47:153–161, 1979.
7. Sebastien Lahaie. An analysis of alternative slot auction designs for sponsored search. In *ACM Conference on Electronic Commerce*, pages 218–227, 2006.
8. Sebastien Lahaie and David M. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *ACM Conference on Electronic Commerce*, 2007.
9. Sandeep Pandey and Christopher Olston. Handling advertisements of unknown quality in search advertising. In *Neural Information and Processing Systems*, 2006.
10. Hal Varian. Position auctions. *Int'l Journal of Industrial Organization*, to appear.

Appendix

A Proofs

A.1 Proof of Existence of SNE

In the existence proof of SNE by Börgers et al. [1], define the mapping $v'_i = w_i v_i$ and $b'_i = w_i b_i$ for every player i. Then, if v'_i are used as valuations and b'_i as corresponding bids, Walrasian equilibrium prices p'_s of each slot exist in our setting by the same argument. Since for each advertiser i, v_i is independent of slot s and $\tilde{c}_i^s > \tilde{c}_i^{s+1}$, $p'_1 \geq p'_2 \geq \cdots \geq p'_K$, also by the same argument as in [1]. Now, index each bidder by the slot they won and define the bids of players $2 \leq i \leq K$ to be the price paid for slot $i-1$: $b'_i = p'_{i-1}$. Define the bids b'_i of advertisers who do not receive a slot to be p'_K. The first bidder can bid anything above p'_1. By the argument presented in [1], this configuration constitutes a symmetric Nash equilibrium.

The inequalities $\tilde{c}_i^i(v'_i - p'_i) \geq \tilde{c}_i^s(v'_i - p'_s)$ characterize SNE in the above construction. Dividing both sides by w_i does not change the direction of inequality. Consequently, if we define $p_i^s = p_s/w_i$, we recover the inequalities $\tilde{c}_i^i(v_i - p_i^i) \geq \tilde{c}_i^s(v_i - p_i^s)$ that characterize SNE in our setting.

A.2 Proof of Theorem 1

In order to prove Theorem 1, we must make sure that when the conditions stated in the theorem hold, no advertiser is happier changing his bid and moving to a different position in the ranking. We break this proof into multiple parts, summarized by the following series of lemmas. Combining the pieces yields a set of conditions that show that no advertiser has incentive to deviate when exploration is added.

The first condition,

$$r_1 \leq \min_{2 \leq i \leq K} \frac{1}{\alpha_i + \alpha_{i-k}},$$

ensures that $r_i + r_{i-k} \leq 1$ for all i. The other two conditions are the subject of Lemmas 1 and 2 below.

Incentives of $1, \cdots, K$ to Switch to Alternate Ranks in $1, \cdots, K$ We first verify that players $i \in \{1, \cdots, K\}$ do not want to switch to alternate ranks $j \in \{1, \cdots, K\}$. The following lemma gives conditions to guarantee this. Recall that we define $c_1^0 = 0$ and $\alpha_0 = 0$.

Lemma 1. *Assume that the condition*

$$(v_i - p_i^i)\hat{c}_i^i > (v_i - p_i^j)\hat{c}_i^j,$$

holds for every $i, j \in \{1, \cdots, K\}$ in equilibrium. Suppose that prior to exploration advertisers' bids are in a symmetric Nash equilibrium. Then there is $r_1 > 0$ such that players in slots $1, \cdots, K$ do not wish to switch to other slots in $1, \cdots, K$ as long as

$$r_1 \leq \min_{1 \leq i,j \leq K; J_{i,j} > 0} \frac{1}{J_{i,j}} \left((v_i - p_i^i)\hat{c}_i^i - (v_i - p_i^j)\hat{c}_i^j \right)$$

where $J_{i,j}$ is as defined in Equation 1.

Proof. To ensure that players in slots $1, \cdots, K$ do not have incentive to switch to other slots within this range, we need the condition

$$\begin{aligned} &\left((1 - r_{i-k} - r_i)\hat{c}_i^i + r_{i-k}\hat{c}_i^{i-k} + r_i\hat{c}_i^{i+k}\right)(v_i - p_i^i) \\ &\geq \left((1 - r_{j-k} - r_j)\hat{c}_i^j + r_{j-k}\hat{c}_i^{j-k} + r_j\hat{c}_i^{j+k}\right)(v_i - p_i^j) \end{aligned}$$

to be satisfied for any $i, j \in \{1, \cdots, K\}$. Using the definition $r_i = \alpha_i r_1$ and rearranging terms, we obtain the equivalent condition

$$r_1 J_{i,j} \leq (v_i - p_i^i)\hat{c}_i^i - (v_i - p_i^j)\hat{c}_i^j. \tag{3}$$

Under our assumption, the right hand side of Equation 3 is always strictly positive. When $J_{i,j} \leq 0$, any $r_1 \in [0, 1]$ satisfies the condition we need. When $J_{i,j} > 0$, we need to guarantee

$$r_1 \leq \frac{(v_i - p_i^i)\hat{c}_i^i - (v_i - p_i^j)\hat{c}_i^j}{J_{i,j}}.$$

