

Robust Advertisement Allocation

Shaojie Tang

Naveen Jindal School of Management
 University of Texas at Dallas
 shaojie.tang@utdallas.edu

Abstract

Internet advertising revenue has surpassed broadcast revenue (including cable televisions) very recently due to the rapid growth of e-commerce and information technology. As online advertising has become a major source of revenue for online publishers, such as Google and Amazon, one problem facing them is to optimize the ads selection and allocation in order to maximize their revenue. Although there is a rich body of work that has been devoted to this field, uncertainty about models and parameter settings is largely ignored in existing algorithm design. To fill this gap, we are the first to formulate and study the *Robust Ad Allocation* problem, by taking into account the uncertainty about parameter settings. We define a Robust Ad Allocation framework with a set of candidate parameter settings, typically derived from different users or topics. Our main aim is to develop robust ad allocation algorithms, which can provide satisfactory performance across a spectrum of parameter settings, compared to the (parameter-specific) optimum solutions. We study this problem progressively and propose a series of algorithms with bounded approximation ratio.

1 Introduction

In this paper, we are the first to study the Robust Ad Allocation problem, in which our goal is to maximize the expected revenue by selecting and allocating a group of ads to an ad space with limited number of slots. We assume that the revenue can be generated only from a click on an ad. Different from most of existing works, we take into account the impact of unreliable estimates. We study our problem under a simple and intuitive model called the *Cascade Model* [Craswell *et al.*, 2008]. The cascade model was originally proposed in organic search, and have been empirically verified using data from commercial search engines [Gomes *et al.*, 2009]. The basic cascade model assumes that the user scans through slots in order π . Having examined ad a_i , the user clicks it with (ad-specific) probability q_i , and continues to scan the next ad with (ad-specific) probability c_i . Each ad is associated with a per-click value λ_i . One common objective of ad allocation

is to maximize the expected revenue. Most existing works assume that the parameters in the click-through models are pre-known and accurate. However, this is generally not true due to incomplete or inaccurate knowledge, and unreliable estimations. Here we list two major contributors to this uncertainty.

Unreliable estimation of user’s behavior. Due to incomplete knowledge of user’s profile, and limitations of the estimation model that has been used, some estimates, such as click-through rate, are often inaccurate, or completely wrong in the worst-case. In addition, user’s behavior can be heavily influenced by environmental factors and therefore is highly unpredictable.

Unreliable estimation of user’s identity. Sharing the same account or device with multiple users has been increasingly popular. [White *et al.*, 2014] found that over half of the machine identifiers comprise the search queries of multiple users leading to disorganized search logs for the personalized search service. As reported in [Giura *et al.*, 2014], half of the survey respondents share their tablets with their spouse, resulting in the interweave of historical records without discriminating different users. All these have made it even more difficult to generate accurate estimates.

In this work, we introduce and study the *Robust Ad Allocation* problem by taking into account the uncertainty of underlying models. We present effective and efficient algorithms with provable performance bound in the presence of uncertain parameter settings. The ultimate goal is to find an ad allocation that can simultaneously maximize the expected revenue across all candidate models. Compared with stochastic guarantees, our adversarial setting is able to generate more robust solutions in practice.

The main contributions are summarized as follows.

1. To best capture all those uncertainties mentioned earlier, we introduce a model space Θ that is composed of a number of candidate models. Each candidate model $\theta \in \Theta$ is derived from a possible parameter setting, including click-through rate, continuation probability, and slate sequence. We define our robust advertising problem based on this model space and aim at finding an ad allocation that maximize the worst-case approximation ratio across all candidate models.
2. The model we analyze evolves progressively, and we

propose a serial of algorithms using Double Oracle method. We prove that all of our approaches achieve optimal or near-optimal performance.

3. We also extend our results to continuous space model and show how to build a finite model space from a continuous model.

2 Related Work

Due to the increasing popularity of e-commerce and World Wide Web, there is huge body of work that has been devoted to the topic of internet advertising (e.g., [Edelman *et al.*, 2005; Lahaie *et al.*, 2007; Varian, 2007]). Some of them [Dembczynski *et al.*, 2008; Attenberg *et al.*, 2009; Xiong *et al.*, 2012; Wang *et al.*, 2013] focus on predicting the click probability of an ad. Meanwhile, several papers focus on the effectiveness measurement of internet ad campaign [Gerken, 2008; Lindsay *et al.*, 2009; Harvey *et al.*, 2010]. The other category of work study internet ad campaign problem [Lai *et al.*, 2010; Baldacci *et al.*, 2013; Hwang *et al.*, 2013] by formulating it as a multi-stage decision problem. Our work is closely related to sponsored search auctions [Aggarwal *et al.*, 2008]. One major challenge in our setting is to deal with the negative externalities among advertisements, e.g., the click probability of all future ads decreases as a user may leave the site before examining them. [Ghosh and Mahdian, 2008] adopted a choice model to describe user click behaviors, and they show that the optimization problem under this model is NP-hard. [Aggarwal *et al.*, 2008; Kempe and Mahdian, 2008] study the social welfare maximization problem under the cascade model. However, majority of existing works assume that the click through rate is known precisely in advance. This assumption may not always hold due to incomplete or inaccurate knowledge, and unreliable estimations. We are the first to study the Robust Ad Allocation problem, taking into account the uncertainty of underlying models. We present a series of effective and efficient algorithms in the presence of uncertain parameter settings.

