

Simple Mechanisms for Profit Maximization in Multi-item Auctions

Yang Cai*
Department of Computer Science
Yale University
yang.cai@yale.edu

Mingfei Zhao†
School of Computer Science
McGill University
mingfei.zhao@mail.mcgill.ca

June 25, 2019

Abstract

We study a classical Bayesian mechanism design problem where a seller is selling multiple items to multiple buyers. We consider the case where the seller has costs to produce the items, and these costs are private information to the seller. How can the seller design a mechanism to maximize her profit? Two well-studied problems, revenue maximization in multi-item auctions and signaling in ad auctions, are special cases of our problem. We show that there exists a simple mechanism whose profit is at least $\frac{1}{44}$ of the optimal profit for multiple buyers with matroid-rank valuation functions. When there is a single buyer, the approximation factor is 11 for general constraint-additive valuations and 6 for additive valuations. Our result holds even when the seller's costs are correlated across items.

We introduce a new class of mechanisms called *permit-selling mechanisms*. For single buyer case these mechanisms are quite simple: there are two stages. For each item j , we create a separate permit that allows the buyer to purchase the item in the second stage. In the first stage, we sell the permits without revealing any information about the costs. In the second stage, the seller reveals all the costs, and the buyer can buy item j by paying the item price if the buyer has purchased the permit for item j in the first stage. We show that either selling the permits separately or as a grand bundle suffices to achieve the constant factor approximation to the optimal profit (6 for additive, and 11 for constrained additive). For multiple buyers, we sell the permits sequentially and obtain the constant factor approximation. Our proof is enabled by constructing a benchmark for the optimal profit by combining a novel dual solution with the existing ex-ante relaxation technique.

*Part of Yang Cai's work is done under the support of NSERC Discovery RGPIN-2015-06127 and FRQNT 2017-NC-198956.

†Supported by Richard H. Tomlinson Doctoral Fellowship and NSERC Discovery RGPIN-2015-06127.

1 Introduction

We study the profit maximization problem for selling multiple items to multiple buyers. Unlike most works in Mechanism Design, we consider the case where the seller has costs for obtaining the items. As these costs usually depend on private information that is only available to the seller, we assume that the costs are private to the seller but are drawn from a distribution also known to the buyers. The goal is to design a mechanism that maximizes the profit, that is, the total revenue less the total cost. Revenue maximization in multi-item auctions, one of the most classical and widely studied problem in Bayesian Mechanism Design [10, 11, 6, 8, 2, 29, 16, 17, 25, 27], is a special case of our problem, where the seller always has cost 0 for each item. Arguably, it is more natural and general to assume that there is a cost associated with each item. It may be production cost or opportunity cost, e.g., there is an outside option to sell the item at a certain price.

Despite being realistic and widely applicable, the profit maximization problem is not well-understood. To the best of our knowledge, the only case with non-zero costs that has been studied is in the context of ad auctions [4, 24, 20, 22, 18]. The problem models the following scenario. The auctioneer is selling an ad displaying slot to an advertiser. There are m types of viewers of the webpage. The advertiser has value t_j for displaying his ad to a type j viewer. In an ad auction, only the auctioneer observes the type of the viewer, and the advertiser only knows a prior distribution π from which the viewer-type is drawn from. What mechanism maximizes the auctioneer’s revenue? We can easily cast this problem as a profit maximization problem. Let there be m items. When the viewer has type j , assign item j with cost 0 and all the other items with cost ∞ . Clearly, maximizing revenue in the ad auction is equivalent to maximizing profit in the corresponding multi-item setting, as the seller will never sell any item with ∞ cost. Many results are known for this special case, which we will discuss in Section 1.4.

In revenue maximization, the optimal mechanism is known to be randomized and complex in multi-item settings. Instead of characterizing the optimal mechanisms, a major and successful research theme in Mechanism Design is devoted to designing simple and approximately optimal mechanisms [10, 11, 13, 2, 8, 27, 28, 31, 7, 9, 29]. As our problem generalizes the multi-item revenue maximization problem, it is clear that the profit-optimal mechanism also requires complex allocation rules and randomization. In this paper, we focus on designing simple and approximately optimal mechanisms for profit maximization.

To facilitate the discussion, we will first focus on the single buyer case. In our model, there are m items for sale, and the seller has cost c_j for parting with item j . The costs (c_1, \dots, c_m) are drawn from a distribution \mathcal{C} that is known to both the seller and the buyer. We allow the seller’s costs to be correlated across items. Consider constrained-additive buyers, that is, the buyer has a downward-closed feasibility constraint $\mathcal{F} \subseteq 2^{[m]}$ that specifies what bundles of items are allowed. The buyer has value t_j for item j , and her value for a bundle S is defined as $\max_{A \in \mathcal{F}, A \subseteq S} \sum_{i \in A} t_j$. Similar to most results in the simple vs. optimal literature, we assume t_j to be drawn from \mathcal{D}_j independently across items.

Before we state our results, let us examine two natural but unsuccessful attempts to solve this problem.

Two unsuccessful attempts: (i) Use a mechanism that (approximately) optimizes the revenue. This is a terrible solution as some of the items sold by the mechanism may have extremely high costs, and as a result the mechanism only generates low if any profit. (ii) After the seller sees the costs, reveal them to the buyer, then use the optimal or approximately optimal mechanism that is tailored to those particular costs. The reason why this mechanism may be far away from optimal is much more subtle. Let us consider the following example in the ad auction setting ¹.

¹The example is similar to the example that shows the revenue of selling the items separately could be $\Omega(\log n)$ -factor worse than the optimal revenue in multi-item auctions with an additive bidder.

Example 1. A random variable X with support $[1, +\infty)$ follows the **equal revenue (ER) distribution** if and only if $\Pr[X \leq x] = 1 - \frac{1}{x}$. Let \mathcal{D}_j be the ER distribution for each item j . Define δ_j to be the n -dimensional vector whose j -th entry is 0 and all the other entries are ∞ . Let the costs $(c_1, \dots, c_m) = \delta_j$ have probability $1/n$ for each j .

The expected profit of the mechanism in (ii): $\sum_{j \in [m]} \frac{1}{m} \cdot \max_p p \cdot \Pr_{t_j \sim \mathcal{D}_j}[t_j \geq p] = 1$. Since for every j , after revealing cost δ_j , the seller can only sell item j to the buyer.

Consider an alternative mechanism which does not reveal the costs, but offers the buyer the following contract: if the buyer pays $\log m/2$ up front, the buyer can take any item that is available, e.g. has cost 0. The chance that the buyer accepts the contract is

$$\Pr \left[\frac{1}{m} \cdot \sum_{j \in [m]} t_j \geq \frac{\log m}{2} \right],$$

and due to a Lemma by [27], we know that it is at least $1/2$. Hence, the mechanism has profit at least $\frac{\log m}{4}$.

Note that in the second mechanism of our example, the seller does not even reveal any information to the buyer and extracts much higher profit.

The two failed attempts highlight two major challenges of our problem: (i) how to balance the revenue and cost; (ii) how to capture the informational rent of the buyer, that is, leveraging the fact that the costs are private information to extract more revenue. We overcome these two challenges by considering what we called *permit-selling* mechanisms. These mechanisms have two stages. For each item j , we create a separate permit that allows the buyer to purchase the item at its cost. In the first stage, we sell the permits without revealing any information about the actual costs. In the second stage, the seller reveals all the costs, and the buyer can buy item j by only paying the cost c_j if the buyer has purchased the permit for item j in the first stage. How does the buyer make a decision in such a mechanism? In the first stage, the buyer needs to choose her favorite bundle of permits to purchase. Since she knows the distribution \mathcal{C} , she can compute her utility for each bundle of permits. In the second stage, the buyer simply picks her favorite set of items based on the permits she own, the costs of the items, and her valuation function. Why do the permit-selling mechanisms help addressing the two challenges? Note that the profit of the permit-selling mechanisms is exactly the revenue from the first stage, so any mechanism that achieves high revenue in the first stage also generates high profit. Moreover, the buyer needs to make a decision on what permits to purchase without learning the costs, therefore, the seller can extract the informational rent by pricing the permits appropriately.

Indeed, we do not even need to use any complex pricing scheme in the first stage. We sell the permits separately or sell them as a grand bundle. In our proof we need one more mechanism, which simply sells the items separately, and the prices change according to the seller's costs. The reason why this class of mechanism is required is more subtle and we only sketch the intuition here. In the permit-selling mechanism, for a fixed buyer type profile \mathbf{t} , the buyer purchases a set of permits P and thus the seller can only extract revenue from items in P , no matter what realized costs she has. However, for different cost vectors, the seller may have different items from which she can extract more revenue. By posting item prices that depend on her cost, the seller is able to target the profitable items based on her realized cost vector. This approach does not capture the informational rent but may generate high profit in certain cases.

Here are the mechanisms we use.

- **sell-items-separately (IS):** for each possible cost vector $\mathbf{c} = (c_1, \dots, c_m)$, sell the items separately, and the price $p_j(\mathbf{c})$ for item j depends on \mathbf{c} .
- **sell-permits-separately (PS):** sell the permits separately, and the price p_j for the j -th permit is independent from the seller's costs.

- **permit-bundling (PB)**: sell all the permits as a grand bundle at a price p that is independent from the seller’s costs.

Since in all these mechanisms, the seller does not even ask the buyer to report her valuation, the mechanism is clear incentive compatible (IC) and individually rational (IR). We show that the best mechanism among these three classes of mechanisms can already achieve a constant fraction of the optimal profit.

Theorem 1. *For any valuation distribution $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_m$, cost distribution \mathcal{C} , and any downward-closed feasibility constraint \mathcal{F} , the best mechanism among all sell-items-separately, sell-permits-separately, and permit-bundling mechanisms is an 11-approximation to the optimal profit.*

When the buyer’s valuation is additive, we can improve the approximation factor to 6.

Theorem 2. *If the buyer has additive valuation, for any valuation distribution $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_m$ and cost distribution \mathcal{C} , the best mechanism among all sell-items-separately, sell-permits-separately, and permit-bundling mechanisms is a 6-approximation to the optimal profit.*

We then generalize the result to accommodate multiple buyers with matroid-rank valuation functions. With multiple buyers, we sell the permits with a sequential mechanism: buyers arrive in some arbitrary order. When a buyer arrives, we first offer the buyer the permits without revealing any information about the seller’s cost or other buyers’ types. Next, the buyer is given the remaining item set as well as an item price for each item. For single buyer case the item price is always chosen as the seller’s cost for this item. Now we allow any prices that may depend on the cost vector.² Again the buyer can purchase any item from the remaining item set by paying the corresponding item price, if she has the permit for it. The mechanism is BIC and interim IR as the buyer needs to purchase the permit without knowing the seller’s cost or what items are still available. In the proof, we use Sequential Permit Posted Price (SPP) mechanisms that sell permits separately and Sequential Permit Bundling (SPB) mechanisms that sell permits as a whole bundle. Similar to the single buyer case, we need another mechanism called Constrained Sequential Item Posted Price mechanism (CSIP). It resembles the Sequential Posted Price mechanism from [11]: the items are sold sequentially with posted prices that depend on the seller’s cost vector. The mechanism also imposes a constraint on the set of items a buyer can purchase. See Section 2 for more details.

We prove that the best of three classes of mechanisms can already achieve a constant fraction of the optimal profit.

Theorem 3. *For any cost distribution and buyers’ valuation distributions, if every buyer has a matroid-rank valuation, the best mechanism among all CSIP, restricted SPP³, and SPB mechanisms is a 44-approximation to the optimal profit.*

1.1 Proof Sketch and Techniques

Since the costs are private, it is a priori not clear that it is sufficient to consider only direct mechanisms. Indeed, signaling mechanisms, a class of indirect mechanisms, are widely studied in the ad auction setting [4, 24, 20, 22, 18]. We first prove a revelation principle for our problem similar to the one proved in [18] for ad auctions. Our revelation principle states that w.l.o.g. we can restrict our attention to direct, BIC, and interim IR mechanisms. Moreover, we can formulate the profit maximization problem as an LP. We next apply the Cai-Devanur-Weinberg duality framework [7]. The framework has become a standard tool for

²In the proof, the item prices is always chosen to be no less than the seller’s cost.

³It’s closed to the Sequential Permit Posted Price mechanism, except that the mechanism may hide some items randomly, preventing the buyer from buying some item even she has permit. See Section 2 for more details.

analyzing the performance of simple mechanisms. In most of the results based on this duality approach, a particular family of dual variables, called the “canonical dual” [7, 9], is used to provide a benchmark for the objective function. However, this set of dual variables does not provide an appropriate benchmark due to the existence of costs. We propose a new set of dual variables that is tailored to handle the costs. Indeed, these dual variables are so informative that they inspired us to introduce the permit-selling mechanisms. In the multi-buyer case, the choice of the dual variables is also inspired by the ex-ante relaxation technique from [13]. A similar set of dual variables are used in [9] to provided a benchmark for the optimal revenue.

The benchmark induced by our dual variables can be easily decomposed into three components – MOST-SURPLUS, PROPHET and LESS-SURPLUS. MOST-SURPLUS can be bounded by the profit of the CSIP mechanism using relatively standard analysis. For PROPHET, we bound the term using the same class of mechanisms, with the help of the Online Contention Resolution Scheme [23].

For LESS-SURPLUS, in order to establish a connection between profit maximization and revenue maximization, we provide a separate and clean proof for the single buyer case. Instead of directly analyzing the term, we construct an auxiliary revenue maximization problem for selling m items to help approximate LESS-SURPLUS. Intuitively, each item in the auxiliary problem corresponds to a permit. We first show that any mechanism in the auxiliary problem can be turned into a permit-selling mechanism in the original problem, such that the revenue in the auxiliary problem is the same as the profit of the permit-selling mechanism. Next, we argue that the buyer has subadditive valuation in the auxiliary problem whenever the buyer has constrained additive valuation in the original problem. Note that the better of selling the items separately and grand bundling is a constant factor approximation of the optimal revenue when the buyer has subadditive valuation [29, 9]. Unfortunately, we cannot use this approximation as a black-box, as it is not yet clear how the revenue in the auxiliary problem relates to the LESS-SURPLUS term. Luckily, Cai and Zhao obtain their result via the CDW duality framework, and in their analysis, they show that a term identical to LESS-SURPLUS can be approximated by the revenue of selling the items separately or grand bundling. Putting everything together, we prove that the profit of a sell-permits-separately or permit-bundling mechanism approximates LESS-SURPLUS, and that completes our proof. For general case, it’s not straightforward build such a connection. We use the standard Core-Tail Decomposition technique [28, 7], dividing LESS-SURPLUS further into two terms TAIL and CORE. TAIL can be approximated using RSPP. For CORE, it can be viewed as all buyers’ truncated welfare with respect to a related fractionally-subadditive valuation and we can bound it using SPB and RSPP, by applying the Talagrand’s concentration inequality [30].

1.2 Our Contributions

Our main contributions in this paper are the followings:

- We introduce the permit-selling mechanisms and demonstrate their ability to approximate the optimal profit.
- We construct a new set of dual variables that can accommodate costs.
- We establish a connection between profit maximization and revenue maximization for the single buyer case.

1.3 Our Model vs. Two-sided Markets

There has been increasing interest in two-sided markets in the Economics and Computation community recently [14, 15, 21, 3, 5, 1]. In a two-sided market, the mechanism should be designed to incentivize both buyers and sellers to reveal their true private information. Our model is related to the two-sided markets

but differs in the following crucial way, that is, we assume that the seller has committing power: the seller commits to a mechanism and follows the mechanism honestly. In other words, the mechanism does not need to satisfy the seller’s Incentive Compatibility (IC) constraint in our model. This is a standard assumption used in both mechanism design for one-sided markets where the seller is also the designer of the mechanism, as well as in information design where the designer commits to a certain information structure. For example, in ad auctions, the seller receives a piece of private information, the type of the item, and sends a signal to the buyer based on this private information using a pre-committed signaling scheme. The model assumes that the seller follows the signaling scheme honestly and does not impose any IC constraints on the seller.