Therefore, we must have for all $i, j \in \{1, \cdots, K\}$ such that $J_{i,j} > 0$

$$r_1 \leq \frac{1}{J_{i,j}} \left((v_i - p_i^i)\hat{c}_i^i - (v_i - p_i^j)\hat{c}_i^j \right).$$

Incentive of Players $1, \ldots, K$ to Move to Ranks $K + 1, \ldots, N$ The previous lemma showed that the K highest ranked advertisers do not have incentive to deviate to another rank between 1 and K when exploration is added under certain conditions. We now consider the conditions that guarantee that these advertisers do not want to move to a rank greater than K.

Lemma 2. *If the players are in a symmetric Nash equilibrium with $v_i > p_i^i$ whenever $Z_i > 0$ before exploration and if*

$$r_1 \leq \min_{1 \leq i \leq K; Z_i > 0} \frac{1}{Z_i} \left(\hat{c}_i^i (v_i - p_i^i) \right)$$

where Z_i is defined as in Equation 2, then the K highest ranking advertisers have no incentive to switch to any slot below K for generic valuations and relevances. Furthermore, it is always possible to set r_1 in such a way such that the above holds and $r_1 > 0$.

Proof. We show that any player $i \in \{1, \cdots, K\}$ has no incentive to deviate to any slot $j \in \{K+1, \cdots, N\}$. For any $j \in \{K+1, \cdots, N\}$, in order to guarantee that player i does not want to switch to slot j, the following condition must hold:

$$\left((1 - r_{i-k} - r_i)\hat{c}_i^i + r_{i-k}\hat{c}_i^{i-k} + r_i \hat{c}_i^{i+k} \right)(v_i - p_i^i) \geq r_K q_j \hat{c}_i^{K-k+1} v_i.$$

As we require that this condition hold simultaneously for all j, it becomes equivalent to

$$r_1 Z_i \leq \hat{c}_i^i (v_i - p_i^i) \tag{4}$$

Since $v_i - p_i^i > 0$, the right hand side of Equation 4 is strictly positive. When $Z_i \leq 0$, the constraint is satisfied trivially. Thus we need only that when $Z_i > 0$,

$$r_1 \leq \frac{1}{Z_i} \left(\hat{c}_i^i (v_i - p_i^i) \right)$$

Incentives of Players $K + 1, \ldots, N$ to Move to Slots $1, \ldots, K$

Lemma 3. *No advertiser in slots $K+1, \cdots, N$ prefers to deviate and move to any slot in $1, \cdots, K$ when exploration is added.*

Proof. Since the set of advertisers $K + 1, \cdots, N$ received no clicks in the pre-exploration symmetric Nash equilibrium, we know that it must be the case that for all i and j such that $K + 1 \leq i \leq N$ and $1 \leq j \leq K$,

$$v_i - p_i^j \leq 0.$$

In other words, these advertisers have a value per click that is lower than the equilibrium price per click of any slot between 1 and K. When exploration is added, the CTRs of these slots might change, but the price is still higher than these advertisers' values. Thus the advertisers still are not interested in these slots when exploration is added.

Note that we have not addressed the incentives for players $K+1, \cdots, N$ to deviate to other ranks $K+1$ or higher. This case will depend on the way in which the distribution q_i is chosen. There are many possible choices for which players $K+1, \cdots, N$ will not have incentive to swap ranks among themselves.

A.3 Proof of Theorem 2

The proof of this theorem relies on several applications of a version of Hoeffding's inequality which is stated here for completeness.

Lemma 4. *(Hoeffding's Inequality) Let $x_1, \cdots, x_n$ be independent random variables bounded in the range $[0,1]$. Let $\hat{x} = (1/n)\sum_{i=1}^{n} x_i$ denote the mean of these variables. Then for any $\delta \in (0,1)$, with probability $1-\delta$,*

$$|\hat{x} - E[\hat{x}]| \leq \sqrt{\frac{\ln(2/\delta)}{2n}}.$$

We apply Hoeffding's inequality first to bound the error of our estimation of the click-through rate of ad i in slot s given that we have observed the ad in this slot $n_{i,s}$ times. The following holds with probability $1-\delta$.

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

To prove the final piece of the theorem, we can apply Hoeffding's inequality once again to bound the deviation of $n_{i,s}$ from its expectation with probably $1-\delta$ as follows.

$$|n_{i,s} - E[n_{i,s}]| \leq \sqrt{Q\ln(2/\delta)/2}$$

Dividing δ by NK and applying the union bound completes the proof.

A.4 Proof of Theorem 3

First, we need a useful lemma that bounds the effective CTRs.

Lemma 5. $|c_i^{\prime s} - \hat{c}_i^{\prime s}| \leq \epsilon$.

