

Simple and Approximately Optimal Pricing for Proportional Complementarities

Yang Cai* Nikhil R. Devanur† Kira Goldner‡ R. Preston McAfee§

May 3, 2019

Abstract

We study a new model of complementary valuations, which we call “proportional complementarities.” In contrast to common models, such as hypergraphic valuations, in our model, we do not assume that the extra value derived from owning a set of items is independent of the buyer’s base valuations for the items. Instead, we model the complementarities as proportional to the buyer’s base valuations, and these proportionalities are known market parameters.

Our goal is to design a simple pricing scheme that, for a single buyer with proportional complementarities, yields approximately optimal revenue. We define a new class of mechanisms where some number of items are given away for free, and the remaining items are sold separately at inflated prices. We find that the better of such a mechanism and selling the grand bundle earns a 12-approximation to the optimal revenue for pairwise proportional complementarities. This confirms the intuition that items should not be sold completely separately in the presence of complementarities. In the more general case, a buyer has a maximum of proportional positive hypergraphic valuations, where a hyperedge in a given hypergraph describes the boost to the buyer’s value for item i given by owning any set of items T in addition. The maximum-out-degree of such a hypergraph is d , and k is the positive rank of the hypergraph. For valuations given by these parameters, our simple pricing scheme is an $O(\min\{d, k\})$ -approximation.

1 Introduction

Consider a setting where multiple items are being sold, and a buyer’s valuations for the items have *complementarities*. That is, the buyer derives some value from owning a combination of items that is not present when owning any of the items individually, as in the following examples.

Microsoft Office Example: A person who values producing documents will value software such as Microsoft Word that helps him in this task. If the person wants to include some charts in his document, then this is made easier and faster by having another piece of software that specializes in making charts such as Microsoft Excel. Thus owning Excel in addition to Word boosts the value of Word for him, since he can then produce more documents in the same amount of time.

*Yale University (yang.cai@yale.edu). The work of Y. Cai was supported NSERC Discovery RGPIN-2015-06127 and FRQNT 2017-NC-198956.

†Microsoft Research (nikdev@microsoft.com)

‡University of Washington (kgoldner@cs.washington.edu). The work of K. Goldner was supported by a Microsoft Research PhD Fellowship as well as NSF grants CCF-1420381 and CCF-1813135.

§Microsoft (preston@mcafee.cc)

Cloud Services Example: A cloud service provider offers multiple heterogeneous items that are both substitutes and complements. You can purchase a general purpose virtual machine (VM) or a special purpose VM such as a “data science VM”; these are substitutes. You can also purchase an upgrade such as a fast solid state disk-drive (SSD) which would be complementary to either of those VMs.

The goal of this work is to understand how a revenue-maximizing seller should price items such as Microsoft Office products or cloud services when facing a buyer with such complementarities. To this end, we introduce a new model of complementarities, design a pricing scheme for this model, and show worst-case approximation guarantees.

In recent years, there has been a surge of research activity on *optimal combinatorial pricing*. This is the problem of determining and pricing bundles of heterogeneous items in order to maximize revenue from selling to a buyer who has a combinatorial valuation function. The theme of the research has been *simple vs. optimal*, where simple pricing schemes are shown to approximate the optimal (possibly randomized) pricing scheme to within a universal constant multiplicative factor, independent of the number of items. E.g., for additive valuations, where the buyer’s valuation for any set of items is just the sum of her valuations for each individual item, Babaioff, Immorlica, Lucier, and Weinberg [2014] show that the revenue from either selling each item separately (SREV), or selling the grand bundle of all the items (BREV) is a 6-approximation. Similar results have been extended to much broader settings, such as one buyer with unit-demand [Chawla, Hartline, and Kleinberg, 2007] and subadditive [Rubinstein and Weinberg, 2015] valuations, and multiple buyers with additive [Yao, 2015], unit-demand [Chawla, Hartline, Malec, and Sivan, 2010], gross-substitutes [Chawla and Miller, 2016], and XOS valuations [Cai and Zhao, 2017].

All of the above valuation classes are complement-free. In contrast, in practice, bundling is most attractive when the items are complementary to each other. The first such result for complementary valuations is by Eden, Feldman, Friedler, Talgam-Cohen, and Weinberg [2017b], who consider the following *positive hypergraphic (PH) valuation* model: each item is a vertex in a given hypergraph. For any *hyperedge* given by a subset of items S , the buyer gets an additional value of $v_S > 0$ if he gets all of the items in S . The value v_S for each hyperedge S is drawn independently from a known prior distribution. Eden et al. show that in this case the approximation ratio of $\max\{\text{SREV}, \text{BREV}\}$ is $\Theta(d)$, where d , the maximum-degree of the hypergraph, is the maximum number of hyperedges that any one item is part of. Further, they show that other natural parameters that have been considered for complementary valuations have very bad lower bounds. The approximation ratio could be exponential in the number of items m , as well as in the *positive rank* k , which is the maximum size of (number of items in) a hyperedge.¹

1.1 Proportional Complementarities Model

The PH valuation model for the Microsoft Office example would have three values, one each for Word, Excel, and the pair (Word, Excel), with each of them drawn independently from a different distribution. While having the values for Word and Excel be independent may be reasonable, that the value for the pair (Word, Excel) be independent of the other two seems unrealistic. Similarly,

¹ This is the choice of parameter used for welfare guarantees for combinatorial auctions, e.g., in Abraham, Babaioff, Dughmi, and Roughgarden [2012]; Feige, Feldman, Immorlica, Izsak, Lucier, and Syrgkanis [2015]; Feldman, Gravin, and Lucier [2015].

for the cloud services example, the PH model would have that the value for the pair (VM, SSD) be independent of the value for the VM alone, which is once again unrealistic.

We introduce what we call a *proportional complementarities* model of valuations; a special case of this model is *proportional pairwise complementarities* (PPC). We illustrate this model through the examples we considered before.

Microsoft Office Example in the ppc model: We still have a value for each of Word and Excel, say v_1 and v_2 respectively, that are independent of each other. Our model differs in how the buyer values the combination of the two by assuming that the additional value derived from having both items is due to a better utilization of either item, and hence is proportional to (rather than independent of) the buyer’s base valuation for Word and for Excel. This is captured in our model by having a multiplier for the pair (Word, Excel), denoted by $\eta_{1,2}$; say Excel always adds 23% to the value of Word, then we would have $\eta_{1,2} = 0.23$. One could get an estimate of this quantity by observing the frequency of activities between the two, such as dragging Excel charts into Word. While not fully general, these proportionalities make intuitive sense, because if a buyer values an item highly, he is likely to care more about its complements too, as they enhance that item. The value for purchasing both items in our model would then be

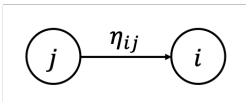
$$v_1(1 + \eta_{1,2}) + v_2.$$

The other assumption we make is that while the seller doesn’t know the exact values, he knows these proportions of complementarities. This is perhaps the least accurate assumption in applications, because such values could reasonably vary across individuals. However, in circumstances where the way products are used together is approximately fixed, such as dragging Excel charts into Word, it is not unreasonable to assume that these values are known. This is especially true when it comes to “digital goods,” where data about interactions between items can be gathered, and the parameters can be estimated from this data, e.g. via estimating cross-price elasticities.

Allowing the proportions (i.e., the η s) to vary across individuals is an interesting direction for future research. We present one possible approach via a common generalization of our model and the PH model in Section 5. This generalization further illustrate the similarities and the differences between the two models.

Proportional pairwise complementarities: We first define the PPC model. A single seller offers m heterogeneous items for sale to a single buyer. (Equivalently, there is a population of buyers, but no supply constraints on the seller, as is the case with digital goods like Microsoft Office products.) We model the structure of the complementarities among the items via the following parameters, which are assumed to be known to the seller:²

$$\eta_{ij} \in \mathbb{R}_+ \quad \forall i, j \in [m], i \neq j.$$



The parameter η_{ij} captures how much having item j boosts the valuation that the buyer derives from item i . The valuation of a buyer is determined by his type t , which is a vector in \mathbb{R}_+^m , and is the private information of the buyer. The i^{th} coordinate of t is t_i , which represents his base valuation

²We use the notation $[m]$ to indicate the set of first m natural numbers, $\{1, 2, \dots, m\}$.

for item i in the absence of any other items. If the buyer also gets item j , then his valuation for item i is boosted by an additional $\eta_{ij}t_i$. From this, we get that for any bundle $S \subseteq [m]$, the buyer's valuation for S is

$$v(t, S) := \sum_{i \in S} \eta_i(S)t_i, \quad \text{where} \quad \eta_i(S) = 1 + \sum_{j \in S \setminus \{i\}} \eta_{ij}.$$

Note that η_{ij} need not be equal to η_{ji} , and asymmetric boosts are only more general. We make the Bayesian assumption that t is drawn from a product distribution $\prod_{i \in [m]} F_i$. The distributions F_i for all $i \in [m]$ (as well as the parameters η_{ij}) are known to the seller.

