Online multiple-item multi-unit auctions

Shipra Agrawal

May 6, 2009

Abstract

We consider the problem of designing competetive algorithms for online multi-item multi-unit auctions. Consider a seller with m non-identical items, k units of each item are available. The bidders arrive online, each bidder demands a subset of items and offers certain bid price. The seller must allocate the available items to a subset of bidders selected in an online manner so as to maximize the total value of their bids. This is a multi-item generalization of secretary problems. Under the assumption that bidders arrive in a random order, we present a truthful, $\log(mk)$-competetive online allocation scheme. Previously, a worst case competetive ratio of $\log(mU/L)$ was known for this problem under the assumption that the bid values are in the range $[L, U]$. In contrast, we do not assume any knowledge regarding the range or distribution of values beyond the fact that the order is random. We demonstrate an application of our algorithm to achieve revenue-competetiveness when the the valuations of agents are drawn independently from known distributions.

1 Introduction

Consider the situation a seller faces when selling multiple units of an item (or multiple items) to bidders who arrive one at a time. Each bidder t has certain type θ_t, and offers a bid price b_t for an allocation that satisfies him. On observing the bidder t, the seller must decide whether to satisfy the bidder or not ($x_t \in \{0, 1\}$), and what price to charge him ($\mathcal{P}_t$). The seller can only satisfy certain "feasible" sets of bidder types simultaneously. The utility of the bidder t is $\pi_t x_t - \mathcal{P}_t$ where π_t is the true valuation of the bidder. The objective of the seller is to pick a feasible subset of bidders in an online manner so that the total valuation of satisfied bidders or the total money collected from the bidders is maximized.

Various models of online auctions considered in the literature can be described in this form. For example, consider an auction where seller has k units of an item and each bidder has a unit demand. This problem is also known as multiple choice secretary problem [1, 7]. Here, the types of bidders are uniform, and any set of atmost k bidders is a feasible set. In an auction where seller has unlimited supply of an item and each bidder has unit demand [3], any set of bidders is a feasible set. In matroid secretary problems [2], the feasible sets of bidders are the independent sets of a matroid. An example is transversal matroid problem, where the seller has m non-identical items. Each bidder specifies a bundle of items $\vec{a}_t \in \{0, 1\}^m$, and is satisfied if *any* item in his bundle is allocated to him. Here, a feasible set is a matching of bidders to items. In this paper, we also consider combinatorial multi-item multi-unit auctions where the seller has k units each of m non-identical items. The type of each bidder t is the bundle of items $\vec{a}_t$ that he desires. A bidder is satisfied if and only if *all* the items in his bundle are allocated to him. A feasible set is any subset of bidders such that total bundle of items they desire is less than the inventory initially available with the seller, i.e. $\sum_t \vec{a}_t x_t \leq k\vec{e}$. Note that in this problem, the feasible set of bidders do not form a matroid, so it is not a special case of matroid secretary problem.

More precisely, let $\mathcal{L}$ denote the collection of feasible sets of bidder types. The online auction problem can be stated as the problem of designing a pricing and allocation mechanism that is

1. **Online:** The price $\mathcal{P}_t$ and allocation x_t for bidder t must be decided before observing the bidders $t+1, \ldots, n$.

2. **Feasible:** At any time, the types of bidders who receive an allocation must form a feasible type set, i.e. at any time t, $\{\theta_s | s \leq t, x_s > 0\} \in \mathcal{L}$

3. **Supports voluntary participation:** For any bidder who receives an allocation, the actual price must be less than or equal his bid price, i.e. $\mathcal{P}_t x_t \leq b_t x_t$.

4. **Incentive compatible:** Bidding true value should be a dominant strategy for bidders. Let π_t denote the true valuation for t^{th} bidder, then in an incentive compatible mechanism a dominant strategy is to set $b_t = \pi_t$. Note that this definition only considers misreports of valuations π_t. In some models, a bidder may also consider mis-reporting his type θ_t. In this general model, a mechanism is incentive compatible if a dominant strategy for bidders is to report both the valuation π_t and type θ_t truthfully.

The goal is to design a mechanism that satisfies above properties and addditionally achieves either good revenue or efficiency.

- **Efficiency**, also known as social utility, is the total valuation of the satisfied bidders. In a truthful mechanism where $b_t = \pi_t$ and the type θ_t is reported truthfully, the social utility is simply the sum of bid values of satisfied bidders, i.e. $\sum_t b_t x_t$.

- **Revenue** for the market organizer is the total money collected from the bidders, that is, the sum of the actual prices offered by the seller to the satisfied bidders $\sum_t \mathcal{P}_t x_t$.

