# The Power of Uncertainty: Bundle-Pricing for Unit-Demand Customers

Patrick Briest[1] and Heiko Röglin[2]

[1] Department of Computer Science, University of Paderborn, Germany
patrick.briest@upb.de
[2] Department of Computer Science, University of Bonn, Germany
heiko@roeglin.org

**Abstract.** We study an extension of the unit-demand pricing problem in which the seller may offer bundles of items. If a customer buys such a bundle she is guaranteed to get one item out of it, but the seller does not make any promises of how this item is selected. This is motivated by the sales model of retailers like hotwire.com, which offers bundles of hotel rooms based on location and rating, and only identifies the booked hotel after the purchase has been made.

As the selected item is known only in hindsight, the buying decision depends on the customer's belief about the allocation mechanism. We study strictly pessimistic and optimistic customers who always assume the worst-case or best-case allocation mechanism relative to their personal valuations, respectively. While the latter model turns out to be equivalent to the pure item pricing problem, the former is fundamentally different, and we prove the following results about it: (1) A revenue-maximizing pricing can be computed efficiently in the uniform version, in which every customer has a subset of items and the same non-zero value for all items in this subset and a value of zero for all other items. (2) For non-uniform customers computing a revenue-maximizing pricing is APX-hard. (3) For the case that any two values of a customer are either identical or differ by at least some constant factor, we present a polynomial time algorithm that obtains a constant approximation guarantee.

## 1 Introduction

Algorithmic pricing deals with the problem of efficiently determining revenue-maximizing ways of selling a collection of items given information about the preferences of the potential customers in the target market. The traditional way of selling items consists of posting a price for each individual item, then letting customers pick the bundle of items they prefer and charging them the sum of prices of items they select. The problem of finding such an *item pricing* under various kinds of customer preferences has received a lot of attention recently [1, 2, 6, 7, 14] and in many cases, its approximation complexity is quite well understood.

Two fundamental classes of customer preferences have been particularly well investigated. Customers are referred to as *single-minded*, if items constitute *strict complements* and each customer is interested in purchasing one particular subset of the items. We say that customers are *unit-demand*, if items are *strict substitutes* and each customer is interested in buying a single item out of some set of alternatives. Assuming unlimited supply of all the items, it is known that in both of these settings the revenue-maximizing item pricing can be approximated within factors that are logarithmic in the number of customers or linear in the number of items [14] and, under appropriate complexity theoretic assumptions, no essential improvement beyond these guarantees is possible [4, 11].

In a recent paper Briest et al. [5] consider the unit-demand pricing problem, but instead of trying to find a revenue-maximizing item pricing, allow to sell items via a so-called *system of lotteries*. Here, rather than posting prices for individual items, the seller may offer a collection of *lottery tickets*, each such ticket representing a probability distribution over items and an associated price. A customer purchasing some given ticket will be asked to pay its price and in turn receive an item randomly sampled from the ticket's distribution. It is shown in [5] that depending on the details of the underlying model of customer behavior this larger class of selling mechanisms can lead to a significant increase in revenue while simultaneously allowing for much better algorithmic solutions.

While, apart from its intrinsic connection to the process of haggling in price negotiations [16, 17], we are not aware of any lottery-like pricing mechanism being applied directly in practice at this point, a related - yet different - method is being employed by several companies. This method also consists of *bundling* subsets of items and pricing bundles of items rather than individual items, and again the understanding is that a unit-demand customer will receive a single item from the bundle she picks. The crucial difference to the lottery-based system described above lies in the fact that the seller does not make any promises as to how the item allocated to the customer will be selected. This might be done according to some probability distribution unknown to the customers, but it might also be done in any other conceivable fashion, e.g. guided by production costs or availability. One prominent example of a retailer employing this kind of pricing scheme is the website hotwire.com, where hotel rooms are bundled according to their location and star rating. Customers can book offers of the form "3 nights in the Philadelphia downtown area, 4 stars and up for $279" and will receive information about the exact hotel they will be staying at only after payment has been made.