This more general asymmetric case corresponds to directed graphs (and hypergraphs). Thus we define the *directed-positive-rank* k of the graph to be the maximum size of (number of items in) the *source* of a (hyper)edge. Thus, for the pairwise case, $k = 1$.

The general case: The general class of valuations we consider is defined formally in Section 2; we give an informal description here. First of all, we allow hyperedges, instead of edges, i.e., each pair of item i and a disjoint set of items T forms a directed hyperedge (T, i) and has a certain boost associated with it, denoted by η_{iT} : this is the boost of having all items in T on item i . The valuation of a set S now includes all possible boosts due to hyperedges (T, i) for $T \sqcup \{i\} \subseteq S$. We call this class of valuations *proportional positive hypergraphic (PPH) valuations*. The other generalization is to allow the boost to be the maximum of the boost from multiple hypergraphs. We call this class of valuations *maximum of proportional positive hypergraphic (MPPH) valuations*. We denote by k the directed-positive-rank and by d the maximum-degree of the hypergraph. We tie this back to the cloud services example to show how such a generalization is useful.

Cloud Services Example: Suppose that we had access to two types of VMs, VM1 and VM2, that are meant for different types of workloads. We can also purchase additional disk drives (DDs) that allow us to run larger workloads. DDs come in two technologies, fast and slow, say DD1 and DD2. Having either of the DDs can boost the value for a VM, and having both of them boosts it even more but less than the sum of the individual boosts. This could be modeled as follows. There are 4 items, 1 and 2 are the VMs, and 3 and 4 are the DDs. For each of $i \in \{1, 2\}$ and $j \in \{3, 4\}$, we have the boosts η_{ij} as well as $\eta_{i\{3,4\}}$. Let x_3 and x_4 be binary variables indicating whether items 3 and 4 were respectively purchased or not. The value derived from item i for $i \in \{1, 2\}$ depending on these choices is

$$t_i \cdot (1 + \max\{\eta_{i3}x_3, \eta_{i4}x_4, \eta_{i\{3,4\}}x_3x_4\}).$$

Thus VM1 can get a boost of η_{13} from having DD1, or η_{14} from DD2, but if you have both DD1 and DD2, the boost is $\eta_{1\{3,4\}}$ rather than $\eta_{13} + \eta_{14}$.

1.2 Pricing scheme

Almost all of the papers in this line of research consider the better of selling each item separately and selling only the grand bundle. Pricing the grand bundle is (conceptually) easy: set the monopoly price for the distribution of the buyer's value for the grand bundle, which can be computed from the given input. For simple valuations such as additive valuations, setting item prices to sell separately is also easy: set the monopoly reserve for each of them separately. In our model, this completely ignores the boost in the valuation on an item from having other items. Not surprisingly, this can

be provably far from optimum when you have complementarities; we therefore need a non-trivial way to price the items in this case. We first illustrate our algorithm for finding these prices via a numerical example.

Numerical Example: Suppose, as shown on the left in Figure 1, that there are 4 items, numbered 1 through 4, and that we have non-zero η s on the pairs (2, 1), (3, 2), (4, 3) and (1, 4). Let all η s be 1. Suppose t_1 and t_3 are distributed identically as follows: the value is 2 w.p. $\frac{1}{2}$ and 0 otherwise; let t_2 and t_4 be distributed identically as follows: the value is 4 w.p. $\frac{1}{2}$ and 0 otherwise. Each t_i is independent of the others.

We denote the monopoly price and the monopoly revenue for item i alone by r_i and R_i respectively. For this example, we have the monopoly prices as $r_1 = r_3 = 2$ and $r_2 = r_4 = 4$; the revenues are $R_1 = R_3 = 1$ and $R_2 = R_4 = 2$. Setting the monopoly prices for each item separately guarantees a revenue of $\sum_i R_i = 6$. The actual revenue would be higher, but in general it is difficult to get a better handle on it than this bound.




Figure 1: Left: The 4-item example described above. Right: The directed graph where the weight of a directed cut corresponds to a lower bound on the revenue of the corresponding SEPARATE/FREE mechanism.

Step 1: Construct a weighted directed graph. We construct a weighted directed graph with 5 vertices, one for each item, and a source node s . The weight on the edge (s, i) is R_i . The weight on the edge (i, j) is $\eta_{ji}R_j$. This graph is shown on the right in Figure 1.

Step 2: Find a max directed cut. We then find a cut in the graph that maximizes the number of directed edges going from the “source” side to the “sink” side. From the figure, it is easy to see that such a cut is given by the vertices $s, 1$ and 3 on the source side, the rest on the sink side, and has weight 8.

Step 3: Set Prices. We set the prices for items on each side of the cut differently.

1. The items on the source side have a price of 0. This set of items, denoted by \mathcal{F} , are “free”. In this case, items 1 and 3 are free.
2. For the items on the sink side, we multiply the monopoly price r_i by $\eta_i(\mathcal{F})$ (1 plus the boost i gets from all the items on the source side). Then items 2 and 4 thus have a price of 8 each.

The weight of the cut, 8, is a lower bound on the revenue of this pricing scheme. Each of items 2 and 4 is bought at the price of 8 with probability $\frac{1}{2}$, giving a lower bound on revenue of 8. In

comparison, the best price for grand bundling is 12, which is bought with probability $\frac{5}{8}$, giving a revenue of $\frac{15}{2}$, which is slightly lower. Both of these are still higher than the revenue lower bound of 6 from setting separate prices of r_i each.

In general, we introduce a class of mechanisms which we call SEPARATE/FREE. Like selling separately, every item is sold separately at some price, and the buyer may take any set of items so long as he pays the sum of their individual prices. However, we partition the items into “free items” \mathcal{F} , where for each item $i \in \mathcal{F}$, the individual price of each such item is \$0, and “priced items” $\bar{\mathcal{F}} = [m] \setminus \mathcal{F}$. Once the free set \mathcal{F} is determined, we use the knowledge that the buyer will take the free items to inflate the monopoly prices of the priced items by the boost on the item from also getting the free set (and only the free set).

Such mechanisms do capture a certain economic intuition that is seen in practice: giving some items away for free in order to charge more for complementary items, e.g., Google sells the Android OS for free since it is complementary to advertising revenue. One can also think of it as a certain form of bundling: there is no reason to give away the free items unless the buyer purchases some priced item. This is equivalent to bundling all the free items with any non-empty subset of paid items. Going back to our cloud services example, such a pricing scheme could determine that one of the two DDs should be free. We would then bundle that DD into the VMs; such bundles are commonly observed in practice.

One difficulty in the above scheme is that in general, finding a max directed cut in a graph is an *NP-Hard* problem. When restricted to polynomial time algorithms, the best worst-case approximation guarantee we can show is by placing each item independently into the free set with some probability α , which is determined by $\min\{d, k\}$. This is a little unsatisfactory since it doesn’t use the specific market parameters η at all. (However, they are used in setting the prices once \mathcal{F} is determined.) An alternative is to use an approximation algorithm for the max directed cut problem, such as the Goemans-Williamson algorithm. The advantage of this method is that it produces a free set that makes use of the structure of the η s; unfortunately, this does not improve the worst case approximation ratio. In fact, we show that no algorithm can improve the approximation ratio when used in conjunction with our current proof technique, but we conjecture that such an algorithm would be better in practice.

1.3 Worst case approximation guarantee

Once again, we begin by illustrating our analysis using the numerical example earlier. For the sake of analysis, we consider an instance of the pricing problem on the same set of items, with additive valuations. The value distribution for item i in this instance, denoted by \hat{F}_i , is just the original distribution F_i multiplied by $\eta_i([m])$, the boost i can obtain from all of the items. In our example, \hat{F}_i for $i = 1$ and 3 is 4 w.p. $\frac{1}{2}$ and 0 otherwise, and for $i = 2$ and 4 is 8 w.p. $\frac{1}{2}$ and 0 otherwise.

