

Location-Based Sponsored Search Advertising

George Trimponias¹, Ilaria Bartolini², and Dimitris Papadias¹

¹ Department of Computer Science and Engineering,
Hong Kong University of Science and Technology
{trimponias,dimitris}@cse.ust.hk

² Department of Computer Science and Engineering,
University of Bologna
i.bartolini@unibo.it

Abstract. The proliferation of powerful mobile devices with built-in navigational capabilities and the adoption in most metropolitan areas of fast wireless communication protocols have recently created unprecedented opportunities for location-based advertising. In this work, we provide models and investigate the market for location-based sponsored search, where advertisers pay the search engine to be displayed in slots alongside the search engine's main results. We distinguish between three cases: (1) advertisers only declare bids but not budgets, (2) advertisers declare budgets but not bids, and (3) advertisers declare both bids and budgets. We first cast these problems as game theoretical market problems, and we subsequently attempt to identify the equilibrium strategies for the corresponding games.

Keywords: Location-based advertising, Game theory, Nash equilibrium.

1 Introduction

The growing popularity of powerful and ubiquitous mobile devices has recently created an immense potential for location-based advertising (LBA) [5]. Smartphone use is rapidly increasing in all parts of the world; in the US only, its penetration is currently approaching 50% of all mobile subscribers, while around 60 percent of the new phones in 2011 were smartphones¹. This development has certainly been facilitated by the adoption of broadband wireless protocols, e.g., 3G/4G networks, and the prevalence of Wi-Fi hotspots. Moreover, modern mobile devices possess built-in navigational functionalities using variations of sophisticated technologies such as triangulation, GPS, and cell-ID [5]. Advertisers can utilize this positional information to send advertising material to relevant consumers, which has in turn created an exciting market for LBA with companies such as AdMob (acquired by Google) and Quattro Wireless (acquired by Apple) leading the charge.

Location-based advertising, especially in its mobile form, is poised for tremendous growth because of its special characteristics [1]. First, it enables personalization:

¹ See <http://www.informationweek.com/news/mobility/business/231602163>

a mobile device is associated with the identity of the user so the advertising material can be individually tailored. For example, users can state their preferences, or even specify the kind of advertising messages they are interested in. Second, it is context-aware, i.e., the advertising messages can take into account the context such as time and location. Third, mobile devices are portable and allow instant access: users carry their device most of the time, and advertisers can target interesting consumers any time of the day. Finally, mobile advertising can be interactive since it is possible to engage the user to discussions with the advertiser; this can also serve as a means of market research. As a result of the aforementioned reasons, marketers can reach their audience of interest in a much more targeted, personal and interactive manner, and thus increase their advertising campaign's success.

On the other hand, currently the most profitable and thriving business model for online advertising is *sponsored search advertising*; Google's total revenue alone in fiscal year 2010 was \$29.3 billion and mainly came from advertising². Sponsored search consists of three parties [12]: (i) *users* pose keyword queries with the goal of receiving relevant material; (ii) *advertisers* aim at promoting their product or service through a properly designed ad, and target relevant users by declaring to the search engine a set of keywords that capture their interest; (iii) the *search engine* mediates between users and advertisers, and facilitates their interaction. As several advertisers may match a given user query, an *auction* is run by the search engine every time a user poses a query to determine the winners as well as the *price per click*. Concretely, each advertiser declares to the engine a priori its *bid* for a given keyword, so the auction assigns ad slots to advertisers based on their bids.

In this work, inspired both by the success of sponsored search advertising and the immense potential for LBA, we study the promising area of *location-based sponsored search advertising*. In particular, we examine how the spatial component can be incorporated into the current sponsored search models, and investigate algorithms for selling advertising opportunities to advertisers. Similar to prior literature on conventional sponsored search advertising [8], in order to model the advertisers we distinguish between the three following cases: (1) the advertisers declare a maximum amount of money that they are willing to pay per click, but are not bounded by a total daily budget, (2) the advertisers have a maximum daily budget at their disposal, but do not have an upper bound on the amount of money that they are willing to pay per click, and (3) the advertisers have both a total daily budget and a maximum amount of money that they are willing to pay per click. We will explicitly show how the introduction of the spatial component affects the underlying sponsored search auction in each of the cases above by using tools and techniques from game theory.

The rest of the paper is organized as follows. Section 2 surveys related work. Section 3 provides a general model for location-based sponsored search. Sections 4-6 investigate three interesting settings for location-based sponsored search: (1) advertisers declare only bids (Section 4), (2) advertisers declare only budgets (Section 5), and (3) advertisers declare both bids and budgets (Section 6). Using tools from game theory, we analyze the three different cases and provide the Nash equilibrium strategies when possible. Finally, Section 7 concludes the paper providing interesting directions for future research.

² See <http://investor.google.com/financial/2010/tables.html>

2 Related Work

2.1 Location-Based Advertising

Location-based advertising (LBA) involves delivering advertising material to users based on their location. It follows two different modes of operation [3]: *pull-based* (also termed *query-based*) and *push-based*. The former provides advertising information only upon specific request by the user, e.g., when a car driver asks for the nearest gas stations. The latter delivers marketing material to users within a specified geographical area, without their explicit request; for instance, shops in a mall seeking to promote their new product may target all shoppers by delivering the corresponding advertising information. Moreover, push-based advertising is further divided into two types: *opt-in*, where users receive relevant advertising material by determining in advance the kind of ads they are interested in, and *opt-out*, where users receive marketing messages until they explicitly declare they do not wish to receive any further material.

LBA presents immense opportunities for higher return on investment compared to other traditional advertising avenues because it enables contextually relevant advertising [5]. Moreover, the ability to instantaneously connect users to places or resources of interest in their immediate vicinity can offer an unrivaled user experience and satisfaction. Interestingly, LBA can also serve as a subtle tool for market research: consumers constantly provide information about their behavior through their mobile activity, which can be subsequently used to increase the effectiveness of a marketing campaign. Despite its obvious benefits, consumer's privacy is still a major cause of concern for LBA. Advertisers need to be very clear about how they utilize, process, and store user information; data breaches, for instance, can be especially detrimental to the advertiser's reputation and long-term success, since they can reveal personal information. A second major concern stems from the intrusive nature of some forms of LBA, in particular push-based that occurs without the user's explicit request. Among the two modes, opt-out is associated with a higher intrusion risk and is thus used more rarely; opt-in, in contrast, is permission-based advertising and may be used to effectively rule out unsolicited marketing messages (i.e., spamming) [18].

2.2 Sponsored Search Advertising

Sponsored search advertising is the most profitable form of online advertising. It constitutes a large and rapidly growing source of revenue for search engines. Currently, the most prominent players in the sponsored search market are Google's AdWords [22], and Bing Ads [21]. In sponsored search advertising, advertisers place properly designed ads to promote their product or service. They target interesting users by declaring to the search engine a list of keywords that a relevant user may search for. For each keyword, they additionally specify their *maximum cost per click* (*maximum CPC*), also known as maximum bid, which corresponds to the maximum amount of money they are willing to spend to appear on the results page for a given keyword. Note that bidding takes place continuously. Moreover, advertisers may be

limited by budget constraints, so they may declare a maximum daily budget as well. Every time a user enters a query, a limited number of *paid* (also, *sponsored*) links appears on top or to the right side of the *unpaid* (also, *organic* or *algorithmic*) search results. In order to determine the winning advertisers as well as the price they need to pay, an auction occurs in an automated fashion. In practice, large search engines also compute a *quality score* (QS) for every advertiser which measures how relevant the keyword, ad text and landing page are to a user.