2 A generalized mechanism for online auctions

In this section, we describe the existing results for online auctions using a common generalized algorithmic framework. This generalized mechanism is described as Algorithm 1 below. A priliminary form of Algorithm 1 (that does not include types of bidders) appears in the Blum, Hartline, 2005. We show that all the existing online mechanisms [1, 3, 9, 7, 2] can be cast into this general mechanism. This online mechanism will be shown to have many nice properties. It naturally satisfies the requirements of feasiblity and voluntary participation. Under some simple conditions, it will guarantee incentive compatibility. Further, as we demonstrate later in Section 3.2, 4, it also allows to extend bounds on efficiency competetivenss to matching bounds on revenue competetiveness. In this paper, we provide a new algorithm for online multi-unit multi-item auctions; our algorithm will take this general form as well.

It is easy to see that Algorithm 1 is online, feasbile and supports voluntary participation. Since the threshold price for bidder t is decided without observing the bid b_t, it is also incentive compatible assuming that the only misreports are of the valuations π_t, i.e. if the bidders always report their types θ_t truthfully.

Algorithm 1 Online auction

Let $F = \phi$ denote the current set of satisfied bidder types. For each bidder t,

1. $\mathcal{P}_t \leftarrow f_t(\theta_1, b_1, \ldots, \theta_{t-1}, b_{t-1}, \theta_t)$
2. If $\mathcal{P}_t \leq b_t$ and if $F \cup \theta_t \in \mathcal{L}$, sell to bidder t at price $\mathcal{P}_t$. That is, set $x_t = 1$ and $F = F \cup \theta_t$.
3. Otherwise, reject bidder t (set $x_t = 0$).

We consider a more general model where bidders may also misreport their types, however we consider only the models where the types have a partial order of the following form:

Definition 1. *There exists a partial order $\succeq$ on space of possible types such that for any sets of bidders L, L', if $L \succeq L'$, then $L \in \mathcal{L}$ implies $L' \in \mathcal{L}$, i.e. L is feasible implies that L' is feasible. Here, $L \succeq L'$ if for every bidder t in L there exists a distinct bidder t' in L' such that $\theta_t \succeq \theta_{t'}$.*

Given the partial order, we will assume that only misreports of types are over-reports.

Assumption 1. *(type misreports) A bidder of type θ_t can only over-report his type as θ' where $\theta' \succeq \theta$.*

Above assumption is natural in many auction models. For example, consider the transversal matroid auction where type of each bidder is a bundle of items $\vec{a}$ and he is satisfied only if *atleast one* item in the bundle is allocated to him. Then, define a partial order $\succeq$ as $\vec{a} \succeq \vec{a}'$ iff $\vec{a} \leq \vec{a}'$. This order satisfies the condition in Definition 1. The bidder with type $\vec{a}$ will not mis-report his type as $\vec{a}'$ for some $\vec{a}' \geq \vec{a}$ because he has no utility for any items in $\vec{a}' - \vec{a}$, however he may consider reporting his type as $\vec{a}' \leq \vec{a}$ if that gets him a better allocation.

Similarly, in the multi-item combinatorial auction where a bidder is satisfied only if *all* the items in the bundle (or any superset of those) are allocated to him, the partial order $\succeq$ is defined as $\vec{a} \succeq \vec{a}'$ iff $\vec{a} \geq \vec{a}'$. He may report $\vec{a}' \geq \vec{a}$ if that gets him a better allocation. However his utility is 0 for any $\vec{a}' \leq \vec{a}$.

Lemma 1. *(Incentive Compatibility) Under Assumption 1 on misreports of types, Algorithm 1 is incentive comptible if f_t is monotone in θ_t i.e. $f_t(\cdot, \theta_t) \geq f_t(\cdot, \theta'_t)$ if $\theta_t \succeq \theta'_t$ (for any fixed setting of $t-1$ bidders and fixed b_t).*

Proof. (Outline) Let the true valuation and type of bidder t is (π_t, θ_t). The bidder can never be worse off by reporting (π_t, θ_t) rather than some (π'_t, θ'_t) where $\pi'_t \neq \pi_t$ and/or $\theta'_t \succeq \theta_t$. This holds because the threshold price depends only on the other bidders' bids and the type of bidder t, further since f_t is monotone in θ_t, the threshold price can only increase as a result of the misreport of θ_t as θ'_t. □

Below we describe some of the existing results using the framework of Algorithm 1.