Proof.

$$\begin{aligned}\left|c_i^{\prime s} - \hat{c}_i^{\prime s}\right| &= \left|(1-r_{s-k}-r_s)(c_i^s - \hat{c}_i^s) + r_{s-k}(c_i^{s-k} - \hat{c}_i^{s-k}) + r_s(c_i^{s+k} - \hat{c}_i^{s+k})\right| \\ &\leq (1-r_{s-k}-r_s)\delta + r_{s-k}\delta + r_s\delta \;=\; \delta.\end{aligned}$$

Now consider a player i deviating to slot (or rank) j:

$$\begin{aligned}
&c_i^{'i}(v_i - p_i^i) \geq \hat{c}_i^{'i}(v_i - p_i^i) - \epsilon(v_i - p_i^i) \\
&\geq \hat{c}_i^{'j}(v_i - p_i^j) - \epsilon(v_i - p_i^i) \geq c_i^{'j}(v_i - p_i^j) - \epsilon(v_i - p_i^i) - \epsilon(v_i - p_i^j) \\
&\geq c_i^{'j}(v_i - p_i^j) - 2\epsilon(v_i - p_i^K).
\end{aligned}$$

Recall that if $j \geq K+1$, $c_i^j = 0$.

A.5 Proof of Theorem 4

Consider incentives for some player i for deviating to slot j and let

$$\mu = r_1(\alpha_i(\hat{c}_i^i - \hat{c}_i^{i+k}) + \alpha_{j-k}(\hat{c}_i^{j-k} - \hat{c}_i^j) + 2\epsilon(\alpha_i + \alpha_{j-k}))(v_i - p_i^K).$$

Then,

$$\begin{aligned}
&[(1 - r_{i-k} - r_i)c_i^i + r_{i-k}c_i^{i-k} + r_i c_i^{i+k}](v_i - p_i^i) + \mu \\
&\geq [(1 - r_{i-k} - r_i)c_i^i + r_{i-k}c_i^{i-k} + r_i c_i^{i+k}](v_i - p_i^i) + \\
&\quad\ r_i(c_i^i - c_i^{i+k})(v_i - p_i^K) + r_{j-k}(c_i^{j-k} - c_i^j)(v_i - p_i^K) \\
&\geq [(1 - r_{i-k} - r_i)c_i^i(v_i - p_i^i) + r_{i-k}c_i^{i-k}(v_i - p_i^i) + \\
&\quad\ r_i c_i^i(v_i - p_i^i) + r_{j-k}(c_i^{j-k} - c_i^j)(v_i - p_i^j) \\
&= c_i^i(v_i - p_i^i) + r_{i-k}(c_i^{i-k} - c_i^i)(v_i - p_i^i) + r_{j-k}(c_i^{j-k} - c_i^j)(v_i - p_i^j) \\
&\geq c_i^j(v_i - p_i^j) + r_{j-k}(c_i^{j-k} - c_i^j)(v_i - p_i^j) \\
&\geq c_i^j(v_i - p_i^j) + r_{j-k}(c_i^{j-k} - c_i^j)(v_i - p_i^j) - r_j(c_i^j - c_i^{j+k})(v_i - p_i^j) \\
&= [(1 - r_{j-k} - r_j)c_i^j + r_{j-k}c_i^{j-k} + r_j c_i^{j+k}](v_i - p_i^j),
\end{aligned}$$

where the inequalities follow from the assumption that $c_i^i(v_i - p_i^i) \geq c_i^j(v_i - p_i^j)$ and the fact that $p_i^s \geq p_i^K$ for any slot s.

A.6 Proof of Theorem 5

$$\begin{aligned}
\Delta_{Revenue} &= \sum_{i=1}^{K} c_i^i p_i^i - \sum_{i=1}^{K} c_i^{'i} p_i^i = \sum_{i=1}^{K} p_i^i(c_i^i - c_i^{'i}) \\
&= \sum_{i=1}^{K} p_i^i(r_i(\hat{c}_i^i - \hat{c}_i^{i+k}) - r_{i-k}(\hat{c}_i^{i-k} - \hat{c}_i^i) + 2\epsilon) \\
&= r_1 \sum_{i=1}^{K} p_i^i \left(\alpha_i(\hat{c}_i^i - \hat{c}_i^{i+k}) - \alpha_{i-k}(\hat{c}_i^{i-k} - \hat{c}_i^i) + 2\epsilon\right).
\end{aligned}$$

A.7 Theorem Statement and Proof for the Factorable Case

Theorem 6 below is mentioned briefly in Section 8.1. For completeness, we state it formally here and provide its proof. Before diving into the proof of the main theorem, note that by defining α_i to be

$$\alpha_i = \prod_{j=2}^{i} \left(\frac{c_{j-1} - c_j}{c_j - c_{j+1}} \right),$$

we obtain the following recursive expression for r_i:

$$r_i = r_{i-1} \left(\frac{c_{i-1} - c_i}{c_i - c_{k+1}} \right). \tag{5}$$

for all $i \in \{2, \cdots, K\}$. This setting of probabilities is convenient because the CTRs of an ad in slots $2, \cdots, K$ do not change when exploration is added if bids do not change, as shown in the following useful lemma .

