Joint Optimization of Bid and Budget Allocation in Sponsored Search*

Weinan Zhang
Shanghai Jiao Tong University
Shanghai, 200240, P. R. China
wnzhang@apex.sjtu.edu.cn

Ying Zhang
Nankai University
Tianjin, 300071, P. R. China
zhangying@dbis.nankai.edu.cn

Bin Gao
Microsoft Research Asia
Beijing, 100080, P. R. China
bingao@microsoft.com

Yong Yu
Shanghai Jiao Tong University
Shanghai, 200240, P. R. China
yyu@apex.sjtu.edu.cn

Xiaojie Yuan
Nankai University
Tianjin, 300071, P. R. China
yuanxj@nankai.edu.cn

Tie-Yan Liu
Microsoft Research Asia
Beijing, 100080, P. R. China
tyliu@microsoft.com

ABSTRACT

This paper is concerned with the joint allocation of bid price and campaign budget in sponsored search. In this application, an advertiser can create a number of campaigns and set a budget for each of them. In a campaign, he/she can further create several ad groups with bid keywords and bid prices. Data analysis shows that many advertisers are dealing with a very large number of campaigns, bid keywords, and bid prices at the same time, which poses a great challenge to the optimality of their campaign management. As a result, the budgets of some campaigns might be too low to achieve the desired performance goals while those of some other campaigns might be wasted; the bid prices for some keywords may be too low to win competitive auctions while those of some other keywords may be unnecessarily high. In this paper, we propose a novel algorithm to automatically address this issue. In particular, we model the problem as a constrained optimization problem, which maximizes the expected advertiser revenue subject to the constraints of the total budget of the advertiser and the ranges of bid price change. By solving this optimization problem, we can obtain an optimal budget allocation plan as well as an optimal bid price setting. Our simulation results based on the sponsored search log of a commercial search engine have shown that by employing the proposed method, we can effectively improve the performances of the advertisers while at the same time we also see an increase in the revenue of the search engine. In addition, the results indicate that this method is robust to the second-order effects caused by the bid fluctuations from other advertisers.

Categories and Subject Descriptors

H.3.5 [**Information Systems**]: Information Storage and Retrieval - Online Information Services; J.0 [**Computer Applications**]: General

General Terms

Algorithms, Experimentation

Keywords

Budget allocation, Bid optimization, Sponsored search

1. INTRODUCTION

Sponsored search is a popular format of online advertising and is also the main revenue source for search engine companies. In sponsored search, a list of ads is displayed along with the organic search results in response to a given query. Sponsored search results are produced by a different mechanism from that of organic search, though they are displayed simultaneously and have similar appearance. Generally, the organic search results are mainly generated based on the relevance of each web page to the query, while the sponsored search results are generated based on an auction process [20, 2, 13].

An advertiser can create a number of campaigns under an account. In each campaign, he/she sets a campaign budget, builds several groups of ad copies (creatives), and bids on some keywords with their match types[1] for each ad group. Each keyword is an auction entry that is supposed to be triggered by some user queries. When a user submits a query, the search engine will first retrieve the most relevant ads as candidates according to a matching function between the bid keywords and the query. Then these candidate ads will participate in an auction, and some ads (e.g., with the largest expected revenue for the search engine) will win and be displayed on the search result page [13]. If an ad is clicked by the user, the corresponding advertiser will be charged by the search engine. Usually, the charged amount is determined by the generalized second price (GSP) [11] auction mechanism, which means that the advertiser's cost of a click depends on the bid price of the next ad in the ranking list of the auction. When a campaign runs out of budget, it will not be

---

[1]The match type might be exact match or broad match.

permitted to participate in any auctions until the budget is increased or the next budget period starts. For example, if the campaign budget is set at a monthly basis, the campaign will be re-involved in the auctions next month.

As can be seen above, besides creating ad groups and selecting bid keywords, an advertiser should also carefully consider two important problems as follows:

a) Bid Price Setting. Different keywords correspond to different opportunities (e.g., search volumes) and different degrees of competition. In many cases, the bids for some keywords are too low to win the auctions while the bids for some other keywords are unnecessarily high. Usually, the optimal bid price setting is very difficult for an individual advertiser to reach since he/she does not have the access to the related information (e.g., the bids of other advertisers) and his/her competitors are also adjusting their bid prices dynamically.

b) Campaign Budget Allocation. Similar to the case of keywords, different campaigns also have different opportunities and competition. As a result, under an account, some campaigns may run out of budget very quickly, some campaigns consume their budgets quite slowly, and the budgets for other campaigns may never be used at all. This will constrain the overall effectiveness for the advertiser to utilize his/her budget.

Both of the above issues are critical for advertisers, however, most advertisers have not been doing well in them according to our statistics (see Section 3). This is because many advertisers are managing hundreds of campaigns and tens of thousands of keywords, which makes it very difficult for them to manually tune the campaign budget allocation and keyword bid prices. There are some third-party tools to help the advertisers tune their bids; some advertisers even build tools themselves to manage the bids automatically. However, the market information they can get is still very limited, which restricts the effectiveness of the tools. In the research community, there have also been some attempts on preforming the task automatically (see Section 2). However, these works are still not sufficient to satisfy the practical requirements. For example, many works on keyword bid price optimization only consider the bid price when ranking the ads, but do not take relevance and position bias into consideration. For another example, although people have investigated keyword bid optimization, to the best of our knowledge, there is no work on campaign budget allocation yet in the literature.

In this paper, we propose a novel method to address the aforementioned issues. In particular, we propose jointly optimizing the campaign budget allocation and bid price setting. That is, for a given advertiser account with multiple campaigns and an account-level budget, we try to find the optimal allocation of the account-level budget into each campaign, and to set the optimal price[2,3] for each bid keyword in the campaign simultaneously. Here we focus on the joint

---

[2]Note that not all clicks are created equal. An advertiser will give different bids for different keywords, for they have different value per click (VPC). In this work, we estimate the VPC of a keyword based on the idea proposed in [4, 20] and use the VPC as the upper bound of the bid price.

[3]Usually an advertiser observes the ad campaign performance and takes reactions to change the bid so as to approach the campaign goal. In this work, we find the optimal bid by maximizing the campaign goal directly, so the optimal bid is not a randomized value.

optimization instead of optimizing campaign budget allocation and keyword bid setting separately due to the following reason. Suppose for some campaign, there are many high-utility keywords (in other words, these keywords contain a lot of opportunities of advertising). In order to achieve significant performance regarding these keywords, one has to put a lot of money on them. However, if we cannot increase the budget for this campaign, we will miss a lot of these opportunities.