1.4 Related Work

There is a large body of beautiful work on simple vs. optimal for revenue maximization in multi-item auctions [10, 11, 12, 26, 8, 28, 2, 31, 7, 13, 29, 9]. They showed that simple mechanisms can extract a constant fraction of the optimal revenue in rich settings. If there is a single buyer, the state-of-the-art results [29, 9] apply to subadditive valuation functions; if there are multiple buyers, the state-of-the-art result [9] applies to XOS valuation functions. However, none of these results considered costs.

The ad auction problem has also been extensively studied in the literature [4, 24, 20, 22, 18]. Signaling mechanisms had been the focus. In a signaling mechanism, the seller first sends a signal to the buyer based on the type of the viewer and according to a signaling scheme known to the buyer. The buyer updates her posterior belief of the viewer type after observing the signal. The seller then uses a mechanism tailored to the buyer’s updated posterior to sell the ad displaying slot. Many results have been obtained regarding the revenue-optimal signaling scheme. Overall, the optimal signaling scheme may be highly complex and hard to pin down. Interestingly, Daskalakis et al. showed that even if we can find the optimal signaling scheme the corresponding mechanism can still be bounded away from the optimum [18]. They showed that the optimal mechanism is direct and does not involve any signaling. Motivated by their result, we focus on simple and direct mechanisms.

In [18], they also showed how to use simple mechanisms to approximate the auctioneer’s profit in an ad auction. They established the result by reducing the problem to revenue maximization in multi-item auctions with an additive buyer. However, their reduction is ad-hoc and heavily relies on a specific property of their cost distribution, that is, the cost is always one of the δ_i s (see Example 1 for the definition). When the cost distribution is general, their reduction no longer holds, and thus is inapplicable to our problem.

2 Preliminaries

We consider the auction where a seller is selling m heterogeneous items to n buyers. We denote buyer i ’s type t_i as $\langle t_{ij} \rangle_{j=1}^m$, where t_{ij} is buyer i ’s value for item j . For each i, j , we assume t_{ij} is drawn independently from the distribution D_{ij} . Let $D_i = \times_{j=1}^m D_{ij}$ be the distribution of buyer i ’s type and $D = \times_{i=1}^n D_i$ be the distribution of the type profile. We use T_{ij} (or T_i, T) and f_{ij} (or f_i, f) to denote the support and density function of D_{ij} (or D_i, D). For notational convenience, we let \mathbf{t} to be the types profile of all buyers, t_{-i} to be the types of all buyers except i and $t_{<i}$ (or $t_{\leq i}$) to be the types of the first $i - 1$ (or i) buyers. Similarly, we define D_{-i}, T_{-i} and f_{-i} for the corresponding distributions, support sets and density functions.

Each buyer has a constraint-additive valuation, which implies that she is additive over the items but is only allowed to receive a set of items that is feasible with respect to a downward-closed family $\mathcal{F}_i \subseteq 2^{[m]}$. In other words, the buyer with type t_i has value $v_i(t_i, S) = \max_{T \subseteq S, T \in \mathcal{F}_i} \sum_{j \in T} t_{ij}$ when receiving set S . For multiple buyer setting we consider a special case of constraint-additive valuation called *matroid-rank*, where each \mathcal{F}_i is a matroid. On the other hand, the seller has a private cost c_j for producing each item j .

Denote \mathbf{c} the cost vector and \mathbf{c} is drawn from distribution \mathcal{C} . Let T^S be the support of \mathcal{C} . We allow correlated costs in our problem.

For any direct⁴ mechanism M and any \mathbf{t}, \mathbf{c} , denote $x_{ij}(\mathbf{t}, \mathbf{c})$ the probability that buyer i is receiving item j , when the buyers has type profile \mathbf{t} and seller has cost \mathbf{c} . Let $\pi_{ij}(t_i, \mathbf{c}) = \mathbb{E}_{t_{-i}}[x_{ij}(\mathbf{t}, \mathbf{c})]$ be the interim allocation probability. Similarly, use $p_i(\mathbf{t}, \mathbf{c})$ to denote the payment for buyer i . For any \mathbf{t} and \mathbf{c} , buyer i 's utility $u_i(\mathbf{t}, \mathbf{c}) = \mathbf{t}_i \cdot x_i(\mathbf{t}, \mathbf{c}) - p_i(\mathbf{t}, \mathbf{c})$. The seller has profit (revenue minus cost) $\sum_i (p_i(\mathbf{t}, \mathbf{c}) - \mathbf{c} \cdot x_i(\mathbf{t}, \mathbf{c}))$. We now define the incentive compatibility and individual rationality for our setting.

- Bayesian Incentive Compatible (BIC): reporting the true value maximizes the buyer's expected utility $\mathbb{E}_{t_{-i}, \mathbf{c}}[u_i(t_i, t_{-i}, \mathbf{c})]$.
- Dominant Strategy Incentive Compatible (DSIC): for every \mathbf{c} and every t_{-i} , reporting the true value maximizes the buyer's utility $u_i(t_i, t_{-i}, \mathbf{c})$.
- interim Individual Rational (interim IR): reporting the true value induces non-negative expected utility. $\mathbb{E}_{t_{-i}, \mathbf{c}}[u_i(t_i, t_{-i}, \mathbf{c})] \geq 0$.
- ex-post Individual Rational (ex-post IR): for every \mathbf{c} and t_{-i} , reporting the true value induces non-negative utility. $u_i(t_i, t_{-i}, \mathbf{c}) \geq 0$.

If the mechanism allocates set S to some buyer, and the buyer is only interested in a feasible subset of items $U \subset S$, the mechanism can simply allocate set U instead. This does not affect the truthfulness for all buyers and increases the seller's profit. In this paper, we will only consider mechanisms that always allocate a feasible set of items $U \in \mathcal{F}_i$ to each buyer i . Denote $\mathcal{P}(\{\mathcal{F}_i\}_{i=1}^n)$ the region for all feasible allocations x .

For every mechanism M , denote $\text{PROFIT}(\mathcal{D}, \mathcal{C}, \{\mathcal{F}_i\}_{i=1}^n, M)$ the seller's expected profit in M . We use $\text{PROFIT}(M)$ for short when $(\mathcal{D}, \mathcal{C}, \{\mathcal{F}_i\}_{i=1}^n)$ is clear and fixed.

$$\text{PROFIT}(M) = \sum_i \mathbb{E}_{\mathbf{t}, \mathbf{c}} [p_i(\mathbf{t}, \mathbf{c}) - \mathbf{c} \cdot x_i(\mathbf{t}, \mathbf{c})]$$

As we will explain in Lemma 1, it is w.l.o.g. to only consider direct, BIC, and interim IR mechanisms. Let $\text{OPT}_{\text{PROFIT}}(\mathcal{D}, \mathcal{C}, \{\mathcal{F}_i\}_{i=1}^n)$ be the optimal profit among all BIC and interim IR mechanisms (use $\text{OPT}_{\text{PROFIT}}$ for short when $(\mathcal{D}, \mathcal{C}, \{\mathcal{F}_i\}_{i=1}^n)$ is clear and fixed). Our goal is to use a simple mechanism to approximate $\text{OPT}_{\text{PROFIT}}$.

2.1 Our Mechanisms

We bound the optimal profit by the following three classes of mechanisms. The first mechanism is a variant of the Sequential Item Posted Price (SIP) mechanism, which is first purposed by [11] in the revenue maximization problem. Here we allow the seller to decide posted prices according to her cost vector. Before the auction starts, the seller decides a posted price $p_{ij}(\mathbf{c})$ for each buyer i and item j , based on her cost vector \mathbf{c} . Then buyers come one by one in an arbitrary order. Each buyer can choose her favorite bundle among all remaining items by paying the posted prices. The mechanism is DSIC and ex-post IR. We call the mechanism Constrained Sequential Item Posted Price (CSIP) if it further adds a sub-constraint on the set of items the buyer can purchase. The mechanism first decides a constraint $\mathcal{J}'(\mathbf{c})$ on the ground set of all buyer-item pairs $J = \{(i, j) | i \in [n], j \in [m]\}$, based on her true cost \mathbf{c} . A (possibly random) set $A \subseteq J$ represents a way of allocating the items. It's feasible if

⁴By Lemma 1, the revelation principle holds in the profit maximization problem. It suffices to consider direct, BIC, and interim IR mechanisms.

- Each item is allocated to at most one buyer: $\forall j, O_j = \{(i, j) \in A\}, |O_j| \leq 1$.
- Each buyer is allocated a feasible set of items: $\forall i, P_i = \{j : (i, j) \in A\}, P_i \in \mathcal{F}_i$.

Let \mathcal{J} be the family of all feasible sets. $\mathcal{J}'(\mathbf{c})$ must satisfy $\mathcal{J}'(\mathbf{c}) \subseteq \mathcal{J}$. When each buyer comes, she is only allowed to take the item that doesn't ruin the constraint $\mathcal{J}'(\mathbf{c})$. See Mechanism 1 for details.

Mechanism 1 Constrained Sequential Item Posted Price Mechanism

Require: $p_{ij}(\mathbf{c})$, the item price for $i \in [n], j \in [m]$; the constraint $\mathcal{J}' \subseteq \mathcal{J}$.

- 1: $A \leftarrow \emptyset$.
 - 2: **for** $i \in [n]$ **do**
 - 3: Reveal item prices $\{p_{ij}(\mathbf{c})\}_{j=1}^m$ to the buyer.
 - 4: i can choose a bundle S_i such that $A \cup \{(i, j)\}_{j \in S_i} \in \mathcal{J}'$.
 - 5: i receives her favorite bundle S_i^* , paying $\sum_{j \in S_i^*} p_{ij}(\mathbf{c})$.
 - 6: $A \leftarrow A \cup \{(i, j)\}_{j \in S_i^*}$.
 - 7: **end for**
-

For the CSIP used in our proof, the corresponding sub-constraint $\mathcal{J}'(\mathbf{c})$ can be computed efficiently. See Section 6.2 for more details. We use CSIP-PROFIT to denote the optimal seller's profit among all CSIP mechanisms.

Next, we define the two Sequential Permit Selling mechanisms used in the proof. The second class of mechanism is called Sequential Permit Posted Price (SPP). Before the auction starts, the seller decides a posted price $p_{ij}(\mathbf{c})$ for each buyer i and item j , based on her cost vector \mathbf{c} . Then buyers come one by one in an arbitrary order. For each buyer i there are two stages: the permit-purchasing stage and item-purchasing stage. In the permit-purchasing stage, instead of selling the items, the seller sells a *permit* for each item. She decides a price l_{ij} for permit j independent from the seller's cost vector \mathbf{c} and buyer type profile \mathbf{t} . The buyer is allowed to purchase any permit $j \in [m]$ by paying l_{ij} . The decision must be made before she sees the remaining item set $S_i(t_{<i}, \mathbf{c})$. In the item-purchasing stage, the seller reveals $S_i(t_{<i}, \mathbf{c})$ and her cost vector \mathbf{c} to the buyer, and the buyer can purchase any remaining item j at a price of $p_{ij}(\mathbf{c})$ if the buyer has permit j . The buyer is not allowed to purchase item j if she does not have the corresponding permit. The buyer chooses her favorite bundle among the items that she is allowed to purchase. Notice that in the second stage, the buyer with set of permits $P \subseteq [m]$ will choose the bundle $S^* = \operatorname{argmax}_{S \subseteq P \cap S_i(t_{<i}, \mathbf{c}), S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - p_{ij}(\mathbf{c}))$. Thus, in the first stage, by knowing her type t_i , all the permit prices l_{ij} s, as well as the cost distribution \mathcal{C} , the buyer is able to calculate her expected surplus in the second stage for any $P \subseteq S_i(t_{<i})$. She will hence choose the best set P^* that maximizes her expected utility in the whole auction and buy all the permits in set P^* . The mechanism is only BIC and Interim IR as buyers have to make decisions before getting any information about other buyers' types and the seller's costs. See Mechanism 2 for details.

In our proof we will use restricted Sequential Permit Posted Price mechanisms (RSPP) by adding the following two changes to the mechanism: Firstly, the buyer is only allowed to purchase at most one permit on the permit-purchasing stage. Secondly, we will further allow the mechanism to hide some items from the buyer on the item-purchasing stage. Formally, the mechanism will choose a (possibly random) set $S'_i(t_{<i}, \mathbf{c}) \subseteq S_i(t_{<i}, \mathbf{c})$ and the buyer is only allowed to purchase item in $S'_i(t_{<i}, \mathbf{c})$. We will now briefly explain how our mechanism used in the proof chooses this set. In the proof, the item price $p_{ij}(\mathbf{c})$ are chosen such that for every i, j, \mathbf{c} , $\Pr_{t_{<i}}[j \in S_i(t_{<i}, \mathbf{c})] \geq \frac{1}{2}$. We define the random set $S'_i(t_{<i}, \mathbf{c})$ as follows: for any $j \in S_i(t_{<i}, \mathbf{c})$, put j in $S'_i(t_{<i}, \mathbf{c})$ with probability $\frac{1}{2} / \Pr_{t_{<i}}[j \in S_i(t_{<i}, \mathbf{c})]$, independently. Now we have $\Pr_{t_{<i}}[j \in S'_i(t_{<i}, \mathbf{c})] = \frac{1}{2}$. This is a crucial property in the proof (See Lemma 19). In the rest of the paper, when we mention RSPP, we refer to the mechanism that hides the item as above. We denote RSPP-PROFIT the optimal profit of these mechanisms.

Mechanism 2 Sequential Permit Posted Price

Require: $l_{ij}, p_{ij}(\mathbf{c})$, the permit and item price for $i \in [n], j \in [m]$.

- 1: $S \leftarrow [m]$
 - 2: **for** $i \in [n]$ **do**
 - 3: Show buyer i the permit price l_{ij} for every j .
 - 4: i chooses a set of permits $P^* \subseteq [m]$ and pays $\sum_{j \in P^*} l_{ij}$.
 - 5: Reveal S and item prices $\{p_{ij}\}_{j=1}^m$ to the buyer.
 - 6: i receives her favorite bundle $S_i^* \subseteq S$, paying $\sum_{j \in S_i^*} p_{ij}(\mathbf{c})$.
 - 7: $S \leftarrow S \setminus S_i^*$.
 - 8: **end for**
-

The third mechanism is Sequential Permit Bundling (SPB). When every buyer i comes, the seller bundles the permit of all items together and sell them as a grand bundle at some price δ_i in the first stage. δ_i is independent from \mathbf{c} . If the buyer refuses to pay the price, then she gets no permit and therefore cannot purchase anything in the second stage. If the buyer buys the permit bundle, the seller then reveals the remaining item set $S_i(t_{<i}, \mathbf{c})$ and the item prices $\{p_{ij}(\mathbf{c})\}_{j=1}^m$ to the buyer. The buyer then chooses her favorite bundle and pays the item prices. The mechanism is also BIC and interim IR due to a similar argument as for SPP. We use SPB-PROFIT to denote the optimal profit for all SPB mechanisms. See Mechanism 3 for details.

Mechanism 3 Sequential Permit Bundling

Require: $p_{ij}(\mathbf{c})$, item price for $i \in [n], j \in [m]$; δ_i , the price for the permit bundle.

- 1: $S \leftarrow [m]$
 - 2: **for** $i \in [n]$ **do**
 - 3: Show buyer i the permit bundle price δ_i .
 - 4: **if** buyer i pays price δ_i **then**
 - 5: Reveal S and item prices $\{p_{ij}\}_{j=1}^m$ to the buyer.
 - 6: i receives her favorite bundle $S_i^* \subseteq S$, paying $\sum_{j \in S_i^*} p_{ij}(\mathbf{c})$.
 - 7: $S \leftarrow S \setminus S_i^*$.
 - 8: **else**
 - 9: The buyer pays nothing and receives nothing.
 - 10: **end if**
 - 11: **end for**
-

When there is a single buyer in the auction, the mechanisms described above becomes:

- **Item Posted Pricing (IP):** for each cost vector \mathbf{c} , sell the items separately at price $p_j(\mathbf{c})$.
- **Permit Posted Pricing (PP):** sell the permits separately at price l_j that is independent from the seller's costs. We consider a specific type of mechanisms in the proof where the item price $p_j(\mathbf{c}) = c_j$.
- **Permit Bundling (PB):** sell all the permits as a grand bundle at a price δ that is independent from the seller's costs. We consider a specific type of mechanisms in the proof where the item price $p_j(\mathbf{c}) = c_j$.