Concretely, sponsored search advertising consists of three stages. (i) *Ad retrieval* returns all ads that are relevant to the user's query, and is usually performed by sophisticated machine learning algorithms. An ad's relevance is measured by several metrics, such as ad-query lexical and semantic similarity. To match an ad against a query, the search engine needs to take into account all ad information, including the bid phrase it is associated with, its title and description, the landing page it leads to, its URL, etc. Moreover, query substitution and query rewriting are frequently used to find relevant ads. As the ad pool may consist of millions of ads, efficient indexing techniques have been proposed to improve the performance of the first stage. (ii) After retrieving relevant ads, the search engine performs *ad ranking*. Ads are sorted in decreasing order of their rank, where the ad rank is determined by both the bid placed by the advertiser on the keyword, and the quality of the ad. The ad with the highest rank appears in the first position, and so on down the page, until all slots have been filled. Google AdWords³, for instance, defines the ad rank as the product $CPC \cdot QS$. (iii) The last stage is *ad pricing* through properly designed auctions to determine the price per click that the advertiser will be charged whenever the user clicks on their ad (pay per click model). The natural method would be to make bidders pay what they bid (i.e., *generalized first-price auction*), but that leads to several instabilities and inefficiencies. Instead, all large search engines currently employ a *generalized second-price auction* (*GSP*) [7][19]. A GSP auction charges an advertiser the minimum amount required to maintain their ad's position in search results, plus a tiny increment. For instance, suppose that ranking is based on Google's AdRank and that K slots are available, and are numbered $1, \dots, K$, starting from the top and going down. Moreover, let the advertiser i at position i have a maximum bid b_i and a quality score QS_i . In GSP, the price for a click for advertiser i is determined by the advertiser $i+1$, and given by $b_{i+1} \cdot QS_{i+1} / QS_i$, which is the minimum that i would have to bid to attain its position. Note that in this pricing scheme, a bidder's payment does not take into consideration its own bid. Also, prices per click can be computed in linear time in the number of advertisers $O(N)$ for a fixed number of slots K .

Despite its prevalence as the standard auction format, GSP is not *truthful* (also known as *incentive-compatible*): advertisers have no incentive to declare their true valuations to the search engine. Stated equivalently, reporting the true bids may not constitute a Nash equilibrium [7]. As a result, advertisers may devote considerable resources to manipulate their bids, potentially paying less attention to ad quality and other campaign goals. Interestingly, we can alleviate this shortcoming by altering the

³ See <https://adwords.google.com/support/aw/bin/answer.py?hl=en&answer=6111>

payment scheme: instead of paying the minimum amount of money required to win its position, an advertiser is requested to pay an amount of money equal to the externalities that it imposes on the others, i.e., the decreases in the valuations of other bidders because of its presence. This payment scheme yields the Vickrey-Clarke-Groves (VCG) auction, named after William Vickrey [20], Edward H. Clarke [4], and Theodore Groves [11]. Contrary to GSP, VCG gives bidders an incentive to bid their true value, and is *socially optimal*, i.e., the bidder with the highest valuation acquires the slot at the highest position, the bidder with the second-highest valuation receives the slot at the second-highest position, etc. Note that GSP rather than VCG is used in practice, even though the latter would (at least theoretically) diminish incentives for strategizing and facilitate the advertisers' task. We believe that the introduction of the ad quality score QS has also played a role in the wide adoption of GSP. Indeed, ad quality scores are now an integral part of both the ranking and pricing protocols; even if advertisers manipulate their bids, it is very difficult to game the system as they have no control over the ad quality scores.

3 Models for Location-Based Sponsored Search

Assume N advertisers and K slots $1, \dots, K$, where 1 is the top slot, 2 the second, and so on. There is ample evidence in the literature that higher slots are associated with higher revenues. There are numerous ways to model this; perhaps the easiest way is to characterize each slot with the click through rate, which denotes the probability that a user will actually click on an ad that is placed in that slot. In this work, we follow the same approach by assuming that whenever an ad is displayed in slot l , $1 \leq l \leq K$, it has a probability c_l , $0 \leq c_l \leq 1$, of being clicked. To incorporate the fact that higher slots are more valuable, we further assume that $c_l > c_{l'}$ whenever $l < l'$. Whether a user clicks on an ad or not depends on numerous factors including the other ads (*ad externalities*), but for the sake of simplicity we do not consider them here; i.e., an ad located at slot l is clicked with probability c_l independent of the rest of the slots [15]. To keep the model simple, we also do not consider quality scores for advertisers.

A salient feature of our work is that advertisers value users according to their location. To model this, we assume that the space is partitioned with a grid of L cells. There are (in expectation) M_j queries per day in cell j , which can be estimated based on historical data. Advertisers have different valuations for the different grid cells. For instance, a typical advertiser would have high valuations for cells nearby and lower valuations for more distant cells. We denote with $w_{i,j}$ the valuation of advertiser i per click inside cell j . Calculating the valuation is a difficult marketing/operational research problem, beyond the scope of our work. Finally, advertiser i may be bounded by a maximum daily budget B_i . We can assume that the advertisers are only aware of their own budget and valuations, which they declare to the search engine. Besides the budgets and valuations per click, the search engine has knowledge of relevant statistical information such as number of queries per cell, or percentage of total clicks that a slot receives, etc. We consider that advertisers are interested in exactly the same (unique) keyword; how keyword interactions affect our market is an interesting research topic in its own right, and can be explored in future work.

Finally, note that the valuation of an advertiser for a given cell is fixed for all points inside the cell. The grid granularity involves an inherent trade-off between valuation expressivity and search engine revenue. On the one hand, small cells allow advertisers to better capture their cells of interest, as opposed to coarse grid granularities that would force an advertiser to declare interest for the entire cell even if they were interested in just a small part. On the other hand, small cells may take a serious toll on the search engine's revenue because the expected number of advertisers expressing interest in a given cell decreases as the grid granularity becomes finer. In the worst case scenario, a cell could attract interest from just a single advertiser and would yield poor income for the search engine. For instance, assume a cell that attracts only one advertiser. The commonly used GSP protocol when advertisers only declare bids will then assign any query inside the cell to that advertiser for a price equal to 0, compromising the search engine's revenue goals. Determining the proper grid granularity is thus a critical factor of success for location-based sponsored search.

In the following sections, we discuss location-based sponsored search advertising focusing on three cases [8], depending on the advertiser input and constraints. In the first *bids-only* case, each advertiser i is not bounded by a daily budget, i.e., $B_i = \infty$, and is willing to pay up to its valuation per click $w_{i,j}$ in cell j , i.e., its maximum bid per click for cell j is equal to $w_{i,j}$. In the second *budgets-only* case, each advertiser i is bounded by a finite daily budget B_i , but is indifferent to the price per click that it is asked to pay, i.e., its maximum bid per click for any cell is unbounded. Finally, in the third *bids-and-budgets* case, each advertiser i is bounded by a finite daily budget B_i , and is willing to pay up to its valuation per click $w_{i,j}$ in cell j . We first cast all three cases as game theoretical problems, and we subsequently attempt to identify the equilibrium strategies for the corresponding games. Table 1 illustrates common symbols used in the rest of the paper.

