# Truthfulness in Advertising? Approximation Mechanisms for Knapsack Bidders

Martin Bichler and Soeren Merting

*Department of Informatics, Technical University of Munich, Germany*
*Email: bichler@in.tum.de, soeren.merting@in.tum.de.*

**Abstract**

Quasilinear utility functions are a standard assumption in auction theory allowing for truthful and welfare-maximizing auction mechanisms. However, the literature on advertising markets suggests that the utility model of bidders rather resembles a knapsack problem, where advertisers try to maximize the sum of item values subject to a budget for a marketing campaign. Non-quasilinear environments rarely allow for truthful mechanisms. However, some characteristics of the market environment might allow for positive results. In particular, markets are large and bidders typically consider prices as exogenous. We introduce a model of advertising markets, and study whether truthful and prior-free approximation mechanisms with good approximation ratios of the maximal welfare are possible. We analyze the offline mechanism design problem and find a truthful and randomized mechanism with an approximation ratio of only 4. This mechanism draws on a fractional deferred acceptance algorithm and randomized rounding, and it illustrates how the relax-and-round principle can be implemented in this non-quasilinear environment. The article highlights the types of assumptions necessary for truthful mechanisms with good approximation ratios in an important class of non-quasilinear utility functions.

*Keywords:* Auctions/bidding, multi-agent systems, approximation mechanisms

*Corresponding author: Martin Bichler, Department of Informatics, Technical University of Munich, Boltzmannstr. 3, 85748 Garching, Germany

[1]The financial support from the Deutsche Forschungsgemeinschaft (DFG) (BI 1057/1-7) is gratefully acknowledged.

## 1. Introduction

In auction theory, bidders are typically modeled as payoff-maximizing individuals via a quasilinear utility function. If bidders have a quasilinear utility function the Vickrey-Clarke-Groves (VCG) mechanism is the unique mechanism to implement maximum welfare in dominant strategies (Green and Laffont, 1979). These positive results do not allow for private budget constraints (Dobzinski et al., 2012). However, there are markets where pure payoff-maximization without budget constraints might just not be the right assumption. For example, a large number of papers on digital advertising auctions suggest that automated bidders in such auctions rather maximize value subject to a budget constraint as in a knapsack problem (Feldman et al., 2008; Berg et al., 2010; Zhou et al., 2008; Lee et al., 2013; Chen et al., 2011; Zhang et al., 2014), which differs from payoff maximization assumed in a quasilinear utility function. The papers are written by academics and by authors working on display ad auctions in firms which are active in this field.

For example, in display ad auctions individual user impressions on a web site are auctioned off. Typically, the advertising firm provides an intermediary, a demand-side platform (DSP), with a budget for the overall campaign and a value or willingness-to-pay for individual impressions. This value might the profit from selling a product and a DSP is not allowed to bid beyond, because he would incur a loss even if the product got sold. The DSP provides autonomous agents bidding on behalf of the advertiser. The advertising firm (or client of the DSP) typically considers the overall budget $c_i$ as sunk cost devoted to a campaign. In other words, the task of the DSP is to invest the campaign budget such that the advertiser $i$'s sum of valuations $\sum_{j=1}^{m} v_{ij}x_{ij}$ for $m$ items is maximized subject to the total payments $\sum_{j=1}^{m} p_j x_{ij} \leq c_i$, where $x_{ij} \in \{0,1\}$ describes the assignment of items, $j \in J$, to bidders, $i \in I$. We refer to such bidders as *knapsack bidders* or bidders having a *knapsack utility function*. The literature above is based on this basic utility model. In this paper, we take the knapsack utility model from the literature serious and ask whether there are truthful mechanisms maximizing welfare if bidders have such utility functions.

Knapsack utility functions are not only adequate for online markets where supply or demand arrive dynamically over time, but also for conventional offline markets where supply and demand are all present at the time when the market is cleared. For example, bidders (advertisers) on a *TV advertising market* have similar characteristics, but ad slots in the program of a TV station are typically sold for certain periods of time such as every week, rather than dynamically whenever a new impression arrives (Goetzendorff et al., 2015; Nisan et al., 2009). Media agencies bidding in such auctions can also be modeled as knapsack bidders. Maximizing value subject to a budget constraint is actually wide-spread and also assumed in classic micro-economic demand theory (Mas-Colell et al., 1995, p. 50). So, our analysis might be relevant for other domains as well, but we motivate the model assumptions primarily from the literature on advertising markets.

Mechanism design theory typically imposes incentive-compatibility and individual rationality as constraints such that the resulting mechanism provides incentives to bid truthful and participants do not make a loss (Borgers et al., 2015). To implement a social choice function means to define a mechanism where truthful bidding is in equilibrium. Ideally, the equilibrium solution concept is prior-free, i.e., an agent does not need assumptions about the type distribution. Dominant strategy equilibria for deterministic mechanisms, and truthfulness-in-expectation (TIE) for randomized mechanisms are two such solution concepts used in this paper. Unfortunately, welfare maximization with quasilinear bidders appears to be an exceptional environment where the social choice function (i.e., welfare maximization) can be implemented in dominant strategies via the VCG mechanism. Also, randomized mechanisms for general preferences which are truthful-in-expectation typically draw on the VCG payment rule (Lavi and Swamy, 2011).

A few papers study mechanism design with non-quasilinear utility models. Kazumura and Serizawa (2015) shows that there is no multi-object auction mechanism for heterogenous goods that is dominant strategy incentive compatible and Pareto-efficient, even if only one bidder has multi-unit demands. Similarly, Baisa (2013) shows that if bidders have multi-dimensional types, there is no mechanism that satisfies (1) individual rationality, (2) dominant strategy incentive compatibility, (3) ex-post Pareto efficiency, and (4) weak budget balance for homogenous goods. The utility functions studied in this new line of mechanism design literature are general. For example, Baisa (2013) only assumes that a bidder's demand for the good increases as her wealth increases for a constant price level and some level of risk aversion.

Utility functions with more specific assumptions might allow for more positive results. Fadaei and Bichler (2017b) introduce a model of *value bidders* who maximize the value of packages of items for which they are given financial limits $v_i(S)$ reflecting their valuation that they must not overbid. A value bidders' utility function is $u_i(S) = v_i(S)$ if $p_i(S) \leq v_i(S)$, and $u_i(S) = -\infty$ otherwise, with $S \subseteq J$. The authors show that with a truthful and deterministic mechanism for *value bidders* only an $n$-approximation can be achieved, where $n$ is the number of bidders. For randomized mechanisms they describe a complex mechanism with a $O(\sqrt(m))$ approximation ratio. These results illustrate that without quasilinearity truthful mechanisms with good approximation ratios might often not be feasible.

*Knapsack bidders* are a specific type of value bidders where bidders only have additive package valuations up to an overall budget constraint they are given. They want to maximize the total value for the budget they invest. The main difference from value bidders to knapsack bidders is that the value for packages of objects is additive, and therefore the results from Fadaei and Bichler (2017b) extend. Let's provide a simple example to illustrate possibilities for manipulation (see Example 1).