Formalizing this kind of unit-demand bundle-pricing problem brings up some modeling issues. Rational unit-demand customers are commonly modeled as assigning a *value* to each of the items and upon observing the item prices selecting the item maximizing their *utility*, defined as the difference between the customer's respective value and its price. While this concept of rationality extends quite naturally to lottery-based pricing (every lottery holds a fixed *expected* utility to a customer), it is not obvious what to do in the bundle pricing setting. A customer's value for a given bundle depends on the actual item she receives and, thus, will only be known in hindsight and, consequently, the buying decision itself has to depend on the customer's *belief* about the seller's allocation mechanism. In this paper, we will investigate the two most basic ways of modeling these beliefs and assume that customers assign to each bundle either the minimum or maximum value of any item contained in it. Intuitively, this corresponds to strictly *pessimistic* or *optimistic* customers who will always assume the worst-case or best-case allocation mechanism relative to their personal valuations, respectively. Before we give an overview of the results presented in this paper, let us introduce the problem more formally.

### 1.1 Preliminaries

The standard *unit-demand pricing problem* (UDP) is defined as follows. Given a set of items $\mathcal{I}$, each available in unlimited supply, and a collection of customers $\mathcal{C}$, each described by a *valuation function* $v_c : \mathcal{I} \to \mathbb{R}_0^+$, we want to assign prices to the items such as to maximize the overall revenue. More precisely, we assume that given prices $p(i)$ for all $i \in \mathcal{I}$, a customer $c \in \mathcal{C}$ will choose to purchase item

$$i_c(p) = \text{argmax}_{i \in \mathcal{I}}(v_c(i) - p(i)),$$

whenever that item's price does not exceed her respective value. To avoid technicalities, we assume that there is a special item $\varnothing$ with $v_c(\varnothing) = 0$, which is always assigned price 0. The quantity $v_c(i) - p(i)$ is termed customer $c$'s *utility* from purchasing item $i$ at price $p(i)$. We will also assume that whenever there are multiple items yielding identical utility, a customer will pick the one with highest price among them.[3] In this way, item $i_c(p)$ is well defined for any set of prices $p$. The revenue of a price assignment $p$ is

$$\text{rev}(p) = \sum_{i \in \mathcal{I}} p(i_c(p)).$$

In the *unit-demand bundle-pricing problem* (UDBP) considered in this paper we are again given a ground set $\mathcal{I}$ of items and a collection $\mathcal{C}$ of unit-demand customers. The output is a collection $\mathcal{B} \subseteq 2^{\mathcal{I}}$ of *bundles* of items and prices $p(B)$ for all bundles $B \in \mathcal{B}$. If a customer decides to purchase a bundle $B$ of items, she is guaranteed to receive an item from $B$. However, a customer does not have any information regarding the details of how the particular item she will receive is selected once the bundle is bought. Consequently, a customer's value for any given bundle has to depend on her belief about the selection procedure.

Most of this paper will be focused on the case of pessimistic customers who will assign to each bundle its worst-case value (UDBP-MIN). Formally, a customer with (unit-demand) valuation function $v_c : \mathcal{I} \to \mathbb{R}_0^+$ will value bundle $B \subseteq \mathcal{I}$ at

$$\bar{v}_c(B) = \min_{i \in B} v_c(i).$$

[3] This assumption is w.l.o.g., since in case of a tie decreasing all prices by a factor of $(1 - \varepsilon)$ for an arbitrary value of $\varepsilon$ ensures that for each customer the utility-maximizing item is one of maximal price.

As in standard unit-demand pricing, given prices $p$ each customer will purchase her utility-maximizing bundle

$$B_c(p) = \text{argmax}_{B \in \mathcal{B}}(\bar{v}_c(B) - p(B)),$$

where we assume that the empty bundle $\emptyset \in \mathcal{B}$ has price $p(\emptyset) = 0$ and is valued at 0 by all customers. Furthermore, ties are again broken in favor of more expensive bundles.

There are of course numerous other ways of extending unit-demand valuation functions to the set of all bundles. In this paper, we will also briefly look at the complementing case of customers assigning each bundle its best-case value (UdBp-Max), formally,

$$\bar{v}_c(B) = \max_{i \in B} v_c(i).$$

Other models, in particular those assuming some kind of probabilistic selection method, are beyond the scope of this paper, but might also be of much interest, particularly as some of them are essentially variations of the *lottery* concept investigated in [5] and might have applications in the design of truthful revenue-maximizing auction mechanisms [3, 12, 15].

