Budget Feasible Mechanisms for Dealers

Hau Chan
Department of Computer Science
Trinity University
San Antonio, TX 78212
hchan@trinity.edu

Jing Chen
Department of Computer Science
Stony Brook University
Stony Brook, NY 11794
jingchen@cs.stonybrook.edu

ABSTRACT

We consider the problem of designing budget feasible mechanisms for a dealer, who aims to maximize revenue by buying items from a seller market and selling them to a buyer market that consists of unit-demand buyers. Different from the related literature, the dealer's "value" for a set of items that he purchased from the seller market is not directly given as a number but it is defined to be the maximum revenue the dealer can obtain from selling the items to the buyers.

We aim to design mechanisms that are dominant-strategy truthful for the sellers to report their costs and envy-free for the buyers to purchase their most preferred items (given their prices) in the final outcome, such that the total payment to the sellers does not exceed the dealer's budget and the dealer's revenue is (approximately) maximized.

First, to understand the structure of the optimal mechanisms, we show that the maximum (envy-free) revenue obtainable by the dealer as a function of the set of purchased items is monotone and subadditive. Thus, existing results on subadditive optimization problems are potentially applicable in solving the mechanism design problem for the dealer.

However, a crucial assumption adopted by all previous studies on subadditive functions is that the mechanism or algorithm has access to the value oracle and/or the demand oracle. In the dealer's problem, instead, we show that (1) the demand oracle can be efficiently simulated by the value oracle and (2) both have efficient $O(\log n)$-approximation algorithms, where n is the number of buyers. This is particularly interesting given the literature, since, in general, the demand oracle can always efficiently simulate the value oracle, and there are cases where the demand oracle is strictly more powerful. Our results show that, for the dealer's problem, the two oracles are as powerful as each other.

Finally, we construct a polynomial-time budget feasible mechanism for the dealer that doesn't use any oracle and provides an $O((\log^2 n)(\log^2 m))$-approximation of the optimal revenue, where m is the number of sellers.

Keywords

budget feasible mechanism; subadditive function; envy-free pricing; revenue; implementing value and demand oracles

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.*

1. INTRODUCTION

The goal of retailers or dealers is to sell products (or items) to consumers to make as much revenue as possible. But first, the dealers must *obtain* items from other sources such as wholesalers and manufacturers. Given the obtained items, the dealers then set the *item prices* so as to maximize the revenue from selling the items to the consumers. It is clear that the dealers often have some budget constraints limiting the amount of money that they can spend in their procurements. For instance, in ancient time, western merchants sailed to Asia and Africa to buy spices to sell in their home market. Their budgets, at that time, are the amount of gold they brought with them to Asia and Africa. In modern time, although electronic transactions have greatly sped up the money flow, the dealer still has a finite amount of money that he can spend at the time of the procurement, say the credit limit of his credit card or other forms of loans he can get. Thus, the dealer's goal can be described more precisely as *to maximize his revenue subject to his budget constraint.*

In hindsight, the dealers can perform the procurement part and the pricing part independently. However, a dealer's revenue not only depends on the consumers' demands/values of all possible items that he could potentially obtain, but also on the set of items that he can obtain under his budget. Hence, in order to make optimal or nearly optimal revenue, a dealer should decide *simultaneously* (1) what items to buy from the suppliers (and how much to pay them) and (2) what are the prices of the purchased items that will maximize revenue from selling them to the consumers.

The literature of mechanism design mainly treats the buying part and the selling part separately, and a lot of progress has been made in designing procurement mechanisms for buying items and auctions for selling the items. In this paper, we are interested in synthesizing these two parts and designing mechanisms for dealers, which allow them to simultaneously decide how to buy and sell the items.

1.1 In a Nutshell: Our Setting and Our Goal

Consider a dealer with a budget B. There is a market of m sellers and m items, where the dealer can purchase item j from seller j. Each seller j has a cost c_j for his item. There is also a market of n unit-demand buyers where the dealer can sell the (purchased) items. Each buyer i has a value v_{ij} for each item j. Ideally, the dealer wants to buy a set of items such that the total cost of the items is no more than the budget, and set the item prices to maximize revenue by selling them to the consumers, where his revenue is the sum of the prices of the sold items.

The sellers and the buyers are *strategic*, but in different ways. Each seller's cost is his private information, and the dealer has to give the sellers incentives to truthfully reveal their costs. On the other hand, the buyers' values are known to the dealer, but each buyer has "free choice" and only buys the item that maximizes his utility (i.e., his value minus the item price) from the dealer—that is, the buyers must be *envy-free* [22] in the final outcome. Intuitively, the buyers are from the "home market" of the dealer, which he knows well, and the sellers are from some "alien market" that the dealer just starts to explore.

Accordingly, a dealer mechanism asks each seller to report a cost, and returns an outcome that specifies

(1) dealer's set of items (i.e., the items he buys from the sellers),

(2) dealer's payment to each seller,

(3) dealer's price for each of his items, and

(4) the allocation of dealer's items to the buyers.

Indeed, this will be the input-output structure of our mechanism M_D constructed in Section 5.

In terms of the desirable properties of our mechanisms, we are interested in mechanisms that are

(a) computationally efficient,

(b) budget feasible (i.e., the total payment made by the dealer does not exceed his budget),

(c) dominant strategy truthful (DST) for the sellers,

(d) producing an envy-free pricing and allocation for the buyers, and

(e) producing an outcome that (approximately) maximizes the dealer's revenue.

When a mechanism is randomized, we require that it is a convex combination of deterministic mechanisms that satisfy the above requirements. Indeed, our mechanism, M_D, will be such a randomized mechanism, as will be shown by Theorem 3 and Corollary 2.

Different Solution Concepts for Sellers and Buyers.

Notice that in our model, the dealer's information about the sellers and the buyers are asymmetric. Therefore, we use different solution concepts to analyze the sellers' and the buyers' incentives. Of course, one can consider a worse scenario for the dealer, where both the sellers' costs and the buyers' values are private information, and ask for mechanisms that are truthful on both sides. However, it is well known that DST mechanisms (e.g., VCG [14, 21, 32]) cannot provide any meaningful revenue guarantee in multi-item multi-buyer auctions. One can consider the Bayesian setting and assume that the buyers' values are drawn from some joint distribution; yet, even in the simplest case of unit-demand single-buyer setting (without the sellers), truthful pricing schemes that approximate the optimal revenue is not well understood (see, e.g., [12] and [8] for recent studies on this topic). Given the literature, we consider the asymmetric "home-alien" market model to be a better starting point for the study of dealer mechanisms.

1.2 Main Contribution

Towards constructing approximately optimal mechanisms for the dealer, we start by identifying important structures of the problem, as shown in Theorems 1 and 2.

Our first result is a characterization of the dealer's revenue under the optimal envy-free pricing and allocation, which is defined as a function of the purchased items.

Theorem 1 (restated) *The dealer's revenue function is monotone and subadditive.*

To the best of our knowledge, this is the first characterization on the structure of the dealer's revenue function and it places the dealer's problem in the literature of mechanism design and optimization based on subadditive functions. In particular, [4, 7, 16] have designed budget feasible mechanisms for subadditive functions, and, assuming access to value and demand oracles (as defined later), their results immediately apply to the dealer's problem.

Our second result is about a crucial assumption widely adopted in mechanism design based on subadditive functions —the existence of oracles. Indeed, since a general subadditive function needs exponentially many numbers to describe, it can only be accessed via certain oracles. All budget feasible mechanisms mentioned above rely on the value oracle and/or the demand oracle, formally defined in Section 4. However, since oracles in general do not have efficient real-life implementations[1], it is unclear how these mechanisms can be implemented. It is important to understand when one can design efficient mechanisms that do not rely on any oracle. Our second theorem shows that the demand oracle for the dealer can be approximated by the value oracle.

Theorem 2 (restated) *Given any polynomial time algorithm that is an α-approximation for the dealer's value oracle, there is a polynomial time algorithm that is an α-approximation for the dealer's demand oracle.*

This result has two important implications. First, following [22], the dealer's value oracle can be efficiently approximated within a logarithmic factor; thus we immediately have that the dealer's demand oracle can also be approximated to the same degree. It eventually leads us to a mechanism for the dealer that runs in polynomial time *without accessing any oracle.* Second, in the literature, it is well known that the demand oracle is more powerful than the value oracle: following [29], a value oracle query can always be simulated by polynomially many demand oracle queries, but there are cases where one needs exponentially many value oracle queries to simulate a single demand oracle query. Our theorem shows that, in the case of the dealer's revenue function, the value oracle is *as powerful as* the demand oracle, and any approximation algorithm for the former can be converted to an approximation algorithm for the latter with the same approximation ratio.