**Example 1.** *Consider two items $A$ and $B$ and two bidders and an auctioneer trying to maximize social welfare, i.e., the sum of bidder valuations and the payments to the auctioneer. Bidder 1 has a value of \$10 for $A$ and \$9 for $B$ and a budget constraint of \$10. Bidder 2 has a value and a budget constraint of \$9 for $A$ only. Bidder 1 wants to maximize his value within budget and will reduce his bid on $B$ to zero. The auctioneer can sell item $A$ to bidder 1 and might leave $B$ unsold. Prices are bounded by the valuations and the overall budget. Independent of the price paid, the sum of bidder val-*

*uations is \$10 and therefore less than \$18, the total of the allocation where bidder 1 gets item B and bidder 2 gets item A.*

Apart from item-level valuations there are a few assumptions where advertising markets in the field differ from the value bidders described in Fadaei and Bichler (2017b): (1) bidders can only bid on objects individually, (2) markets are large with many items and many bidders, and (3) prices are determined by a second-price rule and considered as exogenous random variables by the bidders. Already Roberts and Postlewaite (1976) showed that in large markets the ability of an individual player to influence the market is minimal, so bidders should behave as price-taking agents. This is well reflected in the literature on bidding in display ad auctions (see for example Lee et al. (2013); Chen et al. (2011); Zhang et al. (2014)). Actually, prices for an impression are very volatile across the day and across different types of impressions and therefore very hard to predict in these markets (Ghosh et al., 2009; Cui et al., 2011). We will use the term *advertising model* to refer to large markets with many knapsack bidders and many objects, where bidders consider prices as exogenous and behave as price-takers.

Even in a model where bidders cannot influence prices, truthful mechanisms are difficult to construct as we will show. Bidders have a choice as to which objects they bid on and which objects they do not. This is actually a key strategic decision of DSPs in display ad auctions with millions of impressions per day. Targeting strategies have assumed center stage (Levin and Milgrom, 2010; McAfee, 2011), and they describe this strategic choice of impressions by a bidder (Bergemann and Bonatti, 2011). Cream skimming strategies refer to buying up the best impressions promising the highest value for a particular advertiser, while lemons avoidance refers to strategies avoiding the worst impressions (Abraham et al., 2013). In other words, bidders only bid on high-valued impressions, but they pretend to have no value for low-valued impressions or items, as they promise a lower return on investment. If all bidders only revealed their preferences for a small set of high-valued impressions, this might lead to a few good matches, but would also have a negative impact on efficiency and seller revenues overall (Levin and Milgrom, 2010). Many impressions would remain unsold. Our advertising model is motivated by these real-world observations.

### *1.1. Contributions*

In what follows, we want to study if there are truthful mechanisms with a good approximation ratio for knapsack bidders in the advertising model. We focus on offline markets, because if we cannot find truthful mechanisms with good approximation ratios for this environment, then we cannot hope for respective online mechanisms. Truthful online mechanisms are significantly more challenging, as we will discuss in the conclusions. In the *offline model* we actually get a positive result. We leverage insights from matching theory and use randomization in the allocation rule to incentivize truthful bidding. Interestingly, for knapsack bidders there is a randomized 4-approximate mechanism, which is much better than the $n$-approximation for the general case of value bidders, who have a budget for each package and cannot overbid.

Our paper nicely delineates the assumptions that are necessary for non-dictatorial, truthful, and prior-free mechanisms in markets where quasilinearity cannot be assumed. While the assumptions in our advertising model are realistic, a truthful mechanism with a 4-approximation is randomized and rather complex. The analysis is important for market designers, because it suggests that truthfulness might be too demanding for markets where knapsack bidders are an adequate utility model. This motivates the analysis of equilibrium bidding strategies in simple, but non-truthful auction mechanisms.

## 2. Related Literature

The literature on mechanism design without money is related to this paper. It is well-known that with general valuations, any non-dictatorial mechanism with at least three possible outcomes is not strategy-proof (Gibbard, 1973; Satterthwaite, 1975). Gibbard (1977) showed that every strategy-proof mechanism is a lottery over deterministic mechanisms each of which either has not more than two alternatives or is dictatorial. Even with more specific assumptions on the utility functions, truthful mechanisms appear to be restricted to sequential dictatorships. For example, a number of papers analyze specific assignment problems and bidders with responsive preferences, and show that only sequential dictatorships are strategy-proof and Pareto-optimal (Svensson, 1999; Papai, 2001; Ehlers and Klaus, 2003; Hatfield, 2009).

Procaccia and Tennenholtz (2009) introduced the technique of welfare approximation as a means to derive truthful approximation mechanisms without money. There have been positive results for environments with limited private information. For example, a few related papers analyze truthful mechanisms without money for a strategic variant of the generalized assignment problem (Dughmi and Ghosh, 2010; Chen et al., 2013; Fadaei and Bichler, 2017a). The assumptions in these models are different in a number of details leading to differences in the algorithms, but they also use approaches from matching and approximation as we do in our model.

There have been a number of recent papers about bidders with quasilinear utility functions and a private budget constraint, which is different to the knapsack utility model considered in this paper. For example, Dobzinski et al. (2012) showed that truthful and Pareto-optimal mechanisms without positive transfers are impossible with private budget constraints. Other authors have analyzed approximation mechanisms when bidders have quasilinear utility functions with a budget constraint (see for example Ashlagi et al. (2010); Dütting et al. (2015); Talman and Yang (2014)). A knapsack utility function is different from a quasilinear utility function with a budget constraint, and this difference has ample consequences for auction design.

## 3. The Model

Let us now introduce knapsack bidders formally. First, we draw connections to the value bidders as they have been analyzed in the literature and show that the welfare of truthful approximation mechanisms can be very low in general. Then we make additional assumptions motivated by advertising markets and show that each of these assumptions impacts the approximation ratios of truthful mechanisms.

*3.1. Knapsack Bidders and Value Bidders*

We have a set of $m$ heterogeneous items (or objects) $J$, one seller 0, and $n$ bidders (or agents), $I$. Each bidder $i \in I$ has a cardinal willingness-to-pay or value $v_{ij} \in \mathbb{R}_{\geq 0}$ for any item $j \in J$. There is one copy of each item. We describe the vector of valuations $v_{ij}$ of a bidder $i$ as $v_i \in \mathbb{R}^m_{\geq 0}$. The values $v_{ij}$ are normalized with 0 for the empty set. We denote $\mathcal{V} \subseteq \mathbb{R}^{n\times m}_{\geq 0}$ as the set of all possible valuation matrices. In addition to the willingness to pay or value for one item, bidders face an overall budget constraint $c_i$. In contrast to mechanism design with quasilinear utility functions, where utility is defined as valuation minus price of a bundle, we assume a knapsack utility function.

