Autonomous Agents and Multi-Agent Systems manuscript No.
(will be inserted by the editor)

Pricing Ad Slots with Consecutive Multi-unit Demand

Xiaotie Deng · Paul Goldberg · Yang Sun · Bo Tang · Jinshan Zhang

the date of receipt and acceptance should be inserted later

Abstract We consider the optimal pricing problem for a model of the rich media advertisement market, that has other related applications. Our model differs from traditional position auctions in that we consider buyers whose demand might be multiple *consecutive* slots, which is motivated by modeling buyers who may require these to display a large size ad. We study three major pricing mechanisms, the Bayesian pricing model, the maximum revenue market equilibrium model and an envy-free solution model. Under the Bayesian model, we design a polynomial-time computable truthful mechanism that optimizes the revenue. For the market equilibrium paradigm, we find a polynomial-time algorithm to obtain the maximum revenue market equilibrium solution. In the envy-free setting, an optimal solution is presented for the case where the buyers have the same demand for the number of consecutive slots. We present results of a simulation that compares the revenues from the above schemes.

Keywords Mechanism Design · Revenue · Advertisement auction

1 Introduction

Ever since the pioneering studies on pricing protocols for sponsored search advertisement, especially with the generalized second price auction (GSP), by Edelman,

X. Deng
Department of Computer Science, Shanghai Jiaotong University, China.
E-mail: dengxiaotie@gmail.com

P. Goldberg
Department of Computer Science, Oxford University, UK.
E-mail: Paul.Goldberg@cs.ox.ac.uk

Y. Sun
Department of Computer Science, City University of Hong Kong, Hong Kong.
E-mail: sunyang@live.hk

B. Tang · J.Zhang
Department of Computer Science, University of Liverpool, UK.
E-mail: {Bo.Tang, Jinshan.Zhang}@liverpool.ac.uk

Ostrovsky, and Schwarz [15], as well as Varian [41], market making mechanisms have attracted much attention from the research community in understanding their effectiveness for the revenue maximization task facing platforms providing Internet advertisement services. In the traditional advertisement setting, advertisers negotiate ad presentations and prices with website publishers directly. An automated pricing mechanism simplifies this process by creating a bidding game for the buyers of advertisement space over an IT platform. It creates a complete competition environment for the price discovery process. Accompanying the explosion of the online advertisement business, there is a need for a complete picture of the pricing methods that may be used by advertisers and ad space providers.

In addition to search advertisements, display advertisements have been widely used in commercial webpages. They have a rich format of displays such as text ads and rich media ads. Compared with sponsored search, there is a lack of systematic studies of the associated mechanism design issues. The market maker faces a combinatorial problem of whether to assign a large space to one large rich media ad or multiple small text ads, as well as how to decide on the prices charged for them. We present a study of the allocation and pricing mechanisms for displaying slots in this environment where some buyers would like to have one slot and others may want several consecutive slots in a display panel. In addition to webpage ads, another motivation for our study is TV advertising where inventories of a commercial break are usually divided into slots of a few seconds each, and slots have various qualities measuring their expected number of viewers and the corresponding attractiveness. However, current auction mechanisms might perform poorly under these additional constraints that only consecutive slots can be assigned to any one buyer. For example, GSP auctions cannot be applied to this setting where buyers may have different demand sets. Another example, when applying the famous VCG auctions to this setting, the extracted revenue is relatively low.

We discuss three types of mechanisms and consider the revenue maximization problem under these mechanisms, and compare their effectiveness in revenue maximization under a dynamic setting where buyers may change their bids to improve their utilities. Our results make an important step towards understanding the advantages and disadvantages of their uses in practice. Assume the ad supplier divides the ad space into small enough slots (pieces) such that each advertiser is interested in a position with a fixed number of *consecutive* pieces. In modeling values to the advertisers, we modify the position auction model from the sponsored search market [15,41] where each ad slot is measured by the Click Through Rates (CTR), with users' interest expressed by a click on an ad. Since display advertising is usually sold on a per impression (CPM) basis instead of a per click basis (CTR), the quality factor of an ad slot stands for the expected impression it brings per unit of time. Unlike in the traditional position auctions, people may have varying demands (need different spaces to display their ads) in a rich media ad auction for the market maker to decide on slot allocations and their prices. It should be mentioned that in real markets, bidders might have alternative advertising approaches which are external to the mechanism designer. This competition scenario has been studied in some restricted auction settings [33,11] but few results are obtained for advertising auctions. Therefore, we follow the traditional position auction models and assume that the buyers have no alternatives in the market, i.e. the auctioneer is a monopolist.

1.1 Our Modeling Approach

We have a set of *buyers* (advertisers) and a set of *items* to be sold (the ad slots on a web page). We address the challenge of computing prices that satisfy certain desirable properties. Next we describe the elements of the model in more detail.

Items. Our model considers the geometric organization of ad slots, which commonly has the slots arranged in some sequence (typically, from top to bottom in the right-hand side of a web page). The slots are of variable quality. In the study of sponsored search auctions, a standard assumption is that the quality (corresponding to click-through rate) is higher at the beginning of the sequence and then monotonically decreases. Here we consider a generalization where the quality may go down and up, subject to a limit on the total number of local maxima (which we call *peaks*), corresponding to focal points on the web page. As we will show later, without this limit the revenue maximization problem is NP-hard.

Buyers. A buyer (advertiser) may want to purchase multiple slots, so as to display a larger ad. Such slots should be *consecutive* in the sequence. Each buyer i has a fixed *demand* d_i, the number of slots she needs for her ad. Two important aspects of this are

- *sharp* multi-unit demand, that is, buyer i should be allocated d_i items, or none at all; there is no point in allocating any fewer
- *consecutiveness* of the allocated items, in the pre-existing sequence of items.

These constraints give rise to a new and interesting combinatorial pricing problem.

Valuations. We assume that each buyer i has a parameter v_i representing the value she assigns to a slot of unit quality. Valuations for multiple slots are additive, so that a buyer with demand d_i would value a block of d_i slots to be their total quality, multiplied by v_i. This is an extension of the valuation model considered by Edelman et al. [15] and Varian [41] in their seminal work for keywords advertising where the buyers are unit-demand.

Pricing mechanisms. Given the valuations and demands from the buyers, the market maker decides on a price vector for all slots and an allocation of slots to buyers, as an output of the market. The question is one of which output the market maker should choose to achieve certain objectives. We consider three approaches:

- *Truthful mechanism* whereby the buyers report their demands (publicly known) and values (private) to the market maker; then prices are set so as to ensure that the buyers have no incentive to report incorrect valuations. We give a revenue-maximizing approach (i.e., maximizing the total price paid), within this framework.
- *Competitive equilibrium* whereby we prescribe certain constraints on the prices so as to guarantee certain well-known notions of fairness and envy-freeness.
- *Envy-free solution* whereby we prescribe certain constraints on the prices and allocations so as to achieve envy-freeness, as explained below.

The mechanisms we exhibit are computationally efficient. We also perform experiments to compare the revenues obtained from these three mechanisms.

1.2 Related Work

Sponsored search advertising via auctions is one of the most popular forms of Internet monetization, accounting for a major part of search engines' revenue. Consequently, design and analysis of such auctions has drawn a lot of attention in artificial intelligence and electronic commerce. For example, in the pioneering work by Varian [41] and Edelman et al. [15], the generalized second price (GSP) auctions, used by Google, have been modeled as position auctions having several desirable properties. Variants of this model have been extensively studied from both theoretical and practical points of view [27,22,25]. Lahaie [28] analyzed the incentive, efficiency, and revenue properties of ad slot auction designs and also considered first- and second-price payment rules together with each of these allocation rules. Lucier et al. [31] considered the revenue of the GSP auction at equilibrium and proved that if agent values are drawn from identical regular distributions, then the GSP auction paired with an appropriate reserve price generates a constant fraction of the optimal revenue. Ashlagi et al. [2] investigated the effect of budget constraints in designing sponsored search auctions. They showed that a modification of the Generalized English Auction introduced in [15] is envy-free and Pareto optimal (a weaker condition than social welfare maximizing). In artificial intelligence, there has been also considerable research on ad auctions, with a focus on issues like auction design [6], revenue extraction [25], bidder strategies [5,17], and so on. For a more comprehensive understanding of ad auctions, we refer the reader to recent survey articles (e.g. [32], [37]).

Auctions play an important role in display advertising markets and one challenging problem is to design auctions for online ad exchange platforms. Most relevant is the literature on designing expressive auctions and clearing algorithms for online advertising [36,6]. Ghosh et al. [21] provided a framework to model the advertiser's problem of how to fulfill guaranteed contracts and gave a solution where the publisher bids on behalf of its guaranteed contracts in the spot market. Lahaie et al. [29] studied the bidding languages that admit scalable allocation and pricing algorithms in ad exchange markets. They also presented a discussion on the incentive properties of different pricing approaches. Balseiro et al. [3] studied the interactions among budget-constrained bidders in an Ad Exchange by adopting the notion of Fluid Mean Field Equilibrium. For more challenges and direction of display advertising markets, we refer the reader to the recent survey article by Korula et al. [26].

Our consideration of position auctions in the Bayesian setting fits in the general one-dimensional auction design framework. Our study considers continuous distributions on buyers' values. For discrete distributions, Cai et al. [8] presents an optimal mechanism for budget constrained buyers without demand constraints in multi-parameter settings and subsequently they also give a general reduction from revenue to welfare maximization in [9]; for buyers with both budget constraints and demand constraints, 2-approximate mechanisms [1] and 4-approximate mechanisms [4] exist in the literature.

Envy-free auction design for revenue maximization has been extensively studied in algorithmic game theory. Guruswami et al. [24] showed that it is APX-hard to maximize the revenue among envy-free outcomes in matching markets. After that, both positive [12] and negative results [7] have been proven for this problem without budget constraints. For position auctions, Tang and Zhang [40] studied the revenue maximization problems with budget constrained bidders.