More specifically, for PP and PB, we don't hide items anymore as there is no competition from other buyers. Denote IP-PROFIT, PP-PROFIT, PB-PROFIT the optimal profit of the above mechanisms.

3 Paper Organization

In this section, we provide a roadmap for our paper. In Section 4, we introduce a benchmark of the optimal profit using the CDW duality framework. We formulate the maximization problem as an LP, take the Lagrangian dual (Section 4.1), and then define a new set of dual variables (a flow) to derive our benchmark (Section 4.2).

In Section 5, we prove our result for the single constraint-additive buyer case. We divide the benchmark into two terms – MOST-SURPLUS and LESS-SURPLUS – and bound them separately. For MOST-SURPLUS, we bound it using the sell-items-separately mechanism (Section 5.1). For LESS-SURPLUS, we first construct an auxiliary revenue maximization problem for selling m items (called the revenue setting). We show that any mechanism in the auxiliary problem can be turned into a permit-selling mechanism in the original problem without changing the value of the objective (Lemma 5). Then we point out that in the benchmark of the optimal revenue in the auxiliary problem, one term is identical to LESS-SURPLUS and can be approximated by the revenue of selling the items separately or grand bundling. Thus by converting the two mechanisms to the permit-selling mechanisms, we can bound LESS-SURPLUS using the PS and PB mechanisms.

In Section 6 we study the case with multiple matroid-rank buyers. The benchmark now is divided into three terms MOST-SURPLUS, PROPHET and LESS-SURPLUS. For MOST-SURPLUS, we bound it using the Constrained Sequential Item Posted Price mechanism (Section 6.1). For PROPHET, we bound it with the same class of mechanisms, with the help of Online Contention Resolution Scheme (Section 6.2). In Section 6.3, the last term LESS-SURPLUS is bounded by both permit-selling mechanisms, using standard Core-Tail Decomposition technique.

4 Benchmark for the Maximum Profit

In this section, we construct a benchmark for the optimal profit using the Cai-Devanur-Weinberg duality framework. Before getting into the framework and benchmark, we first show that the revelation principle holds in the profit maximization problem. Therefore, it suffices to find a benchmark for the optimal profit attainable by any direct, BIC, and interim IR mechanisms. The proof is postponed to Appendix B.

Lemma 1. *Any ex-post implementable mechanism in the profit maximization problem can be implemented by a direct, BIC, and interim IR mechanism.*

4.1 Duality Framework

The framework is first developed in [7] and is widely used in mechanism design. Here we apply the framework to our profit maximization problem. We obtain an upper bound of the optimal profit similar to the upper bound of the optimal revenue obtained in [7]. More specifically, the profit of any BIC, interim IR mechanism is upper bounded by the sum of all buyers' virtual welfare minus the seller's total cost for the same allocation, with respect to some virtual value function. We will only show a sketch of the framework in the main body and refer the readers to Appendix A for a complete description.

In the framework, we first formulate the profit maximization problem as an LP. Then take the partial Lagrangian dual of the LP by lagrangifying the BIC and interim IR constraints. Since the buyer's payment is unconstrained in the partial Lagrangian, one can argue that to obtain any finite benchmark, the corresponding dual variables must form a flow. The virtual value function in the benchmark is then defined according to the choice of the dual variables/flow.

Lemma 2. *For any dual solution λ that induces a finite benchmark of the optimal profit and any BIC, interim IR mechanism $M = (x, p)$,*

$$\text{PROFIT}(M) \leq \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\sum_i \pi_i(t_i, \mathbf{c}) \cdot (\Phi_i^{(\lambda)}(t_i) - \mathbf{c}) \right]$$

where

$$\Phi_i^{(\lambda)}(t_i) = t_i - \frac{1}{f_i(t_i)} \cdot \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i)(t'_i - t_i)$$

can be viewed as buyer i 's virtual value function. Here $\pi_i(t_i, \mathbf{c}) = \mathbb{E}_{t_{-i}}[x_i(t_i, t_{-i}, \mathbf{c})]$ is the interim allocation. $\lambda(t'_i, t_i)$ is the Lagrangian dual variable for the BIC/IR constraint that says when the buyer has true type t'_i she does not want to misreport t_i .

4.2 Our Flow

Now we choose the dual variables λ carefully to induce a useful benchmark. First, let us use the single buyer case to provide some intuition behind our flow.

4.2.1 Single Buyer

In [7] and [9], they cleverly choose the canonical flow in the revenue maximization setting. They divide the type space T into m regions R_1, \dots, R_m by finding the largest value t_j among all items (called ‘‘favorite’’ item). It is the item that contributes the most to the buyer’s welfare. Then they let the flow go between two nodes $\mathbf{t}, \mathbf{t}' \in R_j$ only if they differs only on the j -th coordinate. However, the same flow does not give us a useful benchmark in our setting, as the way to divide the type space does not even depend on the information of the seller’s costs (i.e. the realized cost \mathbf{c} or the cost distribution \mathcal{C}). In [7], they also analyze another flow that is considered as a distribution of several canonical flows. We could define our flow similar to theirs: first for any fixed cost vector \mathbf{c}' , divide the region by which item has the largest value $t_i - c'_i$ and use the above flow. Next, define our flow as a distribution of the flow for \mathbf{c}' , over the randomness of \mathbf{c}' . This attempt does take the cost distribution into account. Unfortunately, this flow does not work as the mechanism constructed based on the sampled cost \mathbf{c}' will not represent the seller’s true profit based on \mathbf{c} .

For single buyer, we introduce the following flow. For every $j \in [m]$, let $\bar{v}_j(t_j) = \mathbb{E}_{\mathbf{c}}[(t_j - c_j)^+]$. Define every R_j as follows: R_j contains all types $\mathbf{t} \in T$ such that j is the smallest index among $\text{argmax}_k \bar{v}_k(t_k)$. We route the flow in a similar manner, that is, there is a flow between two nodes $\mathbf{t}, \mathbf{t}' \in R_j$ if they only differ on the j -th coordinate (see Definition 3). Here is the intuition behind our division. Inspired by the canonical flow, we again want to identify the favorite item for the buyer and divide the regions accordingly. However, the favorite item now should be defined as the one that contributes the most to the buyer’s utility instead of the overall welfare. Note that $\bar{v}_j(t_j) = \mathbb{E}_{\mathbf{c}}[(t_j - c_j)^+]$ is exactly the expected utility from item j when the item price is c_j , which is the lowest price that the seller is willing to sell the item. That is why we choose $\bar{v}_j(t_j)$ to represent the contribution of item j to the buyer’s utility. Interestingly, the SPS mechanisms are inspired by our flow, because when there is only one buyer, $\bar{v}_j(t_j)$ can also be viewed as the buyer’s ‘‘value’’ for the j -th permit when the item price $p_j(\mathbf{c}) = c_j$. If we can design a mechanism to extract high revenue from selling the permits, then we have a mechanism that generates high profit. We will make this intuitive connection more concrete in Section 5.2.

4.2.2 Multiple Buyers

Inspired by the single buyer case, we again aim to extract high revenue from selling the permits to make sure our mechanism generates high profit. When there are multiple buyers in the auction, we sell items sequentially to the buyers and our mechanism should satisfy the following two properties:

- The item price should be carefully chosen as the item can not be over-allocated. Usually in the sequential mechanism, the item price should be large enough, to make sure that the item is available to every buyer when she comes to the auction, with certain probability.
- The item price should be at least the seller's cost, to make sure the revenue extracted from selling the items is enough to cover the cost.

Intuitively, how the flow is chosen should also depend on the format of the mechanism we aim to use. To satisfy both properties, we combine our flow in Section 4.2.1 with the ex-ante relaxation technique purposed in [13]. [9] uses the same technique to construct the flow. They divide the type space by comparing the difference between value and the quantile induced from ex-ante allocation probability. Here we involve different quantile thresholds for different cost realization. Furthermore, in order to satisfy the second property, we choose our threshold as the maximum between the quantile and seller's cost.

Definition 1. (Ex-ante relaxation) Fix mechanism $M(\pi, p)$. For every $i \in [n], j \in [m]$ and $\mathbf{c} \in T^S$, define $q_{ij}(\mathbf{c}) = \frac{1}{2} \cdot \mathbb{E}_t[\pi_{ij}(t_i, \mathbf{c})]$, and let

$$\beta_{ij}(\mathbf{c}) = \inf \{a \geq 0 : \Pr[t_{ij} \geq \max\{a, c_j\}] \leq q_{ij}(\mathbf{c})\}$$

$\beta_{ij}(\mathbf{c}) = 0$ if $\Pr[t_{ij} \geq c_j] \leq q_{ij}(\mathbf{c})$. If not, for simplicity we assume that there exists $\beta_{ij}(\mathbf{c})$ such that $\Pr[t_{ij} \geq \beta_{ij}(\mathbf{c})] = q_{ij}(\mathbf{c})$. This is true for continuous distribution D_{ij} . For discrete distributions, our results will hold by dealing with a tie-breaking issue. We refer the readers to Section 5.3 of [9] for more details. In the further proof we will focus on continuous distributions and a same fix will apply for discrete distributions.

We denote β the mappings from \mathbf{c} to $\beta_{ij}(\mathbf{c})$ for all i, j . Before defining the flow, we need the following definition.

Definition 2. Fix β . For every i, t_i and set $P \subseteq [m]$, define

$$\bar{v}_i^{(\beta)}(t_i, P) = \mathbb{E}_{\mathbf{c}} \left[\max_{S \subseteq P, S \in \mathcal{F}_i} \sum_{j \in S} (t_j - \max\{\beta_{ij}(\mathbf{c}), c_j\}) \right]$$

Remark: $\bar{v}_i^{(\beta)}(t_i, P)$ is equal to $\bar{u}_i^p(t_i, P)$ by choosing $p_{ij}(\mathbf{c}) = \max\{\beta_{ij}(\mathbf{c}), c_j\}$.

For notational convenience, let $\bar{v}_{ij}^{(\beta)}(t_{ij}) = \bar{v}_i^{(\beta)}(t_i, \{j\}) = \mathbb{E}_{\mathbf{c}}[(t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\})^+]^5$, which only depends on t_{ij} . It coincides with the definition in Section 4.2.1 with $\beta = 0$.

Now we are ready to define our flow for multiple buyer case.

Definition 3. (Our flow) Fix β . For every $i \in [n], j \in [m]$, $R_{ij}^{(\beta)}$ contains all types $t_i \in T_i$ such that j is the smallest index among $\operatorname{argmax}_k \bar{v}_{ik}^{(\beta)}(t_{ik})$. Define the flow as follows: Each node t_i receives flow of weight $f_i(t_i)$ from the source. For every node $t_i, t'_i \in R_{ij}^{(\beta)}$, $\lambda_i(t'_i, t_i) > 0$ only if $t'_{ik} = t_{ik}$ for all $k \neq j$,

⁵For any value x , denote $x^+ = \max\{x, 0\}$




Figure 1: An example of our flow for two items.

and t'_{ij} is the predecessor type of t_{ij} ⁶. For node $t_i = (t_{ij}, t_{i,-j}) \in R_{ij}^{(\beta)}$, if there does not exist a successor type t'_{ij} of t_{ij} such that $(t'_{ij}, t_{i,-j}) \in R_{ij}^{(\beta)}$, all flow entering node t_i goes to the sink \emptyset . Figure 1 shows an example of our flow for some buyer i when $m = 2$. The curve in the graph contains all (t_{i1}, t_{i2}) such that $\bar{v}_{i1}^{(\beta)}(t_{i1}) = \bar{v}_{i2}^{(\beta)}(t_{i2})$.

Since for all β, i, j , $\bar{v}_{ij}^{(\beta)}(\cdot)$ is non-decreasing, each region $R_{ij}^{(\beta)}$ is upward-closed: for every $t_i = (t_{ij}, t_{i,-j}) \in R_{ij}^{(\beta)}$ and $t'_{ij} > t_{ij}$, $(t'_{ij}, t_{i,-j}) \in R_{ij}^{(\beta)}$. With this property, we have the following Lemma from [7].

Lemma 3. [7] Fix any β . There exists a flow λ such that for every $i \in [n]$, $t_i \in R_{ij}^{(\beta)}$,

$$\Phi_{ik}^{(\lambda)}(t_i) = \begin{cases} t_{ik}, & \text{if } k \neq j \\ \tilde{\varphi}_{ij}(t_{ij}), & \text{if } k = j \end{cases}$$

where $\tilde{\varphi}_{ij}(\cdot)$ is the Myerson's ironed virtual value function w.r.t. D_{ij} .

With Lemma 2 and 3, we have obtained a benchmark for any BIC, interim IR mechanism and divide it into three terms. Note that the benchmark may differ for different mechanisms. The proof of Theorem 4 can be found in Appendix B.

Theorem 4. For any BIC, interim IR mechanism M , let β be the mapping associated with M in Definition 1, then

$$\begin{aligned} \text{PROFIT}(M) &\leq \sum_i \mathbb{E}_{t_i, \mathbf{c}} \left[\sum_j \mathbb{1}[t_i \in R_{ij}^{(\beta)}] \cdot \pi_{ij}(t_i, \mathbf{c}) \cdot (\tilde{\varphi}_{ij}(t_{ij}) - c_j) \right] \\ &\quad + 2 \cdot \sum_i \sum_j \mathbb{E}_{\mathbf{c}} [q_{ij}(\mathbf{c}) \cdot (\max\{\beta_{ij}(\mathbf{c}), c_j\} - c_j)] \\ &\quad + \sum_i \mathbb{E}_{t_i} \left[\sum_j \mathbb{1}[t_i \in R_{ij}^{(\beta)}] \cdot \bar{v}_i^{(\beta)}(t_i, [m] \setminus \{j\}) \right] \end{aligned}$$

⁶In other words, t'_{ij} is the smallest value in the support set T_{ij} that is greater than t_{ij} .

We use $\text{MOST-SURPLUS}(\beta)$, $\text{PROPHET}(\beta)$ and $\text{LESS-SURPLUS}(\beta)$ to denote the three terms accordingly. Note that all three terms depend on β . For the rest of the paper we show that these three terms can be bounded by the profit of simple mechanisms for any β induced by a BIC, interim IR mechanism M . In fact, to prove an approximation of the optimal profit, it is sufficient to consider one specific β induced by the optimal mechanism. In the proof, we fix β and omit it in the notation.

5 Warm-up: Single, Constraint-Additive Buyer

In this section, we bound the benchmark for single, constraint-additive buyer. In this case, all β s are set to 0 and thus $\text{PROPHET} = 0$. We will prove the following theorem.

Theorem 5. *When $n = 1$, for any valuation distribution \mathcal{D} , cost distribution \mathcal{C} and any downward-closed feasibility constraint \mathcal{F} ,*

$$\text{OPT}_{\text{PROFIT}} \leq 2 \cdot \text{IP-PROFIT} + 5 \cdot \text{PP-PROFIT} + 4 \cdot \text{PB-PROFIT}$$

When the buyer's valuation is additive,

$$\text{OPT}_{\text{PROFIT}} \leq \text{IP-PROFIT} + 3 \cdot \text{PP-PROFIT} + 2 \cdot \text{PB-PROFIT}$$

Theorem 5 implies that a simple randomization among the three mechanisms achieves at least $\frac{1}{11}$ the optimal profit for any downward-closed \mathcal{F} . And for additive valuations, a randomization among the three mechanisms is a 6-approximation to the optimal profit.