**Definition 1.** Given an allocation matrix $X \in \mathcal{X} \subseteq \{0,1\}^{n\times m}$, a price matrix $P \in \mathcal{P} \subseteq \mathbb{R}^{n\times m}_{\geq 0}$, and a budget vector $c \in \mathbb{R}^n_{\geq 0}$, the utility function

$$u_i\left(X,P\right) = \begin{cases} v_i^t x_i, & \text{if } p_i^t x_i \leq c_i \wedge \forall j \in J : p_{ij} x_{ij} \leq v_{ij} \\ -\infty, & else \end{cases}$$

is called *knapsack utility function.* Bidders with a respective utility function are called *knapsack bidders.*

In this definition, $x_i$ is a binary vector describing the allocation for bidder $i$ where $x_{ij} = 1$ when bidder $i$ wins item $j$. If we assume linear and anonymous prices, then $p_i$ is the vector of payments $p_{ij} = p_j$ for all items $j \in J$. Similar to the literature on quasilinear bidders with budget-constraints (Dobzinski et al., 2008), we assume the budget $c_i$ as exogenously given and focus on the strategies of knapsack bidders. As in the related literature on display ad auctions (see Introduction), bidder must not bid beyond his value $v_{ij}$ for an item, and the item value can be seen as an item-level budget limit. Suppose, an advertiser is selling cameras which yield a profit of 10\$. The expected profit considering the conversion rate of an impression can now determine the willingness-to-pay or value $v_{ij}$ of the bidder. An advertiser does not want the DSP to bid more than $v_{ij}$ as this would lead to a loss, when selling the product to the end consumer.

Let us now define relevant terms from mechanism design for markets with knapsack bidders. Based on the revelation principle, we focus only on direct revelation mechanisms.

**Definition 2.** A (direct revelation) *mechanism* comprises an allocation function $f : \mathcal{V} \to \mathcal{X}$ and a vector of payment functions $p_1, ..., p_n$, where $p_i : \mathcal{V} \to \mathbb{R}_{\geq 0}$ defines the payments of bidder $i$.

We aim for incentive-compatible and prior-free mechanisms, i.e., truthful bidding is an equilibrium even without the availability of prior value distributions.

**Definition 3.** A mechanism $\mathscr{M} = (f, P)$ is called *truthful* if for every bidder $i$, every $V \in \mathcal{V}$ and every $v_i' \in \mathcal{V}_i$, if we denote $x = f(v_i, V_{-i})$ and $x' = f(v_i', V_{-i})$, then $u_i(x, p) \geq u_i(x', p')$, with $p$ and $p'$ being the prices in the two allocations.

We are only interested in mechanisms which satisfy *individual rationality*, i.e., participation in the mechanism never makes the agent worse off. The second-price rule typically used in display ad auctions satisfies this assumption. Our focus are truthful and prior-free mechanisms. We talk about *strategy-proofness* if we have a deterministic mechanism and truthful bidding is a dominant strategy. For randomized mechanisms, we mainly focus on *truthfulness-in-expectation* (TIE), which will be defined in Section 5.

We assume a neutral auctioneer who provides the exchange and aims for *welfare maximization* for all participants. The buyers are the strategic bidders in this exchange, and, as usual, we assume that the seller has a value of zero for the items. But the knapsack utility function has a significant impact on the social choice function. While in the quasilinear utility model the prices of the items cancel out[2], for knapsack bidders payments (the

[2]with quasilinear utility in case of assignment we have for bidders $u_i(j) = v_{ij} - p_j$ and for the seller $u_s(j) = p_j$. Welfare is calculated as follows: $SW = \sum_{i \in I, j \in J} u_i(j) + \sum_{j \in J} u_s(j) = \sum_{i \in I, j \in j} (v_{ij} - p_j + p_j) x_{ij} = \sum_{ij} v_{ij} x_{ij}$

| | A | B | $c_i$ |
|---|---|---|---|
| Bidder 1 | 6 | 4 | 4 |
| Bidder 2 | 4 | 3 | 4 |
| Bidder 3 | 3 | 0 | 4 |

Table 1: Bidder values and and budgets.

utility of the sellers) need to be considered explicitly in a utilitarian social choice function, such that we get $\sum_{j\in J, i\in I}(v_{ij} + p_j)x_{ij}$.

**Example 2.** *Consider an auction with three bidders with values and budgets like in table* **??**. *A welfare maximizing mechanism that does ignore the sellers revenue at within the allocation would allocate item A to bidder 1 with price of* 4 *and item B to bidder 2 with price* 0. *This leads to a welfare of* 9 *for the bidders and a revenue of* 4 *for the seller, that is an overall welfare of* 13. *If we consider the sellers revenue already for the allocation, we would assign item B to bidder 1 and A to bidder 2, each at a price of* 3. *Now the bidders have a welfare of* 8 *and the seller a revenue of* 6, *that is an overall welfare of* 14. *Hence, it is important for the goal of welfare maximization to consider the sellers revenue as well as the bidders values.*

Unfortunately, we cannot hope for good approximation ratios if value bidders are allowed to submit bids on packages of objects. The extensive proof for the $n$ approximation for general markets in Fadaei and Bichler (2016) yields that a truthful mechanism can only elicit a single package value from each bidder. To avoid arbitrarily low approximation ratios, the auctioneer needs to elicit the valuation for the bundle of all objects from each bidder. The proofs in their model also hold for knapsack bidders who are allowed to bid on packages of items and pay-as-bid.

### *3.2. The Advertising Model*

Several assumptions from the model with value bidders do not carry over to advertising auctions in the field, and the results with value bidders might be too pessimistic. We define the *advertising model*, which adds three assumptions:

1. no package bidding is possible, and
2. the market is large with many items and many bidders, and
3. (linear) prices $p_j$ for each item are considered exogenous by bidders.

All three assumptions are met in display ad auctions, where bidders cannot bid on packages of impressions, and the auctioneer uses a second-price payment rule per impression. This assumption is also made in the economics literature for large markets (Roberts and Postlewaite, 1976). Also the literature on bidding strategies referenced in the introduction considers prices as an exogenous variable (Chen et al., 2011; Zhang et al., 2014).

Let us outline the allocation problem of the auctioneer in an offline auction. This *allocation problem* (AP) can be described as a binary program, where bidders (agents) $i$ with a budget constraint $c_i$ have a value $v_{ij}$ and an anonymous and linear price $p_j$ for each item $j$.

$$\begin{aligned}
\underset{\text{s.t.}}{\text{Max}} \sum_{i \in I} \sum_{j \in J} (v_{ij} + p_j) x_{ij} & & \qquad \textbf{(AP)} \\
\sum_{i \in I} x_{ij} \leq 1 & \qquad \forall j \in J & \qquad \text{(Supply)} \\
\sum_{j \in J} p_j x_{ij} \leq c_i & \qquad \forall i \in I & \qquad \text{(Budget)} \\
x_{ij} \in \{0, 1\} & \qquad \forall i \in I, j \in J & \qquad \text{(Binary)}
\end{aligned}$$

The optimal solution to AP maximizes welfare in a market, if the auctioneer had access to the true valuations and budget constraints of the bidders. A feasible assignment may allocate a subset of items $S$ to bidder $i$ such that $\sum_{j \in S} p_j \leq c_i$ (see constraint *Budget*). A feasible assignment may assign each item at most once (see constraint *Supply*). Note that this is a slightly modified version (the parameter $p_j$ in the objective function) of the generalized assignment problem, which has an integrality gap of 2 (Shmoys and Tardos, 1993). As indicated, we consider price as an exogenous variable for our strategic analysis and study whether truthful and prior-free mechanisms are possible at least in this model. In an online market, price $p_j$ could be determined via a second-price rule as in display ad auctions. In an offline market, the auctioneer might simply determine a fixed price per type of object based on the bid distribution. In order to emphasize the presence of strategic knapsack bidders, we will use $\mathrm{AM_{KB}}$ to describe the strategic version of AP. We refer to $\mathrm{AM_{KB}}$-LP as the LP relaxation of $\mathrm{AM_{KB}}$. The key decision of a knapsack bidder in the advertising model is whether to bid on an item or not at all.

