Location-based Sponsored Search Advertising

George Trimponias[1], Ilaria Bartolini[2], Dimitris Papadias[1]

[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

Technical Report HKUST-CS13-02

Abstract. Recent technological advancements have created unprecedented opportunities for location-based advertising. In this work, we investigate algorithms for location-based sponsored search, assuming a grid structure, where the number of queries per cell is known in advance. Advertisers have different valuations for different cells, and, optionally, daily budget constraints. When no budget constraints exist, we first show how auctions in conventional sponsored search can be properly modified to incorporate location information, and then we discuss the possible use of combinatorial auctions. In the case where budget constraints are present, we distinguish between two cases: advertisers may be indifferent to the price per click, or they may not be willing to pay more than their valuation. For the first case, we utilize a novel probabilistic framework that is inspired by recent progress in resource allocation markets. Under strong competition, a Nash equilibrium always exists. The second case is harder, as the utility functions have a complex form: we show however that a Nash equilibrium exists under certain conditions. In the general case, we can bound a player's most profitable deviation using concave relaxations of the utilities.

Keywords: Location-based advertising, Sponsored search, Game theory.

1. INTRODUCTION

The adoption by millions of users of powerful mobile devices with GPS functionality and the development in metropolitan areas of fast wireless 3G/4G networks have recently created unprecedented opportunities for location-based advertising [Dhar and Varshney 2011, Bruner and Kumar 2007]. Advertisers can utilize users' positional information to send advertising material to relevant consumers, which has in turn created an exciting market for a number of companies such as Google AdMob Ads and Millennial Media.

Location-based advertising, especially in its mobile form, is poised for tremendous growth because of its special characteristics [Banerjee and Dholakia 2008]. First, it enables personalization: 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, e.g., the advertising messages can take into account the 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 in 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 2012 was over $50 billion mainly due to advertising[1]. Sponsored search consists of three parties [Jansen and Mullen 2008]: (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, we investigate *location-based sponsored search advertising*, assuming a grid structure, where the number of queries per cell is estimated based on query logs. To incorporate the location component, we consider that advertisers have different valuations for different cells. Similar to [Feldman and Muthukrishnan 2008], we distinguish between 3 cases depending on the advertisers' constraints: (1) *bids-case*: advertisers declare a maximum amount of money that they are willing to pay per click (*bid*), but are not bounded by budget constraints, (2) *budgets-case-case*: the advertisers have a maximum daily budget, but are indifferent to the price per click, and (3) *bids-and-budgets-case*: the advertisers have both a maximum daily budget and a maximum bid. We investigate the Nash equilibria in all cases, using game theoretical tools. Concretely, the main contributions of the paper can be summarized as follows:

— In the bids-case, we first show how the generalized second-price (GSP) auction [Edelman et al. 2007; Varian 2007] can be properly modified into a location-based GSP procedure (LB-GSP) to incorporate location information. To ensure the search engine revenue is never compromised, we introduce a combinatorial location-based format (LB-COMB-GSP) based on the combinatorial GSP auction [Ghosh et al. 2007], which groups grid cells into bundles and then sells the bundles.

— In the budgets-case, we introduce a novel probabilistic framework inspired by resource allocation markets [Feldman et al. 2009]: the number of ads allocated to an advertiser is proportional to its budget. A Nash equilibrium always exists as long as there are at least two advertisers (or zero) interested in a cell.

— In the bids-and-budgets-case, we use a price-setting mechanism based on [Feldman and Muthukrishnan 2008] to determine the price in each cell. We prove the existence of Nash equilibria for sufficiently small or large budgets. Although we do not have results for the

[1] See http://investor.google.com/financial/2012/tables.html.

general case, we bound a player's most profitable deviation by considering concave relaxations of the utility functions, where it is easier to prove that a Nash equilibrium exists [Rosen 1965].

The rest of the paper is organized as follows. Section 2 surveys related work in sponsored search advertising. Section 3 provides a general model for location-based sponsored search. Sections 4-6 investigate the bids-case, budgets-case, and bids-and-budgets-case-case, respectively. Finally, Section 8 concludes the paper.

2. RELATED WORK

In sponsored search [Jansen and Mullen 2008], advertisers target relevant users by declaring to the search engine a list of keywords of interest. For each keyword, they additionally specify their *cost per click* (or *bid*) which is the maximum amount of money they are willing to spend to appear on the results page for a given keyword. Every time a user enters a query, a small number of *paid* (or, *sponsored*) links appears on top or to the right side of the *organic* search results. In order to determine the winning advertisers and the price they need to pay, an auction occurs in an automated fashion. In practice, large search engines also use a *quality score* (*QS*) for every advertiser to better rank them, but for simplicity we omit such measures here.

We focus on a single keyword and assume that there are K slots to be allocated among $N = K$ advertisers[2]. Slot j has a *click through rate* $c_j \geq 0$, which denotes the probability that an ad placed in that slot is clicked. There is ample evidence in the literature that higher slots are associated with higher revenues. In practice, whether a user clicks on an ad or not depends on numerous factors, including the other ads [Kempe and Mahdian 2008; Ashkan and Clarke 2015], but we do not deal with *ad externalities*. Advertiser i has a *valuation* v_i *per click*, i.e., the expected amount of money it receives per click on the ad. Since c_j clicks per time unit correspond to slot j, the total valuation of i for slot j is equal to $w_{ij} = c_j \cdot v_i$. The sponsored search problem can be represented as a bipartite graph connecting the set of advertisers and the set of slots through edges with weight w_{ij}. The objective is then to determine an *assignment* of advertisers to slots, i.e., a bijection from the set of advertisers to the set of slots, which maximizes the sum of valuations of all advertisers, i.e., the *social welfare*. Essentially, sponsored search is reduced to the *assignment problem* [Shapley and Shubik 1971]. Note that the bipartite model can also incorporate information asymmetries between advertisers and slots (see, e.g., [Dughmi et al. 2014]).

It is possible to extend this to *decentralized* environments by viewing our original assignment problem as a two-sided market problem and introducing a proper price $p_j \geq 0$ for each slot j. Given a set of prices, we assume that bidders have *quasi-linear* utilities. Thus, if advertiser i obtains slot j, its net utility (also called *payoff*) per click is $v_i - p_j$, and its net utility per time unit is $c_j \cdot (v_i - p_j)$. Obviously, advertisers prefer the slots for which their utility is maximized. For any set of valuations, it is always possible to generate a proper set of prices so that each advertiser obtains (one of) its most preferred slots and no contention occurs [Shapley and Shubik 1971]. These prices are called *market-clearing*; progressive auction mechanisms can determine them efficiently [Demange et al. 1986]. Interestingly, for any set of market-clearing prices, the induced assignment is *socially optimal*, i.e., it has the maximum total valuation of any assignment of advertisers to slots.

As the advertiser valuations are private to them, the search engine relies on *sponsored search auctions*[3] where advertisers declare their bids. The result of such an auction is an *assignment rule* that assigns advertisers to slots, and a *payment rule* that determines the price per click. One desirable property of any auction format is to incentivize advertisers to report their valuations *truthfully*. This means that no advertiser has an incentive to provide a bid that is different from its true valuation. Auctions that satisfy this property are called *truthful* or *incentive-compatible*. Furthermore, any auction has to be *individually rational* for all bidders, i.e., a bidder with value 0 per click does not pay anything.

[2] If the advertiser set has larger size, then we can add fictitious slots for which no advertiser is interested. If the slot set has larger size, then we add fictitious advertisers that are not interested in any slot.

[3] A comprehensive review of sponsored search auctions is provided in [Qin et al. 2015].

A central question in the theory of auctions [Milgrom 2004; Krishna 2002; Easley and Kleinberg 2010] and mechanism design [Jackson 2003] is whether an auction format exists that (1) maximizes social welfare, (2) is truthful and (3) is individually rational. The *Vickrey-Clarke-Groves* (*VCG*) *auction* [Vickrey 1961; Clarke 1971; Groves 1973] achieves these goals by assigning the highest slot to the advertiser with the highest bid, the second highest to the advertiser with the second highest bid, and so on. Regarding the payment rule, each advertiser pays the minimum amount equal to *the externalities* that it imposes on the others, i.e., the decreases in the valuations of other bidders because of its presence. The VCG auction is the unique mechanism that satisfies all three properties [Green and Laffont 1977]. Moreover, VCG prices form the unique set of market-clearing prices of minimum total sum in the corresponding market [Leonard 1982; Demange 1983].

Currently, large commercial engines employ the *generalized second-price auction* (*GSP*) [Edelman et al. 2007; Varian 2007], which charges an advertiser the minimum amount required to maintain their ad's position in the sponsored results. For instance, suppose that advertiser with rank i bids b_i. In GSP, the price per click for advertiser i is determined by the bid b_{i+1} of advertiser with rank $i+1$, which is the minimum that the advertiser with rank i would have to bid to maintain its rank. Note that in this pricing scheme, a bidder's payment does not take into consideration its own bid. Despite its prevalence as the standard auction format, GSP is not *truthful*: advertisers have in general no incentive to declare their true valuations to the search engine [Edelman et al. 2007]. A notable exception is $K = 1$; when only one slot is available, GSP is rendered into the truthful *second-price* auction, where the unique winner pays the second-highest bid [Vickrey 1961]. Finally, note that

The above discussion assumes that advertisers are willing to pay any price, provided that it is individually rational to do so. This corresponds to the case where they are not characterized by budget constraints. In practice, however, it is possible that advertisers are bounded by finite budgets (e.g., they are willing to pay up to a certain amount per time unit). We discuss this case later on.

3. MODEL

Assume N advertisers and K slots 1, …, 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. We model this by associating each slot l with its *click through rate* c_l, $0 \le c_l \le 1$, which denotes the probability that an ad that is placed in that slot is clicked. Since higher slots are more valuable, we further assume that $c_l > c_{l'}$ whenever $l < l'$. In fact, 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 exclude them here; i.e., an ad located at slot l is clicked with probability c_l independent of the rest of the slots [Kempe and Mahdian 2008]. To keep the model simple, we also do not consider quality scores for advertisers.

We assume that the space is partitioned with a grid of L cells. Advertisers value users according to their location in the grid. 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. Estimating the valuation is a difficult marketing/operational research problem, beyond the scope of our work. Advertiser i has maximum bids $b_{i,j}$ for every cell j, which may or may not coincide with the true valuations $w_{i,j}$ (this is related to the *truthfulness* discussed in Section 2), and may be bounded by a maximum daily budget B_i. Advertisers are only aware of their own budget and valuations, which they declare to the search engine. We also assume that advertisers have *additive* utilities: their total utility over all cells is the sum of their utilities in each cell. We consider that advertisers are interested in exactly the same (unique) keyword; keyword interactions is an interesting research topic in its own right, and can be explored in future work. Finally, we assume an offline setting, where the expected number M_j of queries per day in cell j is estimated based on statistics. The total expected number of queries per day in all cells is $M = \sum_{j=1}^{L} M_j$.

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 areas 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. However, small cells may take a 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. Determining the proper grid resolution can thus be a critical factor of success for location-based sponsored search. In Section 4.2, we discuss how combinatorial auctions can mitigate this problem.

In the next Sections, we investigate 3 cases depending on the advertisers' constraints. In the *bids-only* case, each advertiser i has finite bids $b_{i,j}$ but is not bounded by a daily budget, i.e., B_i=∞. In the *budgets-only* case, each advertiser i is bounded by a finite daily budget B_i but is indifferent to the price per click, i.e., $b_{i,j}$=∞. Finally, in the *bids-and-budgets* case, each advertiser i has a finite daily budget B_i and finite bids $b_{i,j}$. Table I illustrates common symbols used throughout the paper.