Finally, our third result exploits these structures of the dealer's revenue function and replaces oracle queries in budget feasible mechanisms with approximation algorithms. It is well known that many classic mechanisms (e.g., the VCG mechanism) cease to be truthful when the required computation is only approximated. Accordingly, it is unclear upfront whether existing budget feasible mechanisms are still

[1]In particular, following [22] the dealer's revenue function is APX-hard to compute, and thus both oracles are hard to compute.

truthful when the oracles are replaced by approximation algorithms. We show that this is indeed the case for the mechanism in [31]: it can be revised to use approximation algorithms for the oracles such that the resulting mechanism is still truthful. We emphasize that (1) this is a black-box replacement of the oracles, although other parts of the mechanism need to be changed accordingly; and (2) the revised mechanism works with any approximation algorithm, thus in order to design mechanisms for the dealer, one can focus on designing approximation algorithms for the value oracle. More specifically, we have the following.

Theorem 3 (restated) *Given any polynomial time algorithm that is an α-approximation for the dealer's value oracle, there is a universally truthful, budget feasible, and polynomial time dealer mechanism that gives an $O((\log^2 m)\alpha^2)$-approximation to the optimal revenue.*

Following [22], there is an $O(\log n)$-approximation for the dealer's value oracle, thus we immediately have:

Corollary 2 (restated) *There is a universally truthful, budget feasible, and polynomial-time mechanism that gives an $O((\log^2 m)(\log^2 n))$-approximation to the optimal revenue.*

1.3 Open Problems

It is natural to generalize our model and consider the case where each seller has multiple units of his item for sale (as in [7]). It is easy to see that our results on the monotonicity and subadditivity of the revenue function (Theorem 1) and the approximability of the demand oracle (Theorem 2) hold for the multiple-unit settings. However, we do not know how to generalize our dealer mechanism to the multi-unit settings, and this would be an interesting open problem for future study.

In terms of approximating oracle queries, it is not hard to see that the mechanisms in [4, 31] for single-unit budget feasible procurement and those in [7] for multi-unit procurement remain truthful if the demand oracle is replaced by an ϵ-approximation where ϵ is a small constant. However, it is unknown whether such an approximation algorithm exists for the dealer's revenue function. Also, it is unclear whether these mechanisms are still truthful when the approximation ratio increases (to, say, $\log n$).

In most studies on mechanism design with budgets, the budget is treated as a hard constraint. However, it is conceivable that, in real life, obtaining additional liquidity may be feasible, although at an increasing marginal cost. Modeling and handling such a flexible liquidity will be a very interesting problem for future study.

Moreover, currently we do not have a matching lower bound for our mechanism's approximation ratio. However, even for budget feasible procurements with subadditive valuations, only poly-logarithmic approximations are known. It would be interesting to show matching lower bounds for both the dealer's problem and the procurement problem. Finally, generalizing our model to allow multiple dealers to compete in a market is another important direction worth investigating.

1.4 Additional Prior Work

There is a large amount of literature on the problem of finding the optimal envy-free pricing and allocation that maximizes revenue. The problem was first introduced and considered by [22] where the authors showed that the problem is APX-hard and there is an $O(\log n)$-approximation algorithm for it. Since then, many works have studied different variants of the problem. For example, when there is a single buyer whose values for the items are drawn from independent distributions [8, 12, 15], when the buyers' values have special structures [3, 9, 10], when the buyers have budgets [5, 19, 24], and when the items can only be sold in some special ways [6, 13, 17, 18, 20, 23]. In particular, [12] showed that the pricing problem with a single buyer is NP-complete even when the value distributions all have support size 3. Moreover, [8] showed that when the buyer's value distributions are either discrete or continuous with bounded supports (with proper oracle accesses), and when the distributions satisfy the monotone hazard rate (MHR) condition, there is an algorithm that is a constant approximation for the problem and runs in polynomial time with respect to the size of the largest support (when the distribution is discrete) or the ratio between the upper bound and the lower bound of the supports (when the distribution is continuous).

Budget feasible procurement mechanisms were first introduced by [31]. [31] and [11] provide constant approximation mechanisms for submodular functions, and [16] and [4] respectively provide $O(\log^2 m)$ and $O(\frac{\log^2 m}{\log\log m})$ approximation mechanisms for subadditive functions. All of them study the setting where each seller has a single item. Recently, [7] generalizes the setting to where each seller has multiple units of his item, and gives an $O(\log m)$ approximation for additive functions and an $O(\frac{\log^2 m}{\log\log m})$ approximation for subadditive functions. In [1], the authors study budget feasible mechanisms for additive functions, when each seller's cost is much smaller than the budget.

Another related line of research is double auctions [27, 28, 29]. In a double auction, buyers and sellers submit their private values and costs, respectively, to an intermediary. The goal of the latter is to construct mechanisms that are individually rational and truthful. While a double auction also deals with a two-sided market, the setting and goal are different from ours. Indeed, the goal there is to maximize social welfare and clear the market, instead of maximizing the intermediary's revenue. Also, double auctions treat sellers and buyers symmetrically, using the same solution concept; while we consider an asymmetric setting. Moreover, the intermediary in double auctions has no budget constraint.

Finally, a recent work [30] considers a dealer who has no budget constraint and wants to maximize his revenue minus the total payment to the sellers. However, the settings they consider are highly constrained compared with ours.

2. MODEL: DEFINITIONS AND NOTATIONS

Consider a *dealer* who is buying items from one market and selling them to another market. The market where he is buying is an "alien" market to him, that is, he does not know the sellers' true costs for the items that they sell on that market. The market where he is selling is his "home" market, that is, he knows the buyers' values for the items in that market. The dealer has a budget $B > 0$ specifying the most amount of money he can spend in the alien market.

More specifically, we have the following model. There are m sellers, indexed by $j \in [m] = \{1, 2, ..., m\}$. Each seller $j \in [m]$ has one unit of item j for sale, whose true cost is $c_j > 0$. The cost c_j is the private information of seller j. If the dealer buys item j from seller j with a payment q_j, then the *utility* of seller j, u_j^s, is $q_j - c_j$; and u_j^s is 0 if the dealer does not buy his item and does not pay him.

There are n buyers, indexed by $i \in [n] = \{1, 2, ..., n\}$. Each buyer i has a value v_{ij} for each item $j \in [m]$. Let $V = (v_{ij})_{i\in[n],j\in[m]}$ be the valuation matrix. The dealer knows the buyers' values. Let $\perp$ denote a dummy item such that $v_{i\perp} = 0$ for each $i \in [n]$. When the set of items bought by the dealer is $S \subseteq [m]$, an *assignment* for the buyers consists of a pair (p_S, A_S) where $p_S = (p_j)_{j\in S}$ is the *pricing* vector and $A_S : [n] \to S \cup \{\perp\}$ is the *allocation*, with each item $j \in S$ being allocated to at most one buyer (that is, $|A_S^{-1}(j)| \leq 1$ for each $j \in S$). The price of the dummy item is always $p_\perp = 0$. The *utility* of buyer i, u_i^b, is $v_{iA_S(i)} - p_{A_S(i)}$.

An assignment (p_S, A_S) is *envy-free* if $v_{iA_S(i)} - p_{A_S(i)} \geq v_{ij} - p_j$ for each $i \in [n]$ and $j \in S \cup \{\perp\}$. The *revenue* of an assignment is $Rev(p_S, A_S) = \sum_{i=1}^n p_{A_S(i)}$.

For each $S \subseteq [m]$, the *optimal assignment under* S is an envy-free assignment $(\hat{p}_S, \hat{A}_S)$ such that $(\hat{p}_S, \hat{A}_S) \in \text{argmax}_{(p'_S, A'_S) \text{ is envy-free}} Rev(p'_S, A'_S)$. The dealer's *revenue from* S is the maximum revenue that he can get from any envy-free assignment under S, that is, $R(S) = Rev(\hat{p}_S, \hat{A}_S)$. The *optimal procurement* of the dealer is a subset of items S^* such that $S^* \in \text{argmax}_{S\subseteq[m],\, \sum_{j\in S} c_j \leq B} R(S)$. Notice that $R(S^*) = Rev(\hat{p}_{S^*}, \hat{A}_{S^*})$. Moreover, we will refer R to the *revenue function* of the dealer.

Remark 1. *We have two revenue functions: Rev is the revenue of a particular assignment, and R is the maximum revenue among all assignments for a given set of items.*

The solution concept.

