Online Allocation and Pricing with Economies of Scale

Avrim Blum[*] Yishay Mansour[†] Liu Yang[‡]

Abstract

Allocating multiple goods to customers in a way that maximizes some desired objective is a fundamental part of Algorithmic Mechanism Design. We consider here the problem of offline and online allocation of goods that have economies of scale, or decreasing marginal cost per item for the seller. In particular, we analyze the case where customers have unit-demand and arrive one at a time with valuations on items, sampled iid from some unknown underlying distribution over valuations. Our strategy operates by using an initial sample to learn enough about the distribution to determine how best to allocate to future customers, together with an analysis of structural properties of optimal solutions that allow for uniform convergence analysis. We show, for instance, if customers have $\{0, 1\}$ valuations over items, and the goal of the allocator is to give each customer an item he or she values, we can efficiently produce such an allocation with cost at most a constant factor greater than the minimum over such allocations in hindsight, so long as the marginal costs do not decrease too rapidly. We also give a bicriteria approximation to social welfare for the case of more general valuation functions when the allocator is budget constrained.

1 Introduction

Imagine it is the Christmas season, and Santa Claus is tasked with allocating toys. There is a sequence of children coming up with their Christmas lists of toys they want. Santa wants to give each child some toy from his or her list (all children have been good this year). But of course, even Santa Claus has to be cost-conscious, so he wants to perform this allocation of toys to children at a near-minimum cost to himself (call this the Thrifty Santa Problem). Now if it was the case that every toy had a fixed price, this would be easy: simply allocate to each child the cheapest toy on his or her list and move on to the next child. But here we are interested in the case where goods have economies of scale. For example, producing a millon toy cars might be cheaper than a million times the cost of producing one toy car. Thus, even if producing a single toy car is more expensive than a single Elmo doll, if a much larger number of children want the toy car than the Elmo doll, the minimum-cost allocation might give toy cars to many children, even if some of them also have the Elmo doll on their lists.

The problem faced by Santa (or by any allocator that must satisfy a collection of disjunctive constraints in the presence of economies of scale) makes sense in both offline and online settings. In the offline setting, in the extreme case of goods such as software where all the cost is in the first copy, this is simply weighted

[*]Carnegie Mellon University, Computer Science Department, Email: avrim@cs.cmu.edu. Supported in part by the National Science Foundation under grants CCF-1101215, CCF-1116892, and IIS-1065251.

[†]Tel Aviv University, Computer Science Department, Email: mansour@tau.ac.il. Supported in part by a grant from the Israel Science Foundation, a grant from the United States-Israel Binational Science Foundation (BSF), a grant by Israel Ministry of Science and Technology and and the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11).

[‡]Carnegie Mellon University, Machine Learning Department, Email: liuy@cs.cmu.edu. Supported in part by NSF grant IIS-1065251 and a Google Core AI grant.

set-cover, admitting a $\Theta(\log n)$ approximation to the minimum-cost allocation. We will be interested in the online case where customers are iid from some arbitrary distribution over subsets of item-set $\mathcal{I}$ (i.e., Christmas lists), where the allocator must make allocation decisions online, and where the marginal cost of goods does not decrease so sharply. We show that for a range of cost curves, including the case that the marginal cost of copy t of an item is $t^{-\alpha}$, for some $\alpha \in [0, 1)$, we will be able to get a constant-factor approximation even online so long as the number of customers is sufficiently large compared to the number of items.

A basic structural property we show is that, if the marginal costs are non-increasing, the minimum cost allocation can be compactly described as an ordering of the possible toys, so that as each child comes, Santa simply gives the child the first toy in the ordering that appears on the child's list. We also show that if the marginal costs do not drop too quickly, then if we are given the lists of all the children before determining the allocation, we can efficiently find an allocation that is within a constant factor of the minimum-cost allocation, as opposed to the logarithmic factor required for the set-cover problem. Since, however, the problem we are interested in does not supply the lists before the allocations, but rather requires a decision for each child in sequence, we use ideas from machine learning, as follows: after processing a small initial number of children, we take their wish lists as if they were perfectly representative of the future children, and find an approximately optimal solution based on those, which also will be an ordering over toys. We then take the ordered list of toys from this solution, and and use it to allocate to future children (allocating to each child the earliest toy in the ordering that is also on his or her list). We show that, as long as we take a sufficiently large number of initial children, this procedure will find an ordering that will be near-optimal for allocating to the remaining children, using the fact that these compact representations allow for uniform convergence of the cost estimates to the true costs.

More generally, we can imagine the case where, rather than simple lists of items, the lists also provide valuations for each item, and we are interested in the trade-off between maximizing the total of valuations for allocated items while minimizing the total cost of the allocation. In this case, we might think of the allocator as being a large company with many different projects, and each project has some valuations over different resources (e.g., types of laptops for employees involved in that project), where it could use one or another resource but prefers some resources over others. One natural quantity to consider in this context is the social surplus: the difference between the happiness (total of valuations for the allocation) minus the total cost of the allocation. In this case, it turns out the optimal allocation rule can be described by a pricing scheme. In another words, whatever the optimal allocation is, there always exist prices such that if the buyers purchase what they most want at those prices, they will actually produce that allocation. We note that, algorithmically, this is a harder problem than the list-based problem (which corresponds to binary valuations).

Aside from social surplus, it is also interesting to consider a variant in which we have a budget constraint, and are interested in maximizing the total valuation of the allocation, subject to that budget constraint on the total cost of the allocation. It turns out this latter problem can be reduced to a problem known as the weighted budget maximum coverage problem. Technically, this problem is originally formulated for the case in which the marginal cost of a given item drops to zero after the first item of that type is allocated (as in the set cover reduction mentioned above); however, viewed appropriately, we are able to formulate this reduction for arbitrary decreasing marginal cost functions. What we can then do is run an algorithm for the weighted budget maximum coverage problem, and then convert the solution into a pricing. As before, this strategy will be effective for the offline problem, in which all of the valuations are given ahead of time. However, we can extend it to the online setting with iid valuation functions by generating a pricing based on an appropriately-sized initial sample of valuation functions, and then apply that pricing to sequentially generate allocations for the remaining valuations. As long as the marginal costs are not decreasing too rapidly, we can obtain an allocation strategy for which the sum of valuations of the allocated items will be within a constant factor of the maximum possible, minus a small additive term, subject to the budget constraint on the cost.

1.1 Our Results and Techniques

We consider this problem under two, related, natural objectives. In the first (the "thrifty Santa" objective) we assume customers have binary $\{0, 1\}$ valuations, and the goal of the seller is to give each customer a toy of value 1, but in such a way that minimizes the total cost to the seller. We show that so long as the number of buyers n is large compared to the number of items r, and so long as the marginal costs do not decrease too rapidly (e.g., a rate $1/t^\alpha$ for some $0 \leq \alpha < 1$), we can efficiently perform this allocation task with cost at most a constant factor greater than that of the optimal allocation of items in hindsight. Note that if costs decrease much more rapidly, then even if all customers' valuations were known up front, we would be faced with (roughly) a set-cover problem and so one could not hope to achieve cost $o(\log n)$ times optimal. The second objective we consider, which we apply to customers of arbitrary unit-demand valuation, is that of maximizing total social welfare of customers subject to a cost bound on the seller; for this, we also give a strategy that is constant-competitive with respect to the optimal allocation in hindsight.

Our algorithms operate by using initial buyers to learn enough about the distribution to determine how best to allocate to the future buyers. In fact, there are two main technical parts of our work: the sample complexity and the algorithmic aspects. From the perspective of sample complexity, one key component of this analysis is examining how complicated the allocation rule needs to be in order to achieve good performance, because simpler allocation rules require fewer samples in order to learn. We do this by providing a characterization of what the optimal strategies look like. For example, for the thrifty Santa Claus version, we show that the optimal solution can be assumed wlog to have a simple permutation structure. In particular, so long as the marginal costs are nonincreasing, there is always an optimal strategy in hindsight of this form: order the items according to some permutation and for each bidder, give it the earliest item of its desire in the permutation. This characterization is used inside both our sample complexity results and our algorithmic guarantees. Specifically, we prove that for cost function $\text{cost}(t) = \sum_{\tau=1}^{t} 1/\tau^\alpha$, for $\alpha \in [0, 1)$, running greedy weighted set cover incurs total cost at most $\frac{1}{1-\alpha}\text{OPT}$. More generally, if the average cost is within some factor of the marginal cost, we have a greedy algorithm that achieves constant approximation ratio. To allocate to new buyers, we simply give it the earliest item of its desire in the learnt permutation. For the case of general valuations, we give a characterization showing that the optimal allocation rule in terms of social welfare can be described by a pricing scheme. That is, there exists a pricing scheme such that if buyers purchased their preferred item at these prices, the optimal allocation would result. Algorithmically, we show that we can reduce to a weighted budgeted maximum coverage problem with single-parameter demand for which there is a known constant-approximation-ratio algorithm [8].

1.2 Related Work

In this work we focus on the case of decreasing marginal cost. There have been a large body of research devoted to unlimited supply, which is implicitly constant marginal cost (e.g., [10] Chapter 13), where the goal is to achieve a constant competitive ratio in both offline and online models. The case of increasing marginal cost was studied in [2] where constant competitive ratio was given.

We analyze an online setting where buyers arrive one at a time, sampled iid from some unknown underlying distribution over valuations. Other related online problems with stochastic inputs such as matching problems have been studied in ad auctions [5,9]. Algorithmically, our work is related to the online set cover body of work where [1] gave the first $O(\log m \log n)$ competitive algorithm (here n is the number of elements in the ground set and m is size of a family of subsets of the ground set). The problems we study are also related to online matching problems [3, 4, 7] in the iid setting; however our problem is more like the "opposite" of online matching in that the cumulative cost curve for us is concave rather than convex.

2 Model, Definitions, and Notation

We have a set $\mathcal{I}$ of r items. We have a set $N = \{1, \ldots, n\}$ indexing n unit demand buyers. Our setting can then generally be formalized in the following terms.

2.1 Utility Functions

Each buyer $j \in N$ has a weight $u_{j,i}$ for each item $i \in \mathcal{I}$. We suppose the vectors $u_{j,\cdot}$ are sampled i.i.d. according to a fixed (but arbitrary and unknown) distribution. In the *online* setting we are interested in, the buyers' weight vectors $u_{j,\cdot}$ are observed in sequence, and for each one (before observing the next) we are required to allocate a set of items $T_j \subseteq \mathcal{I}$ to that buyer. The *utility* of buyer j for this allocation is then defined as $u_j(T_j) = \max_{i \in T_j} u_{j,i}$. A few of our results consider a slight variant of this model, in which we are only required to begin allocating goods after some initial $o(n)$ number of customers has been observed (to whom we may allocate items retroactively).