There are extensive studies on multi-unit demand in economics, see for example [10,16]. Chen et al. [13] first considered sharp multi-unit demand, where a buyer with demand d should be allocated d items or none at all, but with no further combinatorial constraint, such as the consecutiveness constraint that we consider here. The sharp demand setting is in contrast with a "relaxed" multi-unit demand (in which one can buy a subset of at most d items), where it is well known that the set of competitive equilibrium prices is non-empty and forms a distributive lattice [23,39]. This immediately implies the existence of an equilibrium with maximum possible prices, that consequently maximizes the revenue. Demange, Gale, and Sotomayor [14] proposed a combinatorial dynamics which always converges to a revenue maximizing (or minimizing) equilibrium for unit demand; their algorithm can be easily generalized to relaxed multi-unit demand. A strongly related work to our consecutive settings is the work of Rothkopf et al. [38], where the authors presented a dynamic programming approach to compute the maximum social welfare of consecutive settings when all the qualities are the same. Hence, our dynamic programming approach for general qualities in Bayesian settings is a non-trivial generalization of their setting.

The auction design problem studied here can be viewed as a special case of the combinatorial auction design setting (c.f. [35]) where the buyers have general utility functions over all subsets of the items. Most work in this line of research focused on designing truthful auctions to maximize social welfare in the prior-free setting. Lehmann et al. [30] showed that the social welfare maximization problem with single-minded bidders is NP-hard. Recently, Feldman et al. [19] showed that, in the Bayesian setting, the optimal social welfare can be approximated within a factor of 2 by using posted prices mechanisms. However, in this work we consider the revenue maximization problem in a special case when buyers have consecutive demand.

1.3 Organization

This paper is organized as follows. In Section 2 we describe the details of our rich media ads model and the related solution concepts. In Section 3, we study the problem under the Bayesian model and provide a Bayesian Incentive Compatible auction with optimal expected revenue for the special case of the single peak in quality values of advertisement positions. Then in Section 4, we extend the optimal auction to the case with limited peaks/valleys and show that it is NP-hard to maximize revenue without this limit. Next, in Section 5, we turn to the full information setting and propose an algorithm to compute the competitive equilibrium with maximum revenue. In Section 6, NP-hardness of envy-freeness for consecutive multi-unit demand buyers is shown. We also design a polynomial time solution for the special case where all advertisers demand the same number of ad slots. The simulation is presented in Section 7.

2 Preliminaries

In our model, a rich media advertisement instance consists of n advertisers and m advertising slots. Each slot $j \in \{1, \dots, m\}$ is associated with a (publicly known) number q_j which can be viewed as the quality or the desirability of the slot. Each advertiser (or buyer) i wants to display her own ad that occupies d_i consecutive slots on the webpage. In addition, each buyer has a private number v_i representing her valuation and thus, the i-th buyer's value for item j is $v_{ij} = v_i q_j$.

In this paper, we will say that slot j is assigned to a set of buyer B, to denote that j is assigned to some buyer in B. We call the set of all slots assigned to B the allocation to B. In addition, a buyer will be called a winner if he succeeds in displaying his ad and a loser otherwise. We use the standard notation $[s]$ to denote the set of integers from 1 to s, i.e. $[s] = \{1, 2, \dots, s\}$. We sometimes use $\sum_i$ instead of $\sum_{i\in[n]}$ to denote the summation over all buyers and $\sum_j$ instead of $\sum_{j\in[m]}$ for items, and the terms $\mathrm{E}_{\mathbf{v}}$ and $\mathrm{E}_{v_{-i}}$ denote, respectively, an expectation over buyer valuation vectors $\mathbf{v}$ sampled from the prior, and an expectation over valuation vectors v_{-i}, representing buyers other than i.

The vector $\mathbf{v}$ of all the buyers' values is sometimes written as $(v_i; v_{-i})$, v_i being the i-th entry of $\mathbf{v}$ and v_{-i} the joint bids of all buyers other than i. We represent a feasible assignment by a vector $\mathbf{X} = (X_1, \dots, X_n)$, where X_i is the set of items allocated to buyer i, i.e. $j \in X_i$ denotes item j is assigned to buyer i. Thus we have $\{X_i\}$ is pairwise disjoint, which can be viewed as a partition of items. Given a fixed assignment $\mathbf{X}$, we use t_i to denote the total quality that buyer i is assigned, precisely, $t_i = \sum_{j\in X_i} q_j$. In general, when $\mathbf{X}$ is a function of buyers' bids $\mathbf{v}$, we define t_i to be a function of $\mathbf{v}$ such that $t_i(\mathbf{v}) = \sum_{j\in X_i(\mathbf{v})} q_j$.

When we say that slot qualities have a single peak, we mean that there exists a peak slot k such that for any slot $j < k$ on the left side of k, $q_j \geq q_{j-1}$ and for any slot $j > k$ on the right side of k, $q_j \geq q_{j+1}$.

2.1 Bayesian Mechanism Design

In Section 3 and 4, we assume that all buyers' values are distributed independently according to publicly known bounded distributions following the work of [34]. The distribution of each buyer i is represented by a Cumulative Distribution Function (CDF) F_i and a Probability Density Function (PDF) f_i. In addition, we assume that the concave and convex closure and integration of those functions can be computed efficiently.

An auction $M = (\mathbf{X}, \mathbf{p})$ consists of an allocation function $\mathbf{X}$ and a payment function $\mathbf{p}$. More precisely, $\mathbf{X}$ specifies the allocation of items to buyers and $\mathbf{p} = (p_i)_i$ specifies the buyers' payments, where both $\mathbf{X}$ and $\mathbf{p}$ are functions of the reported valuations $\mathbf{v}$. Our objective is to maximize the expected revenue of the mechanism $Rev(M) = \mathrm{E}_{\mathbf{v}}\left[\sum_i p_i(\mathbf{v})\right]$ by using Bayesian incentive compatible mechanisms.

Definition 1 A mechanism M is called *Bayesian Incentive Compatible* (BIC) iff the following inequalities hold for all i, v_i, v'_i.

$$\mathrm{E}_{v_{-i}}[v_i t_i(\mathbf{v}) - p_i(\mathbf{v})] \geq \mathrm{E}_{v_{-i}}\left[v_i t_i(v'_i; v_{-i}) - p_i(v'_i; v_{-i})\right] \quad (1)$$

Furthermore, we say M is *Incentive Compatible* if M satisfies a stronger condition that $v_i t_i(\mathbf{v}) - p_i(\mathbf{v}) \geq v_i t_i(v_i'; v_{-i}) - p_i(v_i'; v_{-i})$, for all $\mathbf{v}, i, v_i'$,

To put it in words, in a BIC mechanism, no player can improve her *expected* utility (expectation taken over other players' bids) by misreporting her value. An IC mechanism satisfies the stronger requirement that no matter what the other players declare, no player has incentives to deviate.

2.2 Competitive Equilibrium and Envy-free Solution

In Section 5, we study the revenue maximizing competitive equilibrium and envy-free solution in the full information setting instead of the Bayesian setting. An outcome of the market is a pair $(\mathbf{X}, \mathbf{p})$, where $\mathbf{X}$ specifies an allocation of items to buyers and $\mathbf{p}$ specifies prices paid. Given an outcome $(\mathbf{X}, \mathbf{p})$, recall $v_{ij} = v_i q_j$, let $u_i(\mathbf{X}, \mathbf{p})$ denote the *utility* of i.

Definition 2 A tuple $(\mathbf{X}, \mathbf{p})$ is a *consecutive envy-free pricing* solution if every buyer is consecutive envy-free, where a buyer i is consecutive envy-free if the following conditions are satisfied:

- if $X_i \neq \emptyset$, then (i) X_i is d_i consecutive items. $u_i(\mathbf{X}, \mathbf{p}) = \sum_{j \in X_i} (v_{ij} - p_j) \geq 0$, and (ii) for any other subset of consecutive items T with $|T| = d_i$, $u_i(\mathbf{X}, \mathbf{p}) = \sum_{j \in X_i} (v_{ij} - p_j) \geq \sum_{j \in T} (v_{ij} - p_j)$;
- if $X_i = \emptyset$ (i.e., i wins nothing), then, for any subset of consecutive items T with $|T| = d_i$, $\sum_{j \in T} (v_{ij} - p_j) \leq 0$.

In the literature, there have been two other notions of envy-free allocation, namely, sharp item envy-freeness [13] and bundle envy-freeness [18]. Sharp item envy-free requires that no buyer should envy any bundle of items whose size equals her demand, while bundle envy-free is the (weaker) stipulation that no-one should envy the bundle bought by any other buyer. Note that Definition 2 is about item envy-freeness; a weaker notion of consecutive bundle envy-freeness would require that no buyer should envy a bundle allocated entirely to some other buyer having the same demand.

Example 1 (Three types of envy-freeness) Suppose there are two buyers i_1 and i_2 with per-unit-quality $v_{i_1} = 10$, $v_{i_2} = 8$ and demands $d_{i_1} = 1$, $d_{i_2} = 2$. Three items j_1, j_2, j_3 have quality $q_{j_1} = q_{j_3} = 1$ and $q_{j_2} = 3$. The optimal solution of the three types of envy-freeness are as follows:

- The optimal consecutive envy-free solution, $X_{i_1} = \{j_3\}$, $X_{i_2} = \{j_1, j_2\}$ and $p_{j_1} = p_{j_3} = 6$ and $p_{j_2} = 26$ with total revenue 38;
- Optimal sharp item envy-free solution, $X_{i_1} = \{j_2\}$, $X_{i_2} = \{j_1, j_3\}$ and $p_{j_1} = p_{j_3} = 8$ and $p_{j_2} = 28$ with total revenue 44;
- Optimal (relaxed) bundle envy-free solution, $X_{i_1} = \{j_2\}$, $X_{i_2} = \{j_1, j_3\}$ and $p_{j_1} = p_{j_3} = 8$ and $p_{j_2} = 30$ with total revenue 46;

Definition 3 (Competitive Equilibrium) We say an outcome of the market $(\mathbf{X}, \mathbf{p})$ is a *competitive equilibrium* if it satisfies two conditions.