Table 1. Frequent symbols

Symbol	Meaning
N, L, K	Number of advertisers, grid cells, and ad slots
B_i	Total daily budget of advertiser i
$B_{i,j}$	Part of total budget B_i that advertiser i allocates into cell j
$w_{i,j}$	Valuation per click of advertiser i for cell j
M_i	Expected number of queries per day for cell j
c_l	Probability that an ad located at slot l will get clicked
$U_{i,j} / U_i$	Total daily utility of advertiser i from cell j / from all cells
C_i	Set of cells where advertiser i has the highest valuation per click
\tilde{U}	Upper concave envelope of U
Y_j	Sum of budgets that have been allocated to cell j by all advertisers
p_j	Price per click in cell j (for case 3)
s_j	Permutation of advertisers such that $w_{s_j(i),j}$ is decreasing in i
$w_j^{(2)}$	Second-highest valuation per click in cell j (for case 3)

4 Bids-Only Case

Bids-only is the simplest case, as it constitutes a straightforward generalization of the conventional sponsored search framework. Whenever the search engine receives a query from a cell, it runs an auction where each advertiser is assumed to bid an amount of money equal to its valuation per click for that particular cell. We can utilize any auction format, such as the GSP or the VCG (see Section 2.2), to determine the K winners that will fill the slots, as well as the prices per click that they have to pay. These two auctions have been extensively studied in the literature, and as mentioned earlier truthfully reporting the bids constitutes a Nash equilibrium for the VCG auction, but is in general not an equilibrium for the GSP procedure. Note that the actual number of queries per cell does not matter: every single time a user issues a query, a new auction will play out in an automated way; cells with high workload will simply involve more auctions compared to cells with lower traffic.

Next, we discuss some useful metrics focusing on the GSP framework. A very interesting notion in auction theory concerns an advertiser's payoff, which refers to the net utility the advertiser receives from being advertised. In sponsored search auctions, an advertiser's payoff is defined in terms of a *quasi-linear model*: the payoff per click is equal to the valuation/utility v_i per click the advertiser i gets minus the price per click p_i that it must pay, i.e., $v_i - p_i = v_i - b_{i+1}$, where the bids b_i are in descending order. We can also define the expected payoff per day if we know the average number of queries per day M . Since slot l receives a c_l percentage of the total clicks, the expected payoff per day for slot l is $M \cdot c_l \cdot (v_i - b_{i+1})$. Non-winning advertisers get a payoff equal to 0. Finally, we define the search engine's (cumulative) profit per click simply as $p_1 + \dots + p_K = b_2 + \dots + b_{K+1}$ (similarly for the cumulative profit per day). It is now straightforward to generalize the above metrics in the location-based framework. For instance, the expected payoff per day that advertiser i gets in cell j if it gets assigned to slot l is $M_j \cdot c_l \cdot (w_{i,j} - w_{i+1,j})$ (where $w_{i,j}$ are in decreasing order for cell j).

5 Budgets-Only Case

In case 2, advertiser i declares a maximum daily budget B_i , as well as its valuations per click $w_{i,j}$ for each cell j . As opposed to case 1, where $w_{i,j}$ is the maximum amount that i is willing to pay per click, for case 2 the payments per click are bounded only by B_i , and the cell valuations are used just to determine the relative importance of cells. For simplicity, we initially consider a single slot ($K=1$) with probability of being clicked $c_l=1$, and deal with more slots later. Since case 2 only involves budget constraints, it is convenient to assume a *Fisher market model* [2]: under this model, money does not bear any intrinsic value and every advertiser is willing to burn their entire budget; note that this is different from the quasi-linear model that we assumed in Section 4. Our goal is to assign to every advertiser a probability that their ad will be displayed in any given cell, whenever a user in that cell issues a relevant query. Therefore, no auction takes place and we do not have a winner selection and price determination phase.

Based on the declared budgets and cell valuations, the system computes for each cell the probability that any advertiser will be chosen as a response to user query.

In conventional sponsored search with only one slot, the optimal solution to this problem displays an advertiser with a probability that is proportional to its budget [8][14][13]. So, the advertiser with the highest budget has the highest probability of being displayed, which is equal to its budget divided over the sum of all budgets; and so on for the rest of the advertisers. This rule is called *proportional sharing*, and, intuitively, it guarantees fairness.

In location-based sponsored search, on the other hand, advertisers declare a total daily budget for all cells, but do not specify how this budget should be allocated among the various cells. Now, assume that the advertiser (somehow) decides how to allocate its budget into the cells, so that each cell has a non-negative budget and the sum of budgets over all cells does not exceed the advertiser’s total budget. If such an allocation were known for every advertiser, then we could simply apply the proportional sharing rule: in a given cell, an advertiser is advertised with a probability proportional to its budget for this specific cell. But then a natural question arises: how should every advertiser allocate its budget?

To answer this question, we will resort to the proportional-share allocation market by Feldman et al. [10]. Concretely, assume a budget allocation for advertiser i such that it assigns $B_{i,j} \geq 0$ to cell j and the sum of its allocations over all cells does not exceed B_i . The probability that i will be displayed in cell j is $B_{i,j}/Y_j$, where Y_j is the sum of budgets that have been allocated to cell j by all advertisers. The utility for advertiser i in cell j is then $U_{i,j} = w_{i,j} \cdot M_j \cdot B_{i,j} / Y_j$, since it gets a value $w_{i,j}$ for every query in j when displayed with a probability $B_{i,j}/Y_j$, and there are M_j queries in total in cell j . We assume *additive* utilities, so i ’s total utility U_i is the sum of its utilities $U_{i,j}$ over all cells: $U_i = \sum_j w_{i,j} M_j \frac{B_{i,j}}{\sum_k B_{k,j}}$. Note that the payoff of advertiser i is equal to its utility, because of the Fisher market model assumption (money bears no intrinsic value to the advertiser).

A given set of budget allocations will give rise to different corresponding utilities for the advertisers. So, how should advertisers allocate their budget? Ideally, we would like to allocate every individual budget in a way that maximizes the advertiser’s utility. But since the advertisers compete against each other, one’s gain may translate into another’s loss. To come up with proper budget allocations, we thus utilize the notion of Nash equilibrium. The set of agents consists of the advertisers, while the strategy space for advertiser i is the convex, bounded and closed set $\{(B_{i,1}, \dots, B_{i,L}) \mid B_{i,j} \geq 0 \text{ and } \sum_{1 \leq j \leq L} B_{i,j} = B_i\}$, i.e., the set of all valid budget allocations. From advertiser’s i perspective, a best response strategy is simply a strategy $\mathbf{B}_i = (B_{i,1}, \dots, B_{i,L})$ that maximizes its utility given the other advertisers’ budget allocations, i.e., the solution to the following optimization problem:

$$\begin{aligned} &\text{maximize } U_i(B_{1,1}, \dots, B_{1,L}, \dots, B_{N,1}, \dots, B_{N,L}) \\ &\text{subject to } \sum_{1 \leq j \leq L} B_{i,j} = B_i \text{ and } B_{i,j} \geq 0. \end{aligned}$$

A Nash equilibrium then corresponds to the stable state where no advertiser has an incentive to deviate from their strategy given that the other advertisers stick to their strategy as well. Stated equivalently, every advertiser plays a best response strategy to the rest of the advertisers. Formally, a set of valid strategies $\mathbf{B}_1^*, \dots, \mathbf{B}_N^*$ form a Nash equilibrium if for any other valid strategy $\mathbf{B}_i, 1 \leq i \leq N$, we have:

$$U_i(\mathbf{B}_1^*, \dots, \mathbf{B}_i^*, \dots, \mathbf{B}_N^*) \geq U_i(\mathbf{B}_1^*, \dots, \mathbf{B}_i, \dots, \mathbf{B}_N^*).$$

It turns out that the above game does not always accept a Nash equilibrium. To demonstrate this, consider two advertisers 1 and 2 with budgets $B_1, B_2 > 0$, and two cells 1 and 2 with expected number of queries per day $M_1, M_2 > 0$. Advertiser 1 is interested in both cells, whereas advertiser 2 is only interested in cell 1. For player 2, the best strategy would obviously be to allocate its entire budget B_2 to cell 1 to gain the maximum possible proportion of ads. For advertiser 1, on the other hand, the best strategy would be to allocate a tiny amount $\epsilon > 0$ to cell 2 (and win all advertising opportunities in 2) and spend the rest $B_1 - \epsilon$ on cell 1 (and maximize its share in cell 1 as well). Unfortunately, there is no optimal value of ϵ , since (1) it must be positive to ensure 1 gets all ads in cell 2, and (2) as small as possible so that 1 wins the largest possible share in cell 1. Alternatively, consider the simpler case with a single player 1 with $B_1 > 0$, interested in a single cell 1 with $M_1 > 0$. As before, player 1 should allocate the smallest possible positive $\epsilon > 0$ on cell 1, but such an ϵ does not exist.