This general setting is referred to as the *weighted unit demand* setting. We will also be interested in certain special cases of this problem. In particular, many of our results are for the *uniform unit demand* setting, in which every $j \in N$ and $i \in \mathcal{I}$ have $u_{j,i} \in \{0, 1\}$. In this case, we may refer to the set $S_j = \{i \in \mathcal{I} : u_{j,i} = 1\}$ as the list of items buyer j *wants* (one of).

2.2 Production cost

We suppose there are *cumulative cost functions* $\mathrm{cost}_i : \mathbb{N} \to [0, \infty]$ for each item $i \in \mathcal{I}$, where for $t \in \mathbb{N}$, the value of $\mathrm{cost}_i(t)$ represents the cost of producing t copies of item i. We suppose each $\mathrm{cost}_i(\cdot)$ is nondecreasing.

We would like to consider the case of *decreasing marginal cost*, where $t \mapsto \mathrm{cost}_i(t + 1) - \mathrm{cost}_i(t)$ is nonincreasing for each $i \in \mathcal{I}$.

A natural class of decreasing marginal costs we will be especially interested in are of the form $t^{-\alpha}$ for $\alpha \in [0, 1)$. That is, $\mathrm{cost}_i(t) = c_0 \sum_{\tau=1}^{t} \tau^{-\alpha}$.

2.3 Allocation problems

After processing the n buyers, we will have allocated some set of items T, consisting of $m_i(T) = \sum_{j \in N} \mathbb{I}_{T_j}(i)$ copies of each item $i \in \mathcal{I}$. We are then interested in two quantities in this setting: the *total (production) cost* $\mathrm{cost}(T) = \sum_{i \in \mathcal{I}} \mathrm{cost}_i(m_i(T))$ and the *social welfare* $SW(T) = \sum_{j \in N} u_j(T_j)$.

We are interested in several different objectives within this setting, each of which is some variant representing the trade-off between reducing total production cost while increasing social welfare.

In the *allocate all* problem, we have to allocate to each buyer $j \in N$ one item $i \in S_j$ (in the uniform demand setting): that is, $SW(T) = n$. The goal is to minimize the total cost $\mathrm{cost}(T)$, subject to this constraint.

The *allocate with budget* problem requires our total cost to never exceed a given limit b (i.e., $\text{cost}(T) \leq b$). Subject to this constraint, our objective is to maximize the social welfare $SW(T)$. For instance, in the uniform demand setting, this corresponds to maximizing the number of satisfied buyers (that get an item from their set S_j).

The objective in the *maximize social surplus* problem is to maximize the difference of the social welfare and the total cost (i.e., $SW(T) - \text{cost}(T)$).

3 Structural Results and Allocation Policies

We now present several results about the structure of optimal (and non-optimal but "reasonable") solutions to allocation problems in the setting of decreasing marginal costs. These will be important in our sample-complexity analysis because they allow us to focus on allocation policies that have inherent complexity that depends only on the number of *items* and not on the number of *customers*, allowing for the use of uniform convergence bounds. That is, a small random sample of customers will be sufficient to uniformly estimate the performance of these policies over the full set of customers.

3.1 Permutation and pricing policies

A *permutation policy* has a permutation π over $\mathcal{I}$ and is applicable in the case of uniform unit demand. Given buyer j arriving, we allocate to him the minimal (first) demanded item in the permutation, i.e., $\arg\min_{i \in S_j} \pi(i)$. A *pricing policy* assigns a price price_i to each item i and is applicable to general quasi-linear utility functions. Given buyer j arriving, we allocate to him whatever he wishes to purchase at those prices, i.e., $\arg\max_{T_j} u_j(T_j) - \sum_{i \in T_j} \text{price}_i$.[1]

We will see below that for uniform unit demand buyers, there always exists a permutation policy that is optimal for the allocate-all task, and for general quasilinear utilities there always exists a pricing policy that is optimal for the task of maximizing social surplus. We will also see that for weighted unit demand buyers, there always exists a pricing policy that is optimal for the allocate-with-budget task; moreover, for any even non-optimal solution (e.g., that might be produced by a polynomial-time algorithm) there exists a pricing policy that sells the same number of copies each item and has social welfare at least as high (and can be computed in polynomial time given the initial solution).

3.2 Structural results

Theorem 3.1. *For general quasilinear utilities, any allocation that maximizes social surplus can be produced by a pricing policy. That is, if* $\mathcal{T} = \{T_1, \ldots, T_n\}$ *is an allocation maximizing* $SW(T) - \text{cost}(T)$ *then there exist prices* $\text{price}_1, \ldots, \text{price}_r$ *such that buyers purchasing their most-demanded bundle recovers* $\mathcal{T}$*, assuming that the marginal cost function is strictly decreasing.*[2]

Proof. Consider the optimal allocation OPT. Define price_i to be the marginal cost of the next copy of item i under OPT, i.e., $\text{price}_i = \text{cost}_i(\#_i(\text{OPT}) + 1)$. Suppose some buyer j is assigned set T_j in OPT but

[1] When more that one subset is applicable, we assume we have the freedom to select any such set. Note that such policies are incentive-compatible.

[2] If the marginal cost function is only non-increasing, we can have the same result, assuming we can select between the utility maximizing bundles.

prefers set T'_j under these prices. Then,

$$u_j(T'_j) - \sum_{i \in T'_j} \text{price}_i \geq u_j(T_j) - \sum_{i \in T_j} \text{price}_i,$$

which implies

$$u_j(T'_j) - u_j(T_j) + \sum_{i \in T_j \setminus T'_j} \text{price}_i - \sum_{i \in T'_j \setminus T_j} \text{price}_i \geq 0. \tag{1}$$

Now, consider modifying OPT by replacing T_j with T'_j. This increases buyer j's utility by $u_j(T'_j) - u_j(T_j)$, incurs an extra purchase cost *exactly* $\sum_{i \in T'_j \setminus T_j} \text{price}_i$ and a savings of strictly more than $\sum_{i \in T_j \setminus T'_j} \text{price}_i$ (because marginal costs are decreasing). Thus, by (1) this would be a strictly preferable allocation, contradicting the optimality of OPT. □

Corollary 3.2. *For uniform unit demand buyers there exists an optimal allocation that is a permutation policy, for the* allocate all *task.*

Proof. Imagine each buyer j had valuation v_{max} on items in S_j where v_{max} is greater than the maximum cost of any single item. The allocation OPT that maximizes social surplus would then minimize cost subject to allocating exactly one item to each buyer and therefore would be optimal for the allocate-all task. Consider the pricing associated to this allocation given by Theorem 3.1. Since each buyer j is uniform unit demand, he will simply purchase the cheapest item in S_j. Therefore, the permutation π that orders items according to increasing price according to the prices of Theorem 3.1 will produce the same allocation. □

We now present a structural statement that will be useful for the allocate-with-budget task.

Theorem 3.3. *For weighted unit-demand buyers, for any allocation $\mathcal{T}$ there exists a pricing policy that allocates the same* multiset *of items T (or a subset of T) and has social welfare at least as large as $\mathcal{T}$. Moreover, this pricing can be computed efficiently from $\mathcal{T}$ and the buyers' valuations.*

Proof. Let T be the multiset of items allocated by $\mathcal{T}$. Weighted unit-demand valuations satisfy the gross-substitutes property, so by the Second Welfare Theorem (e.g., see [10] Theorem 11.15) there exists a Walrasian equilibrium: a set of prices for the items in T that clears the market. Moreover, these prices can be computed efficiently from demand queries (e.g., [10], Theorem 11.24), which can be evaluated efficiently for weighted unit-demand buyers. Furthermore, these prices must assign all copies of the *same* item in T the same price (else the pricing would not be an equilibrium) so it corresponds to a legal pricing policy. Thus, we have a legal pricing such that if all buyers were shown only the items represented in T, at these prices, then the market would clear perfectly (breaking any ties in our favor). We can address the fact that there may be items not represented in T (i.e., they had zero copies sold) by simply setting their price to infinity. Finally, by the First Welfare Theorem (e.g., [10] Theorem 11.13), this pricing maximizes social welfare over all allocations of T, and therefore achieves social welfare at least as large as $\mathcal{T}$, as desired. □

The above structural results will allow us to use the following sketch of an online algorithm. First sample an initial set of ℓ buyers. Then, for the allocate-all problem, compute the best (or approximately best) permutation policy according to the empirical frequencies given by the sample. Or, for the allocate-with budget task, compute the best (or approximately best) allocation according to these empirical frequencies and convert it into a pricing policy. Then run this permutation or pricing policy on the remainder of the customers. Finally, using the fact that these policies have low complexity (they are lists or vectors in a space that depends only on the number of items and not on the number of buyers) compute the size of initial sample needed to ensure that the estimated performance is close to true performance uniformly over all policies in the class.

4 Uniform Unit Demand and the Allocate-All problem

Here we consider the allocate-all problem for the setting of uniform unit demand. We begin by considering the following natural class of decreasing marginal cost curves such as $1/\sqrt{t}$.

Definition 4.1. *We say the cost function* $\mathrm{cost}(t)$ *is* α*-poly if the marginal cost of item t is* $1/t^{\alpha}$ *for* $\alpha \in [0, 1)$. *That is,* $\mathrm{cost}(t) = \sum_{\tau=1}^{t} 1/\tau^{\alpha}$.

Theorem 4.2. *If each cost function is* α*-poly, then there exists an efficient offline algorithm that given a set* X *of buyers produces a permutation policy that incurs total cost at most* $\frac{1}{1-\alpha}\mathrm{OPT}$.

Proof. We run the greedy set-cover algorithm. Specifically, we choose the item desired by the most buyers and put it at the top of the permutation π. We then choose the item desired by the most buyers who did not receive the first item and put it next, and so on. For notational convenience assume π is the identity, and let $\mathcal{B}_i$ denote the set of buyers that receive item i, For any set $\mathcal{B} \subseteq X$, let $\mathrm{OPT}(\mathcal{B})$ denote the cost of the optimal solution to the subproblem $\mathcal{B}$ (i.e., the problem in which we are only required to cover buyers in $\mathcal{B}$). Clearly $\mathrm{OPT}(\mathcal{B}_r) = \mathrm{cost}(|\mathcal{B}_r|) = \sum_{\tau=1}^{|\mathcal{B}_r|} 1/\tau^{\alpha} \geq \sum_{t=1}^{|\mathcal{B}_r|} \int_1^{|\mathcal{B}_t|} x^{-\alpha} \mathrm{d}x = \frac{1}{1-\alpha}|\mathcal{B}_r|^{1-\alpha} - 1$, since any solution using more than one set to cover the elements of $\mathcal{B}_r$ has at least as large a cost.