5.1 Bounding MOST-SURPLUS

To bound MOST-SURPLUS , we will consider the Copies Setting from [11], which is a single-dimensional setting in the revenue maximization problem. Here is a sketch of the proof. For any fixed cost vector \mathbf{c} , we first focus on a related revenue maximization problem with a single buyer and multiple items, by simply subtracting the fixed cost \mathbf{c} from the buyer's value \mathbf{t} . Next, we show that the optimal revenue in the copies setting of the related revenue maximization problem is an upper bound of MOST-SURPLUS . According to [11], there exists a posted price mechanism in the multi-item setting whose revenue approximates the optimal revenue in its copies setting. Finally, we show an Item Posted Pricing mechanism whose expected profit is the same as to the expected revenue of the posted price mechanism.

For any fixed \mathbf{c} , we will first focus on the following revenue maximization problem with a single buyer and m items. Buyer has value $t_j - c_j$ for each item j , where t_j is drawn independently from \mathcal{D}_j . Since \mathbf{c} is a fixed vector, the buyer's values are independent across items. The buyer is constraint-additive with respect to the feasibility constraint \mathcal{F} .

The Copies Setting of the above problem is as follows: there are m buyers in the auction and m copies to sell. Buyer j only interests in the j -th copy and has value $t_j - c_j$ for it, where t_j is drawn independently from \mathcal{D}_j . Since \mathbf{c} is a fixed vector, all buyers' values are also independent. The seller has no cost for the copies but has a downward-closed constraint \mathcal{F} that specifies which copies can simultaneously be sold. Denote $\text{OPT-REV}^{\text{COPIES}}(\mathbf{c})$ the optimal revenue for the copies setting. Since it is a single dimensional setting, Myerson's auction achieves the optimal revenue, which equals to the maximum ironed virtual welfare

$$\text{OPT-REV}^{\text{COPIES}}(\mathbf{c}) = \mathbb{E}[\max_{\mathbf{t}} \sum_{j \in S} (\tilde{\varphi}_j(t_j) - c_j)^+].^7$$

⁷Notice that the ironed Myerson's virtual value for buyer j is $\tilde{\varphi}_j(t_j) - c_j$.

Moreover, let $\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})$ be the optimal revenue if we further restrict the seller to sell at most one copy. Similarly $\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c}) = \mathbb{E}_{\mathbf{t}}[\max_j(\tilde{\varphi}_j(t_j) - c_j)^+]$. We have the following lemma.

Lemma 4. $\text{MOST-SURPLUS} \leq \mathbb{E}_{v,c}[\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})] \leq 2 \cdot \text{IP-PROFIT}$. When the buyer is additive, we further have $\text{MOST-SURPLUS} \leq \text{IP-PROFIT}$.

Proof. We first prove the result for arbitrary downward-closed constraint \mathcal{F} . Notice that for every \mathbf{t} , the indicator $\mathbb{1}[\mathbf{t} \in R_j]$ is 1 for only one j . Since $\pi_j(\mathbf{t}, \mathbf{c}) \in [0, 1]$, we have

$$\begin{aligned} \text{MOST-SURPLUS} &= \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\sum_j \mathbb{1}[\mathbf{t} \in R_j] \cdot \pi_j(\mathbf{t}, \mathbf{c}) \cdot (\tilde{\varphi}_j(t_j) - c_j) \right] \\ &\leq \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\max_j (\tilde{\varphi}_j(t_j) - c_j)^+ \right] = \mathbb{E}_{\mathbf{c}} \left[\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c}) \right] \end{aligned}$$

By [11], there exists a posted price mechanism $M(\mathbf{c})$ in the revenue maximization problem whose revenue is at least $\frac{1}{2} \text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})$. Let $\hat{p}_j(\mathbf{c})$ be the posted price for item j .

Now we move back to our profit maximization setting and define the IP mechanism M' as follows: For every cost vector \mathbf{c} , define the posted price for item j as $\hat{p}_j(\mathbf{c}) + c_j$. Notice that for every \mathbf{t} and \mathbf{c} , the buyer in M' will purchase the same bundle $B^*(\mathbf{t}, \mathbf{c})$ as the one in $M(\mathbf{c})$. Here $B^*(\mathbf{t}, \mathbf{c}) = \arg\max_{S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j - \hat{p}_j(\mathbf{c}))$. Thus the seller's profit of M' is

$$\begin{aligned} \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\sum_{j \in B^*(\mathbf{t}, \mathbf{c})} (\hat{p}_j(\mathbf{c}) + c_j - c_j) \right] &= \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\sum_{j \in B^*(\mathbf{t}, \mathbf{c})} \hat{p}_j(\mathbf{c}) \right] \\ &\geq \frac{1}{2} \mathbb{E}_{\mathbf{c}} \left[\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c}) \right] \geq \frac{1}{2} \cdot \text{MOST-SURPLUS} \end{aligned}$$

When the buyer is additive, for any fixed \mathbf{c} , it is not hard to realize that $\text{OPT-REV}^{\text{COPIES}}(\mathbf{c})$ equals to the revenue of selling each item separately using the monopoly reserve in the revenue maximization problem. Let $\hat{p}_j(\mathbf{c})$ be the monopoly reserve for item j in the revenue maximization problem. Following the same proof as above, the IP mechanism M' with price $\hat{p}_j(\mathbf{c}) + c_j$ achieves expected profit at least

$$\begin{aligned} \mathbb{E}_{\mathbf{c}}[\text{OPT-REV}^{\text{COPIES}}(\mathbf{c})] &= \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\sum_j (\tilde{\varphi}_j(t_j) - c_j)^+ \right] \\ &\geq \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\sum_j \mathbb{1}[\mathbf{t} \in R_j] \cdot \pi_j(\mathbf{t}, \mathbf{c}) \cdot (\tilde{\varphi}_j(t_j) - c_j) \right] = \text{MOST-SURPLUS} \end{aligned}$$

□

5.2 Bounding LESS-SURPLUS

Before bounding LESS-SURPLUS, we will first prove a crucial lemma of this section. Consider the revenue maximization problem with a single buyer and m items. The buyer's type $\mathbf{t} \sim \mathcal{D}$. She has valuation function $\bar{v}(\mathbf{t}, \cdot)$ when her type is \mathbf{t} . For simplicity, we will call this revenue maximization problem *the revenue setting*, and the original profit maximization problem *the profit setting*. Recall that in the single buyer case,

$$\bar{v}(\mathbf{t}, P) = \mathbb{E}_{\mathbf{c}} \left[\max_{S \subseteq P, S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j) \right]$$

The following lemma converts any truthful mechanism in the revenue setting into a BIC and interim IR mechanism in the profit setting, without changing the value of the objective (revenue and profit accordingly). The intuition behind the lemma is as follows. For any mechanism in the profit setting that sells the permit before revealing her true cost, the buyer with type \mathbf{t} has expected “value” $\bar{v}(\mathbf{t}, P)$, that is, how much the buyer can make from the second stage if given a set of permits P , for all set of permits P . Thus, the mechanism can be viewed as a corresponding mechanism in the revenue setting where the permits are being sold and the buyer has valuation \bar{v} over the permits.

Lemma 5. *For any truthful mechanism M in the revenue setting, there exists an IC and IR mechanism M' in the profit setting such that, the revenue of M equals to the seller's profit of M' .*

Proof. For any \mathbf{t} , let $X(\mathbf{t})$ be the (possibly random) set of items that the buyer is allocated in mechanism M , when the buyer reports \mathbf{t} . Let $p(\mathbf{t})$ be the payment for the buyer in M . Define M' as follows: in the first stage, the buyer reports her type \mathbf{t} and the seller gives the set of permits $X(\mathbf{t})$ to the buyer and charge $p(\mathbf{t})$. In the second stage, the seller reveals the cost vector \mathbf{c} and the buyer can buy any item that she has a permit by paying item price c_j . To prove M' is an IC and IR mechanism, it suffices to show that the buyer has no incentive to lie in the first stage. If the buyer with type \mathbf{t} reports \mathbf{t}' in M' , she will receive the set of permits $X(\mathbf{t}')$ and purchase her favorite bundle of items under item prices \mathbf{c} . Her expected utility is

$$\mathbb{E}_{\mathbf{c}, X(\mathbf{t}')} \left[\max_{S \in \mathcal{F}, S \subseteq X(\mathbf{t}')} \sum_{j \in S} (t_j - c_j) \right] - p(\mathbf{t}') = \mathbb{E}_{X(\mathbf{t}')} [\bar{v}(\mathbf{t}, X(\mathbf{t}'))] - p(\mathbf{t}')$$

Here the expectation is taken over the randomness of $X(\mathbf{t}')$. Since M is truthful, for any $\mathbf{t} \in T, \mathbf{t}' \in T^+$ ⁸, $\mathbb{E}_{X(\mathbf{t})} [\bar{v}(\mathbf{t}, X(\mathbf{t}))] - p(\mathbf{t}) \geq \mathbb{E}_{X(\mathbf{t}')} [\bar{v}(\mathbf{t}, X(\mathbf{t}'))] - p(\mathbf{t}')$. It states that when the buyer has type \mathbf{t} , reporting \mathbf{t} in the first stage maximizes her expected utility. Thus M' is IC and IR. Notice that in the second stage of M' , the total item prices paid by the buyer is equal to the seller's total cost. Thus the seller's profit is exactly the payment in the first stage. Since M' use $p(\mathbf{t})$ as the payment rule, the seller's profit of M' equals to the revenue of M . \square

Now we are ready to bound the term LESS-SURPLUS. Recall that

$$\text{LESS-SURPLUS} = \mathbb{E}_{\mathbf{t}} \left[\sum_j \mathbf{1}[\mathbf{t} \in R_j] \cdot \bar{v}(\mathbf{t}, [n] \setminus \{j\}) \right].$$

Consider the revenue setting where the buyer has valuation function \bar{v} and let $\text{OPT}_{\text{REV}}(\bar{v})$ be the optimal revenue among all truthful mechanisms. Here we omit \mathcal{D} and \mathcal{C} in the notation as they are fixed. Given Lemma 5, it is tempting to find a simple mechanism that approximates $\text{OPT}_{\text{REV}}(\bar{v})$ and convert it into a permit-selling mechanism. However, since we do not know what class of valuation \bar{v} belongs to, it not a priori clear any simple vs. optimal result applies here. As the original valuation in the profit setting is constrained additive, it is natural to think that \bar{v} is also constrained additive. Unfortunately, we are not able to prove such a claim as there is no clear feasibility that is associated with \bar{v} . The good news is that we are able to relax the class of valuations and show that \bar{v} is indeed a *subadditive function*, which allows us to leverage the result by [29, 9].

Let us first review their results. They bound $\text{OPT}_{\text{REV}}(\bar{v})$ when \bar{v} is *subadditive over independent items* (see Definition 4). In both proofs, they separate the benchmark of the optimal revenue into two terms (called “single” and “non-favorite” in [9]). They then bound the two terms by the optimal revenue of the Selling

⁸Recall that $T^+ = T \cup \{\emptyset\}$ contains the choice of not attending the auction.

Separately mechanism($\text{SREV}(\bar{v})$) and Bundling mechanism($\text{BREV}(\bar{v})$) respectively. The second term “non-favorite” is defined as the expected welfare from all non-favorite items. Here for any fixed \mathbf{t} , the favorite item is defined as the j that maximizes $\bar{v}(\mathbf{t}, \{j\})$. Interestingly, this is how we divide the region into R_j s and the term “non-favorite” is exactly the same as LESS-SURPLUS. We will use $\text{NON-FAV}_{\text{REV}}(\bar{v})$ to denote “non-favorite” here to emphasize that it is from the revenue setting.

In order to apply the result in the revenue setting, we first show that the function $\bar{v}(\cdot, \cdot)$ in Definition 2 is indeed subadditive over independent items. The proof of Lemma 6 is postponed to Appendix C.

Definition 4. [29] Suppose the buyer’s type \mathbf{t} is drawn from a product distribution $D = \prod_j D_j$, her distribution \mathcal{V} of valuation function $v(\mathbf{t}, \cdot)$ is **subadditive over independent items** if:

- $v(\cdot, \cdot)$ **has no externalities**, i.e., for each $\mathbf{t} \in T$ and $S \subseteq [m]$, $v(\mathbf{t}, S)$ only depends on $\langle t_j \rangle_{j \in S}$, formally, for any $\mathbf{t}' \in T$ such that $t'_j = t_j$ for all $j \in S$, $v(\mathbf{t}', S) = v(\mathbf{t}, S)$.
- $v(\cdot, \cdot)$ **is monotone**, i.e., for all $\mathbf{t} \in T$ and $U \subseteq V \subseteq [m]$, $v(\mathbf{t}, U) \leq v(\mathbf{t}, V)$.
- $v(\cdot, \cdot)$ **is subadditive**, i.e., for all $\mathbf{t} \in T$ and $U, V \subseteq [m]$, $v(\mathbf{t}, U \cup V) \leq v(\mathbf{t}, U) + v(\mathbf{t}, V)$.

Lemma 6. $\bar{v}(\cdot, \cdot)$ is monotone, subadditive and has no externalities.

Lemma 7. [9] Suppose \bar{v} is subadditive over independent items, then

$$\text{NON-FAV}_{\text{REV}}(\bar{v}) = \mathbb{E}_{\mathbf{t}} \left[\sum_j \mathbb{1}[\mathbf{t} \in R_j] \cdot \bar{v}(\mathbf{t}, [n] \setminus \{j\}) \right] \leq 5 \cdot \text{SREV}(\bar{v}) + 4 \cdot \text{BREV}(\bar{v})$$

We only provide a sketch of their analysis here and refer the readers to their paper for a formal proof. They first use the Core-Tail Decomposition technique and divide $\text{NON-FAV}_{\text{REV}}(\bar{v})$ into two terms $\text{TAIL}(\bar{v})$ and $\text{CORE}(\bar{v})$. The term $\text{TAIL}(\bar{v})$ can be bounded using $\text{SREV}(\bar{v})$. $\text{CORE}(\bar{v})$ can be viewed as the buyer’s expected welfare under a truncated valuation. They show that the welfare concentrates via a concentration inequality for subadditive functions[30]. In particular, they show that the median of the welfare is comparable to the mean of the welfare. Thus, a grand bundling mechanism that uses the median as the price is able to approximate $\text{CORE}(\bar{v})$.

When the buyer is additive, [7] has an improved bound for $\text{NON-FAV}_{\text{REV}}(\bar{v})$ using $\text{SREV}(\bar{v})$ and $\text{BREV}(\bar{v})$.

Lemma 8. [7] If \bar{v} is an additive function, then

$$\text{NON-FAV}_{\text{REV}}(\bar{v}) \leq 2 \cdot \text{SREV}(\bar{v}) + 3 \cdot \text{BREV}(\bar{v})$$

By Lemma 5, the Selling Separately mechanism in the revenue setting can be converted to the PP mechanism in the profit setting and has profit equals to $\text{SREV}(\bar{v})$. Also the Bundling mechanism can be converted to PB and obtains profit $\text{BREV}(\bar{v})$. Furthermore, when the buyer is additive, there is no constraint \mathcal{F} and for every $\mathbf{t} \in T$, $P \subseteq [m]$,

$$\bar{v}(\mathbf{t}, P) = \mathbb{E}_{\mathbf{c}} \left[\max_{S \subseteq P} \sum_{j \in S} (t_j - c_j) \right] = \mathbb{E}_{\mathbf{c}} \left[\sum_{j \in P} (t_j - c_j)^+ \right] = \sum_{j \in P} \bar{v}(\mathbf{t}, \{j\}).$$

Thus, \bar{v} is an additive function. We have the following Corollary:

Corollary 1. $\text{LESS-SURPLUS} \leq 5 \cdot \text{PP-PROFIT} + 4 \cdot \text{PB-PROFIT}$. When the buyer is additive, $\text{LESS-SURPLUS} \leq 2 \cdot \text{PP-PROFIT} + 3 \cdot \text{PB-PROFIT}$.