We formulate the problem as a constrained optimization, which takes the campaign budgets and the keyword bid prices as variables and finds their optimal values by maximizing the advertiser revenue, with the constraint of the account-level budget. To efficiently solve the optimization problem, we employ the sequential quadratic programming method. Simulation results on the sponsored search log from a commercial search engine show that the proposed technology can effectively help advertisers improve their campaign performance under several metrics like click number, cost per click, and advertiser revenue. At the same time, we can also help the search engine obtain increased revenue. In addition, the proposed method is robust to the second-order effects caused by the advertisers' dynamical bid changes.

To sum up, the contributions of our work are listed as below. (i) We performed a comprehensive data study on the effectiveness of current campaign budget allocation and keyword bid price setting in sponsored search, and pointed out the importance and necessity of jointly optimizing them. (ii) We proposed a novel method for jointly optimizing bid and budget allocation. As far as we know, this is the first work on campaign budget allocation in the literature of sponsored search, and it is also the first work on consider keyword price setting and campaign budget allocation simultaneously.

2. RELATED WORK

As mentioned in the introduction, we focus on the joint optimization of campaign budget allocation and keyword bid price setting in this paper. That is, given an advertiser account with multiple campaigns and an account-level budget, we determine the optimal allocation of the account-level budget into each campaign, and set the optimal bid price for each bid keyword in the campaign. As far as we know, there is no work in the literature solving exactly the same problem. Instead, there is only some related work on keyword bid price optimization. We will briefly review such work in Section 2.1. Besides, there is some other work on budget allocation across different keywords [3], different search engines [8], or different online adverting markets [21], whose problem definitions are totally different with what we concern about in this paper.

2.1 Keyword Bid Price Setting

Chakrabarty *et al* [7] defined the weight and profit of each bid keyword, and proposed a knapsack based algorithm to find the optimal price setting that can maximize the advertiser revenue in sponsored search. The algorithm considered both single-slot and multi-slot auctions. Kitts *et al* [16] proposed a revenue optimization model based on the marketing factors that were related to ad slot positions, in order to solve the problem of keyword bid price setting. In [12, 17], a *budget optimization* problem was defined, in which the target is to find an optimal bid price setting to maximize the campaign performance under a given campaign budget.

In particular, Feldman *et al* [12] defined the cost and click functions considering the position of an ad and the average click-through rate (CTR) on the position, and then developed a landscape algorithm to solve the optimization problem in approximation. Muthukrishnan *et al* [17] extended the above algorithm using a stochastic model, in which each keyword has a click distribution instead of an exactly known click number, in the context of single-slot auction.

In addition, the following works also discuss the problem of keyword bid price optimization, but in different scenarios from the above ones. Dar *et al* [10] studied how to improve the performance of broad match of bid keywords for a given query. They defined the weight and utility function of each bid keyword, and proposed a flow graph based algorithm to work out the optimal setting for keyword bid prices. Broder *et al* [6] pointed out that different matched queries for a bid keyword in broad match had different utilities according to their relevance to the bid keyword, and thus they should have different bid prices. They proposed a statistical approach to generate the corresponding bid prices. Borgs *et al* [5] studied the problem of keyword bid price setting given the budget on each keyword, and proposed an method to optimize the advertiser revenue across all keywords. More parts of the work are discussing the perturbation and convergence in their model.

3. DATA ANALYSIS ON SPONSORED SEARCH

In this section, we report our data analysis on sponsored search. We have used two kinds of data obtained from a mainstream search engine in our study: the auction log that records the detailed auction processes and the advertiser database that includes the bid keywords, bid prices, and the budget for each campaign. The data was collected in half a month, which contains over ten billion auctions and more than one hundred thousand advertiser accounts.

3.1 Campaign Budget

There can be several campaigns under the same advertiser account. In general, each campaign contains a set of ads (or ad groups) with the same campaign goal. Each campaign is assigned a budget, indicating the expected expense in a period of time (e.g., one month). In practice, due to the differences in the advertising settings and the market dynamics, it is quite common that some of the campaigns run out of budget while the other campaigns under the same account do not. We call this kind of accounts *partially-running-out-of-budget accounts* (or *p-accounts* for short). Here we give some statistics about the p-accounts in Table 1.

Table 1: Statistic of partially-running-out-budget accounts.

Number proportion	2.6%
Revenue proportion	31.6%
Potential revenue proportion	164.5%
Average campaign budget use ratio	45.1%
Total budget use ratio	11.5%
Average campaign number	15
Max campaign number	2,423
Average keyword number	10,735
Max keyword number	1,818,285

From Table 1 we have the following observations. (i) Although the number of the p-accounts is relatively small (2.6%), their contribution to the search engine revenue is considerably large (31.6%). It is clear that the individual contribution of each p-account is much larger than that of other accounts. (ii) The potential revenue of the p-accounts[4] is as much as 164.5% of the total revenue of sponsored search. This shows that improving p-accounts can result in a significant impact on the entire advertising system. (iii) The average campaign-level budget use ratio of the p-accounts is about 45.1%, and the total budget use ratio of the p-accounts is 11.5%. The potential of further increasing the budget use ratio for the p-accounts is higher than for the other accounts. This is because the p-accounts have at least one campaign (but not all campaigns) running out of budget, thus we can further improve their budget use ratio by simply reallocating the residual budget to that campaign(s). However, for other accounts this strategy will not help. (iv) For a p-account, the average campaign number is 15 and the maximum campaign number is 2,423. These large numbers suggest that it is difficult to manually adjust the campaign budget.

3.2 Bid Price

Bid price setting is also a non-trivial task. For a p-account, the average number of bid keywords is 10,735 and the maximum is 1,818,285. It is clearly infeasible to manually set keyword bid prices for the p-accounts. Instead, automatic keyword bid price setting is desired. The tools from the third-party usually cannot have adequate market information to aid the bid tuning. For example, we may need to consider the following information. (i) First, according to the commonly-used auction mechanisms in commercial search engines, higher bid prices will increase the probability of winning more auctions and obtaining higher ranking positions. Figure 1 shows the percentage of ads that will get at least one position up if we increase their corresponding keyword bid prices to a certain degree. In particular, 57.36% ads will get position up if their keyword bid prices are increased by 10%, and 78.84% ads will get position up if their keyword bid prices are increased by 40%. (ii) Second, higher ranking positions usually mean more attention and clicks [15], according to Figure 2, which shows the relative CTR[5] of the top ad slots in the commercial search engine. As a result, if an advertiser increases the bid price, the ad will have more opportunities to be clicked and the campaign goal of the advertiser will be better realized.