Now, for the purpose of induction, suppose that some $k \in \{2, \ldots, r\}$ has $\mathrm{OPT}(\bigcup_{t=k}^{r} \mathcal{B}_t) \geq \sum_{t=k}^{r} |\mathcal{B}_t|^{1-\alpha}$. Then, since $\mathcal{B}_{k-1}$ was chosen to be the largest subset of $\bigcup_{t=k-1}^{r} \mathcal{B}_t$ that can be covered by a single item, it must be that the sets used by any allocation for the $\bigcup_{t=k-1}^{r} \mathcal{B}_t$ subproblem achieving $\mathrm{OPT}(\bigcup_{t=k-1}^{r} \mathcal{B}_t)$ have size at most $|\mathcal{B}_{k-1}|$, and thus the marginal costs for each of the elements of $\mathcal{B}_{k-1}$ in the $\mathrm{OPT}(\bigcup_{t=k-1}^{r} \mathcal{B}_t)$ solution is at least $1/|\mathcal{B}_{k-1}|^{\alpha}$.

This implies $\mathrm{OPT}(\bigcup_{t=k-1}^{r} \mathcal{B}_t) \geq \mathrm{OPT}(\bigcup_{t=k}^{r} \mathcal{B}_t) + \sum_{x \in \mathcal{B}_{k-1}} 1/|\mathcal{B}_{k-1}|^{\alpha} = \mathrm{OPT}(\bigcup_{t=k}^{r} \mathcal{B}_t) + |\mathcal{B}_{k-1}|^{1-\alpha}$. By the inductive hypothesis, this latter expression is at least as large as $\sum_{t=k-1}^{r} |\mathcal{B}_t|^{1-\alpha}$. By induction, this implies $\mathrm{OPT}(X) = \mathrm{OPT}(\bigcup_{t=1}^{r} \mathcal{B}_t) \geq \sum_{t=1}^{r} |\mathcal{B}_t|^{1-\alpha}$. On the other hand, the total cost incurred by the greedy algorithm is $\sum_{t=1}^{r} \sum_{\tau=1}^{|\mathcal{B}_r|} 1/\tau^{\alpha} \leq \sum_{t=1}^{r} \int_0^{|\mathcal{B}_t|} x^{-\alpha} \mathrm{d}x = \frac{1}{1-\alpha} \sum_{t=1}^{r} |\mathcal{B}_t|^{1-\alpha}$. By the above argument, this is at most $\frac{1}{1-\alpha}\mathrm{OPT}(X)$. ☐

More general cost curves We can generalize the above result to a broader class of smoothly decreasing cost curves. Define the average cost of item i given to set $\mathcal{B}_i$ of buyers as $AvgC(i, |\mathcal{B}_i|) = \frac{\mathrm{cost}(|\mathcal{B}_i|)}{|\mathcal{B}_i|}$. Define the marginal cost $MarC(i, t) = \mathrm{cost}_i(t) - \mathrm{cost}_i(t-1)$. Here is a greedy algorithm.

Algorithm: $GreedyGeneralCost(\mathcal{B})$
0. $i = \arg\min AvgC(i, |\mathcal{B}_i|)$, where $\mathcal{B}_i = \{j \in \mathcal{B} : i \in S_j\}$
1. Call $GreedyGeneralCost(\mathcal{B} - \mathcal{B}_i)$

We make the following assumption:

Assumption 4.3. $\forall i, t$, $AvgC(i, t) \leq \beta MarC(i, t)$*, for some* $\beta > 0$.

For example, for the case of an α-poly cost, we have: $MarC(t) = \frac{1}{t^{\alpha}}$ and $AvgC = \frac{1}{t}\sum_{\tau=1}^{t} \frac{1}{\tau^{\alpha}} \approx \frac{t^{-\alpha}}{1-\alpha}$; so, therefore we have $\beta = \frac{1}{1-\alpha}$.

Theorem 4.4. *The algorithm GreedyGeneralCost achieves approximation ratio β.*

Proof. Order the elements in the order that GreedyGeneralCost allocates them. Let N_i be the set of consumers that receive item i, and $N = \cup N_i$ in GreedyGeneralCost. For consumer j let $item_{opt}(j)$ be the item that OPT allocates to consumer j. Let $\ell_{opt}(i)$ be the number of consumers that are allocated item i. By Assumption 4.3 we have $MarC(i,l) \leq AvgC(i,l) \leq \beta MarC(i,l)$ (the first inequality is due to having decreasing marginal cost).

We would like to consider the influence of the consumers in N_1 on the cost of OPT. Let

$$\begin{aligned} \mathrm{OPT}(N) - \mathrm{OPT}(N - N_1) \geq & \sum_{j \in N_1} MarC(item_{opt}(j), \ell_{opt}(item_{opt}(j))) \\ \geq & \sum_{j \in N_1} \frac{1}{\beta} AvgC(item_{opt}(j), \ell_{opt}(item_{opt}(j))) \\ \geq & \frac{1}{\beta} |N_1| AvgC(1, |N_1|) = \frac{1}{\beta} GreedyCost(N_1) \end{aligned}$$

The first inequality follows since taking the final marginal cost can only reduce the cost (decreasing marginal cost). The second inequality follows from Assumption 4.3. The third inequality follows since GreedyGeneralCost selects the lowest average cost of any allocated item .

We can now continue inductively. Let $T_0 = N$, $T_1 = N - N_1$, and $T_j = T_{j-1} - N_j$. We can show similarly that,

$$\mathrm{OPT}(T_{j-1}) - \mathrm{OPT}(T_j) \geq \frac{1}{\beta} GreedyCost(N_j)$$

Summing over all j we have

$$OPT(T) - OPT(\emptyset) = \sum_j \mathrm{OPT}(T_{j-1}) - \mathrm{OPT}(T_j) \geq \frac{1}{\beta} \sum_j GreedyCost(N_j) = \frac{1}{\beta} GreedyCost(N)$$

□

Additionally, a property of β-nice cost functions we will need to use later is the following.

Lemma 4.5. *For* cost *satisfying Assumption 4.3,* $\forall x \in \mathbb{N}$, $\forall \epsilon \in (0,1)$, $\forall i \leq r$, $\mathrm{cost}_i(\epsilon x) \leq \epsilon^{\log_2(1+\frac{1}{2\beta})} \mathrm{cost}_i(x)$.

Proof. By the fact that marginal costs are non-negative, $AvgC(2\epsilon x) \geq \mathrm{cost}_i(\epsilon x)/(2\epsilon x)$. Therefore, by Assumption 4.3, $MarC(2\epsilon x) \geq \mathrm{cost}_i(\epsilon x)/(2\epsilon x \beta)$. By the decreasing marginal cost property, we have

$$\mathrm{cost}_i(2\epsilon x) \geq \mathrm{cost}_i(\epsilon x) + \epsilon x MarC(2\epsilon x) \geq \mathrm{cost}_i(\epsilon x) + \mathrm{cost}_i(\epsilon x)/(2\beta) = (1 + \frac{1}{2\beta}) \mathrm{cost}_i(\epsilon x).$$

Applying this argument $\log_2(1/\epsilon)$ times, we have $\mathrm{cost}_i(x) \geq (1+\frac{1}{2\beta})^{\log_2(1/\epsilon)} \mathrm{cost}_i(\epsilon x) = (\frac{1}{\epsilon})^{\log_2(1+\frac{1}{2\beta})} \mathrm{cost}_i(\epsilon x)$. Multiplying both sides by $\epsilon^{\log_2(1+\frac{1}{2\beta})}$ completes the proof. □

4.1 Generalization Result

Say n is the total number of customers; ℓ is the size of subsample we do our estimate on; r is the total number of items; $\alpha \in (0,1]$ is some constant, and the cost is α-poly, so that $\mathrm{cost}(t) = \sum_{\tau=1}^{t} 1/\tau^{\alpha} \simeq \frac{t^{1-\alpha}}{1-\alpha}$. We now show the following uniform convergence over permutation policies, which will justify the use of a near optimal policy for a sample on the larger population; the formal proof of this result is included in Appendix C.

Theorem 4.6. *Suppose $n \geq \ell$ and the cost function is α-poly. With probability at least $1 - \delta^{(\ell)}$, for all permutations Π,*

$$\mathrm{cost}(\Pi, \ell)(1+\epsilon)^{-2}\left(\frac{n}{\ell}\right)^{1-\alpha} \leq \mathrm{cost}(\Pi, n) \leq \mathrm{cost}(\Pi, \ell)(1+\epsilon)^{2(1-\alpha)}\left(\frac{n}{\ell}\right)^{1-\alpha},$$

where $\delta^{(\ell)} = r2^r(\delta_1 + \delta_2 + \delta_3)$ and $\delta_1 = \exp\{-\epsilon^2\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}} n/3\}$, $\delta_2 = \exp\{-\epsilon^2\ell\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}}/3\}$, $\delta_3 = \exp\{-\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}} n\epsilon^2/2\}$. Equivalently, for any given δ, this occurs with probability $1-\delta$ so long as $\ell \gg (\frac{1}{\epsilon^2})(\frac{r}{\epsilon})^{\frac{1}{1-\alpha}}\ln(2^r/\delta)$.

Proof. See Appendix C. □

4.2 Online Performance Guarantees

We define $GreedyGeneralCost(\ell, n)$ as follows. For the first ℓ customers it allocates arbitrary items they desire, and observed their desired sets. Give the sets of the first ℓ customers, it runs $GreedyGeneralCost$ and computes a permutation $\hat{\Pi}$ of the items. For the remaining customers it allocates using permutation $\hat{\Pi}$. Namely, each customer is allocated the first item in the permutation $\hat{\Pi}$ that is in its desired set. The following theorem bounds the performance of $GreedyGeneralCost(\ell, n)$ for α-poly cost functions; the formal proof of this result is included in Appendix C.

Theorem 4.7. *With probability $1-\delta^{(\ell)}$ (for $\delta^{(\ell)}$ as in Theorem 4.6) The cost of $GreedyGeneralCost(\ell, n)$ is at most*

$$\ell + \frac{(1+\epsilon)^{4-2\alpha}}{1-\alpha}\mathrm{OPT}$$

Proof. See Appendix C. □

Corollary 4.8. *For any fixed constant $\delta \in (0,1)$, for any $\ell \geq \frac{3}{\epsilon^2}\left(\frac{r}{\epsilon}\right)^{\frac{1}{1-\alpha}}\ln\left(\frac{3r2^r}{\delta}\right)$, and $n \geq \left(\frac{\ell}{\epsilon}\right)^{\frac{1}{1-\alpha}}$, with probability at least $1-\delta$, $GreedyGeneralCost(n, \ell)$ is at most* $\left(\frac{(1+\epsilon)^{4-2\alpha}}{1-\alpha} + \epsilon\right)\mathrm{OPT}$

4.3 Generalization for β-nice costs

We now consider the case of β-nice costs in the online setting. We begin with a helper lemma.

Lemma 4.9. *For any cost* cost *satisfying Assumption 4.3 with a given β, for any $k \geq 1$, the cost* cost' *with* $\mathrm{cost}'_i(x) = \mathrm{cost}_i(kx)$ *also satisfies Assumption 4.3 with the same β.*

Proof. See Appendix. □

Now the strategy is to run GreedyGeneralCost with the rescaled cost function $\mathrm{cost}'_i(x) = \mathrm{cost}_i(\frac{n}{\ell}x)$. This provides a β-approximation guarantee for the rescaled problem, which, moreover is a permutation policy. The following shows we have uniform convergence of estimates to true costs for permutation policies.