Proof of Theorem 5: It follows from Theorem 4, Lemma 4 and Corollary 1. \square

At the last of this section, we will prove the following lemma that connects the profit maximization problem to the revenue maximization problem. We show that any truthful mechanism in the revenue setting that is an α -approximation to the optimal revenue can be converted to a BIC and interim IR mechanism in the profit setting that is a $(9\alpha + 2)$ -approximation to the optimal profit.

Lemma 9. *Recall that the revenue setting is the revenue maximization problem where the buyer has valuation \bar{v} . Then any truthful mechanism in the revenue setting that is an α -approximation to the optimal revenue $\text{OPT}_{\text{REV}}(\bar{v})$ can be converted to an IC and IR mechanism in the profit setting that is a $(9\alpha + 2)$ -approximation to the optimal profit $\text{OPT}_{\text{PROFIT}}$.*

Proof. By Lemma [9],

$$\text{NON-FAV}_{\text{REV}}(\bar{v}) \leq 5 \cdot \text{SREV}(\bar{v}) + 4 \cdot \text{BREV}(\bar{v}) \leq 9 \cdot \text{OPT}_{\text{REV}}(\bar{v})$$

Let M be the α -approximation mechanism in the revenue setting. Then the revenue of M satisfies:

$$\text{REV}(M) \geq \frac{1}{\alpha} \cdot \text{OPT}_{\text{REV}}(\bar{v}) \geq \frac{1}{9\alpha} \cdot \text{NON-FAV}_{\text{REV}}(\bar{v})$$

By Lemma 5, there exists a BIC and interim IR mechanism M' in the profit setting such that $\text{PROFIT}(M') = \text{REV}(M)$. Notice that $\text{LESS-SURPLUS} = \text{NON-FAV}_{\text{REV}}(\bar{v})$. By Theorem 4,

$$\text{OPT}_{\text{PROFIT}} \leq \text{MOST-SURPLUS} + \text{LESS-SURPLUS} \leq 2 \cdot \text{IP-PROFIT} + 9\alpha \cdot \text{PROFIT}(M')$$

Thus a randomization between M' and the optimal Item Posted Pricing mechanism is a $(9\alpha + 2)$ -approximation. \square

6 Multiple, Matroid-Rank Buyers

In this section, we will bound the benchmark in Theorem 4 for multiple, matroid-rank buyers, using the mechanisms described in Section 2.

Theorem 6. *For any valuation distribution \mathcal{D} , cost distribution \mathcal{C} and any matroid feasibility constraints $\{\mathcal{F}_i\}_{i=1}^m$,*

$$\text{OPT}_{\text{PROFIT}} \leq 14 \cdot \text{CSIP-PROFIT} + 22 \cdot \text{RSPP-PROFIT} + 8 \cdot \text{SPB-PROFIT}$$

Again a simple randomization among the three mechanisms achieves at least $\frac{1}{44}$ the optimal profit.

6.1 Bounding MOST-SURPLUS

Similar to the single buyer case, for every fixed vector \mathbf{c} , we consider the related revenue maximization problem where every buyer i has value $t_{ij} - c_j$ for item j . The corresponding Copies Setting is a revenue maximization problem with mn buyers and m items. Every buyer (i, j) only interests in item j and has value $t_{ij} - c_j$ on it. For every i , at most one (i, j) can be served in the mechanism. We denote $\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})$ the optimal revenue of this setting. In Lemma 10 we first bound MOST-SURPLUS by $\mathbb{E}_{\mathbf{c}}[\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})]$. Then according to [11], $\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})$ can be approximated by the revenue of the optimal sequential posted price mechanism in the related revenue maximization setting. Assume the posted price for buyer i and item j is $\hat{p}_{ij}(\mathbf{c})$. We show that in our setting, an SIP mechanism with $p_{ij}(\mathbf{c}) = \hat{p}_{ij}(\mathbf{c}) + c_j$ has profit the same as the expected (over the randomness of \mathbf{c}) revenue of the above sequential posted price mechanism.

Lemma 10. $\text{MOST-SURPLUS} \leq \mathbb{E}_{\mathbf{c}}[\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})] \leq 6 \cdot \text{CSIP-PROFIT}$.

Proof. Recall that

$$\text{MOST-SURPLUS} = \sum_i \mathbb{E}_{t_i, \mathbf{c}} \left[\sum_j \mathbb{1}[t_i \in R_{ij}^{(\beta)}] \cdot \pi_{ij}(t_i, \mathbf{c}) \cdot (\tilde{\varphi}_{ij}(t_{ij}) - c_j) \right]$$

For every BIC, interim IR mechanism $M = (\pi, p)$ and every \mathbf{c} , consider the following mechanism M' in the Copies Setting: M' serves agent (i, j) if and only if M allocates item j to buyer i and $t_i \in R_{ij}^{(\beta)}$. Since M is feasible, for every j there exist at most one i such that (i, j) is served in M' . Also since every t_i stays in one region, for every i there exists at most one j such that (i, j) is served in M' . Thus M' is feasible and the expected revenue equals to $\sum_i \mathbb{E}_{t_i} \left[\sum_j \mathbb{1}[t_i \in R_{ij}^{(\beta)}] \cdot \pi_{ij}(t_i, \mathbf{c}) \cdot (\tilde{\varphi}_{ij}(t_{ij}) - c_j) \right]$. Thus we have

$$\text{MOST-SURPLUS} \leq \mathbb{E}_{\mathbf{c}}[\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})]$$

By [11], for every \mathbf{c} there exists a sequential posted price mechanism $M(\mathbf{c})$ in the related revenue maximization setting⁹, where every buyer can purchase at most one item, such that its revenue is at least $\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})/6$. Suppose the posted price for buyer i and item j is $\hat{p}_{ij}(\mathbf{c})$. Now let's consider the Constrained Sequential Item Posted Price mechanism with $p_{ij}(\mathbf{c}) = \hat{p}_{ij}(\mathbf{c}) + c_j$ in our profit maximization setting, where every buyer is only allowed to purchase at most one item. For every \mathbf{t}, \mathbf{c} , let $A_i(t_{<i}, \mathbf{c})$ be the remaining item sets when buyer i comes to the auction. Then she will choose her favorite item $\text{argmax}_{j \in A_i(t_{<i}, \mathbf{c})} (t_{ij} - c_j - \hat{p}_{ij}(\mathbf{c}))$ (or choose not to purchase anything). Notice that this is also buyer i 's favorite item in $M(\mathbf{c})$ under the same scenario. Thus the allocation rule for CSIP under \mathbf{c} is the same as the one for $M(\mathbf{c})$. Then profit of the constructed CSIP is equal to the expected revenue of $M(\mathbf{c})$ over the randomness of \mathbf{c} , as the extra item prices just cover the seller's costs. According to [11], the profit is at least $\mathbb{E}_{\mathbf{c}}[\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})]/6$. The proof is done. \square

6.2 Bounding PROPHET

In this section we will bound PROPHET with a Constrained Sequential Item Posted Price mechanism. The proof uses *Online Contention Resolution Scheme(OCRS)* developed by Feldman et al. [23]. It is defined under an online selection problem. For simplicity, we will just describe the setting considered in our proof. Each element e in an ground set is revealed one by one, and an agent has to make a decision whether to take an element before the next one is revealed. In our proof the ground set is J , the set containing all buyer-item pairs. And each element $e = (i, j)$ is one of the buyer-item pair. The agent can only take a feasible set of elements subject to the feasibility constraint \mathcal{J} . The set of *active* elements is random, and the agent can only take the active elements. Let $R(y)$ be the random set of active elements, where vector $y = (y_e)_{e \in J}$. For each $e \in J$, each element is chosen as active or inactive independently and y_e is the probability of e being active. y stays in the polytope $P_{\mathcal{J}}$ corresponding to \mathcal{J} :

$$P_{\mathcal{J}} = \text{conv}(\{\mathbb{1}_A | A \in \mathcal{J}\})$$

An OCRS for $P_{\mathcal{J}}$ is an online algorithm that selects a feasible and active set: $A \subseteq R(y)$ and $A \in \mathcal{J}$. Specifically, a greedy OCRS Π uses a greedy scheme during the selection: for each $y \in P_{\mathcal{J}}$, it determines a subfamily $\mathcal{J}_{\Pi, y} \subseteq \mathcal{J}$. It selects an element e when e arrives if, the set of elements selected still stays in $\mathcal{J}_{\Pi, y}$ after picking e .

In the proof we will mainly consider the following property of a greedy OCRS Π called *selectability*. It is first defined in [23].

⁹Recall that in this setting every buyer i has value $t_{ij} - c_j$ for item j .

Definition 5. [23] Let $b, c \in [0, 1]$. A greedy OCRS Π is (b, c) -selectable if for every $y \in b \cdot P_{\mathcal{J}}$, $e \in J$

$$\Pr[A \cup \{e\} \in \mathcal{J}_{\Pi, y}, \forall A \subseteq R(y), A \in \mathcal{J}_{\Pi, y}] \geq c.$$

The above definition can be described as follows. With probability at least c , over the randomness of $R(y)$, for any subset A of the active elements that is feasible w.r.t. the subfamily $\mathcal{J}_{\Pi, y}$, we can add element e to the set without violating the feasibility. View A as the set of elements being served by the greedy OCRS Π . When e arrives, (b, c) -selectability guarantees that not matter what set of elements has been chosen already (still a subset of A), with probability at least c , adding element e is still feasible. By Definition 5, one can show that by following the OCRS, the agent will select each element e with probability at least $c \cdot y_e$, under any almighty adversary¹⁰.

Lemma 11. ([23]) Consider the online selection setting described above. If there exists a (b, c) -selectable greedy OCRS Π for $P_{\mathcal{J}}$, then for every $y \in b \cdot P_{\mathcal{J}}$, consider the strategy that the agent takes elements greedily subject to the sub-constraint $\mathcal{J}_{\Pi, y}$. Then the agent will select each element e with probability at least $c \cdot y_e$. The result applies for any almighty adversary.

Before getting to the proof, let's first discuss the connection between OCRS and bounding PROPHET. Recall that

$$\text{PROPHET} = 2 \cdot \sum_i \sum_j \mathbb{E}_{\mathbf{c}} [q_{ij}(\mathbf{c}) \cdot (\max\{\beta_{ij}(\mathbf{c}), c_j\} - c_j)]$$

Fix \mathbf{c} . For every $(i, j) \in J$, let $y_{(i, j)} = q_{ij}(\mathbf{c})$. Since $q_{ij}(\mathbf{c})$ is half the ex-ante probability that a feasible mechanism M serves the pair (i, j) when the true cost is \mathbf{c} , thus $y = (y_{(i, j)})_{(i, j) \in J} \in \frac{1}{2} \cdot \mathcal{J}$. Now consider the CSIP with item prices $\max\{\beta_{ij}(\mathbf{c}), c_j\}$. Then by Definition 1, each buyer i can afford item j with probability $q_{ij}(\mathbf{c})$ ¹¹, i.e. the element (i, j) is active with probability $q_{ij}(\mathbf{c})$. In Lemma 12 we show that for one specific almighty adversary, the set of element (i, j) chosen by the agent following the greedy OCRS is exactly same as the set of buyer-item pair served in the mechanism, for every type profile. Then (b, c) -selectability guarantees that every buyer i will purchase every item j in the mechanism with probability at least c given the fact that she can afford this item. This gives a lower bound of CSIP-PROFIT.

Lemma 12. Fix seller's cost vector \mathbf{c} . Suppose there exists a (b, c) -selectable greedy OCRS Π for polytope $P(\mathcal{J})$, for some constant $c \in (0, 1)$. For every l_{ij} s such that $l \in b \cdot P(\mathcal{J})$, consider the CSIP under the specific cost profile \mathbf{c} , with posted price $p_{ij}(\mathbf{c}) = F_{ij}^{-1}(1 - l_{ij})$ and sub-constraint $\mathcal{J}_{\Pi, l}$. Then the mechanism will gain profit at least

$$c \cdot \sum_{i, j} l_{ij} \cdot (p_{ij}(\mathbf{c}) - c_j)$$

under cost \mathbf{c} .

Proof. Under cost \mathbf{c} , consider the CSIP with posted price $p_{ij}(\mathbf{c})$, associated with the constraint $\mathcal{J}_{\Pi, y}$. When every buyer i comes, let A_i be the set of buyer-item pairs that have already been served. And let B_i^* be her favorite bundle among the remaining items, such that after taking those items, the sub-constraint $\mathcal{J}_{\Pi, y}$ is not violated. Now consider the online selection setting with the following almighty adversary: Ground set is

¹⁰The adversary can determine the order of elements shown to the agent. An almighty adversary has all the information it needs to decide the order, including the agent's type and strategies, and the realization of all possible randomness. In other words, the adversary will choose the worst order for the agent.

¹¹It's true when $\beta_{ij}(\mathbf{c}) \geq c_j$. For those (i, j) such that $\beta_{ij}(\mathbf{c}) < c_j$, the corresponding term in PROPHET is 0. We could simply never serve those pairs.

J . The agent has value t_{ij} for each element (i, j) . Each element (i, j) is active if $t_i \geq p_{ij}(\mathbf{c})$, i.e. is active with probability l_{ij} . The adversary divides the whole item-revealing process into n stages. For each stage i , let B_i be the set of elements that have been selected in the past. The adversary first reveals all (i, j) s where $j \in B_i^*$, one after another. Then it reveals the remaining (i, j) s.

Notice that by following the greedy OCRS Π , the agent will follow the constraint $\mathcal{J}_{\Pi, l}$ and choose all the element (i, j) where $j \in B_i^*$ on each stage, as taking those elements won't violate the constraint by the definition of B_i^* . It's equivalent to the buyer-item pair selection process in the CSIP. Thus under the above adversary, the set of element (i, j) chosen by the agent is exactly same as the set of buyer-item pair served in the mechanism. Since each element (i, j) is active with probability l_{ij} and $l \in b \cdot P(\mathcal{J})$, by Lemma 11, each element is chosen by the agent with probability at least $c \cdot l_{ij}$. In other words, in CSIP, each buyer i purchases item j with probability at least $c \cdot l_{ij}$, under cost \mathbf{c} . Thus the obtained profit is at least

$$c \cdot \sum_{i,j} l_{ij} \cdot (p_{ij} - c_j)$$

□

Now it's sufficient to show that there exists a $(\frac{1}{2}, c)$ -selectable greedy OCRS Π for $P_{\mathcal{J}}$. Recall that every $A \in \mathcal{J}$ satisfies:

- Each item is allocated to at most one buyer: $\forall j, O_j = \{i : (i, j) \in A\}, |O_j| \leq 1$.
- Each buyer is allocated a feasible set of items: $\forall i, P_i = \{j : (i, j) \in A\}, P_i \in \mathcal{F}_i$.

Let \mathcal{J}_1 (or \mathcal{J}_2) be the subfamily that contains all set A that satisfies the first(or second) bullet point. It's straightforward to see that \mathcal{J}_1 forms a partition matroid. We will show that \mathcal{J}_2 is also a matroid, given the fact that every \mathcal{F}_i is a matroid.

Lemma 13. \mathcal{J}_2 is a matroid.

Proof. Consider any $A, A' \in \mathcal{J}_2$ such that $|A| > |A'|$. For every i , let $P_i = \{j : (i, j) \in A\}$ and $P'_i = \{j : (i, j) \in A'\}$. We have $P_i \in P'_i \in \mathcal{F}_i$. Notice that $|A| = \sum_i |P_i| > |A'| = \sum_i |P'_i|$, there must exist i_0 such that $|P_{i_0}| > |P'_{i_0}|$. Since \mathcal{F}_{i_0} is a matroid, there exists some $j_0 \in P_{i_0} \setminus P'_{i_0}$ such that $P'_{i_0} \cup \{j_0\} \in \mathcal{F}_{i_0}$. By definition of \mathcal{J}_2 , we also have $(i_0, j_0) \in A \setminus A'$ and $A' \cup (i_0, j_0) \in \mathcal{J}_2$. Thus \mathcal{J}_2 is a matroid. □

Note that $\mathcal{J} = \mathcal{J}_1 \cap \mathcal{J}_2$. \mathcal{J} is an intersection of two matroids. We can show that there exists a $(\frac{1}{2}, \frac{1}{4})$ -selectable greedy OCRS for $P(\mathcal{J})$ by the following two facts from [23].