Finally, we define a few important terms.

Definition 1. Let S be a subset of R^n. S is called *convex* if for any $x, y \in S$ and any $\lambda \in [0, 1]$, the point $\lambda x + (1 - \lambda)y$ is also in S.

Definition 2. Let $f: S \rightarrow R$ be a function defined on a convex subset S of R^n. We call the function f *concave* (resp., *convex*), if for any $x, y \in S$ and any $\lambda \in [0, 1]$, we have that $f(\lambda x + (1 - \lambda)y) \geq \lambda f(x) + (1 - \lambda)f(y)$ (resp., $f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$).

Definition 3. Let $f: S \rightarrow R$ be a function defined on a convex subset S of R^n. We call the function f *quasi-concave* (resp., *quasi-convex*), if for any $x, y \in S$ and any $\lambda \in [0, 1]$, we have that $f(\lambda x + (1 - \lambda)y) \geq min\{f(x), f(y)\}$ (resp., $f(\lambda x + (1 - \lambda)y) \leq max\{f(x), f(y)\}$). A function that is both quasi-concave and quasi-convex is called *quasi-linear*.

Definition 4. Let $f: S \rightarrow R$ be a function defined on a convex subset S of R^n. The *upper concave envelope* of f, denoted $\tilde{f}$, is defined as the smallest (point-wise) concave function such that $\tilde{f}(x) \geq f(x)$ for every $x \in S$.

TABLE I. 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 of advertiser i that is allocated into cell j
$w_{i,j}$	Valuation per click of advertiser i for cell j
$b_{i,j}$	Maximum bid per click of advertiser i for cell j (can be the same or different from $w_{i,j}$)
M_j	Expected number of queries per day for cell j
M	Total expected number of queries per day
c_l	Probability 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
$\widetilde{U}_{i,j}$	Upper concave envelope of $U_{i,j}$
$U_{i,j}^{rel}/U_i^{rel}$	Relaxation of $U_{i,j}$ / U_i
p_j	Price per click in cell j (for case 3)
s_j	Permutation of advertisers in cell j 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)
$S_{i,j}^k$	$\sum_{i'=0}^{i-1} B_{i,j} + \sum_{i'=i+1}^{k-1} B_{i,j}$ (for k=i, $S_{i,j}^i = \sum_{i'=0}^{i-1} B_{i,j}$) (for case 3)

4. BIDS-CASE

Section 4.1 discusses the adaptation of conventional sponsored search to account for location information. Section 4.2 presents an alternative approach that builds upon the theory of combinatorial auctions.

4.1 Location-based GSP Auction

In sponsored search advertising, whenever the search engine receives a query from a cell, it runs an auction where each advertiser bids an amount of money equal to its valuation per click for that particular cell. Assuming M queries per day, this would imply that a total of M auctions take place, one for every query, in order to determine the K winners that will fill the slots, as well as the prices per click that they have to pay. However, since the result of an auction is the same for a given cell, no matter how many queries are issued inside that cell, L auctions would suffice. Indeed, the actual number of queries per cell does not matter for the auction; cells with high workload will simply involve more auctions compared to cells with lower traffic. Any auction format, such as the GSP or the VCG (see Section 2), can be utilized. These two auctions have been extensively studied in the literature; 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.

Given the prevalence of the GSP auction in the industry, we next focus on the GSP framework. First we discuss the complexity of deciding the winners and the price they need to pay. For a given cell, we need to determine the first $K + 1$ out of N advertisers with the highest bid, which costs $O(N)$ for $K \ll N$. Computing the prices for the K winners takes $O(K)$ operations. Finally, since we need to repeat this process for al L cells, the total complexity can be bounded by $O(L \cdot (N + K)) = O(L \cdot N)$, for small K. Regarding the metrics for the GSP auction, it is straightforward to generalize the metrics of Section 2 in the location-based framework. For instance, if advertiser i has the l^{th} highest bid $b_l^{(j)}$ in cell j, $1 \le l \le K$, i will be allocated to slot l for a price $b_{l+1}^{(j)}$ per click, where $b_{l+1}^{(j)}$ is the $(l+1)^{th}$ highest bid in j (or 0 if i=K=N). Then, the expected payoff per query for advertiser i in cell j is $c_l(w_{i,j} - b_{l+1}^{(j)})$, and its expected payoff per day for cell j is $M_j c_l(w_{i,j} - b_{l+1}^{(j)})$. Since we assumed additive utilities, an advertiser's total expected utility per day is the sum of the expected utilities per day per cell.

4.2 Location-based Combinatorial GSP Auction

The previous approach suffers from a significant drawback, demonstrated in Figure 1a. Consider 4 advertisers 1, 2, 3, 4, and 4 cells c_1, c_2, c_3, c_4. The bids per click for the four advertisers are shown in a vector form: the first element corresponds to advertiser 1, the second to advertiser 2, etc. Assume there is only one slot per query for any cell, i.e., K=1, and only one query per cell. If the GSP protocol is employed, then advertiser 1 wins the slot for cell c_1 and pays 0.1 per click (the second highest bid); advertiser 2 wins c_2 and pays 0.2; advertiser 3 wins c_3 and pays 0.2; and advertiser 4 wins c_4 and pays 0.3 per click. The cumulative profit for the search engine will then be 0.1+0.2+0.2+0.3 = 0.8; this is significantly lower than the sum of the highest bids per cell, which is equal to 4. The reason for this discrepancy is that in every cell there is only one advertiser with high bid, while the rest of the advertisers bid low.

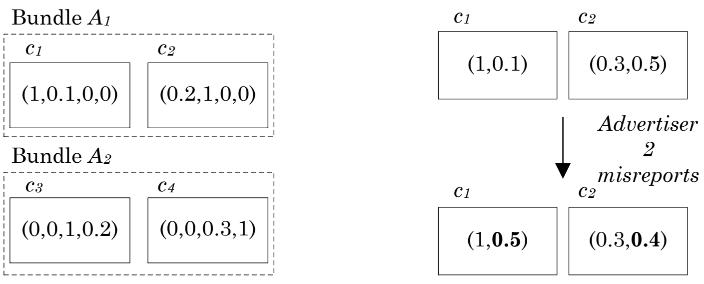

(a) Motivation. (b) GNS is non-truthful.

Fig. 1. Combinatorial GSP auction.

In general, it may well be the case that for some cells, there are only few advertisers with non-negligible bids. In that case, the standard GSP procedure can seriously compromise the search engine's revenue, since it may lead to low prices per click. To alleviate this problem, the search engine can instead utilize a combinatorial auction to allocate cell slots to advertisers. To illustrate the idea, consider first that there is only one slot per query, i.e., K=1. While standard auctions assign a query to the winning advertiser, a combinatorial auction assigns *bundles* of cells and their associated queries to advertisers. A bundle is defined as a set of cells whose queries are sold together as a single unit. Due to the additive valuations, the valuation of any advertiser for a bundle is equal to the sum of its valuations for all queries in the bundle.

For instance, imagine that the search engine creates two bundles A_1 and A_2; the former for cells c_1 and c_2, and the latter for cells c_3 and c_4. Since we assumed additive utilities, the valuation of advertiser 1 for A_1 is 1.2 while the valuation of advertiser 2 is 1.1 (advertisers 3 and 4 have 0 valuations for bundle A_1); the GSP protocol will then assign A_1 to advertiser 1 for a price of 1.1. Thus, the profit for the search engine for cells c_1, c_2 increases from 0.1+0.2=0.3 to 1.1. Similar arguments can be made for the bundle A_2 of cells c_3 and c_4 that increases the search engine's profit from 0.2+0.3 = 0.5 to 1.2 (with advertiser 3 the winner for A_2).

A search engine is naturally interested in an *optimal bundling*, i.e., a partitioning of the cell slots into bundles that maximizes its revenue. Interestingly, finding an optimal bundling is strongly NP-hard, which can be proved by reduction from the 3-partition problem [Ghosh et al. 2007]. However, the same authors provide an approximate algorithm which yields revenue at least as high as half the optimal one. Algorithm 1 properly adjusts their algorithm, which we term LB-COMB-GSP. We assume temporarily that advertisers bid truthfully, i.e., $b_{i,j} = w_{i,j}$; we discuss truthfulness later.

ALGORITHM 1. Combinatorial Location-based Algorithm for Bundling Cell Queries

Input: N advertisers and their valuations for L cells, number of queries M_j per cell j per day.
Output: Bundles of cell queries.
for *each advertiser i* **do**
    $W_i = \sum_{j\in C_i} w_{i,j}\, M_j$;
end
Renumber advertisers so that $W_1 \geq \cdots \geq W_N$;
$r_1 = \sum_{j\in C_1} w_j^{(2)} M_j + \sum_{i=1}^{\lfloor N/2 \rfloor} W_{2i+1}$;
$r_2 = \sum_{i=1}^{\lfloor N/2 \rfloor} W_{2i}$;
if $r_1 \geq r_2$ **then**
    $\forall cell\ j \in C_1$, sell its M_j queries separately from other cells;
    for i = 1 **to** $\lfloor N/2 \rfloor$ **do**
        bundle together queries for cells in C_{2i} and C_{2i+1};
else
    for i = 1 **to** $\lfloor N/2 \rfloor$ **do**
        bundle together queries for cells in C_{2i-1} and C_{2i};

LB-COMB-GSP views the combined M_j queries in cell j rather than individual queries as an item to be sold. In particular, it first decomposes the space of cells into N subsets $C_1, \ldots, C_N$, where C_i is the set of cells where advertiser i, $1 \leq i \leq N$, has the highest valuation per click. The highest total valuation W_i for subset C_i will then depend by definition on advertiser i's valuations for cells in C_i, i.e., $W_i = \sum_{j\in C_i} w_{i,j}\, M_j$. Without loss of generality, we assume that $W_1 \geq \cdots \geq W_N$. Next, the algorithm attempts to bundle consecutive subsets of the decomposition; these have either the form (C_{2i}, C_{2i+1}) or (C_{2i-1}, C_{2i}) (even-indexed or odd-indexed first, respectively). The reason why this operation is so important is that by bundling together (C_{2i}, C_{2i+1}) (or, (C_{2i-1}, C_{2i})), the search engine can guarantee a minimal profit of W_{2i+1} (or, W_{2i}). Indeed, when bundling (C_{2i}, C_{2i+1}) together, the highest advertiser has a valuation for that bundle of at least W_{2i} (since advertiser $2i$ has a valuation of no less than W_{2i}) and the second-highest advertiser a valuation of at least W_{2i+1} (since advertiser $2i$+1 has a valuation of no less than W_{2i+1}); but then selling the bundle with the GSP pricing will result in a price of at least W_{2i+1}. Similar reasoning can be applied to bundles of the form (C_{2i-1}, C_{2i}). One caveat concerns the subset C_1, when the bundles are of the form (C_{2i}, C_{2i+1}). In that case, C_1 is not bundled with any other subset, so its queries are sold separately according to the GSP format: a query in cell $j \in C_1$ will then be sold for $w_j^{(2)}$, i.e., the second-highest valuation per click for cell j. To decide on the bundling form, LB-COMB-GSP computes the quantities $r_1 = \sum_{j\in C_1} w_j^{(2)} M_j + \sum_{i=1}^{\lfloor N/2 \rfloor} W_{2i+1}$ and $r_2 = \sum_{i=1}^{\lfloor N/2 \rfloor} W_{2i}$. If the former is greater, it sells queries in C_1 separately and creates bundles of the form (C_{2i}, C_{2i+1}); otherwise, it sells bundles of the form (C_{2i-1}, C_{2i}). Based on the aforementioned observation, this operation guarantees search engine revenue of at least $\max\{r_1, r_2\}$. On the other hand, $r_1 + r_2$ is the highest possible revenue that the search engine can achieve (corresponding to the ideal case where the highest advertiser pays its full valuation). The optimal bundling can then guarantee revenue at least as high as $r_1 + r_2$. Moreover, since $\max\{r_1, r_2\} \geq \frac{r_1+r_2}{2}$, this means that LB-COMB-GSP guarantees revenue that is at least as high as half the optimal one. The algorithm has a time complexity of $O(L \cdot N^2)$, since it creates $\lfloor N/2 \rfloor$ bundles and the GSP auction in each bundle costs $O(L \cdot N)$ as we discussed above.