Theorem 4.10. *Suppose $n \geq \ell$ and the cost function satisfies Assumption 4.3, and that $\forall i$, $\mathrm{cost}_i(1) \in [1, B]$, where $B \geq 1$ is constant. Let $\mathrm{cost}'_i(x) = \mathrm{cost}_i(\frac{n}{\ell}x)$. With probability at least $1-\delta^{(\ell)}$, for any permutations Π,*

$$\mathrm{cost}'(\Pi, \ell)\frac{1-\epsilon}{1+2\epsilon-\epsilon^2} \leq \mathrm{cost}(\Pi, n) \leq \mathrm{cost}'(\Pi, \ell)\frac{(1+\epsilon)^2}{1-\epsilon},$$

where $\delta^{(\ell)} = r^2 2^{r+1}(\delta_1+\delta_2)$ and $\delta_1 = \exp\{-\frac{\epsilon^3}{3rB(1+\epsilon)} n^{\log_2(1+\frac{1}{2\beta})}\}$, $\delta_2 = \exp\{-\ell\frac{\epsilon^3}{rB(1+\epsilon)} n^{\log_2(1+\frac{1}{2\beta})-1}/3\}$.

Proof. See Appendix. □

Using the above uniform convergence result, we find that if we run GreedyGeneralCost on the sample of $\ell = o(n)$ initial buyers and apply it to the entire population, we achieve near optimal cost.

Theorem 4.11. *If* cost *satisfies Assumption 4.3, and has* $\mathrm{cost}_j(1) \in [1, B]$ *for every* $j \leq r$*, with probability at least* $1 - \delta$*, the cost of applying the policy found by* $GreedyGeneralCost(\{1, \ldots, \ell\})$ *to all* n *customers is at most*

$$\beta \frac{(1+\epsilon)^2(1+2\epsilon-\epsilon^2)}{(1-\epsilon)^2}\mathrm{OPT}(n),$$

where $\ell = \left\lceil n^{1-\log_2(1+\frac{1}{2\beta})} \frac{3rB(1+\epsilon)}{\epsilon^3} \ln\left(\frac{r^2 2^{r+2}}{\delta}\right)\right\rceil = o(n)$.

Proof. By Theorem 4.4, Lemma 4.9, and Theorem 4.10, with probability at least $1 - \delta$, the cost of applying the policy $\hat{\Pi}$ found by $GreedyGeneralCost(\{1, \ldots, \ell\})$ to customers $1, \ldots, n$ is at most

$$\begin{aligned}
\mathrm{cost}'(\hat{\Pi}, \ell)\frac{(1+\epsilon)^2}{1-\epsilon} &\leq \beta \min_{\Pi} \mathrm{cost}'(\Pi, \ell)\frac{(1+\epsilon)^2}{1-\epsilon} \\
&\leq \beta \min_{\Pi} \mathrm{cost}(\Pi, n)\frac{(1+\epsilon)^2(1+2\epsilon-\epsilon^2)}{(1-\epsilon)^2} \\
&= \beta \frac{(1+\epsilon)^2(1+2\epsilon-\epsilon^2)}{(1-\epsilon)^2}\mathrm{OPT}(n).
\end{aligned}$$

□

Note that Theorem 4.11 assumes the initial $\ell = o(n)$ buyers can be "previewed" before allocations are made and need not themselves be allocated online.

5 General Unit Demand Utilities

In this section we show how to give a constant approximation for the case of general unit demand buyers in the offline setting in the case when we have a budget B to bound the cost we incur and we would like to maximize the buyers social welfare given this budget constraint. The main tool would be a reduction of our problem to the budgeted maximum coverage problem.

Definition 5.1. *An instance of the* budgeted maximum coverage problem *has a universe* X *of* m *elements where each* $x_i \in X$ *has an associated weight* w_i*; there is a collection of* m *sets* $\mathcal{B}$ *such that each sets* $S_j \in \mathcal{B}$ *has a cost* c_j*; and there is a budget* L*. A feasible solution is a collection of sets* $\mathcal{B}' \subset \mathcal{B}$ *such that* $\sum_{S_j \in \mathcal{B}'} c_j \leq L$*. The goal is to maximize the weight of the elements in* $\mathcal{B}'$*, i.e.,* $w(\mathcal{B}') = \sum_{x_i \in \cup_{S \in \mathcal{B}'} S} w_i$.

While the budgeted maximum coverage problem is NP-complete there is a $(1 - 1/e)$ approximation algorithm [8]. Their algorithm is a variation of the greedy algorithm, where on the one hand it computes the greedy allocation, where each time a set which maximizes the ratio between weight of the elements covered and the cost of the set is added, as long as the budget constraint is not violated. On the other hand the single best set is computed. The output is the best of the two alternative (either the single best set of the greedy allocation).

Before we show the reduction from a general unit demand utility to the budgeted maximum coverage problem, we show a simpler case where for each buyer j has a value v_j such that of any item i either $v_j = u_{j,i}$ or $u_{j,i} = 0$, which we call *buyer-uniform unit demand.*

Lemma 5.2. *There is a reduction from the budgeted buyer-uniform unit demand buyers problem to the budgeted maximum coverage problem. In addition the greedy algorithm can be computed in polynomial time on the resulting instance.*

Proof. For each buyer j we create an element x_j with weight v_j. For each item k and any subsets of buyers S we create a set $T_{S,k} = \{x_j : j \in S\}$ and has cost $cost_k(|S|)$. The budget is set to be $L = B$. Clearly any feasible allocation of the budgeted maximum coverage problem $T_{S_1,k_1}, \ldots T_{S_r,k_r}$ can be translated to a solution of the budgeted buyer-uniform unit demand buyers by simply producing item k_i for all the buyers in T_{S_i,k_i}. The welfare is the sum of the weight of the elements covered which is the social welfare, and the cost is exactly the production cost.

Note that the reduction generates an exponential number of sets, if we do it explicitly. However,we can run the Greedy algorithm easily, without generating the sets explicitly. Assume we have m' remaining buyers. For each item i and any $\ell \in [1, m']$ we compute the cost $cost_i(\ell)/gain_i(\ell)$, where $gain_i(\ell)$ is the weight of the ℓ buyers with highest valuation for item i. Greedy select the item i and number of buyers ℓ which have the highest ratio and adding this set still satisfies the budget constraint. Note that given that greedy selects $T_{S,k}$ where $|S| = \ell$ then its cost is $cost_k(\ell)$ and its weigh is $w(T_{S,k}) \leq gain_k(\ell)$, and hence Greedy will always select one of the sets we are considering. □

In the above reduction we used very heavily the fact that each buyer j has a single valuation v_j regardless of which desired item it gets. In the following we show a slightly more involved reduction which handles the general unit demand buyers.

Lemma 5.3. *There is a reduction from the budgeted general unit demand buyers problem to the budgeted maximum coverage problem. In addition the greedy algorithm can be computed in polynomial time on the resulting instance.*

Proof. For each buyer j we sort its valuations $u_{j,i_1} \leq \cdots \leq u_{j,i_m}$. We set $v_{j,i_1} = u_{j,i_1}$ and $v_{j,i_r} = u_{j,i_r} - u_{j,i_{r-1}}$. Note that $\sum_{s=1}^{r} v_{j,i_s} = u_{j,i_r}$. For each buyer j we create m elements $x_{j,r}$, $1 \leq r \leq m$. For a buyer j and item k let $X_{j,k}$ be all the elements that represent lower valuation than $u_{j,k}$, i.e., $X_{j,k} = \{x_{j,r} : u_{j,i_r} \leq u_{j,k}\}$. For each item k and any subsets of buyers S we create a set $T_{S,k} = \cup_{j \in S} X_{j,k}$ and has cost $cost_k(|S|)$. The budget is set to be $L = B$.

Any feasible allocation of the budgeted maximum coverage problem $T_{S_1,k_1}, \ldots T_{S_l,k_r}$ can be translated to a solution of the budgeted general unit demand buyers producing item k_i for all the buyers in T_{S_i,k_i}. We call buyer j as *winner* if there exists some b such that $x_{j,b} \in \cup_{i=1}^{r} T_{S_i,k_i}$. Let $Winners$ we the set of all winner buyers. For any winner buyer $j \in Winner$ let $item(j) = s$ such that $s = \max\{b : x_{j,b} \in \cup_{i=1}^{r} T_{S_i,k_i}\}$.

The cost of our allocation is by definition at most $L = B$. The social welfare is

$$\sum_{x_{j,b} \in \cup_{i=1}^{r} T_{S_i,k_i}} v_{j,b} = \sum_{j \in Winner} u_{j,item(j)}$$

Again, note that the reduction generates an exponential number of sets, if we do it explicitly. However, we can run the Greedy algorithm easily, without generating the sets explicitly. For each item i and any $\ell \in [1, m]$ we compute the cost $cost_i(\ell)/gain_i(\ell)$, where $gain_i(\ell)$ is the weight of the ℓ buyers with highest valuation for item i. Greedy selects the item i and number of buyers ℓ which have the highest ratio which still satisfies the budget constraint. Note that given that greedy selects $T_{S,k}$ where $|S| = \ell$ then its production cost is $cost_k(\ell)$ and its weight is $w(T_{S,k}) \leq gain_k(\ell)$, and hence Greedy will always select one of the sets we are considering. Once the Greedy selects a set $T_{S,k}$ we need to update the utility of any buyer $j \in S$

for any other item i, by setting $u_{j,i} = \max\{u_{j,i} - u_{j,k}, 0\}$, which is the residual valuation buyer j has for getting item i in addition to item k. □

Combining our reduction with approximation algorithm of [8] we have the following theorem.

Theorem 5.4. *There exists a poly-time algorithm for the budgeted general unit demand buyers problem which achieves social welfare at least* $(1 - 1/e)\mathrm{OPT}$.

5.1 Generalization

To extend these results to the online setting, we will use Theorem 3.3 to represent allocations by pricing policies, and then use the results from above to learn a good pricing policy based on an initial sample.

Theorem 5.5. *Suppose every* $u_{j,i} \in [0, C]$. *With* $\ell = O((1/\epsilon^2)(r^3 \log(rC/\epsilon) + \log(1/\delta)))$ *random samples, with probability at least* $1 - \delta$*, the empirical per-customer social welfare is within* $\pm\epsilon$ *of the expected per-customer social welfare, uniformly over all price vectors in* $[0, C]^r$.

Proof. We will show that, for any distribution P and value $\epsilon > 0$, there exist $N = 2^{O(r^3 \log(rC/\epsilon))}$ functions $f_1, \ldots, f_N$ such that, for every price vector $\mathrm{price} \in [0, C]^r$, the function $g(x) = x_{\arg\max_{i \le r} x_i - \mathrm{price}_i}$ has $\min_{k \le N} \int |f_k - g| \mathrm{d}P \le \epsilon$. This value N is known as the *uniform* ϵ*-covering number*. The result then follows from standard uniform convergence bounds (see e.g., [6]).