However, given the constraint of campaign budget, usually we cannot increase the bid prices for all the keywords. Instead, we have to balance between different keywords. Note that the utility of keywords are different from each other, and the return of lifting prices on some low-utility keywords might be very small. Therefore, we should consider the keyword utility when adjusting the keyword bid prices. For those keywords with low utility, the best choice may be to decrease their bid prices and instead put more money on the high utility keywords.

3.3 Joint Optimization

The data statistics reported in the previous section show that both campaign budget allocation and bid price setting are important to advertisers, but they are non-trivial to op-

---

[4]The potential revenue is calculated by summing up all the unused budgets of the p-accounts.

[5]The relative CTR is normalized by dividing the CTR on each position by the CTR on the first position.
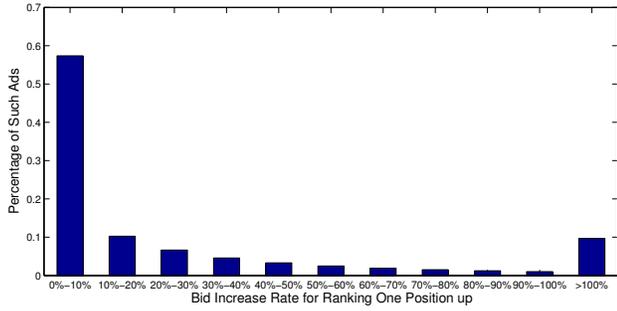
Figure 1: The proportion of ad volume ranking one position up with the increased bid price.
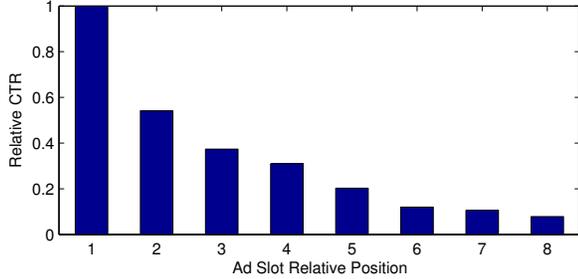

Figure 2: The relative CTR of each ad slot.

timize. If we want to jointly optimize them, the task may become even more difficult. However, we would like to point out that it is necessary to perform the joint optimization.

On one hand, if we only perform bid price optimization, each campaign will be optimized independently. For campaigns that have run out of budget, the help from bid price optimization will be limited because the potential capacity of these campaigns is constrained. For campaigns that have a lot of unused budget, bid price optimization can help achieve better performance, but it is hard to take full use of the budget due to the constraint of the value per click (VPC) of the bid keywords. On the other hand, if we only perform budget optimization, the unused budget will be reallocated to the campaigns that have already or tend to run out of budget, and thus their performance will be improved. However, the campaigns (no matter they run out of budget or not) will lose the opportunity to further improve their performance since their current bid price settings might not be optimal. Therefore, we had better consider both bid price optimization and campaign budget allocation simultaneously. In this way, the unused budget can be moved to the appropriate campaigns and can be effectively used on the best keyword candidates.

Note that in terms of optimizing the performances of advertisers, other efforts such as ad copy improvement, bid keyword selection, behavior and demography targeting, and landing page optimization are also important and effective. However, we will not discuss them as they are not in the scope of this paper.

4. JOINT OPTIMIZATION OF BID AND BUDGET ALLOCATION

In this section, we introduce the proposed method for joint optimization of bid price setting and campaign budget allocation. The key idea is to maximize the account-level ad-

vertiser revenue subject to the constraints of account-level budget. To better illustrate this idea, we first give some necessary notations including the definition of winning price interval, which serves as a basis of the following discussions. Then we adopt a probabilistic model to approximate the probability of winning a certain ad position given a bid price. After that, we define an optimization problem based on the probability model, and convert the problem to be a sequential quadratic programming problem. By solving the problem, we can get the optimal solution to campaign budget allocation and bid price setting.

4.1 Notations and Definitions

In the joint optimization of bid price setting and campaign budget allocation the input is an advertiser account $A = \{C_1, C_2, \ldots, C_m\}$, where m is the number of campaigns under account A and $C_i (i = 1, 2, \ldots, m)$ is the i-th campaign. For simplicity, we will not distinguish ad group and ad in the following discussions.[6] Accordingly, we can denote campaign C_i as $C_i = \{g_i^{(0)}, D_i, K_i\}$ $(i = 1, \ldots, m)$, where $g_i^{(0)}$ denotes the original periodical (e.g., monthly) budget set by the advertiser, D_i denotes the set of ads, and K_i denotes the set of bid keywords[7] in campaign C_i, respectively.

The ad set D_i can be written as $D_i = \{d_{i,1}, d_{i,2}, \ldots, d_{i,l_i}\}$, where l_i is the number of ads in campaign C_i and $d_{i,s}$ ($s = 1, 2, \ldots, l_i$) denotes the s-th ad in the campaign. The bid keyword set K_i can be written as $K_i = \{(k_{i,1}, b_{i,1}^{(0)}, v_{i,1}), (k_{i,2}, b_{i,2}^{(0)}, v_{i,2}), \ldots, (k_{i,n_i}, b_{i,n_i}^{(0)}, v_{i,n_i})\}$ where n_i is the number of bid keywords[8] in campaign C_i, $k_{i,t}$ ($t = 1, 2, \ldots, n_i$) denotes the t-th bid keyword, $b_{i,t}^{(0)}$ ($t = 1, 2, \ldots, n_i$) denotes the original bid price for $k_{i,t}$, and $v_{i,t}$ ($t = 1, 2, \ldots, n_i$) denotes the VPC that $k_{i,t}$ brings. We approximately estimate the VPC based on the idea proposed in [4, 20], and regard it as the upper bound of $b_{i,t}^{(0)}$ [2, 11]. In addition, we denote the minimum reserve price as ϵ_b, which serves as the lower bound of the valid bids.

In sponsored search, the advertiser can associate several keywords to his/her ad. When a query is issued, an auction might be triggered. If one of the associated keywords of an ad matches the query by the corresponding match functions to the match types of the keywords, the ad (together with the matched keyword) will be involved in the auction. Therefore, the candidates in the auction is actually a tuple as $\omega_{i,s,t} = (d_{i,s}, (k_{i,t}, b_{i,t}^{(0)}, v_{i,t}))$, where $s = 1, 2, \ldots, l_i$ and $t = 1, 2, \ldots, n_i$. We call the tuple *order item*. For ease of reference, for an order item ω, we also use $(\cdot)_w$ to denote the attribute associate with it, such as its keyword k_ω, bid price b_ω, and VPC v_ω.

We use Φ to denote the maximum number of ad slots in each search result page of the sponsored search system. Suppose the ads are ranked in the auction according to their rank scores, then we have the following definitions.