The LB-COMB-GSP algorithm for one slot has two desirable properties. First, its revenue is at least half the optimal revenue, as we just showed; second, its efficiency is at least half the maximum efficiency. Furthermore, it can be easily generalized to multiple slots, provided the click-through rate decreases geometrically from higher to lower slots [Ghosh et al. 2007]. Unfortunately, even in the simple case of K=1, LB-COMB-GSP is not truthful, as we show in Figure 1b. In this example, there are 2 advertisers, 2 cells (c_1, c_2), and just 1 query per cell. Advertiser 1 has a valuation vector (1,0.1), while advertiser 2 has a valuation vector (0.3,0.5). If both advertisers bid truthfully, c_1 and c_2 are bundled, and advertiser 1 wins them for a price of 0.1+0.5 = 0.6. But if advertiser 2 changes its bid to (0.5,0.4), the queries for c_1 and c_2 are sold separately: advertiser 1 wins the query for c_1 for a price of 0.5, and advertiser 2 wins the query for c_2 for a price of 0.3. Since advertiser 2 gets a higher payoff by misreporting, LB-COMB-GSP is not truthful, even for a single slot. This is in sharp contrast with GSP, which is truthful for the case K=1 (see Section 2). One may then wonder whether it is possible to design truthful combinatorial auctions with good approximation ratios on the efficiency. It turns out this is a hard question, well beyond the scope of this work. For the sake of completeness, we just mention that a polynomial in the number of advertisers and items algorithm exists when valuations are submodular (as our additive utilities) with approximation ratio $O(\sqrt{m})$, where m the number of items [Dobzinski et al. 2005]. This is arguably a poor ratio. Moreover, it is impossible to do better in polynomial time under certain conditions [Dobzinski 2011].

5. BUDGETS-CASE

Each 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 the bids-case, where $b_{i,j}$ is the maximum amount that i is willing to pay per click, the payments per click are now bounded only by B_i, and the cell bids are only used to determine the relative importance of cells. For simplicity, we initially consider a single slot (K=1) with probability of being clicked c_1=1, and deal with several slots later. Moreover, although an advertiser does not necessarily declare its true valuations, i.e., it may be that $b_{i,j} \neq w_{i,j}$, we will assume for simplicity of exposition that $b_{i,j} = w_{i,j}$. At the end of the Section, we discuss truthfulness in more detail. Since budget constraints are now involved, it is convenient to assume a *Fisher market* rather than a quasi-linear model [Brainard and Scarf 2000]. In the latter, the price an advertiser pays is subtracted from the utility; in the former, money does not bear any intrinsic value and advertisers are willing to burn their entire budget.

In the absence of budget constraints, our goal was to design auctions for selecting the winners and determining the price. Although it is possible to extend conventional GSP to deal with budget constraints, we propose a different framework, where no auction takes place. Instead 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.

5.1 A Proportional Framework

Our proportional framework computes for each cell the probability that any advertiser will be chosen as a response to a 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 [Feldman and Muthukrishnan 2008; Kelly 1997; Johari and Tsitsiklis 2004]. Thus, 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. If 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. 2009]. 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. We say i's budget allocation $\boldsymbol{B_i} = (B_{i,1}, \dots, B_{i,L})$ is *valid* if $\sum_{j=1}^{L} B_{i,j} = B_i$ and $B_{i,j} \geq 0$. The probability that i will be displayed in cell j is $\frac{B_{i,j}}{\sum_k B_{k,j}}$ (assuming $\sum_k B_{k,j} \neq 0$; when $\sum_k B_{k,j} = 0$, the utility may be discontinuous, as we show later), where $\sum_k B_{k,j}$ 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 $U_{i,j}= w_{i,j} M_j \frac{B_{i,j}}{\sum_k B_{k,j}}$ (assuming $\sum_k B_{k,j} \neq 0$), since it gets a value $w_{i,j}$ for every query in j when displayed with a probability $\frac{B_{i,j}}{\sum_k B_{k,j}}$, and there are M_j queries in total in cell j. We have assumed *additive* utilities, so i's total utility U_i is the sum of its utilities $U_{i,j}$ over all cells: $U_i = \sum_j U_{i,j} = \sum_j w_{i,j} M_j \frac{B_{i,j}}{\sum_k B_{k,j}}$ (again, assuming $\sum_k B_{k,j} \neq 0, \forall j$). 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). Moreover, the utility functions U_i and $U_{i,j}$ are concave in advertiser i's arguments.

Regarding the form of the utility function, we observe that the utility of any advertiser i in any cell j only depends on the budget allocations in that particular cell c, i.e., $U_{i,j} = U_{i,j}(B_{1,j}, \dots, B_{i,j}, \dots, B_{N,j})$ for any $1\leq i\leq N$, $1\leq j\leq L$. The total utility is then function of all budget allocations, i.e.,$U_i = U_i(\boldsymbol{B_1}, \dots, \boldsymbol{B_i}, \dots, \boldsymbol{B_N})$ for any $1\leq i\leq N$.

In order to obtain a proper budget allocation, we will utilize the notion of Nash equilibrium. The strategy space for advertiser i is the convex, bounded and closed set $\{(B_{i,1}, \dots, B_{i,L}) \mid B_{i,j} \geq 0$ and $\sum_{j=1}^{L} B_{i,j} = B_i\}$, i.e., an L-simplex. 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 $\boldsymbol{B_1^*}, \dots, \boldsymbol{B_N^*}$ (where $\boldsymbol{B_i} = (B_{i,1}, \dots, B_{i,L})$) form a Nash equilibrium if for any other valid strategy $\boldsymbol{B_i}$, $1\leq i\leq N$, we have:

$$U_i(\boldsymbol{B_1^*}, \dots, \boldsymbol{B_i^*}, \dots \boldsymbol{B_N^*},)\geq U_i(\boldsymbol{B_1^*}, \dots, \boldsymbol{B_i}, \dots, \boldsymbol{B_N^*}).$$

It turns out that the above game does not always accept a Nash equilibrium. To see why, consider two advertisers 1 and 2 with budgets B_1, $B_2 > 0$, and two cells c_1 and c_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 c_1. For player 2, the best strategy would obviously be to allocate its entire budget B_2 to c_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 $\varepsilon > 0$ to c_2 (and win all advertising opportunities in 2) and spend the rest $B_1 - \varepsilon$ on cell 1 (and maximize its share in cell 1 as well). Unfortunately, there is no optimal value for ε, since it must be (1) positive to ensure 1 gets all ads in c_2, and (2) as small as possible so that advertiser 1 wins the largest possible share in c_1. Such ε does not exist.

The root of the non-existence of a Nash Equilibrium[4] in the examples above lies in the discontinuity of the utility functions at point 0. It turns out, however, that a sufficient (but not necessary) condition for existence of a Nash equilibrium is the game to be *strongly competitive* [Feldman et al. 2009], i.e., for a given cell there are at least two advertisers with positive valuations. Moreover, the equilibrium has good efficiency properties. This condition is in general expected to hold in a big market with several participants, such as the market for sponsored search advertising. In the rare case that the condition does not hold, we can still guarantee a

[4] More accurately, the non-existence of Nash equilibria in the examples is due to the fact that the utility functions are not upper semi-continuous at 0 [Kakutani 1941].

Nash equilibrium by enforcing a minimum price per click that advertisers need to pay (*reserve price*) [Feldman et al. 2009] [5].

The next issue concerns the computation of the Nash equilibrium. We first describe the optimization problem that an advertiser faces. Consider advertiser i with budget $\boldsymbol{B_i} = (B_{i,1}, \dots, B_{i,L})$. Now, assume the other advertisers have their budgets *fixed* and equal to $\boldsymbol{B_{-i}} = (\boldsymbol{B_1}, \dots, \boldsymbol{B_{i-1}}, \boldsymbol{B_{i+1}}, \dots, \boldsymbol{B_N})$. To find an optimal $\boldsymbol{B_i}$, advertiser i must solve the problem *OPT*:

$$OPT\colon \max_{\boldsymbol{B_i}} U_i(\boldsymbol{B_i}; \boldsymbol{B_{-i}})$$

$$s.t. B_{i,j} \geq 0 \forall j \in \{1, \dots, L\} \; and \; \sum_{j=1}^{L} B_{i,j} = B_i.$$

In the above $U_i(\boldsymbol{B_i}; \boldsymbol{B_{-i}})\colon \boldsymbol{R}^L \to \boldsymbol{R}$ is a function of only $\boldsymbol{B_i}$, considering $\boldsymbol{B_{-i}}$ fixed; on the other hand $U_i(\boldsymbol{B_i}, \boldsymbol{B_{-i}}) : \boldsymbol{R}^{L \cdot N} \to \boldsymbol{R}$ denotes the function of all arguments $(\boldsymbol{B_1}, \dots, \boldsymbol{B_i}, \dots, \boldsymbol{B_N})$. Now assume that the game is strongly competitive so that a Nash equilibrium exists. We can then make the following observation. In cells j for which all advertisers have valuations equal to 0, all advertisers will have valuations $U_{i,j} = 0$ no matter what the budget allocations on these cells. This function is trivially continuously differentiable. On the other hand, consider a cell j where at least one advertiser has positive valuation. Then, at equilibrium the sum of budgets $\sum_k B_{k,j}$ on that cell must be strictly positive, i.e., $\sum_k B_{k,j} > 0$. Otherwise, we can use similar reasoning as in the above example to show that an advertiser would have an incentive to allocate a tiny budget to j to win all ads, thus contradicting that is an equilibrium. But then the functions $U_{i,j} = w_{i,j} M_j \frac{B_{i,j}}{\sum_k B_{k,j}}, 1 \leq i \leq N$, are defined at equilibrium, and it is easy to see they are continuously differentiable. Moreover, at equilibrium $U_i = \sum_j U_{i,j}$ is also continuously differentiable as the sum of continuously differentiable functions. Also, everywhere *near* the equilibrium U_i will be continuously differentiable.