Lemma 6. *When r_i are computed recursively by (5), then $c'_i = c_i$ for all $i \in \{2, \cdots, K\}$.*

Proof. For any $i \in \{2, \cdots, K\}$,

$$\begin{aligned} c'_i &= (1 - r_{i-1} - r_i)c_i + r_{i-1}c_{i-1} + r_i c_{i+1} \\ &= c_i + r_{i-1}(c_{i-1} - c_i) - r_i(c_i - c_{i+1}) \\ &= c_i + r_{i-1}(c_{i-1} - c_i) - r_{i-1} \left(\frac{c_{i-1} - c_i}{c_i - c_{i+1}} \right) (c_i - c_{i+1}) = c_i. \end{aligned}$$

Now we are ready to state and prove the main theorem.

Theorem 6. *Consider the setting in which CTRs are factorable into the product of advertiser relevance and a slot-specific CTR factor. Let r_i be defined as in Equation 5 for all $i \in \{2, \cdots, K\}$. Assume that advertiser 1 strictly prefers his current slot to all others in equilibrium, i.e. the condition $\left(v_1 - p_1^1\right) c_1 > \left(v_1 - p_1^j\right) c_j$ holds for all $2 \leq j \leq K$ whenever $J_{1,j} > 0$. Then for generic valuations and relevances there exists an $r_1 > 0$ such that no advertiser has an incentive to deviate from the pre-exploration SNE bids once exploration is added. Any r_1 for which the following conditions hold is sufficient:*

$$r_1 \leq \min \left\{ \min_{2 \leq i \leq K} \frac{1}{\alpha_i + \alpha_{i-1}}, \quad \min_{1 \leq i \leq K; Z_i > 0} \frac{1}{Z_i}(v_i - p_i^i)c_i^i \right.$$
$$\left. \min_{2 \leq j \leq K; J_{1,j} > 0} \frac{1}{J_{1,j}}((v_1 - p_1^1)c_1 - (v_1 - p_1^j)c_j). \right\}$$

where

$$J_{1,j} = \left(v_1 - p_1^1\right)(c_1 - c_2) + \left(v_1 - p_1^j\right)(\alpha_{j-1}(c_{j-1} - c_j) - \alpha_j(c_j - c_{j+1}))$$

and Z_i is defined according to Equation 2.

First, we demonstrate a simple and intuitive result that $w_i v_i > w_{i+1} b_{i+1}$ for $1 \le i \le K$ for generic valuations and relevances.

Lemma 7. *In a symmetric Nash equilibrium with $w_i v_i \ge w_{i+1} v_{i+1}$, $w_i v_i > w_{i+1} b_{i+1}$ for $1 \le i \le K$ for generic valuations and relevances.*

Proof. Suppose $w_i v_i \le w_{i+1} b_{i+1}$ for some $i \in \{1, 2, \cdots, K\}$.

For $j > 1$, the following upper bound on $w_j b_j$ holds in a symmetric Nash equilibrium (a simple extension of Varian [10, 8]):

$$w_j b_j \le w_{j-1} v_{j-1} (1 - \beta_j) + w_{j+1} b_{j+1} \beta_j \tag{6}$$

where $\beta_j = c_j / c_{j-1} < 1$. From the above bound with $j = i + 1$ and the assumption that $w_i v_i \le w_{i+1} b_{i+1}$, it follows that $w_i v_i \le w_i v_i (1 - \beta_{i+1}) + w_{i+2} b_{i+2} \beta_{i+1}$ and, therefore $w_i v_i \le w_{i+2} b_{i+2}$.

Applying the bound in Equation 6 again with $j = i + 2$, we see that $w_{i+2} b_{i+2} \le w_{i+1} v_{i+1} (1 - \beta_{i+2}) + w_{i+3} b_{i+3} \beta_{i+2}$. When CTRs are factorable, $w_i v_i \ge w_{i+1} v_{i+1}$ [1, 3, 10]. Since for generic values and relevances, $w_{i+1} v_{i+1} < w_i v_i$, we have $w_i v_i \beta_{i+2} < w_{i+3} b_{i+3} \beta_{i+2}$ and, consequently, $w_i v_i < w_{i+3} b_{i+3} \le w_{i+1} b_{i+1}$. Thus, $w_i v_i < w_{i+1} b_{i+1}$ and $v_i - \frac{w_{i+1} b_{i+1}}{w_i} < 0$. But this is a contradiction, since player i would then want to switch to slot $K+1$, whereas we assumed that all bidders were in a Nash equilibrium.

Thus, $v_i > p_i^i$ for all players i when CTRs are factorable.