DEFINITION 1 (WINNING SCORE). *For an auction θ, its winning score at position ρ_ϕ (denoted by $\mu_{\phi,\theta}$, $\phi = 1, 2, \ldots, \Phi$)*

---

[6]In practice, the most relevant ad in an ad group will participate in the auction.

[7]A keyword may have different match types and different bid prices accordingly. For simplicity, we regard them as different keywords.

[8]The same keyword with different match types are regarded as different keywords in K_i.

is the least rank score that can make an order item get the ϕ-th ad slot ρ_ϕ in the auction. Let $\mu_{0,\theta} = +\infty$, then we have $\mu_{0,\theta} \geq \mu_{1,\theta} \geq \ldots \geq \mu_{\Phi,\theta}$.

DEFINITION 2 (WINNING SCORE INTERVAL). *For an auction θ, its winning score interval at position ρ_ϕ is $[\mu_{\phi,\theta}, \mu_{\phi-1,\theta}]$ $(\phi = 1, 2, \ldots, \Phi)$, which is the range of the rank score that can make an order item exactly get the ϕ-th ad slot ρ_ϕ in the auction.*

Mainstream search engines use the product of the bid price and the quality score as the rank score in their auctions [13]. Suppose the quality score of an order item ω in an auction θ is $r_{\omega,\theta}$, which can be calculated based on a group of features such as the query-ad similarity, semantic similarity, taxonomy, user query time, user query location and so on [14, 15, 19]. As indicated by the subscript, the quality score $r_{\omega,\theta}$ of an order item can be different in different auctions, due to some contextual information related to time, location, and user for the triggering query [14]. Usually such a quality score indicates the probability that an ad will be clicked after it is noticed by users. In this context, we have the following definitions of winning price and winning price interval.

DEFINITION 3 (WINNING PRICE). *Given an order item ω in an auction θ with its quality score $r_{\omega,\theta}$, its winning price at position ρ_ϕ (denoted by $\beta_{\omega,\phi,\theta}$, $\phi = 1, 2, \ldots, \Phi$) is $\frac{\mu_{\phi,\theta}}{r_{\omega,\theta}}$, which is the least bid price that can make ω get the ϕ-th ad slot ρ_ϕ in the auction. Let $\beta_{\omega,0,\theta} = +\infty$, and we have $\beta_{\omega,0,\theta} \geq \beta_{\omega,1,\theta} \geq \ldots \geq \beta_{\omega,\Phi,\theta}$.*

DEFINITION 4 (WINNING PRICE INTERVAL). *Given an order item ω in an auction θ with its quality score $r_{\omega,\theta}$, its winning price interval at position ρ_ϕ is $[\beta_{\omega,\phi,\theta}, \beta_{\omega,\phi-1,\theta}]$ $(\phi = 1, 2, \ldots, \Phi)$, which is the range of the bid price that can make ω exactly get the ϕ-th ad slot ρ_ϕ in the auction.*

4.2 Probabilistic Model for Ad Ranking

In order to compute the expected advertiser revenue, we need to get the probability for an order item ω with bid price b_ω to be ranked at position ρ_ϕ in the auctions.[9] More specifically, we define a probability distribution $P_\omega(b_\omega)$ as,

$$P_\omega(b_\omega) = (p_\omega(\rho_1|b_\omega), p_\omega(\rho_2|b_\omega), \ldots, p_\omega(\rho_\Phi|b_\omega), p_\omega(\rho_{\Phi+1}|b_\omega)),$$

where $p_\omega(\rho_\phi|b_\omega)$ $(\phi = 1, 2, \ldots, \Phi)$ denotes the probability for ω to be ranked in slot ρ_ϕ when its bid price is b_ω, and $p_\omega(\rho_{\Phi+1}|b_\omega)$ denotes the probability for ω to lose the auction (i.e., to be ranked lower than ρ_Φ). It is clear that,

$$\sum_{\phi=1}^{\Phi+1} p_\omega(\rho_\phi|b_\omega) = 1.$$

To get the above probability distribution, we apply the Bayes theorem to each element of it.

$$p_\omega(\rho_\phi|b_\omega) = \frac{p_\omega(b_\omega|\rho_\phi)p_\omega(\rho_\phi)}{\sum_{j=1}^{\Phi+1} p_\omega(b_\omega|\rho_j)p_\omega(\rho_j)}$$

Here $p_\omega(\rho_\phi)$ is the probability of any ad being displayed at position ρ_ϕ in the auctions that ω participates in, which can be approximately obtained by simple counting in the historical auction log. $p_\omega(b_\omega|\rho_\phi)$ is the probability of observing b_ω in the winning price interval at position ρ_ϕ in the auction

[9]Note that we compute this probability at the order item level but not for each individual auction, mainly because of the concern of data sparseness.

θ that ω participates in. A straightforward way is also to obtain this value by simple counting in the historical auction log. That is, for each auction θ of ω ($\theta = 1, 2, \ldots, \Theta_\omega$, where Θ_ω denotes the number of auctions ω participates in.), we calculate the winning price interval $[\beta_{\omega,\phi,\theta}, \beta_{\omega,\phi-1,\theta}]$ for position ρ_ϕ from the auction log. If $b_\omega \in [\beta_{\omega,\phi,\theta}, \beta_{\omega,\phi-1,\theta}]$, we say that there is an observation of price b_ω. However, the problem with this approach is that we need to go through the entire log for every possible value of b_ω, which will be too costly when we performing the optimization. Therefore we propose a new approach as described below that can be much more efficient without revisiting the entire auction logs during the optimization process.
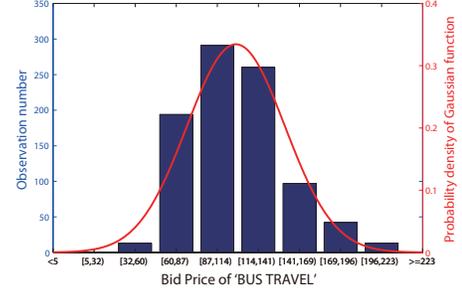

Figure 3: A case of Gaussian fitting for the bound of the winning price intervals.