3 Click-Through Model and Problem Formulation

We consider the problem of selecting and allocating a set of n ads \mathcal{A} to display in an ad space with m slots. Each ad $a_i \in \mathcal{A}$ is associated with a per-click revenue λ_{a_i} , which specifies the amount of advertising revenue that is generated from a click on that ad. In order to maximize the expected revenue that is generated, the ad network needs to optimize their selection of ads and carefully decide the sequencing of selected ads in an ad space. Next we introduce the click-through model and important notations adopted in this paper.

3.1 Click-Through Model

Basic Model: We adopt *Cascade Model* [Craswell *et al.*, 2008] in this paper and assume that the user will view the ads sequentially. After examining an ad, say a_i , in the sequence, the user clicks a_i with probability q_{a_i} . This click probability is decided by the intrinsic quality or relevance of ad a_i . Independently of whether ad a_i was clicked or not, the user

continues to examine the next ad with probability c_{a_i} ; otherwise, terminates the scanning process. Let $\sigma \in \Sigma$ denote one ad allocation, where Σ is the (ad allocation) strategy space, we use $\sigma(i)$ to denote the ad placed in slot i by σ . Given any ad allocation $\sigma \in \Sigma$, the user will see a particular slot k with probability $\prod_{i=1}^{k-1} c_{\sigma(i)}$, and the click-through rate of ad $a_{\sigma(k)}$ is therefore $q_{\sigma(k)} \prod_{i=1}^{k-1} c_{\sigma(i)}$. As a result, the expected revenue of σ can be calculated as:

$$\gamma(\sigma) = \sum_{k=1}^m (\lambda_{\sigma(k)} q_{\sigma(k)} \prod_{i=1}^{k-1} c_{\sigma(i)}) \quad (1)$$

3.2 Robust Ad Allocation

Our work is motivated by the following observation: Due to the unreliable estimation of user's preference and identity, the parameters used in Equation (1), including the click probability and continuation probability of each ad, are not known precisely. As a result, there may exist a number of candidate click through models due to different parameter settings. In addition, the sponsored search ads are often placed in multiple different slates by search engines [Kempe and Mahdian, 2008], e.g., the mainline ads and the sidebar ads. It is possible that different users may scan those slates in different orders.

To best capture these uncertainties, we introduce a model space Θ that is composed of possible candidate models. Each candidate model $\theta \in \Theta$ can be represented by

$$\theta \triangleq \langle \mathbf{q}^\theta, \mathbf{c}^\theta, \pi^\theta \rangle$$

- $\mathbf{q}^\theta = \{q_{a_i}^\theta | a_i \in \mathcal{A}\}$ is the click probabilities under model θ , where $q_{a_i}^\theta$ represents the click probability of a_i under θ ;

- $\mathbf{c}^\theta = \{c_{a_i}^\theta | a_i \in \mathcal{A}\}$ is the continuation probabilities under model θ , where $c_{a_i}^\theta$ represents the continuation probability of a_i under θ ;

- π^θ is the slate sequence of model θ . Assume there are L of slates $S = \{s_1, \dots, s_L\}$, where L is a small constant in practise. Each slate $s \in S$ contains m_s slots, we use $\pi^\theta(l)$ to denote the l -th slate that is examined by the user under θ .

In the rest of this paper, we assume that the model space is much smaller than the allocation space: $|\Theta| \ll |\Sigma|$. We will relax this assumption in Section 8 by extending our results to continuous space model. Given a candidate model θ and allocation σ , let σ_θ^* denote the optimal allocation under θ , we define the approximation ratio of σ under θ as

$$\rho_\theta(\sigma) = \frac{\gamma_\theta(\sigma)}{\gamma_\theta(\sigma_\theta^*)}.$$

Our objective is to find a mixed strategy of ad allocation that maximizes the worst-case approximation ratio across all candidate models. Next we formally define our *Robust Ad Allocation* problem (RA).

RA: Maximize $\min_{\theta \in \Theta} \sum_{\sigma \in \Sigma} x_\sigma \rho_\theta(\sigma)$
subject to:

$$\begin{cases} \sum_{\sigma \in \Sigma} x_\sigma = 1 \\ x_\sigma \geq 0, \forall \sigma \in \Sigma \end{cases}$$

In the above formulation, x_σ indicates the probability that σ will be selected. Our goal is to find a mixed strategy $x_\Sigma = \{x_\sigma | \sigma \in \Sigma\}$ that performs well in the worst-case. Due to space constraints, the missing proofs are deferred to the full version.

4 Optimal Policy when $\theta \triangleq \langle q^\theta, c, \pi \rangle$.

To motivate our ideas, we first study a special case by assuming both continuation probability and slate sequence are model-independent, i.e., all models share the same continuation probability c and slate sequence π .

After introducing an additional variable Z , below is the standard formulation for computing a maximin strategy.

Primal: Maximize $Z_{\sigma \in \Sigma}$

subject to:

$$\begin{cases} \sum_{\sigma \in \Sigma} x_\sigma = 1 \\ x_\sigma \geq 0, \forall \sigma \in \Sigma \\ Z \leq \sum_{\sigma \in \Sigma} x_\sigma \rho_\theta(\sigma), \forall \theta \in \Theta \end{cases}$$

This LP has $|\Sigma|$ variables which could easily be exponential in the number of ads. One possible approach to handle this case is ellipsoid algorithm [Grötschel *et al.*, 1981]. However, this method lacks of numerical stability and suffers from poor performance. We decide to adopt Double Oracle [McMahan *et al.*, 2003] which is capable of solving certain LP problems with exponential number of variables more efficiently in practice.