Given Lemma 6, note that if none of the bidders in slots $2, \cdots, N$ wanted to move up to slot 1 in equilibrium before exploration, they have even less incentive to do so once exploration is added since the effective CTR for slot 1 is now lower and the effective CTR of their own slots is the same by our definition of α_i. Furthermore, none of the bidders in slots $2, \cdots, N$ want to to switch to alternate slots in $2, \cdots, K$ since the effective CTRs are now the same for all of these slots and do not depend on the bidder's identity due to the factorization assumption. Consequently, we need only examine whether or not the top bidder wants to move down, or whether any bidder might like to move into a slot below K. The analysis of these cases is directly analogous to the analysis in the proof of Theorem 1, and the sufficient conditions on r_1 are derived in the same manner.

A.8 Proof of Player 1's Strict Preference at SNE

Below is a formal statement and proof of the theorem mentioned in Section 8.1.

Theorem 7. *Suppose that the players are playing a minimum symmetric Nash equilibrium. Then for generic valuations and relevances $(v_1 - p_1^1) c_1 > (v_1 - p_1^j) c_j$.*

$$\begin{aligned} c_1 w_2 b_2 - c_j w_{j+1} b_{j+1} &= \sum_{t=1}^{K} (c_t - c_{t+1}) w_{t+1} v_{t+1} - \sum_{t=j}^{K} (c_t - c_{t+1}) w_{t+1} v_{t+1} \\ &= \sum_{t=1}^{j-1} (c_t - c_{t+1}) w_{t+1} v_{t+1} \\ &\leq w_2 v_2 \sum_{t=1}^{j-1} (c_t - c_{t+1}) = w_2 v_2 (c_1 - c_j). \end{aligned}$$

For generic valuations and relevances, $w_2 v_2 (c_1 - c_s) < w_1 v_1 (c_1 - c_s)$ and, consequently, $c_1 w_2 b_2 - c_j w_{j+1} b_{j+1} < w_1 v_1 c_1 - w_1 v_1 c_j$ for every $2 \leq j \leq K$. Rewriting, we get $c_1 (w_1 v_1 - w_2 b_2) > c_j (w_1 v_1 - w_{j+1} b_{j+1})$ and we recover the desired strict inequality.

A.9 Learning Bounds in the Factorable Case

The proof of the following theorem, which is mentioned in Section 8.1, is given below. As before, let $n_{i,s}$ denote the number of times we have observed advertiser i in slot s at the current fixed CTR estimates, and let $z_{i,s,j}$ be a random variable indicating whether or not the ad i was clicked on the jth time it appeared in slot s.

Theorem 8. *Suppose that CTRs can be factored into advertiser-dependent and slot-dependent components. In other words, for all i and s, $c_i^s = e_i c_s$ where c_s is known. Suppose we have observed $n_{i,s}$ instances of ad i at slot s with a fixed set of broadcasted CTR estimates. Let $\hat{c}_i^s$ be our* new *estimate of CTR, defined as:*

$$\hat{c}_i^s = \frac{c_s}{c_{s_i} n_{i,s_i}} \sum_{j=1}^{n_{i,s_i}} z_{i,s_i,j}$$

for all advertisers i and slots s; let $s_i = \arg\max_s c_s \sqrt{n_{i,s}}$. Then for any $\delta \in (0,1)$, with probability $1 - \delta$, the following holds for all i and s:

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

Now, for each advertiser i, we base our estimate of click-through rate on data from the slot s_i maximizing $c_s \sqrt{n_s}$. Let

$$\hat{c}_i^{s_i} = \frac{1}{n_{i,s_i}} \sum_{j=1}^{n_{i,s_i}} z_{i,s_i,j}$$

be the estimate of click-through rate in this slot. By Hoeffding's inequality, for any $\delta' \in (0,1)$, with probability $1-\delta'$,

$$|\hat{c}_i^{s_i} - c_i^{s_i}| \le \sqrt{\frac{\ln(2/\delta')}{2n_{i,s_i}}}.$$

Since for any s, $c_i^s = c_i^{s-i}(c_s/c_{s_i})$ and $\hat{c}_i^s = \hat{c}_i^{s-i}(c_s/c_{s_i})$, we thus have for all s

$$|\hat{c}_i^s - c_i^s| \le \frac{c_s}{c_{s_i}}\sqrt{\frac{\ln(2/\delta')}{2n_{i,s_i}}}.$$

We want this claim to hold for all N advertisers. Setting $\delta' = \delta/N$ and applying the union bound completes the proof.

A.10 Theorem Statement and Proof for Constant Slot Ratios

The following Theorem is mentioned in Section 8.2. For completeness, the theorem and proof are formally stated here.