For all auctions of ω, we can calculate their winning price intervals at position ρ_ϕ. As mentioned above, the lower bound and upper bound of the winning price interval are actually in fluctuation in different auctions. For simplicity, we use the following Gaussian distributions[10] to model the fluctuation of the bounds.

$$q_\omega^{(L)}(x|\rho_\phi) = \frac{1}{\sqrt{2\pi\sigma_{\omega,\phi}^2}}e^{-(x-\bar{\beta}_{\omega,\phi})^2/2\sigma_{\omega,\phi}^2}$$
$$q_\omega^{(U)}(y|\rho_\phi) = \frac{1}{\sqrt{2\pi\sigma_{\omega,\phi-1}^2}}e^{-(y-\bar{\beta}_{\omega,\phi-1})^2/2\sigma_{\omega,\phi-1}^2}$$

Here x and y are the random variables for the lower bound and upper bound of the winning price interval of ω at position ρ_ϕ, and the superscripts L and U stand for lower bound and upper bound respectively. In addition, $\bar{\beta}_{\omega,\phi}$ and $\sigma_{\omega,\phi}$ are the mean and standard deviation for the lower bounds of the winning price intervals at position ρ_ϕ for all auctions of ω. That is,

$$\bar{\beta}_{\omega,\phi} = \frac{1}{\Theta_\omega}\sum_{\theta=1}^{\Theta_\omega}\beta_{\omega,\phi,\theta}$$
$$\sigma_{\omega,\phi}^2 = \frac{1}{\Theta_\omega}\sum_{\theta=1}^{\Theta_\omega}(\beta_{\omega,\phi,\theta} - \frac{1}{\Theta_\omega}\sum_{\theta=1}^{\Theta_\omega}\beta_{\omega,\phi,\theta})^2$$

Similarly, $\bar{\beta}_{\omega,\phi-1}$ and $\sigma_{\omega,\phi-1}$ are the mean and standard deviation for the upper bounds of the winning price intervals at position ρ_ϕ for all auctions of ω.

[10]We have tested several possible distributions including Gaussian, Beta, and Gamma distributions, and found Gaussian is one of the best choices. Due to space limit, we will not show the parameter fitting for the distributions, but we can show an running example like Figure 3. As this is an approximation for the real data, we will ignore the negative values from the Gaussian distribution, just like what we usually do when we assume the heights of a group of people are sampled from a Gaussian distribution.

Thus $p_\omega(b_\omega|\rho_\phi)$ can be computed as below.

$$
\begin{aligned}
p_\omega(b_\omega|\rho_\phi) &= p_\omega(x \le b_\omega < y|\rho_\phi) \\
&= p_\omega(x \le b_\omega|\rho_\phi)p_\omega(b_\omega < y|\rho_\phi) \\
&= \int_{-\infty}^{b_\omega} q_\omega^{(L)}(x|\rho_\phi)dx \int_{b_\omega}^{\infty} q_\omega^{(U)}(y|\rho_\phi)dy \\
&= \emptyset(\tfrac{b_\omega - \bar{\beta}_{\omega,\phi}}{\sigma_{\omega,\phi}})(1 - \emptyset(\tfrac{b_\omega - \bar{\beta}_{\omega,\phi-1}}{\sigma_{\omega,\phi-1}})) \\
&\quad (\phi = 2, \ldots, \Phi)
\end{aligned}
$$

Here $\emptyset(\cdot)$ represents the cumulative distribution function of the standard normal distribution. In particular, for the first ad slot ρ_1, the upper bound y is infinity. Hence $p_\omega(b_\omega < y|\rho_1) \equiv 1$, and then,

$$
p_\omega(b_\omega|\rho_1) = p_\omega(x \le b_\omega|\rho_1) = \emptyset(\tfrac{b_\omega - \bar{\beta}_{\omega,1}}{\sigma_{\omega,1}}).
$$

Similarly, for $\rho_{\Phi+1}$, the lower bound x is zero. Hence $p_\omega(x \le b_\omega|\rho_{\Phi+1}) \equiv 1$, and then,

$$
p_\omega(b_\omega|\rho_{\Phi+1}) = p_\omega(b_\omega < y|\rho_{\Phi+1}) = 1 - \emptyset(\tfrac{b_\omega - \bar{\beta}_{\omega,\Phi}}{\sigma_{\omega,\Phi}}).
$$

Figure 4 uses an example to explain the calculation of the probability $p_\omega(b_\omega|\rho_\phi)$. Suppose the bid of an observation is 28 (cents), then the probability $p_\omega(b_\omega|\rho_\phi)$ equals to the product of the area of the left shadow (probability of lower bound < 28) and the area of the right shadow (probability of upper bound > 28).
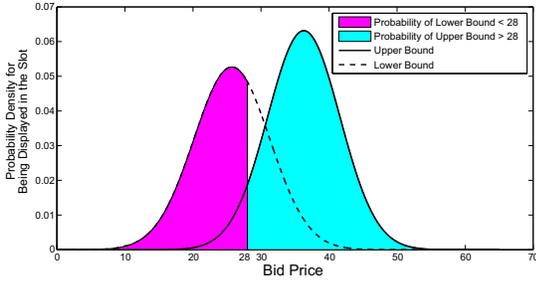

Figure 4: An example of the probability density distribution of the upper bound and lower bound of a winning price interval.

So far we have discussed our probabilistic model for ad ranking based on the ad auction log data. Compared with the conventional ranking models, the probabilistic model is smooth and easy to be directly optimized. Note that this model can be designed in other forms with different data formats in different scenarios.

4.3 Optimization Problem

Given the definition of the probabilistic model described in the previous subsection, we can define the expected advertiser revenue. To ease the illustration, we further introduce two notations as below. (i) τ_ϕ - the position bias at slot ρ_ϕ. As discussed in Section 3.2, the relative CTR (shown in Figure 2) indicates the probability of ads in the position ρ_ϕ being noticed. Further considering the definition of quality score $r_{\omega,\theta}$, the actual probability of an ad being clicked when ranked on the slot ρ_ϕ will be $\tau_\phi r_{\omega,\theta}$ [15]. (ii) $c_{\omega,\phi,\theta}$ - the cost for a click on ω in an auction θ where it is ranked on position ρ_ϕ. According to the GSP system, the cost can be calculated as $c_{\omega,\phi,\theta} = \frac{b_{\omega'} r_{\omega',\theta}}{r_{\omega,\theta}}$, where ω' is the order item that is ranked one slot lower than ω in the auction θ, and $b_{\omega'}$ is its bid price.

The objective of the optimization problem is to maximize the total revenue of the advertiser account, which reflects the final profit the advertiser can make from the sponsored search. The constraints are the bounds of the bid prices and the campaign budget.

For all the campaigns in an advertiser account, the total expected click number can be written as

$$
\sum_{i=1}^m \sum_{\omega \in C_i} \sum_{\theta=1}^{\Theta_\omega} \sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}.
$$

where the factor $\sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}$ is the probability of click on ω in one auction when the bid price is b_ω.