Unlimited supply auctions In this model, the seller is selling multiple units of an item to bidders who arrive one at a time and each desires one unit of the item. It is assumed that the number of units for sale is more than the number of bidders, thus effectively there is an unlimited supply of the good for sale. Blum and Hartline [3] present an expert-advice learning based algorithm for this problem under an assumption that the bidders valuations are drawn from independent identical but unknown distributions. The idea is to discretize the possible price levels and consider each price level as an expert. Given a new bidder t, an expert generates a payoff if the corresponding price level is less than the bid b_t. Their online mechanism takes the form of Algorithm 1. Since each bidder desires one unit of the item, the types of bidders are fixed, and the set $\mathcal{L}$ contains all possible subsets of bidders. At time t, the function f_t takes as input the bids till time $t-1$, and picks an expert based on the payoff generated for $t-1$ bids using an expert-learning algorithm like Follow-the-perturbed-leader. The price $\mathcal{P}_t$ is the price level advocated by the expert. The authors show a constant factor competetive revenue when compared to the best fixed price offline auction.

Single item limited supply auctions These are also known as generalized secretary problems. The simplest problem in this category is that of selecting best secretary out of candidates arriving online. In the auction setting, this is equivalent to selling single unit of an item to bidders arriving online so that the valuation of the bidder receiving the item is maximized. The more general problem is that of selling k units of an item, also known as multiple choice secretary problem. Under the assumption that the bidders arrive in a random order, the general approach for solving these problems is to use the first r bidders as a sample to generate a threshold price for the remaining bidders, usually $r = \lceil n/e \rceil$. In the multiple-choice secretary problems, the threshold may be updated as units are sold. In the framework of Algorithm 1, this is equivalent to $f_t(\cdot) = \infty$ for $t \leq r$. In the e-competetive efficiency maximizing algorithm proposed in [6] for single secretary problem, the threshold for remaining bidders is the maximum bid in the sample, i.e. $f_t(\cdot) = \max_{s\in[1,r]} b_s$ for $t > r$. In the e-competetive algorithm proposed in [1] for multiple-choice secretary problem, the threshold price for t^{th} bidder is the k^{th} maximum in $t-1$ bidders, $f_t(\cdot) = \max^k_{s\in[1,t-1]} b_s$ where $\max^k$ denotes the k^{th} maximum.

Matroid auctions Matroid auctions further generalize the secretary problems. Here, the bidders are elements of an underlying matroid of rank k, the weight of an element refers to the true valuation of the bidder, and the independent sets of the matroid form the feasible sets of bidders. Thus, the goal of finding efficiency maximizing feasible allocation is equivalent to find maximum weighted independent set of the matroid. [2] propose a sampling based algorithm for the general matroid auctions, which is similar to the mechanism discussed above for secretary problems. They set $r = \lceil n/2 \rceil$, and the threshold price function $f_t(\cdot) = \max_{s\in[1,r]} b_s/2^j$ for $t > r$, where j is a random number between 0 and $\lceil \log k \rceil$. The competetive ratio (for efficiency) of this algorithm is shown to be $\log(k)$.

In this paper, we extend the results on secretary problems to the non-matroid problem of multi-item multi-unit auctions. We provide a logarithmic bound on efficiency of our online algorithm under the same assumption of random order as in the above literature. Our algorithm also takes the form of general online auction described in Algorithm 1. We also show a matching revenue bound. Also, as a result of formulating the existing algorithms into the framework of Algorithm 1, we show that a matching revenue competetiveness result can be similarly achieved for each of them.

3 Multi-item multi-unit auctions

We consider the problem of allocating m non-identical items to n bidders, k_i units of item are i available. Each bidder j specifies a subset of m items that he is interested in, denoted by vector

$\vec{a}_j \in \{0,1\}^m$. The bidder is satisfied if *all* the items in $\vec{a}_j$ are allocated to him. A feasible set of satisfied bidders is any set A such that $\sum_{j \in A} \vec{a}_j \leq \vec{k}$.

Note that the feasible sets in this problem do not form independent sets of a matroid. For example, consider the problem with $k_i = k$. Consider a feasible set A of k bidders, each wants one unit of all the m items. Another feasible set B contains mk bidders, each of the bidder wants exactly one unit of one item. Here, $|A| < |B|$, however, none of the bidder in B can be added to A to get a feasible set of larger size. Thus, this problem is *not* a special case of matroid secretary problem.

We design an efficiency competetive online auction mechanism for this problem under the assumption that bidders arrive in random order. Let $x_t \in \{0,1\}$ denote the decision whether the goods in bundle $\vec{a}_t$ were allocated to bidder t or not. Let π_t be the valuation of bidder t for his demand. Then, under truthful bidding $b_t = \pi_t$, and the efficiency or social utility is given by $\sum_t \pi_t x_t$. The offline problem of finding the efficiency maximizing allocation is given as:

$$\begin{array}{llll} \text{(P):} & \max_{\{x_j\}} & \sum_{j=1}^n \pi_j x_j & \\ & s.t. & \sum_{j=1}^n a_{ij} x_j \leq k_i & i = 1, \ldots, m \\ & & x_j \in \{0,1\} & \forall j \end{array} \tag{1}$$

The authors in [4] gave a $\log(mU/L)$ online algorithm for this problem, assuming that $\pi_t \in [L, U]$. We design an $O(\log(mk))$ truthful online algorithm for this problem, where $k = \max_i k_i$. Our algorithm gets rid of the lower and upper bounds L, U using the enabling assumption that bids arrive in random order. In Section 3.2, we demonstrate an application of the new algorithm for designing revenue competetive online auctions.