The root of the non-existence of a Nash Equilibrium in the examples above is due to the discontinuity of the utility functions at point 0. This problem can be circumvented in two different ways. First, we can enforce a *reserve price*, which is defined as the minimum possible price that an advertiser must pay per click. Indeed, a reserve price means that the advertiser cannot buy any click with an arbitrarily small budget, and the discontinuity at 0 ceases to exist. Second, we can restrict our attention to so called *strongly competitive* games [10], i.e., games where for a given cell there are at least two advertisers with positive valuations. Indeed, strong competition implies that if only one advertiser would allocate a tiny budget on a given cell, then any other advertiser who has non-zero valuation for that cell will have an incentive to also allocate (a tiny) budget in that cell to guarantee a percentage of ads [10].

Computing the Nash equilibrium is the next source of concern. There are 2 classes of algorithms for this purpose. The *best response algorithm* iteratively updates the budget allocations of every player to reflect the other players' current strategies. This algorithm simulates the best response dynamics of the game and thus has a very natural interpretation. We describe it in Figure 1; the interested reader is referred to [10] for further details. Note that its time complexity is dominated by the sorting procedure, so it is $O(N \log N)$. Theoretically, the best-response dynamics does not necessarily converge to a Nash equilibrium of the game; nevertheless, in practice the algorithm performs very well.

Repeat for each advertiser $i, 1 \leq i \leq N$

1. Sort the cells according to $\frac{w_{i,j}}{\sum_{i' \neq i} B_{i',j}}$ in decreasing order, where $1 \leq j \leq L$

2. Compute the largest k such that

$$\frac{\sqrt{w_{i,k} M_k \sum_{i' \neq i} B_{i',k}}}{\sum_{m=1}^k w_{i,m} M_m \sum_{i' \neq i} B_{i',m}} \left(B_i + \sum_{m=1}^k \sum_{i' \neq i} B_{i',m} \right) - \sum_{i' \neq i} B_{i',k} > 0$$

3. Set $B_{i,j}=0$ for $j > k$, and for $1 \leq j \leq k$ set

$$B_{i,j} = \frac{\sqrt{w_{i,j} M_j \sum_{i' \neq i} B_{i',j}}}{\sum_{m=1}^k w_{i,m} M_m \sum_{i' \neq i} B_{i',m}} \left(B_i + \sum_{m=1}^k \sum_{i' \neq i} B_{i',m} \right) - \sum_{i' \neq i} B_{i',j}$$

until convergence.

Fig. 1. Best-response dynamics for $K=1$ and strong competition

The alternative to best response dynamics is the *local greedy adjustment method* [10]. Under this algorithm, we first identify for every advertiser the two cells that provide the highest and lowest marginal utilities. We then move a fixed small amount of money from the cell with the lowest marginal utility to the cell with the highest one. This strategy aims to adjust the budget allocations so that the marginal values in each cell are the same. For concave utility functions (as ours), this is a sufficient condition for an optimal allocation. However, the method suffers from lower convergence rates.

As a last remark, note that contrary to case 1, the actual query distribution is now important. To understand why, assume the advertiser has a high valuation for cell 1 and a low valuation for cell 2. However, a small number of queries are issued in cell 1, whereas several queries are issued in cell 2. In bids-only sponsored search, a separate auction occurs every time a query is issued, so the advertiser can bid high for cell 1 and low for cell 2; since there are far more queries in cell 2 the advertiser will obviously participate in the auction for cell 2 far more times, but has no reason not to bid high for cell 1 and low for cell 2. In the budgets-only setting, however, query distribution has a profound effect on the budget allocation. In the above example, the advertiser may have to allocate a large part of its budget to cell 2 just because there are far too many queries in that cell.

Multiple Slots: We can generalize the above discussion in the case of several slots, by assuming for simplicity that a given advertiser may appear with non-zero probability in more than one slots (as opposed to the bids-only case). This assumption is necessary for a straightforward and simple generalization. Indeed, the idea is that every advertiser allocates part of its budget into all slots in every cell. The utility that advertiser i extracts from being advertised at slot l in cell j is $w_{i,j}c_l M_j \frac{B_{i,j,l}}{\sum_k B_{k,j,l}}$, where $B_{i,j,l}$ the amount of money that i allocates in slot l of cell j . Similar to before, we can assume additive utilities, so that the total utility of advertiser i the sum of its utilities over all slots and over all cells. Using the above techniques, we can then find budget allocations that constitute a Nash equilibrium.

6 Bids-and-Budgets Case

In this setting, advertiser i declares a maximum daily budget B_i as before, but contrary to case 2, i is now not willing to spend more than $w_{i,j}$ per click in cell j . Stated equivalently, the price that advertiser i pays per click in a given cell j cannot exceed its declared valuation $w_{i,j}$ for that cell. The valuations thus act as maximum bids per click, and we also refer to case 3 as bids-and-budgets case. We only deal with the case of a single slot, i.e., $K=1$ with $c_j=1$, and we assume again that money bears no intrinsic value to the advertisers (Fisher market model). The case of several slots is more complex, and can be investigated in future work.

Before dealing with the location-based setting, we first explore how conventional sponsored search addresses the case where both budgets and maximum bids per click are declared. In particular, we will attempt to highlight how this setting is inherently more complex than the budgets-only case. We focus on cell j with M_j queries per day

and budget allocations in it $B_{1,j}, \dots, B_{N,j}$. First, assume that every advertiser receives a share of the total ads proportional to its budget. Then, the price per click would be equal to $p_j = (B_{1,j} + \dots + B_{N,j})/M_j$. As long as this quantity is not greater than all valuations per click $w_{1,j}, \dots, w_{N,j}$, no problem occurs. But if an advertiser i exists with $w_{i,j} < p_j$, this advertiser would not be willing to pay as much as p_j per click, so the proportional allocation framework of Section 5 cannot be directly applied. To alleviate this problem, we need to come up with a price p_j^* such that all advertisers who can afford that price have sufficient budgets to purchase all the advertising opportunities. Figure 2 presents the price-setting mechanism by Feldman et al. [9][8] that determines that price p_j^* . It is essentially a price-descending mechanism: the price keeps falling until p_j^* is reached. Moreover, it has the desired property of being truthful.