– $(\mathbf{X}, \mathbf{p})$ must be consecutive envy-free (Definition 2).
– The unsold items must be priced at zero.

We are interested in the revenue maximizing competitive equilibrium and envy-free solutions. It is well known that a competitive equilibrium always exists for unit demand buyers (even for general v_{ij} valuations) [39]. For our consecutive multi-unit demand model, however, a competitive equilibrium may not always exist as the following example shows.

Example 2 (Competitive equilibrium may not exist) There are two buyers i_1, i_2 with values $v_{i_1} = 10$ and $v_{i_2} = 9$, respectively. Let their demands be $d_{i_1} = 1$ and $d_{i_2} = 2$, respectively. There are two items j_1, j_2, both with unit quality $q_{j_1} = q_{j_2} = 1$. If i_1 wins an item, without loss of generality, say j_1, then j_2 is unsold and $p_{j_2} = 0$; by envy-freeness of i_1, we have $p_{j_1} = 0$ as well. Thus, i_2 envies the bundle $\{j_1, j_2\}$. On the other hand, if i_2 wins both items, then $p_{j_1} + p_{j_2} \leq v_{i_2 j_1} + v_{i_2 j_2} = 18$, implying that $p_{j_1} \leq 9$ or $p_{j_2} \leq 9$. Therefore, i_1 is not envy-free. Hence, there is no competitive equilibrium in the given instance.

In the unit demand case, it is well-known that the set of equilibrium prices forms a distributive lattice; hence, there exist extremes which correspond to the maximum and the minimum equilibrium price vectors. In our consecutive demand model, however, even if a competitive equilibrium exists, maximum equilibrium prices may not exist.

Example 3 (Maximum equilibrium need not exist) There are two buyers i_1, i_2 with values $v_{i_1} = 1$, $v_{i_2} = 10$ and demands $d_{i_1} = 1, d_{i_2} = 2$, and two items j_1, j_2 with unit quality $q_{j_1} = q_{j_2} = 1$. It can be seen that allocating the two items to i_2 at prices $(19, 1)$ or $(1, 19)$ are both revenue maximizing equilibria; but there is no equilibrium price vector that is at least both $(19, 1)$ and $(1, 19)$.

Because of the consecutive multi-unit demand, it is possible that some items are 'over-priced'; this is a significant difference between consecutive multi-unit and unit demand models. Formally, in a solution $(\mathbf{X}, \mathbf{p})$, we say an item j is *over-priced* if there is a buyer i such that $j \in X_i$ and $p_j > v_i q_j$. That is, the price charged for item j is larger than its contribution to the utility of its winner.

Example 4 (Over-priced items) There are two buyers i_1, i_2 with values $v_{i_1} = 20, v_{i_2} = 10$ and demands $d_{i_1} = 1$ and $d_{i_2} = 2$, and three items j_1, j_2, j_3 with qualities $q_{j_1} = 3, q_{j_2} = 2, q_{j_3} = 1$. We can see that the allocations $X_{i_1} = \{j_1\}, X_{i_2} = \{j_2, j_3\}$ and prices $(45, 25, 5)$ constitute a revenue maximizing envy-free solution with total revenue 75, where item j_2 is over-priced. If no items are over-priced, the maximum possible prices are $(40, 20, 10)$ with total revenue 70.

3 Optimal Auction for the Single Peak Case

The goal of this section is to present our optimal auction for the single peak case that serves as an elementary component in the general case later. En route, several principal techniques are examined exhaustively to the extent that they can be applied directly in the next section. With these techniques, we show that the optimal

Bayesian Incentive Compatible auction can be represented by a simple Incentive Compatible one. Furthermore, this optimal auction can be implemented efficiently. Given some mechanism, let $T_i(v_i) = \mathrm{E}_{v_{-i}}[t_i(\mathbf{v})]$ and $P_i(v_i) = \mathrm{E}_{v_{-i}}[p_i(\mathbf{v})]$. We also use $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$ to denote the virtual value of buyer i when his valuation is v_i. Similar to Myerson [34], we suppose $P_i(\underline{v}_i) = \underline{v}_i T_i(\underline{v}_i)$ otherwise the revenue can be improved by increasing $P_i(\underline{v}_i)$. From Myerson [34], we obtain the following three lemmas.

Lemma 1 (From [34]) *A mechanism $M = (\mathbf{X}, \mathbf{p})$ is Bayesian Incentive Compatible if and only if, letting $\underline{v}_i$ be a lower bound on values taken by v_i:*
a) $T_i(v_i)$ is monotone non-decreasing for any agent i,
b) $P_i(v_i) = v_i T_i(v_i) - \int_{\underline{v}_i}^{v_i} T_i(z)dz$.

Lemma 2 (From [34]) *For any BIC mechanism $M = (\mathbf{X}, \mathbf{p})$, the expected revenue $\mathrm{E}_{\mathbf{v}}[\sum_i P_i(v_i)]$ is equal to the virtual surplus $\mathrm{E}_{\mathbf{v}}[\sum_i \phi_i(v_i)t_i(\mathbf{v})]$.*

We assume $\phi_i(\cdot)$ is monotone increasing, i.e. the distribution is regular. If not, Myerson's ironing technique can be applied to make $\phi_i(\cdot)$ monotone — it is here that we invoke our assumption that we can efficiently compute the convex closure of a continuous function and integration. Lemma 3 follows directly from Lemmas 1 and 2.

Lemma 3 *Suppose that $\mathbf{X}$ is the allocation function that maximizes $\mathrm{E}_{\mathbf{v}}[\phi_i(v_i)t_i(\mathbf{v})]$ subject to the constraints that $T_i(v_i)$ is monotone non-decreasing for any fixed profile v_{-i} of the other bidders, and any agent i is assigned either d_i consecutive slots or nothing. Suppose also that*

$$p_i(\mathbf{v}) = v_i t_i(\mathbf{v}) - \int_{\underline{v}_i}^{v_i} t_i(v_{-i}, s_i) ds_i \tag{2}$$

Then $(\mathbf{X}, \mathbf{p})$ represents an optimal mechanism for the consecutive multiple-slot ad auction problem.

We will use dynamic programming to maximize the virtual surplus in Lemma 2. First we show the following useful lemma which states that all the allocated slots are consecutive (see Figure 1).

Lemma 4 *Suppose that ad slot qualities have the single-peak property. There exists an optimal allocation $\mathbf{X}$ that maximizes $\sum_i \phi_i(v_i)t_i(\mathbf{v})$, having the following property. For any unassigned slot j, it must be that either $\forall j' > j$, slot j' is unassigned or $\forall j' < j$, slot j' is unassigned.*

Proof Let $\mathbf{X}$ be an allocation maximising the sum of the virtual values. If $\mathbf{X}$ does not satisfy the property in the statement of the lemma, we show how to modify $\mathbf{X}$ so as to satisfy the property while preserving optimality. Let j be a slot that violates the property, in that there are allocated slots to either side of j. Letting k be a maximum-quality slot, assume $j \geq k$ (where the case $j \leq k$ is similar). Modify $\mathbf{X}$ by taking all buyers allocated slots to the right of j, and moving their allocation one position to the left. The sum of virtual valuations cannot decrease, since any buyer's allocation is to slots having at least as high quality. Also, the modification

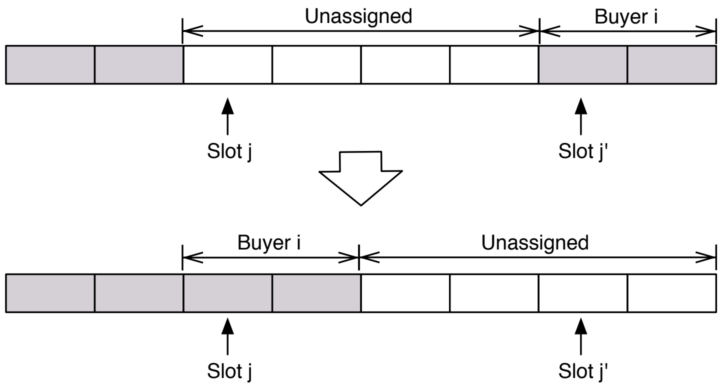

Fig. 1 By reassigning i the slots from j, we make the set of assigned slots consecutive.

reduces the number of slots violating the condition, so if applied repeatedly, we end up with an optimal allocation having the claimed property. □

By renumbering the bidders in the order of the virtual values, we can assume all the buyers are sorted in a non-increasing order according to their virtual values. Next, we prove that the consecutiveness of allocated slots holds for any buyers sets $[s] \subseteq [n]$. That is, Lemma 5 says that there exists an optimal allocation that always allocates the first s buyers consecutive slots, for $s \in [n]$.