4.1 Double Oracle

The basic idea of Double Oracle is to model RA as a zero-sum game between the defender and the attacker. The pure strategies of the defender correspond to the ad allocation space Σ , and the pure strategies of the attacker correspond to the candidate model space Θ . Our goal is to find a maximin strategy for the defender. The detailed description of Double Oracle is listed in Algorithm 1. $\bar{\Sigma}$ is the set of ad allocations generated so far and $\bar{\Theta}$ is the set of attacker models generated so far. *Primal*($\bar{\Sigma}, \bar{\Theta}$) finds an equilibrium of the two-player zero-sum game composed of the sets of pure strategies, $\bar{\Sigma}$ and $\bar{\Theta}$. The formal definition of *Primal*($\bar{\Sigma}, \bar{\Theta}$) is listed as follows.

Primal ($\bar{\Sigma}, \bar{\Theta}$): Maximize $Z_{\sigma \in \bar{\Sigma}}$

subject to:

$$\begin{cases} \sum_{\sigma \in \bar{\Sigma}} x_\sigma = 1 \\ x_\sigma \geq 0, \forall \sigma \in \bar{\Sigma} \\ Z \leq \sum_{\sigma \in \bar{\Sigma}} x_\sigma \rho_\theta(\sigma), \forall \theta \in \bar{\Theta} \end{cases}$$

Primal($\bar{\Sigma}, \bar{\Theta}$) and its dual return $x_{\bar{\Sigma}}$ and $y_{\bar{\Theta}}$, which are the current equilibrium mixed strategies over $\bar{\Sigma}$ and $\bar{\Theta}$ respectively. The Defender Oracle (Section 4.2) generates an ad allocation $\sigma \in \Sigma$ that is a best response for the defender against $y_{\bar{\Theta}}$. (This is a best response from Σ , not restricted to $\bar{\Sigma}$). By enumerating all models in Θ , the Attacker Oracle (Section 4.3) generates a attacker model θ that is a best response for the

attacker against $x_{\bar{\Sigma}}$. (This is the best response from Θ , not just $\bar{\Theta}$.)

The starting point of double oracle is a small set of pure strategies of each player, and we will expand this set in iterations by applying the best-response oracles from both players to the current solution. The convergence is achieved when $\bar{\Sigma}$ and $\bar{\Theta}$ can not be expanded further. A more detailed description and analysis of this approach can be found in [McMahan *et al.*, 2003]. As shown in [McMahan *et al.*, 2003], if both Defender Oracle and Attacker Oracle can be solved optimally, the double oracle approach returns an optimal mixed strategy. Then the following theorem follows immediately from Lemma 3 and the optimality of the Attacker Oracle.

Theorem 1 *The double oracle approach (Algorithm 1) returns an optimal mixed strategy of ad allocation $x_{\bar{\Sigma}}$.*

Algorithm 1 Double Oracle for Robust Advertising

- 1: Initialize $\bar{\Sigma}$ by selecting an arbitrary ad allocation.
 - 2: Initialize $\bar{\Theta}$ by selecting an arbitrary model.
 - 3: **repeat**
 - 4: $(x_{\bar{\Sigma}}, y_{\bar{\Theta}}) \leftarrow \text{Primal}(\bar{\Sigma}, \bar{\Theta})$;
 - 5: $\sigma \leftarrow \mathbf{DO}(y_{\bar{\Theta}})$; ▷ Algorithm 2
 - 6: $\bar{\Sigma} = \bar{\Sigma} \cup \{\sigma\}$;
 - 7: $\theta \leftarrow \mathbf{AO}(x_{\bar{\Sigma}})$; ▷ Section 4.3
 - 8: $\bar{\Theta} = \bar{\Theta} \cup \{\theta\}$;
 - 9: **until** convergence
 - 10: Return $(x_{\bar{\Sigma}}, y_{\bar{\Theta}})$
-

4.2 Defender Oracle

In this section, we focus on solving the Defender Oracle problem and compute the best pure strategy for the defender.

Definition 1 (Defender Oracle Problem) *Given an attacker mixed strategy $y_{\bar{\Theta}} = \{y_\theta | \theta \in \bar{\Theta}\}$, i.e., $\theta \in \bar{\Theta}$ happens with probability y_θ , find $\arg \max_\sigma \sum_{\theta \in \bar{\Theta}} y_\theta \rho_\theta(\sigma)$, i.e., an ad allocation that maximizes the expected approximation ratio.*

Given a fixed model θ , its optimal allocation σ_θ^* and revenue γ_θ^* can be computed using dynamic programming [Kempe and Mahdian, 2008]. To simplify the notation, we use γ_θ^* to represent $\gamma_\theta(\sigma_\theta^*)$.

Lemma 1 *Consider any ad allocation σ , the expected approximation ratio of σ under $y_{\bar{\Theta}}$ is*

$$\rho(\sigma) = \sum_{k=1}^m \left(\lambda_{\sigma(k)} \sum_{\theta \in \bar{\Theta}} (y_\theta \frac{q_{\sigma(k)}^\theta}{\gamma_\theta^*} \prod_{i=1}^{k-1} c_{\sigma(i)}) \right).$$

Proof: Since the click through rate of ad $\sigma(k)$ under θ is $q_{\sigma(k)}^\theta \prod_{i=1}^{k-1} c_{\sigma(i)}$, then the expected revenue of $\sigma(k)$ under θ is $\lambda_{\sigma(k)} q_{\sigma(k)}^\theta \prod_{i=1}^{k-1} c_{\sigma(i)}$. Thus, the expected approximation ratio of σ under $y_{\bar{\Theta}}$ is $\sum_{k=1}^m \left(\lambda_{\sigma(k)} \sum_{\theta \in \bar{\Theta}} (y_\theta \frac{q_{\sigma(k)}^\theta}{\gamma_\theta^*} \prod_{i=1}^{k-1} c_{\sigma(i)}) \right)$. □

The next lemma reveals an important structural property of any optimal ad allocation.