-
1. Assume w.l.o.g. that $w_{1,j} > w_{2,j} > \dots > w_{N,j} \geq 0$.
 2. Let k^* be the first bidder such that $w_{k^*+1,j} \leq \frac{\sum_{i=1}^{k^*} B_{i,j}}{M_j}$. Set price $p_j^* = \min \left\{ \frac{\sum_{i=1}^{k^*} B_{i,j}}{M_j}, w_{k^*,j} \right\}$.
 3. Allocate $B_{i,j}/p_j^*$ ads to each advertiser $i \leq k^* - 1$. Allocate $M_j - \sum_{i=1}^{k^*} B_{i,j}/p_j^*$ ads to advertiser k^* . Allocate 0 ads to the rest of the bidders.
-

Fig. 2. The price-setting mechanism in cell j for $K=1$ slot in the bids-and-budgets case

Now, recall that in the case where only budgets are available, the price per query in cell j would be equal to $p_j = \frac{\sum_{i=1}^N B_{i,j}}{M_j}$. Obviously, p_j is linear in its arguments $B_{i,j}$ ($1 \leq i \leq N$) and continuous. On the other hand, the price-setting mechanism in Figure 2 yields prices that are clearly more complex. First, we notice the price p_j for a given cell j will again be an argument of only the budget allocations for that cell $B_{1,j}, \dots, B_{N,j}$. However, it does not have the simple linear form as in the case of only budgets. To get a flavor of the price function, consider a setting with only 2 advertisers 1 and 2 with maximum bids $w_{1,j}$ and $w_{2,j}$ (with $w_{1,j} > w_{2,j}$) for cell j that has M_j queries per day. Figure 3 depicts how the price varies according to the budgets $B_{1,j}$ and $B_{2,j}$ that the advertisers allocate in cell j . In particular, if $B_{1,j} \geq M_j \cdot w_{2,j}$, then $k^*=1$ and the price is determined as the minimum of $B_{1,j}/M_j$ and $w_{1,j}$. When $B_{1,j} \geq M_j \cdot w_{1,j}$ then the price is equal to $w_{1,j}$ (region I), while when $B_{1,j} < M_j \cdot w_{1,j}$, the price is equal to $B_{1,j}/M_j$ (region II). On the other hand, when $B_{1,j} < M_j \cdot w_{2,j}$, then $k^*=2$ and the price is the minimum of $w_{2,j}$ and $(B_{1,j} + B_{2,j})/M_j$; for $B_{1,j} + B_{2,j} \geq M_j \cdot w_{2,j}$ the price is $w_{2,j}$ (region III), while for $B_{1,j} + B_{2,j} < M_j \cdot w_{2,j}$, the price is $(B_{1,j} + B_{2,j})/M_j$ (region IV). Inside a region, the price can be either constant or linear. We first observe that the price function is everywhere continuous; the boundaries of the regions are carefully chosen so that the price is continuous as we move from one region to the other. Note also that the price function for the price-setting mechanism is bounded: it achieves a minimum value of 0 at the origin (0,0), and it can never get larger than $w_{1,j}$. On the contrary, the price per click in the budgets-only case is unbounded: it can get arbitrarily large as the budgets that the advertisers allocate grow larger.

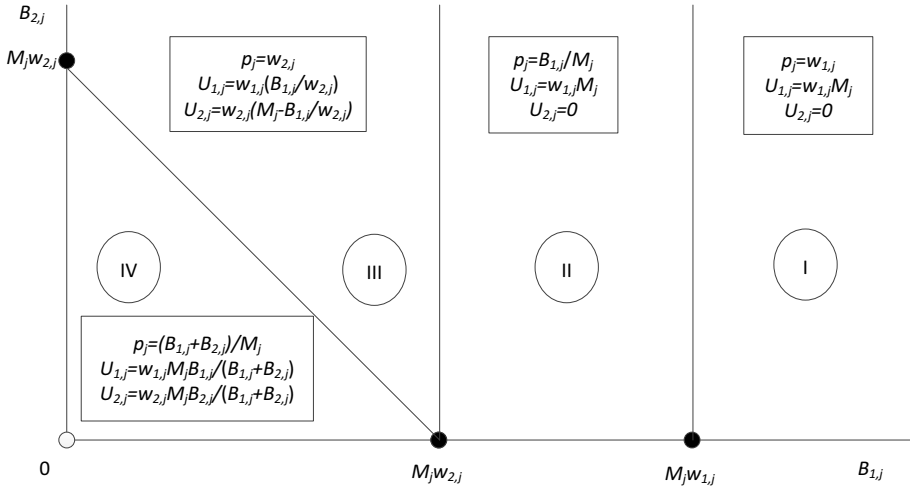


Fig. 3. Price p_j and utilities $U_{1,j}$ and $U_{2,j}$ in cell j when the number of advertisers is $N=2$

The above example captures some important properties of the price function in the case of both maximum bids and budgets. The price setting mechanism decomposes the budget space into N regions (one for each of the N possible k^*), and then further divides that region into two subregions: the price is constant inside one of them and linear in the other. In the following, we state several results. The proofs of all results are available in our technical report [17]. For every cell j , we consider the permutation s_j that reorders the bids in decreasing order, i.e., $w_{s_j(1),j} > w_{s_j(2),j} > \dots > w_{s_j(N),j} > 0$ for every cell j . Moreover, all budget allocations $B_{i,j}$ are non-negative: $B_{i,j} \geq 0$.

Lemma 1: The price function $p_j(B_{1,j}, \dots, B_{N,j})$ is continuous in $(B_{1,j}, \dots, B_{N,j})$.

Let’s now try to formalize the location-based setting where advertisers have valuations over the various cells. Similar to the previous case, we will be looking for a budget allocation $\mathbf{B}_i = (B_{i,1}, \dots, B_{i,L})$ for every advertiser i , $1 \leq i \leq N$. For a given allocation, denote $z_{i,j}$ the share of ads that advertiser i gets in cell j . Then, its utility from cell j is $w_{i,j} z_{i,j}$; its total utility from all cells simply is $U_i = \sum_j w_{i,j} z_{i,j}$. In order to compute the share $z_{i,j}$, we will exploit the price-setting mechanism, assuming that $w_{s_j(1),j} > w_{s_j(2),j} > \dots > w_{s_j(N),j} > 0$. If $w_{s_j(2),j} \leq B_{s_j(1),j} / M_j$, then $k_j^* = 1$ and the price p_j^* is $\min\{B_{s_j(1),j} / M_j, w_{s_j(1),j}\}$. On the other hand, if $w_{s_j(2),j} > B_{s_j(1),j} / M_j$, we continue by checking whether $w_{s_j(3),j} \leq (B_{s_j(1),j} + B_{s_j(2),j}) / M_j$. If the latter is true, then $k_j^* = 2$ and the price $p_j^* = \min\{(B_{s_j(1),j} + B_{s_j(2),j}) / M_j, w_{s_j(2),j}\}$. If it is false, we proceed in exactly the same way, until we come up with the proper k_j^* , and subsequently compute the price p_j^* . Figure 3 depicts the utility functions in cell j in the case of $N=2$ advertisers.

Next, we compare the above utility function with the simpler utility function in the case where only budgets are declared: $U_{i,j} = w_{i,j} M_j \frac{B_{i,j}}{\sum_k B_{k,j}}$. Clearly, the latter function is concave in $B_{i,j}$, gets a minimum value of 0 for $B_{i,j} = 0$, and asymptotically converges to

$w_{i,j} \cdot M_j$ as $B_{i,j}$ tends to infinity. In other words, the advertiser will get allocated all ads in cell j as its budget gets infinitely large, given that the other advertisers' budgets for this cell are fixed. But can we say something similar for the utility function in the more complex setting when both budgets and maximum bids per click are declared? It turns out that the answer to that question is negative: the utility function $U_{i,j}$ for the price-setting mechanism is not concave in $B_{i,j}$ anymore, as we later show (e.g., see Figure 4). However, $U_{i,j}$ is monotonically increasing in $B_{i,j}$:

Lemma 2: $U_{i,j}(B_{i,j})$ is monotonically increasing in $B_{i,j}$.