We relate the revenue from selling separately and selling the grand bundle on the given instance to the corresponding mechanisms for the additive instance. It is easy to see that the bundle revenue remains the same in both instances, as the complements buyer receives the boost $\eta_i([m])$ on every item:

$$\text{BREV-PPC}(F) = \text{BREV-ADDITIVE}(\hat{F}).$$

As we computed earlier, a lower bound on selling separately with our pricing scheme for the given instance is 8. Selling separately for the additive instance gives a revenue of 12, which is $3/2$ times

8. So for this example, we have that

$$\text{SREV-PPC}(F) \geq \frac{2}{3} \text{SREV-ADDITIVE}(\hat{F}).$$

More generally, $\text{SREV-ADDITIVE}(\hat{F})$ is equal to the total the weight of all the edges in the digraph that we construct. If you place each item on either side of the cut with equal probability, then each edge is cut with probability $\frac{1}{4}$, which results in a factor of 4 between the two SREVs. This is indeed tight: consider a complete unweighted digraph; any cut can only cut a $\frac{1}{4}$ th fraction of edges.

We can now use a slight generalization of the result of Babaioff et al. [2014] to bound the optimum revenue for the additive instance, denoted by $\text{OPT-ADDITIVE}(\hat{F})$, in terms of SREV and BREV.

$$\text{OPT-ADDITIVE}(\hat{F}) \leq 2 \text{SREV-ADDITIVE}(\hat{F}) + 4 \text{BREV-ADDITIVE}(\hat{F}).$$

Finally, we show that the optimum revenue for the additive instance is only higher.

$$\text{OPT-PPC}(F) \leq \text{OPT-ADDITIVE}(\hat{F}),$$

which gives an approximation ratio of 7 for this example, and 12 in general, working through the inequalities above. Note that even for additive valuations, 5.2 is the best known approximation ratio.

This last step may seem obvious, but it turns out to be quite tricky. One might expect a direct argument, that given a mechanism M for the original instance, we construct a mechanism M' for the additive instance, with a larger revenue. Such approaches are inherently difficult, as evidenced by “Revenue non-monotonicity” in Hart and Reny [2012]. We instead argue the upper bound by covering the dual of the smaller setting with the dual of the larger setting, a novel use of the [CDW16] Lagrangian duality framework.

We show the following approximation guarantee more generally.

Theorem 1 (Informal). *The better of BREV and the revenue from a mechanism of type SEPARATE/FREE is an $O(\min\{d, k\})$ -factor approximation to the optimal revenue for valuations in the class MPPH. When $k = 1$, i.e., the boosts are the maximum over directed graphs, the approximation factor is at most 12.*

Recall that d is the maximum-degree of the hypergraph, and k is the directed-positive-rank of the hypergraph.

We also show that our analysis of Theorem 1 is tight up to a constant factor via the following lower bound. A crucial step in our analysis is to upper bound the optimal revenue for MPPH valuations by the optimal revenue for an instance of additive valuations. Further, the actual revenue of a mechanism from a buyer with proportional complements is extremely difficult to analyze. Instead, we analyze a lower bound on the revenue we deem the “proxy revenue,” and we show that with respect to our upper bound, no mechanism of the following type can give an $o(k)$ -approximation to the proxy revenue. The mechanisms we consider first partition the set of items into bundles, designating one bundle as the free set. Each of the other bundles is priced separately. The buyer always gets the free set for free. Specifically, the price for a bundle is its monopoly reserve price inflated by the boosts of only the other items in its own bundle and by the free set, and not by anything else. The proxy revenue undercounts the revenue in the same way, by assuming that the buyer’s boosted values match the way prices are set in these mechanisms: only within bundles and from the free set. We elaborate on motivation for using this proxy in Section 3.

2 Preliminaries

We now give the formal description of the MPPH valuation model. There is a single seller offering m heterogeneous items for sale to a single buyer. The following parameters determine the structure of complementarities among items via boosts to base valuations. There is a hypergraph with the set of items $[m]$ as vertices whose edges (T, i) correspond to a combination of items T and a disjoint item i to which the combination gives a boost. Moreover, there could be several possible boosts out of which only the highest is activated. For each item $i \in [m]$, for each hyperedge (T, i) , and for each $\ell \in [K]$ for some integer K , we have the parameter $\eta_{iT}^\ell \in \mathbb{R}_+$.




Figure 2: A directed graph representation of the η parameters.

The buyer's valuation for any bundle $S \subseteq [m]$ is

$$v(t, S) = \sum_{i \in S} \eta_i(S) t_i, \quad \text{where} \quad \eta_i(S) = 1 + \max_{\ell \in [K]} \sum_{T \subseteq S \setminus \{i\}} \eta_{iT}^\ell.$$

We refer to the case where the boosts are simply the sum (i.e. $K = 1$) as additive boosts, and the general case ($K > 1$) as XOS boosts. Note that $\eta_i(S)$ always includes the base valuation for item i (the +1) so it is not entirely comprised of boosts, but we overload and call this term the boost anyway. Observe that the boosts are always monotone in the set, that is, if $\ell(S) \in \operatorname{argmax}_{\ell \in [K]} \sum_{T \subseteq S \setminus \{i\}} \eta_{iT}^\ell$, then it is always the case that for all $S \subseteq S'$,

$$\eta_i(S) = 1 + \sum_{T \subseteq S \setminus \{i\}} \eta_{iT}^{\ell(S)} \leq 1 + \sum_{T \subseteq S' \setminus \{i\}} \eta_{iT}^{\ell(S)} \leq 1 + \max_{\ell \in [K]} \sum_{T \subseteq S' \setminus \{i\}} \eta_{iT}^\ell = \eta_i(S') \quad (1)$$

We assume that t is drawn from a product distribution $F = \prod_{i=1}^m F_i$. The distributions F_i for all $i \in [m]$ and the η s are all known to the seller. However, the type realization t is private information of the buyer.

Our approximation ratios depend on the parameters k and d of the underlying hypergraph. The parameter k , the directed-positive-rank, is an upper bound on the size of the set in any hyperedge, i.e., $|T| \leq k$ for each hyperedge (T, i) . The parameter d , the maximum-out-degree, is an upper bound on the number of hyperedges that contain a particular vertex, i.e., for each $i \in [m]$, $|\{\text{hyperedge } (T, j) : i \in T\}| \leq d$. We suppress the dependence on the hypergraph in our notation, since it should always be clear from the context. For the special case of pairwise complementarities (PPC) we follow the notation in Section 1.1.

2.1 Optimal Mechanisms in Various Settings

From the revelation principle, we can restrict our attention to direct revelation mechanisms, where the buyer reports his type. A mechanism is therefore defined by the allocation and the payment

functions. We allow randomized allocation rules, with the assumption that the buyer is risk neutral. Let $x_S(t)$ denote the probability that the bundle $S \subseteq [m]$ is allocated to the buyer of type t ; let $p(t)$ be his payment. The incentive-compatibility (IC) constraints require that for each buyer type, the buyer maximizes utility by reporting his true type.³ Among all IC mechanisms, the optimal mechanism maximizes the expected revenue

$$\mathbf{E}_t[p(t)].$$

Notation: We use the following convention to denote the revenue from a particular mechanism for a given class of valuations, for a particular distribution over types:

$$[\text{Mechanism name}]\text{-}[\text{Valuation Class}](\text{[Distribution]}).$$

For example, the optimal mechanism for PPC valuations with types drawn from F is denoted by $\text{OPT-PPC}(F)$. We drop the distribution when it is clear from the context. We also drop the valuation class when it is additive (ADDITIVE) and it is clear from the context: e.g., the revenue from selling the grand bundle for additive valuations on types drawn from the distribution F is just BREV.

3 A Constant-Factor Approximation via a Random Free Set

We begin with the case of pairwise complementarities and show a 12-approximation for this setting.

Recall that the two standard mechanisms considered in previous work are selling the grand bundle and selling each item separately. Selling the grand bundle only gets better with complements, since we are certain that the buyer will receive all possible boosts, and we can price accordingly. It is selling the items separately that is problematic. A conservative way to set the prices while selling separately is to ignore the complementarities, and sell them as if they are just additive; this could clearly be quite suboptimal. We can price an item more aggressively in order to capture some of the boost from complementarities, but this will decrease its probability of sale, which can further decrease the probabilities of sale for other items that receive a boost from this item. The pricing must get the right tradeoff between capturing more of the boost from complementarity while making sure that sufficient quantity of items are sold in the first place in order for the boosts to accrue. Overall, it is difficult to characterize the behavior of the buyer, which makes optimizing the prices extremely challenging.