Theorem 9. *Suppose that CTRs are of the form $c_i^s = e_i(\hat{\gamma}_i)^{s-1}$ for all i and s. Assume that advertisers $i = 1, \cdots, K$ strictly prefer their current slots to all others in equilibrium, i.e. the following condition holds for $i = 1, K$ and all $j \in \{1, \cdots, K\}, j \ne i$ whenever $J_{i,j} > 0$:*

$$(v_i - p_i^i)\hat{c}_i^i > (v_i - p_i^j)\hat{c}_i^j.$$

Furthermore, assume that for all i such that $Z_i > 0$, $v_i - p_i^i > 0$. Then for generic valuations and relevances there exists $r_1 > 0$ such that no advertiser has an incentive to deviate from the pre-exploration symmetric Nash equilibrium bids once exploration is added. Any r_1 satisfying the following set of conditions is sufficient:

$$r_1 \le \min\left\{ \frac{1}{2}, \quad \min_{i=1,K;Z_i>0} \frac{\hat{c}_i^1}{Z_i}(v_i - p_i^i)(\hat{\gamma}_i)^{i-1}, \right.$$
$$\left. \min_{i=1,K;1\le j\le K;J_{i,j}>0} \frac{\hat{c}_i^1}{J_{i,j}}\left((v_i - p_i^i)(\hat{\gamma}_i)^{i-1} - (v_i - p_i^j)(\hat{\gamma}_i)^{j-1}\right) \right\}$$

The incentives for player 1 and players $K, K+1, \cdots, N$ can be analyzed exactly as in the proof of Theorem 1, and the conditions on r_1 follow from this analysis as before. Thus here we focus only on the incentives of players $2, \cdots, K-1$.

Incentives of Players $2, \ldots, K-1$ to Move to Slots $1, \ldots, K$

Lemma 8. *In the constant ratio setting, when $r_i = r_1$ for $i \in 2, \cdots, K$, players $2, \ldots, K-1$ do not have incentive to deviate to other slots $1, \ldots, K$ during adjacent-ad swapping exploration for any r_1.*

Proof. To ensure that players in slots $2, \cdots, K-1$ do not have incentive to switch to other slots within this range, we need the condition

$$\begin{aligned}&\left((1 - r_{i-1} - r_i)\hat{c}_i^i + r_{i-1}\hat{c}_i^{i-1} + r_i\hat{c}_i^{i+1}\right)(v_i - p_i^i) \\ &\geq \left((1 - r_{j-1} - r_j)\hat{c}_i^j + r_{j-1}\hat{c}_i^{j-1} + r_j\hat{c}_i^{j+1}\right)(v_i - p_i^j)\end{aligned}$$

to be satisfied for any $i, j \in \{1, \cdots, K\}$. Setting $r_i = r_1$ for all i and plugging in $\hat{c}_i^{s-1} = \hat{c}_i^s/\hat{\gamma}_i$ and $\hat{c}_i^{s+1} = \hat{c}_i^s\hat{\gamma}_i$, we get

$$\begin{aligned}(1 - r_{s-1} - r_s)\hat{c}_i^s + r_{s-1}\hat{c}_i^{s-1} + r_s\hat{c}_i^{s+1} &= \hat{c}_i^s + \left(\hat{\gamma}_i - 2 + \frac{1}{\hat{\gamma}_i}\right) r_1\hat{c}_i^s \\ &= \left(1 + \frac{r_1(\hat{\gamma}_i - 1)^2}{\hat{\gamma}_i}\right)\hat{c}_i^s\end{aligned}$$

for every player i and for all $s = 2, \ldots, K-1$. Thus, the condition that needs to be satisfied is

$$\hat{c}_i^i\left(1 + \frac{r_1(\hat{\gamma}_i - 1)^2}{\hat{\gamma}_i}\right)(v_i - p_i^i) \geq \hat{c}_i^j\left(1 + \frac{r_1(\hat{\gamma}_i - 1)^2}{\hat{\gamma}_i}\right)(v_i - p_i^j)$$

which is equivalent to the symmetric equilibrium condition before exploration and thus holds by assumption. For deviation to slot 1, note that $(1 + \frac{r_1(\hat{\gamma}_i - 1)^2}{\hat{\gamma}_i}) \geq 1$ and thus $\hat{c}_i^{'i} \geq \hat{c}_i^i$ for every $i = 2, \ldots, K-1$, whereas $\hat{c}_1^{'1} \leq \hat{c}_1^1$. Consequently, if deviations to 1 were unprofitable prior to exploration, they are certainly still unprofitable with exploration. Finally, for deviations to slot K, note that the effective CTR for slot K is $\hat{c}_i^K(1 + r_1(\hat{\gamma}_i - 2)) \leq \hat{c}_i^K(1 + r_1(\hat{\gamma}_i + 1/\hat{\gamma}_i - 2))$. Thus,

$$\begin{aligned}\hat{c}_i^i\left(1 + \frac{r_1(\hat{\gamma}_i - 1)^2}{\hat{\gamma}_i}\right)(v_i - p_i^i) &\geq \hat{c}_i^K\left(1 + \frac{r_1(\hat{\gamma}_i - 1)^2}{\hat{\gamma}_i}\right)(v_i - p_i^j) \\ &\geq \hat{c}_i^K\left(1 + r_1(\hat{\gamma}_i - 2)\right)(v_i - p_i^j)\end{aligned}$$