Since $U_{i,j}$ is monotonically increasing in $B_{i,j}$, it will also be *quasi-concave* in $B_{i,j}$. On the other hand, $U_i = \sum_{j=1}^L U_{i,j}$. It turns out that when $U_{i,j}$ are quasi-concave, but not concave in $B_{i,j}$, then their sum is not quasi-concave in $(B_{i,1}, \dots, B_{i,L})$ [6]. This is a worrisome result, in the sense that existence theorems for Nash equilibria usually assume concave or, at least, quasi-concave utility functions.

There are, however, two special cases where we can easily show that a Nash equilibrium exists. First, assume that the sum of the advertisers' budgets is sufficiently small, i.e., $\sum_{i=1}^N B_i \leq M_j w_{N,j}$, for every cell j . In this case, independent of the budget allocation, we have in any cell j that $\sum_{i=1}^N B_{i,j} \leq M_j w_{N,j}$, so the price setting mechanism will allocate to every advertiser a percentage of advertising opportunities proportional to the budget that they allocate in every cell. But this is identical to case 2, and it thus always admits a Nash equilibrium if 1) there is a reserve price, or 2) there is strong competition. Second, assume that every advertiser has sufficiently large budget, and that there is strong competition in every cell. For any advertiser i , consider the set of cells C_i where i has the highest valuation per click among all advertisers, i.e., $C_i = \{j | w_{i,j} = \max_{1 \leq i' \leq N} \{w_{i',j}\}\}$ (for some advertisers this set may be empty). For advertiser i , we then define the following budget allocation strategy: allocate 0 to cell j if $j \in C_i$, else allocate an amount of money equal to or greater than $M_j w_j^{(2)}$, where $w_j^{(2)} > 0$ the second highest valuation per click in cell j (it is positive because of the strong competition assumption). This is always possible if $B_i \geq \sum_{j \in C_i} M_j w_j^{(2)}$, for every advertiser i . It is easy to verify that the above sets of budget allocations correspond to Nash equilibria, since any advertiser cannot increase its utility by deviating to a different budget allocation. Indeed, with the previous budget allocation every advertiser i wins all ads for the cells that belong to C_i . Obviously, i cannot gain a higher utility by changing its budget allocation for cells $j \in C_i$. On the other hand, even if i allocates a positive budget in cells $j \notin C_i$, it will still gain 0 advertising opportunities, since the first advertiser has adequate budget and valuation to buy all ads in that cell. In fact, a Nash equilibrium in the case of sufficiently large budgets can be given by the following rule: in every cell the advertiser with the highest valuation per click pays a price per click equal to the valuation per click of the second highest advertiser, and wins all ads for that cell. But what we have just described is the GSP procedure. Stated equivalently, the GSP auction for sufficiently large budgets results in a Nash equilibrium.

We have thus observed how the bids-and-budgets case encompasses the simpler bids-only and budgets-only cases for sufficiently small or large budgets, respectively.

On the other hand, when only one advertiser has a positive valuation for a cell j , then using the same line of arguments as in Section 5 we can see that its utility function is discontinuous at 0, and the game accepts no Nash equilibrium. It is however possible to slightly modify the game in a way that makes the discontinuity at 0 disappear, similar to [10][13]. In this direction, we will introduce a *fictitious* advertiser $N+1$ who allocates a tiny budget $B_\epsilon > 0$ in every cell, but has an arbitrarily large valuation per click for every cell. We call the perturbed game with the additional player G . So, what is the impact of the additional player $N+1$ on the game structure? Essentially, the arbitrarily large valuation per click for every cell implies that advertiser $N+1$ will have the highest valuation per click in every cell and will thus be able to pay any price that the price mechanism sets. On the other hand, we set B_ϵ to be very small so that player $N+1$ has a negligible impact. Note that the introduction of the fictitious player serves the same purpose as the reserve price of the budgets-only case, namely to smooth out the utility function and tackle the discontinuity at 0.

In the general case, we are currently not aware whether game G always accepts a Nash equilibrium since each advertiser's utility function is not quasi-concave. Although we cannot answer whether a Nash equilibrium exists, we can nevertheless find a budget allocation such that the maximum utility that an advertiser can gain by deviating is known.

In this direction, we will consider the *upper concave envelope* $\tilde{U}_{i,j}$ of the utility $U_{i,j}$, for any advertiser i and any cell j . Formally, we will be looking for the infimum of all functions that are concave and are greater than or equal to $U_{i,j}$ for any $B_{i,j}$. This is, in general, not an easy task, but as we shall see the upper concave envelope for the utility functions that arise in the bids-and-budgets setting has a relatively simple form.

We focus on advertiser i and cell j , $1 \leq i \leq N$ and $1 \leq j \leq L$. Assume the rest of the advertisers' budgets for cell j are fixed and equal to $B_{1,j}, \dots, B_{i-1,j}, B_{i+1,j}, \dots, B_{N,j}$. Also, w.l.o.g. assume that $w_{1,j} > \dots > w_{N,j}$. We are interested in the first advertiser k^* such that $w_{k^*+1,j} \leq \frac{\sum_{i=1}^{k^*} B_{i,j}}{M_j}$ as $B_{i,j}$ varies. Let $k^* = k^0$ when $B_{i,j} = 0$. If $k^0 < i$, then no matter how much budget i allocates, the price setting mechanism allocates no advertising opportunities to them, because advertisers $1, \dots, k^0$ have sufficient budget to buy all ads at a price that is higher than what i can afford; thus $U_{i,j} = 0$ and, subsequently, $\tilde{U}_{i,j} = U_{i,j} = 0$. If, on the other hand, $k^0 \geq i$, then the utility function $U_{i,j}$ will have the form that we depict in Figure 4. In particular, we can form the $i - k^0 + 1$ regions R_k , $i \leq k \leq k^0$, such that that the first advertiser in region R_k with the property that $w_{k+1,j} \leq \frac{\sum_{i=1}^k B_{i,j}}{M_j}$ is k . In particular, when $B_{i,j} = 0$ then $k^* = k^0$ and we get the leftmost region R_{k^0} ; as $B_{i,j}$ grows larger k^* eventually becomes i and remains so thereafter. Points P, P_1, P_2 , and P_3 in Figure 4 correspond to budget allocations $B_{i,j}$ equal to $B, B_j = w_{k+1,j} \cdot M_j - S_{k,j} - B_{k,j}, B_2 = w_{k,j} \cdot M_j - S_{k,j} - B_{k,j}$, and $B_3 = w_{k,j} \cdot M_j - S_{k,j}$, respectively.