Our approach is to shift the focus away from optimizing prices. We do this by giving some items away for free, and then just selling the remaining items individually as if they are additive, but accounting the boost from the items that are given for free. The free items make sure that sufficient boosts accrue; the priced items extract the value thus generated. The problem now becomes one of choosing the set of free items, but in fact we show that a random choice suffices. The analysis compares the revenue to a seemingly crude upper bound, where every item receives the fullest boost that an item could possibly receive—the boost on the item if the buyer were to receive all of the items, that is, the grand bundle.

³We do not formally define IC constraints since we can bypass it due to Lemma 3, but our mechanisms will be clearly IC.

We now formally describe our mechanism SEPARATE/FREE. For each item $i \in [m]$, let r_i^* be the monopoly reserve for the distribution F_i , i.e.,

$$r_i^* = \arg \max_{p \in \mathbb{R}_+} p \cdot (1 - F_i(p)),$$

and let R_i be the revenue of the monopoly reserve for the distribution F_i ,

$$R_i := r_i^* \cdot (1 - F_i(r_i^*)).$$

Mechanism SEPARATE/FREE(\mathcal{F}) : Partition the items into “free items” \mathcal{F} and “priced items” $\bar{\mathcal{F}} = [m] \setminus \mathcal{F}$. The price of a priced item $i \in \bar{\mathcal{F}}$ is

$$p_i = \eta_i(\mathcal{F}) \cdot r_i^*.$$

The buyer gets all of the items in \mathcal{F} for free, that is, they are priced each at 0. We denote by SEPARATE/FREE(\mathcal{F}) the expected revenue from the mechanism with (potentially random) free set \mathcal{F} , and we overload notation slightly to use SEPARATE/FREE = $\max_{\mathcal{F} \subseteq [m]}$ SEPARATE/FREE(\mathcal{F}).

Theorem 2. *The better of selling the grand bundle and Mechanism SEPARATE/FREE is a 12-approximation for PPC valuations:*

$$\text{OPT-PPC} \leq 12 \max\{\text{BREV-PPC}, \text{SEPARATE/FREE-PPC}\}.$$

3.1 Proof of Theorem 2

The proof of this theorem is largely along the lines of the analysis described in Section 1.3. We first relate OPT-PPC to the optimal revenue for an instance of additive valuations; where the buyer’s valuation for each item is inflated as if he receives the boosts from owning every possible item in addition to this one, even if he receives no additional items. Then, the buyer’s new (much larger) valuations are additive. We refer to this setting as the fully-boosted additive setting, where we call the values t multiplied by the full boosts as drawn from the distribution \hat{F} , even though t is drawn identically as from F . We show that the revenue from this setting is only larger than from the proportional complements setting.

Lemma 1.

$$\text{OPT-PPC}(F) \leq \text{OPT-ADDITIVE}(\hat{F}).$$

This is a very loose upper bound and intuitively it should be true: for every type t , the buyer’s value for every set in the fully-boosted additive setting is only larger than in the proportional complements setting. However, due to revenue non-monotonicities, the proof requires more care, and is deferred to Subsection 3.2.

In Appendix A, we prove improve the analysis of the 6-approximation by Babaioff et al. [2014] to allow a parameterization in the bound⁴. Then our Theorem 6 with $a = 1$ gives that

$$\text{OPT-ADDITIVE}(\hat{F}) \leq 2 \text{SREV}(\hat{F}) + 4 \text{BREV}(\hat{F}).$$

⁴This analysis also improves the 6-approximation to 5.382. The state of the art coefficient is 5.2 Ma and Simchi-Levi [2015], but our proof uses the [CDW16] framework and is more modular.

It is easy to see that the revenue from grand bundling in the complements setting on the original distribution is the same as the grand bundling in the fully-boosted additive setting, i.e., $\text{BREV-PPC}(F) = \text{BREV-ADDITIVE}(\hat{F})$, as the buyer receives the full boosts in both cases. It now remains to show that Mechanism SEPARATE/FREE on F is a 4-approximation to $\text{SREV}(\hat{F})$, despite the fact that the prices in the fully-boosted additive setting are each inflated by *full* boost of getting the grand bundle.

Lemma 2.

$$\text{SREV}(\hat{F}) \leq 4 \text{SEPARATE/FREE-PPC}(F).$$

Proof. First, we derive a lower bound on the revenue from Mechanism SEPARATE/FREE for any partition of the items into free and priced. What revenue do we yield for the partition $(\mathcal{F}, \bar{\mathcal{F}})$? Recall that for every item $i \in \bar{\mathcal{F}}$, the price posted is $\eta_i(\mathcal{F}) \cdot r_i^*$. The probability that the buyer purchases item i is at least $\Pr[t_i \geq r_i^*] = 1 - F_i(r_i^*)$, because the buyer receives the boost $\eta_i(\mathcal{F})$ from all the free items with certainty. If the buyer also purchases other items, it will only increase the buyer's value for buying item i , so the probability of purchasing item i can only increase. Hence, the revenue of mechanism SEPARATE/FREE under this particular partition $(\mathcal{F}, \bar{\mathcal{F}})$ is at least

$$\sum_{i \in \bar{\mathcal{F}}} \eta_i(\mathcal{F}) \cdot r_i^* \cdot (1 - F_i(r_i^*)) = \sum_{i \in \bar{\mathcal{F}}} \eta_i(\mathcal{F}) R_i.$$

Now we construct a graph and show that the revenue of Mechanism SEPARATE/FREE under any partition $(\mathcal{F}, \bar{\mathcal{F}})$ of the items is at least the weight of a corresponding directed cut in the following graph. Consider the graph with vertices $[m]$ corresponding to the m items, where directed edge (j, i) has weight $w_{j,i} := \eta_{ij} \cdot R_i$, where R_i is the optimal revenue for selling only item i . The graph also contains a source node s , where for all items $i \in [m]$, the edge (s, i) has weight $w_{s,i} = R_i$. (This will account for the coefficient 1 for the base valuation of the item.) The weight of the directed cut from $\mathcal{F} + \{s\}$ to $\bar{\mathcal{F}}$ is precisely:

$$\sum_{i \notin \mathcal{F}} \sum_{j \in \mathcal{F} + \{s\}} w_{j,i} = \sum_{i \in \bar{\mathcal{F}}} \left(1 + \sum_{j \in \mathcal{F}} \eta_{ij} \right) R_i = \sum_{i \in \bar{\mathcal{F}}} \eta_i(\mathcal{F}) R_i.$$

Hence, for any partition of free and priced items $(\mathcal{F}, \bar{\mathcal{F}})$, the weight of the directed cut from $\mathcal{F} + \{s\}$ to $\bar{\mathcal{F}}$ gives a lower bound on the revenue yielded by Mechanism SEPARATE/FREE for this partition.

We construct our free set by placing each item independently and uniformly at random into \mathcal{F} or $\bar{\mathcal{F}}$. The expected weight of the corresponding random cut from $\mathcal{F} + \{s\}$ to $\bar{\mathcal{F}}$ is at least $\frac{1}{4} \sum_{i \in [m]} \eta_i([m]) \cdot R_i = \frac{1}{4} \text{SREV}_{\eta_{[m]}\text{ot}}$. To see this, observe that for every pair of items (j, i) , the cut gets the weight of $\eta_{ij} R_i$ from this edge whenever $j \in \mathcal{F}$ and $i \notin \mathcal{F}$, which occurs with probability $\frac{1}{4}$. The cut also gets a weight of R_i whenever $i \in \bar{\mathcal{F}}$, which happens with probability $\frac{1}{2}$. \square

Theorem 2 now follows from Lemmas 1 and 2, and Theorem 6 with $a = 1$:

$$\begin{aligned} \text{OPT-PPC}(F) &\leq \text{OPT-ADDITIVE}(\hat{F}) \\ &\leq 2 \text{SREV}(\hat{F}) + 4 \text{BREV}(\hat{F}) \\ &\leq 8 \text{SEPARATE/FREE-PPC}(F) + 4 \text{BREV-PPC}(F). \end{aligned}$$

3.2 Proof of the Benchmark

We now prove Lemma 1: that the optimal revenue from the proportional complements setting is bounded by the optimal revenue from the fully-boosted additive setting. Again, while this is intuitive, revenue non-monotonicities make it unclear how to execute a direct proof. Instead, we use the machinery from the Lagrangian duality framework of Cai et al. [2016] to give a “dual-covering” argument. While the argument is simple and easy-to-see for those familiar with the machinery, the machinery itself is not easy.