## 4. Deterministic Approximation Mechanisms

Let us first analyze deterministic approximation mechanisms. Example 2 is sufficient to show that a truthful and welfare maximizing mechanism is not possible independent of the payment rule that we use in the presence of knapsack bidders.

**Example 3.** *Consider a market with two bidders and two items. Table 1 describes bidder valuations and their budget constraints. The allocation maximizing value is to assign item B to bidder 1 and item A to bidder 2 with a total bidder value of* $1+3\varepsilon$*. However, bidder 1 can increase his utility to* $1$*, by bidding on A only and pretending that his value for B is null. This way the auctioneer allocates A to bidder 1 and B to bidder 2 with a total value of* $1+2\varepsilon$*. Thus, there cannot be a deterministic and welfare maximizing mechanism that is also truthful.*

| | **Valuations** | | |
|---|---|---|---|
| Items | A | B | $c_i$ |
| Bidder 1 | 1 | $4\varepsilon$ | 1 |
| Bidder 2 | $1-\varepsilon$ | $2\varepsilon$ | 1 |

Table 2: Bidder valuations, where bidder 1 can shade his value for $B$ to win $A$.

Since we cannot hope for deterministic, truthful and welfare maximizing mechanisms, we relax the the goal of maximizing welfare and look at deterministic approximation mechanisms. A mechanism returns (at most) an $\alpha$-approximation of the optimal if its value is always greater than or equal to $^1/_\alpha$ times the optimal value ($\alpha \geq 1$). We first give a strategyproof mechanism in Algorithm 1 and analyze afterwards, how well it performs in terms of welfare. This serial dictatorship (SD) mechanism sorts the bidders by decreasing value of $min \left\{ c_i, \sum_j v_{ij} \right\}$. Then the algorithm maximizes the sum of valuations for each bidder subject to the budget constraint, removes the items allocated and the bidder and repeats until there is no bidder left.

**Algorithm 1:** The SD Algorithm

**Input** : Knapsack bidders $I$ with private budgets $c_i$ and values $V$ for items $J$, prices $p$

$\mathcal{L}$ is a list of bidders sorted by decreasing value of $min \left\{ c_i, \sum_j v_{ij} \right\}$.
$s_j = 1, \forall j \in J$
**for** *all* $i \in \mathcal{L}$ **do**
  Compute $argmax \left\{ v_i^t x_i \mid \sum_j p_j x_{ij} \leq c_i,\ x_{ij} \leq s_j,\ x_{ij} \in \{0, 1\} \right\}$.
  **for** $j$ *with* $x_{ij} = 1$ **do** $s_j = 0$

**Output**: assignment $X$

**Lemma 1.** *Algorithm 1 is a strategy-proof mechanism for knapsack bidders.*

*Proof.* The mechanism is truthful because a knapsack bidder cannot claim a higher budget $c_i$. These knapsack bidders can also not overbid. Shading the budget $c_i$, or hiding valuations for items to zero, only decreases their ranking in the algorithm and cannot improve their allocation.

□

Unfortunatly, Algorithm SD can lead to low welfare in worst case.

**Lemma 2.** *Algorithm SD achieves an approximation ratio not better than $m$.*

*Proof.* Consider a market with $m+1$ bidders and $m$ items with $p_j \leq 1$ for all $j \in J$. Table 2 describes bidder valuations and their budget constraints. In this market, a deterministic and value-maximizing auction would allocate item $j_i$ to bidder $i$. The welfare is $m^2 - \sum_{i=1}^{m} \varepsilon_i$. However, Algorithm 1 assigns first the most valued feasible set of items to bidder 0, since his sum of valuations as well as his budget is greater than those of the other bidders. That is, Algorithm 1 allocates all items to bidder 0 with a total utility of $m$ which implies an approximation ratio $\geq m$. $\square$

| | **Valuations** | | | | |
|---|---|---|---|---|---|
| Items | $j_1$ | $j_2$ | $\cdots$ | $j_m$ | $c_i$ |
| Bidder 0 | 1 | 1 | $\cdots$ | 1 | $m$ |
| Bidder 1 | $m-\varepsilon_1$ | 0 | $\cdots$ | 0 | $m-\varepsilon_1$ |
| Bidder 2 | 0 | $m-\varepsilon_2$ | $\cdots$ | 0 | $m-\varepsilon_2$ |
| $\vdots$ | $\vdots$ | $\vdots$ | $\ddots$ | $\vdots$ | $\vdots$ |
| Bidder m | 0 | 0 | $\cdots$ | $m-\varepsilon_m$ | $m-\varepsilon_m$ |

Table 3: Bidder valuations.

**Theorem 1.** *Algorithm SD is strategy-proof and it achieves an approximation ratio within* $\Theta(m)$.

*Proof.* Lemma 1 shows that SD is strategy-proof. Based on Lemma 2, we know that the approximation ratio for Algorithm 1 is within $\Omega(m)$. We only have to show that the approximation ratio is within $\mathcal{O}(m)$.

We relax the algorithm in line 4 to allow for fractional solutions $x_{ij} \in [0,1]$ and construct a feasible dual solution with a value at most $(m+1)$ times the value obtained by the relaxed algorithm. By calling the weak duality theorem together with the integrality gap of 2 Shmoys and Tardos (1993), the claim follows. Assume $x$ is the outcome of the relaxed Algorithm 1. Using $x$ we can construct a feasible solution to the dual of AM$_{\text{KB}}$-LP.

$$\begin{aligned} \underset{\text{s.t.}}{\text{Min}} \sum_{j=1}^{m} \rho_j + \sum_{i=1}^{n} c_i d_i & \qquad \textbf{(AM}_{\textbf{KB}}\textbf{-LPD)} \\ \rho_j + p_j d_i \geq v_{ij} + p_j,\ \forall i \in I, j \in J & \\ \rho, d \geq 0 & \end{aligned}$$