First, note that if we suppose the valuations and price vectors are augmented by one dimension, setting $\mathrm{price}_{r+1} = u_{j,r+1} = 0$ for all price vectors and all customers j, then we can represent the per-customer social welfare as $x \mapsto \mathrm{sw}(x; \mathrm{price}) = x_{\arg\max_i (x_i - \mathrm{price}_i)}$.

The function $x \mapsto x_i - \mathrm{price}_i$ is a hyperplane with slope 1 in coordinate i and slope 0 in all other coordinates. So the subgraph (i.e., the set of $(r + 2)$-dimensional points (x, y) for which $\max_{i \le r+1} x_i - \mathrm{price}_i \ge y$ is a union of r halfspaces in $r + 2$ dimensions. The space of unions of $r + 1$ halfspaces in $r + 2$ dimensions has VC dimension $(r + 1)(r + 3)$, so this upper bounds the pseudo-dimension of the space of functions $\max_{i \le r+1} x_i - \mathrm{price}_i$, parametrized by the price vector price. Therefore, the uniform ϵ-covering number of this class is $2^{O(r^2 \log(C/\epsilon))}$.

For each $i \le r+1$, the set of vectors $x \in [0, C]^{r+1}$ such that $i = \arg\max_k x_k - \mathrm{price}_k$ is an intersection of $r + 1$ halfspaces in $r + 1$ dimensions. Thus, the function $x \mapsto \mathrm{price}_{\arg\max_i x_i - \mathrm{price}_i}$ is contained in the family of linear combinations of $r + 1$ disjoint intersections of $r + 1$ halfspaces. The VC dimension of an intersection of $r + 1$ halfspaces in $r + 1$ dimensions is $(r + 1)(r + 2)$. So assuming the prices are bounded in a range $[0, C]$, the uniform ϵ-covering number for linear combinations (with weights in $[0, C]$) of $r + 1$ disjoint intersections of $r + 1$ halfspaces is $2^{O(r^3 \log(rC/\epsilon))}$. To prove this, we can take an $\epsilon/(2(r + 1)C)$ cover (of $\{0, 1\}$-valued functions) of intersections of $r + 1$ halfspaces, which has size $((r + 1)C/\epsilon)^{O(r^2)}$, and then take an $\epsilon/(2(r + 1))$ grid in $[0, C]$ and multiply each function in the cover by each of these values to get a space of real-valued functions; there are $((r + 1)C/\epsilon)^{O(r^2)}$ total functions in this cover, and for each term in the linear combination of $r + 1$ disjoint intersections of $r + 1$ halfspaces, at least one of these real-valued functions will be within $\epsilon/(r + 1)$ of it. Thus, taking the set of sums of $r + 1$ functions from this cover forms an ϵ-cover of the space of linear combinations of $r + 1$ disjoint intersections of $r + 1$ halfspaces, with size $((r + 1)C/\epsilon)^{O(r^3)}$.

Now note that $x_{\arg\max_i (x_i - \mathrm{price}_i)} = \max_i (x_i - \mathrm{price}_i) + \mathrm{price}_{arg\max_i (x_i - \mathrm{price}_i)}$. So the uniform ϵ-covering number for the space of possible functions $x_{\arg\max_i (x_i - \mathrm{price}_i)}$ is at most the product of the uniform $(\epsilon/2)$-covering number for the space of functions $x \mapsto \max_i (x_i - \mathrm{price}_i)$ and the uniform $(\epsilon/2)$-covering number for the space of functions $x \mapsto \mathrm{price}_{\arg\max_i (x_i - \mathrm{price}_i)}$; by the above, this product is $2^{O(r^3 \log(rC/\epsilon))}$. □

We also make use of the following result.

Theorem 5.6. *With* $\ell \geq O((1/\epsilon^2)(r^2 + \log(1/\delta)))$ *random samples, with probability at least* $1 - \delta$*, the empirical probability of a customer buying item* j *is within* $\pm\epsilon$ *of the actual probability, uniformly over all price vectors and all* j.

Proof. For a given price vector, the region of customers j purchasing item i is delineated by at most r hyperplanes (corresponding to i tying with another index for maximizing $u_{j,i} - \mathrm{price}_i$, or else $u_{j,i} - \mathrm{price}_i = 0$ and all the other indices make it at most 0). Thus it is an intersection of r hyperplanes in r dimensions. The VC dimension of this is class is $r(r+1)$, so the VC dimension of the collection of sets of customers buying item i as we vary price is at most $r(r+1)$. So by the standard uniform convergence results for convergence of frequencies to probabilities, $O((1/\epsilon^2)(r^2 + \log(r/\delta))$ samples suffice so that the empirical probability of getting a customer that buys item i is within $\pm\epsilon$ of the actual probability, uniformly over all price vectors, with probability at least $1 - \delta/r$. A union bound implies this holds simultaneously for all i with probability at least $1 - \delta$. Noting that $(1/\epsilon^2)(r^2 + \log(r/\delta)) = O((1/\epsilon^2)(r^2 + \log(1/\delta)))$ completes the proof. □

Consider an algorithm that does not allocate anything to the first $\ell = O((C^2/\epsilon^2)(r^3 \log(rC/\epsilon) + \log(1/\delta)))$ customers, then finds a $(1 - 1/e)$-approximate solution to the offline budgeted general unit demand problem on these ℓ customers, with budget B, and cost functions $\mathrm{cost}'_i(x) = \mathrm{cost}_i(x \cdot ((n-\ell)/\ell))$, via the reduction to the budgeted maximum coverage problem. The algorithm then finds a pricing policy price providing at least as good of a social welfare on these ℓ customers, within this budget B. Let ℓ_i denote the number of copies of item i this pricing policy allocates among the ℓ customers. The algorithm then proceeds to allocate to the remaining stream of $n - \ell$ customers using this pricing policy, but if at any time the item i this pricing policy determines should be allocated to the next customer has already had $\ell_i((n-\ell)/\ell)$ copies allocated to customers in the past, then the algorithm does not allocate any item to that customer and simply moves on to the next customer. (As stated, this is not incentive-compatible: we are assuming that if a buyer enters the store and finds his most-desired item is sold-out, he just leaves rather than buying some other item; however, we rectify this in Corollary 5.8 below.) We have the following result on the performance of this algorithm.

Theorem 5.7. *The allocation given by the above algorithm does not exceed the budget* B*, and if* $n \geq O((1/\epsilon)\ell)$*, with probability at least* $1 - 4\delta$*, achieves a social welfare at least*

$$(1 - 1/e)\mathrm{OPT} - (2(2 - 1/e)(1 + Cr) + C)\epsilon n.$$

Proof. See Appendix D. □

To make the above procedure incentive-compatible, if at any time the pricing policy attempts to allocate more than $\ell_i((n-\ell)/\ell)$ copies of item i, then for that customer j we can just allocate the item i' that has the next-highest $u_{j,i'} - \mathrm{price}_{i'}$ among those i' for which the number of copies of item i' this policy has attempted to allocate previously is less than $\ell_{i'}((n-\ell)/\ell)$ (or nothing, if all remaining i' have $u_{j,i'} - \mathrm{price}_{i'} < 0$). A simple modification of the above proof yields the following result on the performance of this algorithm.

Corollary 5.8. *The allocation given by the above algorithm does not exceed the budget* B*, and if* $n \geq O((1/\epsilon)\ell)$*, with probability at least* $1 - 4\delta$*, the allocation achieves a social welfare at least*

$$(1 - 1/e)\mathrm{OPT} - O(Cr^2\epsilon n).$$

Proof Sketch. The full details of this proof are provided in Appendix D; here, we present a brief sketch. As above, on the $(1-4\delta)$-probability event of Theorem 5.7, we also have that for each i, $n_i \le \ell_i \frac{n-\ell}{\ell} + 2\epsilon(n-\ell) \le \ell_i \frac{n-\ell}{\ell} + 2\epsilon n$. Let us now define T_i to be the time at which item i runs out; i.e., the index of the buyer who purchases copy $\ell_i \frac{n-\ell}{\ell}$ of that item (or ∞ if the item never runs out) and wlog assume $T_1 \le T_2 \le \cdots \le T_r$. We now argue that there are at most $2\epsilon n(i-1)$ buyers who were given item i in the procedure of Theorem 5.7 but receive some different item now. Specifically, consider one such buyer. The copy of item i that this buyer originally received must have been purchased by some other buyer whose preferred item was some $i' \le i$ (because items of index greater than i still remain). This buyer, under the procedure of Theorem 5.7 either received item i' or nothing. If it received item i' we continue back up the chain, examining the buyer who took that copy of item i' and must have a preferred item $i'' \le i'$, and so on until we reach a buyer who received nothing under the procedure of Theorem 5.7. Because each buyer only purchases one item, these chains cannot merge, and so the total number of buyers who were given item i in the procedure of Theorem 5.7 but receive some different item now, is at most the number of possible endpoints of such chains, which is $2\epsilon n(i-1)$. Summing this over all r items yields the claimed bound. □

References

[1] N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor. The online set cover problem. *SIAM Journal on Computing*, 39(2):361–370, 2009.

[2] A. Blum, A. Gupta, Y. Mansour, and A. Sharma. Welfare and profit maximization with production costs. In *FOCS*, pages 77–86, 2011.

[3] N. R. Devanur and T. P. Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *Proc. ACM EC*, EC '09, pages 71–78, 2009.

[4] N. R. Devanur and K. Jain. Online matching with concave returns. In *Proc. STOC*, pages 137–144, 2012.

[5] G. Goel and A. Mehta. Online budgeted matching in random input models with applications to adwords. In *Proc. SODA*, pages 982–991, 2008.

[6] D. Haussler. Decision theoretic generalizations of the PAC model for neural net and other learning applications. *Information and Computation*, 100:78–150, 1992.

[7] R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proc. STOC*, pages 352–358, 1990.

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

[9] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5), 2007.

[10] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.

A Properties of β-nice cost

Let $cost(n)$ be a β-nice cost function. We show a few properties of it.