So, consider in *OPT* that $\boldsymbol{B_{-i}} = \boldsymbol{B^*_{-i}}$, i.e., players other than i play the Nash equilibrium allocations. Then by the definition of a Nash equilibrium $\boldsymbol{B_i} = \boldsymbol{B^*_i}$ will be a solution to *OPT*. Since (1) U_i is continuously differentiable at $(\boldsymbol{B^*_i}, \boldsymbol{B^*_{-i}})$ and (2) the constraints are affine functions, by the Karush-Kuhn-Tucker (KKT) conditions [Bertsekas 2009] there exist a real constant λ and non-negative constants $\mu_j \geq 0, 1 \leq j \leq L$, that satisfy the two conditions:

$$\left.\frac{\partial U_i(\boldsymbol{B_i}; \boldsymbol{B^*_{-i}})}{\partial B_{i,j}}\right|_{\boldsymbol{B_i} = \boldsymbol{B^*_i}} = \lambda - \mu_j \text{ and } \mu_j B_{i,j} = 0, 1 \leq j \leq L$$

Consider all cells for which advertiser i has a non-zero allocation at equilibrium, i.e., $B^*_{i,j} > 0$. Then, the second condition implies that $\mu_j = 0$ for these cells and thus $\left.\frac{\partial U_i(\boldsymbol{B_i}; \boldsymbol{B^*_{-i}})}{\partial B_{i,j}}\right|_{\boldsymbol{B_i} = \boldsymbol{B^*_i}} = \lambda$. The term $\frac{\partial U_i(\boldsymbol{B_i}; \boldsymbol{B^*_{-i}})}{\partial B_{i,j}}$ corresponds to what we call the marginal utility of advertiser i for cell j. Thus, we just showed that at equilibrium the marginal utilities of advertiser i for all cells with $B^*_{i,j} > 0$ are the same. Regarding cells with $B^*_{i,j} = 0$, the second condition yields $\mu_j \neq 0$, and since μ_j is non-negative it will be $\mu_j > 0$ for these cells. Consequently, by the first condition, the marginal utility at these cells will be lower than λ, i.e., the marginal utility at cells with positive budget allocation.

In the original game with the discontinuous utilities it may occur that some advertiser i has positive valuation for a context j and $\sum_k B_{k,j} = 0$. In that case, the marginal utility of i for j is infinite, since any allocation $\varepsilon > 0$ means that i can win all ads in j. We call such contexts *irregular*, to distinguish them from contexts where marginal utilities are bounded. Unfortunately, irregular contexts complicate the update rules of our methods, which are largely based on the marginal utilities. Indeed, an infinite marginal utility implies that any "tiny" $\varepsilon > 0$ from a regular to an irregular context yields higher total utility.

To tackle this issue, we introduce a continuous approximation of the original utilities based on the notion of *perturbation*. Specifically, assume a *fictitious* advertiser (the new number of

---

[5] One potential challenge with reserve prices is that they may render our framework computationally intractable: an advertiser has the choice to either allocate zero, or the reserve price, or more than the reserve price in any cell, which may in turn lead to an exponential number of combinations.

advertisers becomes $N+1$) that allocates a very small and *fixed* amount $\varepsilon > 0$ in each context. We call the new game an *ε-perturbation* of the original game. As shown in Figure 2b, the new game is everywhere continuous; indeed, there is now always a budget of at least ε in every context, so it is not possible for the denominator in the marginal utilities to be equal to 0, i.e., $\sum_k B_{k,j} > 0$ for every context j. Moreover, any advertiser will always have bounded marginal utilities for all contexts. For $\varepsilon = 0$, the ε-perturbation coincides with the original game. For $\varepsilon > 0$, the ε-perturbation accepts a pure Nash equilibrium by Rosen's theorem [Rosen 1965]. As long as the game is strongly competitive, this equilibrium will tend to the equilibrium of the original game as $\varepsilon \to 0$ [Feldman et al. 2009].

An immediate question is how to compute the aforementioned equilibrium, which is guaranteed to exist in strongly competitive games. One approach would be to numerically solve the KKT conditions. However, this is a very cumbersome task, further complicated by the fact that the utility functions are not continuous at 0. For this reason, we use alternative approaches that implement various classes of *distributed dynamics*. The idea behind distributed dynamics is that if the current budget allocation is not in equilibrium, then at least one player gains by deviating. So, instead of trying to calculate the equilibrium numerically, we simulate the fact that in each iteration every player will update its individual budget allocation to increase its total utility [Feldman et al. 2009].

5.2 Multiple Slots and Truthfulness

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 slot (as opposed to the bids-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.

The above generalization has the problem that the assignment of slots to advertisers may not be a bijection anymore. Indeed, if advertiser i is assigned more than M_j slots in total for context j, then by the pigeonhole principle there will be a query (out of all M_j queries of context j) where i is assigned two slots. This, in turn, violates the general assumption of the sponsored search model that one advertiser is assigned to at most one slot for any given query. If users click on ads independently and we ignore the ad externalities, then it should not make a difference whether an advertiser appears just once or several times on the ad list. In general, however, this may not be the case.

One way to address this issue is to add a constraint that a user can never get allocated more than M_j slots in total for context j. In slot l of context j, advertiser i appears $\frac{B_{i,j,l}}{\sum_k B_{k,j,l}} M_j$ times; thus, in all slots it appears $\sum_{l=1}^{K} \frac{B_{i,j,l}}{\sum_k B_{k,j,l}} M_j$ times. Since we do not want this number to exceed M_j, we will have $\sum_{l=1}^{K} \frac{B_{i,j,l}}{\sum_k B_{k,j,l}} M_j \leq M_j$, or equivalently, $\sum_{l=1}^{K} \frac{B_{i,j,l}}{\sum_k B_{k,j,l}} \leq 1$. By adding one such constraint per user-context pair (i,j), we can make sure that the resulting solution will correspond to a bijection. Unfortunately, this constraint is not linear, which makes the problem significantly harder. This can be investigated in future work.

Finally, with a proper example we can show that our proportional framework is not truthful, even for 1 slot (K=1). Assume 3 advertisers 1, 2, 3 with budgets B_1=B_2=B_3=100, and 2 cells c_1, c_2 with M_1=M_2=100. All advertisers have a valuation equal to 10 per click for both cells, with one exception: advertiser 1 has valuation 6 per click for c_1. Observe that advertisers 2 and 3 are symmetric, so there must be an equilibrium where they have identical allocations and utilities. Now, if all advertisers declare their valuations truthfully, we can compute that in the Nash equilibrium advertiser 1 approximately allocates 71.2 in c_1 and 28.8 in c_2 for a total utility of 1051.5. (In equilibrium, each of the other 2 advertisers approximately allocates 45.8 in c_1 and 54.2 in c_2 for a total utility of 1343.2). If, however, 1 chooses to misreport its valuation for c_2, and pretend it has a 10 valuation for c_2, then in the Nash equilibrium all advertisers' bids and budgets are identical, so advertiser 1 allocates 50 in c_1 and 50 in c_2 for a total utility of 1066.7, which is better than reporting truthfully. (Each of the other 2 advertisers also allocates 50 in c_1 and 50 in c_2 for a total utility of 1333.3). Note, however, that in order to manipulate the bids, an advertiser needs to have information that is private to the other players, such as their budget and bids. In practice this information is very hard to obtain, which may well be one of the reasons why GSP is also used so widely, even though in theory it is not truthful.

6. BIDS-AND-BUDGETS-CASE

In the bids-and-budgets-case, advertiser i declares a maximum daily budget B_i, but contrary to the budgets-case, i is now not willing to exceed a declared bid $b_{i,j}$ per click in cell j. Similar to Section 5, we assume $b_{i,j} = w_{i,j}$, and we discuss non-truthfulness later. We only deal with the case of a single slot, i.e., K=1 with c_1=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.

6.1 The Price-Setting Mechanism

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-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 $p_j = (B_{1,j} + \cdots + B_{N,j})/ M_j$. 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 derive a price p_j^* such that all advertisers who can afford that price have sufficient budgets to purchase all the advertising opportunities.

Algorithm 5 presents the price-setting mechanism [Feldman et al. 2008; Feldman and Muthukrishnan 2008] that determines that price p_j^* in cell j. It is essentially a price-descending mechanism: the price keeps falling until p_j^* is reached. Moreover, it is truthful. We state it w.l.o.g. in the simple case where the bids are decreasingly ordered, i.e., $w_{1,j} > w_{2,j} > \cdots > w_{N,j} > 0$; in the general case, we can define for every cell a permutation $s_j: \{1, \dots, N\} \to \{1, \dots, N\}$ that reorders the bids in j in decreasing order, i.e., $w_{s_j(1),j} > w_{s_j(2),j} > \cdots > w_{s_j(N),j} > 0$.

ALGORITHM 5. Price-Setting Mechanism in Cell j for One Slot.

Input: N advertisers, their budget allocations in cell j and their valuations for cell j. Number of queries M_j in cell j per day.
Output: Price per click.
Assume w.l.o.g. that $w_{1,j} > w_{2,j} > \cdots > w_{N,j} > 0$;
$k^* \leftarrow \min_k \ w_{k+1,j} \le \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$;
$p_j^* \leftarrow \min\left\{\frac{\sum_{i=1}^{k^*} B_{i,j}}{M_j}, w_{k^*,j}\right\}$;
for $i = 1$ **to** $k^* - 1$ **do**
    allocate $\frac{B_{i,j}}{p_j^*}$ ads to advertiser i;
allocate $M_j - \sum_{i=1}^{k^*} B_{i,j}/p_j^*$ ads to advertiser i;
for $i > k$ **do**
    allocate 0 ads to advertiser i;

Recall that in the budgets-case, the price per query in cell j would be $p_j = \frac{\sum_{i=1}^{N} B_{i,j}}{M_j}$. Obviously, p_j is linear in its arguments $B_{i,j}$ (1≤i≤N) and continuous. On the other hand, the price-setting mechanism of Algorithm 5 yields prices that are clearly more complex. First, the price p_j for cell j is 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 demonstrate the price function, consider a setting with 2 advertisers 1 and 2 with maximum bids $w_{1,j}$ and $w_{2,j}$ (with $w_{1,j} > w_{2,j}$), and cell j with M_j queries per day. Figure 2 depicts how the price varies according to the budgets $B_{1,j}$ and $B_{2,j}$ that the advertisers allocate in cell j.

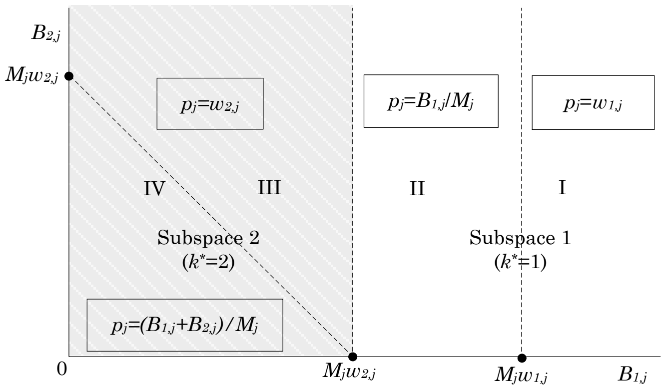

Fig. 2. Price p_j in cell j in the case of 2 advertisers with $w_{1,j} > w_{2,j}$.

In particular, if $B_{1,j} \geq M_j w_{2,j}$, then k^*=1 (Subspace 1) and the price is determined as the minimum of $B_{1,j}/M_j$ and $w_{1,j}$. Concretely, when $B_{1,j} \geq M_j w_{1,j}$ then the price is equal to $w_{1,j}$ (region I), and when $B_{1,j} < M_j w_{1,j}$, the price is equal to $B_{1,j}/M_j$ (region II). On the other hand, when $B_{1,j} < M_j w_{2,j}$, then k^*=2 (Subspace 2) and the price is the minimum of $w_{2,j}$ and $(B_{1,j} + B_{2,j})/M_j$; for $B_{1,j} + B_{2,j} \geq M_j w_{2,j}$ the price is $w_{2,j}$ (region III), while for $B_{1,j} + B_{2,j} > M_j w_{2,j}$, the price is $(B_{1,j} + B_{2,j})/M_j$ (region IV). Note that inside a region, the price can be either constant or linear. Moreover, 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-case $p_j = (B_{1,j} + \cdots + B_{N,j})/M_j$ is unbounded: it can get arbitrarily large as the budgets that the advertisers allocate grow larger.