Lemma 2 Assume σ is an optimal ad allocation under $y_{\bar{\Theta}}$. Consider any two ads $\sigma(k_1)$ and $\sigma(k_2)$ where $k_1 \leq k_2$, we have

$$\frac{\lambda_{\sigma(k_1)} \sum_{\theta \in \bar{\Theta}} (y_{\theta} \frac{q_{\sigma(k_1)}^{\theta}}{\gamma_{\theta}^*})}{1 - c_{\sigma(k_1)}} \geq \frac{\lambda_{\sigma(k_2)} \sum_{\theta \in \bar{\Theta}} (y_{\theta} \frac{q_{\sigma(k_2)}^{\theta}}{\gamma_{\theta}^*})}{1 - c_{\sigma(k_2)}} \quad (2)$$

Given Lemma 2 in hand, we next show how to compute an optimal ad allocation using dynamic programming. This part is similar to the approach developed for traditional ad allocation in [Kempe and Mahdian, 2008]. We first sort all ads

in non-decreasing order of $\frac{\lambda_{a_i} \sum_{\theta \in \bar{\Theta}} (y_{\theta} \frac{q_{a_i}^{\theta}}{\gamma_{\theta}^*})}{1 - c_{a_i}}$. In the following recurrence, $\rho[a_i, t]$ stores the optimum value that can be obtained from ads a_i, \dots, a_n in slots t, \dots, T . According to Lemma 2, once a_i has been selected, it should be placed in slot t . The conditional optimal value gained from placing a_i in slot t can be calculated as $\lambda_{a_i} \sum_{\theta \in \bar{\Theta}} (y_{\theta} \frac{q_{a_i}^{\theta}}{\gamma_{\theta}^*}) + c_{a_i} \rho[a_{i-1}, t-1]$. Otherwise, the optimal value is $\rho[a_{i-1}, t]$, conditioned on a_i is not selected. Thus we adopt the following recurrence:

$$\rho[a_i, t] = \max\{\lambda_{a_i} \sum_{\theta \in \bar{\Theta}} (y_{\theta} \frac{q_{a_i}^{\theta}}{\gamma_{\theta}^*}) + c_{a_i} \rho[a_{i-1}, t-1], \rho[a_{i-1}, t]\}.$$

Algorithm 2 DO ($y_{\bar{\Theta}}$)

- 1: Sort all ads in non-decreasing order of $\frac{\lambda_{a_i} \sum_{\theta \in \bar{\Theta}} (y_{\theta} \frac{q_{a_i}^{\theta}}{\gamma_{\theta}^*})}{1 - c_{a_i}}$
- 2: Adopt dynamic programming to find the optimal ad allocation

$$\rho[a_i, t] = \max\{\lambda_{a_i} \sum_{\theta \in \bar{\Theta}} (y_{\theta} \frac{q_{a_i}^{\theta}}{\gamma_{\theta}^*}) + c_{a_i} \rho[a_{i-1}, t-1], \rho[a_{i-1}, t]\}$$

Then we obtain the following lemma.

Lemma 3 Given an attacker mixed strategy $y_{\bar{\Theta}}$, $\text{DO}(y_{\bar{\Theta}})$ outputs an optimal ad allocation.

4.3 Attacker Oracle

The Attacker Oracle problem can be formulated as follows.

Definition 2 (Attacker Oracle Problem) Given a defender mixed strategy $x_{\bar{\Sigma}} = \{x_{\sigma} | \theta \in \bar{\Sigma}\}$, i.e., $\sigma \in \bar{\Sigma}$ happens with probability x_{σ} , find $\arg \min_{\theta \in \Theta} \sum_{\sigma \in \bar{\Sigma}} x_{\sigma} \rho_{\theta}(\sigma)$, i.e., a model that minimizes the expected approximation ratio.

Since we are assuming $|\Theta| \ll |\Sigma|$, a best attacker's response $\arg \min_{\theta \in \Theta} \sum_{\sigma \in \bar{\Sigma}} x_{\sigma} \rho_{\theta}(\sigma)$ can be found through enumeration.

5 PTAS when $\theta \triangleq \langle \mathbf{q}, \mathbf{c}, \pi^{\theta} \rangle$

In this section, we discuss another special case when all models share the same click-through rate and continuation probability $\theta \triangleq \langle \mathbf{q}, \mathbf{c}, \pi^{\theta} \rangle$. Notice that the only difference between different candidate models come from the slate sequence π^{θ} . It turns out that the Defender Oracle problem under this case

has already be studied in [Kempe and Mahdian, 2008] where they provide a polynomial-time approximation scheme (PTAS). As a result, we can use their approach to build a Defender Oracle and find a PTAS for the original problem.

6 Approximate Policy when $\theta \triangleq \langle \mathbf{q}^{\theta}, \mathbf{c}, \pi^{\theta} \rangle$

We next study a generalized model by assuming all models share only the same continuation probability \mathbf{c} . We still leverage the double oracle method to tackle this case.

6.1 Double Oracle

We first notice that previous Attacker Oracle can still find the best response for the attacker in each iteration through enumeration. However, as shown in Corollary 3, our Defender Oracle (**a-DO**) can only find a $(1 - \delta)$ approximate best response in each iteration, where $\delta \in (0, 1]$ is a control parameter. Based on Theorem 2 in [McMahan and Gordon, 2003], we have the following performance bound when the Defender Oracle is only approximate best response oracle.