We will now determine the upper concave envelope of $U_{i,j}$ by focusing on regions R_k , with $i \leq k \leq k^0$. Define $S_{k,j} = \sum_{i=1}^{i-1} B_{i,j} + \sum_{i'=i+1}^{k-1} B_{i',j}$ (for $k=i$ this expression gives $S_{i,j} = \sum_{i'=1}^{i-1} B_{i',j}$). Region R_i (rightmost region in Figure 4) consists of a concave part which corresponds to the utility function $U_{i,j}(B_{i,j}) = w_{i,j} M_j \frac{B_{i,j}}{\sum_{i'=1}^i B_{i',j}}$ for $B_{i,j} \in [w_{i+1,j} \cdot M_j - S_{i,j} - B_i, w_{i,j} \cdot M_j - S_{i,j}]$, followed by a constant part for $B_{i,j} \geq w_{i,j} \cdot M_j - S_{i,j}$

(the constant part corresponds to the maximum possible advertising opportunities that advertiser i may get); the utility function in region R_i is thus already concave so we do not need to focus more on it. Every other region R_k , $i < k \leq k^0$, will consist of the concave part $w_{i,j}M_j \frac{B_{i,j}}{\sum_{l=1}^k B_{l,j}}$ for $B_{i,j} \in [w_{k+1,j}M_j - S_{k,j} - B_{k,j}, w_{k,j}M_j - S_{k,j} - B_{k,j}]$, followed by the linear part $w_{i,j} \frac{B_{i,j}}{w_{k,j}}$ for $B_{i,j} \in [w_{k,j}M_j - S_{k,j} - B_{k,j}, w_{k,j}M_j - S_{k,j}]$. Of course $B_{i,j} \geq 0$, so if any of the endpoints of the aforementioned intervals is negative we simply replace it with 0. From Lemma 1, we can easily derive that $U_{i,j}(B_{i,j})$ is continuous in the domain $B_{i,j} \geq 0$. It is also differentiable everywhere except for the points where the utility function transitions from the concave part to the linear part, and vice versa.




Fig. 4. Utility function $U_{i,j}$ when $k^0 \geq i$ and its upper concave envelope

We now examine the derivatives in regions R_k , $i \leq k \leq k^0$. For region R_i , i.e., when $k=i$, the derivative in $(w_{i+1,j}M_j - S_{i,j} - B_{i,j}, w_{i,j}M_j - S_{i,j})$ is $w_{i,j}M_j \frac{S_{i,j}}{(S_{i,j} + B_{i,j})^2}$, while it is 0 for $B_{i,j} > w_{i,j}M_j - S_{i,j}$. For region R_k , with $i \leq k \leq k^0$, the derivative in $(w_{k+1,j}M_j - S_{k,j} - B_{k,j}, w_{k,j}M_j - S_{k,j} - B_{k,j})$ is $w_{i,j}M_j \frac{S_{k,j} + B_{k,j}}{(S_{k,j} + B_{k,j} + B_{i,j})^2}$, while the derivative in $(w_{k,j}M_j - S_{k,j} - B_{k,j}, w_{k,j}M_j - S_{k,j})$ is $w_{i,j}/w_{k,j}$. Although $U_{i,j}$ is not differentiable at the transition points, the left $\partial_- U_{i,j}$ and right $\partial_+ U_{i,j}$ derivatives obviously exist. We now state the following two results.

Lemma 3: $\partial_- U_{i,j}(w_{k+1,j}M_j - S_{k,j} - B_{k,j}) > \partial_+ U_{i,j}(w_{k+1,j}M_j - S_{k,j} - B_{k,j})$, for $i \leq k < k^0$.

Lemma 4: $\partial_- U_{i,j}(w_{k,j}M_j - S_{k,j} - B_{k,j}) < \partial_+ U_{i,j}(w_{k,j}M_j - S_{k,j} - B_{k,j})$, for $i < k \leq k^0$.

Lemma 3 implies that whenever we make a transition from the linear to the concave part (e.g., point P_j in Figure 4) the first derivative gets lower, and concavity is

maintained. In contrast, Lemma 4 suggests that when we move from the concave to the linear part (e.g., point P_2), the first derivative gets higher; this in turn violates concavity. We will show how to tackle this by considering region R_k , $i \leq k \leq k^0$, in Figure 4. The idea is to draw a line ε_l from P_3 to the point P in the concave part of region R_k so that the line ε_l is tangent to the curve. Based on our previous discussion, the derivative at P is $w_{i,j}M_j \frac{S_{k,j}+B_{k,j}}{(S_{k,j}+B_{k,j}+B)^2}$. On the other hand, the slope of ε_l is

$$\frac{U_{i,j}(B_3) - U_{i,j}(B)}{B_3 - B} = w_{i,j} \frac{\frac{B_3}{w_{k,j}} M_j \frac{B}{S_{k,j}+B_{k,j}+B}}{B_3 - B}. \text{ Thus, we are looking for a } B \text{ such that}$$

$w_{i,j}M_j \frac{S_{k,j}+B_{k,j}}{(S_{k,j}+B_{k,j}+B)^2} = w_{i,j} \frac{\frac{B_3}{w_{k,j}} M_j \frac{B}{S_{k,j}+B_{k,j}+B}}{B_3 - B}$. But $B_3 = w_{k,j} \cdot M_j - S_{k,j}$, so the previous equation becomes after some algebraic manipulations:

$$S_{k,j}B^2 - 2(M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j})B - (M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j})(S_{k,j} + B_{k,j} - M_j w_{k,j}) = 0 \quad (1)$$

Equation (1) is a quadratic equation, which accepts the two solutions $\frac{(M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j}) \pm \sqrt{(M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j})S_{k,j}M_j w_{k,j}}}{S_{k,j}}$. First, note that $M_j \cdot w_{k,j} > S_{k,j}$ (since $B_3 > 0$), so the solutions are real numbers. Second, we keep the solutions with the minus because it is lower than $B_3 = M_j \cdot w_{k,j} - S_{k,j}$ and even $B_2 = M_j \cdot w_{k,j} - S_{k,j} - B_{k,j}$. Indeed, after performing some algebraic manipulations we get

$$\frac{(M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j}) - \sqrt{(M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j})S_{k,j}M_j w_{k,j}}}{S_{k,j}} < M_j w_{k,j} - S_{k,j} - B_{k,j} \Leftrightarrow M_j w_{k,j} > S_{k,j} +$$

$B_{k,j}$, which is true. Now, there are 2 cases. If the solution is greater than $w_{k+1,j}M_j - S_{k,j} - B_{k,j}$ (see point P_l in Figure 4), then we draw the line ε_l from P to P_3 as we show in Figure 4. Else, we draw the line from P_l to P_3 (we illustrate such a scenario with line ε_2 in region R_{k^0} in Figure 4). We summarize the two cases by writing

$$B = \max \left\{ \frac{(M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j}) - \sqrt{(M_j w_{k,j} - S_{k,j})(S_{k,j} + B_{k,j})S_{k,j}M_j w_{k,j}}}{S_{k,j}}, w_{k+1,j}M_j - S_{k,j} - B_{k,j} \right\}.$$

We will now prove that the slope of ε_l is greater than the right derivative at P_3 . Indeed, the slope of ε_l is $w_{i,j}M_j \frac{S_{k,j}+B_{k,j}}{(S_{k,j}+B_{k,j}+B)^2}$. The right derivative at P_3 , on the other hand, is $w_{i,j}M_j \frac{S_{k,j}}{(w_{k,j}M_j)^2}$. But then $w_{i,j}M_j \frac{S_{k,j}+B_{k,j}}{(S_{k,j}+B_{k,j}+B)^2} > w_{i,j}M_j \frac{S_{k,j}}{(S_{k,j}+B_{k,j}+B)^2} >$

$w_{i,j}M_j \frac{S_{k,j}}{(S_{k,j}+B_{k,j}+(M_j w_{k,j} - S_{k,j} - B_{k,j}))^2} = w_{i,j}M_j \frac{S_{k,j}}{(M_j w_{k,j})^2}$, which proves our claim. Moreover, it is easy to see that the slope of ε_l is lower than the left derivative at P_l , since the opposite would imply that the line segment P_2 - P_3 has a slope $w_{i,j}/w_{k,j}$ that is greater than the left derivative at P_l $w_{i,j}/w_{k+1,j}$, which is untrue given that $w_{k,j} > w_{k+1,j}$.