A *dealer mechanism* M asks each seller j to report a cost $c'_j \geq 0$. Given the cost profile $c' = (c'_1, ..., c'_m)$, the mechanism outputs a subset of items $S \subseteq [m]$, a payment profile $q = (q_1, ..., q_m)$, and an envy-free assignment (p_S, A_S), such that $q_j = 0$ for all $j \in [m]\backslash S$. That is $M(c') = (S, q, p_S, A_S)$. The utility of each seller j under the strategy profile c' is $u_j^s(c') = q_j - c_j$ if $j \in S$, and $u_j^s(c') = 0$ otherwise.

A deterministic mechanism M is *truthful* if for each seller j, c'_j, and c'_{-j}, $u_j^s(c_j, c'_{-j}) \geq u_j^s(c'_j, c'_{-j})$. Let $M(c) = (S, q, p_S, A_S)$. Mechanism M is *individually rational* if for each seller j, $u_j^s(c) \geq 0$; and M is *budget feasible* if the dealer's total payment is no more than the budget, that is, $\sum_{j=1}^m q_j \leq B$. A randomized mechanism is *universally truthful* (respectively, *individually rational* and *budget feasible*) if it is a convex combination of deterministic truthful mechanisms (respectively, *individually rational*, and *budget feasible*).

Definition 1. *Let $f(n) > 0$. A universally truthful dealer mechanism M is an $f(n)$-approximation to the value of the optimal procurement if the mechanism is individually rational and budget feasible, and the revenue under the true cost profile c, $Rev(p_S, A_S)$, has expected value at least $\frac{R(S^*)}{f(n)}$.*

3. PROPERTIES OF THE DEALER'S REVENUE FUNCTION

The structure of the dealer's revenue function R has never been studied. In this section, we show, for the first time, that R is subadditive, that is, $R(S \cup T) \leq R(S) + R(T)$ for any subsets of items S and T. Thus, existing techniques on subadditive optimization could potentially be applied to compute/approximate it. We first show the following lemma.

Lemma 1. *The dealer's revenue function R is monotone. That is, for any two subsets S and T such that $S \subseteq T \subseteq [m]$, $R(S) \leq R(T)$.*

Proof. Let $(\hat{p}_S, \hat{A}_S)$ be the optimal assignment under the set S. We define an assignment (p_T, A_T) for T as follows. Let $p_j = \hat{p}_j$ for $j \in S$, $p_j = \infty$ for $j \in T \setminus S$ and $A_T = A_S$. For any $i \in [n]$ and $j \in S \cup \{\perp\}$, we have $u_i^b(p_T, A_T) = u_i^b(\hat{p}_S, \hat{A}_S) = v_{i\hat{A}_S(i)} - \hat{p}_{\hat{A}_S(i)} \geq v_{ij} - \hat{p}_j = v_{ij} - p_j$, where the inequality is because of the envy-freeness of $(\hat{p}_S, \hat{A}_S)$. For any $i \in [n]$ and $j \in T \setminus S$, $u_i^b(p_T, A_T) \geq v_{ij} - p_j$ as $p_j = \infty$. Thus the assignment (p_T, A_T) is envy-free. Moreover, by definition, $R(S) = Rev(\hat{p}_S, \hat{A}_S) = \sum_{i=1}^n \hat{p}_{\hat{A}_S(i)} = \sum_{i=1}^n \hat{p}_{A_T(i)} = \sum_{i=1}^n p_{A_T(i)} = Rev(p_T, A_T)$. Since $Rev(p_T, A_T) \leq R(T)$ by definition, we have $R(S) \leq R(T)$. $\square$

We have the following theorem.

Theorem 1. *The function R is subadditive.*

Before proving Theorem 1, we recall the following notations, definitions, and results that we will use. Given a valuation matrix V, when the set of available items is S, we let V_S be the sub-matrix of V that contains only the items in S. Given the set of available items S, and a vector $r = (r_j)_{j\in S}$ of non-negative real numbers, a *Walrasian equilibrium with reserve prices* r [22] is an envy-free assignment (p_S, A_S) such that (1) $p_j \geq r_j$ for all $j \in S$, (2) if item j is not assigned to any buyer, then $p_j = r_j$, and (3) if item j is a most-preferred item of buyer i and j is not assigned, then buyer i is assigned an item. In [22], the authors constructed a polynomial time algorithm, MaxWEQ_r, that computes a Walrasian equilibrium with reserve prices, see the following lemma.

Lemma 2. *[22] Given any inputs S, V_S, and r, the algorithm $MaxWEQ_r$ outputs a Walrasian equilibrium (p_S, A_S) with reserve prices r in polynomial time.*

Below we prove a key step in proving Theorem 1.

Lemma 3. *For any set of items $S \subseteq [m]$, any envy-free assignment (p_S, A_S), and any subset $T \subseteq S \cap A_S([n])$, we have $\sum_{j\in T} p_j \leq R(T)$.*

Proof. Given (p_S, A_S), we construct an envy-free assignment under T, whose revenue is at least $\sum_{j\in T} p_j$. To do so, we run the following algorithm Alg_{Assign} with inputs S, V_S, (p_S, A_S), and T.

Notice that for each $j \in T$, the price output by MaxWEQ_r in step 1 is $\bar{p}_j \geq r_j = p_j$, by Lemma 2. However, under $(\bar{p}_T, \bar{A}_T)$, some items in T might not be assigned to any buyer. Therefore, $Rev(\bar{p}_T, \bar{A}_T)$ may be strictly less than $\sum_{j\in T} \bar{p}_j$. That is why we augment $\bar{A}_T$ to A_T^* which assigns all items in T. More precisely, we have the following claim.

Claim 1. Alg_{Assign} *returns an envy-free assignment $(\bar{p}_T, A_T^*)$ that assigns all items in T.*

Proof. Notice that, when the algorithm terminates, A_T^* allocates all items in T. Thus, we need to show that the algorithm will not loop forever in step 3. To see why this

Algorithm 1: Alg_{Assign} - An Algorithm to Produce an Envy-Free Assignment where All Items Are Sold

Input : A set of items S, a valuation matrix V_S, an envy-free assignment (p_S, A_S), and a set $T \subseteq S \cap A_S([n])$.

Output: An envy-free assignment $(\bar{p}_T, A^*_T)$

1 Run MaxWEQ_r on inputs T, V_T, and reserve prices $r = (p_j)_{j \in T}$ and denote the outcome by $(\bar{p}_T, \bar{A}_T)$.
2 Initialize $A^*_T = \bar{A}_T$.
3 **while** *there exists an unassigned item j in T according to A^*_T* **do**
4 Let $i \in [n]$ be such that $A_S(i) = j$, and set $A^*_T(i) = j$.
5 **end while**
6 **return** $(\bar{p}_T, A^*_T)$

is true, notice that for each buyer $i \in [n]$, there is at most one item $j \in T$ such that $A_S(i) = j$. Accordingly, buyer i gets reassigned in A^*_T at step 3 for at most once. Since in each execution of step 3, one buyer gets reassigned, and since there are only n buyers, step 3 is executed at most n times and the algorithm will not loop forever.

Next, we show that after each execution of step 3, (1) $(\bar{p}_T, A^*_T)$ is envy-free and (2) for each item j that is not assigned by the current A^*_T, $\bar{p}_j = p_j$. To see why this is true, notice that the initial $(\bar{p}_T, A^*_T)$ satisfies (1) and (2) by the definition of Walrasian equilibrium with reserve prices. Accordingly, if item j is not assigned and $A_S(i) = j$, we have $v_{iA^*_T(i)} - \bar{p}_{A^*_T(i)} \geq v_{ij} - \bar{p}_j = v_{ij} - p_j$ and $v_{ij} - p_j \geq v_{iA^*_T(i)} - p_{A^*_T(i)}$, where the first inequality is because (1) and (2), and the second is because (p_S, A_S) is envy-free. Since $\bar{p}_{A^*_T(i)} \geq r_{A^*_T(i)} \geq p_{A^*_T(i)}$, we have $v_{ij} - p_j \geq v_{iA^*_T(i)} - p_{A^*_T(i)} \geq v_{iA^*_T(i)} - \bar{p}_{A^*_T(i)} \geq v_{ij} - \bar{p}_j = v_{ij} - p_j$, which implies $v_{ij} - \bar{p}_j = v_{ij} - p_j = v_{iA^*_T(i)} - p_{A^*_T(i)} = v_{iA^*_T(i)} - \bar{p}_{A^*_T(i)}$. Thus, item j maximizes i's utility under $\bar{p}_T$ and $p_{A^*_T(i)} = \bar{p}_{A^*_T(i)}$.