Incentives of Players $2, \ldots, K-1$ to Move to Ranks $K+1, \ldots, N$

Lemma 9. *In the constant ratio setting, when $r_i = r_1$ for $i \in 2, \cdots, K$, players in slots $2, \ldots, K-1$ have no incentive to move to ranks $K+1, \ldots, N$ during adjacent-ad swapping exploration for any r_1.*

Proof. Since $\alpha_K = \min\{\frac{(\hat{\gamma}_i - 1)^2}{q_{max}}, 1\}$, we have that $\alpha_K q_{max} \leq (\hat{\gamma}_i - 1)^2$ or $\frac{(\hat{\gamma}_i - 1)^2}{\hat{\gamma}_i} \geq \alpha_K q_{max}(\frac{1}{\hat{\gamma}_i})^{K-i}$.

In order to eliminate incentives to deviate to slots $K+1,\ldots,N$, we need to satisfy

$$\left(1+\frac{r_1(\hat{\gamma}_i-1)^2}{\hat{\gamma}_i}\right)(v_i-p_i^i)\geq r_1\alpha_K q_{max}(\frac{1}{\hat{\gamma}_i})^{K-i}v_i,$$

or, alternatively,

$$(v_i-p_i^i)\geq r_1\left(v_i\left(\alpha_K q_{max}\left(\frac{1}{\hat{\gamma}_i}\right)^{K-i}-\frac{(\hat{\gamma}_i-1)^2}{\hat{\gamma}_i}\right)-\frac{p_i^i(\hat{\gamma}_i-1)^2}{\hat{\gamma}_i}\right).$$

But since $(\hat{\gamma}_i-1)^2/\hat{\gamma}_i\geq\alpha_K q_{max}(1/\hat{\gamma}_i)^{K-i}$, we know that $\alpha_K q_{max}(1/\hat{\gamma}_i)^{K-i}-(\hat{\gamma}_i-1)^2/\hat{\gamma}_i\leq 0$, and the right-hand side is at most 0. Since the left-hand side is at least 0 (otherwise our assumption of equilibrium prior to exploration does not hold), any r_1 satisfies the condition.

A.11 Learning Bounds for Constant Slot Ratios

The following theorem giving learning bounds for the constant slot ratio setting is mentioned in Section 8.2. Again, let $n_{i,s}$ denote the number of times we have observed advertiser i in slot s at the current fixed CTR estimates, and let $z_{i,s,j}$ be a random variable indicating whether or not the ad i was clicked on the jth time it appeared in slot s.

Theorem 10. *Suppose that CTRs are of the form $c_i^s=e_i(\gamma_i)^{s-1}$ for all i and s, where γ_i is unknown and $e_i=c_i^1$. Suppose also that the adjacent-ad swapping algorithm has been running for Q queries with a fixed set of broadcasted CTR estimates. Suppose we have observed $n_{i,s}$ instances of ad i at slot s with a fixed set of broadcasted CTR estimates. Let $\hat{c}_i^s$ be our* new *estimate of CTR, defined as:*

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

where

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

Then for any $\delta\in(0,1)$, with probability $1-\delta$, the following bound holds for all advertisers i and slots s:

$$|\hat{c}_i^s-c_i^s|\leq\left(s(s+1)+s_i^2+1\right)\sqrt{\frac{\ln(4N/\delta)}{n_i}}+\frac{(s(s+1)+s_i^2)\ln(4N/\delta)}{\hat{c}_i^{s_i}n_i}.$$

where $n_i=\min\{n_{i,s_i},n_{i,s_i+1}\}$.

The proof is divided into a sequence of lemmas. The first one, stated below, follows from a direct application of Hoeffding's inequality and the union bound. Recall that s_i is a chosen slot to explore for ad i.

Lemma 10. *For each i, with probability at least $1-\delta$, we have*

$$\left|\hat{c}_i^{s_i} - c_i^{s_i}\right| \le \Delta_i, \quad \left|\hat{c}_i^{s_i+1} - c_i^{s_i+1}\right| \le \Delta_i, \tag{7}$$

where

$$\Delta_i = \max\left\{\sqrt{\frac{\ln\frac{4}{\delta}}{2n_{i,s_i}}}, \sqrt{\frac{\ln\frac{4}{\delta}}{2n_{i,s_i+1}}}\right\}.$$