Lemma 5 *Suppose that ad slot qualities have the single-peak property. There exists an optimal allocation* $\mathbf{X}$ *having the following property. For any slot* j *not allocated to buyers in* $[s]$*, we either have* $\forall j' > j$*, slot* j' *is not allocated to any member of* $[s]$*, or* $\forall j' < j$*, slot* j' *is not allocated to any member of* $[s]$*.*

Proof The idea is to pick an arbitrary optimal allocation $\mathbf{X}$ and modify it to the desired one. Suppose $\mathbf{X}$ does not satisfy the property on a subset $[s]$. By Lemma 4, there is no unassigned slots in the middle of allocations to set $[s]$. Then there must be a slot assigned to a buyer i not in the set $[s]$ that separates the allocations to $[s]$. We use W_i to denote the allocated slots of buyer i. Suppose slot k is the peak. There are two cases to be considered:

Case 1. $k \notin W_i$. Let j and j' be the leftmost and rightmost slot in W_i respectively. We consider two cases $q_j \geq q_{j'}$ and $q_j < q_{j'}$. We only prove for the first case and the proof for the other case is symmetric. If $q_j \geq q_{j'}$, we find the leftmost slot $j_1 > j'$ assigned to $[s]$ and the rightmost slot $j_2 < j_1$ not assigned to $[s]$. In addition, let $i_1 \in [s]$ be the buyer that j_1 is assigned to and $i_2 > s$ be the buyer that j_2 is assigned to. In single peak case, it is easy to check $q_j \geq q_{j'}$ implies that all the slots assigned to i_2 have higher quality than i_1's. Thus swapping the positions of i_1 and i_2 will always increase the virtual surplus, $\sum_i \phi_i(v_i)t_i(\mathbf{v})$ as illustrated in Figure 2. By continuing to do this, we can eliminate all slots not allocated to $[s]$ in the middle of allocation to $[s]$ and attain the desired optimal solution.

Case 2. $k \in W_i$. Suppose $W_i = \{j_1^i, j_2^i, \cdots, j_{u_i}^i\}$ with $j_1^i < j_2^i < \cdots < j_{u_i}^i$ and there exists $1 \leq e \leq u_i$ such that $k = j_e^i$. Let a and b be the left and right neighbour buyers of i winning slots next to W_i. As we know $a, b \in [s]$, hence, $\phi_a(v_a) \geq \phi_i(v_i)$ and $\phi_b(v_b) \geq \phi_i(v_i)$. Let $W_a = \{j_1^a, j_2^a, \cdots, j_{u_a}^a\}$ and $W_b = \{j_1^b, j_2^b, \cdots, j_{u_b}^b\}$ denote the allocated slots of buyer a and b respectively, where $j_1^a < j_2^a < \cdots < j_{u_a}^a$ and $j_1^b < j_2^b < \cdots < j_{u_b}^b$. As $k \in W_i$, then $q_{j_1^i} \geq q_{j_{u_a}^a}$ and $q_{j_{u_i}^i} \geq q_{j_1^b}$ (note that $j_{u_a}^a$ and j_1^b are the indices of slots with the largest qualities in W_a and W_b respectively). We will show that either swapping winning slots of i with a or with b will increase the virtual surplus. To prove this, there four cases needed to be considered: (1). $u_i \geq u_a$ and $u_i \geq u_b$; (2). $u_i \geq u_a$ and $u_i < u_b$; (3). $u_i < u_a$ and $u_i \geq u_b$; (4). $u_i < u_a$ and $u_i < u_b$. We only prove the case (1) since the other cases can be proved similarly. Now, suppose $u_i \geq u_a$ and $u_i \geq u_b$, then we must have either (i). $\sum_{k=1}^{u_b} q_{j_k^i} \geq \sum_{k=1}^{u_b} q_{j_k^b}$ or (ii). $\sum_{k=1}^{u_a} q_{j_{u_i-k+1}^i} \geq \sum_{k=1}^{u_a} q_{j_k^a}$. Suppose (i) is not true, that is $\sum_{k=1}^{u_b} q_{j_k^i} < \sum_{k=1}^{u_b} q_{j_k^b}$, if $u_b \leq e$, then we have $q_{j_1^i} \leq q_{j_{u_b}^i}$, as a result

$$u_b q_{j_1^i} \leq \sum_{k=1}^{u_b} q_{j_k^i} < \sum_{k=1}^{u_b} q_{j_k^b} \leq u_b q_{j_1^b} \leq u_b q_{j_{u_i}^i},$$

thus, $q_{j_1^i} < q_{j_{u_i}^i}$; otherwise $u_b > e$, then it must also hold that $q_{j_1^i} \leq q_{j_{u_i}^i}$ (otherwise, for any $1 \leq \ell \leq u_b$, $q_{j_\ell^i} \geq q_{j_{u_i}^i} \geq q_{j_1^b}$ implying that $\sum_{k=1}^{u_b} q_{j_k^i} \geq u_b q_{j_1^b} \geq \sum_{k=1}^{u_b} q_{j_k^b}$, a contradiction). In both cases, it is obtained that $q_{j_1^i} \leq q_{j_{u_i}^i}$, therefore,

$$\sum_{k=1}^{u_a} q_{j_{u_i-k+1}^i} \geq u_a q_{j_1^i} \geq \sum_{k=1}^{u_a} q_{j_k^a}$$

implying (ii) is true. Thus, if (i) is true, by simple calculations, swapping winning slots of i with b will increase the virtual value (since $\phi_b(v_b) \geq \phi_i(v_i)$), otherwise swapping winning slots of i with a will increase the virtual surplus (since $\phi_a(v_a) \geq \phi_i(v_i)$). Then keep doing it by the method of Case 1 until eliminating all slots not allocated to $[s]$ in the middle of allocation to $[s]$ and attaining the desired optimal solution. □

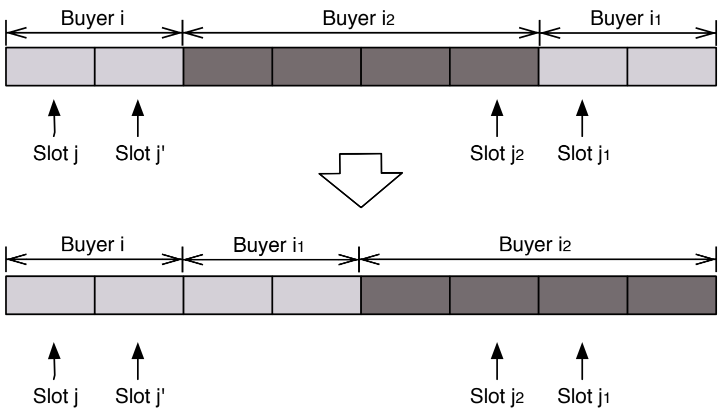

Fig. 2 Slots with light color are assigned to $[s]$. By swapping the positions of i_1 and i_2, we make the allocations to $[s]$ consecutive.

Since there exists an optimal solution that assigns slots to $[s]$ consecutively (Lemma 5), we can express the slots allocated to $[s]$ as an interval $[\ell, r]$. Let $g[s, \ell, r]$ denote the maximized value of our objective function $\sum_i \phi_i(v_i) t_i(\mathbf{v})$ when we only consider the first s buyers and the allocation of $[s]$ is exactly the interval $[\ell, r]$. Then we have the following recurrence.

$$g[s, \ell, r] = \max \begin{cases} g[s-1, \ell, r] \\ g[s-1, \ell, r-d_s] + \phi_s(v_s) \sum_{j=r-d_s+1}^{r} q_j \\ g[s-1, \ell+d_s, r] + \phi_s(v_s) \sum_{j=\ell}^{\ell+d_s-1} q_j \end{cases} \tag{3}$$

Our summary statement is as follows.

Theorem 1 *The mechanism that applies the allocation rule according to Dynamic Programming (3) and payment rule according to Equation (2) is an optimal mechanism for the banner advertisement problem with single peak qualities.*

Proof To complete the proof, it suffices to prove that $T_i(v_i)$ is monotone non-decreasing. More specifically, we prove a stronger fact, that $t_i(v_i, v_{-i})$ is non-decreasing as v_i increases. Given other buyers' bids v_{-i}, the monotonicity of t_i is equivalent to $t_i(v_i, v_{-i}) \le t_i(v_i', v_{-i})$ if $v_i' > v_i$. Assuming that $v_i' > v_i$, the regularity of ϕ_i implies that $\phi_i(v_i) \le \phi_i(v_i')$. If $\phi_i(v_i) = \phi_i(v_i')$, then $t_i(v_i, v_{-i}) = t_i(v_i', v_{-i})$ and we are done.

Consider the case that $\phi_i(v_i) < \phi_i(v_i')$. Let Q and Q' denote the total quantities obtained by all the other buyers except buyer i in the mechanism when buyer i bids v_i and v_i' respectively.

$$\phi_i(v_i') t_i(v_i', v_{-i}) + Q' \ge \phi_i(v_i') t_i(v_i, v_{-i}) + Q$$
$$\phi_i(v_i) t_i(v_i, v_{-i}) + Q \ge \phi_i(v_i) t_i(v_i', v_{-i}) + Q'.$$

The above inequalities are due to the optimality of allocations when i bids v_i and v_i' respectively. It follows that

$$\phi_i(v_i')(t_i(v_i, v_{-i}) - t_i(v_i', v_{-i})) \le Q' - Q$$
$$\phi_i(v_i)(t_i(v_i, v_{-i}) - t_i(v_i', v_{-i})) \ge Q' - Q$$

By the fact that $\phi_i(v_i) < \phi_i(v_i')$, it must be $t_i(v_i, v_{-i}) \le t_i(v_i', v_{-i})$. □

4 Multiple Peaks Case

Suppose now that there are only h peaks (local maxima) in the qualities. Thus, there are at most $h-1$ valleys (local minima). Since h is a constant, we can enumerate all the buyers occupying the valleys. After this enumeration, we can divide the sequence of slots into at most h consecutive pieces, each of which is single-peaked. Theorem 2 shows, using similar properties as those in Lemma 4 and 5, how we can obtain a larger size (polynomial-time) dynamic program similar to dynamic program (3) to solve the problem.