Notice that after the execution of step 3, buyer i is the only buyer whose assigned item changes, and the original item $A^*_T(i)$ is the only item that changed from assigned to unassigned. Accordingly, the new assignment $(\bar{p}_T, A^*_T)$ is still envy-free and any unassigned item $j' \in T$ has price $\bar{p}_{j'} = p_{j'}$. That is, $(\bar{p}_T, A^*_T)$ still satisfies (1) and (2) after the first execution of step 3. By induction, $(\bar{p}_T, A^*_T)$ satisfies (1) and (2) after each execution of step 3, as desired.

In sum, the final assignment $(\bar{p}_T, A^*_T)$ returned by Alg_{Assign} is envy-free and assigns all items in T, and Claim 1 holds. $\square$

By the construction of Alg_{Assign} and Claim 1, $Rev(\bar{p}_T, A^*_T) = \sum_{j \in T} \bar{p}_j \geq \sum_{j \in T} r_j = \sum_{j \in T} p_j$, which implies $\sum_{j \in T} p_j \leq R(T)$. Thus, Lemma 3 holds.

Now we are ready to prove Theorem 1.

Proof of Theorem 1. For any subsets of items S and T, let $(\hat{p}_{S \cup T}, \hat{A}_{S \cup T})$ be the optimal assignment under $S \cup T$ and let $W = \hat{A}_{S \cup T}([n])$ be the set of assigned items under $\hat{A}_{S \cup T}$. We have $R(S \cup T) = Rev(\hat{p}_{S \cup T}, \hat{A}_{S \cup T}) = \sum_{i=1}^{n} \hat{p}_{A_{S \cup T}(i)} = \sum_{j \in W} \hat{p}_j \leq \sum_{j \in S \cap W} \hat{p}_j + \sum_{j \in T \cap W} \hat{p}_j$, where the first two equalities are by the definitions of functions R and Rev; the third equality is because the total revenue collected from the players is the same as the total revenue obtained by selling the items; and the inequality is because $W \setminus \{\perp\} \subseteq (S \cap W) \cup (T \cap W)$.

Since $\sum_{j \in S \cap W} \hat{p}_j \leq R(S \cap W)$ by Lemma 3 and $R(S \cap W) \leq R(S)$ by Lemma 1, we have $\sum_{j \in S \cap W} \hat{p}_j \leq R(S)$. Similarly, $\sum_{j \in T \cap W} \hat{p}_j \leq R(T)$. Therefore, we have $R(S \cup T) \leq R(S) + R(T)$, and Theorem 1 holds. $\square$

Remark 2. *A function $f : 2^{[m]} \rightarrow \mathbb{R}$ is submodular if for any subsets of items S and T, $f(S \cup T) + f(S \cap T) \leq f(S) + f(T)$. It is unknown whether R is submodular or not.*

4. APPROXIMATING THE DEMAND ORACLE FOR THE DEALER

Mechanism design for the dealer requires optimization over his revenue function. In general, optimization over subadditive functions is NP-hard [25, 26] [2], and oracles of various forms are used to help solving the mechanism design problem. In particular, two oracles have been widely considered — the *value* oracle and the *demand* oracle [4, 7, 11, 16, 29, 31]. In our setting, the value oracle takes as input a subset $S \subseteq [m]$ and returns $R(S)$; and the demand oracle takes as input a cost vector $\bar{c} = (\bar{c}_1, ..., \bar{c}_m)$ and returns $S^*_{\bar{c}} \in \text{argmax}_{T \subseteq [m]} R(T) - \sum_{j \in T} \bar{c}_j$. For any subset $T \subseteq [m]$, the quantity $R(T) - \sum_{j \in T} \bar{c}_j$ is referred to as the *net revenue* of T given $\bar{c}$, denoted by $NR(T, \bar{c})$.

Given access to the value and demand oracles, the mechanism design problem for the dealer can be solved under the framework of budget feasible mechanisms for subadditive valuation functions [4, 7, 16]. However, how can a polynomial-time dealer answer oracle queries? There is little literature about oracle implementation, and most mechanisms over subadditive functions rely on oracle queries and thus do not have polynomial-time implementations.

As part of our main contribution, we show that the demand oracle for the dealer can be efficiently approximated. Using this approximation, in Section 5, we provide a polynomial time mechanism for the dealer *without oracle assesses.*

Our approximation algorithm of the demand oracle uses an approximation algorithm of the value oracle. More specifically, we say that an algorithm Alg_{value} is an α-approximation of the value oracle with $\alpha \geq 1$ if it takes as inputs a set S and a valuation matrix V_S, and returns an envy-free assignment (p_S, A_S) such that $Rev(p_S, A_S) \geq \frac{R(S)}{\alpha}$. We say that an algorithm Alg_{demand} is an α-approximation of the demand oracle with $\alpha \geq 1$ if it takes as inputs a cost profile $\bar{c}$ and a valuation matrix V, and returns a set $S_{\bar{c}}$ and an envy-free assignment $(p_{S_{\bar{c}}}, A_{S_{\bar{c}}})$ such that $Rev(p_{S_{\bar{c}}}, A_{S_{\bar{c}}}) - \sum_{j \in S_{\bar{c}}} \bar{c}_j \geq \frac{NR(S^*_{\bar{c}}, \bar{c})}{\alpha}$. We have the following theorem.

Theorem 2. *Given any polynomial time algorithm that is an α-approximation for the value oracle, there is a polynomial time algorithm that is an α-approximation for the demand oracle.*

As shown in [22], the value oracle is APX-hard to compute, and there is a polynomial-time $O(\log n)$-approximation for it. Thus, we immediately have the following corollary.

Corollary 1. *There exists a polynomial-time $O(\log n)$-approximation algorithm for the demand oracle.*

[2]Since submodular functions belong in the class of subadditive functions, and optimization over submodular is NP-hard in general, we have that optimization over subadditive functions is also NP-hard.

Interestingly, the demand oracle is in general more powerful and computationally harder than the value oracle. In [29], the authors show that a value query can be simulated by polynomially many demand queries, but an exponential number of value queries may be required for simulating a single demand query. Theorem 2 shows that this is not the case for the mechanism design problem for the dealer: here, the two oracles are as powerful as each other. The remaining part of this section is devoted to proving Theorem 2.

Let Alg_{value} be an arbitrary polynomial time algorithm that is an α-approximation of the value oracle, with $\alpha \geq 1$. Our algorithm Alg_{demand} is defined as follows.

Algorithm 2: Alg_{demand} - An Algorithm to Approximate the Demand Oracle

Input : A valuation matrix V and a cost profile $\bar{c}$.
Output: A set S and an envy-free assignment $(\bar{p}_S, A_S)$.

1 Let V' be a valuation matrix such that $v'_{ij} = v_{ij} - \bar{c}_j$ for each $i \in [n]$ and $j \in [m]$.
2 Run Alg_{value} with inputs $[m]$ and V', and let (p, A) be the output.
3 Let $S = A([n]) \setminus \{\perp\}$ and $\tilde{A}_S = A$.
4 Let $\tilde{p}_S = (\tilde{p}_j)_{j \in S}$ where $\tilde{p}_j = p_j + \bar{c}_j$ for $j \in S$.
5 **return** S and $(\tilde{p}_S, \tilde{A}_S)$.

Below, we show that Alg_{demand} is an α-approximation of the demand oracle. Because we deal with two valuation matrices V and V', we use $f(\cdot; V')$ to denote a function f evaluated when the valuation matrix is V'. If no valuation matrix is specified, then the computation is always with respect to V. Recall that $S^*_{\bar{c}}$ is the output of the demand oracle given inputs V and $\bar{c}$. Let $(\hat{p}_{S^*_{\bar{c}}}, \hat{A}_{S^*_{\bar{c}}})$ be the envy-free assignment that gives the highest net revenue with respect to $S^*_{\bar{c}}$. We want to show that $(\hat{p}_{S^*_{\bar{c}}}, \hat{A}_{S^*_{\bar{c}}})$ is also an (not necessarily optimal) envy-free assignment for the buyers whose values for the items are according to the valuation matrix V' when the set of available items is $[m]$. We begin by showing some properties of $(\hat{p}_{S^*_{\bar{c}}}, \hat{A}_{S^*_{\bar{c}}})$.