Lemma 11. *Assuming Equation 7 holds. Then for each i we have*

$$|\hat{\gamma}_i^s - \gamma_i^s| \le \frac{s(s+1)\Delta_i}{\hat{c}_i^{s_i}}.$$

Proof. First we can see that for any i,

$$\begin{aligned}|\hat{\gamma}_i - \gamma_i| &= \left|\frac{\hat{c}_i^{s_i+1}}{\hat{c}_i^{s_i}} - \frac{c_i^{s_i+1}}{c_i^{s_i}}\right| = \frac{\left|\hat{c}_i^{s_i+1}c_i^{s_i} - \hat{c}_i^{s_i}c_i^{s_i+1}\right|}{\hat{c}_i^{s_i}c_i^{s_i}} \\ &= \frac{\left|(\hat{c}_i^{s_i+1} - c_i^{s_i+1})c_i^{s_i} - (\hat{c}_i^{s_i} - c_i^{s_i})c_i^{s_i+1}\right|}{\hat{c}_i^{s_i}c_i^{s_i}} \\ &\le \frac{\left|\hat{c}_i^{s_i+1} - c_i^{s_i+1}\right|c_i^{s_i} + \left|\hat{c}_i^{s_i} - c_i^{s_i}\right|c_i^{s_i+1}}{\hat{c}_i^{s_i}c_i^{s_i}} \\ &\le \frac{\Delta_i(c_i^{s_i} + c_i^{s_i+1})}{\hat{c}_i^{s_i}c_i^{s_i}} \le \frac{2\Delta_i}{\hat{c}_i^{s_i}}.\end{aligned}$$

Now, by Taylor's theorem, we have

$$(\hat{\gamma}_i)^s = (\gamma_i)^s + s(\hat{\gamma}_i - \gamma_i)\left((\hat{\gamma}_i)^{s-1} - (\gamma_i)^{s-1}\right) + \frac{1}{2}s(s-1)(\hat{\gamma}_i - \gamma_i)^2(\tilde{\gamma})^{s-2},$$

for some $\tilde{\gamma}$ between γ_i and $\hat{\gamma}_i$. Since $\left|(\hat{\gamma}_i)^{s-1} - (\gamma_i)^{s-1}\right| \le 1$ and $0 \le \gamma_i, \hat{\gamma}_i, \tilde{\gamma} \le 1$,[10] we have

$$\begin{aligned}|(\hat{\gamma}_i)^s - (\gamma_i)^s| &= \left|s(\hat{\gamma}_i - \gamma_i)\left((\hat{\gamma}_i)^{s-1} - (\gamma_i)^{s-1}\right) + \frac{1}{2}s(s-1)(\hat{\gamma}_i - \gamma_i)^2(\tilde{\gamma})^{s-2}\right| \\ &\le |\hat{\gamma}_i - \gamma_i|\left(s + \frac{s(s-1)}{2}\right) \le \frac{s(s+1)\,|\hat{\gamma}_i - \gamma_i|}{2} \le \frac{s(s+1)\Delta_i}{\hat{c}_i^{s_i}}.\end{aligned}$$

With these lemmas, we can prove Theorem 10. For each i and s,

$$\begin{aligned}|\hat{c}_i^s - c_i^s| &= \left|(\hat{\gamma}_i)^{s-s_i}\hat{c}_i^{s_i} - (\gamma_i)^{s-s_i}c_i^{s_i}\right| \\ &= \left|\left((\hat{\gamma}_i)^{s-s_i} - (\gamma_i)^{s-s_i}\right)\hat{c}_i^{s_i} + (\gamma_i)^{s-s_i}\left(\hat{c}_i^{s_i} - c_i^{s_i}\right)\right| \\ &\le \left|(\hat{\gamma}_i)^{s-s_i} - (\gamma_i)^{s-s_i}\right|\hat{c}_i^{s_i} + (\gamma_i)^{s-s_i}\left|\hat{c}_i^{s_i} - c_i^{s_i}\right| \\ &\le \left((s-s_i)^2 + (s-s_i)\right)\Delta_i + (\gamma_i)^{s-s_i}\Delta_i \\ &\le \Delta_i\left(s^2 + s + s_i^2 + (\gamma_i)^{s-s_i}\right).\end{aligned}$$

[10] If $\hat{\gamma}_i$ happens to be greater than 1 (which is possible), then we can safely set it to 1. This change can only make the estimate more accurate, since we know $\gamma_i \in (0, 1]$.

By the previous lemma,

$$(\gamma_i)^{s-s_i} \leq (\hat{\gamma}_i)^{s-s_i} + \frac{((s-s_i)^2 + (s-s_i))\Delta_i}{\hat{c}_i^{s_i}} \leq (\hat{\gamma}_i)^{s-s_i} + \frac{s^2 + s + s_i^2 \Delta_i}{\hat{c}_i^{s_i}},$$

and we obtain

$$|\hat{c}_i^s - c_i^s| \leq \Delta_i \left(s^2 + s + s_i^2 + (\hat{\gamma}_i)^{s-s_i} + \frac{(s^2 + s + s_i^2)\Delta_i}{\hat{c}_i^{s_i}} \right).$$

A simple application of the union bound results in Theorem 10.