To design an online algorithm for this problem, we use a sampling approach similar to the works of [1, 2], combined with the primal dual approach of [4]. The algorithm described below as Algorithm 2.

Algorithm 2 Online multi-item multi-unit auction

---

1. Observe $r = \lceil n/2 \rceil$ bidders without picking any bidder, and let R be the set of observed bidders (R is called the sample), and S be the set of remaining $n - r$ bidders. Let M be the sample maximum, $M = \max_{j \in R} b_j$.

2. Let vector $\vec{s}_t$ denote the remaining units of items at time t, initially $s_{r,i} = k_i$.
   At time $t = r+1, \ldots, n$,

   - Set threshold price
     $$\mathcal{P}_t = \min\{ \ p_t^T \vec{a}_t, \ \ M \ \ \}$$
     where price vector $\vec{p}_t$ is defined as $p_{t,i} = M \exp\left(-\theta_i (s_{t-1,i} - a_{t,i})/k_i\right)$, $\theta_i = \log(mk_i)$.
   - If $\mathcal{P}_t \leq b_t$ and if $\vec{a}_t \leq \vec{s}_{t-1}$, sell to bidder t at price $\mathcal{P}_t$, i.e. set $x_t = 1$ and $\vec{s}_t = \vec{s}_{t-1} - \vec{a}_t$. Otherwise, reject bidder t.

Algorithm 2 takes the same general form as Algorithm 1. Specifically, for $t \leq r$, $f_t(\cdot) = \infty$, and for $t > r$,

$$\mathcal{P}_t = f_t(\vec{s}_{t-1}, \vec{a}_t) = \min\left\{ \sum_{i=1}^m M e^{-\theta_i (s_{t-1,i} - a_{t,i})/k_i} a_{t,i}, \ \ M \right\}$$

Thus, it inherits the properties of being online, feasible and voluntary participation from Algorithm 1. Further, f_t is monotone in $\vec{a}_t$, therefore by Lemma 1, the algorithm is incentive compatible. Next, we demonstrate its competetiveness with respect to efficiency and revenue.

3.1 Efficiency

Efficiency of our online algorithm is given by $\sum_t \pi_t x_t$, where $\pi_t = b_t$ due to incentive compatibility. We prove its competetive ratio against efficiency of optimal offline algorithm. For simplicity, we assume unique bids/valuations $\{\pi_t\}$. In case of non-unique valuations, a slight modification to the algorithm and the proof is required, which is discussed in the appendix.

Theorem 1. *Let OPT be the optimal value obtained from offline allocation problem (P) (refer (1)). Assume that that the bidders arriving online are randomly ordered, and random variable $ALGO$ denotes the value of bids accepted by the online solution for a given order. Then, assuming $k_i \geq \log(mk_i)$ for all i,*

$$OPT \leq O(\log(mk))\mathbb{E}[ALGO]$$

where the expectation is taken over the random order, and $k = \max_i k_i$.

Proof. Let h denote the index of highest bid, i.e. $h = \arg\max_{j=1,\dots,n} \pi_j$ and h' denote the index of second highest bid, $h = \arg\max_{j=1,\dots,n, j\neq h} \pi_j$.

As defined in Algorithm 2, R denotes the sample set of size $r = \lceil n/2 \rceil$ and S denotes the remaining $n - r$ bidders. Let ε denote the event that h' was sampled and h was not sampled, i.e. $h' \in R$, $h \in S$. Then:

$$\Pr(\varepsilon) = \frac{r}{n} \cdot \frac{n-r}{n-1} > 1/4$$

Given non-sampled set S, define $\overline{P}(S)$ as the offline problem on the bidders in $S \backslash h$.

$$\begin{array}{llll} \overline{P}(S): & \max_{\{x_j\}} & \sum_{j\in S, j\neq h} \pi_j x_j & \\ & s.t. & \sum_{j\in S, j\neq h} \vec{a}_{ij} x_j \leq k_i & i = 1,\dots,m \\ & & x_j \in \{0,1\} & \forall j \end{array}$$