Lemma 4.9. *For any cost* cost *satisfying Assumption 4.3 with a given β, for any $k \geq 1$, the cost* cost' *with* $\mathrm{cost}'_i(x) = \mathrm{cost}_i(kx)$ *also satisfies Assumption 4.3 with the same β.*

Proof.

$$\frac{\mathrm{cost}_i(kx)}{x} = k\frac{\mathrm{cost}_i(kx)}{kx} \leq \beta k(\mathrm{cost}_i(kx) - \mathrm{cost}_i(kx-1)).$$

Also, the property of nonincreasing marginal costs implies $\forall t \in \{1, \ldots, k\}$,

$$\mathrm{cost}_i(kx) - \mathrm{cost}_i(kx-1) \leq \mathrm{cost}_i(kx-(t-1)) - \mathrm{cost}_i(kx-t),$$

so that

$$k(\mathrm{cost}_i(kx) - \mathrm{cost}_i(kx-1)) \leq \sum_{t=1}^{k}(\mathrm{cost}_i(kx-(t-1)) - \mathrm{cost}_i(kx-t)) = \mathrm{cost}_i(kx) - \mathrm{cost}_i(k(x-1)).$$

Therefore,

$$\frac{\mathrm{cost}_i(kx)}{x} \leq \beta(\mathrm{cost}_i(kx) - \mathrm{cost}_i(k(x-1))).$$

□

Claim A.1.

$$cost(2n) \geq cost(n)\left(1 + \frac{1}{2\beta}\right)$$

Proof. Let $a = cost(n)/n$ be the average cost of the first n items. Then the cost of the first $2n$ items is at least an, and has an average cost of at least $a/2$. The marginal cost of item $2n$ is at least $a/(2\beta)$. Therefore the cost of the items $n+1$ to $2n$ is at least $an/(2\beta)$. □

We can get a better bound by a more refine analysis.

Claim A.2. *Let $a_n = cost(n)/n$ be the average cost of the first n items. Then,*

$$a_{n+1} \geq a_n \frac{n}{n+1}\left(1 + \frac{1}{\beta(n+1)}\right)$$

and

$$a_n \geq a_1 \frac{1}{n}\prod_{t=1}^{n}\left(1 + \frac{1}{\beta(t+1)}\right) \geq e^{1/\beta^2} \cdot a_1 n^{-1+(1/\beta)}$$

Proof. The marginal cost of item $n+1$ is at least a_n/β. Therefore the cost of the first items $n+1$ is at least $na_n + a_n/(\beta)$, which gives the first expression.

We get the expression of a_n as a function of a_1 by repeatedly using the recursion. The approximation follows from,

$$\begin{aligned}\ln(a_n) &\geq \ln(a_1) - \ln(n) + \sum_{t=1}^{n} \ln(1 + \frac{1}{\beta(n+1)}) \\ &\geq \ln(a_1) - \ln(n) + \sum_{t=1}^{n} \frac{1}{\beta(t+1)} - \frac{1}{(\beta(t+1))^2} \\ &\geq \ln(a_1) - \ln(n) + \frac{1}{\beta}\ln(n) - \frac{1}{\beta^2}\end{aligned}$$

where we used the identity $x - x^2 \leq \ln(1+x)$. □

B Additional Proofs

Theorem 4.10. *Suppose $n \geq \ell$ and the cost function satisfies Assumption 4.3, and that $\forall i$, $\mathrm{cost}_i(1) \in [1, B]$, where $B \geq 1$ is constant. Let $\mathrm{cost}'_i(x) = \mathrm{cost}_i(\frac{n}{\ell}x)$. With probability at least $1 - \delta^{(\ell)}$, for any permutations Π,*

$$\mathrm{cost}'(\Pi, \ell)\frac{1-\epsilon}{1+2\epsilon-\epsilon^2} \leq \mathrm{cost}(\Pi, n) \leq \mathrm{cost}'(\Pi, \ell)\frac{(1+\epsilon)^2}{1-\epsilon},$$

where $\delta^{(\ell)} = r^2 2^{r+1}(\delta_1 + \delta_2)$ and $\delta_1 = \exp\{-\epsilon^3 n^{\log_2(1+\frac{1}{2\beta})}/(3rB(1+\epsilon))\}$, $\delta_2 = \exp\{-\epsilon^2 \ell \frac{\epsilon}{rB(1+\epsilon)} n^{\log_2(1+\frac{1}{2\beta})-1}/3\}$.

Proof. Fix a permutation Π. Let π_j denote the event that a customer buys item Π_j and not covered by items Π_1 through Π_{j-1}. Namely, the probability that the consumer set of desired items include j and none of the items $1, \ldots, j-1$. Let q_j denote $Pr[\pi_j]$, and let $\hat{q}_j$ denote the fraction of Π_j on the initial ℓ-sample.

Let $q^* = \frac{\epsilon}{rB(1+\epsilon)} n^{c-1}$, where $c = \log_2(1 + \frac{1}{2\beta})$. Item j is a "Low probability item" if $q_j < q^*$, and is called a "High probability item" if $q_j \geq q^*$. Let the set "Low" include all "Low probability items"; and the set "High" include all "High probability items".

First we address the case of item j of low probability. By a Chernoff bound, the quantity of item j that we will sell when applying Π to n customers is at most $q^* n(1+\epsilon)$, with probability at least $1 - \exp\{-\epsilon^2 q^* n/3\} = 1 - \delta_1$. By a union bound, this holds for all low probability items j with probability at least $1 - |\mathrm{Low}|\delta_1$.

Next, suppose j has high probability. In this case, the quantity of item j we will sell when applying Π to n customers is at most $q_j n(1+\epsilon)$, with probability at least $1 - \exp\{-\epsilon^2 q_j n/3\} \geq 1 - \delta_1$. Again, a union bound implies this holds for all high probability j with probability at least $1 - |\mathrm{High}|\delta_1$.

We have that (by Chernoff bounds), with probability at least $1 - \exp\{-\epsilon^2 \ell q_j/3\} \geq 1 - \delta_2$, we have $q_j/\hat{q}_j \leq (1+\epsilon)$. A union bound implies this holds for all high probability j with probability $1 - r\delta_2$.

Furthermore, noting that $q_j n(1+\epsilon) = \hat{q}_j n(1+\epsilon)\frac{q_j}{\hat{q}_j}$, and upper bounding $\frac{q_j}{\hat{q}_j}$ by $1+\epsilon$, we get that

$q_j n(1+\epsilon) \le (1+\epsilon)^2 \hat{q}_j n$, with probability at least $1-\delta_2$. Thus, with probability at least $1 - r\delta_1 - r\delta_2$,

$$\begin{aligned}
\mathrm{cost}(\Pi, n) &\le \mathrm{cost}(\text{Low}) + \mathrm{cost}(\text{High}) \\
&\le \sum_{j \in \text{Low}} \mathrm{cost}_j\left(q^* n(1+\epsilon)\right) + \sum_{j \in \text{High}} \mathrm{cost}_j\left((1+\epsilon)^2 \hat{q}_j n\right) \\
&\le rBq^* n(1+\epsilon) + (1+\epsilon)^2 \sum_{j \in \text{High}} \mathrm{cost}_j\left(\hat{q}_j n\right) \\
&= rBq^* n(1+\epsilon) + (1+\epsilon)^2 \sum_{j \in \text{High}} \mathrm{cost}'_j(\Pi, \ell).
\end{aligned}$$

Note that Lemma 4.5 (with $\epsilon = 1/x$) implies that on n customers,

$$\mathrm{OPT} \ge \min_j \mathrm{cost}_j(n) \ge n^{\log_2(1+\frac{1}{2\beta})} \min_j \mathrm{cost}_j(1) \ge n^{\log_2(1+\frac{1}{2\beta})} = n^c,$$

where the third inequality is by the assumption on the range of $\mathrm{cost}_i(1)$. Thus, $rBq^* n(1+\epsilon) = \epsilon n^c \le \epsilon \mathrm{OPT}$.

We showed that

$$\mathrm{cost}(\Pi, n) \le \epsilon \mathrm{OPT} + (1+\epsilon)^2 \sum_{j \in \text{High}} \mathrm{cost}'_j(\Pi, \ell)$$

$$\le \epsilon \mathrm{cost}(\Pi, n) + (1+\epsilon)^2 \sum_{j \in \text{High}} \mathrm{cost}'_j(\Pi, \ell).$$

Therefore,

$$\begin{aligned}
\mathrm{cost}(\Pi, n) &\le \frac{(1+\epsilon)^2}{1-\epsilon} \sum_{j \in \text{High}} \mathrm{cost}'_j(\Pi, \ell) \\
&\le \frac{(1+\epsilon)^2}{1-\epsilon} \mathrm{cost}'(\Pi, \ell).
\end{aligned}$$

The lower bound is basically similar. For $j \in$ Low, a Chernoff bound implies we have $\hat{q}_j < q^*(1+\epsilon)$ with probability at least $1 - \exp\{-\epsilon^2 q^* \ell/3\} \ge 1 - \delta_2$. So we have

$$\begin{aligned}
\sum_{j \in \text{Low}} \mathrm{cost}_j(\hat{q}_j n) &\le \sum_{j \in \text{Low}} \mathrm{cost}_j(q^*(1+\epsilon)n) \\
&\le rB(1+\epsilon)q^* n \\
&= \epsilon n^c \\
&\le \epsilon \mathrm{OPT} \\
&\le \epsilon \mathrm{cost}(\Pi, n).
\end{aligned}$$

For $j \in$ High, again by a Chernoff bound, we have $\hat{q}_j / q_j \le (1+\epsilon)$ with probability at least $1 - \exp\{-\epsilon^2 q_j \ell/3\} \ge 1 - \delta_2$. Thus, by a union bound, with probability at least $1 - r\delta_2$,

$$\begin{aligned}
\mathrm{cost}'(\Pi, \ell) &= \sum_{j \in \text{Low}} \mathrm{cost}_j(\hat{q}_j n) + \sum_{j \in \text{High}} \mathrm{cost}_j(\hat{q}_j n) \\
&\le \epsilon \mathrm{cost}(\Pi, n) + \sum_{j \in \text{High}} \mathrm{cost}_j(q_j n(1+\epsilon)).
\end{aligned}$$

By another application of Chernoff and union bounds, with probability at least $1-\sum_{j\in\text{High}}\exp\{-\epsilon^2 q_j n/2\} \geq 1-r\delta_1$, for every $j \in$ High, the number of j we will sell when applying Π to n customers is at least $q_i n(1-\epsilon)$. Thus,

$$\sum_{j\in\text{High}} \text{cost}_j(q_j n(1+\epsilon)) = \sum_{j\in\text{High}} \text{cost}_j(q_j n(1-\epsilon)\frac{1+\epsilon}{1-\epsilon}) \leq \frac{1+\epsilon}{1-\epsilon}\sum_{j\in\text{High}} \text{cost}_j(q_j n(1-\epsilon)) \leq \frac{1+\epsilon}{1-\epsilon}\text{cost}(\Pi, n).$$

Altogether, we have proven that with probability at least $1-r(\delta_1+\delta_2)$,

$$\begin{aligned}\text{cost}'(\Pi,\ell) &\leq \left(\epsilon + \frac{1+\epsilon}{1-\epsilon}\right)\text{cost}(\Pi,n)\\ &= \frac{1+2\epsilon-\epsilon^2}{1-\epsilon}\text{cost}(\Pi,n),\end{aligned}$$

which implies

$$\frac{1-\epsilon}{1+2\epsilon-\epsilon^2}\text{cost}'(\Pi,\ell) \leq \text{cost}(\Pi,n).$$

A naive union bound can be done over all the permutations, which will add a factor of $r!$; we can reduce the factor to $r2^r$ by noticing that we are only interested in events of the type π_j, namely a given item (say, j) is in the set of desired items, and another set (say, $\{1,\ldots,j-1\}$) is not in that set. This has only $r2^r$ different events we need to perform the union over. Thus, the above inequalities hold for all permutations with probability at least $1-r^2 2^{r+1}(\delta_1+\delta_2)$. □

C Proofs of Theorems 4.6 and 4.7

Proof of Theorem 4.6. Fix a permutation Π. Let π_j denote the event that a customer buys item Π_j and not covered by items Π_1 through Π_{j-1}. Namely, the probability that the consumer set of desired items include j and none of the items $1,\ldots,j-1$. Let q_j denote $Pr[\pi_j]$, and let $\hat{q}_j$ denote the fraction of Π_j on the initial ℓ-sample.