Considering the cost of click and the VPC of each bid keyword, we can get the expected advertiser revenue as follows,

$$
\sum_{i=1}^m \sum_{\omega \in C_i} \sum_{\theta=1}^{\Theta_\omega} \sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}(v_\omega - c_{\omega,\phi,\theta}).
$$

Given the above objective function, we can formulate the joint optimization as below, where g_i $(i = 1, 2, \ldots, m)$ and b_ω denote the variables of campaign budgets and keyword bid prices respectively. The minimum campaign budget is ϵ_g, for the advertiser might not like to entirely close a campaign by letting $g_i = 0$.

$$
\begin{aligned}
\max_{g_i,b_\omega} \quad & \sum_{i=1}^m \sum_{\omega \in C_i} \sum_{\theta=1}^{\Theta_\omega} \sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}(v_\omega - c_{\omega,\phi,\theta}) \\
s.t. \quad & \sum_i^m g_i = \sum_i^m g_i^{(0)} \\
& 0 \le \sum_{\omega \in C_i} \sum_{\theta=1}^{\Theta_\omega} \sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}c_{\omega,\phi,\theta} \le g_i \\
& (i = 1, 2, \ldots, m) \\
& \epsilon_g \le g_i \quad (i = 1, 2, \ldots, m) \\
& \epsilon_b \le b_\omega \le v_\omega \quad (\omega \in C_i, \ i = 1, 2, \ldots, m)
\end{aligned}
$$

4.4 Efficient Solution

The above optimization problem is a typical constrained optimization problem. It can be approximately solved by means of sequential quadratic programming (SQP) [9] in an efficient manner.

Suppose $\xi_1 = \{b_\omega\}$ $(\omega \in C_i, \ i = 1, 2, \ldots, m)$ denotes the vector of bid prices for all the order items in an advertiser account, and $\xi_2 = \{g_i\}$ $(i = 1, 2, \ldots, m)$ denotes the vector of campaign budgets. Then $\xi = (\xi_1^T, \xi_2^T)^T$ is the vector of all variables. We can rewrite the optimization problem as the following form.

$$
\begin{aligned}
\min_\xi \quad & f(\xi) \\
s.t. \quad & h_0(\xi) = 0 \\
& h_i(\xi) \ge 0 \quad (i = 1, 2, \ldots, 5)
\end{aligned}
$$

The functions $f(\xi)$ and $h_i(\xi)$ $(i = 0, 1, \ldots, 5)$ are written in the following forms,

$$
\begin{aligned}
f(\xi) &= -\sum_{i=1}^m \sum_{\omega \in C_i} \sum_{\theta=1}^{\Theta_\omega} \sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}(v_\omega - c_{\omega,\phi,\theta}) \\
h_0(\xi) &= \sum_i^m g_i - \sum_i^m g_i^{(0)} = 0 \\
h_1(\xi) &= \sum_{\omega \in C_i} \sum_{\theta=1}^{\Theta_\omega} \sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}c_{\omega,\phi,\theta} \ge 0 \\
h_2(\xi) &= g_i - \sum_{\omega \in C_i} \sum_{\theta=1}^{\Theta_\omega} \sum_{\phi=1}^{\Phi} p_\omega(\rho_\phi|b_\omega)\tau_\phi r_{\omega,\theta}c_{\omega,\phi,\theta} \ge 0 \\
& \quad (i = 1, 2, \ldots, m) \\
h_3(\xi) &= g_i - \epsilon_g \ge 0 \quad (i = 1, 2, \ldots, m) \\
h_4(\xi) &= b_\omega - \epsilon_b \ge 0 \quad (\omega \in C_i, \ i = 1, 2, \ldots, m) \\
h_5(\xi) &= v_\omega - b_\omega \ge 0 \quad (\omega \in C_i, \ i = 1, 2, \ldots, m)
\end{aligned}
$$

The above problem can be approximately solved by SQP. In SQP, a quadratic programming (QP) subproblem is solved in each iteration which is obtained by linearizing the constraints and approximating the following Lagrangian function $L(\xi, \eta)$ quadratically.

$$L(\xi, \eta) = f(\xi) - \sum_{i=0}^{5} \eta_i h_i(\xi)$$

Here $\eta = \{\eta_i\}$ is the Lagrangian multiplier. The optimization may start from any initial $\xi^{(0)}$. Suppose $\xi^{(j)}$ is the solution in the j-th iteration, $\eta^{(j)}$ is the corresponding Lagrangian multipliers, and $H^{(j)} = H(\xi^{(j)}, \eta^{(j)})$ is the Hessian of the Lagrangian function, then the following QP subproblem should be solved in the $(j+1)$-th iteration.

$$\begin{aligned}\min_z \quad & \tfrac{1}{2}z^T H^{(j)} z + \nabla L(\xi, \eta)^T z \\ s.t. \quad & \nabla h_0(\xi^{(j)})^T z + h_0(\xi^{(j)}) = 0 \\ & \nabla h_i(\xi^{(j)})^T z + h_i(\xi^{(j)}) \geq 0 \ (i = 1, \ldots, 5)\end{aligned}$$

Here ∇ denotes gradient calculus. If $z^{(j)}$ is the solution of the above QP subproblem and $w^{(j)}$ is the corresponding multiplier of this subproblem, then we use the following formulas to update the solution of the SQP problem.

$$\begin{aligned}\xi^{(j+1)} &= \xi^{(j)} + z^{(j)} \\ \eta^{(j+1)} &= w^{(j)}\end{aligned}$$

The above iterations lead to a local approximation of the solution. The algorithm can be stabilized by line search. By solving this SQP problem, we can get the optimal campaign budgets as well as the optimal bid price for each order item.

5. EXPERIMENTAL EVALUATION

In this section, we first introduce our experimental settings, including the datasets, baseline algorithms, and evaluation measures. Then we report the experimental results on our proposed algorithm and make analysis and discussions.

5.1 Experimental Settings

5.1.1 Datasets

The data used in our experiments came from a mainstream commercial search engine. There are basically two types of data: auction log and advertiser database. The auction log was collected in a month of 2011, containing over one billion auction events. The log was partitioned into two parts, one for training and the other for testing. The training data was used as historical data to get the probability model for ad rank and the empirical values used in our algorithms. The test data was used for simulation [1] in order to evaluate the performance of each algorithm. The advertiser database is a snapshot in the same month, which contains about 150 thousand sponsored search accounts.

We sampled 400 p-accounts to study in our experiments. Accordingly, over 1.2 million related auction events were extracted from the auction log. We also got the bid price for each keyword and the budget for each campaign in these accounts from the advertiser database.