Let $\overline{OPT}(S)$ denote the optimal solution of $\overline{P}(S)$. Let $\vec{x}^*$ denote the optimal solution to the original offline problem (P), i.e. $OPT = \sum_j \pi_j x_j^*$. Since the only difference between the problems (P) and $\overline{P}(S)$ is that the latter does not consider the bidders in set R and the bidder h, therefore

$$\overline{OPT}(S) \geq \sum_{j=1}^{n} \pi_j x_j^* I(j \in S) - \pi_h$$

Then, conditional expectation:

$$\mathbb{E}[\overline{OPT}(S)|\varepsilon] \geq \sum_{j=1}^{n} \pi_j x_j^* \Pr(j \in S | h' \in R, h \in S) - \pi_h \geq \frac{1}{2}\sum_{j=1}^{n} \pi_j x_j^* - \pi_{h'} - \pi_h = \frac{OPT}{2} - \pi_{h'} - \pi_h$$

Let $ALGO(S)$ denote the value of bids accepted by the online solution given that the non-sampled set is S. In Lemma 2, we show that conditional on event ε,

$$AlGO(S) \geq \pi_h$$

And, in Lemma 3, we prove that conditional expectation $\mathbb{E}[ALGO(S)|\varepsilon]$ is within $(e\log(mk)+e)$ factor of $\mathbb{E}[\overline{OPT}(S)|\varepsilon]$. Therefore,

$$3\mathbb{E}[ALGO(S)|\varepsilon] \geq \frac{1}{(e\log(mk)+e)}\mathbb{E}[\overline{OPT}(S)|\varepsilon] + 2\pi_h \geq \frac{OPT}{2(e\log(mk)+e)}$$

Then, since the event ε has probability atleast 1/4, it follows that

$$\mathbb{E}[ALGO(S)] \geq \frac{1}{24(e\log(mk)+e)}OPT$$

□

Lemma 2. *Conditional on event ε, $ALGO(S) \geq \pi_h$.*

Proof. Conditional on event ε, $M = \pi_{h'}$. Since $\pi_h \geq M$ and $\mathcal{P}_h \leq M$, the only reason the online algorithm may not pick bidder h is because the items bidder h desires are already allocated, i.e. $s_{h-1,i} = 0$ for some i such that $a_{h,i} = 1$. However, this implies that a previously picked bidder (say at time $t < h$) paid price greater than or equal to $p_{ti} = Me^0 = \pi_{h'}$, which is a contradiction since $\pi_{h'}$ is the second highest bid price and bids are assumed to be unique[1]. Thus, we conclude that the bidder h must be picked by the online algorithm, giving $ALGO(S) \geq \pi_h$. □

Lemma 3.

$$\mathbb{E}[ALGO|\varepsilon] \geq \frac{1}{e\log(mk)+e}\mathbb{E}[\overline{OPT}|\varepsilon]$$

Proof. Consider a given instance of random set R, S conditional on event ε, that is $h' \in R, h \in S$. In this case, $M = \pi_{h'}$. The online solution

$$ALGO = \sum_{t=r+1}^{n} \pi_t x_t$$

where $x_t = 1$ if $\mathcal{P}_t \geq \pi_t$ and $\vec{s}_{t-1} \geq \vec{a}_t$.

Now, we use a primal dual approach similar to [4] to prove bounds on suboptimality of the online algorithm. Consider the relaxed offline problem $\overline{P}$ (relax the primal constraints $x_j \in \{0,1\}$ to $0 \leq x_t \leq 1$) and its dual:

$$\begin{array}{llll} \text{Primal } (\overline{P}): & \max_{\{x_j\}} & \sum_{j=r+1,j\neq h}^{n} \pi_j x_j & \\ & s.t. & \sum_{j=r+1,j\neq h}^{n} \vec{a}_{ij} x_j \leq k_i & i = 1,\ldots,m \\ & & 0 \leq x_j \leq 1 & j = r+1,\ldots,n, j \neq h \end{array}$$

$$\begin{array}{llll} \text{Dual } (\overline{P}): & \min_{\vec{p},\vec{y}} & \sum_{i=1}^{m} p_i k_i + \sum_{j=r+1,j\neq h}^{n} y_j & \\ & s.t. & \vec{p}^T \vec{a}_j + y_j \geq \pi_j & j = r+1,\ldots,n, j \neq h \\ & & y_j \geq 0 & j = r+1,\ldots,n, j \neq h \\ & & \vec{p} \geq 0 & \end{array}$$

Let $\overline{OPT}_{LP}$ denote the optimal value of the above LP relaxation of $\overline{P}$, clearly $\overline{OPT} \leq \overline{OPT}_{LP}$. We will upper bound $\overline{OPT}_{LP}$ by the value of a dual feasible solution. Constructed a dual solution as follows. Let $y_t = \pi_t - \vec{p}_t^T a_t$ if in the online solution $x_t = 1$, set $y_t = 0$ otherwise.