Theorem 2 The double oracle approach with **a-DO** (Algorithm 3) as an approximate Defender Oracle returns a $(1 - \delta)$ approximate solution.

6.2 Defender Oracle

We first introduce the outline of **a-DO** (Algorithm 3). Our main idea is to convert the original Defender Oracle problem to an new problem with smaller size, which enables us to find an approximate solution through enumeration. Without loss of generality, assume $a_1, \dots, a_{|\mathcal{A}|}$ are sorted in non-increasing order of their continuation probabilities, let g be the largest number that satisfies $\prod_{i=1}^g c_{a_i} \geq \delta$. Recall that m_s is the size of slate s in the original setting, we now introduce a restricted Defender Oracle problem by restricting s 's size to $\min\{g, m_s\}$ for every slate $s \in S$. We next show that the gap between the restricted optimal solution and the actual optimal solution can be bounded.

Lemma 4 Given any attacker mixed strategy $y_{\bar{\Theta}}$, let σ^r denote the restricted optimal ad allocation with slate size $|s| = \min\{g, m_s\}$ for all slates $s \in S$, and σ^o denote the optimal allocation under the original setting, it holds that

$$\sum_{\theta \in \bar{\Theta}} y_{\theta} \rho_{\theta}(\sigma^r) \geq (1 - \delta) \sum_{\theta \in \bar{\Theta}} y_{\theta} \rho_{\theta}(\sigma^o).$$

Since the size of each slate is upper bounded by g , we can find the best ad allocation through enumeration and its time complexity is $O(N^{gL})$. The parameter g can be fine tuned, giving the tradeoff curve with time complexity on the one end and the approximation ratio on the other end. Based on Lemma 4, we have

Corollary 3 Given any attacker mixed strategy $y_{\bar{\Theta}}$, **a-DO**($y_{\bar{\Theta}}$) achieves approximation ratio $(1 - \delta)$.

6.3 Improving the Time Complexity

In fact, we can further reduce the time complexity of Algorithm 3. In Lemma 5, we prove that a result similar to

Algorithm 3 a-DO ($y_{\overline{\Theta}}$)

- 1: Assume $a_1, \dots, a_{|A|}$ are sorted in non-increasing order of their continuation probabilities;
- 2: Find the largest g that satisfies $\prod_{i=1}^g c_{a_i} \geq \delta$;
- 3: Return the best restricted solution σ^T ;

Lemma 2 also holds within each slate for this case. Given θ , let $S^\theta(s)$ denote all slates examined before s under slate sequence π^θ . We write an ad allocation as $\sigma = \{\sigma_1, \dots, \sigma_L\}$ where σ_s represents the ad allocation within each slate s . Consider any model θ and allocation σ , the probability that slate s can be examined by the user is $C_s^\theta(\sigma) \triangleq \prod_{s \in S^\theta(s)} \prod_{a \in \sigma_s} c_a$.

Lemma 5 Assume $\sigma = \{\sigma_1, \dots, \sigma_L\}$ is an optimal allocation under $y_{\overline{\Theta}}$. Consider any two ads $\sigma_s(k_1)$ and $\sigma_s(k_2)$ in slate s with $k_1 \leq k_2$, we have

$$\frac{\lambda_{\sigma_s(k_1)} \sum_{\theta \in \overline{\Theta}} (y_\theta \frac{C_s^\theta(\sigma)}{\gamma_\theta^*} q_{\sigma_s(k_1)}^\theta)}{1 - c_{\sigma_s(k_1)}} \geq \frac{\lambda_{\sigma_s(k_2)} \sum_{\theta \in \overline{\Theta}} (y_\theta \frac{C_s^\theta(\sigma)}{\gamma_\theta^*} q_{\sigma_s(k_2)}^\theta)}{1 - c_{\sigma_s(k_2)}} \quad (3)$$

Lemma 5 implies that given a group of candidate ads in each slate, we do not really need to enumerate all possible orderings. Instead, we can simply sort them in non-increasing order of $\frac{\lambda_a \sum_{\theta \in \overline{\Theta}} (y_\theta \frac{C_s^\theta(\sigma)}{\gamma_\theta^*} q_a^\theta)}{1 - c_a}$.

7 Approximate Policy when $\theta \triangleq \langle \mathbf{q}^\theta, \mathbf{c}^\theta, \pi^\theta \rangle$

For the ultimate model when $\theta \triangleq \langle \mathbf{q}^\theta, \mathbf{c}^\theta, \pi^\theta \rangle$, we can still use **a-DO** ($y_{\overline{\Theta}}$) as a Defender Oracle to get an approximate solution. The only difference is that we need to sort all ads in non-increasing order of their continuation probabilities for each model, and find the largest g that satisfies $\max_{\theta \in \overline{\Theta}} \prod_{i=1}^g c_{a_i}^\theta \geq \delta$. However, the enhanced approach developed in Section 6.3 can not apply to this case. This is because the property established in Lemma 5 does not hold when the continuation probabilities are different across different models.

Lemma 6 Given any attacker mixed strategy $y_{\overline{\Theta}}$, **a-DO** ($y_{\overline{\Theta}}$) outputs an ad allocation with approximation ratio $(1 - \delta)$.

Then we have the following theorem.

Theorem 4 The double oracle approach with **a-DO** ($y_{\overline{\Theta}}$) as an approximate Defender Oracle returns a $(1 - \delta)$ approximate solution.