LEMMA 4. *$\hat{A}_{S^*_{\bar{c}}}([n]) = S^*_{\bar{c}}$ or $\hat{A}_{S^*_{\bar{c}}}([n]) = S^*_{\bar{c}} \cup \{\perp\}$.*

PROOF. For ease of notations, let $S = S^*_{\bar{c}}$, $\hat{p}_S = \hat{p}_{S^*_{\bar{c}}}$, and $\hat{A}_S = \hat{A}_{S^*_{\bar{c}}}$. First notice that all of the available items are in S. Since $\hat{A}_S$ does not assign buyers to any item that is not in S, $\hat{A}_S([n]) \subseteq S \cup \{\perp\}$. Let $U = S \cup \{\perp\} - \hat{A}_S([n])$ be the set of unassigned items in S. To prove the lemma, it is suffice to show that U is either empty or contains only $\perp$. For the sake of contradiction, suppose U contains some item that is not $\perp$. The net revenue of the set S is

$$NR(S, \bar{c}) = R(S) - \sum_{j \in S} \bar{c}_j = \sum_{i=1}^{n} \hat{p}_{\hat{A}_S(i)} - \sum_{j \in S} \bar{c}_j \geq NR(\bar{S}, \bar{c}), \quad (1)$$

for any $\bar{S} \subseteq [m]$. Letting $\bar{S} = S - U$, $\bar{p}_{\bar{S}} = (\bar{p}_i)_{i \in \bar{S}}$ such that $\bar{p}_i = \hat{p}_i$ for $i \in \bar{S}$, and $\bar{A}_{\bar{S}} = \hat{A}_S$. It is easy to see that the assignment $(\bar{p}_{\bar{S}}, \bar{A}_{\bar{S}})$ is envy-free when the set of available items is $\bar{S}$, that is, $v_{i\bar{A}_{\bar{S}}(i)} - p_{\bar{A}_{\bar{S}}(i)} = v_{i\hat{A}_S(i)} - \hat{p}_{\hat{A}_S(i)} \geq v_{ij} - \hat{p}_j = v_{ij} - \bar{p}_j$ for all buyers $i \in [n]$ and all items $j \in \bar{S}$. It follows that $NR(\bar{S}, \bar{c}) = R(\bar{S}) - \sum_{j \in \bar{S}} \bar{c}_j \geq Rev(\bar{p}_{\bar{S}}, \bar{A}_{\bar{S}}) - \sum_{j \in \bar{S}} \bar{c}_j = \sum_{i=1}^{n} \bar{p}_{\bar{A}_{\bar{S}}(i)} - \sum_{j \in \bar{S}} \bar{c}_j = \sum_{i=1}^{n} \hat{p}_{\hat{A}_S(i)} - \sum_{j \in \bar{S}} \bar{c}_j > \sum_{i=1}^{n} \hat{p}_{\hat{A}_S(i)} - \sum_{j \in \bar{S} \cup U} \bar{c}_j = \sum_{i=1}^{n} \hat{p}_{\hat{A}_S(i)} - \sum_{j \in S} \bar{c}_j = NR(S, \bar{c})$, which is a contradiction to Equation 1. □

The above proof relies on the fact that the costs of the items are strictly greater than zero. If we allow the items' costs to be zero and if $\hat{A}_{S^*_{\bar{c}}}$ has some unallocated items with cost zero, then removing those items, we obtain another envy-free assignment with the (same) optimal net revenue that satisfies Lemma 4. Next, we show that the price of each item in $S^*_{\bar{c}}$ is at least its cost.

LEMMA 5. *For any $j \in S^*_{\bar{c}}$, $\hat{p}_j - \bar{c}_j \geq 0$.*

PROOF. Again for ease of notations, let $S = S^*_{\bar{c}}$, $\hat{p}_S = \hat{p}_{S^*_{\bar{c}}}$, and $\hat{A}_S = \hat{A}_{S^*_{\bar{c}}}$. Let $N = \{j \in S | \hat{p}_j - \bar{c}_j < 0\}$ be the set of items with prices less than costs. Assume $N \neq \emptyset$ and let $\bar{S} = S - N$. We show that $NR(\bar{S}, \bar{c}) > NR(S, \bar{c})$, which contradicts the optimality of S.

To do so, we run Alg_{Assign} with inputs S, V_S, $(\hat{p}_S, \hat{A}_S)$, and $\bar{S} = S - N$. Let $(\bar{p}_{\bar{S}}, \bar{A}_{\bar{S}})$ be the output. We have $NR(\bar{S}, \bar{c}) = R(\bar{S}) - \sum_{j \in \bar{S}} \bar{c}_j \geq Rev(\bar{p}_{\bar{S}}, \bar{A}_{\bar{S}}) - \sum_{j \in \bar{S}} \bar{c}_j = \sum_{i=1}^{n} \bar{p}_{\bar{A}_{\bar{S}}(i)} - \sum_{j \in \bar{S}} \bar{c}_j = \sum_{j \in \bar{S}} \bar{p}_j - \bar{c}_j > \sum_{j \in \bar{S}} \bar{p}_j - \bar{c}_j + \sum_{j \in N} \hat{p}_j - \bar{c}_j \geq \sum_{j \in \bar{S}} \hat{p}_j - \bar{c}_j + \sum_{j \in N} \hat{p}_j - \bar{c}_j = \sum_{j \in S} \hat{p}_j - \bar{c}_j = \sum_{i=1}^{n} \hat{p}_{\hat{A}_S(i)} - \sum_{j \in S} \bar{c}_j = NR(S, \bar{c})$, where the first inequality is because $R(\bar{S}) \geq Rev(\bar{p}_{\bar{S}}, \bar{A}_{\bar{S}})$, the third equality is because all items in $\bar{S}$ are sold according to $\bar{A}_{\bar{S}}$ (Claim 1), the second inequality is because $\hat{p}_j - \bar{c}_j < 0 \, \forall j \in N$, the third inequality is because $\bar{p}_j \geq \hat{p}_j \, \forall j \in \bar{S}$ (by the construction of Alg_{Assign}), and the fifth equality is because all items in S are sold according to $\hat{A}_S$ (Lemma 4). □

Next, we show that there is an assignment which is envy-free with respect to the valuation matrix V' and has revenue equals to $NR(S^*_{\bar{c}}, \bar{c})$. More precisely, let $\bar{p} = (\bar{p}_j)_{j \in S^*_{\bar{c}}}$ be such that $\bar{p}_j = \hat{p}_j - \bar{c}_j$ for all $j \in S^*_{\bar{c}}$ and $\bar{p}_j = \infty$ for all $j \in [m] \setminus S^*_{\bar{c}}$, and let $\bar{A} = \hat{A}_{S^*_{\bar{c}}}$. By Lemma 5, $\bar{p}_j \geq 0$ for each $j \in [m]$ and $\bar{p}$ is a well-defined price profile. We have the following two lemmas.

LEMMA 6. *$Rev(\bar{p}, \bar{A}; V') = NR(S^*_{\bar{c}}, \bar{c})$.*

PROOF. $Rev(\bar{p}, \bar{A}; V') = \sum_{i=1}^{n} \bar{p}_{\bar{A}(i)} = \sum_{i=1}^{n} \hat{p}_{\bar{A}(i)} - \bar{c}_{\bar{A}(i)} = \sum_{i=1}^{n} \hat{p}_{\hat{A}_{S^*_{\bar{c}}}(i)} - \bar{c}_{\hat{A}_{S^*_{\bar{c}}}(i)} = \sum_{i=1}^{n} \hat{p}_{\hat{A}_{S^*_{\bar{c}}}(i)} - \sum_{j \in S^*_{\bar{c}}} \bar{c}_j = NR(S^*_{\bar{c}}, \bar{c})$, where the first, second, and last equalities are by definition, the third is because $\bar{A} = \hat{A}_{S^*_{\bar{c}}}$, and the fourth is because all items in $S^*_{\bar{c}}$ are sold according to $\hat{A}_{S^*_{\bar{c}}}$ (by Lemma 4). □

LEMMA 7. *$(\bar{p}, \bar{A})$ is envy-free with respect to V'.*

PROOF. Recall the value of buyer i for item j with respect to V' is $v'_{ij} = v_{ij} - \bar{c}_j$. First, notice that for each buyer i, $u^b_i(\bar{p}, \bar{A}; V') = (v_{i\bar{A}(i)} - \bar{c}_{\bar{A}(i)}) - \bar{p}_{\bar{A}(i)} = (v_{i\bar{A}(i)} - \bar{c}_{\bar{A}(i)}) - (\hat{p}_{\bar{A}(i)} - \bar{c}_{\bar{A}(i)}) = v_{i\bar{A}(i)} - \hat{p}_{\bar{A}(i)} = v_{i\hat{A}_{S^*_{\bar{c}}}(i)} - \hat{p}_{\hat{A}_{S^*_{\bar{c}}}(i)} \geq 0$, where the last inequality is because $(\hat{p}_{S^*_{\bar{c}}}, \hat{A}_{S^*_{\bar{c}}})$ is envy-free under valuation matrix V. (Surely, since $\bar{p}_j \geq 0$ for each $j \in [m]$, we have $v_{i\bar{A}(i)} - \bar{c}_{\bar{A}(i)} \geq 0$ for each buyer $i \in [n]$.)