Theorem 2 *There is a polynomial algorithm to compute revenue maximization problem in Bayesian settings where the qualities of slots have a constant number of peaks.*

Proof Our proof is based on the single peak algorithm. Assume there are h peaks, thus there must be $h-1$ valleys. Suppose these valleys are indexed $j_1, j_2, \cdots, j_{h-1}$. In an optimal allocation, for any j_k, $k = 1, 2, \cdots, h-1$, j_k must be allocated to a buyer or unassigned to any buyer. If j_k is assigned to a buyer, say, buyer i, since i would buy d_i consecutive slots, j_k may appear in the ℓth position of these d_i consecutive slots. Hence, by this brute force, each j_k will at most have $\sum_i d_i + 1 \leq mn + 1$ possible positions to be allocated. In all, all the valleys have $(mn+1)^h$ possible allocated positions. For each of this allocation, the slots is broken into h single peak slots. We can obtain similar properties to those in Lemma 4 and 5. Without loss of generality, suppose the rest of the buyers are still the set $[n]$, with non-increasing virtual value. Since the optimal solution always assigns to $[s]$ consecutively, we can express the allocations to $[s]$ as intervals denoted by $[\ell_i, r_i]$, $i = 1, 2, \cdots, h$, where $[\ell_i, r_i]$ lies in the i-th single peak slot. Let $g[s, \ell_1, r_1, \cdots, \ell_h, r_h]$ denote the maximized value of our objective function $\sum_i \phi_i(v_i) t_i(\mathbf{v})$ when we only consider the first s buyers and the allocations of $[s]$ are exactly intervals $[\ell_i, r_i]$, $i = 1, 2, \cdots, h$. Then we have the following dynamic program: $g[s, \ell_1, r_1, \cdots, \ell_h, r_h]$ can be evaluated as

$$\max_{i \in [d]} \begin{cases} g[s-1, \ell_1, r_1, \cdots, \ell_h, r_h] \\ g[s-1, \ell_1, r_1, \cdots, \ell_i, r_i - d_s, \cdots, \ell_h, r_h] + \phi_s(v_s) \sum_{j=r_i-d_s+1}^{r_i} q_j \\ g[s-1, \ell_1, r_1, \cdots, \ell_i + d_s, r_i, \cdots, \ell_h, r_h] + \phi_s(v_s) \sum_{j=\ell_i}^{\ell_i+d_s-1} q_j \end{cases}$$

Note that the dynamic programming runs in polynomial time provided that the number of peaks is constant. ☐

Now we consider the case without the constant peak assumption and prove the following hardness result.

Theorem 3 *The revenue maximization problem for allocating consecutive ad slots (with unrestricted qualities) to buyers with fixed demands, is NP-hard.*

Proof We prove the NP-hardness by reducing from the 3-PARTITION problem, which is to decide whether a given multi-set of integers can be partitioned into triplets that all have the same sum. More precisely, given a multi-set S of $3n$ positive integers, can S be partitioned into n triplets $S_1, \ldots, S_n$ such that the sum of the numbers in each subset is equal? We use the fact that 3-PARTITION remains NP-complete in a strong sense in [20], meaning that it remains NP-complete even when the integers in S are bounded above by a polynomial in n.

Given an instance of 3-PARTITION $(a_1, a_2, \ldots, a_{3n})$, we construct an instance of the advertising problem with $3n$ advertisers and $m = n + \sum_i a_i$ slots. Note that m is polynomial in n due to the fact that all a_i are bounded by a polynomial in n. In the advertising instance, the valuation v_i for each advertiser i is 1 and his demand d_i is defined as a_i. Moreover, for any advertiser, his valuation distribution is that $v_i = 1$ with probability 1. Then everyone's virtual value is exactly 1. Maximizing revenue is equivalent to maximizing the simplified function $\sum_i \sum_{j \in X_i} q_j$.

Let $B = \sum_i a_i / n$. We define the quality of slot j to be 0 if j is a multiple of $B + 1$, otherwise $q_j = 1$. That can be illustrated as follows.

$$\underbrace{1\,1 \cdots 1}_{B}\; 0\; \underbrace{1\,1 \cdots 1}_{B}\; 0 \ldots \underbrace{1\,1 \cdots 1}_{B}\; 0$$

It is not hard to see that the optimal revenue is $\sum_i a_i$ iff there is a solution to this 3-PARTITION instance. ☐

5 Competitive Equilibrium

In this section, we study the revenue maximizing competitive equilibrium in the full information setting. To simplify the following discussions, we sort all buyers in non-increasing order of their values, i.e., $v_1 \geq v_2 \geq \cdots \geq v_n$. (By contrast, in Section 3, we sorted them in non-increasing order of virtual values, which are not relevant to competitive equilibrium.) The revenues obtained are compared experimentally in Section 7.

We say that an allocation $\mathbf{Y} = (Y_1, Y_2, \cdots, Y_n)$ is *efficient* if $\mathbf{Y}$ maximizes the total social welfare, that is, $\sum_i \sum_{j \in Y_i} v_{ij}$ is maximized over all the possible allocations. We call $\mathbf{p} = (p_1, p_2, \cdots, p_m)$ an *equilibrium price* if there exists an allocation $\mathbf{X}$ such that $(\mathbf{X}, \mathbf{p})$ is a competitive equilibrium. The following lemma is implicitly stated in [23]; for completeness, we give a proof below.

Lemma 6 *If allocation $\mathbf{Y}$ is efficient, then for any equilibrium price $\mathbf{p}$, $(\mathbf{Y}, \mathbf{p})$ is a competitive equilibrium.*

Proof Since $\mathbf{p}$ is an equilibrium price, there exists an allocation $\mathbf{X}$ such that $(\mathbf{X}, \mathbf{p})$ is a competitive equilibrium. As a result, by envy-freeness, $u_i(\mathbf{X}, \mathbf{p}) \geq u_i(\mathbf{Y}, \mathbf{p})$ for any $i \in [n]$. Let $T = [m] \setminus \cup_i Y_i$, then we have

$$\begin{aligned}
&\sum_i \sum_{j \in Y_i} v_{ij} - \sum_{j=1}^{m} p_j \geq \sum_i \sum_{j \in X_i} v_{ij} - \sum_{j=1}^{m} p_j = \sum_i \sum_{j \in X_i} v_{ij} - \sum_i \sum_{j \in X_i} p_j \\
&= \sum_i u_i(\mathbf{X}, \mathbf{p}) \geq \sum_i u_i(\mathbf{Y}, \mathbf{p}) = \sum_i \sum_{j \in Y_i} v_{ij} - \sum_i \sum_{j \in Y_i} p_j \\
&= \sum_i \sum_{j \in Y_i} v_{ij} - \sum_{j=1}^{m} p_j + \sum_{j \in T} p_j
\end{aligned} \tag{4}$$

where the first inequality is due to $\mathbf{Y}$ being efficient and first equality due to $u_i(\mathbf{X}, \mathbf{p})$ being competitive equilibrium (unallocated item priced at 0). Therefore, $\sum_{j \in T} p_j = 0$ and the above inequalities are all equalities. $\forall i : u_i(\mathbf{X}, \mathbf{p}) = u_i(\mathbf{Y}, \mathbf{p})$. Further, because the price is the same, we have

- For every loser i and every set Z of consecutive items with $|Z| = d_i$, we have $u_i(Z) \leq 0$.
- For every winner i and every set Z of consecutive items with $|Z| = d_i$, we have

$$u_i(Y_i) = u_i(X_i) \geq u_i(Z).$$

Therefore, $(\mathbf{Y}, \mathbf{p})$ is a competitive equilibrium. □

By Lemma 6, to find a revenue maximizing competitive equilibrium, we can first find an efficient allocation and then use linear programming to settle the prices. We develop the following dynamic programming to find an efficient allocation. We first only consider there is one peak in the quality order of items. The case with constant peaks is similar to the above approaches, for general peak case, as shown in above Theorem 3, finding one competitive equilibrium is NP-hard if the competitive equilibrium exists, and determining existence of competitive equilibrium is also NP-hard. This is because that considering the instance in the proof of Theorem 3, it is not difficult to see the constructed instance has an equilibrium if and only if 3 partition has a solution.

Recall that all the values are sorted in non-increasing order, $v_1 \geq v_2 \geq \cdots \geq v_n$. $g[s,\ell,r]$ denotes the maximized value of social welfare when we only consider first s buyers and the allocation of s is exactly the interval $[\ell, r]$. Then we have the following recurrence.

$$g[s,l,r] = \max \begin{cases} g[s-1,\ell,r] \\ g[s-1,\ell,r-d_s] + v_s \sum_{j=r-d_s+1}^{r} q_j \\ g[s-1,\ell+d_s,r] + v_s \sum_{j=\ell}^{\ell+d_s-1} q_j \end{cases} \tag{5}$$

By tracking procedure 5, an efficient allocation denoted by $\mathbf{X}^* = (X_1^*, X_2^*, \cdots, X_n^*)$ can be found. The price $\mathbf{p}^*$ such that $(\mathbf{X}^*, \mathbf{p}^*)$ is a revenue maximization competitive equilibrium can be determined from the following linear programming. Let T_i be any consecutive number of d_i slots, for all $i \in [n]$.

$$\begin{aligned} \max \quad & \sum_{i\in[n]} \sum_{j\in X_i^*} p_j \\ s.t. \quad & p_j \geq 0 && \forall\, j \in [m] \\ & p_j = 0 && \forall\, j \notin \cup_{i\in[n]} X_i^* \\ & \sum_{j\in X_i^*} (v_i q_j - p_j) \geq \sum_{j'\in T_i} (v_i q_{j'} - p_{j'}) && \forall\, i \in [n] \\ & \sum_{j\in X_i^*} (v_i q_j - p_j) \geq 0 && \forall i \in [n] \end{aligned}$$

Clearly there is only a polynomial number of constraints. The constraints in the first line represent that all the prices are non negative (no positive transfers). The constraint in the second line means unallocated items must be priced at zero (market clearance condition). And the constraint in the third line contains two aspects of information. First, for any loser k with $X_k = \emptyset$, the utility that k gets from any consecutive number of d_k is no more than zero, which makes all the losers envy-free. The second aspect is that any winner i with $X_i \neq \emptyset$ must receive a bundle with d_i consecutive slots maximizing its utility over all d_i consecutive slots, which together with the constraint in the fourth line (winner's utilities are non negative) guarantees that all winners are envy-free.