Observe that in Algorithm 1 the bidders are sorted in decreasing order of $min\left\{c_i, \sum_j v_{ij}\right\}$. Wlog. let $i_1, i_2, i_3, ..., i_n$ be the order of bidders computed by Algorithm 1. Furthermore, the algorithm assigns the current bidder available (fractions of) items $j$ in decreasing order of the density ${v_{ij}}/{p_j}$ to maximize his utility. Note, that the price $p_j$ for item $j$ is the same for all bidders $i \in I$ and therefore the ordering with respect to this density is the same as to ${(p_j+v_{ij})}/{p_j} = {v_{ij}}/{p_j} + 1$. Initially, let $\rho = \vec{0}$ and $d = \vec{0}$. If item $j$ gets exhausted by assigning $j$ to bidder $i_k$, set $\rho_j = p_j + max\ \{ v_{i_l j} \mid l \geq k \}$, the highest valuation for items $j$ over all bidders with a lower rank than $i_k$ plus the price of item $j$. Furthermore, for all bidders $i$ with exhausted budget, set $d_i = 1 + min\ \{ {v_{ij}}/{p_j} \mid \forall j \in J : x_{ij} > 0 \}$, that is, one plus the density of the last to $i$ assigned item. This satisfies the dual constraint in AM$_{\text{KB}}$-LPD. In particular, if the budget of bidder $i$ is exhausted, then for each item $j$ (fractionally) assigned to bidder $i$, either $j$ gets exhausted with this assignment or does not. If $j$ is exhausted we have $\rho_j \geq v_{ij} + p_j$ and therefore the constraint holds. If $j$ is not exhausted, we have either that $j$ is the last assigned item and therefore $d_i = {v_{ij}}/{p_j} + 1$, or $j$ is not assigned to $i$ and therefore $d_i \geq {v_{ij}}/{p_j} + 1$; and the constraint holds as well. If bidder $i$ has residual budget, every item $j$ which is assigned to it is exhausted by this assignment. That is, we have $\rho_j \geq v_{ij} + p_j$ and the constraint thus holds. For every item $j$ which is not assigned to the bidder but for which the bidder has a positive valuation, we have $\rho_j \geq v_{ij} + p_j$, since the item is exhausted due to the assignment to a bidder $i_k$ with higher rank than $i$, and we have $v_{ij} \leq max\ \{ v_{i_l j} \mid l \geq k \} = \rho_j - p_j$. In sum, we have constructed a feasible dual solution using $x$, the allocation resulting from the relaxed Algorithm 1.

Now, we bound the value of the dual solution with respect to the primal solution. First, we observe that $\sum_{i,j}(v_{ij} + p_j)x_{ij} \geq \sum_{i,j} d_i p_j x_{ij} = \sum_i d_i \sum_j p_j x_{ij}$, because $d_i$ is only non-zero if the budget of bidder $i$ is exhausted, $d_i p_j = \left(1 + min\ \left\{ {v_{ij'}}/{p_{j'}} \mid \forall j' \in J : x_{ij'} > 0 \right\}\right) p_j \leq ({v_{ij}}/{p_j} + 1)p_j = v_{ij} + p_j$ for items with $x_{ij} > 0$, and $(p_j + v_{ij})x_{ij} = d_i p_j x_{ij} = 0$ for items $j$ with $x_{ij} = 0$. Second, we show that $m \sum_{i,j}(p_j + v_{ij})x_{ij} \geq \sum_j \rho_j \sum_i x_{ij}$.

Now we sum up both inequations and obtain

$$\begin{aligned}(m+1)\sum_{i,j}(v_{ij}+p_j)x_{ij} & \\ &\geq \sum_j \rho_j \sum_i x_{ij} + \sum_i d_i \sum_j p_j x_{ij} \\ &= \sum_j \rho_j + \sum_i d_i c_i\end{aligned}$$

Notice, only for items $j$ which get exhausted ($\sum_i x_{ij} = 1$) we have $\rho_j > 0$ and only for full bidders ($\sum_j p_j x_{ij} = c_i$) we have $d_i > 0$. The final term is the value of the dual, the desired conclusion.

It remains to show that $m\sum_{i,j}(v_{ij} + p_j)x_{ij} \geq \sum_j \rho_j \sum_i x_{ij}$. Since $v_{ij} \leq c_i$ for all $i \in I$ and $j \in J$ we can observe, that $(\rho_j - p_j)x_{i_k j} = max\,\{\, v_{i_l j} \mid l \geq k \,\}\, x_{i_k j} \leq c_{i_k}$. This follows from the ordering of the bidders in Algorithm 1. For all bidders $i_l$ with a lower rank than bidder $i_k$ we have $min\left\{ c_{i_l}, \sum_j v_{i_l j} \right\} \leq min\left\{ c_{i_k}, \sum_j v_{i_k j} \right\} \leq c_{i_k}$ . That means either $v_{i_l j} \leq c_{i_l} \leq c_{i_k}$ or $v_{i_l j} \leq \sum_j v_{i_l j} \leq c_{i_k}$ for all items $j$. That is, the assignment of one item to bidder $i_k$ with dual price $\rho_j - p_j$ is always primal feasible as well.[3] The worst case is, that bidder $i_k$ gets all $m$ items, but all these items $j$ have a high (dual) price $\rho_j - p_j$ such that $i_k$ can buy at most one of them without violating his budget constraint. Since for all those items $j$ of bidder $i$ $\rho_j - p_j \leq c_i$, we have

$$\begin{aligned}\sum_{i,j} v_{ij}x_{ij} &\geq {}^1\!/_m \sum_j(\rho_j - p_j)\sum_i x_{ij} \\ &= {}^1\!/_m \sum_j \rho_j \sum_i x_{ij} - {}^1\!/_m \sum_j p_j \sum_i x_{ij}\end{aligned}$$

and therefore

$$\begin{aligned}\sum_{i,j}(v_{ij} + p_j)x_{ij} &\geq \sum_{i,j}(v_{ij} + {}^{p_j}\!/_m)x_{ij} \\ &= \sum_{i,j} v_{ij}x_{ij} + {}^1\!/_m \sum_j p_j \sum_i x_{ij} \\ &\geq {}^1\!/_m \sum_j \rho_j \sum_i x_{ij},\end{aligned}$$

the desired conclusion.

---

[3] Since prices $p_j$ do not depend on the allocation and are the same for all bidders, a loss in social welfare can only occur due to a sub-optimal assignment of the items to bidders if all items are assigned. Hence, the dual prices $\rho_j - p_j$ can bee interpreted as opportunity costs for the bidders.

□

Algorithm 1 suggests that deterministic and truthful mechanisms yield low welfare. It is straightforward to show that the ratio is tight for a market with two knapsack bidders and two items only.

**Proposition 1.** *No strategy-proof and deterministic mechanism for knapsack bidders exists with an approximation ratio better than* 2 *of the optimal welfare for markets with two bidders and two items.*

*Proof.* Consider a market with two bidders and two items. Table 3 describes bidder valuations and their budget constraints. In this market, a deterministic and value-maximizing auction would allocate item 2 to bidder 1 and item 1 to bidder 2. The welfare is $2 - 2\varepsilon$. Bidder 1 can shade his bid for item 2 to zero, such that only 1 would be allocated to him. That is, any strategy-proof mechanism has to assign item 1 to bidder 1 before trying to assign item 2, leading to an approximation ratio of 2. This is true for any price $p_j \leq 1 - \varepsilon$.

| | **Valuations** | | |
|---|---|---|---|
| Items | 1 | 2 | $c_i$ |
| Bidder 1 | 1 | $1 - \varepsilon$ | 1 |
| Bidder 2 | $1 - \varepsilon$ | 0 | $1 - \varepsilon$ |

Table 4: Bidder valuations.

We leave open a lower bound on the approximation ratio for truthful and general markets with arbitrary numbers of bidders and items, which is challenging to show as the general result for value bidders suggests (Fadaei and Bichler, 2016). However, a simple example suggests that also large markets with knapsack bidders are susceptible to similar types of manipulation as in the proof to Proposition 1 and this can lead to significant welfare losses. Note that the example is a stylized form of cream skimming, which was described as a wide-spread bidding strategy in display ad auctions in the introduction.