The above example captures some important properties of the price. The price setting mechanism decomposes the budget space into N *subspaces* (one for each of the N possible k^*), and then further divides that region into two regions: the price is constant inside one and linear in the other. In the example of Figure 2, subspace 1 for k^*=1 comprises of regions I and II, while subspace 2 (for k^*=2) comprises of regions III and IV.

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

PROOF. We assume w.l.o.g. that $w_{1,j} > w_{2,j} > \cdots > w_{N,j} > 0$. For a region that corresponds to a given k^*, the price function is given by the minimum of the two continuous functions $\frac{\sum_{i=1}^{k^*} B_{i,j}}{M_j}$ and $w_{k^*,j}$, so it is continuous. So, any potential discontinuity would occur on the boundary as we move from k^* to k^*+1 (or from k^* to k^*−1). The N boundaries are given by the N equalities $w_{k+1,j} = \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$, k=1, …, N. Consider any boundary $w_{k+1.j} = \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$. On the boundary, we have p_k=min{$\frac{\sum_{i=1}^{k} B_{i,j}}{M_j}, w_{k,j}$}. But $w_{k,j} > w_{k+1,j} = \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$, so $p_k = \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$. Let's now consider any point in the vicinity of the boundary (sufficiently close to the boundary). There are two cases, either $\frac{\sum_{i=1}^{k} B_{i,j}}{M_j} < w_{k+1,j}$, or $\frac{\sum_{i=1}^{k} B_{i,j}}{M_j} > w_{k+1,j}$. If $w_{k+1,j} < \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$, then these points belong to the same region as the boundary points and thus continuity is ensured from the discussion above. If $w_{k+1,j} > \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$, then it holds that $w_{k+2,j} < \frac{\sum_{i=1}^{k} B_{i,j}}{M_j}$, given that bids are in decreasing order and we consider sufficiently close points. But since$B_{i,j}$>0, this also implies that $w_{k+2,j} < \frac{\sum_{i=1}^{k+1} B_{i,j}}{M_j}$. Hence those points will belong to the region defined by k^*=k+1. The price for these points is min{$\frac{\sum_{i=1}^{k+1} B_{i,j}}{M_j}$, $w_{k+1,j}$}. Now, as we approach any point in the boundary it will be that $\frac{\sum_{i=1}^{k} B_{i,j}}{M_j} \to w_{k+1,j}$, so min{$\frac{\sum_{i=1}^{k+1} B_{i,j}}{M_j}$, $w_{k+1,j}$}→ $w_{k+1,j}$, since the function min{$\frac{\sum_{i=1}^{k+1} B_{i,j}}{M_j}$, $w_{k+1,j}$} is continuous and $B_{k+1,j}$≥0. So, no matter how we approach the points in the boundary, the price will tend to k+1, which is the value at any point at the boundary so the price function is continuous at the boundary points as well. □

Next, we discuss the utilities. For a given allocation, denote $z_{i,j}$ the number of ads (i.e., number of queries) 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 $U_i = \sum_j w_{i,j}z_{i,j}$. We can compute $z_{i,j}$ directly from the price-setting mechanism, assuming that $w_{1,j} > w_{2,j} > \cdots > w_{N,j} > 0$. If $w_{2,j} \le B_{1,j}/M_j$, then k_j^*=1 and the price p_j^* is min{$B_{1,j}/M_j$,$w_{1,j}$}. On the other hand, if $w_{2,j} > B_{1,j}/M_j$ we continue by checking whether $w_{3,j} \le (B_{1,j} + B_{2,j})/M_j$. If the latter is true, then k_j^*=2 and the price p_j^* = min{$(B_{1,j} + B_{2,j})/M_j$, $w_{2,j}$}. 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^*.

How does the above utility function compare to the simpler utility function $U_{i,j} = w_{i,j}M_j \frac{B_{i,j}}{\sum_k B_{k,j}}$ of the budgets-case? 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 $w_{i,j}M_j$ as $B_{i,j}$ tends to infinity. In other words, the advertiser will get 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 state something similar for the utility function in the more complex setting when both budgets and maximum bids per click are declared? We begin by showing the monotonicity of $U_{i,j}$.

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

PROOF. Consider that the budget allocations for cell c_j are fixed by all advertisers except for i. Let $b_1 > b_2 \ge 0$. We will prove that $U_{i,j}(b_1) \ge U_{i,j}(b_2)$. We distinguish between the following cases:

1. If $U_{i,j}(b_2)$=0, then it follows immediately that $U_{i,j}(b_1) \ge U_{i,j}(b_2)$, since the utility is always non-negative.

2. If $U_{i,j}(b_2)$>0, then it can be that a) k^*=i, or b) k^*>i. We examine both cases.

—2(a). After i increases its bid to b_1, k^* will remain equal to i, since it must still hold that $w_{k^*+1,j} \le \frac{\sum_{i'=1}^{k^*} B_{i',j}}{M_j}$, and moreover all $B_{i,j}$ for i'<i are unchanged, so k^* cannot become smaller. But then the new price $p^* = \min\left\{\frac{\sum_{i=1}^{k^*} B_{i,j}}{M_j}, w_{k^*,j}\right\}$ will be at least as high as the previous price, so the number of queries that i wins $M_j - \frac{\sum_{i=1}^{k^*-1} B_{i,j}}{p^*}$ will be at least as high as what it got when its bid was lower.

— 2(b). There are 2 cases about the new k^* when bidder i increases its bid to b_1: the new $\tilde{k}^*$ is (a) still greater than i, or (b) becomes equal to i. In case (a), the share of queries that i wins will be $\frac{b_1}{\sum_{i'=1}^{\tilde{k}^*} B_{i,j}} \ge \frac{b_1}{\sum_{i'=1}^{k^*} B_{i,j}}$. Moreover, because the function $\frac{B_{i,j}}{\sum_{i'=1}^{k^*} B_{i,j}}$ is monotonically increasing in $B_{i,j}$, we will have that $\frac{b_1}{\sum_{i'=1}^{k^*} B_{i,j}} > \frac{b_2}{\sum_{i'=1}^{k^*} B_{i,j}}$, so when bidding $b_1 >\ b_2$ user i will get a larger share of queries. In case (b), the new share of queries that i wins is 1−$\frac{\sum_{i'=1}^{i-1} B_{i',j}}{\min\left\{\sum_{i'=1}^{i} B_{i',j}, b_i\right\}}$. If $\min\left\{\frac{\sum_{i'=1}^{i} B_{i',j}}{M_j}, w_{i,j}\right\} = \frac{\sum_{i'=1}^{i} B_{i',j}}{M_j}$, then the share of queries that i wins is 1−$\frac{\sum_{i'=1}^{i-1} B_{i',j}}{\sum_{i'=1}^{i} B_{i',j}} = \frac{B_{i,j}}{\sum_{i'=1}^{i} B_{i',j}} > \frac{B_{i,j}}{\sum_{i'=1}^{k^*} B_{i',j}}$. The last expression is monotonically increasing in $B_{i,j}$, so the advertiser ends up with a higher share when bidding more. If $\min\left\{\frac{\sum_{i'=1}^{i} B_{i',j}}{M_j}, w_{i,j}\right\} = w_{i,j}$, 1−$\frac{\sum_{i'=1}^{i-1} B_{i',j}}{\min\left\{\sum_{i'=1}^{i} B_{i',j}, w_i\right\}}$>1−$\frac{\sum_{i'=1}^{i-1} B_{i',j}}{\sum_{i'=1}^{i} B_{i',j}}$, so by using the same argument as above we conclude that i will win a higher percentage of queries. □

As we show later, $U_{i,j}$ is not concave in $B_{i,j}$. But since it is monotonically increasing in $B_{i,j}$, it must also be *quasi-concave* in $B_{i,j}$. It turns out, however, that when $U_{i,j}$ are quasi-concave, but not concave in $B_{i,j}$, then their sum $U_i = \sum_{j=1}^{L} U_{i,j}$ is not quasi-concave in $(B_{i,1}, \dots, B_{i,L})$ [Debreu and Koopmans 1982]. This is a worrisome result, in the sense that standard existence theorems for Nash equilibria usually assume concave or, at least, quasi-concave utility functions. For sufficiently small or large budgets, we can however prove that a Nash equilibrium exists. Note that for every cell j, s_j denotes the permutation $s_j\colon \{1, \dots, N\} \to \{1, \dots, N\}$ that reorders the bids in cell j in decreasing order, i.e., $w_{s_j(1),j} > w_{s_j(2),j} > \cdots > w_{s_j(N),j} > 0$.

PROPOSITION 1. Assume $\sum_{i=1}^{N} B_i \leq M_j w_{N,j}$ for every cell j, and that there is strong competition. Then a Nash equilibrium exists, and coincides with the Nash equilibrium of the budgets-case.

PROOF. Since $B_{i,j} \leq B_i$ for any cell j, we have that $\sum_{i=1}^{N} B_{i,j} \leq \sum_{i=1}^{N} B_i \leq M_j w_{N,j}$. Since $\sum_{i=1}^{N-1} B_{s_j(i),j} \leq \sum_{i=1}^{N} B_{i,j}$ (given that $B_{i,j}$ non-negative), we then have $\frac{\sum_{i=1}^{N-1} B_{s_j(i),j}}{M_j} \leq w_{s_j(N),j}$. Similarly and using the definition of s_j, we can prove that $\frac{\sum_{i=1}^{i^*} B_{s_j(i),j}}{M_j} \leq w_{s_j(i^*+1),j}$ for any$i^* \in \{1, \dots, N-1\}$. But then in any cell it holds that k^*=N, and the price $p_j^* = \min\left\{\frac{\sum_{i=1}^{N} B_{i,j}}{M_j}, w_{N,j}\right\} = \frac{\sum_{i=1}^{N} B_{i,j}}{M_j}$, since $\sum_{i=1}^{N} B_{i,j} \leq M_j w_{N,j}$. 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 the proportional framework of Section 5, and it thus always admits a Nash equilibrium if there is strong competition. ☐

PROPOSITION 2. Assume $B_i \geq \sum_{j \in C_i} M_j w_j^{(2)}$ for every advertiser i, and that there is strong competition. Then a set of Nash equilibria always exists. Moreover, the GSP auction of the bids-case yields a Nash equilibrium.

PROOF. For any advertiser i, consider the set of cells C_i where i has the highest valuation per click among all advertisers, i.e., $C_i = \{j \mid w_{i,j} = \max_{1 \leq i' \leq N}\{w_{i',j}\}\}$ (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 since $B_i \geq \sum_{j \in C_i} M_j w_j^{(2)}$, for every advertiser i. We now show that the above sets of budget allocations correspond to Nash equilibria. 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 the case where i allocates exactly $M_j w_j^{(2)}$ in cells $j \in C_i$, then in the resulting Nash equilibrium 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 yields a Nash equilibrium. ☐

6.2 Relaxation of the Utility Function and Nash Equilibria

We observed how the bids-and-budgets-case encompasses the simpler bids-case and budgets-case for sufficiently small or large budgets, respectively. Since both bids-case and budgets-case are non-truthful, the bids-and-budgets-case is also not truthful in general. Furthermore, 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 [Johari and Tsitsiklis 2004; Feldman et al. 2009]. In this direction, we will introduce a *fictitious* advertiser with i=0 with two properties. First, 0 allocates a fixed tiny budget $B_\varepsilon > 0$ in every cell, i.e., $B_{0,j} = B_\varepsilon$, for every cell j. Second, 0 has an arbitrarily large valuation per click for every cell, i.e., $w_{0,j} = w_{max}$ for every cell j. We call the perturbed game with the additional player G^{ε}. The arbitrarily large valuation per click for every cell implies that advertiser 0 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 0 has a negligible impact. Essentially, the introduction of player 0 removes the discontinuity at point 0. From now on, we will only consider the new game G^{ε} instead of the original setting.