Theorem 4 *Under the condition of a constant number of peaks in the qualities of slots, there is a polynomial time algorithm to decide whether there exists a competitive equilibrium or not and to compute a revenue maximizing revenue market equilibrium if one does exist. If the number of peaks in the qualities of the slots is unbounded, both the problems are NP-complete.*

Proof Clearly the above linear programming and procedure (5) run in polynomial time. If the linear programming output a price $\mathbf{p}^*$, then by its constraint conditions, $(\mathbf{X}^*, \mathbf{p}^*)$ must be a competitive equilibrium. On the other hand, if there exist a competitive equilibrium $(\mathbf{X}, \mathbf{p})$ then by Lemma 6, $(\mathbf{X}^*, \mathbf{p})$ is a competitive equilibrium, providing a feasible solution of above linear programming. By the objective of the linear programming, it must be a revenue maximizing one.

The NP-hardness follows from Theorem 3. □

6 Consecutive Envy-freeness

We first prove a negative result on computing the revenue maximization problem in general demand case. We show it is NP-hard even if all the qualities are the same.

Theorem 5 *The revenue maximization problem for allocating consecutive ad slots to envy-free buyers is NP-hard, even if all slot qualities are the same.*

Proof Using similar ideas in the proof of Theorem 3, we prove the NP-hardness by reducing the 3-PARTITION problem. Given an instance of 3-PARTITION $(a_1, a_2, \ldots, a_{3n})$. Let $B = \sum_i a_i/n$. We construct an instance for advertising problem with $3n+1$ advertisers and $m = B + 1 + n + \sum_i a_i$ slots. It should be mentioned that m is polynomial of n due to the fact that all a_i are bounded by a polynomial of n. In the advertising instance, the valuation v_i for each advertiser i is 1 and his demand d_i is defined as a_i and there is another buyer with valuation 2 for each slot and with demand $B+1$. The quality of each slot j is 1. It is not hard to see that the optimal revenue is $nB + 2(B+1)$ if and only if there is a solution to this 3-PARTITION instance, the optimal solution is illustrated as follows.

$$\underbrace{1\,1\cdots1}_{B+1}\ \underbrace{1}_{\text{unassigned}}\ \underbrace{1\,1\cdots1}_{B}\ \underbrace{1}_{\text{unassigned}}\ \underbrace{1\,1\cdots1}_{B}\ \underbrace{1}_{\text{unassigned}}\ \cdots\underbrace{1\,1\cdots1}_{B}$$

□

Although the hardness in Theorem 5 indicates that finding the optimal revenue for general demand in polynomial time is impossible, it does not however rule out the important special case where the demand is uniform, that is, $d_i = d$. If in addition we have slots that are arranged are in non-increasing order of their qualities, that is, $q_1 \geq q_2 \geq \cdots \geq q_m$, then we have the following computational positive result.

Theorem 6 *There is a polynomial-time algorithm to compute the consecutive envy-free solution in the case where all the buyers have the same demand and slots are arranged in non-increasing order of quality.*

The proof of Theorem 6 is based on bundle envy-free solutions; in fact we will prove the bundle envy-free solution is also a consecutive envy-free solution by defining price of items properly. Thus, we first need to give the result on bundle envy-free solutions. Suppose d is the uniform demand for all the buyers. Let T_i be the slot set allocated to buyer i, $i = 1, 2, \cdots, n$. Let P_i be the total payment of buyer i and p_j be the price of slot j. Let t_i denote the total qualities obtained by buyer i, $t_i = \sum_{j \in T_i} q_j$, and $\alpha_i = iv_i - (i-1)v_{i-1}$, $\forall i \in [n]$.

Theorem 7 *The revenue maximization problem of bundle envy-freeness is equivalent to solving the following LP.*

$$\begin{aligned} \textit{Maximize:} \quad & \sum_{i=1}^{n} \alpha_i t_i \\ \textit{s.t.} \quad & t_1 \geq t_2 \geq \cdots \geq t_n \\ & T_i \subset [m], \quad T_i \cap T_k = \emptyset \quad \forall i, k \in [n] \end{aligned} \tag{6}$$

Proof Recall P_i denotes the payment of buyer i, and we next prove that the linear program (6) actually gives an optimal solution of bundle envy-free pricing. By the definition of bundle envy-free, where buyer i would not envy buyer $i+1$ and vice versa, we have

$$v_i t_i - P_i \geq v_i t_{i+1} - P_{i+1} \tag{7}$$

$$v_{i+1} t_{i+1} - P_{i+1} \geq v_{i+1} t_i - P_i \tag{8}$$

By summing up the above two inequalities, we have $(v_i - v_{i+1})(t_i - t_{i+1}) \geq 0$. Hence, if $v_i > v_{i+1}$, then $t_i \geq t_{i+1}$. From (7), we get $P_i \leq v_i(t_i - t_{i+1}) + P_{i+1}$. The maximum payment of buyer i is

$$P_i = v_i(t_i - t_{i+1}) + P_{i+1}, \tag{9}$$

Together with $t_i \geq t_{i+1}$, the above equation implies (7) and (8). Furthermore, the maximum payment of n is $P_n = t_n v_n$. (9) together with $t_i \geq t_{i+1}$ and $P_n = t_n v_n$ would make everyone bundle envy-free, the arguments are as follows.

All the buyers must be bundle envy free. By (9), we have $P_i - P_{i+1} = v_i(t_i - t_{i+1})$, hence $P_i = \sum_{k=i}^{n-1} v_k(t_k - t_{k+1}) + P_n$. Noticing that if $t_i = 0$, then $P_i = 0$, which means i is a loser. For any buyer $j < i$, we have

$$P_j - P_i = \sum_{k=j}^{i-1} v_k(t_k - t_{k+1}) \leq \sum_{k=j}^{i-1} v_j(t_k - t_{k+1}) = v_j(t_j - t_i)$$

By rearranging the terms, we have $v_j t_i - P_i \leq v_j t_j - P_j$, that means buyer j would not envy buyer i. Similarly,

$$P_j - P_i = \sum_{k=j}^{i-1} v_k(t_k - t_{k+1}) \geq \sum_{k=j}^{i-1} v_i(t_k - t_{k+1}) = v_i(t_j - t_i)$$

We have $v_i t_i - P_i \geq v_i t_j - P_j$ that implies i would not envy buyer j.

Now we are ready to calculate $\sum_{i=1}^{n} P_i$ based on (9) and $t_{n+1} = 0$.

$$\begin{aligned}
\sum_{i=1}^{n} P_i &= \sum_{i=1}^{n} \left[\sum_{k=i}^{n-1} v_k(t_k - t_{k+1}) + P_n \right] = \sum_{i=1}^{n} \sum_{k=i}^{n} v_k(t_k - t_{k+1}) \\
&= \sum_{k=1}^{n} \sum_{i=1}^{k} v_k(t_k - t_{k+1}) = \sum_{k=1}^{n} k v_k(t_k - t_{k+1}) \\
&= \sum_{k=1}^{n} k v_k t_k - \sum_{k=1}^{n} (k-1) v_{k-1} t_k = \sum_{i=1}^{n} \alpha_i t_i
\end{aligned}$$

We know the revenue maximizing problem of bundle envy-freeness can be formalized as (6). □

Since consecutive envy-free solutions are a subset of (sharp) bundle envy-free solutions, hence the optimal value of optimization (6) gives an upper bound of optimal objective value of consecutive envy-free solutions. Note that the optimization LP (6) can be solved by dynamic programming. Let $g[s, j]$ denote the optimal objective value of the following LP with some set in $[j]$ allocated to all the buyers in $[s]$:

$$\begin{aligned} \text{Maximize:} \quad & \sum_{i=1}^{s} \alpha_i t_i \\ \text{s.t.} \quad & t_1 \geq t_2 \geq \cdots \geq t_s \\ & T_i \subset [j], \quad T_i \cap T_k = \emptyset \quad \forall i, k \in [s] \end{aligned}$$

Then

$$g[s,j] = \max \begin{cases} g[s, j-1] \\ g[s-1, j-d] + \alpha_s \sum_{u=j-d+1}^{j} q_u \end{cases}$$

Next, we show how to modify the (sharp) bundle envy-free solution to consecutive envy-free solutions by properly defining the slot price of T_i, for all $i \in [n]$. Suppose the optimal winner set of bundle envy-free solution is $[L]$. Assume the optimal allocation and price of bundle envy-free solution are $T_i = \{j_1^i, j_2^i, \cdots, j_d^i\}$ with $j_1^i \geq j_2^i \geq \cdots \geq j_d^i$ and P_i respectively, for all $i \in [L]$.

Proof of Theorem 6 Define the price of T_i iteratively: $p_{j_k^L} = v_L q_{j_k^L}$, for all $k \in [d]$; $p_{j_k^i} = v_i(q_{j_k^i} - q_{j_k^{i+1}}) + p_{j_k^{i+1}}$ for $k \in [d]$ and $i \in [n]$. Now we observe that the price defined by the above procedure is still a bundle envy-free solution. This is because by induction, we have $P_i = \sum_{k=1}^{d} p_{j_k^i}$. Hence, we need only to check the prices defined as above and allocations T_i constitute a consecutive envy-free solution. In fact, we prove a strong version: suppose T_i's are consecutive from top to bottom in a line S, we will show each buyer i would not envy any consecutive sub line of S comprising d slots. For any i, we consider two cases.