[1]This is the only place where we use uniqueness assumption, in appendix we show how to eliminate this assumption.

By construction $y_t x_t = y_t$, and $\vec{p}_t^T \vec{a} x_t + y_t x_t = \pi_t x_t$. Let T denote the last bidder who was accepted. In Lemma 4, we show that $(e\vec{p}_T, \{y_t\})$ forms a feasible solution for the dual problem Dual $(\overline{P})$. Now, since optimal primal value of a convex program is less than the value of any dual feasible solution:

$$\overline{OPT} \leq \overline{OPT}_{LP} \quad \leq \quad e\vec{p}_T \cdot \vec{k} + \sum_{t=r+1, t\neq h}^{n} y_t$$

For t such that $x_t = 1$, $\vec{p}_t = g(\vec{s}_t)$, where vector valued function $g(\vec{s})$ is defined as $g(\vec{s})_i = Me^{-\theta_i s_i/k_i}$ for $i = 1, \ldots, m$. Therefore,

$$\vec{p}_T = g(\vec{s}_T) = \sum_{t=r+1, x_t=1}^{n} (g(\vec{s}_t) - g(\vec{s}_t + \vec{a}_t)) + g(\vec{k})$$

Substituting above:

$$\begin{aligned}
\overline{OPT}_{LP} &\leq e \sum_{t=r+1, x_t=1}^{n} (g(\vec{s}_t) - g(\vec{s}_t + \vec{a}_t)) \cdot \vec{k} + eg(\vec{k}) \cdot \vec{k} + \sum_{t=r+1, t\neq h}^{n} y_t x_t \\
&= e \sum_{t=r+1, x_t=1}^{n} \sum_{i=1}^{m} Me^{-\theta_i s_{t,i}/k_i}(1 - e^{-\theta_i a_{t,i}/k_i})k_i + eg(\vec{k}) \cdot \vec{k} + \sum_{t=r+1, t\neq h}^{n} y_t x_t \\
\text{(since } \tfrac{\theta_i a_{it} x_t}{b_i} \leq 1\text{)} \quad &\leq e \sum_{t=r+1, x_t=1}^{n} \sum_{i=1}^{m} g(\vec{s}_t)_i (\frac{\theta_i a_{t,i}}{k_i}) k_i + eg(\vec{k}) \cdot \vec{k} + \sum_{t=r+1}^{n} y_t x_t \\
&\leq \theta e \sum_{t=r+1, x_t=1}^{n} g(\vec{s}_t) \cdot \vec{a}_t + e \sum_{i=1}^{m} \frac{M}{mk_i} k_i + \sum_{t=r+1}^{n} y_t x_t \\
&= e\theta \sum_{t=r+1}^{n} (\vec{p}_t^T \vec{a}_t x_t + y_t x_t) + eM \\
&= e\theta \sum_{t=r+1}^{n} \pi_t x_t + e\pi_{h'} \\
\text{(using Lemma 2)} \quad &\leq e\theta ALGO + eALGO
\end{aligned}$$

where $\theta = \log(mk)$. □

Lemma 4. *Let $\{x_t\}$ be the online solution, and $\vec{p}_t$ be the threshold price vector at time t. Let $y_t = \pi_t - \vec{p}_t^T \vec{a}_t$ if $x_t = 1$, and 0 otherwise. Let $T = \max\{t | x_t = 1\}$. Then, $(e\vec{p}_T, \{y_t\})$ forms a feasible solution for the dual problem Dual $(\overline{P})$.*

Proof. Consider any $t \in \{r + 1, \ldots, n\}, t \neq h$.

First consider the case $s_{T,i} > 0$ for all $i \in \vec{a}_t$, then $s_{t,i} > 0$, which means the bidder t could be rejected only because his bid was less than the threshold price. That is, at any time t, either $x_t = 1$ and $\vec{p}_t^T \vec{a}_t + y_t = \pi_t$, or $\vec{p}_t^T \vec{a} \geq \pi_t$ and $x_t = 0$. Thus, $(\vec{p}_t, y_t)$ satisfies the t^{th} constraint in dual $(\overline{P})$. Since $\vec{s}_{t-1} \geq \vec{s}_T$ for all t,

$$\vec{p}_{t,i} = \exp\left(-\frac{\theta_i(\vec{s}_{t-1,i} - \vec{a}_{t,i})}{k_i}\right) \leq \exp\left(-\frac{\theta_i(\vec{s}_{T,i} - \vec{a}_{t,i})}{k_i}\right) \leq e\vec{p}_{T,i}$$

The last inequality follows since $\theta_i \le k_i$. Therefore, $(e\vec{p}_T, y_t)$ satisfies the t^{th} dual constraint. Now, consider the case when $s_{T,i} = 0$ for some $i \in \vec{a}_t$, then

$$e\vec{p}_T \cdot \vec{a}_t + y_t x_t \ge \vec{p}_{T,i} = Me^0 = M = \pi_{h'} \ge \pi_t$$

The last inequalities follow from the fact that conditional on event ε, $M = \pi_{h'}$, and since we did not include the largest value bidder h, $\pi_t \le \pi_{h'}$.