In the general case, we cannot answer whether game G^{ε} always accepts a Nash equilibrium since each advertiser's utility function is not quasi-concave. Nevertheless, we can find a budget allocation such that the maximum utility that an advertiser can gain by deviating is known. Since U_i is a difficult function that we cannot directly deal with, we will instead consider a suitable *relaxation* $\widetilde{U}_i$ of U_i in G^{ε}, where we can show that a Nash equilibrium always exists. In this direction, we will first construct for every advertiser i the *upper-concave envelope* of $U_{i,j}(B_{i,j};\boldsymbol{B_{-i,j}})$ considering that the budgets of the rest of the advertisers in cell j are fixed ($\boldsymbol{B_{-i,j}}$ is the vector of budget allocations in cell j of all advertisers but i). Formally, we will be looking for the infimum $\widetilde{U}_{i,j}$ over all functions that are concave and are greater than or equal to $U_{i,j}$ for any $B_{i,j}$. But first, we describe in more detail the form of utility functions $U_{i,j}$ in G^{ε}.

6.2.1. Form of utility function $U_{i,j}(B_{i,j})$. We focus on advertiser i and cell j, $1 \le i \le N$ and $1 \le j \le L$. (Our analysis takes into consideration the fictitious player 0, but recall that its budget allocation is fixed.) Assume the rest of the advertisers' budgets for cell j are fixed and equal to $B_{0,j} = B_{\varepsilon}, B_{1,j}, \dots, B_{i-1,j}, B_{i+1,j}, \dots, B_{N,j}$. Also, w.l.o.g. assume that $w_{1,j} > w_{2,j} > \cdots > w_{N,j} > 0$. We are interested in the first advertiser k^* such that $w_{k^*+1,j} \le \frac{\sum_{i=0}^{k^*} B_{i,j}}{M_j}$ as $B_{i,j}$ varies; this corresponds to the advertiser of the price-setting mechanism (Algorithm 5).

Let $k^*=k^0$ when $B_{i,j} = 0$.

The case when $k^0<i$. Independently of the budget of i, the price setting mechanism allocates no advertising opportunities to them, because advertisers 1, …, k^0 have sufficient budget to buy all ads at a price that is higher than what i can afford; thus $U_{i,j} \equiv 0$ and, subsequently, $\widetilde{U}_{i,j} = U_{i,j} = 0$, for all $B_{i,j} \ge 0$.

The case when $k^0=i$. The utility function $U_{i,j}$ consists of a concave part, followed by a constant part. Indeed, the price starts increasing as $B_{i,j}$ increases from 0 until it reaches a value $w_{i,j} = w_{k^0,j}$ for a budget $B_{i,j}$ that satisfies $w_{i,j} = \frac{\sum_{i'=0}^{i-1} B_{i',j} + B_{i,j}}{M_j}$. No matter how much advertiser i increases its budget thereafter this will have no effect on the ads it wins. The form of $U_{i,j}$ is similar to region R_i in Figure 3. Since $U_{i,j}$ is already concave, its upper concave envelope trivially coincides with it, i.e., $\widetilde{U}_{i,j} = U_{i,j}$, for all $B_{i,j} \ge 0$.

The case when $k^0>i$. The utility function $U_{i,j}$ has a complex form. Figures 3 and 4 depict two possible types of $U_{i,j}$ that we use in our analysis. In particular, we can form the $i-k^0+1$ *regions* R_k, $i \le k \le k^0$, such that the first advertiser in region R_k with the property that $w_{k+1,j} \le \frac{\sum_{i=0}^{k} B_{i,j}}{M_j}$ is k. In particular, when $B_{i,j}=0$ then $k^*=k^0$ and we obtain the leftmost region R_{k^0}; as $B_{i,j}$ grows larger k^* eventually becomes i and remains so thereafter. Now, define $S_{i,j}^{k} = \sum_{i'=0}^{i-1} B_{i',j} + \sum_{i'=i+1}^{k-1} B_{i',j}$ (for $k=i$ this expression gives $S_{i,j}^{i} = \sum_{i'=0}^{i-1} B_{i',j}$). For $i<k<k^0$, we can identify 3 points of interest for region R_k. Point P_1 corresponds to the budget $B_{i,j} = B_1 = w_{k+1,j}M_j - S_{i,j}^{k} - B_{k,j}$ (the value of B_1 can be easily derived from the price-setting mechanism) where the k^* of the price-setting mechanism changes from k+1 to k. Point P_2 corresponds to the budget $B_{i,j} = B_2 = w_{k,j}M_j - S_{i,j}^{k} - B_{k,j}$ where the price of the price-setting mechanism when $k^*=k$, changes from $\frac{\sum_{i=1}^{k^*} B_{i,j}}{M_j}$ to $w_{k^*,j}$. Point P_3 corresponds to the budget $B_{i,j} = B_3 = w_{k,j}M_j - S_{i,j}^{k}$ where the k^* of the price-setting mechanism changes from k to k-1.

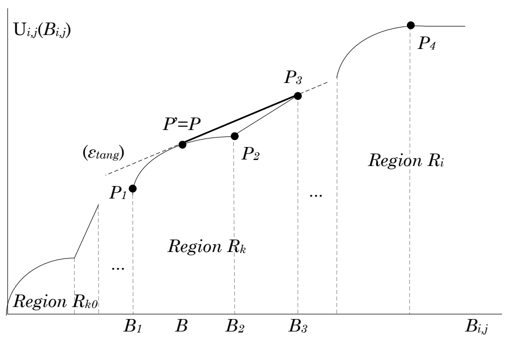

Fig. 3. Example 1 of $U_{i,j}$ and $\widetilde{U}_{i,j}$ when $k^0>i$.

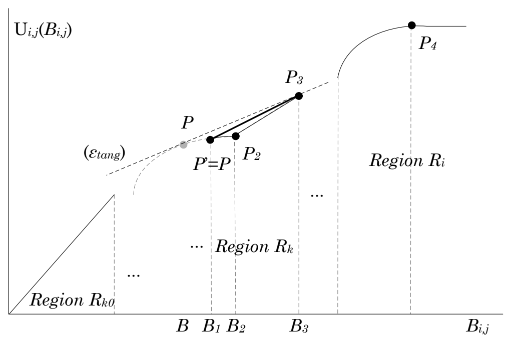

Fig. 4. Example 2 of $U_{i,j}$ and $\widetilde{U}_{i,j}$ when $k^0>i$.

We will now describe regions R_k, for $i\leq k\leq k^0$.

The case when $k=i$. Region R_i (rightmost region in Figures 3 and 4) consists of a concave part which corresponds to the utility function $U_{i,j}(B_{i,j}) = w_{i,j}M_j \frac{B_{i,j}}{\sum_{i'=0}^{i} B_{i',j}}$ for $B_{i,j} \in [w_{i+1,j}M_j - \sum_{i'=0}^{i-1} B_{i',j}, w_{i,j}M_j - \sum_{i'=0}^{i-1} B_{i',j}]$, followed by a constant part for $B_{i,j} \geq w_{i,j}M_j - \sum_{i'=0}^{i-1} B_{i',j}$ (the constant part corresponds to the maximum possible advertising opportunities that advertiser i may get).

The case when $i<k<k^0$. Region R_k, $i<k\leq k^0$, consists of the concave part $w_{i,j}M_j \frac{B_{i,j}}{\sum_{i'=0}^{k} B_{i',j}}$ for $B_{i,j} \in [w_{k+1,j}M_j - S_{i,j}^k - B_{k,j}, w_{k,j}M_j - S_{i,j}^k - B_{k,j}]$, followed by the linear part $w_{i,j} \frac{B_{i,j}}{w_{k,j}}$ for $B_{i,j} \in [w_{k,j}M_j - S_{i,j}^k - B_{k,j}, w_{k,j}M_j - S_{i,j}^k]$. $U_{i,j}$ is differentiable everywhere except for the points where its form changes from concave to linear, or linear to concave (points P_1, P_2, and P_3 in Figures 3 and 4).

The case when $k=k^0>i$. There are two possibilities for region R_{k^0}. (1) It consists of the concave part of the concave part $w_{i,j}M_j \frac{B_{i,j}}{\sum_{i'=0}^{k^0} B_{i',j}}$ for $B_{i,j} \in [0, w_{k^0,j}M_j - S_{i,j}^{k^0} - B_{k^0,j}]$, followed by the linear part $w_{i,j} \frac{B_{i,j}}{w_{k^0,j}}$ for $B_{i,j} \in [w_{k^0,j}M_j - S_{i,j}^{k^0} - B_{k^0,j}, w_{k^0,j}M_j - S_{i,j}^{k^0}]$, as in Figure 3. (2) It consists of only the linear part $w_{i,j} \frac{B_{i,j}}{w_{k^0,j}}$ for $B_{i,j} \in [w_{k^0,j}M_j - S_{i,j}^{k^0} - B_{k^0,j}, w_{k^0,j}M_j - S_{i,j}^{k^0}]$, as in Figure 4.

6.2.2. Derivatives of utility function $U_{i,j}(B_{i,j})$ when $k^0>i$. $U_{i,j}$ is differentiable everywhere except for the points where its form changes from concave to linear, linear to concave, and concave to constant (e.g., points P_2, P_1, P_3, and P_4 in Figures 3 and 4).

In particular, for region R_i, i.e., when $k=i$, the derivative in $(w_{i+1,j}M_j - \sum_{i'=0}^{i-1} B_{i,j}, w_{i,j}M_j - \sum_{i'=0}^{i-1} B_{i,j}$ is $w_{i,j}M_j \frac{\sum_{i'=0}^{i-1} B_{i,j}}{(\sum_{i'=0}^{i} B_{i,j})^2}$, while it is 0 for $B_{i,j} \geq w_{i,j}M_j - \sum_{i'=0}^{i-1} B_{i,j}$. For region R_k, with $i<k<k^0$, the derivative in $(w_{k+1,j}M_j - S_{i,j}^k - B_{k,j}, w_{k,j}M_j - S_{i,j}^k - B_{k,j})$ is $w_{i,j}M_j \frac{S_{i,j}^k + B_{k,j}}{(S_{i,j}^k + B_{k,j} + B_{i,j})^2}$, while the derivative in $(w_{k,j}M_j - S_{i,j}^k - B_{k,j}, w_{k,j}M_j - S_{i,j}^k)$ is $\frac{w_{i,j}}{w_{k,j}}$. 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. $\partial_- U_{i,j}(w_{k+1,j}M_j - S_{i,j}^k - B_{k,j}) > \partial_+ U_{i,j}(w_{k+1,j}M_j - S_{i,j}^k - B_{k,j})$, for $i\leq k<k^0$.