Case 1. Buyer i would not envy the slots below his slots.
for any consecutive line T below i with size d, suppose T comprises of slots won by buyer k (denoted such slot set by U_k) and $k+1$ (denoted such slot set by U_{k+1} and let $\ell = |U_{k+1}|$) where $k \geq i$. Recall that $t_i = \sum_{j \in T_i} q_j$, then

$$\begin{aligned} & \sum_{j \in T_i} p_j - \sum_{j \in T} p_j = v_i(t_i - t_{i+1}) + P_{i+1} - \sum_{j \in U_k \cup U_{k+1}} p_j \\ & = v_i(t_i - t_{i+1}) + v_{i+1}(t_{i+1} - t_{i+2}) + \cdots + P_k - \sum_{j \in U_k \cup U_{k+1}} p_j \\ & = v_i(t_i - t_{i+1}) + v_{i+1}(t_{i+1} - t_{i+2}) + \cdots + \sum_{j \in T_k \setminus U_k} p_j - \sum_{j \in U_{k+1}} p_j \\ & = v_i(t_i - t_{i+1}) + v_{i+1}(t_{i+1} - t_{i+2}) + \cdots + \sum_{u=1}^{\ell} v_k(q_{j_u^k} - q_{j_u^{k+1}}) \\ & \leq v_i(t_i - t_{i+1}) + v_i(t_{i+1} - t_{i+2}) + \cdots + \sum_{u=1}^{\ell} v_i(q_{j_u^k} - q_{j_u^{k+1}}) \\ & = v_i t_i - v_i \sum_{j \in T} q_j. \end{aligned}$$

By rewriting $\sum_{j \in T_i} p_j - \sum_{j \in T} p_j \leq v_i t_i - v_i \sum_{j \in T} q_j$ as $v_i t_i - \sum_{j \in T_i} p_j \geq v_i \sum_{j \in T} q_j - \sum_{j \in T} p_j$, we get the desired result.

Case 2. Buyer i would not envy the slots above his slots. For any consecutive line T above i with size d, suppose T comprises of slots won by buyer k (denoted such slot set by U_k) and $k-1$ (denoted such slot set by U_{k-1} and let $\ell = |U_{k-1}|$) where $k \leq i$. Recall that $t_i = \sum_{j \in T_i} q_j$, then

$$\begin{aligned}
&\sum_{j \in T} p_j - \sum_{j \in T_i} p_j = \sum_{j \in U_{k-1} \cup U_k} p_j - \sum_{j \in T_i} p_j \\
&= \sum_{u=d-\ell+1}^{d} v_{k-1}(q_{j_u^{k-1}} - q_{j_u^k}) + \sum_{j \in T_k} p_j - \sum_{j \in T_i} p_j \\
&= \sum_{u=d-\ell+1}^{d} v_{k-1}(q_{j_u^{k-1}} - q_{j_u^k}) + v_k(t_k - t_{k+1}) + \cdots + v_{i-1}(t_{i-1} - t_i) \\
&\geq \sum_{u=d-\ell+1}^{d} v_i(q_{j_u^{k-1}} - q_{j_u^k}) + v_i(t_k - t_{k+1}) + \cdots + v_i(t_{i-1} - t_i) \\
&= v_i \sum_{j \in T} q_j - v_i t_i.
\end{aligned}$$

By rewriting $\sum_{j \in T} p_j - \sum_{j \in T_i} p_j \geq v_i \sum_{j \in T} q_j - v_i t_i$ as $v_i t_i - \sum_{j \in T_i} p_j \geq v_i \sum_{j \in T} q_j - \sum_{j \in T} p_j$ we get the desired result. □

7 Simulation

Since the consecutive model has a direct application for rich media advertisement, the simulation for comparing the schemes: Bayesian optimal mechanism (Bayesian for simplicity in this section), consecutive CE (CE for simplicity in this section), consecutive EF (EF for simplicity in this section), generalized GSP, will be presented in this section. We did a simulation to compare the expected revenue among those pricing schemes. The sampling method is applied to the competitive equilibrium, envy-free solution, Bayesian truthful mechanism, as well as generalized GSP, which is the widely used pricing scheme for text ads in most advertisement platforms nowadays.

The value samples v come from the same uniform distribution $U[20, 80]$. With a random number generator, we produced 200 group samples $\{V_1, V_2, \cdots, V_{200}\}$, that are used as the input of our simulation. Each group contains n samples, $V_k = \{v_k^1, v_k^2, \cdots, v_k^n\}$, where each v_k^i is sampled from the uniform distribution $U[20, 80]$. For the parameters of slots, we assume there are 6 slots to be sold, and fix their position qualities:

$$Q = \{q_1,\ \ q_2,\ \ q_3,\ \ q_4,\ \ q_5,\ \ q_6\} = \{0.8,\ 0.7,\ 0.6,\ 0.5,\ 0.4,\ 0.3\}$$

The actual ads auction is complicated, and simplified in our simulation. We do not consider richer conditions, such as set all bidders' budgets unlimited, and there are no reserve prices in the mechanisms. We vary the group size n from 5 to 12, and observe their expected revenue variation. From $j = 1$ through 200, at each j, invoke the function EF (V_j, D, Q), GSP (V_j, D, Q), CE (V_j, D, Q) and Bayesian (V_j, D, Q) respectively. Thus, those mechanisms use the same data from the same distribution as inputs and compare their expected revenue fairly. Finally, we average those results from different mechanisms respectively, and compare their expected revenue at sample size n.

The generalized GSP mechanism for rich ads in the simulation was not introduced in the previous sections. Here, in our simulation, it is a simple generalization of the standard GSP which is used in keywords auction. In our generalization of GSP, the allocations of bidders are given by maximizing the total social welfare, which is compatible with GSP in keywords auction, and each winner's price per quality is given by the next highest bidder's bid per quality. Since the real generalization of GSP for rich ads is unknown and the generalization form may be various, our generalization of GSP for rich ads may not be a revenue maximizing one, however, it is a natural one.

Incentive analysis is also considered in our simulation, except Bayesian mechanism (it is truthful bidding, $b_i = v_i$). Since the bidding strategies in other mechanisms (GSP, CE, EF) are unclear, we present a simple bidding strategy for bidders to converge to an equilibrium. We try to find the equilibrium bids by searching the bidder's possible bids ($b_i < v_i$) one by one, from top rank bidders to lower rank bidders iteratively, until reaching an equilibrium where no one would like to change his bid. If any equilibrium exists, we count the expected revenue for this sample; if not, we ignore this sample.

Since the Envy-Free solution in our paper only works for the condition that all the bidders have the same demand, thus, we did the simulation in 2 separate ways:

1. Simulation I is for bidders with fixed demands, $d_i = 2$ for all i, and compares expected revenues obtained by GSP, CE, EF, Bayesian.
2. Simulation II is for bidders with different demands and compares expected revenues obtained by GSP, CE, Bayesian.

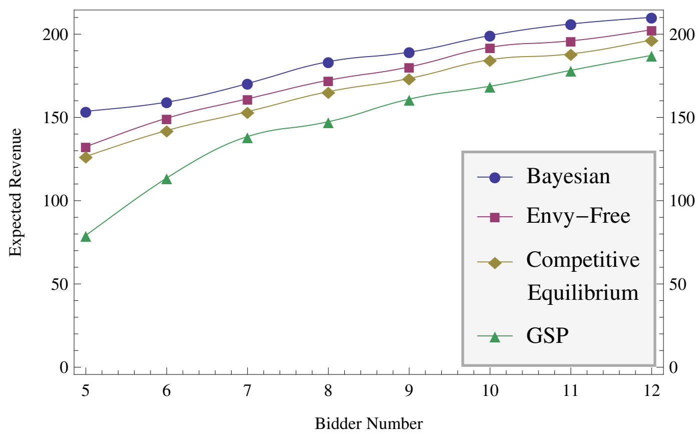

Fig. 3 Simulation results from different mechanisms, all bidders' demand fixed at $d_i = 2$

Figure 3 shows I's results when all bidders' demand fixed at 2. Obviously, the expected revenue is increasing when more bidders involved. When the bidders' number rises, the rank of expected revenue of different mechanisms remains the same in the order Bayesian > EF > CE > GSP.

Simulation II is for bidders with various demands. With loss of generality, we assume that bidder's demand $D = \{d_1, d_2, \cdots, d_i\}, d_i \in \{1, 2, 3\}$, we assign those bidders' demand randomly, with equal probability.

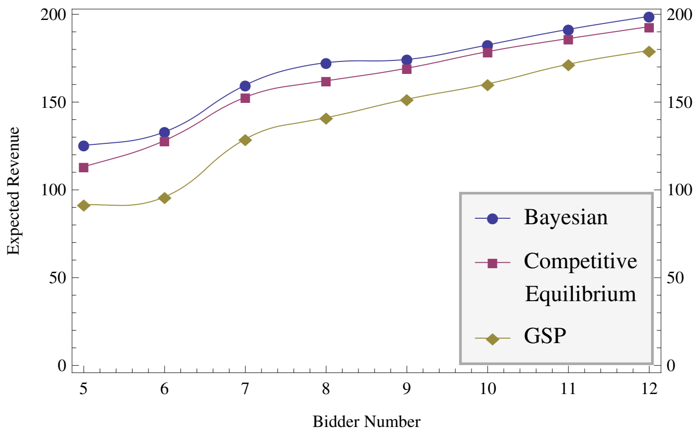

Fig. 4 Simulation results from different mechanisms, bidders' demand varies in {1,2,3}

Figure 4 shows our simulation results for II when bidders' demand varies in {1,2,3}, the rank of expected revenue of different mechanisms remains the same as simulation I, From this chart, we can see that Bayesian truthful mechanism and competitive equilibrium get more revenues than generalized GSP.

8 Conclusion

The rich media pricing models for consecutive demand buyers in the context of Bayesian truthfulness, competitive equilibrium and envy-free solution paradigm are investigated in this paper. As a result, an optimal Bayesian incentive compatible mechanism is proposed for various settings such as single peak and multiple peaks. In addition, to incorporate fairness e.g. envy-freeness, we also present a polynomial-time algorithm to decide whether or not there exists a competitive equilibrium and to compute a revenue maximized market equilibrium if one does exist. For envy-free settings, though the revenue maximization of general demand case is shown to be NP-hard, we still provide optimal solution of common demand case. Besides, our simulation shows a reasonable relationship of revenues among these schemes plus a generalized GSP for rich media ads.