Therefore $(e\vec{p}_T, \{y_t\})$ is a feasible dual solution to the relaxed problem $(\overline{P})$. □

3.2 Revenue

In this section, we design a revenue-competetive online algorithm under the assumption that bids are drawn independently from some known (may be non-identical) distributions, and arrive in random order. It is important to note that the distribution of bids may be non-identical. For example, in the multi-unit multi-item auction, it is unreasonable to assume that the bidders of different types $(\vec{a}_t)$ have valuations drawn from identical distributions. However, one can imagine that bidders who demand same bundle $\vec{a}_t$ have some known identical distribution on valuations.

Denote by F_t, the distribution of the bid of t^{th} bidder, and by f_t its density function. The *virtual valuation* of bidder t with valuation π_t is defined as:

$$\phi_t(\pi_t) = \pi_t - \frac{1 - F_t(\pi_t)}{f_t(\pi_t)}$$

We prove the following result in this section:

Theorem 2. *A slight modification of Algorithm 2 that uses virtual valuations $\phi(\pi_t)$ instead of π_t is $O(\log(mk))$-revenue competetive against any offline truthful mechanism for multi-item multi-unit auction.*

We use a well known result (refer [8] for details) that under independent distribution of bids, the expected revenue of any truthful mechanism is equal to its expected virtual valuation. That is, let $\{x_t^*\}$ is the allocation made by a truthful mechanism given bid vector $\vec{\pi}$, then:

$$\text{Expected Revenue} = \mathbb{E}[\sum_t \phi_t(\pi_t) x_t^*(\vec{\pi})]$$

Moreover, for single parameter domains, it is known that any truthful mechanism is monotone, i.e. there is a threshold price $\mathcal{P}_t$ such that $x_t = 1$ for $b_t \ge \mathcal{P}_t$, and $x_t = 0$ otherwise. The only truthful pricing scheme is the one that charges threshold price $\mathcal{P}_t$ from the accepted bidders. Since our online algorithm is monotone, truthful, and charges the threshold price $\mathcal{P}_t$ from the accepted bidders, its expected revenue

$$\mathbb{E}[\mathcal{P}_t(\vec{\pi}) x_t(\vec{\pi})] = \mathbb{E}[\sum_t \phi_t(\pi_t) x_t(\vec{\pi})]$$

And, using the efficiency competetiveness result derived in the previous section, for any allocation $x^*(\vec{\pi})$:

$$\mathbb{E}[\sum_t \phi_t(\pi_t) x_t(\vec{\pi})] \ge O(\log(mk)) \mathbb{E}[\sum_t \phi_t(\pi_t) x_t^*(\vec{\pi})]$$

This proves the theorem.

Remark: The main idea in proving the revenue theorem was to bound the efficiency of an alternate online auction where valuations are replaced by virtual valuations. Thus, it is important here that the competetive ratio for efficiency is independent of the bid values. Our bound of $\log(mk)$ on competetive ratio satisfies this property. The earlier competetive ratio of $\log(nU/L)$ [4] cannot be directly extended to achieve revenue competetiveness in this manner since the lower bound L may be very small (close to 0) for virtual valuations, even if only the accepted bids are considered.

4 Discussion

We surveyed the existing online auction mechanisms using a common algorithmic framework of Algorithm 1. We extended the results on online efficiency-competetive algorithms to multi-item multi-unit auctions. The competetive ratio provided by our online algorithm is independent of range or distribution of the valuations, and only assumes that bids arrive in a random order. Due to this property, we show that our algorithm can be naturally extended to achieve revenue-competetiveness using virtual valuations. Infact, it is easy to see that the discussion in Section 3.2 holds true for any mechanism of the general form described in Algorithm 1, as long as the competetive ratio for efficiency does not depend on the range or distribution of bid values. Thus, the results on secretary problems [1, 2, 5] can be extended to achieve matching revenue bounds as well.