Moreover, for each buyer $i \in [n]$ and each item $j \in [m]$, $u^b_i(\bar{p}, \bar{A}; V') = v_{i\hat{A}_{S^*_{\bar{c}}}(i)} - \hat{p}_{\hat{A}_{S^*_{\bar{c}}}(i)} \geq v_{ij} - \hat{p}_j = (v_{ij} - \bar{c}_j) - (\hat{p}_j - \bar{c}_i) = (v_{ij} - \bar{c}_j) - \bar{p}_j = v'_{ij} - \bar{p}_j$, where the inequality is again because $(\hat{p}_{S^*_{\bar{c}}}, \hat{A}_{S^*_{\bar{c}}})$ is envy-free under V.

In sum, $(\bar{p}, \bar{A})$ is envy-free with respect to V'. □

Similarly, we have the following lemma for $(\tilde{p}_S, \tilde{A}_S)$, the assignment output by Alg_{demand}. Recall that (p, A) is the assignment output by Alg_{value} with inputs $[m]$ and V'.

Lemma 8. *$Rev(\tilde{p}_S, \tilde{A}_S) - \sum_{j\in S} \bar{c}_j = Rev(p, A; V')$, and $(\tilde{p}_S, \tilde{A}_S)$ is envy-free with respect to V.*

Proof. First of all, $Rev(\tilde{p}_S, \tilde{A}_S) - \sum_{j\in S} \bar{c}_j = \sum_{i=1}^n \tilde{p}_{\tilde{A}_S(i)} - \sum_{j\in S} \bar{c}_j = \sum_{i=1}^n (p_{\tilde{A}_S(i)} + \bar{c}_{\tilde{A}_S(i)}) - \sum_{j\in S} \bar{c}_j = \sum_{i=1}^n p_{A(i)} + \sum_{i=1}^n \bar{c}_{A(i)} - \sum_{j\in S} \bar{c}_j = \sum_{i=1}^n p_{A(i)} = Rev(p, A; V')$, where the fourth equality is because $S = A([n]) \setminus \{\perp\}$ and all the others are by definition.

Second, to show $(\tilde{p}_S, \tilde{A}_S)$ is envy-free, notice that for each buyer $i \in [n]$ and for each item $j \in S$, $u_i^b(\tilde{p}_S, \tilde{A}_S) = v_{i\tilde{A}_S(i)} - \tilde{p}_{\tilde{A}_S(i)} = v_{i\tilde{A}_S(i)} - (p_{\tilde{A}_S(i)} + \bar{c}_{\tilde{A}_S(i)}) = (v_{i\tilde{A}_S(i)} - \bar{c}_{\tilde{A}_S(i)}) - p_{\tilde{A}_S(i)} = (v_{iA(i)} - \bar{c}_{A(i)}) - p_{A(i)} \geq (v_{ij} - \bar{c}_j) - p_j = v_{ij} - (p_j + \bar{c}_j) = v_{ij} - \tilde{p}_j$, where the inequality is because (p, A) is envy-free with respect to V'. $\square$

We are now ready to prove the main theorem.

Proof of Theorem 2. Notice that $Rev(\tilde{p}_S, \tilde{A}_S) - \sum_{j\in S} \bar{c}_j = Rev(p, A; V') \geq \frac{R([m];V')}{\alpha} \geq \frac{Rev(\bar{p},\bar{A};V')}{\alpha} = \frac{NR(S^*_{\bar{c}},\bar{c})}{\alpha}$, where the first equality is due to Lemma 8, the first inequality is because Alg_{value} is an α-approximation of the value oracle for the dealer, the second inequality is due to Lemma 7, and the second equality is due to Lemma 6. Therefore, Alg_{demand} is an α-approximation for the dealer and Theorem 2 holds. $\square$

5. BUDGET FEASIBLE MECHANISMS FOR DEALERS

Recall that our main goal is to design individually rational, truthful, and budget feasible mechanisms for a dealer. Moreover, we want the mechanisms to return envy-free assignments that maximize the dealer's revenue (i.e., $R(S^*)$).

As shown earlier, the dealer's revenue function is subadditive. Previous work on designing budget feasible mechanisms for subadditive functions have relied on value and demand oracle queries. However, for envy-free revenue maximization, the existence of those oracles is questionable, as the optimal revenue is hard to compute. Thus, we are interested in mechanisms that do not access any oracle.

Given an α-approximation algorithm, Alg_{value}, of the revenue function R (i.e., the dealer's value oracle) and our α-approximation algorithm, Alg_{demand}, of the optimal net revenue (i.e., the demand oracle), we construct a truthful mechanism for the dealer, which is a variant of the one in [16].

For simplicity, let $R^\alpha(S, V_S) = Rev(Alg_{value}(S, V_S))$. When the valuation matrix is clear from the context, we write $R^\alpha(S)$ instead of $R^\alpha(S, V_S)$. Let $Val = \{1, 2, ..., 2^{\log(\alpha R^\alpha([m]))}\}$. As in [16], we assume without loss of generality that the smallest non-zero value of the revenue function R is 1, and we start with an algorithm, Alg_t, for finding a set of items of size t whose revenue is approximately optimal. Here, we only need to run Alg_{demand} a polynomial number of times.

Claim 2. *Let S and $(\tilde{p}_S, \tilde{A}_S)$ be the output of Alg_{demand} with inputs V and a cost vector $\bar{c}$ such that $\bar{c}_1 = \bar{c}_2 = \cdots = \bar{c}_m$. For any $S' \subseteq S$, we have $Rev(\bar{p}_{S'}, \bar{A}_{S'}) \geq |S'|\bar{c}_1$, where $(\bar{p}_{S'}, \bar{A}_{S'})$ is the envy-free assignment output by Alg_{Assign} given S, V_S, $(\tilde{p}_S, \tilde{A}_S)$, and S'.*

Algorithm 3: Alg_t - Finding a Set of Items of Size t with Approximately Optimal Revenue

Input : A positive integer t.
Output: A value $v \in Val$, a set S_v, and an envy-free assignment $(\bar{p}_{S_v}, A_{S_v})$.

```
Let S̄ = (S_v)_{v∈Val}.
for v ∈ Val do
    Let c̄ = (c̄_j)_{j=1}^m where c̄_j = v/(2t) for all j ∈ [m]
    and run Alg_demand with V and c̄ to obtain a
    subset of items S ⊆ [m] and an envy-free
    assignment (p̃_S, Ã_S).
    Let S_v = ∅.
    if Rev(p̃_S, Ã_S) − |S| v/(2t) ≥ v/(2α) then
        Let S_v = S.
        if |S| > t then
            Let S_v be some set of items of size t
            such that S_v ⊆ S.
        end if
    end if
end for
Run Alg_Assign with inputs S, V_S, (p̃_S, Ã_S), and
S_v and let (p̄_{S_v}, Ā_{S_v}) be the output.
return (v, S_v, p̄_{S_v}, Ā_{S_v}) for the maximum
v ∈ Val such that S_v is not empty from S̄
```

Proof. We have $Rev(\bar{p}_{S'}, \bar{A}_{S'}) - |S'|\bar{c}_1 = \sum_{j\in S'} \bar{p}_j - |S'|\bar{c}_1 = \sum_{j\in S'}(\bar{p}_j - \bar{c}_1) \geq \sum_{j\in S'}(\tilde{p}_j - \bar{c}_1) \geq 0$, where the first equality is because all items in S' are sold according to Alg_{Assign}, the first inequality is because $\bar{p}_j \geq \tilde{p}_j$ for all $j \in S'$ according to Alg_{Assign}, and the second inequality is because $\tilde{p}_j \geq \bar{c}$ for all $i \in S$ by Alg_{demand}. Thus, $Rev(\bar{p}_{S'}, \bar{A}_{S'}) \geq |S'|\bar{c}$ and Claim 2 holds. $\square$