Lemma 14. ([23]) For every $b \in [0, 1]$, there exists a $(b, 1 - b)$ -selectable greedy OCRS for matroid polytopes.

Lemma 15. ([23]) Suppose there exists a (b, c^1) -selectable greedy OCRS for $P(\mathcal{J}_1)$ and a (b, c^2) -selectable greedy OCRS for $P(\mathcal{J}_2)$. Then there exists a $(b, c^1 \cdot c^2)$ -selectable greedy OCRS for $P(\mathcal{J}_1) \cap P(\mathcal{J}_2)$. Moreover, since $P(\mathcal{J}_1 \cap \mathcal{J}_2) \subseteq P(\mathcal{J}_1) \cap P(\mathcal{J}_2)$, there exists a $(b, c^1 \cdot c^2)$ -selectable greedy OCRS for $P(\mathcal{J}_1 \cap \mathcal{J}_2)$.

Put everything together, we are able to bound PROPHET using CSIP-PROFIT.

Lemma 16. PROPHET $\leq 8 \cdot$ CSIP-PROFIT.

Proof. First for those (i, j) such that $\beta_{ij}(\mathbf{c}) < c_j$, the corresponding term in PROPHET is 0. We could simply never serve those pairs. Thus without loss of generality, we assume that $\beta_{ij}(\mathbf{c}) \geq c_j$ for every (i, j) . For every \mathbf{c} , since $q_{ij}(\mathbf{c})$ is half the ex-ante probability that a feasible mechanism M serves the pair (i, j) when the true cost is \mathbf{c} , thus $q(\mathbf{c}) = (q_{ij}(\mathbf{c}))_{(i,j) \in J} \in \frac{1}{2} \cdot \mathcal{J}$. By Lemma 14 and 15, there exists a $(\frac{1}{2}, \frac{1}{4})$ -selectable greedy OCRS Π for $P(\mathcal{J})$. We thus consider the CSIP with posted price $p_{ij}(\mathbf{c}) = \max\{\beta_{ij}(\mathbf{c}), c_j\}$, associated with the constraint $\mathcal{J}_{\Pi, q(\mathbf{c})}$. By Lemma 12, the profit of the mechanism is at least

$$\frac{1}{4} \cdot \sum_i \sum_j \mathbb{E}_{\mathbf{c}} [q_{ij}(\mathbf{c}) \cdot (\max\{\beta_{ij}(\mathbf{c}), c_j\} - c_j)]$$

□

6.3 Bounding LESS-SURPLUS

In this Section we will bound LESS-SURPLUS. As discussed in Section 4, we will fix β and omit it in the notation. We first give an informal proof by reducing the multiple buyer problem to single buyer problems with a new valuation \bar{v}_i in Definition 2. This shows a connection to the single buyer setting as well as the ex-ante relaxation by Chawla and Miller[13]. In their paper they solve the revenue maximization problem for multiple matroid-rank buyers, bounding the benchmark by the sum of optimal revenue for the single buyer problem under an ex-ante constraint.

Recall that

$$\text{LESS-SURPLUS} = \sum_i \mathbb{E}_{t_i} \left[\sum_j \mathbb{1}[t_i \in R_{ij}] \cdot \bar{v}_i(t_i, [m] \setminus \{j\}) \right]$$

This is the sum of all buyer's welfare contributed by those non-favorite "items"¹², under a new valuation \bar{v}_i . Consider the Sequential Permit Selling mechanism with posted price $p_{ij}(\mathbf{c}) = \max\{\beta_{ij}(\mathbf{c}), c_j\}$. Since $p_{ij}(\mathbf{c}) \geq c_j$ for every \mathbf{c} , the profit of the mechanism will be at least the revenue extracted from the permit copies.

Notice that for every i, j, \mathbf{c} , buyer i can afford the item price for j with probability $q_{ij}(\mathbf{c})$. Thus by union bound, each item j is still available when buyer i comes with probability at least $1 - \sum_j q_{ij}(\mathbf{c}) \geq \frac{1}{2}$. By Lemma 23, buyer i 's expected utility in the second stage after purchasing a set of permit copies P , is at least $\frac{1}{2} \cdot \bar{v}_i(t_i, P)$. Now we have reduced the multiple buyer problem to n single buyer problems where buyer i has valuation $\bar{v}_i(\cdot, \cdot)$ for the set of permit copies. Thus from Section 5, the term $\mathbb{E}_{t_i} \left[\sum_j \mathbb{1}[t_i \in R_{ij}] \cdot \bar{v}_i(t_i, [m] \setminus \{j\}) \right]$ can be extracted from selling the permit copies separately and as a whole bundle.

The above argument doesn't give a formal proof because the buyer's expected utility on the copies does not exactly equal to $\frac{1}{2} \cdot \bar{v}_i(t_i, P)$ and thus a reduction like Lemma 5 cannot be directly obtained. Now we provide a formal and separate proof, bounding LESS-SURPLUS with RSPP and SPB mechanisms. First we decompose the term using a standard Core-Tail decomposition technique [28, 7], according to $\bar{v}_{ij}(t_{ij})$. For every $i \in [n]$, define $\tau_i = \inf\{a \geq 0 : \sum_j \Pr_{t_{ij}}[\bar{v}_{ij}(t_{ij}) \geq a] \leq \frac{1}{2}\}$. For every t_i , let $C_i(t_i) = \{j \in [m] : \bar{v}_{ij}(t_{ij}) \leq \tau_i\}$.

¹²Here we use quotations on the word 'item' as in the corresponding single buyer problem, the goods sold to the buyers are permits, not real items.

Lemma 17.

$$\begin{aligned} \text{LESS-SURPLUS} &\leq \sum_i \sum_j \mathbb{E}_{t_{ij}:\bar{v}_{ij}(t_{ij})>\tau_i} [\bar{v}_{ij}(t_{ij}) \cdot \Pr_{t_{i,-j}} [\exists k \neq j \text{ s.t. } \bar{v}_{ik}(t_{ik}) \geq \bar{v}_{ij}(t_{ij})]] \quad (\text{TAIL}) \\ &\quad + \sum_i \mathbb{E}_{t_i} [\bar{v}_i(t_i, C_i(t_i))] \quad (\text{CORE}) \end{aligned}$$

Proof.

$$\begin{aligned} \text{LESS-SURPLUS} &= \sum_i \mathbb{E}_{t_i} \left[\sum_j \mathbb{1}[t_i \in R_{ij}] \cdot \bar{v}_i(t_i, [m] \setminus \{j\}) \right] \\ &\leq \sum_i \mathbb{E}_{t_i} \left[\sum_j \mathbb{1}[t_i \in R_{ij}] \cdot (\bar{v}_i(t_i, C_i(t_i) \setminus \{j\}) + \bar{v}_i(t_i, [m] \setminus \{j\} \setminus C_i(t_i))) \right] \\ &\leq \sum_i \mathbb{E}_{t_i} [\bar{v}_i(t_i, C_i(t_i))] + \sum_i \mathbb{E}_{t_i} \left[\sum_j \mathbb{1}[t_i \in R_{ij}] \cdot \sum_{k \in [m] \setminus \{j\} \setminus C_i(t_i)} \bar{v}_{ik}(t_{ik}) \right] \\ &= \sum_i \mathbb{E}_{t_i} [\bar{v}_i(t_i, C_i(t_i))] + \sum_i \sum_k \mathbb{E}_{t_{ik}:\bar{v}_{ik}(t_{ik})>\tau_i} [\bar{v}_{ik}(t_{ik}) \cdot \Pr_{t_{i,-k}} [(t_{ik}, t_{i,-k}) \notin R_{ij}]] \\ &= \text{CORE} + \text{TAIL} \end{aligned}$$

□

6.3.1 TAIL

We will bound TAIL using RSPP mechanisms. For every i, j , let $r_{ij} = \max_{a \geq \tau_i} a \cdot \Pr_{t_{ij}} [\bar{v}_{ij}(t_{ij}) \geq a]$, which is the optimal revenue from selling permit j to buyer i . Let $r = \sum_i \sum_j r_{ij}$. We first show that $\text{TAIL} \leq \frac{1}{2} \cdot r$ and then bound r using a RSPP.

Lemma 18. $\text{TAIL} \leq \frac{1}{2} \cdot r$.

Proof.

$$\begin{aligned} \text{TAIL} &= \sum_i \sum_j \mathbb{E}_{t_{ij}:\bar{v}_{ij}(t_{ij})>\tau_i} [\bar{v}_{ij}(t_{ij}) \cdot \Pr_{t_{i,-j}} [\exists k \neq j \text{ s.t. } \bar{v}_{ik}(t_{ik}) \geq \bar{v}_{ij}(t_{ij})]] \\ &\leq \sum_i \sum_j \mathbb{E}_{t_{ij}:\bar{v}_{ij}(t_{ij})>\tau_i} [\bar{v}_{ij}(t_{ij}) \cdot \sum_{k \neq j} \Pr_{t_{ik}} [\bar{v}_{ik}(t_{ik}) \geq \bar{v}_{ij}(t_{ij})]] \\ &\leq \sum_i \sum_j \mathbb{E}_{t_{ij}:\bar{v}_{ij}(t_{ij})>\tau_i} \left[\sum_{k \neq j} r_{ik} \right] \\ &\leq \sum_i \sum_j \Pr[\bar{v}_{ij}(t_{ij}) > \tau_i] \cdot r_i \\ &= \frac{1}{2} \cdot r \end{aligned}$$

□

The following lemma bounds r using the RSPP.

Lemma 19. For any positive $\{\xi_{ij}\}_{i,j}$ such that $\sum_j \Pr_{t_{ij}} [\bar{v}_{ij}(t_{ij}) \geq \xi_{ij}] \leq \frac{1}{2}$, we have

$$\sum_i \sum_j \xi_{ij} \cdot \Pr_{t_{ij}} [\bar{v}_{ij}(t_{ij}) \geq \xi_{ij}] \leq 4 \cdot \text{RSPP-PROFIT}$$

Proof. Consider the RSPP mechanism with permit price $\frac{1}{2}\xi_{ij}$ and item price $\max\{\beta_{ij}(\mathbf{c}), c_j\}$. Notice that for every buyer i , her expected utility for purchasing each permit j is $\frac{1}{2} \cdot \bar{v}_{ij}(t_{ij})$. She will purchase every permit j for sure if both of the events happen:

1. She is willing to purchase permit j , i.e., $\bar{v}_{ij}(t_{ij}) \geq \xi_{ij}$.
2. She is not willing to purchase other permits, i.e., $\bar{v}_{ik}(t_{ik}) < \xi_{ik}, \forall k \neq j$.

(1) happens with probability $\Pr[\bar{v}_{ij}(t_{ij}) \geq \xi_{ij}]$; By union bound, (2) happens with probability at least $\frac{1}{2}$ as $\sum_j \Pr_{t_{ij}}[\bar{v}_{ij}(t_{ij}) \geq \xi_{ij}] \leq \frac{1}{2}$. Furthermore, both events are independent and thus buyer i will purchase permit j and pay the permit price with probability at least $\frac{1}{2} \cdot \Pr_{t_{ij}}[\bar{v}_{ij}(t_{ij}) \geq \xi_{ij}]$. \square

We point out that in the above lemma, it's necessary to make every buyer i 's expected utility for purchasing each permit j to be exactly $\frac{1}{2} \cdot \bar{v}_{ij}(t_{ij})$. This is the reason the RSPP mechanism needs to hide each item randomly to make each item available with probability exactly $\frac{1}{2}$ (See Section 2). If the mechanism doesn't hide the item, we only know that her expected utility for each permit is at least that much. We are not able to lower bound the probability that (2) happens using union bound.

Lemma 20. $\text{TAIL} \leq 2 \cdot \text{RSPP-PROFIT}$.

Proof. It directly follows from Lemma 18 and 19 by applying $\operatorname{argmax}_{a \geq \tau_i} a \cdot \Pr_{t_{ij}}[\bar{v}_{ij}(t_{ij}) \geq a]$ as ξ_{ij} (Notice that it satisfies the constraint in Lemma 19 by the definition of τ_i). \square

6.3.2 CORE

In this section we bound CORE using RSPP and SPB.

Theorem 7. $\text{CORE} \leq 8 \cdot \text{SPB-PROFIT} + 20 \cdot \text{RSPP-PROFIT}$.

Recall that $\text{CORE} = \sum_i \mathbb{E}_{t_i}[\bar{v}_i(t_i, C_i(t_i))]$. In the proof we will consider the SPB mechanism with item prices $\max\{\beta_{ij}(\mathbf{c}), c_j\}$ and permit bundle price $\delta_i = \frac{1}{2} \cdot \text{median}_{t_i}(\bar{v}_i(t_i, C_i(t_i)))$. In order to show that each buyer will accept this bundle price with at least half probability, we will prove the expected utility for the item-purchasing stage is at least $\frac{1}{2} \cdot \bar{v}_i(t_i, [m])$. We need the following definition.

Definition 6. Consider the above SPB mechanism. For every i, t_i, \mathbf{c} and $P \subseteq [m]$, let

$$u_i(t_i, \mathbf{c}, P) = \max_{S \subseteq P, S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\})$$

By Definition 6, buyer i 's expected utility for the item purchasing stage is $\mathbb{E}_{t_{<i}, \mathbf{c}}[u_i(t_i, \mathbf{c}, S_i(t_{<i}, \mathbf{c}))]$. Recall that $S_i(t_{<i}, \mathbf{c})$ is the set of available items in the above SPB mechanism. We notice that for every i, j, \mathbf{c} , buyer i can afford the item price for j with probability $q_{ij}(\mathbf{c})$. Thus by union bound, each item j is still available when buyer i comes with probability at least $1 - \sum_j q_{ij}(\mathbf{c}) \geq \frac{1}{2}$, i.e. $\Pr_{t_{<i}}[j \in S_i(t_{<i}, \mathbf{c})] \geq \frac{1}{2}$. Then by showing that all u_i s are XOS valuations, we prove that every buyer has expected utility at least $\frac{1}{2} \cdot \bar{v}_i(t_i, [m])$ to enter the auction.

Definition 7. A function $v : 2^{[m]} \rightarrow R_+$ is XOS (or fractionally-subadditive) if for every $S \subseteq [m]$, $v(S) = \max_{k \in [K]} v^{(k)}(S)$ for some finite K and additive functions $v^{(k)}(\cdot)$.

Lemma 21. ([19]) A function $v(\cdot)$ is XOS if and only if for every $S \subseteq [m]$, there exist prices $\{p_j\}_{j \in S}$ (called supporting prices) such that

- $v(S') \geq \sum_{j \in S'} p_j$ for all $S' \subseteq S$.
- $\sum_{j \in S} p_j \geq v(S)$.

Lemma 22. For every i, t_i, \mathbf{c} , $u_i(t_i, \mathbf{c}, \cdot)$ is an XOS function.

Proof. Fix i, t_i, \mathbf{c} . For every $P \subseteq [m]$, let $S^* = \operatorname{argmax}_{S \subseteq P, S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\})$. Define supporting prices for set P as follows: $p_j^P = (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) \cdot \mathbb{1}[j \in S^*]$. It's easy to check that p_j^P satisfy both constraints in Lemma 7. Thus $u_i(t_i, \mathbf{c}, \cdot)$ is an XOS function. \square

Lemma 23. Consider the above SPB mechanism. For every i , buyer i will accept the bundle price δ_i with at least $\frac{1}{2}$ probability.