By *uniform* UdBp we refer to the restricted problem version in which customers' valuation functions assign identical (positive) values to some subset of the items and value 0 to all items in the complement of this subset. Formally, every customer is characterized by the set $S_c \subseteq \mathcal{I}$ of items she desires and her value $v_c \in \mathbb{R}^+$ for receiving any such item.

### 1.2 Contributions

We will first consider UdBp-Min and present a number of algorithmic and complementing hardness results. In Section 2 we present a polynomial time algorithm for uniform UdBp-Min, which is essentially based on two main ingredients. First, we observe that the number of bundles that might be part of an optimal bundle-pricing is small and, in fact, we can derive the set of bundles we need to consider immediately from the given set of customers. We then show that the extension of the valuation functions to this collection of bundles is very nicely structured, as a consequence of which one can apply techniques from [8] to solve the problem. More precisely, if we define a relation between bundles depending on whether there exists a customer who strictly prefers one of them to the other, this relation turns out to be transitive, as a consequence of which we can reduce the bundle-pricing problem to solving a weighted independent set problem in a perfect graph.

We proceed by considering the general (i.e., non-uniform) case of UdBp-Min which turns out to be significantly more complex. In Section 3 we show that general UdBp-Min is APX-hard. This is true even if all customers have non-zero values for at most 2 items and there are only 3 distinct values among all of them. The main distinction of our reduction from previous hardness results for unit-demand pricing problems stems from the fact that because of the enlarged solution space (containing all possible bundles of items) we need to argue about a significantly larger set of potential solutions to prove that the reduction is indeed approximation preserving.

On the algorithmic side, we introduce the concept of $\alpha$-*coarse* instances, in which any two values of a single customer must be identical or differ by a factor of at least $\alpha$. We present a polynomial time algorithm that obtains a constant approximation guarantee for any given constant value of $\alpha > 1$. This is an interesting distinction from the several related item pricing problems, where the known inapproximability results suggest that coarse instances in particular seem to form the hard core of the problem [4]. The algorithm is based on a novel reduction of the general to the uniform problem, in the process of which each general valuation function is simulated by a carefully tailored collection of uniform valuation functions yielding similar revenue under all relevant pricings. This reduction is also interesting in its own right, as it can be applied to other unit-demand pricing problems as well, yielding similar algorithmic results as long as the uniform problem version allows for a good approximation. In particular, we can obtain constant factor approximation algorithms for $\alpha$-coarse instances of UdP with *price-ladder constraint* [1], i.e., when the relative order of item prices is predetermined, as the uniform version of this problem is known to be solvable in polynomial time via dynamic programming. These results are found in Section 4.

Finally, we briefly turn to UdBp-Max and show that this problem behaves fundamentally different from UdBp-Min. In Section 5 we argue that the problem turns out to be equivalent to the pure item pricing problem and, thus, all results known for UdP carry over in this case.

## 2 A Polynomial-Time Algorithm for Uniform UdBp-Min

The first main ingredient for our polynomial-time algorithm for uniform UdBp-Min are the following observations regarding the structure of the optimal collection of bundles and their prices. Note, that Definition 1 and Proposition 2 also apply to non-uniform UdBp-Min.

**Definition 1.** *For a customer $c$ with valuation function $v_c : \mathcal{I} \to \mathbb{R}_0^+$ we let*

$$L_c^v = \{i \in \mathcal{I} \,|\, v_c(i) \geq v\}.$$

*We say that $L_c^v$ is customer $c$'s* level-$v$ *set.*

**Proposition 2.** *Let $(\mathcal{I}, \mathcal{C})$ be an instance of* UdBp-Min. *Then there exists a revenue-maximizing collection $\mathcal{B}$ of bundles with corresponding prices $p$, such that $\mathcal{B} \subseteq \{L_c^v \,|\, c \in \mathcal{C}, v \in \mathbb{R}_0^+\}$.*

**Proposition 3.** *Let $(\mathcal{I}, \mathcal{C})$ be an instance of uniform* UdBp-Min *and $(\mathcal{B}, p)$ a revenue-maximizing solution. It holds w.l.o.g. that $\mathcal{B} \subseteq \{S_c \,|\, c \in \mathcal{C}\}$ and $p(i) \in \mathcal{P}$ for all $i \in \mathcal{I}$, where $\mathcal{P} = \{v_c \,|\, c \in \mathcal{C}\}$.*

The proofs of Propositions 2 and 3 are left for the full version of this paper. Proposition 3 states that in the uniform case of UdBp-Min both the set of possible bundles and the set of possible prices that can appear as part of an optimal solution are quite managable. In particular, the problem of computing an optimal bundle pricing reduces to deciding which customers should purchase their respective bundles $S_c$ and at which price from $\mathcal{P}$.