First, we formulate the (primal) optimization problem: maximize revenue subject to incentive-compatibility, individual rationality, and feasibility. We have Lagrangian dual variables, denoted by λ , corresponding to each IC constraint, i.e., corresponding to each pair of types (t, t') . Then the Lagrangian duality framework states that, via strong duality, optimal revenue is equal to the optimal dual minimization problem, and upper bounded by any feasible dual.

Of the vast array of works that use the Lagrangian duality framework to achieve an upper bound for approximation [Cai et al., 2016; Cai and Zhao, 2017; Brustle, Cai, Wu, and Zhao, 2017; Eden et al., 2017b; Eden, Feldman, Friedler, Talgam-Cohen, and Weinberg, 2017a; Fu, Liaw, Lu, and Tang, 2017; Liu and Psomas, 2017], the standard approach used by almost all of them is to select dual variables for the setting at hand that naturally split the upper bound into terms that can be bounded by a few simple mechanisms. Then, the bulk of the work remains in bounding the unique terms with the correct mechanisms. Here, however, it is not even clear how to choose a set of dual variables that induces a good upper bound due to the complementarities across items. We take a different path. We first create a new proxy additive setting, where buyers’ valuations are fully-boosted. We then argue that the optimal revenue in our setting is upper bounded by the optimal revenue in the boosted additive setting. As the buyers’ valuations in the boosted setting “dominate” the original buyers’ valuations, the claim is intuitively true. However, due to revenue non-monotonicities, this intuition does not directly translate to a proof. We rely on duality to prove the claim. We show that the optimal dual in the original setting is at most the optimal dual in the fully-boosted additive setting, which by strong duality, is equal to the optimal revenue. This step is the only place we use duality and the rest of the analysis all happens in the primal/mechanism space.

We use $\phi_i(t) := t_i - \frac{1}{f(t)} \sum_{t'} (t'_i - t_i) \lambda(t', t)$ as the “virtual value function” given by λ . Let $f(t)$ denote the probability that the type t is realized. (We assume discrete distributions for simplicity of notation.) We denote the set of feasible allocations by \mathcal{P} —this is just the set that allocates at most one unit of each good. The following lemma is a direct application of Theorem 4.4 of Cai and Zhao [2017] to our setting and gives the optimal revenue in terms of these dual variables.

Lemma 3.

$$\text{OPT-MPPH} = \min_{\lambda \geq 0} \max_{x \in \mathcal{P}} \sum_i \sum_t f(t) \phi_i(t) \sum_{S: i \in S} x_S(t) \eta_i(S).$$

This lemma allows us to move back and forth between the revenue in the primal space and a bound in the dual space.

Proof. Theorem 4.4 of Cai and Zhao [2017] states that the optimal revenue from a buyer with type $t \in T$ and *any* valuation $v(t, S)$ for the set S is as follows, where $x(t, S)$ is the primal variable for the probability that the buyer receives exactly set S when he reports type t :

$$\text{OPT-}v(\cdot, \cdot) = \min_{\lambda \geq 0} \max_{x \in \mathcal{P}} \sum_t f(t) \Phi(t, S) x_S(t)$$

where

$$\Phi(t, S) = v(t, S) - \frac{1}{f(t)} \sum_{t' \in T} \lambda(t', t)(v(t', S) - v(t, S)).$$

In our setting, we have that $v(t, S) = \sum_{i \in S} \eta_i(S)t_i$. Thus

$$\begin{aligned} \Phi(t, S) &= v(t, S) - \frac{1}{f(t)} \sum_{t' \in T} \lambda(t', t)(v(t', S) - v(t, S)) \\ &= \sum_{i \in S} \eta_i(S)t_i - \frac{1}{f(t)} \sum_{t' \in T} \lambda(t', t) \left(\sum_{i \in S} \eta_i(S)t'_i - \sum_{i \in S} \eta_i(S)t_i \right) \\ &= \sum_{i \in S} \eta_i(S) \left(t_i - \frac{1}{f(t)} \sum_{t' \in T} \lambda(t', t) (t'_i - t_i) \right) \\ &= \sum_{i \in S} \eta_i(S) \phi_i(t) \end{aligned}$$

and the above claim holds.⁵ Note that this also applies to the additive setting, where for all i , $\eta_{ij} = 0$ for all j and $\eta_i(S) = 1$. \square

We first relate OPT-PPC to the optimal revenue for an instance of additive valuations; in essence we just multiply the value t_i by $\eta_i([m])$. We set up some notation first. Define $\eta_{[m]}$ to be the vector whose i^{th} coordinate is $(\eta_{[m]})_i = \eta_i([m])$, and let $\eta_{[m]} \circ t$ be the Hadamard product of the vector $\eta_{[m]}$ and the vector t . Let \hat{F} be the distribution where $\eta_{[m]} \circ t$ is drawn identically to t in $F = \Pi_i F_i$, i.e., $\hat{f}(\eta_{[m]} \circ t) = f(t)$. We refer to this setting as the fully-boosted additive setting.

Proof of Lemma 1. For each i and allocation rule x , by the monotonicity in (1), the boost from $[m]$ is larger than that from any set S , i.e., $\eta_i(S) \leq \eta_i([m])$. Thus, we have that

$$\sum_{S: i \in S} x_S(t) \eta_i(S) \leq \eta_i([m]) \sum_{S: i \in S} x_S(t) = \eta_i([m]) \pi_i(t), \quad (2)$$

where we define $\pi_i(t) := \sum_{S: i \in S} x_S(t)$ to be the probability that item i is allocated to a buyer of type t . We now have the following sequence of equalities and inequalities. The first line uses Lemma 3 to move to the dual space. We would like to replace $\eta_i(S)$ by $\eta_i([m])$ everywhere (using (2)), but this is not possible since the virtual value function may be negative on some types. Lines 2 and 3 do this by using only non-negative virtual valuations as an upper bound. We use z^+ to denote $\max\{z, 0\}$ for any real number z . In line 4 we can bring back the original (possibly negative) virtual value function because in order to maximize this quantity, the optimal π must set $\pi_i(t) = 0$ when $\phi_i(t) < 0$. Line 5 then moves to the dual space for the fully-boosted additive setting, by suitably defining the dual variables there. (The exact duals are defined below.) Line 6 uses Lemma

⁵The theorem from Cai and Zhao [2017] also holds for multiple buyers, as does a restatement of Lemma 3; we only state it for a single buyer for simplicity.

3 once again to come back to the primal, OPT-ADDITIVE(\hat{F}).

$$\begin{aligned}
\text{OPT-PPC}(F) &= \min_{\lambda \geq 0} \max_{x \in \mathcal{P}} \sum_i \sum_t f(t) \phi_i(t) \sum_{S:i \in S} x_S(t) \eta_i(S) && \text{by Lemma 3} \\
&\leq \min_{\lambda \geq 0} \max_{x \in \mathcal{P}} \sum_i \sum_t f(t) (\phi_i(t))^+ \sum_{S:i \in S} x_S(t) \eta_i(S) \\
&\leq \min_{\lambda \geq 0} \max_{\pi} \sum_i \sum_t f(t) (\phi_i(t))^+ \cdot \eta_i([m]) \pi_i(t) && \text{by (2)} \\
&= \min_{\lambda \geq 0} \max_{\pi} \sum_i \sum_t f(t) \phi_i(t) \eta_i([m]) \pi_i(t) \\
&= \min_{\lambda \geq 0} \max_{\pi} \sum_i \sum_{\eta_{[m]} \circ t} \hat{f}(\eta_{[m]} \circ t) \hat{\phi}_i(\eta_{[m]} \circ t) \pi_i(\eta_{[m]} \circ t) && \text{by (3)} \\
&= \text{OPT-ADDITIVE}(\hat{F}) && \text{by Lemma 3.}
\end{aligned}$$