Lemma 9. *Let $S^* \in \operatorname{argmax}_{|S|=t} R(S)$. The algorithm Alg_t finds a subset S_v such that $Rev(\bar{p}_{S_v}, \bar{A}_{S_v}) \geq \frac{R(S^*)}{4\alpha}$.*

Proof. From the definition of Val, the value $v \in Val$ is such that either $v \leq R(S^*)$ or $v > R(S^*)$. We first consider the case where $v \leq R(S^*)$. Given V and a cost $\bar{c}_j = \frac{v}{2t}$ for all items $j \in [m]$, Alg_{demand} returns an item-set S and an envy-free assignment $(\tilde{p}_S, \tilde{A}_S)$. Recall that $S^*_{\bar{c}}$ is the output of the dealer's demand oracle. We have $Rev(\tilde{p}_S, \tilde{A}_S) - |S|\bar{c}_1 \geq \frac{NR(S^*_{\bar{c}},\bar{c})}{\alpha} = \frac{R(S^*_{\bar{c}})-|S^*_{\bar{c}}|\bar{c}_1}{\alpha} \geq \frac{R(S^*)-|S^*|\bar{c}_1}{\alpha} = \frac{R(S^*)-t\frac{v}{2t}}{\alpha} = \frac{R(S^*)-\frac{v}{2}}{\alpha} \geq \frac{R(S^*)-\frac{R(S^*)}{2}}{\alpha} = \frac{R(S^*)}{2\alpha} \geq \frac{v}{2\alpha}$, where the first inequality is by Theorem 2, the second inequality is because $S^*_{\bar{c}}$ is the output of the dealer's demand oracle, and the third and last inequalities are because $v \leq R(S^*)$. Therefore, $S_v \neq \emptyset$ and $Rev(\tilde{p}_S, \tilde{A}_S) \geq \frac{R(S^*)}{2\alpha}$.

If $|S| \leq t$, then $S_v = S$ and $Rev(\bar{p}_{S_v}, \bar{A}_{S_v}) \geq Rev(\tilde{p}_{S_v}, \tilde{A}_{S_v})$; thus $Rev(\bar{p}_{S_v}, \bar{A}_{S_v}) \geq \frac{R(S^*)}{2\alpha} \geq \frac{R(S^*)}{4\alpha}$, as desired. If $|S| > t$, then $Rev(\bar{p}_{S_v}, \bar{A}_{S_v}) \geq |S_v|\bar{c}_1 = t\frac{v}{2t} = \frac{v}{2}$ by Claim 2. Since v is the maximal $v \in Val$ such that $v \leq R(S^*)$, we have that $v \geq \frac{R(S^*)}{2}$, therefore $Rev(\bar{p}_{S_v}, \bar{A}_{S_v}) \geq \frac{R(S^*)}{4} \geq \frac{R(S^*)}{4\alpha}$.

Now consider the case of $v > R(S^*)$. When step b is passed, if $|S| \leq t$, then $S_v = S$ and $Rev(\bar{p}_{S_v}, \bar{A}_{S_v}) \geq Rev(\tilde{p}_{S_v}, \tilde{A}_{S_v}) = Rev(\tilde{p}_S, \tilde{A}_S) \geq \frac{v}{2\alpha} \geq \frac{R(S^*)}{2\alpha} \geq \frac{R(S^*)}{4\alpha}$; if $|S| > t$, by Claim 2, $Rev(\bar{p}_{S_v}, \bar{A}_{S_v}) \geq |S_v|\bar{c}_1 = t\frac{v}{2t} = \frac{v}{2} \geq \frac{R(S^*)}{2} \geq \frac{R(S^*)}{4\alpha}$. In sum, Lemma 9 holds. $\square$

Using algorithm Alg_t, we construct a mechanism M_D for the dealer. Let $\beta = 2\lceil \log m \rceil$ and let $j^* \in [m]$ such that $j^* \in \text{argmax}_{j\in[m]} R(\{j\})$. It is relatively easy to compute j^*: look at the item that has the highest value among all buyers and assign that item to the buyer with the highest value and charge the buyer that value and set the price of all other items to be sufficiently high so that nobody will be assigned to any other item. Let (p^*_j, A^*_j) be the optimal envy-free assignment of $\{j^*\}$.

Mechanism 1: M_D - A Randomized Budget Feasible Mechanism for a Dealer

Input : Each seller $j \in [m]$ reports a cost $c_j > 0$.
Output: $(S, (q_j)_{j\in S}, \bar{p}_S, \bar{A}_S)$ or $(j^*, q_{j^*}, p^*_{j^*}, A^*_{j^*})$

1 Let $T = \{1, 2, ..., 2^{\lceil \log m \rceil}\}$.
2 **for** $t \in T$ *in decreasing order* **do**
3   Let N' be the set of items with cost at most $B/(\beta t)$ that are different from j^*.
4   Using the algorithm Alg_t with input t, find $(v_t, S_t, \bar{p}_{S_t}, \bar{A}_{S_t})$ among items in N'. Let $S = \cup_t S_t$, and let $(\bar{p}_S, \bar{A}_S)$ be the output of Alg_{value} on S.
5   **for** $j \in S$ **do**
6   let q_j be the threshold for seller j: that is, the highest cost j can announce so that he is still selected, given c_{-j}.
7   **end for**
8 **end for**
9 Flip a fair coin.
10 **if** *coin* $=$ *heads* **then**
11   The dealer buys items in S, pays q_j to each seller $j \in S$, pays 0 to each seller $j \notin S$, and sells the items to the buyers according to $(\bar{p}_S, \bar{A}_S)$.
12 **else**
13   The dealer buys item j^*, pays B to seller j^*, and sells the item according to $(p^*_{j^*}, A^*_{j^*})$.
14 **end if**

Theorem 3. *The mechanism M_D is universally truthful, budget feasible, and is an $O((\log^2 m)\alpha^2)$-approximation to the optimal procurement.*

Proof. To show truthfulness of M_D, we note that in single-parameter settings, a mechanism is truthful if and only if it satisfies monotonicity and threshold payment (see [2] for the definition). Monotonicity comes from steps 2-3 of M_D and threshold payment is explicitly required in steps 5-6. Finally, it can be shown that the total payment is budget feasible in expectation. The analysis of the universal truthfulness and budget feasibility of M_D follows from the same reasoning as in [16] and thus is omitted. Below, we prove the claimed approximation ratio.

Recall the optimal procurement of the dealer is a subset of items S^* such that $S^* \in \text{argmax}_{S\subseteq[m],\ \sum_{j\in S} c_j \le B} R(S)$. Let $(\hat{p}_{S^*}, \hat{A}_{S^*})$ be the optimal envy-free assignment such that $R(S^*) = Rev(\hat{p}_{S^*}, \hat{A}_{S^*})$. Let $\bar{S}^* = \{i | i \in S^*, c_i > \frac{B}{\beta}\}$, and $\underline{S}^* = S^* - \bar{S}^*$. By Theorem 1, R is subadditve and $R(S^*) \le R(\bar{S}^*) + R(\underline{S}^*)$. Thus, $R(\bar{S}^*) \ge \frac{R(S^*)}{2}$ or $R(\underline{S}^*) \ge \frac{R(S^*)}{2}$.

We first consider the case that $R(\bar{S}^*) \ge \frac{R(S^*)}{2}$. Because the costs of the items in $\bar{S}^*$ is at least $\frac{B}{\beta}$, the payment for each of the items in $\bar{S}^*$ is at least $\frac{B}{\beta}$. Thus, $B \ge |\bar{S}^*|\frac{B}{\beta}$ and $\beta \ge |\bar{S}^*|$. By subaddtivity, we have $\sum_{j\in\bar{S}^*} R(\{j\}) \ge R(\bar{S}^*)$. Let $\bar{j} \in \text{argmax}_{j\in\bar{S}^*} R(\{j\})$, we have $R(\{\bar{j}\}) \ge \frac{R(\bar{S}^*)}{\beta}$. Moreover, $R(\{j^*\}) \ge R(\{\bar{j}\}) \ge \frac{R(\bar{S}^*)}{\beta} \ge \frac{R(S^*)}{2\beta}$. Since $Rev(p^*_j, A^*_j) = R(\{j^*\})$, with probability 1/2, j^* will be selected and we have a $O(\log m)$-approximation.