8 Extensions to Continuous Model Space

So far we assume a discrete model space, we next discuss the case when the model space is continuous. One important special case of this model is the Perturbation Interval (PI) model [He and Kempe, 2016]. In their model, the uncertainty about parameters are characterized as an interval. For each ad a_i , we have two intervals $I_{a_i}^q = [l_i^q, r_i^q]$ and $I_{a_i}^c = [l_i^c, r_i^c]$, and we only know that the actual click probability of a_i lies in interval $I_{a_i}^q$ and the actual continuation probability of a_i lies

in interval $I_{a_i}^c$; the exact value can be decided by an adversary. It seems that Θ is infinite, e.g., the number of candidate models is infinite under this model.

Our solution to this difficulty is to discretize the space model as follows: We partition each interval into a group of points with step-size ϵ . Thus, for each ad a_i , interval $I_{a_i}^q$ is mapped to $\lfloor \frac{r_i^q - l_i^q}{\epsilon} \rfloor$ points: $Q_{a_i} = \{l_i^q + \epsilon, l_i^q + 2\epsilon, \dots, r_i^q\}$, and interval $I_{a_i}^c$ is mapped to $\lfloor \frac{r_i^c - l_i^c}{\epsilon} \rfloor$ points: $C_{a_i} = \{l_i^c + \epsilon, l_i^c + 2\epsilon, \dots, r_i^c\}$. Then we can build an approximate discrete space model as $\Theta = \times_{a_i \in A} (Q_{a_i} \times C_{a_i})$. We next prove that the solution derived under the approximate discrete space model is close to the optimal solution under the PI model.

Lemma 7 The gap between the worst case approximation ratio achieved by the optimal solution under PI model and the one achieved under the approximate discrete space model is upper bounded by $\frac{m(m+3)}{2} \frac{\lambda_{\max} \epsilon}{\underline{\gamma}}$, where $\underline{\gamma} = \min_{\theta \in \Sigma} \gamma_\theta(\sigma_\theta^*)$.

Proof: Given any model $\hat{\theta}$ in PI model, we round down $q_{a_i}^{\hat{\theta}}$ and $c_{a_i}^{\hat{\theta}}$ to the nearest point in Q_{a_i} and C_{a_i} , respectively. Now consider an ad allocation σ and any ad $\sigma(k)$ at the k -slot, the probability that $\sigma(t)$ is clicked is $\prod_{t=1}^{k-1} c_{\sigma(t)}$. After rounding down $c_{a_i}^{\hat{\theta}}$, this probability is decreased to $\prod_{t=1}^{k-1} (c_{\sigma(t)} - \min\{\epsilon, c_{\sigma(t)}\})$. We next prove that the gap between the ‘‘reachability’’ of $\sigma(k)$ before and after rounding down is bounded by $k\epsilon$:

$$\prod_{t=1}^{k-1} (c_{\sigma(t)} - \min\{\epsilon, c_{\sigma(t)}\}) \geq \prod_{t=1}^{k-1} c_{\sigma(t)} - k\epsilon.$$

We prove this through induction on k . The base case when $k = 1$ is obvious since the probability to examine the first ad is always 1. Next we assume that this relation holds when $k = i$, i.e., $(\prod_{t=1}^{i-1} (c_{\sigma(t)} - \min\{\epsilon, c_{\sigma(t)}\})) \geq \prod_{t=1}^{i-1} c_{\sigma(t)} - (i-1)\epsilon$. It follows that

$$\begin{aligned} & \prod_{t=1}^i (c_{\sigma(t)} - \min\{\epsilon, c_{\sigma(t)}\}) \\ &= \left(\prod_{t=1}^{i-1} (c_{\sigma(t)} - \min\{\epsilon, c_{\sigma(t)}\}) \right) (c_{\sigma(i)} - \min\{\epsilon, c_{\sigma(i)}\}) \\ &\geq \left(\prod_{t=1}^{i-1} c_{\sigma(t)} - (i-1)\epsilon \right) (c_{\sigma(i)} - \min\{\epsilon, c_{\sigma(i)}\}) \text{ (by induction)} \\ &= \prod_{t=1}^i c_{\sigma(t)} - (i-1)\epsilon(c_{\sigma(i)} + \min\{\epsilon, 1 - c_{\sigma(i)}\}) \\ &\quad - \prod_{t=1}^{i-1} c_{\sigma(t)} \min\{\epsilon, c_{\sigma(i)}\} \\ &\geq \prod_{t=1}^i c_{\sigma(t)} - (i-1)\epsilon - \epsilon = \prod_{t=1}^i c_{\sigma(t)} - i\epsilon \end{aligned}$$

By similar approach, we can prove that rounding down $q_{a_i}^{\hat{\theta}}$ will add another ϵ , thus the click through rate of $\sigma(k)$ is decreased by at most $(k+1)\epsilon$. Therefore, after rounding down, the revenue of any given ad allocation is decreased by at most

$\lambda_{\max} \sum_{t=1}^m (t+1)\epsilon = \lambda_{\max} \frac{m(m+3)\epsilon}{2}$. It follows that, the approximation ratio of any model $\hat{\theta}$ is changed to

$$\frac{\gamma_{\hat{\theta}}(\sigma) - \frac{m(m+3)}{2} \lambda_{\max} \epsilon}{\gamma_{\hat{\theta}}(\sigma_{\hat{\theta}}^*)} \geq \frac{\gamma_{\hat{\theta}}(\sigma)}{\gamma_{\hat{\theta}}(\sigma_{\hat{\theta}}^*)} - \frac{m(m+3)}{2} \frac{\lambda_{\max} \epsilon}{\underline{\gamma}}$$

This finishes the proof. \square

8.1 Pure Strategy Setting

We next reveal another interesting property of PI model under pure strategy setting. If only pure strategy is allowed, PI model can be reduced to an equivalent discrete space model.