**Example 4.** *Consider bidder 1 who has a value of \$1 for 100 impressions of type A and a value of \$0.9 for 100 impressions of type B. His budget is \$50. Bidders 2 and 3 have a value of \$0.5 for type A valuations and a budget of \$50. The welfare maximizing allocation would be to allocate A impressions to either bidder 2 or 3 and the B impressions to bidder 1. Welfare would be \$190, the total of \$90 from bidder 1 plus \$50 from bidder 2 plus \$50 for the auctioneer, which is the payment of bidder 2 or 3, if we assume a second-price rule. However, if bidder 1 does not bid on B impressions, he will be awarded the A impressions and the welfare will only be \$150, i.e., \$100 for bidder 1 and \$50 for the auctioneer, which is the payment of bidder 1 with a second-price rule. Bidder 1 would increase his utility from \$90 for the B impressions to \$100 for the A impressions.*

## 5. Randomized Approximation Mechanisms

We now resort to randomization as a means to improve the approximation guarantees, but first introduce some definitions. Let $\mathcal{A}$ denote a randomized algorithm which takes instance $(V, c)$ of $\text{AM}_{\text{KB}}$ and computes $X \in \mathcal{X}$, an assignment of items to bidders. The assignment is deterministic (each bidder receives a set of items), but algorithm $\mathcal{A}$ is randomized, i.e., $\mathcal{A}$ returns a solution which is randomly chosen according to a probability distribution over feasible assignments. Truthfulness-in-expectation is a solution concept for randomized approximation mechanisms.

**Definition 4.** A randomized algorithm $\mathcal{A}$ is said to be *truthful-in-expectation (TIE)* if we have for $X \sim \mathcal{A}(V, c)$

1. (feasibility) $\forall j \in J : \mathbb{P}\left(\sum_{i \in I} x_{ij} \leq 1\right) = 1$ and
   $\forall i \in I : \mathbb{P}\left(\sum_{j \in J} p_{ij} x_{ij} \leq c_i \wedge \forall j \in J : p_{ij} \leq v_{ij}\right) = 1$
2. (truthfulness) for any $i$ and $v_i' \in \mathbb{R}^m_{\geq 0}$, we have
   $\mathbb{E}\left[\sum_{j \in J} v_{ij} x_{ij}\right] \geq \mathbb{E}\left[\sum_{j \in J} v_{ij} x_{ij}'\right]$, where $X' \sim \mathcal{A}\left(v_i' \cup V_{-i}, c\right)$.

Note, the expected value of the bidder in both cases is computed with respect to his true valuations. Random Serial Dictatorship (RSD) is a well-known example of a randomized mechanism, but it is easy to see that the welfare of RSD can be arbitrary low. For example, RSD might first assign all items to a bidder with low valuations and sufficient budget, although other bidders have a much higher valuation. Therefore, our objective is to propose a randomized algorithm $\mathcal{A}$ for $\mathrm{AM}_{\mathrm{KB}}$ which is truthful and always returns a feasible assignment whose value approximates the optimal total value as well as possible.

Our technique is a relax-and-round approach with oblivious rounding as it has been used recently in the literature on approximation mechanisms (Lavi and Swamy, 2011; Dughmi and Ghosh, 2010; Fadaei and Bichler, 2017a). We design a *fractionally truthful* approximation algorithm which returns a feasible solution to $\mathrm{AM}_{\mathrm{KB}}$-LP. A fractionally truthful algorithm allocates fractional assignments to bidders $i$, and no bidder can improve its fractional value by an untruthful report. In particular, a fractionally truthful algorithm $\mathcal{A}^F$ takes $(V, c)$ and returns $X \in \mathcal{X}$, a feasible solution to $\mathrm{AM}_{\mathrm{KB}}$-LP with the following property. For each bidder $i$, if the bidder reports $v'_i \in \mathbb{R}^m_{\geq 0}$, we have $\sum_{j \in J} v_{ij} x_{ij} \geq \sum_{j \in J} v_{ij} x'_{ij}$, where $x' = \mathcal{A}^F(v'_i \cup V_{-i}, c)$. Next, we round the fractional solution using a special rounding technique which makes sure that each bidder obtains a fixed fraction of its fractional value in expectation. The *randomized meta-rounding* (Carr and Vempala, 2000) is capable of maintaining this fixed fraction.

To use the randomized meta-rounding, we have to scale down the fractional solution by factor 2, which is the integrality gap of the $\mathrm{AM}_{\mathrm{KB}}$-LP (Shmoys and Tardos, 1993). Assuming $X^* = \mathcal{A}^F(V, c)$, the randomized meta-rounding represents $X^*/2$ as a convex combination of polynomially-many feasible integer solutions. Looking at the provided convex combination as a probability distribution over integer solutions, we sample a randomized solution $X \in \mathcal{X}$ which is always feasible, and its expected value is $1/2$ of the fractional value of $X^*_F$. This is confirmed by Theorem 2, which has been proved in related work by Dughmi and Ghosh (2010).

**Theorem 2.** *(Dughmi and Ghosh, 2010) If there exists a fractionally truthful $\alpha$-approximation algorithm for $\mathrm{AM}_{\mathrm{KB}}$, then there exists a TIE $(2\alpha)$-approximation solution for $\mathrm{AM}_{\mathrm{KB}}$.*

Even though, their model is different, we can draw on this theorem. We propose Algorithm 2, an adaptation of the Deferred Acceptance algorithm by Gale and Shapley (1962). In Algorithm 2 all bidders with a positive budget submit a bid for the item with the highest density ${v_{ij}}/{p_j}$ from their preference list (*Pref*) and delete this item from this list. Afterwards, all items $j \in J$ sort their list of not rejected bids ($B_j$) according to the value of the bids, and accept all (fractional) bids until their supply ($s_j$) of 1 is exhausted. If the current item $j$ has sufficient residual supply to store the whole amount of the current bid $b_{ij}$, we assign (temporary) $j$ to $i$ but only that fraction that the bidder is still able to pay (line 13).

If the supply of $j$ is still not zero, but not high enough to store the whole bid, we assign $i$ only the available fraction of $j$ (line 19). If the supply of $j$ is exhausted, we reject the bid finally (line 18), since the threshold for assignment weakly increases each iteration. If no bidder can submit a new bid, the algorithm terminates and returns the assignment matrix $X$.