Even though our main motivation arises from the rich media advert pricing problem, our models have other potential applications. For example TV ads can also be modeled under our consecutive demand adverts where inventories of a commercial break are usually divided into slots of fixed sizes, and slots have various qualities measuring their expected number of viewers and corresponding attractiveness. With an extra effort to explore the periodicity of TV ads, we can extend our multiple peak model to one involved with cyclic multiple peaks.

Besides single consecutive demand where each buyer only have one demand choice, the buyer may have more options to display his ads, for example select a large picture or a small one to display them. Our dynamic programming algorithm (3) can also be applied to this case (the transition function in each step selects maximum value from $2k+1$ possible values, where k is the number of choices of the buyer).

Another reasonable extension of our model would be to add budget constraints for buyers, i.e., each buyer cannot afford the payment that is more than his budget. By relaxing the requirement of Bayesian incentive compatible (BIC) to one of approximate BIC, this extension can be obtained by the recent milestone work of Cai et al. [9]. It remains an open problem how to do it under the exact BIC requirement. It would also be interesting to handle it under the market equilibrium paradigm for our model.

Acknowledgements Xiaotie Deng is supported by the National Science Foundation of China (Grant No. 61173011) and a Project 985 grant of Shanghai Jiaotong University. Paul W. Goldberg and Jinshan Zhang are supported by EPSRC project EP/K01000X/1.

References

1. Alaei, S.: Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. SIAM Journal on Computing **43**(2), 930–972 (2014)
2. Ashlagi, I., Braverman, M., Hassidim, A., Lavi, R., Tennenholtz, M.: Position auctions with budgets: Existence and uniqueness. The BE Journal of Theoretical Economics **10**(1) (2010)
3. Balseiro, S.R., Besbes, O., Weintraub, G.Y.: Auctions for online display advertising exchanges: Approximations and design. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC '13, pp. 53–54. ACM, New York, NY, USA (2013). DOI 10.1145/2482540.2482576. URL `http://doi.acm.org/10.1145/2482540.2482576`
4. Bhattacharya, S., Goel, G., Gollapudi, S., Munagala, K.: Budget-constrained auctions with heterogeneous items. Theory of Computing **8**, 429–460 (2012)
5. Borgs, C., Chayes, J., Immorlica, N., Jain, K., Etesami, O., Mahdian, M.: Dynamics of bid optimization in online advertisement auctions. In: Proceedings of the 16th international conference on World Wide Web, pp. 531–540. ACM (2007)
6. Boutilier, C., Parkes, D.C., Sandholm, T., Walsh, W.E.: Expressive banner ad auctions and model-based online optimization for clearing. In: Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 1, AAAI'08, pp. 30–37. AAAI Press (2008). URL `http://dl.acm.org/citation.cfm?id=1619995.1620002`
7. Briest, P.: Uniform budgets and the envy-free pricing problem. In: Automata, Languages and Programming, pp. 808–819. Springer (2008)
8. Cai, Y., Daskalakis, C., Weinberg, S.M.: An algorithmic characterization of multi-dimensional mechanisms. In: Proceedings of the forty-fourth annual ACM Symposium on Theory of Computing, pp. 459–478. ACM (2012)
9. Cai, Y., Daskalakis, C., Weinberg, S.M.: Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In: Proceedings of the fifty-third annual IEEE Symposium on Foundations of Computer Science, pp. 130–139. IEEE (2012)
10. Cantillon, E., Pesendorfer, M.: Combination bidding in multi-unit auctions. CEPR Discussion Papers **6083** (2007)
11. Caragiannis, I., Chatzigeorgiou, X., Kanellopoulos, P., Krimpas, G.A., Protopapas, N., Voudouris, A.A.: Efficiency and complexity of price competition among single-product vendors. In: IJCAI, pp. 25–31. AAAI Press (2015)
12. Chen, N., Deng, X.: Envy-free pricing in multi-item markets. In: Automata, Languages and Programming, vol. 6199, pp. 418–429. Springer (2010). DOI 10.1007/978-3-642-14162-1_35. URL `http://dx.doi.org/10.1007/978-3-642-14162-1_35`

13. Chen, N., Deng, X., Goldberg, P., Zhang, J.: On revenue maximization with sharp multi-unit demands. Journal of Combinatorial Optimization (2014). DOI 10.1007/s10878-014-9817-y. URL http://dx.doi.org/10.1007/s10878-014-9817-y
14. Demange, G., Gale, D., Sotomayor, M.: Multi-item auctions. The Journal of Political Economy pp. 863–872 (1986)
15. Edelman, B., Ostrovsky, M., Schwarz, M.: Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. The American Economic Review **97**(1), 242–259 (2007)
16. Engelbrecht-Wiggans, R., Kahn, C.M.: Multi-unit auctions with uniform prices. Economic theory **12**(2), 227–258 (1998)
17. Feldman, J., Muthukrishnan, S., Pal, M., Stein, C.: Budget optimization in search-based advertising auctions. In: Proceedings of the 8th ACM conference on Electronic commerce, pp. 40–49. ACM (2007)
18. Feldman, M., Fiat, A., Leonardi, S., Sankowski, P.: Revenue maximizing envy-free multi-unit auctions with budgets. In: Proceedings of the 13th ACM Conference on Electronic Commerce, pp. 532–549. ACM (2012)
19. Feldman, M., Gravin, N., Lucier, B.: Combinatorial auctions via posted prices. In: SODA '15. SIAM (2015)
20. Garey, M.R., Johnson, D.S.: Complexity results for multiprocessor scheduling under resource constraints. SIAM Journal on Computing **4**(4), 397–411 (1975)
21. Ghosh, A., McAfee, P., Papineni, K., Vassilvitskii, S.: Bidding for representative allocations for display advertising. In: Internet and Network Economics, pp. 208–219. Springer (2009)
22. Graepel, T., Candela, J.Q., Borchert, T., Herbrich, R.: Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft's bing search engine. In: Proceedings of the 27th International Conference on Machine Learning (ICML-10), pp. 13–20 (2010)
23. Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. Journal of Economic Theory **87**(1), 95–124 (1999)
24. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, pp. 1164–1173. Society for Industrial and Applied Mathematics (2005)
25. He, D., Chen, W., Wang, L., Liu, T.Y.: A game-theoretic machine learning approach for revenue maximization in sponsored search. In: Proceedings of the Twenty-Third international joint conference on Artificial Intelligence(IJCAI), pp. 206–212. AAAI Press (2013)
26. Korula, N., Mirrokni, V., Nazerzadeh, H.: Optimizing display advertising markets: Challenges and directions. IEEE Internet Computing (2016)
27. Kuminov, D., Tennenholtz, M.: User modeling in position auctions: Re-considering the gsp and vcg mechanisms. In: Autonomous Agents and Multiagent Systems, AAMAS '09, pp. 273–280. Richland, SC (2009). URL http://dl.acm.org/citation.cfm?id=1558013.1558050
28. Lahaie, S.: An analysis of alternative slot auction designs for sponsored search. In: Proceedings of the 7th ACM Conference on Electronic Commerce, pp. 218–227. ACM (2006)
29. Lahaie, S., Parkes, D.C., Pennock, D.M.: An expressive auction design for online display advertising. In: AAAI, pp. 108–113 (2008)
30. Lehmann, D., Oćallaghan, L.I., Shoham, Y.: Truth revelation in approximately efficient combinatorial auctions. J. ACM **49**(5), 577–602 (2002). DOI 10.1145/585265.585266. URL http://doi.acm.org/10.1145/585265.585266
31. Lucier, B., Paes Leme, R., Tardos, E.: On revenue in the generalized second price auction. In: Proceedings of the 21st International Conference on World Wide Web, WWW '12, pp. 361–370. ACM, New York, NY, USA (2012). DOI 10.1145/2187836.2187886. URL http://doi.acm.org/10.1145/2187836.2187886
32. Maillé, P., Markakis, E., Naldi, M., Stamoulis, G.D., Tuffin, B.: Sponsored search auctions: An overview of research with emphasis on game theoretic aspects. Electronic Commerce Research **12**(3), 265–300 (2012)
33. Meir, R., Lu, T., Tennenholtz, M., Boutilier, C.: On the value of using group discounts under price competition. Artificial Intelligence **216**, 163–178 (2014)
34. Myerson, R.B.: Optimal auction design. Mathematics of Operations Research **6**(1), 58–73 (1981)
35. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic game theory, vol. 1. Cambridge University Press Cambridge (2007)

36. Parkes, D.C., Sandholm, T.: Optimize-and-dispatch architecture for expressive ad auctions. 1st Workshop on Ad Auctions (2005)
37. Qin, T., Chen, W., Liu, T.Y.: Sponsored search auctions: Recent advances and future directions. ACM Trans. Intell. Syst. Technol. **5**(4), 60:1–60:34 (2015). DOI 10.1145/2668108. URL `http://doi.acm.org/10.1145/2668108`
38. Rothkopf, M.H., Pekeč, A., Harstad, R.M.: Computationally manageable combinational auctions. Management Science **44**(8), 1131–1147 (1998)
39. Shapley, L.S., Shubik, M.: The assignment game i: The core. International Journal of Game Theory **1**(1), 111–130 (1971)
40. Tang, B., Zhang, J.: Envy-free sponsored search auctions with budgets. In: Proceedings of the Twenty-Fourth international joint conference on Artificial Intelligence (IJCAI '15). AAAI Press (2015)
41. Varian, H.R.: Position auctions. International Journal of Industrial Organization **25**(6), 1163–1178 (2007)