We repeat the process described above in all regions. At the end of this process, we derive a utility function $\tilde{U}_{i,j}$ that is continuous everywhere, differentiable everywhere except for the points where it changes slope, and the left and right derivatives (which exist for all $B_{i,j} \geq 0$) are monotonically non-increasing in the allocated budget $B_{i,j}$. But then $\tilde{U}_{i,j}$ will be concave in terms of $B_{i,j}$. Now, recall that $U_i = \sum_{j=1}^L U_{i,j}$. If we repeat the above process for every $U_{i,j}$, $1 \leq j \leq L$, we can eventually form the function $\tilde{U}_i(\mathbf{B}_i; \mathbf{B}_{-i}) = \sum_{j=1}^L \tilde{U}_{i,j}(B_{i,j}; B_{-i,j})$ (where \mathbf{B}_{-i} denotes the vector of budget allocations

of all advertisers but i). The function $\tilde{U}_i(\mathbf{B}_i; \mathbf{B}_{-i})$ is the sum of concave functions, so it is also concave in i 's strategy \mathbf{B}_i . In the end, the new utility functions $\tilde{U}_i(\mathbf{B}_i; \mathbf{B}_{-i})$, $1 \leq i \leq N$, possess two important properties: (1) each $\tilde{U}_i(\mathbf{B}_i; \mathbf{B}_{-i})$ is continuous in $(\mathbf{B}_i; \mathbf{B}_{-i})$, and (2) each $\tilde{U}_i(\mathbf{B}_i; \mathbf{B}_{-i})$ is concave in \mathbf{B}_i for any fixed value of \mathbf{B}_{-i} . Moreover, the strategy space of every advertiser is convex, closed and bounded. Consequently, based on Rosen's theorem [16] we can immediately derive that a Nash equilibrium exists. We denote that equilibrium by $\tilde{\mathbf{B}} = (\tilde{\mathbf{B}}_1, \dots, \tilde{\mathbf{B}}_N)$. Moreover, we call \tilde{G} the new game when the utility functions are replaced by their upper-concave envelopes.

Note that $\tilde{\mathbf{B}}$ may not be an equilibrium of game G . This means that there may be players in game G who have an incentive to deviate if the strategy vector $\tilde{\mathbf{B}}$ is chosen. However, the following lemma shows that we can bound the maximum utility that a player can gain by deviating.

Lemma 5: Let the strategy vector $\tilde{\mathbf{B}}$ be a Nash equilibrium of game \tilde{G} . Then the maximum utility that player i can gain by deviating from $\tilde{\mathbf{B}}$ in game G is $\tilde{U}_i(\tilde{\mathbf{B}}_i; \tilde{\mathbf{B}}_{-i}) - U_i(\tilde{\mathbf{B}}_i; \tilde{\mathbf{B}}_{-i})$.

Essentially, the above result says that we can find a set of budget allocations such that we can know exactly the maximum utility that an advertiser may gain by deviating. Note that in the special case where the Nash equilibrium of game \tilde{G} falls into the parts of \tilde{U}_i that are equal to U_i , then the Nash equilibria of game \tilde{G} are also Nash equilibria of game G .

7 Conclusion

The market for location-based advertising is set to witness an unprecedented growth over the next years. The massive proliferation of modern mobile phones with embedded geo-positioning functionality and the development of fast wireless communication protocols have created exciting opportunities for advertisers to reach the user base that is most relevant to them. On the other hand, sponsored search advertising has been a thriving market in the last decade for advertisers who want to advertise their product or service to online users posing relevant queries. Inspired by the enormous success of sponsored search and the immense potential for LBA, we address the market for location-based sponsored search advertising. We provide models that build on prior work in sponsored search advertising, but we additionally consider that advertisers are characterized by location-dependent valuations. We distinguish between three cases: (1) bids-only case, (2) budgets-only case, and (3) bids-and-budgets case, and analyzed the equilibrium strategies in the corresponding markets using game theoretical tools.

There are several research directions that we would like to pursue with regard to the market for location-based sponsored search advertising. First, we would like to extend our model so that it takes into account the more subtle issues that are involved in the sponsored search market such as the externalities between the displayed ads, or the more realistic scenario of advertisers who are interested in several keywords.

Second, our model assumed offline ad slot scheduling [9], where we estimate the number of queries in every cell, and then allocate to every advertiser a percentage of the ads in every cell. It would be interesting to deal with the more challenging problem of online ad slot scheduling, where the number of expected queries per cell is not available in advance. Finally, we would like to fully explore the equilibrium strategies in the bids-and-budgets case, as our current work provides equilibrium strategies only for the case where advertisers have sufficiently small or large budgets.

Acknowledgements. This work was supported by grant RPC11EG01 from HKUST. We thank the anonymous reviewers for their valuable comments.

References

1. Banerjee, S., Dholakia, R.: Does location-based advertising work? *International Journal of Mobile Marketing* 3(1) (2008)
2. Brainard, W., Scarf, H.: How to compute equilibrium prices in 1891. Cowles Foundation Discussion Paper 1270 (2000)
3. Bruner, G., Kumar, A.: Attitude toward location-based advertising. *Journal of Interactive Advertising* 7(2) (2007)
4. Clarke, E.H.: Multipart pricing of public goods. *Public Choice* 11(1), 17–33 (1971)
5. Dhar, S., Varshney, U.: Challenges and business models for mobile location-based services and advertising. *Communications of the ACM* 54(5), 121–129 (2011)
6. Debreu, G., Koopmans, T.C.: Additively decomposed quasiconvex functions. *Mathematical Programming* 24(1), 1–38 (1982)
7. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review* 9(1), 242–259 (2007)
8. Feldman, J., Muthukrishnan, S.: Algorithmic methods for sponsored search advertising. In: *Performance Modeling and Engineering*, Springer, Heidelberg, pp. 91–124. Springer, Heidelberg (2008)
9. Feldman, J., Muthukrishnan, S., Nikolova, E., Pál, M.: A truthful mechanism for offline ad slot scheduling. In: Monien, B., Schroeder, U.-P. (eds.) *SAGT 2008*. LNCS, vol. 4997, pp. 182–193. Springer, Heidelberg (2008)
10. Feldman, M., Lai, K., Zhang, L.: The proportional-share allocation market for computational resources. *IEEE Transactions on Parallel and Distributed Systems* 20(8), 1075–1088 (2009)
11. Groves, T.: Incentives in teams. *Econometrica* 41, 617–631 (1973)
12. Jansen, B., Mullen, T.: Sponsored search: an overview of the concept, history, and technology. *Int. J. Electronic Business* 6(2), 114–131 (2008)
13. Johari, R., Tsitsiklis, J.N.: Efficiency loss in a network resource allocation game. *Mathematics of Operations Research* 29(3), 407–435 (2004)
14. Kelly, F.: Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8, 33–37 (1997)
15. Kempe, D., Mahdian, M.: A cascade model for externalities in sponsored search. In: *4th International Workshop on Internet and Network Economics*, pp. 585–596 (2008)
16. Rosen, J.B.: Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica* 33(3), 520–534 (1965)

17. Trimponias, G., Papadias, D., Bartolini, I.: Location-based sponsored search advertising. Technical Report, <http://www.cse.ust.hk/~dimitris/PAPERS/SSTD13-TR.pdf>
18. Unni, R., Harmon, R.: Perceived effectiveness of push vs. pull mobile location-based advertising. *Journal of Interactive Advertising* 7(2), 28–40 (2007)
19. Varian, H.: Position auctions. *International Journal of Industrial Organization* 25(6), 1163–1178 (2007)
20. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. *Finance* 16, 8–27 (1961)
21. <http://advertise.bingads.microsoft.com/en-us/home>
22. <http://adwords.google.com/support/aw/?hl=en>