Lemma 8 *Under PI model with a fixed slate sequence, the worst case for the ratio in ρ for any ad allocation σ is achieved by making each q_{a_i} equal to l_i^q or r_i^q , and each c_{a_i} equal to l_i^c or r_i^c .*

Proof: Our analysis is inspired by the one conducted in [He and Kempe, 2016]. We first consider the selection of the click probability. Pick any ad \bar{a} and fix the parameters of the other ads $a \neq \bar{a}$. Consider any ads allocation σ and define $\gamma_{\sigma}(x)$ as the expected revenue of σ when the click probabilities of all ads $a \neq \bar{a}$ are q_a and the click probability of \bar{a} is x . In addition, we assume that the continuation probabilities of all ads a , including \bar{a} , are c_a . Then $\gamma_{\sigma}(x)$ can be calculated as follows: if \bar{a} is not selected in σ , $\gamma_{\sigma}(x) = \sum_{k=1}^m (\lambda_{\sigma(k)} q_{\sigma(k)} \prod_{i=1}^{k-1} c_{\sigma(i)})$; otherwise, if \bar{a} is located in slot z ,

$$\begin{aligned} \gamma_{\sigma}(x) = & \sum_{k=1}^{z-1} (\lambda_{\sigma(k)} q_{\sigma(k)} \prod_{i=1}^{k-1} c_{\sigma(i)}) + \lambda_{\bar{a}} x \prod_{i=1}^{z-1} c_{\sigma(i)} \\ & + \sum_{k=z+1}^m (\lambda_{\sigma(k)} q_{\sigma(k)} \prod_{i=1}^{k-1} c_{\sigma(i)}). \end{aligned}$$

It follows that $\gamma_{\sigma}(x)$ is a linear function of x .

Let $\gamma_{\sigma}^*(x) = \arg \max_{\sigma} \gamma_{\sigma}(x)$ denote the maximum revenue gained from this model. Since $\gamma_{\sigma}^*(x)$ is a maximum of linear functions of x , it is convex and piecewise linear. Thus using level set argument, we can prove that the approximation ratio $\frac{\gamma_{\sigma}(x)}{\gamma_{\sigma}^*(x)}$ is quasi-concave. As a result, the minimum value of this function lies at one of the endpoints of the interval.

We next consider the selection of the continuation probability. Pick any ad \bar{a} and fix the parameters of the other ads $a \neq \bar{a}$. Consider any ads allocation σ and define $\gamma_{\sigma}(y)$ as the expected revenue of σ when the continuation probabilities of all ads $a \neq \bar{a}$ are c_a and the continuation probability of \bar{a} is y . In addition, we assume that the continuation probabilities of all ads a , including \bar{a} , are c_a . Then $\gamma_{\sigma}(x)$ can be calculated as follows: if \bar{a} is not selected in σ , $\gamma_{\sigma}(x) = \sum_{k=1}^m (\lambda_{\sigma(k)} q_{\sigma(k)} \prod_{i=1}^{k-1} c_{\sigma(i)})$; otherwise, if \bar{a} is located in slot z ,

$$\begin{aligned} \gamma_{\sigma}(x) = & \sum_{k=1}^z (\lambda_{\sigma(k)} q_{\sigma(k)} \prod_{i=1}^{k-1} c_{\sigma(i)}) \\ & + \sum_{k=z+1}^m (\lambda_{\sigma(k)} q_{\sigma(k)} \prod_{i=1}^{z-1} c_{\sigma(i)} y \prod_{i=z+1}^{k-1} c_{\sigma(i)}). \end{aligned}$$




Figure 1: Worst-case approximation ratio vs. case number.

In either case, $\gamma_{\sigma}(y)$ is a linear function of y . Let $\gamma_{\sigma}^*(y) = \arg \max_{\sigma} \gamma_{\sigma}(y)$ denote the maximum revenue gained from this model. Since $\gamma_{\sigma}^*(y)$ is a maximum of linear functions of y , it is convex and piecewise linear. Similar to the proof for the selection of click probability, we can prove that the minimum value of this function lies at one of the endpoints of the interval.

This lemma can be proved by repeating the same arguments on all ads one by one. \square

Based on Lemma 8, we can easily construct a discrete model space Θ by enumerating all slate sequences, and parameters chosen from $q_{a_i} \in \{l_i^q, r_i^q\}$, and $c_{a_i} \in \{l_i^c, r_i^c\}$ for each ad a_i .

9 Numerical Evaluation

We generate multiple sets of candidate click-through models as follows. For each candidate model $\theta \in \Theta$, we set the number of ads $n = 100$, the number of slates $L = 3$, the number of ad slots in each slate $m_s = 5$, the continuation probability of each ad $c_{a_i}^{\theta} = 0.9$. The click through probability of each ad, $q_{a_i}^{\theta}$, is randomly sampled from $[0, 1]$, and the revenue of clicking each ad is randomly selected from $[1, 10]$. A slate sequence, π^{θ} , is randomly generated for each candidate model. We run the simulations for each set of candidate models multiple rounds and report the worst-case approximation ratio in Figure 1. As shown in the figure, the approximation ratio stays above 0.71 across all the test cases, which empirically demonstrates the effectiveness of our algorithm.