5.1.2 Baseline Algorithms

The proposed algorithm jointly optimizes the bid price and campaign budget. In order to understand the benefit of performing joint optimization, we will also investigate how it works if we only optimize bid price or campaign budget. This leads to two natural simplified methods for our algorithm. We also compare our algorithm with two state-of-the-art algorithms for bid price optimization. In addition, using the original bid prices and campaign budget is the pure baseline. Furthermore, we adopt a bidding strategy to mimic the second order effect by the bid fluctuations, and test the robustness of our proposed method. Therefore, as a whole, we have the following seven algorithms in our experiments.

Original Bid Price and Campaign Budget (ORI). This algorithm uses the original keyword bid prices and campaign budgets in the experiment. This is the pure baseline in the experiments.

Joint Optimization of Bid and Budget (JO). This is our proposed algorithm. Note that we approximately calculate the VPC according to the idea in [4], i.e., we ran a simulation on the auction log to get the incremental cost per click (ICPC) [4] for each keyword and use it as the estimated value for the VPC of the keyword.

Bid Price Only (BID). This algorithm uses our proposed algorithm, but only optimizes the keyword bid prices, without changing the campaign budgets. By comparing this method with JO, we can see the impact of joint optimization. This method can also be used to directly compare with other algorithms for bid optimization.

Campaign Budget Only (BGT). This algorithm uses our proposed algorithm, but only optimizes the campaign budget, without changing the keyword bid prices. By comparing this method with JO, we can also see the impact of joint optimization.

Knapsack Problem (KS). This algorithm is the multi-slot auction model proposed in [7], which adapts the problem of keyword bid price setting into a multi-choice knapsack problem. We set the t in the model as one day. That is, we will change the bid price daily according to remaining budget of each campaign.

Market Optimal Bid (MOB). This algorithm refers to the optimal bid model proposed in [16]. This model is an optimization framework that maximizes the advertiser revenue. Position is also considered in this model. The difference between this model and the BID model lies in that MOB captures the best position with fixed bid prices in each position to approach the observation instead of using a probability distribution for positions.

Joint Optimization with Advertiser Modeling (JO-AM). In the real online traffic, the advertisers would dynamically tune their bid prices in response to the changes of their campaign performances, and thus the changes of the bid prices will lead to the changes of the action outputs. Therefore, the optimal budget allocation and bid prices calculated from the historical log might not be so effective as expected in the real online experiment. We adopt a bidding strategy [18] by the advertisers to mimic the dynamic bid changes, and check whether the proposed JO method can still be robust with the possible bid changes by the advertisers.

Note that we do not compare with the algorithms in [12, 17], for their objective is to maximize the expected click number. The number of clicks can be a good measure for the campaign performance, but it may not be a good objective to optimize for the advertisers. They may care more about the quality of clicks or the value they desire from each click. In other words, the revenue might be more meaningful for the advertisers to optimize. We do not compare with the algorithms in [5, 6, 10] either, for they are in different scenarios with our setting.

5.1.3 Evaluation Measures

We used the following evaluation measures in our experiments.

Ad Impressions. After simulation on the test data, we can get the impressions for an advertiser account.

Expected Clicks. We use the adPredictor model [14] to calculate the probability of click, and transform the impressions to expected ad clicks.[11] More expected clicks mean better performance of an advertiser.

Advertiser Revenue. Since there is a cost and a VPC associated with each click, the advertiser revenue for each click can be calculated as the difference between the VPC and the cost. By summing up the advertiser revenue for all the clicks obtained by an advertiser account, we can get the overall advertiser revenue.

Search Engine Revenue. While the above measures mainly describe the performances of advertisers, this measure corresponds to the performance of the search engine. Note that the search engine revenue does not equal the expense of the target advertiser accounts. If the keyword auction is highly competitive and there is no empty ad slot, the impression of a new ad might not increase the expected search engine revenue since there is another ad losing the auction.

5.2 Experimental Results

We study the performances of the algorithms ORI, BID, BGT, KS, MOB, JO, and JO-AM which all take the advertiser revenue as the objective function. We report the performance of these algorithms with respect to ad impressions, expected clicks, advertiser revenue, and search engine revenue. Note that we normalized the values under each evaluation metric by dividing them by the maximum value in the corresponding results, to protect the business secrets of the search engine.
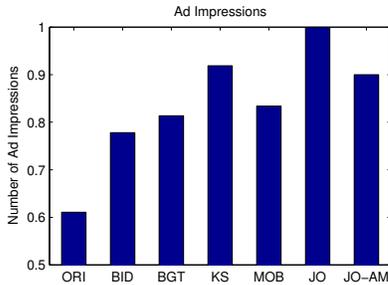
5.2.1 Ad Impressions

Figure 5: The average ad impressions among p-accounts.

The average ad impressions of each algorithm are shown in Figure 5. From the figure we have the following observations. (i) All the algorithms successfully achieve a lot more impressions as compared to ORI. This demonstrates that the p-accounts have large potential to improve their performance by tuning the bid prices or budget allocation. (ii) JO performs better than BID and BGT. Therefore, although BID and BGT can help improve the ad impressions, jointly optimizing both bid and budget will help achieve more. This shows the necessity of adopting our proposed method. (iii)

---

[11]The adPredictor model performs quite well on our data and has a predict error of less than 10%.

The performance of JO-AM is hurt by the advertiser dynamics on bid changes. However, it is still comparable with the second highest performer KS. It shows the proposed method is robust against the second-order effects from the advertiser dynamics.
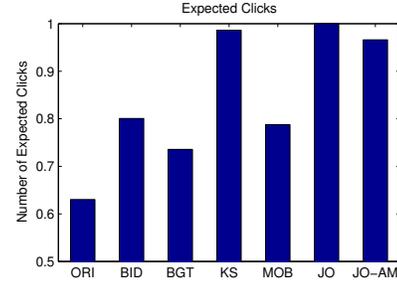
5.2.2 Expected Clicks

Figure 6: The average expected clicks among p-accounts.

Figure 6 depicts the number of expected clicks in the simulation period for the seven algorithms. From the figure we have the following observations. (i) BID, BGT, KS, MOB, JO and JO-AM all improve the click number as compared to ORI. This shows that using advertiser revenue as the optimization objective can also help improve the click number. (ii) BID performs better than MOB, indicating that our proposed ways of optimizing bid prices is more effective than some previous work. (iii) Further considering the budget allocation, JO performs the best. (iv) JO-AM is comparable with the second highest performer KS, again showing the robustness of the proposed method against the second-order effects.
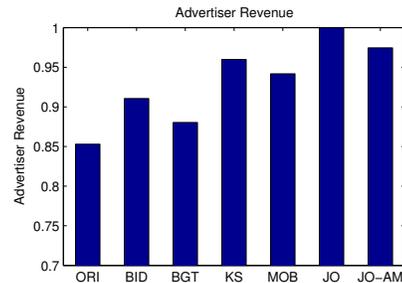
5.2.3 Advertiser Revenue

Figure 7: The average advertiser revenue among p-accounts.