The main open question in the area of online mechanism design seems to be that of finding an optimal mechanism for general matroids. [2] gave a $\log(k)$-factor algorithm and conjectured that there exists a constant factor efficiency competetive algorithm for this problem . For specific cases of transversal and graphic matroids, constant factor algorithms have been given by [5] and [2]. We plan to investigate the applications of the techniques introduced in this paper for the matroid problem. Another potential direction of research is to investigate more non-matroid generalizations of the multi-item multi-unit auctions, for example, the case when a bidder may demand multiple units of an item, i.e. $\vec{a}$ is any non-negative integer vector, or the case when $\vec{a}$ could even take negative values. The latter would represent the interest of bidders in both buying and selling items.

References

[1] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. A knapsack secretary problem with applications. In *APPROX '07/RANDOM '07: Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 16–28, Berlin, Heidelberg, 2007. Springer-Verlag.

[2] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA '07: Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 434–443, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.

[3] A. Blum and J. D. Hartline. Near-optimal online auctions. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1156–1163, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[4] N. Buchbinder and J. S. Naor. Improved bounds for online routing and packing via a primal-dual approach. In *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 293–304, Washington, DC, USA, 2006. IEEE Computer Society.

[5] N. B. Dimitrov and G. C. Plaxton. Competitive weighted matching in transversal matroids. In *ICALP (1)*, volume 5125 of *Lecture Notes in Computer Science*, pages 397–408. Springer, 2008.

[6] E. B. Dynkin. The optimum choice of the instant for stopping a markov process. *Soviet Mathematics Doklady*, 4:627–629, 1963.

[7] R. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 630–631, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[8] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[9] M. Pai and R. Vohra. Optimal dynamic auctions. Technical Report, 2008. http://www.kellogg.northwestern.edu/research/math/papers/1461.pdf.

A Online auction with non-unique valuations

Lemma 2 in Section 3.1 required the assumption of unique valuations/bids. Here we show how to handle non-unique bids. The only scenario that the current proof cannot handle in case of non-unique bids is when $\pi_h \neq \pi_{h'}$ and $\pi_{h'} = \pi_t$ for some $t \in S$. Intuitively, in this case many bidders may have the same valuation as the second highest bid. Since $\mathcal{P}_t \leq M = \pi_{h'}$, we must accept these bidders, and it may happen that till the time we reach the highest bid, all the items are already sold. Then, if the highest bid is much higher than all the other bids, our approximation ratio can be arbitrarily bad.

To handle such non-unique valuations and make sure that $ALGO(S) \geq \pi_h$, we modify Algorithm 2 slightly as described below in Algorithm 3. The modified algorithm works with only $\ell_i = k_i - 1$ items and leaves one unit of each item for the highest bidder. The left units are sold to a bidder t if and only if $\pi_t > M$. Since conditional event ε, $M = \pi_{h'}$, the only bidder that can have this property is the highest bidder h. Thus, the highest bidder is always selected and $ALGO(S) \geq \pi_h$. Assume $k_i \geq 2$, then as a result of reducing sellable items from k_i to $k_i - 1$, we may reduce the optimal solution by atmost a factor of 2, and Theorem 3 follows.

Theorem 3. *Let OPT be the optimal value obtained from offline allocation problem (P) (refer (1)). Assume that that the bidders arriving online are randomly ordered, and random variable $ALGO$ denotes the value of bids accepted by the online Algorithm 3 for a given order. Then, assuming $k_i \geq \log(mk_i), k_i \geq 2$ for all i,*

$$OPT \leq O(\log(mk))\mathbb{E}[ALGO]$$

where expectation is taken over the random order, and $k = \max_i k_i$.

Algorithm 3 Online multi-item multi-unit auction for non-unique bids

1. Observe $r = \lceil n/2 \rceil$ bidders without picking any bidder, and let R be the set of observed bidders (R is called the sample), and S be the set of remaining $n - r$ bidders. Let $M = \max_{j \in R} b_j$.

2. Let $\vec{s}_t$ denote the remaining units of items at time t, initially $s_{r,i} = \ell_i (= k_i - 1)$.
At time $t = r + 1, \ldots, n$,

   - Set threshold price
   $$\mathcal{P}_t = \min\{\ \ p_t^T \vec{a}_t, \quad M\ \ \}$$
   where price vector $\vec{p}_t$ is defined as $p_{t,i} = M \exp\left(-\theta_i(s_{t-1,i} - a_{t,i})/\ell_i\right)$, $\theta_i = \log(m\ell_i)$.
   - If $M < b_t$, sell to bidder t at price M, and skip this step in future.
   - Otherwise, if $\mathcal{P}_t \leq b_t$ and if $\vec{a}_t \leq \vec{s}_{t-1}$, sell to bidder t at price $\mathcal{P}_t$, i.e. set $x_t = 1$ and $\vec{s}_t = \vec{s}_{t-1} - \vec{a}_t$.
   - Otherwise, reject bidder t.