10 Conclusion

In this paper, we study the Robust Ad Allocation problem with unreliable estimates. The input of our problem is a set of candidate parameter settings, typically derived from different users or topics, our goal is to develop a robust ad allocation scheme, which can provide satisfactory performance across a spectrum of parameter settings, compared to the (parameter-specific) optimum solutions. We study this problem progressively and propose a series of algorithms based on Double Oracle approach with bounded approximation ratio. We also extend our results to continuous space model. One limitation of this work is that we assume the size of model space is relatively small, and the Attacker Oracle problem is solved using the brute force approach. We leave it as future work to find scalable solutions for the Attacker Oracle problem.

References

- [Aggarwal *et al.*, 2008] Gagan Aggarwal, Jon Feldman, S Muthukrishnan, and Martin Pál. Sponsored search auctions with markovian users. In *International Workshop on Internet and Network Economics*, pages 621–628. Springer, 2008.
- [Attenberg *et al.*, 2009] Josh Attenberg, Sandeep Pandey, and Torsten Suel. Modeling and predicting user behavior in sponsored search. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1067–1076. ACM, 2009.
- [Baldacci *et al.*, 2013] Roberto Baldacci, Aristide Mingozzi, Roberto Roberti, and Roberto Wolfler Calvo. An exact algorithm for the two-echelon capacitated vehicle routing problem. *Operations Research*, 61(2):298–314, 2013.
- [Craswell *et al.*, 2008] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. An experimental comparison of click position-bias models. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*, pages 87–94. ACM, 2008.
- [Dembczynski *et al.*, 2008] Krzysztof Dembczynski, Wojciech Kotłowski, and Dawid Weiss. Predicting ads click-through rate with decision rules. In *Workshop on targeting and ranking in online advertising*, volume 2008, 2008.
- [Edelman *et al.*, 2005] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Technical report, National Bureau of Economic Research, 2005.
- [Gerken, 2008] David A Gerken. System and method for selectively acquiring and targeting online advertising based on user ip address, May 20 2008. US Patent 7,376,714.
- [Ghosh and Mahdian, 2008] Arpita Ghosh and Mohammad Mahdian. Externalities in online advertising. In *Proceedings of the 17th international conference on World Wide Web*, pages 161–168. ACM, 2008.
- [Giura *et al.*, 2014] Paul Giura, Ilona Murynets, Roger Piqueras Jover, and Yevgeniy Vahlis. Is it really you?: user identification via adaptive behavior fingerprinting. In *Proceedings of the 4th ACM conference on Data and application security and privacy*, pages 333–344. ACM, 2014.
- [Gomes *et al.*, 2009] Renato Gomes, Nicole Immorlica, and Evangelos Markakis. Externalities in keyword auctions: An empirical and theoretical assessment. In *International Workshop on Internet and Network Economics*, pages 172–183. Springer, 2009.
- [Grötschel *et al.*, 1981] Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.
- [Harvey *et al.*, 2010] William Morris Harvey, Gerald Leo Despain, Mark Lieberman, Brian P Canning, and Pavel Bochman. Analyzing return on investment of advertising campaigns by matching multiple data sources, June 1 2010. US Patent 7,729,940.
- [He and Kempe, 2016] Xinran He and David Kempe. Robust influence maximization. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 885–894, New York, NY, USA, 2016. ACM.
- [Hwang *et al.*, 2013] Hark-Chin Hwang, Hyun-Soo Ahn, and Philip Kaminsky. Basis paths and a polynomial algorithm for the multistage production-capacitated lot-sizing problem. *Operations Research*, 61(2):469–482, 2013.
- [Kempe and Mahdian, 2008] David Kempe and Mohammad Mahdian. A cascade model for externalities in sponsored search. In *International Workshop on Internet and Network Economics*, pages 585–596. Springer, 2008.
- [Lahaie *et al.*, 2007] Sébastien Lahaie, David M Pennock, Amin Saberi, and Rakesh V Vohra. Sponsored search auctions. *Algorithmic game theory*, pages 699–716, 2007.
- [Lai *et al.*, 2010] Guoming Lai, François Margot, and Nicola Secomandi. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. *Operations research*, 58(3):564–582, 2010.
- [Lindsay *et al.*, 2009] Robert Taaffe Lindsay, Thomas Carriero, and Yun-Fang Juan. Measuring impact of online advertising campaigns, May 26 2009. US Patent App. 12/472,318.
- [McMahan and Gordon, 2003] H Brendan McMahan and Geoffrey J Gordon. Planning in cost-paired markov decision process games. In *NIPS Workshop: Planning for the Real-World*, volume 3, 2003.
- [McMahan *et al.*, 2003] H Brendan McMahan, Geoffrey J Gordon, and Avrim Blum. Planning in the presence of cost functions controlled by an adversary. In *ICML*, pages 536–543, 2003.
- [Varian, 2007] Hal R Varian. Position auctions. *international Journal of industrial Organization*, 25(6):1163–1178, 2007.
- [Wang *et al.*, 2013] Taifeng Wang, Jiang Bian, Shusen Liu, Yuyu Zhang, and Tie-Yan Liu. Psychological advertising: exploring user psychology for click prediction in sponsored search. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 563–571. ACM, 2013.
- [White *et al.*, 2014] Ryen W White, Ahmed Hassan, Adish Singla, and Eric Horvitz. From devices to people: Attribution of search activity in multi-user settings. In *Proceedings of the 23rd international conference on World wide web*, pages 431–442. ACM, 2014.
- [Xiong *et al.*, 2012] Chenyan Xiong, Taifeng Wang, Wenkui Ding, Yidong Shen, and Tie-Yan Liu. Relational click prediction for sponsored search. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 493–502. ACM, 2012.