PROOF. From the discussion above, we can easily derive that $\partial_- U_{i,j}(w_{k+1,j}M_j - S_{k,j} - B_{k,j}) = \frac{w_{i,j}}{w_{k+1,j}}$ and $\partial_+ U_{i,j}(w_{k+1,j}M_j - S_{i,j}^k - B_{k,j}) = w_{i,j}M_j \frac{S_{i,j}^k + B_{k,j}}{(w_{k+1,j}M_j)^2} = w_{i,j} \frac{S_{i,j}^k + B_{k,j}}{w_{k+1,j}^2 M_j}$. So, we need to prove that $\frac{w_{i,j}}{w_{k+1,j}} > w_{i,j} \frac{S_{i,j}^k + B_{k,j}}{w_{k+1,j}^2 M_j}$, or, equivalently, $w_{k+1,j}M_j > S_{i,j}^k + B_{k,j}$, which is true, since the point with

$B_1 = w_{k+1,j}M_j - S_{i,j}^k - B_{k,j}$ (point P_1 in Figures 3 and 4) satisfies $B_1>0$. Note that it cannot be $w_{k+1,j}M_j = S_{i,j}^k + B_{k,j}$ because we assumed that $k<k^0$ so $B_1>0$. □

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

PROOF. We have that $\partial_- U_{i,j}(w_k, jM_j - S_{i,j}^k - B_{k,j}) = w_{i,j}M_j \frac{S_{i,j}^k + B_{k.j}}{(w_{k,j}M_j)^2} = w_{i,j} \frac{S_{i,j}^k + B_{k,j}}{w_{k,j}^2 M_j}$, and $\partial_+ U_{i,j}(w_{k,j}M_j - S_{i,j}^k - B_{k,j}) = \frac{w_{i,j}}{w_{k,j}}$. So, we need to prove that $w_{i,j} \frac{S_{i,j}^k + B_{k,j}}{w_{k,j}^2 M_j} < \frac{w_{i,j}}{w_{k,j}}$, or, equivalently, $w_{k,j} > \frac{S_{i,j}^k + B_{k,j}}{M_j}$. But in Lemma 3 we proved that $w_{k+1,j} > \frac{S_{i,j}^k + B_{k,j}}{M_j}$. Since $w_{k,j} > w_{k+1,j}$, the inequality we want to prove follows immediately. □

Lemma 3 implies that whenever we make a transition from the linear to the concave part (e.g., point P_1 in Figures 3 and 4) the first derivative decreases, 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 in Figures 3 and 4), the first derivative increases, violating concavity.

6.2.3. Constructing the upper-concave envelope $\widetilde{U}_{i,j}$ when $k^0>i$. Based on the observation that the transition points are the reason why $U_{i,j}$ is not concave, we now show how to construct the upper-concave envelope $\widetilde{U}_{i,j}$.

The case when $i<k<k^0$. First, consider region R_k, $i<k<k^0$. The idea is to draw a line ε_{tang} from P_3 to the point P *to the right* of P_3 in the concave part of region R_k so that the line ε_{tang} is tangent to the curve $w_{i,j}M_j \frac{B_{i,j}}{\sum_{i'=0}^k B_{i',j}}$. Given on our previous discussion, the derivative at P is $w_{i,j}M_j \frac{S_{i,j}^k + B_{k,j}}{(S_{i,j}^k + B_{k,j} + B)^2}$. On the other hand, the slope of ε_{tang} is $\frac{U_{i,j}(B_3) - U_{i,j}(B)}{B_3 - B} = w_{i,j} \frac{\frac{B_3}{w_{k,j}} - M_j \frac{B}{S_{i,j}^k + B_{k,j} + B}}{B_3 - B}$. Thus, we are looking for a B such that $w_{i,j}M_j \frac{S_{i,j}^k + B_{k,j}}{(S_{i,j}^k + B_{k,j} + B)^2} = w_{i,j} \frac{\frac{B_3}{w_{k,j}} - M_j \frac{B}{S_{i,j}^k + B_{k,j} + B}}{B_3 - B}$. But $B_3 = w_{k,j}M_j - S_{i,j}^k$, so, after some algebraic manipulations, the previous equation becomes:

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

Equation $(*)$ is a quadratic equation, which accepts the two solutions:

$$B = \frac{(M_j w_{k,j} - S_{i,j}^k)(S_{i,j}^k + B_{k,j}) \pm \sqrt{(M_j w_{k,j} - S_{i,j}^k)(S_{i,j}^k + B_{k,j}) S_{i,j}^k M_j w_{k,j}}}{S_{i,j}^k}.$$

First, note that $w_{k,j}M_j > S_{i,j}^k$ (since $B_3>0$), so the solutions are real numbers. Second, we keep the solution with the minus because it is lower than $B_2 = w_{k,j}M_j - S_{i,j}^k - B_{k,j}$. Indeed, after performing some algebraic manipulations we get: $\frac{(M_j w_{k,j} - S_{i,j}^k)(S_{i,j}^k + B_{k,j}) - \sqrt{(M_j w_{k,j} - S_{i,j}^k)(S_{i,j}^k + B_{k,j}) S_{i,j}^k M_j w_{k,j}}}{S_{i,j}^k} < M_j w_{k,j} - S_{i,j}^k - B_{k,j}$. The last expression is equivalent to $M_j w_{k,j} > S_{i,j}^k + B_{k,j}$, which is true. Note that the solution with the plus corresponds to the point *to the left* of P_3, so that the line from P_3 to that point is tangent to the curve $w_{i,j}M_j \frac{B_{i,j}}{\sum_{i'=0}^k B_{i',j}}$.

Note that there are 2 cases. If point P is greater than B_1 (see Figure 3), then we draw the line *segment* from P to P_3. Else (see Figure 4), we draw the line segment from P_1 to P_3. To summarize, we always take the line segment from P' to P_3, where the x-coordinate B' of P' satisfies $B' = \max\{B, B_1\}$.

Finally, we form $\widetilde{U}_{i,j}$ in region R_k as follows: For $B_{i,j} \in [B_1, B']$ (which may be a singleton set if $B' = B_1$) we set $\widetilde{U}_{i,j} = U_{i,j}$. And for $B_{i,j} \in [B', B_3]$, we set $\widetilde{U}_{i,j}$ to be equal to the line segment P'-P_3.

We now show that the slope of the line segment P'-P_3 is greater than the right derivative at P_3 in both aforementioned cases. To prove this, note first that the slope of the line segment can never be lower than the derivative at P: if $B'=B$ they are obviously equal; else, $B'=B_1>B$ and the slope of P_3-P_1 is greater than that of P_3-P. So, to prove the claim it suffices to show that the derivative at P is greater than the right derivative at P_3. Indeed, the derivative at P is $w_{i,j}M_j \frac{S_{i,j}^k+B_{k,j}}{\left(S_{i,j}^k+B_{k,j}+B\right)^2}$, while the right derivative at P_3 is $w_{i,j}M_j \frac{S_{i,j}^k}{\left(w_{k,j}M_j\right)^2}$. But then $w_{i,j}M_j \frac{S_{i,j}^k+B_{k,j}}{\left(S_{i,j}^k+B_{k,j}+B\right)^2} > w_{i,j}M_j \frac{S_{i,j}^k}{\left(S_{i,j}^k+B_{k,j}+B\right)^2} > w_{i,j}M_j \frac{S_{i,j}^k}{(S_{i,j}^k+B_{k,j}+(M_j w_{k,j}-S_{i,j}^k-B_{k,j}))^2} = w_{i,j}M_j \frac{S_{i,j}^k}{(M_j w_{k,j})^2}$, which proves our claim.

To conclude the case of $i<k<k^0$, we show that the slope of the segment P'-P_3 is lower than the left derivative at P'. (1) if $P'=P>P_1$, then the slope of the segment is equal to the right derivative at P, which is equal to the left derivative. (2) If $P'=P_1$, then first observe that the slope of the segment P'-P_3 is lower than the slope of the line segment P_2-P_3. Indeed, consider a moving point P_0 in the segment P_2-P_1. As P_0 moves from P_1 to P_2, the slope of P_3-P_0 decreases until P_0 reaches point P; from thereafter, it starts increasing again. But since in case (2) $P<P_1=P'$, it must be that the slope of P'-P_3 is lower than the slope of P_2-P_3. Moreover, from Lemma 3 the left derivative at P_1 is higher than the slope of P_2-P_3, which together with the previous observation proves our claim.

The case when $k=k^0>i$. (1) If region R_{k^0} only consist of a linear part, then in region R_{k^0}, we set $\widetilde{U}_{i,j}$ to be the same as $U_{i,j}$ (2) If it consists of both a concave and a linear part, we proceed exactly as in the case $i<k<k^0$, but set B' as $B' = \max\{B, 0\}$.

The case when $k=i$. In region R_i, we set $\widetilde{U}_{i,j}$ to be the same as $U_{i,j}$.

We can now prove the following:

PROPOSITION 3. The function $\widetilde{U}_{i,j}(B_{i,j}; \boldsymbol{B_{-i,j}})$ that we constructed for $k^0>i$ is an upper-concave envelope of $U_{i,j}(B_{i,j}; \boldsymbol{B_{-i,j}})$.

PROOF. The utility function $\widetilde{U}_{i,j}$ is continuous everywhere, and the left and right derivatives exist for all $B_{i,j}\geq 0$ and are monotonically non-increasing in the allocated budget $B_{i,j}$. Furthermore, the left derivative at any $B_{i,j}\geq 0$ is equal to or greater than the right derivative. But then $\widetilde{U}_{i,j}$ will be concave in terms of $B_{i,j}$ [Gordon 2001]. Moreover, $\widetilde{U}_{i,j} \geq U_{i,j}$ for every $B_{i,j} \geq 0$. Finally, by construction $\widetilde{U}_{i,j}$ either coincides with $U_{i,j}$, or consists of line segments that connect the endpoints of consecutive parts where $\widetilde{U}_{i,j}$ coincides with $U_{i,j}$. Now, assume $\widetilde{U}_{i,j}$ is not an upper-concave envelope, then there will be a function $U_{i,j}^{up}$ with $U_{i,j}^{up}(B) < \widetilde{U}_{i,j}(B)$ for some $B\geq 0$. Obviously, that B cannot correspond to the parts of $\widetilde{U}_{i,j}$ that coincide with $U_{i,j}$. Thus, it must correspond to a point in the interior of the line segments that connect the parts that coincide. But a real-valued function is concave if it never lies beneath the line that connects any two points in the function [Bertsekas 2009], which leads to a contradiction. Thus, $\widetilde{U}_{i,j}$ is an upper-concave envelope of $U_{i,j}$. ☐

6.2.4. Defining relaxations of $U_{i,j}$ and U_i. In the above construction, we considered budgets $\boldsymbol{B_{-i,j}}$ fixed and obtained the upper-concave envelope of $U_{i,j}(B_{i,j}; \boldsymbol{B_{-i,j}})$. Consider now a set of non-negative allocations $B_{1,j}, \dots, B_{i,j}, \dots, B_{N,j}$ in cell j, and define

$$U_{i,j}^{rel}\left(B_{1,j}, \dots, B_{i,j}, \dots, B_{N,j}\right) = \widetilde{U}_{i,j}\left(B_{i,j}; \boldsymbol{B_{-i,j}}\right).$$

The new function has two useful properties. First, by its own definition for any fixed value of $\boldsymbol{B}_{-i}$ it is concave in $B_{i,j}$. Second, it is continuous in $B_{1,j}, \dots, B_{i,j}, \dots, B_{N,j}$. To understand why this is true, note that for two points $(B_{1,j}, \dots, B_{i,j}, \dots, B_{N,j})$ and $B'_{1,j}, \dots, B'_{i,j}, \dots, B'_{N,j}$ that are arbitrarily close, the transition points of $U_{i,j}(B_{i,j}; \boldsymbol{B_{-i,j}})$ and $U_{i,j}(B'_{i,j}; \boldsymbol{B'_{-i,j}})$ must be arbitrarily close as well. This is a consequence of the fact that transitions are continuous, as we showed in Lemma 1. Now, recall that we constructed the upper-concave envelope through a finite number of continuous operations on function $U_{i,j}(B_{i,j}; \boldsymbol{B_{-i,j}})$ and the transition points. But then, since $U_{i,j}(B_{i,j}; \boldsymbol{B_{-i,j}})$ and the transition points are continuous, so is $U_{i,j}^{rel}\left(B_{1,j}, \dots, B_{i,j}, \dots, B_{N,j}\right)$.