**Lemma 3.** *Algorithm 2 is truthful.*

*Proof.* Knapsack bidders only bid 0 or their true value. Since bidders propose to items in decreasing order of the density ${v_{ij}}/{p_j}$, hiding values can only lead to an assignment to items with a lower density, and therefore decrease the utility of a bidder. Therefore, knapsack bidders have no incentives to hide valuations. Next, we discuss the budget $c_i$. A knapsack bidder must not spend more than $c_i$. If a bidder reports a lower $\hat{c}_i < c_i$ there are two possible cases. If $\hat{c}_i$ is not exhausted during the process, it does not have an impact on the allocation computed by the algorithm. If $\hat{c}_i$ becomes binding, it can only restrict the assignment of additional items to a bidder, which would otherwise have been possible, and therefore decrease the total sum of valuations of a bidder. That is, Algorithm 2 is truthful for knapsack bidders. □

**Lemma 4.** *Algorithm 2 returns a* 2*-approximation solution to* $AM_{KB}$*-LP.*

*Proof.* Based on the proof of Proposition 1, we know that an approximation ratio better than 2 is impossible. We only have to show that it is not worse

**Algorithm 2:** Fractional Deferred Acceptance Algorithm for Knapsack Bidders.

**Input** : Knapsack bidders $I$ with private budgets $c_i$ and values $V$ for items $J$, prices $p$

$x_{ij} = 0,\ \forall i \in I, j \in J$
$Pref_i = \{ j \mid v_{ij} > 0, j \in J \},\ \forall i \in I$
$B_j = \emptyset,\ \forall j \in J$
**while** $(\forall i \in I, c_i > 0 : Pref_i \neq \emptyset)$ **do**
  **for** *all* $i \in I$ *with* $c_i > 0$ **do**
    $j = argmax \{ v_{ij}/p_j \mid j \in Pref_i \}$
    $b_{ij} = v_{ij}$
    $Pref_i = Pref_i \setminus j,\ B_j = B_j \cup b_{ij}$
  **for** *all* $j \in J$ **do**
    sort $B_j$ according to $b_{ij}$, $s_j = 1$
    **for** *all* $b_{ij} \in B_j$ **do**
      $c_i = c_i + p_j x_{ij},\ x_{ij} = 0$
      **if** $(s_j - min \{ 1, c_i/p_j \} \geq 0)$ **then**
        $x_{ij} = min \{ 1, c_i/p_j \}$
        $s_j = s_j - min \{ 1, c_i/p_j \}$
        $c_i = c_i - min \{ p_j, c_i \}$
      **else**
        **if** $(s_j = 0)$ **then** $B_j = B_j \setminus b_{ij}$
        $x_{ij} = s_j$
        $c_i = c_i - p_j s_j,\ s_j = 0$

**Output**: assignment $X$

than 2. For this, we construct a feasible dual solution with a value at most twice the value obtained by the algorithm, then by calling the weak duality theorem, the claim follows. Assume $x$ is the outcome of Algorithm 2. Using $x$ we can construct a feasible solution to the dual of AM$_{\text{KB}}$-LP.

Observe that in Algorithm 2 a bidder $i$ proposes items in decreasing order of the density $v_{ij}/p_j$, and item $j$ examines bids in decreasing order of $v_{ij}$. Since $p_j$ is the same for all bidders $i$, the ordering of the bids at an item $j$ according to $v_{ij}/p_j$ is the same as to $v_{ij}$ and $v_{ij} + p_j$. Initially, let $\rho = \vec{0}$ and $d = \vec{0}$. If item $j$ gets exhausted, set $\rho_j = v_{i_k j} + p_j$, where $i_k = argmin \{ v_{ij} \mid \forall i \in I : x_{ij} > 0 \}$. Furthermore, for all bidder $i$ with exhausted budget, set $d_i = v_{ij_k}/p_{j_k} + 1$, where $j_k = argmin \{ v_{ij}/p_j \mid \forall j \in J : x_{ij} > 0 \}$ is the last item assigned to bidder $i$. This satisfies the dual constraint in AM$_{\text{KB}}$-LPD. In particular, if the budget of bidder $i$ is exhausted, then for each $j$ of bidder $i$, either $j$ gets exhausted with this assignment or does not. If $j$ is exhausted we have $\rho_j = v_{ij} + p_j$ and therefore the constraint holds. If $j$ is not exhausted, we have $d_i \geq v_{ij}/p_j + 1$ and thus the constraint holds as well. If bidder $i$ has residual budget, every item $j$ which is assigned to it is exhausted by this assignment. That is, we have $\rho_j = v_{ij} + p_j$ and the constraint thus holds. For every item $j$ which is not assigned to the bidder but the bidder has a positive valuation, we have $\rho_j \geq v_{ij} + p_j$, since the item is exhausted due to another assignment $v_{i'j} \geq v_{ij}$ which is equivalent to $v_{i'j} + p_j \geq v_{ij} + p_j$. In sum, we have constructed a feasible dual solution using $x$, the allocation resulting from Algorithm 2.

Now, we bound the value of the dual solution with respect to the primal solution. First, we observe that $\sum_{i,j}(v_{ij} + p_j)x_{ij} \geq \sum_j \rho_j \sum_i x_{ij}$, since $\rho_j$ is non-zero only if $j$ is fully assigned. Second, $\sum_{i,j}(v_{ij} + p_j)x_{ij} \geq \sum_i d_i \sum_j (p_j x_{ij})$, because $d_i$ is only non-zero if the budget of bidder $i$ is exhausted, and then it holds $d_i p_j \leq (v_{ij}/p_j + 1)p_j = v_{ij} + p_j$ for all $j$ with $x_{ij} > 0$. Now we sum up both inequalities and obtain

$$\begin{aligned} 2\sum_{i,j}(v_{ij} + p_j)x_{ij} &\geq \sum_j \rho_j \sum_i x_{ij} + \sum_i d_i \sum_j p_j x_{ij} \\ &= \sum_j \rho_j + \sum_i d_i c_i \end{aligned}$$

Notice, only for item $j$ which get exhausted ($\sum_i x_{ij} = 1$) we have $\rho_j > 0$ and only for full bidder ($\sum_j p_j x_{ij} = c_i$) we have $d_i > 0$. The final term is the value of the dual, the desired conclusion. □

Finally, we use Theorem 2 together with Lemma 3 and 4 and obtain the following.

**Theorem 3.** *There exists a* TIE *4-approximation mechanism for the* AM$_{\text{KB}}$ *with strategic bidders.*

## 6. Conclusions

Arguments for second-price auctions in advertising markets are often based on the standard assumption of quasilinear utility functions. Based on literature about bidding strategies, quasilinearity does not appear to be an adequate model for bidding agents in these markets. Rather agents try to maximize the total sum of valuations subject to an overall budget that is devoted to a campaign. Mechanism design tries to devise auctions where truthful bidding is an equilibrium. This is typically hard to achieve in environments, where bidders do not have quasi-linear utility functions and their types are multi-dimensional. We have explored possibilities for truthful approximation mechanisms with knapsack utility functions in large advertising markets. An offline version of this environment allows for truthful 4-approximate mechanisms. The paper highlights the types of assumptions which allow for truthful mechanisms in an offline market.

It is interesting to understand the approximation ratio of truthful online mechanisms as, for example, display ads need to be sold one at a time and allocation and pricing need to be decided dynamically. Online mechanisms are often analyzed in the adversarial setting with items arriving in the worst possible order. The approximation ratios achieved in this paper for the offline environment can be seen as lower bounds on what can be achieved with truthful online mechanisms. Online mechanism design is, however, much less well understood even with quasilinear utility functions, and only recently truthful mechanisms were designed. Typically, the literature assumes unit-demand bidders and only a few papers deal with multi-unit demand assuming homogeneous goods and quasilinear utility functions. Devanur et al. (2015) analyzed online mechanisms for quasilinear bidders who have preferences for multiple heterogeneous items arriving over time. They show that there is no deterministic truthful and individually rational mechanism that gets any finite approximation factor to the optimal social welfare, even if payments can be computed after all items arrived. We conjecture that the design of truthful online mechanisms will be as challenging for knapsack bidders, but leave this question for future research.