The equality in line 5 is true because if we set the dual variable $\hat{\lambda}(\eta_{[m]} \circ t', \eta_{[m]} \circ t) = \lambda(t', t)$ in the fully-boosted additive setting, $\hat{\lambda}$ still corresponds to a feasible dual variable⁶. Therefore, it induces the following virtual value function:

$$\begin{aligned}
\hat{\phi}_i(\eta_{[m]} \circ t) &= \eta_i([m]) \circ t_i - \frac{1}{\hat{f}(\eta_{[m]} \circ t)} \sum_{\eta_{[m]} \circ t'} (\eta_i([m]) t'_i - \eta_i([m]) t_i) \hat{\lambda}(\eta_{[m]} \circ t', \eta_{[m]} \circ t) \\
&= \eta_i([m]) t_i - \frac{1}{f(t)} \sum_{t'} \eta_i([m]) (t'_i - t_i) \lambda(t', t) \\
&= \eta_i([m]) \phi_i(t).
\end{aligned} \tag{3}$$

□

3.3 XOS Complementarities

For simplicity, our analysis is written for additive boosts. However, the extension to XOS boosts is fairly straight-forward. As shown in (1), XOS boosts are also monotone, so the upper bound from using $\eta_i([m])$ holds. We modify our graph construction from the proof of Lemma 2 as follows. Define $\ell_i^* \in \arg\max_{\ell \in [K]} \sum_{j \in [m] \setminus \{i\}} \eta_{ij}^\ell$; then $\eta_i([m]) = 1 + \sum_{j \in [m] \setminus \{i\}} \eta_{ij}^{\ell_i^*}$. Then in the XOS analysis, the directed edge (j, i) has weight $w_{j,i} := \eta_{ij}^{\ell_i^*} \cdot R_i$. A cut from $\{s\} \cup \mathcal{F}$ to $\bar{\mathcal{F}}$ will have thus have weight

$$\sum_{i \notin \mathcal{F}} \sum_{j \in \mathcal{F} + \{s\}} w_{j,i} = \sum_{i \in \bar{\mathcal{F}}} \left(1 + \sum_{j \in \mathcal{F}} \eta_{ij}^{\ell_i^*} \right) R_i \leq \sum_{i \in \bar{\mathcal{F}}} \eta_i(\mathcal{F}) R_i.$$

That is, the weight of the cut is a lower bound on the revenue of the mechanism with free set \mathcal{F} and items in $\bar{\mathcal{F}}$ priced accordingly, using the actual $\eta_i(\mathcal{F})$'s. Since a uniformly random \mathcal{F} guarantees a cut of weight $\frac{1}{4} \sum_i \eta_i([m]) \cdot R_i$ in expectation, then the expected revenue is again at least as high.

Similarly, in Lemmas 5 and 6, the same modification of using $w_{T,i} := \eta_{iT}^{\ell_i^*}$ on edges (T, i) will guarantee that the weight of any cut is again a lower bound on the corresponding SEPARATE/FREE revenue, so our random cut constructions give the same guarantees under XOS boosts as well.

⁶For readers familiar with CDW16, $\hat{\lambda}$ still corresponds to a flow.

Finally, it is not hard to see that even when the boosts are XOS functions, the revenue of selling the grand bundle is still the same as the fully-boosted additive $\text{BREV}(\hat{F})$.

Theorem 3. *The better of selling the grand bundle and Mechanism SEPARATE/FREE is a 12-approximation to the optimal revenue for XOS complementarities.*

4 Extension to MPPH

In this section, we show how to extend the mechanism and the analysis to the more general proportional positive hypergraphic (PPH) valuation class. Recall that η_{iT} may be defined for any subset $T \in [m] \setminus \{i\}$, and that $\eta_i(S) = 1 + \sum_{T \subseteq S \setminus \{i\}} \eta_{iT}$. Also recall that k is the directed-positive-rank of the hypergraph, and d is the maximum-out-degree. The extensions to the boosts being a maximum over many hypergraphs (MPPH) is covered in our analysis in Section 3.3.

The distribution for the fully-boosted additive setting \hat{F} is defined as before, except with $\eta_i([m])$ defined according to the PPH valuations.

Lemma 4 states that the fully-boosted additive setting is again a crude upper bound on revenue; it is the analog of Lemma 1 for PPH valuations and can be proven similarly.

Lemma 4.

$$\text{OPT-PPH}(F) \leq \text{OPT-ADDITIVE}(\hat{F}).$$

Next, we prove an analog of Lemma 2 which shows that we can obtain a $4k$ -approximation to $\text{SREV}(\hat{F})$.

Lemma 5.

$$\text{SREV}(\hat{F}) \leq 4k \text{SEPARATE/FREE-PPH}(F).$$

Proof. We use a random construction of the free set, and we show that the expected revenue of our mechanism is at least a $1/4k$ -fraction of $\text{SREV}(\hat{F})$. Each item independently is free (in \mathcal{F}) with probability $(1 - \frac{1}{2k})$, and otherwise it is priced. By definition of the directed-positive-rank, for every given η_{iT} , $|T| \leq k$. Then for any such T , all items in T appear simultaneously in \mathcal{F} with probability $(1 - \frac{1}{2k})^{|T|} \geq (1 - \frac{1/2}{k})^k \geq \frac{1}{2}$. In addition, every item $i \in \bar{\mathcal{F}}$ with probability $\frac{1}{2k}$.

Consider the graph construction where a directed edge (T, i) has weight $w_{T,i} := \eta_{iT} \cdot R_i$ and we have an edge (s, i) for every item i with weight $w_{s,i} := R_i$. Every edge is cut from $\{s\} \cup \mathcal{F}$ to $\bar{\mathcal{F}}$ with probability $\geq \frac{1}{2} \cdot \frac{1}{2k} = \frac{1}{4k}$. The expected weight of the cut from $\{s\} \cup \mathcal{F}$ to $\bar{\mathcal{F}}$ is then $\geq \frac{1}{4k} \sum_i \eta_i([m]) \cdot R_i = \frac{1}{4k} \cdot \text{SREV}(\hat{F})$.

Again, as in the proof of Lemma 2, we observe that the expected revenue of Mechanism SEPARATE/FREE with partition $(\mathcal{F}, \bar{\mathcal{F}})$ achieves at least as much revenue as the directed cut from $\{s\} + \mathcal{F}$ to $\bar{\mathcal{F}}$, and thus the mechanism obtains the $\frac{1}{4k}$ -approximation. \square

We prove in the next Lemma that there is a different way to choose the free set to obtain a $4d$ -approximation to $\text{SREV}(\hat{F})$.

Lemma 6.

$$\text{SREV}(\hat{F}) \leq 4d \text{SEPARATE/FREE-PPH}(F).$$

Proof. When the hypergraph has maximum-out-degree d , that is, d is the largest number of edges directed out of any item, a slightly different random construction of the free set gives a $4d$ -approximation to $\text{SREV}(\hat{F})$. For each hyperedge (T, i) , with probability $\frac{1}{2d}$, we place *all* items $j \in T$ into the free set. We run this process for every hyperedge (in some arbitrary order). If, after this process, an item j is not assigned to the free set, then item j is priced (placed into $\bar{\mathcal{F}}$). For any item i , the item is priced when none of the (at most d) edges that are directed from a set which contains i are placed into the free set, which occurs with probability at least $(1 - \frac{1}{2d})^d \geq \frac{1}{2}$. The probability of i being free is of course at least $\frac{1}{2d}$.

Then any edge (T, i) crosses the cut from $\{s\} + \mathcal{F}$ to $\bar{\mathcal{F}}$ with probability at least $\frac{1}{2d} \cdot \frac{1}{2} = \frac{1}{4d}$. Then by the same analysis as in the proof of Lemma 5, the expected weight of the cut from $\{s\} + \mathcal{F}$ to $\bar{\mathcal{F}}$ is at least $\frac{1}{4d} \sum_i \eta_i([m]) \cdot R_i = \frac{1}{4d} \cdot \text{SREV}(\hat{F})$, which is again a lower bound on the expected revenue of Mechanism SEPARATE/FREE with partition $(\mathcal{F}, \bar{\mathcal{F}})$. \square

Together, this gives

$$\begin{aligned} \text{OPT-PPH} \leq \text{OPT}(\hat{F}) &\leq 2 \text{SREV}(\hat{F}) + 4 \text{BREV}(\hat{F}) && \text{Lemma 4 and Theorem 6} \\ &\leq 8 \min\{k, d\} \text{SEPARATE/FREE-PPH} + 4 \text{BREV-PPH}. && \text{Lemmas 5 and 6} \end{aligned}$$

Theorem 4. *The better of selling the grand bundle and Mechanism SEPARATE/FREE for PPH valuations, with directed-positive-rank k and maximum-out-degree d , is an $(8 \min\{d, k\} + 4)$ -approximation to the optimal revenue.*

The analysis in Section 3.3 generalizes the guarantees to MPPH (from additive to XOS boosts).