Now we consider the second case that $R(\underline{S}^*) \ge \frac{R(S^*)}{2}$. If $R(\{j^*\}) \ge \frac{R(\underline{S}^*)}{2} \ge \frac{R(S^*)}{4}$, then with probability 1/2, we have $R(\{j^*\}) \ge \frac{R(S^*)}{8}$ and we are done. Otherwise, let $\underline{S}^{*\prime} = \underline{S}^* \setminus \{j^*\}$. It follows that, by subadditivity, $R(\underline{S}^{*\prime}) + R(\{j^*\}) \ge R(\underline{S}^*)$, that is, $R(\underline{S}^{*\prime}) \ge R(\underline{S}^*) - R(\{j^*\})$, which implies $R(\underline{S}^{*\prime}) \ge R(\underline{S}^*) - \frac{R(\underline{S}^*)}{2} \ge \frac{R(\underline{S}^*)}{2} \ge \frac{R(S^*)}{4}$.

In what follows, we partition the items in $\underline{S}^{*\prime}$ based on their costs into $O(\log m)$ bins. Let $\text{Bin}(i) = \{j \in \underline{S}^{*\prime} \mid B/(\beta 2^{i+1}) < c_j \le B/(\beta 2^i)\}$ for $i \in \{0, 1, ..., \lceil \log m \rceil - 1\}$ and $\text{Bin}(\lceil \log m \rceil) = \{j \in \underline{S}^{*\prime} \mid 0 < c_j \le B/(\beta 2^{\lceil \log m \rceil})\}$. By subadditivity, there is a bin k such that $R(\text{Bin}(k)) \ge R(\text{Bin}(i))$ for all $i \in \{0, 1, ..., \lceil \log m \rceil\}$ and $R(\text{Bin}(k)) \ge \frac{R(\underline{S}^{*\prime})}{O(\log m)} = \frac{R(\underline{S}^*)}{O(\log m)}$, that is, the optimal revenue of the items in $\text{Bin}(k)$ is at least $O(\log m)$-fraction of $R(\underline{S}^*)$. By construction, $\text{Bin}(k)$ has at most $\beta 2^{k+1}$ items, it follows that the optimal solution of size $\beta 2^{k+1}$ of items with cost at most $B/(\beta 2^k)$ has revenue at least $\frac{R(\underline{S}^*)}{O(\log m)}$. Consider the iteration at $t = 2^k$, algorithm Alg_t with $t = 2^k$ gives us a set S_t that is a 4α-approximation to the optimal revenue of 2^k items (Lemma 9). Notice that by subadditivity, the optimal revenue of 2^k items is an 2β-approximation to the optimal revenue of $\beta 2^{k+1}$ items. It follows that $R(S_t)$ is an $8\alpha\beta$-approximation to the optimal solution of size $\beta 2^{k+1}$ with cost at most $B/(\beta 2^k)$. Thus, $R(S_t) \ge \frac{R(\underline{S}^*)}{O(\alpha\beta \log m)}$, and, by monotonicity, $R(\cup_t S_t) \ge R(S_t) \ge \frac{R(\underline{S}^*)}{O(\alpha \log^2 m)} \ge \frac{R(S^*)}{O(\alpha \log^2 m)}$. However, we do not know $R(\cup_t S_t)$, and we run Alg_{value} to get a α-approximation of it. Therefore, with probability $\frac{1}{2}$, we have the desired $O((\log^2 m)\alpha^2)$-approximation. □

Remark 3. *Notice that the dealer will buy all items in S, but Alg_{value} may not sell all of them, thus the dealer's revenue may be less than his payment to the sellers. This has to be the case because the requirement of truthfulness. Indeed, if the dealer only buys items that are actually sold by Alg_{value}, the mechanism is not monotone and thus is not truthful. How to design truthful mechanisms where the dealer always sells all his purchased items is an interesting open problem.*

From Corollary 1, there is $O(\log n)$-approximation for the demand oracle. Thus, we have the following result.

Corollary 2. *There is a mechanism that is universally truthful, budget feasible, and is an $O((\log^2 m)\ (\log^2 n))$- approximation to the optimal procurement.*

Acknowledgments

We thank several anonymous reviewers for their comments.

REFERENCES

[1] N. Anari, G. Goel, and A. Nikzad. Mechanism design for crowdsourcing: An optimal 1-1/e competitive budget-feasible mechanism for large markets. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, FOCS '14, pages 266–275, 2014.

[2] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS'01)*, pages 482–491, 2001.

[3] M.-F. Balcan, A. Blum, and Y. Mansour. Item pricing for revenue maximization. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, EC '08, pages 50–59, 2008.

[4] X. Bei, N. Chen, N. Gravin, and P. Lu. Budget feasible mechanism design: From prior-free to Bayesian. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 449–458, 2012.

[5] P. Briest. Uniform budgets and the envy-free pricing problem. In *the 35th International Colloquium on Automata, Languages and Programming, Part I*, ICALP '08, pages 808–819, 2008.

[6] P. Briest and H. Röglin. The power of uncertainty: Bundle-pricing for unit-demand customers. In *Proceedings of the 8th International Conference on Approximation and Online Algorithms*, WAOA'10, pages 47–58, 2011.

[7] H. Chan and J. Chen. Truthful multi-unit procurements with budgets. In T.-Y. Liu, Q. Qi, and Y. Ye, editors, *Web and Internet Economics*, volume 8877, pages 89–105. Springer, 2014.

[8] S. Chawla, J. D. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pages 243–251, 2007.

[9] N. Chen and X. Deng. Envy-free pricing in multi-item markets. *ACM Trans. Algorithms*, 10(2):7:1–7:15, Feb. 2014.

[10] N. Chen, A. Ghosh, and S. Vassilvitskii. Optimal envy-free pricing with metric substitutability. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, EC '08, pages 60–69, 2008.

[11] N. Chen, N. Gravin, and P. Lu. On the approximability of budget feasible mechanisms. In *the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 685–699, 2011.

[12] X. Chen, I. Diakonikolas, D. Paparas, X. Sun, and M. Yannakakis. The complexity of optimal multidimensional pricing. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, pages 1319–1328, 2014.

[13] M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 35–44, 2008.

[14] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:pp. 17–33, 1971.

[15] C. Daskalakis, A. Deckelbaum, and C. Tzamos. Optimal pricing is hard. In *Proceedings of the 8th International Conference on Internet and Network Economics*, WINE'12, pages 298–308, 2012.

[16] S. Dobzinski, C. H. Papadimitriou, and Y. Singer. Mechanisms for complement-free procurement. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, pages 273–282, 2011.

[17] K. Elbassioni. A QPTAS for ϵ-envy-free profit-maximizing pricing on line graphs. In *Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming - Volume Part II*, ICALP'12, pages 513–524, 2012.

[18] K. Elbassioni, R. Raman, S. Ray, and R. Sitters. On profit-maximizing pricing for the highway and tollbooth problems. In *Proceedings of the 2Nd International Symposium on Algorithmic Game Theory*, SAGT '09, pages 275–286, 2009.

[19] M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *the 13th ACM Conference on Electronic Commerce*, EC '12, pages 532–549, 2012.

[20] A. Grigoriev, J. van Loon, R. Sitters, and M. Uetz. How to sell a graph: Guidelines for graph retailers. In *Proceedings of the 32nd International Conference on Graph-Theoretic Concepts in Computer Science*, WG'06, pages 125–136, 2006.

[21] T. Groves. Incentives in teams. *Econometrica*, 41(4):617–31, July 1973.

[22] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 1164–1173, 2005.

[23] S. Im, P. Lu, and Y. Wang. Envy-free pricing with general supply constraints. In *Proceedings of the 6th International Conference on Internet and Network Economics*, WINE'10, pages 483–491, 2010.

[24] D. Kempe, A. Mu'Alem, and M. Salek. Envy-free allocations for budgeted bidders. In *Proceedings of the 5th International Workshop on Internet and Network Economics*, WINE '09, pages 537–544, 2009.

[25] S. Khuller, A. Moss, and J. S. Naor. The budgeted maximum coverage problem. *Inf. Process. Lett.*, 70(1):39–45, Apr. 1999.

[26] A. Krause and C. Guestrin. A note on the budgeted maximization of submodular functions. Technical report, CMU-CALD-05-103, 2005.

[27] P. McAfee. A dominant strategy double auction. *J. of Economic Theory*, 56(2):434–450, April 1992.

[28] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *J. of Economic Theory*, 29(2):265–281, April 1983.

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

[30] A. Ray, D. Mandal, and Y. Narahari. Profit maximizing prior-free multi-unit procurement auctions with capacitated sellers. In *the 14th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1753–1754, 2015.

[31] Y. Singer. Budget feasible mechanisms. In *Proceedings of the 51st IEEE Annual Symposium on Foundations of Computer Science*, FOCS '10, pages 765–774, 2010.

[32] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 03 1961.