## 7. References


Abraham, I., Athey, S., Babaioff, M., Grubb, M., 2013. Peaches, lemons, and cookies: designing auction markets with dispersed information., in: EC, pp. 7–8.

Ashlagi, I., Braverman, M., Hassidim, A., Lavi, R., Tennenholtz, M., 2010. Position auctions with budgets: Existence and uniqueness. The BE Journal of Theoretical Economics 10.

Baisa, B., 2013. Auction design without quasilinear preferences. Available at SSRN 2311753 .

Berg, J., Greenwald, A., Naroditskiy, V., Sodomka, E., 2010. A knapsack-based approach to bidding in ad auctions., in: ECAI, pp. 1013–1014.

Bergemann, D., Bonatti, A., 2011. Targeting in advertising markets: implications for offline versus online media. The RAND Journal of Economics 42, 417–443.

Borgers, T., Krahmer, D., Strausz, R., 2015. An introduction to the theory of mechanism design. Oxford University Press, USA.

Carr, R., Vempala, S., 2000. Randomized metarounding, in: Proceedings of the thirty-second annual ACM symposium on Theory of computing, ACM. pp. 58–62.

Chen, N., Gravin, N., Lu, P., 2013. Truthful generalized assignments via stable matching. Mathematics of Operations Research 39, 722–736.

Chen, Y., Berkhin, P., Anderson, B., Devanur, N.R., 2011. Real-time bidding algorithms for performance-based display ad allocation, in: Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM. pp. 1307–1315.

Cui, Y., Zhang, R., Li, W., Mao, J., 2011. Bid landscape forecasting in online ad exchange marketplace, in: Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM. pp. 265–273.

Devanur, N.R., Sivan, B., Syrgkanis, V., 2015. Multi-parameter auctions with online supply. arXiv preprint arXiv:1511.03699 .

Dobzinski, S., Lavi, R., Nisan, N., 2008. Multi-unit auctions with budget limits, in: Foundations of Computer Science, Philadelphia, PA. pp. 260–269.

Dobzinski, S., Lavi, R., Nisan, N., 2012. Multi-unit auctions with budget limits. Games and Economic Behavior 74, 486 – 503.

Dughmi, S., Ghosh, A., 2010. Truthful assignment without money, in: Proceedings of the 11th ACM conference on Electronic commerce, ACM. pp. 325–334.

Dütting, P., Henzinger, M., Starnberger, M., 2015. Auctions for heterogeneous items and budget limits. ACM Transactions on Economics and Computation 4, 4.

Ehlers, L., Klaus, B., 2003. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. Social Choice and Welfare 21, 265–280.

Fadaei, S., Bichler, M., 2016. Truthfulness and approximation with value-maximizing bidders, in: Symposium on Algorithmic Game Theory. Springer, pp. 566–576.

Fadaei, S., Bichler, M., 2017a. Generalized assignment problem: Truthful mechanism design without money, in: Operations Research Letters. Elsevier. volume 45, pp. 72–76.

Fadaei, S., Bichler, M., 2017b. Truthfulness with value-maximizing bidders: On the limits of approximation in combinatorial auctions. European Journal of Operational Research 14, 1–11.

Feldman, J., Muthukrishnan, S., Nikolova, E., Pal, M., 2008. A truthful mechanism for offline ad slot scheduling, in: Monien, B., Schroeder, U. (Eds.), Algorithmic Game Theory. Springer Berlin Heidelberg. volume 4997 of *Lecture Notes in Computer Science*, pp. 182–193.

Gale, D., Shapley, L., 1962. College admissions and the stability of marriage. American Mathematical Monthly 69, 9–14.

Ghosh, A., Rubinstein, B.I., Vassilvitskii, S., Zinkevich, M., 2009. Adaptive bidding for display advertising, in: Proceedings of the 18th international conference on World wide web, ACM. pp. 251–260.

Gibbard, A., 1973. Manipulation of voting schemes: a general result. Econometrica 41, 587–601.

Gibbard, A., 1977. Manipulation of schemes that mix voting with chance. Econometrica: Journal of the Econometric Society , 665–681.

Goetzendorff, A., Bichler, M., Shabalin, P., Day, R.W., 2015. Compact bid languages and core pricing in large multi-item auctions. Management Science 61, 1684–1703.

Green, J., Laffont, J., 1979. Incentives in public decision making. North Holland, Amsterdaml.

Hatfield, J., 2009. Strategy-proof, efficient, and nonbossy quota allocations. Social Choice and Welfare 33, 505–515.

Kazumura, T., Serizawa, S., 2015. Efficiency and strategy-proofness in object assignment problems with multi demand preferences .

Lavi, R., Swamy, C., 2011. Truthful and near-optimal mechanism design via linear programming. Journal of the ACM (JACM) 58, 25.

Lee, K.C., Jalali, A., Dasdan, A., 2013. Real time bid optimization with smooth budget delivery in online advertising, in: Proceedings of the Seventh International Workshop on Data Mining for Online Advertising, ACM. p. 1.

Levin, J., Milgrom, P., 2010. Online advertising: Heterogeneity and conflation in market design. The American Economic Review 100, 603–607.

Mas-Colell, A., Whinston, M.D., Green, J.R., et al., 1995. Microeconomic theory. volume 1. Oxford university press New York.

McAfee, R.P., 2011. The design of advertising exchanges. Rev Ind Organ 39, 169–185.

Nisan, N., Bayer, J., Chandra, D., Franji, T., Gardner, R., Matias, Y., Rhodes, N., Seltzer, M., Tom, D., Varian, H., et al., 2009. Googles auction for tv ads. Automata, Languages and Programming , 309–327.

Papai, S., 2001. Strategyproof and nonbossy multiple assignments. Journal of Public Economic Theory 3, 257–271.

Procaccia, A.D., Tennenholtz, M., 2009. Approximate mechanism design without money, in: Proceedings of the 10th ACM conference on Electronic commerce, ACM. pp. 177–186.

Roberts, D.J., Postlewaite, A., 1976. The incentives for price-taking behavior in large exchange economies. Econometrica: Journal of the Econometric Society , 115–127.

Satterthwaite, M.A., 1975. Strategy-proofness and arrow's conditions. existence and correspondence theorems for voting procedures and social welfare functions. Journal of Economic Theory 10.

Shmoys, D.B., Tardos, É., 1993. An approximation algorithm for the generalized assignment problem. Mathematical Programming 62, 461–474.

Svensson, L.G., 1999. Strategy-proof allocation of indivisible goods. Social Choice and Welfare 16, 557–567.

Talman, A., Yang, Z., 2014. An Efficient Multi-Item Dynamic Auction with Budget Constrained Bidders. volume 44.

Zhang, W., Yuan, S., Wang, J., 2014. Optimal real-time bidding for display advertising, in: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, NY, USA. pp. 1077–1086.

Zhou, Y., Chakrabarty, D., Lukose, R., 2008. Budget constrained bidding in keyword auctions and online knapsack problems, in: Internet and Network Economics. Springer, pp. 566–576.