4.1 Lower Bound of $O\left(\frac{1}{k} \text{OPT}(\hat{F})\right)$

In our analysis, we make two relaxations. First, we relax our benchmark from OPT-MPPH to the upper bound of $\text{OPT}(\hat{F})$. Second, we lower bound the revenue of our SEPARATE/FREE mechanism by undercounting the probabilities of sale.

It is extremely difficult to reason about the probability that a buyer will be interested in buying an item (or a set of items): her value may only be high enough if she buys multiple bundles simultaneously, or she may purchase a bundle even though her value for it is low because it improves her value for other bundles. Instead, we undercount this probability in the following manner: when the buyer is deciding whether to take a priced bundle of items B , we suppose that she only counts the boosts between items within that bundle B and the boost from the free items in \mathcal{F} . We refer to this lower bound on revenue as the *proxy revenue*.

In this section, we show that with respect to these two relaxations, for a reasonable class of simple mechanisms which includes ours, there exists an instance such that the proxy revenue of every mechanism from the class is a factor of k off from $\text{OPT}(\hat{F})$. Note that this does not imply that the proxy revenue of these mechanisms is far from OPT-MPPH, as we do not know how far OPT-MPPH is from the benchmark of $\text{OPT}(\hat{F})$; we also do not know how far the proxy revenue may be from the actual revenue.

Definition 1. A mechanism is from the class of *Bundle Pricing Mechanisms* \mathcal{B} if it computes prices as follows. The mechanism determines a partition of items into y priced bundles of size n_1, \dots, n_y and one free set \mathcal{F} . The j^{th} bundle B_j is priced at its monopoly reserve when counting (1) the boosts of the complementarities within the bundle and (2) the boosts to the items in B_j from the free set \mathcal{F} .

Theorem 5. *Among Bundle Pricing Mechanisms \mathcal{B} , no mechanism has proxy revenue better than $O\left(\frac{1}{k} \text{OPT}(\hat{F})\right)$, and SEPARATE/FREE with a random free set \mathcal{F} achieves this.*

Proof. Consider the following instance. There are m items, and the buyer's type for item $i \in [m]$ is

$$t_i = \begin{cases} 2^i & \text{w.p. } 2^{-i} \\ 0 & \text{otherwise.} \end{cases}$$

For every size- k set $T \in \binom{[m]}{k}$, for all items $i \notin T$, we have that $\eta_{iT} = c := \frac{m}{2^{\binom{m-1}{k}}}$. That is, the market structure is the directed complete graph of hyperedges of size exactly k . Any other hyperedge (T, i) where $|T| \neq k$ has weight $\eta_{iT} = 0$. In total, there are $\binom{m-1}{k}$ edges of weight $\frac{m}{2^{\binom{m-1}{k}}}$ into each item i , thus $\eta_i([m]) = 1 + \frac{m}{2}$.

Under these valuations and market parameters, for the random free set construction described in the previous section (pricing any item with probability $\frac{1}{2k}$), we get proxy revenue at least $(1 + \frac{m}{2})m\frac{1}{4k} = \frac{m+m^2}{4k}$.

We now show that the proxy revenue of every mechanism from \mathcal{B} is $O(\frac{m^2}{k})$, and is thus no better than a constant factor times the proxy revenue of our mechanism.

Lemma 7. *In the above construction, for every bundle of n items and a free set \mathcal{F} of size $|\mathcal{F}| = w$ items, the proxy revenue of the bundle is $O\left(\left(\binom{n}{k} + \binom{w}{k}\right) \cdot c\right)$, where $c = \frac{m}{2^{\binom{m-1}{k}}}$.*

Proof. We count the boosts that are incorporated into the proxy revenue for any item: from within its bundle, and from the free set. First, for any item i within some bundle B of size n , the boosts from within the bundle are exactly $\sum_{T \subseteq B \setminus \{i\}: |T|=k} \eta_{iT} = \binom{n-1}{k} \cdot c$. Then, the boosts that i gets from the free set are $\sum_{T \subseteq \mathcal{F}: |T|=k} \eta_{iT} = \binom{w}{k} \cdot c$. Together, i 's boosts accounted for in the proxy revenue are

$$\eta(B + \mathcal{F}) := \eta_i(B + \mathcal{F}) = \left(\binom{n-1}{k} + \binom{w}{k} \right) \cdot c.$$

We now show that for any bundle B of n items, the proxy revenue is at most $4 \cdot \eta(B + \mathcal{F})$. According to the way we undercount probability for the proxy revenue, for any ℓ , the (undercounted) probability that the buyer's value for bundle B is greater than $\eta(B + \mathcal{F})2^\ell$ is at most the probability that he has value at least 2^ℓ in base valuations, which is $\sum_{i=\ell}^m 2^{-i} \leq 1/2^{\ell-1}$.

Therefore, for any price for this bundle $p_B \in [\eta(B + \mathcal{F})2^\ell, \eta(B + \mathcal{F})2^{\ell+1}]$, the expected proxy revenue for this bundle is no more than $\eta(B + \mathcal{F})2^{\ell+1} \cdot 2^{-(\ell-1)} = 4 \cdot \eta(B + \mathcal{F})$. \square

1. The proxy revenue for selling separately is m . This is the proxy revenue earned from optimally selling the m items separately, without giving any item out for free. Posting a price of 2^i for each item i earns expected proxy revenue 1 for each of the m items.
2. $\text{BREV} \leq 2m$, by the proof of Lemma 7 when $n = m$ and $w = 0$.
3. For any mechanism $\mathcal{M} \in \mathcal{B}$, the proxy revenue of \mathcal{M} -PPH is no more than $O(m^2/k)$. Consider the mechanism that offers a free set of size w to the buyer, and then splits the remaining

$m - w$ items into y bundles where the j^{th} bundle is of size n_j . According to Lemma 7, the mechanism's proxy revenue is

$$O \left(\left(\binom{w}{k} \cdot y + \sum_{j=1}^y \binom{n_j}{k} \right) \cdot c \right).$$

Clearly, $\sum_{j=1}^y \binom{n_j}{k} \leq \binom{\sum_{j=1}^y n_j}{k} = \binom{m-w}{k}$. By definition of c , $O \left(\binom{m-w}{k} \cdot c \right) = O(m)$.

Next, we bound $\binom{w}{k} \cdot y$ by $\frac{w^k}{k!} \cdot (m - w)$. Then, by the AM-GM inequality,

$$w^k \cdot (m - w) = \left(\frac{w}{k} \right)^k \cdot (m - w) \cdot k^k \leq \left(\frac{m}{k+1} \right)^{k+1} \cdot k^k = O \left(\frac{m^{(k+1)}}{k+1} \right).$$

Combining everything, we have that

$$c \cdot \binom{w}{k} \cdot y \leq O \left(\frac{c}{k!} \cdot \frac{m^{(k+1)}}{k+1} \right) = O \left(\frac{m^2}{k} \right),$$

where again the definition of c kills the factor of $\frac{m^k}{k!}$.

□

5 A common generalization

As observed earlier, our model captures scenarios where the additional value from a combination of items depends on the base values for the items, whereas the common PH model captures scenarios where this is independent. We now present a common generalization of these two models. Consider the hypergraphic representation of a valuation function, i.e., where the valuation function is represented by

$$v(S) = \sum_{T \subseteq S} v_T$$

for some values v_T ; the T for which $v_T > 0$ are the hyperedges of the underlying hypergraph. Our model can be thought of as a special case where v_T is a linear combination of the base values for the items in T :

$$v_T = \sum_{i \in T} \eta_{iT} v_i.$$

More generally, one could have an arbitrary linear transformation from the type space to the hypergraphic representation: let $t = (t_1, t_2, \dots, t_d)$ be the type, for some dimension d , and

$$v_T = \sum_{i \in [d]} \eta_{iT} t_i.$$

An interpretation of this model is that, for each $i \in [d]$, t_i represents the buyer's value for some activity, and η_{iT} is the additional boost for that activity made possible by the buyer owning the combination of items in T . Assume that each t_i is independent of the others. This generalizes the PH model with independent v_T s: each hyperedge corresponds to a different activity, and is boosted only by itself. The model can be further extended to XOS boosts, i.e., a maximum of many linear combinations (as in MPH). We now give an example where such a model is useful.