Proof. As stated above, for every buyer i with type t_i , her expected utility on the item-purchasing stage is $\mathbb{E}_{t_{<i}, \mathbf{c}}[u_i(t_i, \mathbf{c}, S_i(t_{<i}, \mathbf{c}))]$. For every i, j, \mathbf{c} , buyer i can afford the item price for j with probability $q_{ij}(\mathbf{c})$. Thus by union bound, $\Pr_{t_{<i}}[j \in S_i(t_{<i}, \mathbf{c})] \geq 1 - \sum_j q_{ij}(\mathbf{c}) \geq \frac{1}{2}$.

By Lemma 21 and 22, let $p_j^P(t_i, \mathbf{c})$ be the supporting price for $u_i(t_i, \mathbf{c}, \cdot)$ and set P . We have

$$\begin{aligned} \mathbb{E}_{t_{<i}, \mathbf{c}}[u_i(t_i, \mathbf{c}, S_i(t_{<i}, \mathbf{c}))] &\geq \mathbb{E}_{t_{<i}, \mathbf{c}} \left[\sum_{j \in S_i(t_{<i}, \mathbf{c})} p_j^{[m]}(t_i, \mathbf{c}) \right] \\ &= \mathbb{E}_{\mathbf{c}} \left[\sum_{j \in [m]} p_j^{[m]}(t_i, \mathbf{c}) \cdot \Pr_{t_{<i}}[j \in S_i(t_{<i}, \mathbf{c})] \right] \\ &\geq \frac{1}{2} \cdot \mathbb{E}_{\mathbf{c}}[u_i(t_i, \mathbf{c}, [m])] = \frac{1}{2} \bar{v}_i(t_i, [m]) \geq \frac{1}{2} \cdot \bar{v}_i(t_i, C_i(t_i)) \end{aligned}$$

Thus buyer i will pay $\delta_i = \frac{1}{2} \cdot \operatorname{median}_{t_i}(\bar{v}_i(t_i, C_i(t_i)))$ with probability at least $\frac{1}{2}$. \square

Now it's sufficient to show that $\mathbb{E}_{t_i}[\bar{v}_i(t_i, C_i(t_i))]$ is comparable to δ_i for every i . This is obtained by applying the Talagrand's concentration inequality on $\bar{v}_i(t_i, C_i(t_i))$. Let $\mu_i(t_i, S) = \bar{v}_i(t_i, C_i(t_i) \cap S)$. We show that μ_i is subadditive and has small Lipschitz constant. The proof in Lemma 24 is postponed to Appendix D.

Definition 8. A function $v(\cdot, \cdot)$ is a -Lipschitz if for any type $t, t' \in T$, and set $X, Y \subseteq [m]$,

$$|v(t, X) - v(t', Y)| \leq a \cdot (|X \Delta Y| + |\{j \in X \cap Y : t_j \neq t'_j\}|),$$

where $X \Delta Y = (X \setminus Y) \cup (Y \setminus X)$ is the symmetric difference between X and Y .

Lemma 24. $\mu_i(t_i, \cdot)$ is monotone, subadditive, no externalities and has Lipschitz constant τ_i .

Lemma 25. [30] Let $g(t, \cdot)$ with $t \sim D = \prod_j D_j$ be a function drawn from a distribution that is subadditive over independent items of ground set I . If $g(\cdot, \cdot)$ is c -Lipschitz, then for all $a > 0, k \in \{1, 2, \dots, |I|\}, q \in \mathbb{N}$,

$$\Pr_t[g(t, I) \geq (q+1)a + k \cdot c] \leq \Pr_t[g(t, I) \leq a]^{-q} q^{-k}.$$

The following corollary comes from [9]. It's a corollary of Lemma 25. It shows that for every i , $\mathbb{E}_{t_i}[\bar{v}_i(t_i, C_i(t_i))]$ is bounded by δ_i and the Lipschitz constant τ_i .

Corollary 2. [9]

$$\mathbb{E}_{t_i}[\bar{v}_i(t_i, C_i(t_i))] = \mathbb{E}_{t_i}[\mu_i(t_i, [m])] \leq 4 \cdot \delta_i + \frac{5}{2} \cdot \tau_i$$

For the last step, $\sum_i \tau_i$ can be bounded using the RSPP.

Lemma 26. $\sum_i \tau_i \leq 8 \cdot \text{RSPP-PROFIT}$.

Proof. By definition, $\sum_j \Pr_{t_{ij}}[\bar{v}_{ij}(t_{ij}) \geq \tau_i] = \frac{1}{2}$ for every i . By Lemma 19,

$$\text{RSPP-PROFIT} \geq \frac{1}{4} \sum_i \tau_i \cdot \sum_j \Pr_{t_{ij}}[\bar{v}_{ij}(t_{ij}) \geq \tau_i] = \frac{1}{8} \cdot \sum_i \tau_i$$

□

Proof of Theorem 7: Consider the SPB mechanism with item prices $\max\{\beta_{ij}(\mathbf{c}), c_j\}$ and permit bundle price $\delta_i = \frac{1}{2} \cdot \text{median}_{t_i}(\bar{v}_i(t_i, C_i(t_i)))$. According to Lemma 23, 26 and Corollary 2,

$$\text{SPB-PROFIT} \geq \frac{1}{2} \cdot \sum_i \delta_i \geq \frac{1}{8} \left(\sum_i \mathbb{E}[\bar{v}_i(t_i, C_i(t_i))] - \frac{5}{2} \tau_i \right) \geq \frac{1}{8} \cdot (\text{CORE} - 20 \cdot \text{RSPP-PROFIT})$$

□

References

- [1] Moshe Babaioff, Yang Cai, Yannai A. Gonczarowski, and Mingfei Zhao. The best of both worlds: Asymptotically efficient mechanisms with a guarantee on the expected gains-from-trade. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, page 373, 2018. 1.3
- [2] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A Simple and Approximately Optimal Mechanism for an Additive Buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014. 1, 1.4
- [3] Liad Blumrosen and Shahar Dobzinski. (almost) efficient mechanisms for bilateral trading. *CoRR*, abs/1604.04876, 2016. 1.3
- [4] Peter Bro Miltersen and Or Sheffet. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 234–247. ACM, 2012. 1, 1.1, 1.4
- [5] Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 589–590, 2017. 1.3
- [6] Yang Cai and Constantinos Daskalakis. Extreme-Value Theorems for Optimal Multidimensional Pricing. In *the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011. 1
- [7] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *the 48th Annual ACM Symposium on Theory of Computing (STOC)*, 2016. 1, 1.1, 1.4, 4.1, 4.2.1, 4.2.2, 3, 5.2, 8, 6.3, A
- [8] Yang Cai and Zhiyi Huang. Simple and Nearly Optimal Multi-Item Auctions. In *the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2013. 1, 1.4

- [9] Yang Cai and Mingfei Zhao. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 170–183, 2017. [1](#), [1.1](#), [1.4](#), [4.2.1](#), [4.2.2](#), [4.2.2](#), [5.2](#), [7](#), [5.2](#), [6.3.2](#), [2](#), [E](#)
- [10] Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic Pricing via Virtual Valuations. In *the 8th ACM Conference on Electronic Commerce (EC)*, 2007. [1](#), [1.4](#)
- [11] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*, 2010. [1](#), [1](#), [1.4](#), [2.1](#), [5.1](#), [5.1](#), [6.1](#), [6.1](#), [E](#)
- [12] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015. [1.4](#)
- [13] Shuchi Chawla and J. Benjamin Miller. Mechanism design for subadditive agents via an ex-ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 579–596, 2016. [1](#), [1.1](#), [1.4](#), [4.2.2](#), [6.3](#)
- [14] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1424–1443, 2016. [1.3](#)
- [15] Riccardo Colini-Baldeschi, Paul W. Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, and Stefano Turchetta. Approximately efficient two-sided combinatorial auctions. *CoRR*, abs/1611.05342, 2016. [1.3](#)
- [16] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *ACM Conference on Electronic Commerce, EC '13, Philadelphia, PA, USA, June 16-20, 2013*, pages 269–286, 2013. [1](#)
- [17] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014. [1](#)
- [18] Constantinos Daskalakis, Christos Papadimitriou, and Christos Tzamos. Does information revelation improve revenue? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 233–250. ACM, 2016. [1](#), [1.1](#), [1.4](#)
- [19] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC*, pages 610–618, 2005. [21](#)
- [20] Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. Constrained signaling in auction design. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1341–1357. Society for Industrial and Applied Mathematics, 2014. [1](#), [1.1](#), [1.4](#)
- [21] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. Modularity and greed in double auctions. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 241–258, 2014. [1.3](#)
- [22] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato PaesLeme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation*, 2(2):5, 2014. [1](#), [1.1](#), [1.4](#)

- [23] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1014–1033, 2016. [1.1](#), [6.2](#), [5](#), [11](#), [6.2](#), [14](#), [15](#)
- [24] Hu Fu, Patrick Jordan, Mohammad Mahdian, Uri Nadav, Inbal Talgam-Cohen, and Sergei Vassilvitskii. Ad auctions with data. In *Algorithmic Game Theory*, pages 168–179. Springer, 2012. [1](#), [1.1](#), [1.4](#)
- [25] Yiannis Giannakopoulos and Elias Koutsoupias. Duality and optimality of auctions for uniform distributions. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 259–276, 2014. [1](#)
- [26] Sergiu Hart and Noam Nisan. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*, 2012. [1.4](#)
- [27] Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017. [1](#), [1](#)
- [28] Xinye Li and Andrew Chi-Chih Yao. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences*, 110(28):11232–11237, 2013. [1](#), [1.1](#), [1.4](#), [6.3](#)
- [29] Aviad Rubinstein and S. Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 377–394, 2015. [1](#), [1.1](#), [1.4](#), [5.2](#), [4](#)
- [30] Gideon Schechtman. Concentration, results and applications. *Handbook of the geometry of Banach spaces*, 2:1603–1634, 2003. [1.1](#), [5.2](#), [25](#)
- [31] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *SODA*, 2015. [1](#), [1.4](#)

A Duality Framework

The seller aims to maximize her profit among all direct, BIC, and interim IR mechanisms. This maximization problem can be captured by the following LP (see Figure 2). Here we use type \emptyset to represent the choice of not participating in the mechanism. Now the IR constraint can be described as another BIC constraint that the buyer won't report type \emptyset . Let $T_i^+ = T_i \cup \{\emptyset\}$.

Variables:

- $\pi_i(t_i, \mathbf{c})$, for all $i \in [n]$, $t_i \in T_i$, $\mathbf{c} \in T^S$, denotes the interim probability vector that buyer i with type t_i receives each item, when the seller has cost \mathbf{c} .
- $p_i(t_i, \mathbf{c})$, for all $i \in [n]$, $t_i \in T_i$, $\mathbf{c} \in T^S$, denoting the buyer i 's interim payment when she has type t_i and the seller has cost \mathbf{c} .

Constraints:

- $\mathbb{E}_{\mathbf{c}}[t_i \cdot \pi_i(t_i, \mathbf{c}) - p_i(t_i, \mathbf{c})] \geq \mathbb{E}_{\mathbf{c}}[t_i \cdot \pi_i(t'_i, \mathbf{c}) - p_i(t'_i, \mathbf{c})]$, for all $i \in [n]$, $t_i \in T_i$, $t'_i \in T_i^+$, guaranteeing that the mechanism is BIC and interim IR.
- $\pi \in P(\{\mathcal{F}_i\}_{i=1}^n)$, guaranteeing the allocation is implementable.

Objective:

- $\max \sum_i \mathbb{E}_{t_i, \mathbf{c}}[p_i(t_i, \mathbf{c}) - \mathbf{c} \cdot \pi_i(t_i, \mathbf{c})]$, the expected seller's profit.

Figure 2: A Linear Program (LP) for Maximizing Profit.

We then take the partial Lagrangian dual of the LP in Figure 2 by lagrangifying the BIC and interim IR constraints. Let $\lambda_i(\mathbf{t}, \mathbf{t}')$ be the Lagrangian multiplier. The dual problem is described in Figure 3.

Variables:

- $\pi_i(t_i, \mathbf{c})$ and $p_i(t_i, \mathbf{c})$.
- $\lambda_i(\mathbf{t}, \mathbf{t}')$ for all $i \in [n]$, $t_i \in T_i$, $t'_i \in T_i^+$, the Lagrangian multiplier for buyer i 's BIC and interim IR constraints.

Constraints:

- $\lambda_i(\mathbf{t}, \mathbf{t}') \geq 0$ for all $i \in [n]$, $t_i \in T_i$, $t'_i \in T_i^+$.
- $\pi \in P(\{\mathcal{F}_i\}_{i=1}^n)$.

Objective:

- $\min_{\lambda} \max_{\pi, p} \mathcal{L}(\lambda, \pi, p)$.

Figure 3: Partial Lagrangian of the LP for Maximizing Profit.

$$\begin{aligned}
& \mathcal{L}(\lambda, \pi, p) \\
&= \sum_i \mathbb{E}_{t_i, \mathbf{c}} [p_i(t_i, \mathbf{c}) - \mathbf{c} \cdot \pi_i(t_i, \mathbf{c})] \\
&\quad + \sum_i \sum_{t_i, t'_i} \lambda_i(t_i, t'_i) \cdot \mathbb{E}_{\mathbf{c}} [(t_i \cdot \pi_i(t_i, \mathbf{c}) - p_i(t_i, \mathbf{c})) - (t_i \cdot \pi_i(t'_i, \mathbf{c}) - p_i(t'_i, \mathbf{c}))] \\
&= \sum_i \sum_{t_i} \mathbb{E}_{\mathbf{c}} [p_i(t_i, \mathbf{c})] \cdot \left(f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i) - \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \right) \\
&\quad + \sum_i \sum_{t_i} \mathbb{E}_{\mathbf{c}} \left[\pi_i(t_i, \mathbf{c}) \cdot \left(\sum_{t'_i \in T_i^+} t_i \cdot \lambda_i(t_i, t'_i) - \sum_{t'_i \in T_i} t'_i \cdot \lambda_i(t'_i, t_i) - f_i(t_i) \cdot \mathbf{c} \right) \right]
\end{aligned} \tag{1}$$

Definition 9. A feasible dual solution λ is useful if $\max_{\pi \in P(\{\mathcal{F}_i\}_{i=1}^n), p} \mathcal{L}(\lambda, \pi, p) < \infty$.

Similar to [7], we show that every useful dual solution forms a flow.

Lemma 27. A dual solution λ is useful if and only if it forms the following flow:

- Nodes: For every $i \in [n]$ and $t_i \in T_i$ a node t_i . A source s and a sink \emptyset .
- For every $i \in [n]$ and $t_i \in T_i$, a flow of weight $f_i(t_i)$ from s to t_i .
- For every $i \in [n]$ and $t_i \in T_i, t'_i \in T_i^+$, a flow of weight $\lambda_i(t_i, t'_i)$ from t_i to t'_i .

Proof. Suppose there exists $i \in [n], t_i \in T_i, t'_i \in T_i^+$ such that

$$f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i) - \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) \neq 0$$

Without loss of generality, suppose it's positive. Notice that $p_i(t_i, \mathbf{c})$ is unconstrained. Thus when $\mathbb{E}_{\mathbf{c}} [p_i(t_i, \mathbf{c})] \rightarrow +\infty$, the Lagrangian also goes to $+\infty$ (see Equation 1). Hence for every $i \in [n], t_i \in T_i, t'_i \in T_i^+$,

$$f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i) - \sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i) = 0$$

It's essentially the flow conservation equation for node t_i . Thus λ forms a flow. On the other hand, if λ forms a flow, the Lagrangian only depends on π and thus bounded since π is bounded. \square

For any useful dual solution λ , by Lemma 27, we can replace $\sum_{t'_i \in T_i^+} \lambda_i(t_i, t'_i)$ by $f_i(t_i) + \sum_{t'_i \in T_i} \lambda_i(t'_i, t_i)$ in Equation (1) and simplify $\mathcal{L}(\lambda, \pi, p)$. For any BIC and interim IR mechanism $M = (\pi, p)$, both $\lambda_i(t_i, t'_i)$ and $\mathbb{E}_{\mathbf{c}} [(t_i \cdot \pi_i(t_i, \mathbf{c}) - p_i(t_i, \mathbf{c})) - (t_i \cdot \pi_i(t'_i, \mathbf{c}) - p_i(t'_i, \mathbf{c}))]$ are non-negative for all $i \in [n], t_i \in T_i, t'_i \in T_i^+$. Thus by Equation (1), $\mathcal{L}(\lambda, \pi, p) \geq \text{PROFIT}(M)$. We have the following lemma.