Item j to is a "Low probability item" if $q_j < \left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}}$; and "High probability items" if $q_j \geq \left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}}$. Let the set "Low" include all "Low probability items"; and the set "High" include all "High probability items".

First we address the case of item j of low probability. The quantity of item j that we will sell is at most $\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}} n(1+\epsilon)$ (Chernoff bound) with probability at least $1-\delta_1$ with $\delta_1 = \exp\{-\epsilon^2\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}} n/3\}$. By a union bound, this holds for all low probability item j, with probability at least $1-|\text{Low}|\delta_1$.

Next, we suppose j has high probability. In this case, the quantity of item j we will sell is at most $q_j n(1+\epsilon)$, with probability at least $1-\exp\{-\epsilon^2 q_j n/3\} \geq 1-\delta_1$. Again, a union bound implies this holds for all high probability j with probability at least $1-|\text{High}|\delta_1$.

We have that (by Chernoff bounds), with probability at least $1-\exp\{-\epsilon^2\ell q_j/3\} \geq 1-\delta_2$, we have $q_j/\hat{q}_j \leq (1+\epsilon)$. A union bound implies this holds for all high probability j with probability $1-r\delta_2$.

Furthermore, noting that $q_j n(1+\epsilon) = \hat{q}_j n(1+\epsilon)\frac{q_j}{\hat{q}_j}$, and upper bounding $\frac{q_j}{\hat{q}_j}$ by $1+\epsilon$, we get that $q_j n(1+\epsilon) \leq (1+\epsilon)^2\hat{q}_j n$, with probability $1-\delta_2$. Thus,

$$\begin{aligned}\text{cost}(\Pi,n) &\leq \text{cost}(\text{Low}) + \text{cost}(\text{High})\\ &\leq r\left(\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}} n(1+\epsilon)\right)^{1-\alpha} + \sum_{j\in\text{High}}\left((1+\epsilon)^2\hat{q}_j n\right)^{1-\alpha}\\ &\leq \epsilon(1+\epsilon)^{1-\alpha}n^{1-\alpha} + (1+\epsilon)^{2(1-\alpha)}n^{1-\alpha}\sum_{j\in\text{High}}(\hat{q}_j)^{1-\alpha}.\end{aligned}$$

Note that the total cost of all low probability items is at most ϵ-fraction of OPT which is at least $\frac{n^{1-\alpha}}{1-\alpha}$. Also,

$$\begin{aligned}(1+\epsilon)^{2(1-\alpha)}n^{1-\alpha}\sum_{j\in\text{High}}(\hat{q}_j)^{1-\alpha} &= (1+\epsilon)^{2(1-\alpha)}\left(\frac{n}{\ell}\right)^{1-\alpha}\sum_j(\hat{q}_j\ell)^{1-\alpha}\\ &= (1+\epsilon)^{2(1-\alpha)}\left(\frac{n}{\ell}\right)^{1-\alpha}\text{cost}(\Pi,\ell)\end{aligned}$$

by definition of $\text{cost}(\Pi,\ell)$.
Therefore we showed,

$$\begin{aligned}\text{cost}(\Pi,n) &\le \epsilon(1+\epsilon)^{1-\alpha}\ell^{1-\alpha}\left(\frac{n}{\ell}\right)^{1-\alpha}+(1+\epsilon)^{2(1-\alpha)}\left(\frac{n}{\ell}\right)^{1-\alpha}\text{cost}(\Pi,\ell)\\ &\le (1+5\epsilon)\left(\frac{n}{\ell}\right)^{1-\alpha}\text{cost}(\Pi,\ell)\end{aligned}$$

The lower bound is basically similar. For $j\in Low$, we have $q_j<\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}}$ and $\hat{q}_j<\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}}(1+\epsilon)$ (by Chernoff bounds). So we have:

$$\begin{aligned}\sum_j(\hat{q}_j\ell)^{1-\alpha} &\le \sum_j\left(\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}}(1+\epsilon)\ell\right)^{1-\alpha}\\ &= r\frac{\epsilon}{r}(1+\epsilon)^{1-\alpha}\ell^{1-\alpha}\\ &= \epsilon(1+\epsilon)^{1-\alpha}n^{1-\alpha}\left(\frac{\ell}{n}\right)^{1-\alpha}\\ &\le \epsilon(1+\epsilon)^{1-\alpha}\text{cost}(\Pi,n)\left(\frac{\ell}{n}\right)^{1-\alpha}\end{aligned}$$

Thus,

$$\begin{aligned}\text{cost}(\Pi,\ell) &= \sum_{j\in\text{Low}}(\hat{q}_j\ell)^{1-\alpha}+\sum_{j\in\text{High}}(\hat{q}_j\ell)^{1-\alpha}\\ &\le \text{cost}(\Pi,n)\epsilon\left(\frac{\ell}{n}\right)^{1-\alpha}(1+\epsilon)^{1-\alpha}+\sum_{j\in\text{High}}(q_jn)^{1-\alpha}\left(\frac{\ell}{n}\right)^{1-\alpha}\left(\frac{\hat{q}_j}{q_j}\right)^{1-\alpha}\\ &\le \text{cost}(\Pi,n)\epsilon\left(\frac{\ell}{n}\right)^{1-\alpha}(1+\epsilon)+\sum_{j\in\text{High}}(q_jn)^{1-\alpha}\left(\frac{\ell}{n}\right)^{1-\alpha}(1+\epsilon)\\ &\le (1+\epsilon)^2\text{cost}(\Pi,n)\left(\frac{\ell}{n}\right)^{1-\alpha}\end{aligned}$$

with probability at least $1-\exp\left\{-q_jn\epsilon^2/2\right\}\ge 1-\delta_3$. For low-probability j, the number of item j sold is $\ge\left(\frac{\epsilon}{r}\right)^{\frac{1}{1-\alpha}}n(1-\epsilon)$ with probability at least $1-\delta_3$. A union bound extends these to all j with combined probability $1-r\delta_3$.

Thus we obtain the upper bound: $\text{cost}(\Pi,n)\le\text{cost}(\Pi,\ell)(1+\epsilon)^{2(1-\alpha)}\left(\frac{n}{\ell}\right)^{1-\alpha}$ and the lower bound: $\text{cost}(\Pi,n)\ge\text{cost}(\Pi,\ell)(1+\epsilon)^{-2}\left(\frac{n}{\ell}\right)^{1-\alpha}$, with probability at least $1-r2^r(\delta_1+\delta_2+\delta_3)$.

A naive union bound can be done over all the permutations, which will add a factor of $r!$, we can reduce the factor to $r2^r$ by noticing that we are only interested in events of the type π_j, namely a given item (say, j) is in the set of desired items, and another set (say, $\{1, \ldots, j-1\}$) is not in that set. This has only $r2^r$ different events we need to perform the union over. □

Proof of Theorem 4.7. Let $\hat{\Pi}$ be the permutation policy produced by GreedyGeneralCost, after the ℓ first customers. By Theorem 4.4,

$$\mathrm{cost}(\hat{\Pi}, \ell) \leq \frac{1}{1-\alpha} \min_{\Pi} \mathrm{cost}(\Pi, \ell).$$

By Theorem 4.6, with probability $1 - \delta^{(\ell)}$,

$$\min_{\Pi} \mathrm{cost}(\Pi, \ell) \leq \min_{\Pi} \mathrm{cost}(\Pi, n)(1+\epsilon)^2 \left(\frac{\ell}{n}\right)^{1-\alpha}.$$

Additionally, on this same event,

$$\mathrm{cost}(\hat{\Pi}, n) \leq \mathrm{cost}(\hat{\Pi}, \ell)(1+\epsilon)^{2(1-\alpha)} \left(\frac{n}{\ell}\right)^{1-\alpha}.$$

Altogether, this implies

$$\begin{aligned} \mathrm{cost}(\hat{\Pi}, n) &\leq \frac{(1+\epsilon)^{2(1-\alpha)}}{1-\alpha} \left(\frac{n}{\ell}\right)^{1-\alpha} \min_{\Pi} \mathrm{cost}(\Pi, n)(1+\epsilon)^2 \left(\frac{\ell}{n}\right)^{1-\alpha} \\ &= \frac{(1+\epsilon)^{4-2\alpha}}{1-\alpha} \min_{\Pi} \mathrm{cost}(\Pi, n). \end{aligned}$$

□

D Proof of Theorem 5.7 and Corollary 5.8

Proof of Theorem 5.7. By Theorem 5.6, with probability at least $1-\delta$, both the price vector price produced by this method and the price vector price^* corresponding to the OPT solution have, for every item i, the probabilities p_i and p_i^* a random customer purchases item i given prices price adn price^*, respectively, satisfy $|p_i - \ell_i/\ell| \leq \epsilon$ and $|p_i^* - \ell_i^*/\ell| \leq \epsilon$, where ℓ_i^* is the number of customers among $1, \ldots, \ell$ that purchase item i given prices price^*. Furthermore, letting n_i and n_i^* denote the number of customers among $\ell+1, \ldots, n$ that would purchase item i given the price vectors price and price^*, respectively, Theorem 5.6 also implies that, with probability at least $1-\delta$, every item i has $|p_i - n_i/(n-\ell)| \leq \epsilon$ and $|p_i^* - n_i^*/(n-\ell)| \leq \epsilon$. Therefore, by a union bound, with probability at least $1-2\delta$,

$$\left|\frac{\ell_i}{\ell} - \frac{n_i}{n-\ell}\right| \leq 2\epsilon \tag{2}$$

and

$$\left|\frac{\ell_i^*}{\ell} - \frac{n_i^*}{n-\ell}\right| \leq 2\epsilon. \tag{3}$$