Finally, we define for any set of valid budget allocations $(\boldsymbol{B_1}, \dots, \boldsymbol{B_N})$ the relaxation of $U_i = \sum_{j=1}^{L} U_{i,j}$ as:

$$U_i^{rel}(\boldsymbol{B_1}, \dots, \boldsymbol{B_N}) = \sum_{j=1}^{L} U_{i,j}^{rel}\left(B_{1,j}, \dots, B_{i,j}, \dots, B_{N,j}\right).$$

The new function is continuous $(\boldsymbol{B_1}, \ldots, \boldsymbol{B_N})$ in as a sum of continuous functions. Moreover, it is concave in $\boldsymbol{B}_i$ for any fixed value of $\boldsymbol{B}_{-i}$ as a sum of functions that are concave in $\boldsymbol{B}_i$ for any fixed value of $\boldsymbol{B}_{-i}$.

6.2.5. Bounding the maximum deviation in game G^{ε}. Now, let $\tilde{G}^{\varepsilon}$ be the new game when the utility functions are replaced by their relaxations $U_i^{rel}(\boldsymbol{B_1}, \ldots, \boldsymbol{B_N})$.

PROPOSITION 4. The game $\tilde{G}^{\varepsilon}$ always accepts a Nash equilibrium.

PROOF. The utility functions $U_i^{rel}(\boldsymbol{B}_i; \boldsymbol{B}_{-i})$, $1 \le i \le N$, possess two properties: (1) each $U_i^{rel}(\boldsymbol{B}_i; \boldsymbol{B}_{-i})$ is continuous in $(\boldsymbol{B}_i; \boldsymbol{B}_{-i})$, and (2) each $U_i^{rel}(\boldsymbol{B}_i; \boldsymbol{B}_{-i})$ is concave in $\boldsymbol{B}_i$ for any fixed value of $\boldsymbol{B}_{-i}$. Moreover, the strategy space of every advertiser is convex, closed and bounded. Consequently, based on Rosen's theorem [Rosen 1965] we can immediately derive that a Nash equilibrium exists. □

Let $\widetilde{\boldsymbol{B}} = (\widetilde{\boldsymbol{B}}_1, \ldots, \widetilde{\boldsymbol{B}}_N)$ be a Nash equilibrium of game $\tilde{G}^{\varepsilon}$. Note that $\widetilde{\boldsymbol{B}}$ may not be an equilibrium of game G^{ε}. This means that there may be players in game G^{ε} who have an incentive to deviate if the strategy vector $\widetilde{\boldsymbol{B}}$ is chosen. However, the following lemma shows that we can bound the maximum utility that a player can gain by deviating.

PROPOSITION 5. Let the strategy vector $\widetilde{\boldsymbol{B}}$ be a Nash equilibrium of game $\tilde{G}^{\varepsilon}$. Then the maximum utility that player i can gain by deviating from $\widetilde{\boldsymbol{B}}$ in game G^{ε} is $U_i^{rel}(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i}) - U_i(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i})$.

PROOF. Assume that advertisers play according to the strategy vector $\widetilde{\boldsymbol{B}}$ in game G^{ε}. Furthermore, assume that there as a deviation strategy $\boldsymbol{B}_i$ for player i that will result in a utility gain higher than $U_i^{rel}(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i}) - U_i(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i})$, i.e., $U_i(\boldsymbol{B}_i; \widetilde{\boldsymbol{B}}_{-i}) - U_i(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i}) > U_i^{rel}(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i}) - U_i(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i})$, or equivalently, $U_i(\boldsymbol{B}_i; \widetilde{\boldsymbol{B}}_{-i}) > U_i^{rel}(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i})$. But since $U_i^{rel} \ge U_i$, we will have that $U_i^{rel}(\boldsymbol{B}_i; \widetilde{\boldsymbol{B}}_{-i}) \ge U_i(\boldsymbol{B}_i; \widetilde{\boldsymbol{B}}_{-i}) > U_i^{rel}(\widetilde{\boldsymbol{B}}_i; \widetilde{\boldsymbol{B}}_{-i})$. But then player i has an incentive to deviate to strategy $\boldsymbol{B}_i$ in game $\tilde{G}^{\varepsilon}$, thus contradicting the assumption that $\widetilde{\boldsymbol{B}}$ is a Nash equilibrium of game $\tilde{G}^{\varepsilon}$. □

Essentially, according to the above result, we can find a set of budget allocations such that we can know exactly the maximum utility that an advertiser may gain by deviating. An immediate consequence of Proposition 5 is that if the maximum profitable deviation from $\widetilde{\boldsymbol{B}}$ in game G^{ε} is 0 for all advertisers, then $\widetilde{\boldsymbol{B}}$ is also a Nash equilibrium of game G^{ε}.

Regarding the computation of the equilibrium in $\tilde{G}^{\varepsilon}$, we can use the local greedy adjustment method of Section 5, albeit with one caveat. Because utilities are not everywhere differentiable, marginal utility can be either left or right, depending on whether they are defined in terms of the left or the right partial derivative discussed before. In particular, in every round we identify the cell c_h with the highest right marginal utility and the cell c_l with the lowest left marginal utility. If the right marginal utility of cell c_h is greater than the left marginal utility of cell c_l, with the lowest left marginal utility, we then move a fixed small amount of money δ from c_l to c_h. Given that all U_i^{rel} is concave and continuous, the method always converges to an equilibrium within a ratio that depends on δ.

An important line of future work is to investigate the conditions under which the perturbed game G^{ε} has an equilibrium. Now, assume for a given set of advertisers' bids and budgets, G^{ε} has an equilibrium for every $\varepsilon > 0$. We can then consider a sequence of $\varepsilon \to 0$, and take the sequence of corresponding equilibria. Since the strategy space is compact, the sequence of equilibria has an infinite subsequence that converges to a limit point [Gordon 2001]. Unfortunately, this does not imply that the original game without the fictitious advertiser accepts the limit point as an equilibrium. The reason is that the utilities in G^{ε} are continuous everywhere, while the utilities in the original game are not continuous at 0. In the budgets-case, we discussed how strong competition helps us circumvent this problem. It would be interesting to investigate in future work whether strong competition guarantees the existence of a Nash equilibrium when both bids and budgets are present.

7. CONCLUSION

In this work, we investigate location-based sponsored search advertising, assuming a grid structure with cells and an offline setting where the number of queries per cell is known in advance. Advertisers have different valuations for different cells, and they optionally have daily budget constraints. We explored 3 cases: (1) bids-case, (2) budgets-case, and (3) bids-and-budgets-case, and analyzed the Nash Equilibrium in the corresponding markets using game theoretical tools.

There are several directions for future research. Important issues for the budget-case include the extension of our techniques to the case of multiple slots, and the design of truthful procedures. It would also be interesting to deal with the more challenging online setting, where the expected number of queries per context and time unit is not available in advance. Finally, we would like to deploy the methods on a real search engine, or conduct extensive simulations using real valuation data, in order to gain a better insight into their real world performance.

REFERENCES

Syagnik Banerjee and Ruby R. Dholakia. 2008. Does location-based advertising work? *International Journal of Mobile Marketing*, 3(1).

J. Bai, J.-Y. Nie, G. Cao, and H. Bouchard. 2007. Using query contexts in information retrieval. *ACM SIGIR Conference on Research and Development in Information Retrieval*, 15-22.

Dimitri P. Bertsekas. 2009. *Convex optimization theory*. Athena Scientific.

William C. Brainard and Herbert E. Scarf. 2000. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper*, No. 1272.

Gordon C. Bruner and Anand Kumar. 2007. Attitude toward location-based advertising. *Journal of Interactive Advertising*, 7(2).

Edward H. Clarke. 1971. Multipart pricing of public goods. *Public Choice*, 11(1):17–33.

G. Demange. 1982. Strategyproofness in the assignment market game. Laboratoire d'Econometrie de l'Ecole Polytechnique.

G. Demange, D. Gale, and M. Sotomayor. 1986. Multi-item auctions. *Journal of Political Economy*, 94(4):863-872.

Subhankar Dhar and Upkar Varshney. 2011. Challenges and business models for mobile location-based services and advertising. *Communications of the ACM*, 54(5):121–129.

Shahar Dobzinski, Noam Nisan, and Michael Schapira. 2005. Approximation algorithms for combinatorial auctions with complement-free bidders. *ACM Symposium on Theory of Computing*.

Shahar Dobzinski. 2011. An impossibility result for truthful combinatorial auctions with submodular valuations. *ACM Symposium on Theory of Computing*.

David Easley and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, Cambridge.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. 2007. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 9(1):242–259.

Jon Feldman and S. Muthukrishnan. 2008. Algorithmic methods for sponsored search advertising. *SIGMETRICS Tutorial*, *Performance Modeling and Engineering*.

Jon Feldman, S. Muthukrishnan, Evdokia Nikolova, and Martin Pál. 2008. A truthful mechanism for offline ad slot scheduling. *Symposium on Algorithmic Game Theory*.

Michal Feldman, Kevin Lai, and Li Zhang. 2009. The proportional-share allocation market for computational resources. *IEEE Transactions on Parallel and Distributed Systems*, 20(8):1075–1088.

Arpita Ghosh, Hamid Nazerzadeh, and Mukund Sundararajan. 2007. Computing optimal bundles for sponsored search. *International Workshop on Internet and network economics*.

Jerry Green and Jean-Jacques Laffont. 1977. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2):427-438.

Theodore Groves. 1973. Incentives in teams. *Econometrica*, 41:617–631.

Matthew O. Jackson. 2003. *Mechanism Theory*. Encyclopedia of Life Support Systems. EOLLS Publishers, Oxford.

Bernard J. Jansen, and Tracy Mullen. 2008. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business*, 6(2):114–131.

Bernard J. Jansen and Lauren Solomon. 2010. Gender demographic targeting in sponsored search. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.

Ramesh Johari and John N. Tsitsiklis. 2004. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435.

Shizuo Kakutani. 1941. A generalization of Brouwer's fixed point theorem. *Duke Mathematical Journal*, 8:457–459.

Frank Kelly. 1997. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37.

David Kempe and Mohammad Mahdian. 2008. A cascade model for externalities in sponsored search. *International Workshop on Internet and Network Economics*.

Vijay Krishna. 2002. *Auction theory*. Academic Press.

Herman B. Leonard. 1983. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461-479.

Paul Milgrom. 2004. *Putting Auction Theory to Work*. Cambridge University Press, Cambridge.

Tao Qin, Wei Chen, and Tie-Yan Liu. 2015. Sponsored search auctions: recent advances and future directions. *ACM Transactions on Intelligent Systems and Technology*, 5(4).

J. B. Rosen. 1965. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica*, 33(3):520–534.

Lloyd S. Shapley and Martin Shubik. 1971/72. The assignment game I: The core. *International Journal of Game Theory*, 1(1):111-130.

George Trimponias, Ilaria Bartolini, and Dimitris Papadias. 2013. Location-based sponsored search advertising. *Symposium on Spatial and Temporal Databases*.

Hal R. Varian. 2007. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178.

William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Finance*, 16:8–27.