Lemma 28. (Restatement of Lemma 2) For any useful dual solution λ and any BIC, interim IR mechanism $M = (x, p)$,

$$\text{PROFIT}(M) \leq \mathbb{E}_{\mathbf{t}, \mathbf{c}} \left[\sum_i \pi_i(t_i, \mathbf{c}) \cdot (\Phi_i^{(\lambda)}(t_i) - \mathbf{c}) \right]$$

where

$$\Phi_i^{(\lambda)}(t_i) = t_i - \frac{1}{f_i(t_i)} \cdot \sum_{t'_i \in T_i} \lambda(t'_i, t_i)(t'_i - t_i)$$

can be viewed as buyer i 's virtual value function.

B Missing Proofs from Section 4

Proof of Lemma 1: Consider a mechanism M that is ex-post implementable. For every i, t_i , let $A_i(t_i)$ be buyer i 's (possibly randomized) equilibrium strategy, when her type is t_i . It specifies all the actions that the buyer takes in mechanism M . For every \mathbf{c} , let $X_i(\vec{A}, \mathbf{c})$ be the vector of (possibly randomized) indicator variables that indicate whether buyer i gets each item j when buyers choose strategies $\vec{A} = (A_1, \dots, A_n)$ and the seller's realized cost vector is \mathbf{c} ; let $P_i(\vec{A}, \mathbf{c})$ be the payment for the buyer. For every i and t_{-i} , denote $A_{-i}(t_{-i}) = (A_1(t_1), \dots, A_{i-1}(t_{i-1}), A_{i+1}(t_{i+1}), \dots, A_n(t_n))$.

Since A is an equilibrium strategy for the buyers, for every $i, t_i, t'_i \in T_i$, acting as $A_i(t_i)$ induces more utility than $A_i(t'_i)$, when the buyer's type is t_i and other buyers follow strategy A_{-i} . We have

$$\begin{aligned} \mathbb{E}_{\mathbf{c}, t_{-i}} [t_i \cdot X_i(A_i(t_i), A_{-i}(t_{-i}), \mathbf{c}) - P_i(A_i(t_i), A_{-i}(t_{-i}), \mathbf{c})] \geq \\ \mathbb{E}_{\mathbf{c}, t_{-i}} [t_i \cdot X_i(A_i(t'_i), A_{-i}(t_{-i}), \mathbf{c}) - P_i(A_i(t'_i), A_{-i}(t_{-i}), \mathbf{c})] \end{aligned} \quad (2)$$

We now define the direct mechanism $M' = (x, p)$ as follows: for every profile (\mathbf{t}, \mathbf{c}) , let $x_i(\mathbf{t}, \mathbf{c}) = X_i(A(\mathbf{t}), \mathbf{c})$ and $p_i(\mathbf{t}, \mathbf{c}) = P_i(A(\mathbf{t}), \mathbf{c})$ for all i . It's the allocation and payment rule when the reported type profile is \mathbf{t} and the seller's realized cost vector is \mathbf{c} . Then Inequality (2) is equivalent to: for every $i, t_i, t'_i \in T_i$

$$\mathbb{E}_{\mathbf{c}, t_{-i}} [t_i \cdot x_i(t_i, t_{-i}, \mathbf{c}) - p_i(t_i, t_{-i}, \mathbf{c})] \geq \mathbb{E}_{\mathbf{c}, t_{-i}} [t_i \cdot x_i(t'_i, t_{-i}, \mathbf{c}) - p_i(t'_i, t_{-i}, \mathbf{c})]$$

It's exactly the BIC constraint for M' . Thus, M' is BIC.

Moreover, each buyer can choose not to participate in M , so $\mathbb{E}_{\mathbf{c}, t_{-i}} [t_i \cdot X_i(A_i(t_i), A_{-i}(t_{-i}), \mathbf{c}) - P_i(A_i(t_i), A_{-i}(t_{-i}), \mathbf{c})] \geq 0$, which implies that $\mathbb{E}_{\mathbf{c}, t_{-i}} [t_i \cdot x_i(t_i, t_{-i}, \mathbf{c}) - p_i(t_i, t_{-i}, \mathbf{c})] \geq 0$. Hence, M' is also interim IR. \square

Proof of Theorem 4: Let $\Phi^{(\lambda)}(\cdot)$ be the virtual value function induced by the above canonical flow λ . By

Lemma 2 and Lemma 3,

$$\begin{aligned}
\text{PROFIT}(M) &\leq \sum_i \mathbb{E}_{t_i, \mathbf{c}} \left[\sum_j \pi_{ij}(t_i, \mathbf{c}) \cdot (\Phi_{ij}^{(\lambda)}(t_{ij}) - c_j) \right] \\
&= \sum_i \mathbb{E}_{t_i, \mathbf{c}} \left[\sum_j \mathbf{1}[t_i \in R_{ij}^{(\beta)}] \cdot \pi_{ij}(t_i, \mathbf{c}) \cdot (\tilde{\varphi}_{ij}(t_{ij}) - c_j) \right] \\
&\quad + \sum_i \mathbb{E}_{t_i, \mathbf{c}} \left[\sum_j \mathbf{1}[t_i \notin R_{ij}^{(\beta)}] \cdot \pi_{ij}(t_i, \mathbf{c}) \cdot (\max\{\beta_{ij}(\mathbf{c}), c_j\} - c_j) \right] \\
&\quad + \sum_i \mathbb{E}_{t_i, \mathbf{c}} \left[\sum_j \mathbf{1}[t_i \notin R_{ij}^{(\beta)}] \cdot \pi_{ij}(t_i, \mathbf{c}) \cdot (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) \right] \\
&\leq \sum_i \mathbb{E}_{t_i, \mathbf{c}} \left[\sum_j \mathbf{1}[t_i \in R_{ij}^{(\beta)}] \cdot \pi_{ij}(t_i, \mathbf{c}) \cdot (\tilde{\varphi}_{ij}(t_{ij}) - c_j) \right] \quad (\text{MOST-SURPLUS}) \\
&\quad + 2 \cdot \sum_i \sum_j \mathbb{E}_{\mathbf{c}} [q_{ij}(\mathbf{c}) \cdot (\max\{\beta_{ij}(\mathbf{c}), c_j\} - c_j)] \quad (\text{PROPHET}) \\
&\quad + \sum_i \mathbb{E}_{t_i} \left[\sum_j \mathbf{1}[t_i \in R_{ij}^{(\beta)}] \cdot \bar{v}_i^{(\beta)}(t_i, [m] \setminus \{j\}) \right] \quad (\text{LESS-SURPLUS})
\end{aligned}$$

The first inequality is due to Lemma 2, and the first equality is due to Lemma 3. The second inequality is because: For the second term, notice that $\max\{\beta_{ij}(\mathbf{c}), c_j\} - c_j \geq 0$, we bound the indicator by 1 and use the fact that $\mathbb{E}_{t_i}[\pi_{ij}(t_i, \mathbf{c})] = 2 \cdot q_{ij}(\mathbf{c})$ for every \mathbf{c} ; For the third term, notice that for every $i, t_i \in R_{ij}^{(\beta)}$ and \mathbf{c} , since π is feasible, we have

$$\sum_{k \neq j} \pi_{ik}(t_i, \mathbf{c}) \cdot (t_{ik} - \max\{\beta_{ik}(\mathbf{c}), c_k\}) \leq \max_{S \in \mathcal{F}_i, j \notin S} \sum_{k \in S} (t_{ik} - \max\{\beta_{ik}(\mathbf{c}), c_k\}).$$

Taking expectation over \mathbf{c} , the RHS equals to $\bar{v}_i^{(\beta)}(t_i, [m] \setminus \{j\})$. Thus the inequality holds. \square

C Missing Proofs from Section 5

Proof of Lemma 6:

Monotonicity: Fix any $\mathbf{t} \in T, U \subseteq V \subseteq [m]$. For all \mathbf{c} we have $\max_{S \subseteq U, S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j) \leq \max_{S \subseteq V, S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j)$. Taking expectation over \mathbf{c} on both sides proves the monotonicity.

Subadditivity:

Fix any $\mathbf{t} \in T$ and $U, V \subseteq [m]$. For every \mathbf{c} , let $S^*(\mathbf{c}) = \operatorname{argmax}_{S \subseteq U \cup V, S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j)$. Clearly, $t_j - c_j \geq 0$ for all $j \in S^*(\mathbf{c})$. Notice that $S^*(\mathbf{c}) \cap U \subseteq U$ and $S^*(\mathbf{c}) \cap U \in \mathcal{F}$; also $S^*(\mathbf{c}) \cap V \subseteq V$ and

$S^*(\mathbf{c}) \cap V \in \mathcal{F}$. We have

$$\begin{aligned} \max_{S \subseteq U \cup V, S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j) &\leq \sum_{j \in S^*(\mathbf{c}) \cap U} (t_j - c_j) + \sum_{j \in S^*(\mathbf{c}) \cap V} (t_j - c_j) \\ &\leq \max_{S \subseteq U, S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j) + \max_{S \subseteq V, S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j) \end{aligned}$$

Taking expectation over \mathbf{c} on both sides, we have $\bar{v}(\mathbf{t}, U \cup V) \leq \bar{v}(\mathbf{t}, U) + \bar{v}(\mathbf{t}, V)$.

No externalities: fix any $\mathbf{t} \in T$, $S \subseteq [m]$ and any $\mathbf{t}' \in T$ such that $t'_j = t_j$ for all $j \in S$. To prove $\bar{v}(\mathbf{t}', S) = \bar{v}(\mathbf{t}, S)$, it suffices to show that for any \mathbf{c} ,

$$\max_{U \subseteq S, U \in \mathcal{F}} \sum_{j \in U} (t_j - c_j) = \max_{U \subseteq S, U \in \mathcal{F}} \sum_{j \in U} (t'_j - c_j)$$

It follows directly from the fact that $t'_j = t_j$ for all $j \in S$. \square

D Missing Proofs from Section 6

Proof of Lemma 24:

Monotonicity: Fix any $t_i, U \subseteq V \subseteq [m]$. For all \mathbf{c} we have

$$\max_{S \subseteq U \cap C_i(t_i), S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) \leq \max_{S \subseteq V \cap C_i(t_i), S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}).$$

Taking expectation over \mathbf{c} on both sides proves the monotonicity.

Subadditivity:

Fix any t_i and $U, V \subseteq [m]$. For every \mathbf{c} , let $S^*(\mathbf{c}) = \operatorname{argmax}_{S \subseteq (U \cup V) \cap C_i(t_i), S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\})$.

Clearly, $t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\} \geq 0$ for all $j \in S^*(\mathbf{c})$. Notice that $S^*(\mathbf{c}) \cap U \subseteq U \cap C_i(t_i)$ and $S^*(\mathbf{c}) \cap U \in \mathcal{F}_i$; also $S^*(\mathbf{c}) \cap V \subseteq V \cap C_i(t_i)$ and $S^*(\mathbf{c}) \cap V \in \mathcal{F}_i$. We have

$$\begin{aligned} &\max_{S \subseteq (U \cup V) \cap C_i(t_i), S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) \\ &\leq \sum_{j \in S^*(\mathbf{c}) \cap U} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) + \sum_{j \in S^*(\mathbf{c}) \cap V} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) \\ &\leq \max_{S \subseteq U \cap C_i(t_i), S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) + \max_{S \subseteq V \cap C_i(t_i), S \in \mathcal{F}_i} \sum_{j \in S} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) \end{aligned}$$

Taking expectation over \mathbf{c} on both sides, we have $\mu_i(t_i, U \cup V) \leq \mu_i(t_i, U) + \mu_i(t_i, V)$.

No externalities: fix any $t_i, S \subseteq [m]$ and any t_i such that $t'_{ij} = t_{ij}$ for all $j \in S$. To prove $\mu_i(t'_i, S) = \mu_i(t_i, S)$, it suffices to show that for any \mathbf{c} ,

$$\max_{U \subseteq S \cap C_i(t_i), U \in \mathcal{F}_i} \sum_{j \in U} (t_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\}) = \max_{U \subseteq S \cap C_i(t_i), U \in \mathcal{F}_i} \sum_{j \in U} (t'_{ij} - \max\{\beta_{ij}(\mathbf{c}), c_j\})$$

It follows directly from the fact that $t'_{ij} = t_{ij}$ for all $j \in S$.

Now we prove that $\mu_i(t_i, \cdot)$ has Lipschitz constant τ_i . For any t_i, t'_i , and set $X, Y \subseteq [m]$, define set $H = \{j \in X \cap Y : t_{ij} = t'_{ij}\}$. Since $\mu_i(\cdot, \cdot)$ has no externalities, $\mu_i(t_i, H) = \mu_i(t'_i, H)$. Therefore,

$$\begin{aligned}
|\mu_i(t_i, X) - \mu_i(t'_i, Y)| &= \max \{ \mu_i(t_i, X) - \mu_i(t'_i, Y), \mu_i(t'_i, Y) - \mu_i(t_i, X) \} \\
&\leq \max \{ \mu_i(t_i, X) - \mu_i(t'_i, H), \mu_i(t'_i, Y) - \mu_i(t_i, H) \} \quad (\text{Monotonicity}) \\
&\leq \max \{ \mu_i(t_i, X \setminus H), \mu_i(t'_i, Y \setminus H) \} \quad (\text{Subadditivity}) \\
&\leq \tau_i \cdot \max \{ |X \setminus H|, |Y \setminus H| \} \\
&\leq \tau_i \cdot (|X \Delta Y| + |X \cap Y| - |H|)
\end{aligned}$$

The second last inequality is because $\mu_i(t_i, \cdot)$ is subadditive and for any item $j \in \mathcal{C}_j(t_i)$ ($\mathcal{C}_j(t'_i)$) the single-item valuation $\bar{v}_{ij}(t_{ij})$ ($\bar{v}_{ij}(t'_{ij})$) is less than τ_i .

□

E Computing the Simple Mechanisms for the Single Buyer Case

In this section, we briefly discuss how to compute the simple mechanisms used in Section 5 in polynomial time. For the sell-items-separately mechanism, for every \mathbf{c} , let $\hat{p}_j(\mathbf{c})$ be the posted price of the mechanism constructed in [11], which approximates $\text{OPT-REV}^{\text{COPIES-UD}}(\mathbf{c})$. According to Lemma 4, selling the items separately with price $\hat{p}_j(\mathbf{c}) + c_j$ when the seller's cost is \mathbf{c} can achieve profit at least half of the MOST-SURPLUS.

For the permit-selling mechanisms, we first point out that value oracle for the valuation \bar{v} can be (approximately) implemented efficiently, given an algorithm for the following maximization problem over the feasibility constraint \mathcal{F} : $\max_{S \in \mathcal{F}} \sum_{j \in S} (t_j - c_j)$, for every (\mathbf{t}, \mathbf{c}) . Given \mathbf{t} as an input, for every \mathbf{c} the oracle uses the algorithm to obtain the favorite bundle and then calculates the expected utility. The computation of the oracle requires sampling from the cost distribution T^S , but it's not hard to argue that a polynomial number of samples suffices. Besides, the conversion in Lemma 5 is also polynomial time, since the permit-selling mechanism follows the same allocation and payment rule in the first stage as the mechanism in the revenue setting. According to [9], both the Selling Separately mechanism and Bundling mechanism used to bound $\text{NON-FAV}_{\text{REV}}(\bar{v})$ can be computed efficiently given the value oracle for \bar{v} . And the two mechanisms can be converted to the PS and PB mechanism accordingly in polynomial time.