In Figure 7, we evaluate the advertiser revenue, which is just the optimization objective of these algorithms. From the figure we have the following observations. (i) All the algorithms improve the advertiser revenue significantly. Though BID and BGT perform slightly worse than KS and MOB, by combining the advantages of BID and BGT, JO achieves better performance as compared to KS and MOB. It shows that the advertisers can benefit a lot from the joint optimization of bid price setting and budget allocation. (ii) JO-AM gets the second highest position in advertiser revenue. It even outperforms KS, showing that the proposed method is a good choice for the advertisers in practice when facing the bid dynamics.
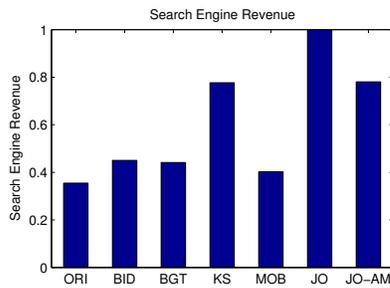
Figure 8: The average search engine revenue among p-accounts.

5.2.4 Search Engine Revenue

The performances in terms of search engine revenue are shown in Figure 8. From the figure, we can see that search engine can earn more money by using all the algorithms as compared to ORI. Among these algorithms, and JO still performs the best. Therefore, the joint optimization method can in return help the search engine improve the income.

To sum up all the aforementioned experimental results, we can get the following conclusions. (i) Bid price optimization (BID) can lead to high improvement on the ad impressions and expected clicks, and thus can largely increase both the advertiser revenue and the search engine revenue. (ii) Campaign budget optimization itself (BGT) seems not to perform as well as bid prices optimization (BID) in terms of the evaluation measures like expected clicks and advertiser revenue, however, it can ensure that only a very small number of advertisers will have worse performance than before after the optimization. (iii) By jointly optimizing both bid prices and budget allocation, JO perform the best among all the algorithms in all evaluation measures. This clearly demonstrates the effectiveness of our proposal for both the advertisers and the search engine.

6. CONCLUSION AND FUTURE WORK

In this paper, we studied the joint optimization of campaign budget allocation and bid price setting. For this purpose, we proposed a probability model for ad ranking, and proposed an integer optimization problem to maximize the expected advertiser revenue. We show that the problem can be approximately solved by means of alternately working on a binary integer optimization problem and another constraint optimization problem. Experimental results have shown that our proposed algorithm can increase the performance of advertisers in terms of several evaluation measures, as well as the search engine revenue.

For future work, we plan to further study like the following problems. For example, we will refine our model to perform the optimization for multiple accounts simultaneously. For this purpose, we need to consider the inter-competition between accounts, which will make the optimization problem much harder to solve. Furthermore, we want to apply the proposed method in online traffic to test its performance.

7. REFERENCES

[1] S. Acharya, P. Krishnamurthy, K. Deshp, T. W. Yan, and C. chao Chang. A simulation framework for evaluating designs for sponsored search markets. In *WWW 2007 Workshop on Sponsored Search Auctions*. ACM, 2007.

[2] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM conference on Electronic commerce*, 2006.

[3] N. Archak, V. Mirrokni, and M. Muthukrishnan. Budget optimization for online advertising campaigns with carryover effects. In *Sixth Ad Auctions Workshop.*, 2010.

[4] S. Atheya and D. Nekipelov. A structural model of sponsored search advertising auctions. Available at `http://groups.haas.berkeley.edu/marketing/sics/pdf_2010/paper_athey.pdf`.

[5] C. Borgs, J. Chayes, N. Immorlica, K. Jain, O. Etesami, and M. Mahdian. Dynamics of bid optimization in online advertisement auctions. In *Proceedings of the 16th international conference on World Wide Web*, WWW '07, pages 531–540, New York, NY, USA, 2007. ACM.

[6] A. Broder, E. Gabrilovich, V. Josifovski, G. Mavromatis, and A. Smola. Bid generation for advanced match in sponsored search. In *Proceedings of WSDM'11*.

[7] D. Chakrabarty, Y. Zhou, and R. Lukose. Budget constrained bidding in keyword auctions and online knapsack problems. In *Proceedings of the 16th World Wide Web Conference*, 2007.

[8] L. Chen and Y. Li. Allocating budget across portals in search engine advertising. In *Management Science and Engineering, 2009. ICMSE 2009. International Conference on*, pages 679 –685, 2009.

[9] R. W. Cottle, J.-S. Pang, and R. E. Stone. *The linear complementarity problem.* Academic Press, Inc., 1992.

[10] E. Dar, V. Mirrokni, S. Muthukrishnan, Y. Mansour, and U. Nadav. Bid optimization for broad match ad auctions. In *Proceedings of WWW'09*, 2009.

[11] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, 2007.

[12] J. Feldman, S. Muthukrishnan, M. Pĺćl, and C. Stein. Budget optimization in search-based advertising auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, 2007.

[13] J. Feng, H. Bhargava, and D. Pennock. Comparison of allocation rules for paid placement advertising in search engines. In *Proceedings of the 5th international conference on Electronic commerce*, pages 294–299. ACM ICEC, 2003.

[14] T. Graepel, J. Q. Candela, T. Borchert, and R. Herbrich. Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine. In *Proceedings of the 27th International Conference on Machine Learning*. ACM ICML, 2010.

[15] D. Hillard, S. Schroedl, and E. Manavoglu. Improving ad relevance in sponsored search. In *Proceedings of WSDM'10*.

[16] B. Kitts and B. Leblanc. Optimal bidding on keyword auctions. *Electronic Markets*, 14(3):186–201, 2004.

[17] S. Muthukrishnan, M. Pĺćl, and Z. Svitkina. Stochastic models for budget optimization in search-based advertising. In *Proceedings of WWW'07*, 2007.

[18] C. Nittala and Y. Narahari. Optimal equilibrium bidding strategies for budget constrained bidders in sponsored search auctions. In *Operational Research: An International Journal*. Springer, 2011.

[19] F. Radlinski, A. Broder, P. Ciccolo, E. Gabrilovich, V. Josifovski, and L. Riedel. Optimizing relevance and revenue in ad search: A query substitution approach. In *proceeding of SIGIR'08*, 2008.

[20] H. R. Varian. Position auctions. In *International Journal of Industrial Organization*, 2006.

[21] N. Virk. Is your search marketing budget allocation right? Available at `http://www.internetseer.com/services/article.xtp?id=37080`.