Example 3: Consider a computing device such as a tablet, which has multiple uses, such as browsing the web, and taking notes. A buyer’s valuation for such a device can be modeled as a linear combination of his value for each of the activities it enables. Now consider an accessory such as a stylus. This makes some of the activities faster, such as taking notes. The additional value it provides can be modeled as a linear combination of values for the corresponding activities. Similarly a note-taking app also makes the note-taking activity more valuable. Moreover, it could be that a combination of a stylus and a compatible app has further added boost to the valuation for that activity.

This perspective is similar in spirit to the ‘subadditive with independent items’ model of Rubinstein and Weinberg [2015]. The types (t_1, t_2, \dots, t_m) are drawn from a product distribution of m spaces, one for each item; the space corresponding to each item itself can be multi-dimensional. The valuation function for a set $v(S)$ can be an arbitrary function that depends only on t_i for $i \in S$, subject to subadditivity.

What is the point of a model even more general than PH when we have seemingly strong lower bounds for PH? These lower bounds are for $\max\{\text{SREV}, \text{BREV}\}$, which are (by now) the standard pricing mechanisms for which upper bounds have been shown. While it makes sense to consider the simplest of the pricing schemes when it comes to upper bounds, lower bounds against such utterly simple pricing schemes are much less compelling. When it comes to items that are complements, where such pricing schemes may not be the most natural, such lower bounds are more of an indication that we need to study alternate pricing schemes, rather than a sign of hopelessness. A take away from our results is that suitably simple pricing schemes *could* give constant factor approximations for reasonably general valuation models with complements. It is too early to discard the hope for such results for PH and other generalizations.

References

- Ittai Abraham, Moshe Babaioff, Shaddin Dughmi, and Tim Roughgarden. 2012. Combinatorial auctions with restricted complements. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 3–16. 2
- Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. 2014. A simple and approximately optimal mechanism for an additive buyer. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. IEEE, 21–30. 2, 7, 10, 22
- Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. 2017. 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*. 589–590. <https://doi.org/10.1145/3033274.3085148> 12
- Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. 2016. A duality based unified approach to Bayesian mechanism design. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 926–939. 7, 10, 12, 14, 22
- Yang Cai and Mingfei Zhao. 2017. Simple Mechanisms for Subadditive Buyers via Duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*. ACM, New York, NY, USA, 170–183. <https://doi.org/10.1145/3055399.3055465> 2, 12, 13

- Shuchi Chawla, Jason D Hartline, and Robert Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM conference on Electronic commerce*. ACM, 243–251. 2
- Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, 311–320. 2
- Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation (EC '16)*. ACM, New York, NY, USA, 579–596. <https://doi.org/10.1145/2940716.2940756> 2
- Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017a. The Competition Complexity Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. *Proceedings of the Eighteenth ACM Conference on Economics and Computation, EC '17* (2017). 12
- Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. 2017b. A Simple and Approximately Optimal Mechanism for a Buyer with Complements: Abstract. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC '17)*. ACM, New York, NY, USA, 323–323. <https://doi.org/10.1145/3033274.3085116> 2, 12
- Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. 2015. A Unifying Hierarchy of Valuations with Complements and Substitutes. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*. 2
- Michal Feldman, Nick Gravin, and Brendan Lucier. 2015. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 123–135. 2
- Hu Fu, Chris Liaw, Pinyan Lu, and Zhihao Gavin Tang. 2017. The Value of Information Concealment. *arXiv preprint arXiv:1707.05875* (2017). 12
- S. Hart and P.J. Reny. 2012. *Maximal revenue with multiple goods: Nonmonotonicity and other observations*. 7
- Siqi Liu and Christos-Alexandros Psomas. 2017. On the Competition Complexity of Dynamic Mechanism Design On the Competition Complexity of Dynamic Mechanism Design On the Competition Complexity of Dynamic Mechanism Design On the Competition Complexity of Dynamic Mechanism Design. *arXiv preprint arXiv:1709.07955* (2017). 12
- Will Ma and David Simchi-Levi. 2015. Reaping the benefits of bundling under high production costs. *arXiv preprint arXiv:1512.02300* (2015). 10
- Aviad Rubinstein and S Matthew Weinberg. 2015. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 377–394. 2, 19

Andrew Chi-Chih Yao. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 92–109. 2

A Improved Additive Bound

We now give a proof which improves the 6-approximation by Babaioff et al. [2014] to 5.382.

Theorem 6. *For any $a > 0$,*

$$\text{OPT} \leq \left(2 + \frac{2}{a^2}\right) \cdot \text{BREV} + (a + 1) \cdot \text{SREV}.$$

In particular, if we choose $a = \sqrt[3]{4}$, then

$$\text{OPT} \leq \left(3 + \frac{3}{2}\sqrt[3]{4}\right) \cdot \max\{\text{SREV}, \text{BREV}\} \leq 5.382 \cdot \max\{\text{SREV}, \text{BREV}\}.$$

Proof of Theorem 6. We improve the analysis used in Cai et al. [2016], where they obtain an upper bound on OPT using duality. They further partition the upper bound into three parts:

$$\text{OPT} \leq \text{SINGLE} + \text{TAIL} + \text{CORE}.$$

The first term SINGLE is upper bounded by SREV. The second term TAIL is also upper bounded by SREV, but the first thing we show is that it can also be upper bounded by BREV.

Let item j 's value t_j be drawn from F_j independently, and $f_j(v_j)$ be the probability that $t_j = v_j$. Following the notation of Cai et al. [2016], we use R to denote SREV, and TAIL is defined as follows.

$$\text{TAIL} = \sum_{j \in [m]} \sum_{t_j > R} f_j(t_j) \cdot t_j \cdot \Pr_{t_{-j} \sim F_{-j}}[\exists \ell \neq j, t_\ell \geq t_j].$$

This quantity is the expected value above r from all but the highest item. Note that for any j and any t_j , selling the grand bundle at a price of t_j earns revenue at least $t_j \cdot \Pr_{t_{-j} \sim F_{-j}}[\exists \ell \neq j, t_\ell \geq t_j]$. Hence,

$$\text{TAIL} \leq \text{BREV} \cdot \left(\sum_{j \in [m]} \sum_{t_j > R} f_j(t_j) \right) \leq \text{BREV}.$$

The second inequality is because R_j is the optimal revenue for selling only item j , and $R_j \geq R \cdot \Pr[t_j \geq R]$, thus $\sum_{t_j > R} f_j(t_j) \leq R_j/R$; also, $R = \sum_j R_j$.

Next, we improve the analysis of the term CORE. In Cai et al. [2016], CORE is upper bounded by $2\text{BREV} + 2\text{SREV}$. They make use of Chebyshev's inequality to obtain this bound. We improve their analysis using a tighter inequality due to Cantelli.

The CORE is defined as follows.

$$\text{CORE} = \sum_{j \in [m]} \sum_{t_j \leq R} f_j(t_j) \cdot t_j = \mathbb{E}_{t \sim F} \left[\sum_{j \in [m]} t_j \cdot \mathbf{1}[t_j \leq R] \right]$$

It is shown in Cai et al. [2016] that $\text{Var}_{t \sim F} \left[\sum_{j \in [m]} t_j \cdot \mathbf{1}[t_j \leq R] \right] \leq 2R^2$. Now we state Cantelli's inequality:

Theorem 7 (Cantelli's Inequality). *For any real valued random variable X and any positive number τ ,*

$$\Pr[X \geq \mathbb{E}[X] - \tau] \geq \frac{\tau^2}{\tau^2 + \text{Var}[X]}.$$

We define the random variable $V = \sum_{j \in [m]} t_j \cdot \mathbf{1}[t_j \leq R]$ and apply Cantelli's inequality to V with $\tau = aR$.

$$\Pr[V \geq \text{CORE} - aR] \geq \frac{a^2 R^2}{\text{Var}[V] + a^2 R^2} \geq \frac{a^2}{2 + a^2}.$$

The last inequality is because $\text{Var}[V] \leq 2R^2$. Therefore, $\text{BREV} \geq (\text{CORE} - aR) \cdot \frac{a^2}{2+a^2}$, which implies $\text{CORE} \leq (1 + \frac{2}{a^2}) \cdot \text{BREV} + a \cdot \text{SREV}$. Combining our new analysis for the TAIL and the CORE, we obtain the new bound. \square