Next, let $\mathrm{sw}(x, \mathrm{price})$ and $\mathrm{sw}(x, \mathrm{price}^*)$ denote the per-customer social welfares for allocating to a customer with utility vector x using prices price and price^*, respectively. Theorem 5.5 implies that with probability at least $1-\delta$,

$$\left|\frac{1}{\ell}\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price})-\mathbb{E}[\mathrm{sw}(x,\mathrm{price})|\mathrm{price}]\right|\leq\epsilon$$

and

$$\left|\frac{1}{\ell}\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)-\mathbb{E}[\mathrm{sw}(x,\mathrm{price}^*)|\mathrm{price}^*]\right|\leq\epsilon,$$

where x is sampled independently from the $u_{j,\cdot}$ according to the same distribution. Furthermore, Theorem 5.5 also implies that with probability at least $1-\delta$,

$$\left|\frac{1}{n-\ell}\sum_{j=\ell+1}^{n}\mathrm{sw}(u_{j,\cdot},\mathrm{price})-\mathbb{E}[\mathrm{sw}(x,\mathrm{price})|\mathrm{price}]\right|\leq\epsilon$$

and

$$\left|\frac{1}{n-\ell}\sum_{j=\ell+1}^{n}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)-\mathbb{E}[\mathrm{sw}(x,\mathrm{price}^*)|\mathrm{price}^*]\right|\leq\epsilon.$$

Therefore, by a union bound, with probability at least $1-2\delta$,

$$\left|\frac{1}{\ell}\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price})-\frac{1}{n-\ell}\sum_{j=\ell+1}^{n}\mathrm{sw}(u_{j,\cdot},\mathrm{price})\right|\leq 2\epsilon \tag{4}$$

and

$$\left|\frac{1}{\ell}\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)-\frac{1}{n-\ell}\sum_{j=\ell+1}^{n}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)\right|\leq 2\epsilon. \tag{5}$$

Now consider the allocation that allocates item i to the first $n_i^*\frac{\ell}{n-\ell}$ customers in $1,\ldots,\ell$ that would purchase item i given prices price^*, and let SW^* denote the social welfare achieved by this policy on these customers. Since the allocation induced by price^* on the customers $\ell+1,\ldots,n$ satisfies the budget constraint, we know that the above allocation for customers $1,\ldots,\ell$ achieving social welfare SW^* has total cost (under cost') at most $\sum_i \mathrm{cost}'_i\left(n_i^*\frac{\ell}{n-\ell}\right)=\sum_i\mathrm{cost}_i(n_i^*)\leq B$, so that this allocation respects the budget on this ℓ-sized problem. In particular, this implies $\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price})\geq(1-1/e)SW^*$. By (3), we have $\forall i$, $n_i^*\frac{\ell}{n-\ell}\geq\ell_i^*-2\epsilon\ell$, so that $SW^*\geq\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)-2\epsilon\ell Cr$. Furthermore, (5) implies $\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)\geq\frac{\ell}{n-\ell}\sum_{j=\ell+1}^{n}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)-2\epsilon\ell$. Together, we have

$$\sum_{j=1}^{\ell}\mathrm{sw}(u_{j,\cdot},\mathrm{price})\geq(1-1/e)\left(\frac{\ell}{n-\ell}\sum_{j=\ell+1}^{n}\mathrm{sw}(u_{j,\cdot},\mathrm{price}^*)-2\epsilon\ell-2\epsilon\ell Cr\right). \tag{6}$$

Next, consider the allocation used by the algorithm above, which allocates each item i to the first $\ell_i\frac{n-\ell}{\ell}$ customers in $\ell+1,\ldots,n$ that would purchase item i given the prices price, and let SW denote the social welfare achieved by this policy on these customers. Since the allocation induced by price on the customers

$1,\ldots,\ell$ satisfies the budget constraint under cost', we know that this allocation for customers $\ell+1,\ldots,n$ achieving social welfare SW has total cost (under cost) at most $\sum_i \mathrm{cost}_i\left(\ell_i \frac{n-\ell}{\ell}\right) = \sum_i \mathrm{cost}'_i(\ell_i) \leq B$, so that this allocation respects the budget on this $(n-\ell)$-sized problem. Also, by (2), we have $\forall i$, $\ell_i \frac{n-\ell}{\ell} \geq n_i - 2\epsilon(n-\ell)$, so that $SW \geq \sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}) - 2\epsilon(n-\ell)Cr$. Furthermore, (4) implies $\sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}) \geq \frac{n-\ell}{\ell}\sum_{j=1}^{\ell} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}) - 2\epsilon(n-\ell)$. Together, we have

$$SW \geq \frac{n-\ell}{\ell}\sum_{j=1}^{\ell} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}) - 2\epsilon(n-\ell) - 2\epsilon(n-\ell)Cr.$$

Combined with (6), this implies

$$\begin{aligned} SW &\geq (1-1/e)\sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}^*) - 2(2-1/e)(1+Cr)\epsilon(n-\ell) \\ &\geq (1-1/e)\sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}^*) - 2(2-1/e)(1+Cr)\epsilon n. \end{aligned}$$

The result then follows by a union bound (combining the two $1-2\delta$ probability events) and the fact that

$$\sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}^*) \geq \mathrm{OPT} - C\ell \geq \mathrm{OPT} - C\epsilon n.$$

□

Proof of Corollary 5.8. As in the proof of Theorem 5.7, on the $(1-4\delta)$-probability event of Theorem 5.7, we also have that for each i, $n_i \leq \ell_i \frac{n-\ell}{\ell} + 2\epsilon(n-\ell) \leq \ell_i \frac{n-\ell}{\ell} + 2\epsilon n$. Therefore, the number of customers j for which $i = \arg\max_{i'} u_{j,i'} - \mathrm{price}_{i'}$ and $u_{j,i} - \mathrm{price}_i \geq 0$, which appear in the sequence of customers *after* the initial $\ell_i \frac{n-\ell}{\ell}$ such customers is at most $2\epsilon n$.

For each t, let i_t be the t^{th} item i to reach a number of copies of item i allocated equal $\left\lfloor \ell_i \frac{n-\ell}{\ell} \right\rfloor$. Consider the following notion of a "chain". Consider any given customer j for which there exists an item $i = \arg\max_{i'} u_{j,i'} - \mathrm{price}_{i'}$ with $u_{j,i} - \mathrm{price}_i \geq 0$, and yet the customer j is not allocated item i. Suppose customer j is the q^{th} customer j' in the sequence to have $i = \arg\max_{i''} u_{j',i''} - \mathrm{price}_{i''}$ and $u_{j',i} - \mathrm{price}_i \geq 0$. If $q > \ell_i \frac{n-\ell}{\ell}$, then we define the "chain" rooted at customer j to be simply $\{j\}$. On the other hand, if $q \leq \ell_i \frac{n-\ell}{\ell}$, then denoting by q'_i the number of customers j' in the sequence with $i = \arg\max_{i''} u_{j',i''} - \mathrm{price}_{i''}$ and $u_{j',i} - \mathrm{price}_i \geq 0$ who actually *do* receive a copy of item i in the allocation, and denote by $j_{i,q}$ the index of the $(q-q')^{\text{th}}$ customer j' to receive a copy of item i while $i \neq \arg\max_{i''} u_{j',i''} - \mathrm{price}_{i''}$; there must exist such a value $j_{i,q} < j$, since $q \leq \ell_i \frac{n-\ell}{\ell}$, but the number of copies of item i allocated has reached $\left\lfloor \ell_i \frac{n-\ell}{\ell} \right\rfloor$ prior to reaching customer j. Since $j_{i,q} < j$, and customer $j_{i,q}$ did not receive item $i''' = \arg\max_{i''} u_{j_{i,q},i''} - \mathrm{price}_{i''}$ even though $u_{j_{i,q},i'''} - \mathrm{price}_{i'''} \geq 0$, we can inductively suppose there is a chain R rooted at customer $j_{i,q}$. Then we define the chain rooted at customer j to be $\{j\} \cup R$. Furthermore, note that $i = i_t$ for some t (since the number of copies of item i allocated reaches $\left\lfloor \ell_i \frac{n-\ell}{\ell} \right\rfloor$ prior to reaching customer j), and in the latter case above we must have that $i''' = i_{t'}$ for some $t' < t$, since customer $j_{i,q}$ did not receive item i''' (indicating its number of copies had reached $\left\lfloor \ell_{i'''} \frac{n-\ell}{\ell} \right\rfloor$), but *did* receive item i (indicating its number of samples had not yet reached $\left\lfloor \ell_i \frac{n-\ell}{\ell} \right\rfloor$ at that time). Therefore, by induction, each time we augment a chain in this way, we increase the largest value of t'' for which one of the customers j' in the chain has $i_{t''} = \arg\max_{i''} u_{j',i''} - \mathrm{price}_{i''}$ and $u_{j',i_{t''}} - \mathrm{price}_{i_{t''}} \geq 0$; this implies that all chains have length at most r. Thus, we have inductively defined a notion of a "chain" for every customer j that does not receive the item that would be allocated by the pricing policy price.

Note that, since each customer j' is allocated at most one item, we can have $j' = j_{i,q}$ for at most one i and q in the above construction; this implies that, if R is the chain rooted at j', at most one customer $j \notin R$ has $\{j\} \cup R$ as the chain rooted at j. Therefore, any two chains are either disjoint or one is a subset of the other. Furthermore, note that the union of the maximal chains (the chains R such that there are no customers $j \notin R$ with $\{j\} \cup R$ as the chain rooted at j) contains every customer j that does not receive the item that would be allocated to j by the pricing policy price. Also, each chain contains some customer j for which the chain rooted at j is $\{j\}$, which means that for $i = \arg\max_{i'} u_{j,i'} - \mathrm{price}_{i'}$, if j was the q^{th} customer j' to have $i = \arg\max_{i'} u_{j',i'} - \mathrm{price}_{i'}$ and $u_{j',i} - \mathrm{price}_i \geq 0$, where $q > \ell_i \frac{n-\ell}{\ell}$. Since the number of such customers is at most $\sum_{i'} n_{i'} - \left\lfloor \ell_{i'} \frac{n-\ell}{\ell} \right\rfloor \leq r(2\epsilon n + 1)$, and since any two maximal chains are disjoint, the total number of maximal chains is at most $r(2\epsilon n + 1)$. Thus, since any maximal chain has size at most r, the total number of customers j *not* allocated the item $i = \arg\max_{i'} u_{j,i'} - \mathrm{price}_{i'}$ when $u_{j,i} - \mathrm{price}_i \geq 0$ is at most $r^2(2\epsilon n + 1)$.

Thus, the social welfare achieved by this allocation is at least

$$\sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}) - r^2(2\epsilon n + 1)C.$$

Furthermore, as in the proof of Theorem 5.7, we have

$$\begin{aligned}
&\sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}) \geq \frac{n-\ell}{\ell} \sum_{j=1}^{\ell} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}) - 2\epsilon n \\
&\geq (1 - 1/e) \left(\sum_{j=\ell+1}^{n} \mathrm{sw}(u_{j,\cdot}, \mathrm{price}^*) - 2\epsilon n - 2\epsilon n C r \right) - 2\epsilon n \\
&\geq (1 - 1/e) \left(\mathrm{OPT} - C\epsilon n - 2\epsilon n - 2\epsilon n C r \right) - 2\epsilon n,
\end{aligned}$$

and the result follows. □