Similar to the approach first introduced in [8], we will transform the problem of computing the optimal bundle pricing into a weighted independent set problem and argue that the resulting graph is perfect, which allows us to solve the independent set problem in polynomial time [13]. We use $(c, S_c, p)$ to denote the fact that customer $c$ purchases bundle $S_c$ at price $p$. We create a vertex with label $(c, S_c, p)$ and weight $p$ for every $c \in \mathcal{C}$ and $p \in \mathcal{P}$ with $p \leq v_c$. Then we create a directed edge from the vertex with label $(c, S_c, p)$ to the vertex with label $(d, S_d, q)$, if and only if $S_d \subseteq S_c$ and $q < p$. Let us refer to the resulting directed graph as $G$ and let $\tilde{G}$ refer to the same graph but with undirected edges.

**Lemma 4.** *Graph $\tilde{G}$ as constructed above is perfect.*

The proof of Lemma 4, which is an application of the strong perfect graph theorem [9] and essentially similar to the proof given in [8], is omitted due to space limitations. Lemma 4 immediately yields a polynomial time algorithm for uniform UdBp-Min.

**Algorithm 1:** Poly-Time Algorithm for Uniform UdBp-Min.

**(1)** Given instance $(\mathcal{I}, \mathcal{C})$, construct the perfect weighted graph $\tilde{G}$ containing a vertex with label $(c, S_c, p)$ and weight $p$ for all $c \in \mathcal{C}$, $p \in \mathcal{P}$ with $p \leq v_c$, and an edge between vertices with labels $(c, S_c, p)$ , $(d, S_d, q)$ iff either $S_c \subseteq S_d$ and $p < q$ or $S_d \subseteq S_c$ and $q < p$.

**(2)** Find a maximum weight independent set in $\tilde{G}$.

**(3)** For each vertex in the independent set, if it has label $(c, S_c, p)$, offer bundle $S_c$ at price $p$.

**Theorem 5.** *Algorithm 1 returns a revenue-maximizing bundle pricing in polynomial time.*

*Proof.* By Proposition 3 there always exists a revenue-maximizing bundle pricing in which every customer $c \in \mathcal{C}$ buys bundle $S_c$ or nothing at all and all prices are chosen from the set $\mathcal{P}$. Clearly, for any customer purchasing bundle $S_c$ at price $p$ it must be the case that no bundle $B \subseteq S_c$ is offered at a price below $p$, as buying this bundle would yield higher utility for customer $c$. Consequently, the bundle pricing corresponds to an independent set in $\tilde{G}$ of total weight equal to the revenue obtained by the bundle pricing. Similarly, setting prices according to an independent set in $\tilde{G}$ ensures that customer $c$ purchase bundle $S_c$ at price $p$ whenever the vertex with label $(c, S_c, p)$ is part of the independent set. Thus, we have a one-to-one correspondence between bundle pricings and weighted independent sets in $\tilde{G}$ and it follows that Algorithm 1 returns a revenue-maximizing pricing.

Polynomial running time follows from the observation that graph $\tilde{G}$ has at most a polynomial number $|\mathcal{C}| \cdot |\mathcal{P}|$ of vertices and the fact that it is perfect by Lemma 4, so finding a maximum weight independent set can be done in polynomial time. □

## 3 Hardness of Approximation of UdBp-Min

In this section we show that UdBp-Min is APX-hard. In particular, we show that there is no polynomial time approximation algorithm achieving a revenue of at least $(428/429 + \epsilon)$ times the revenue of the optimal bundle pricing, unless P=NP. This is even true if every customer has non-zero value for at most two items and there are only three different values among all of them.

**Theorem 6.** *It is NP-hard to approximate* UdBp-Min *within* $c = \frac{428}{429} + \epsilon$ *for any* $\epsilon > 0$.

*Proof.* Our reduction is from the unweighted Max Dicut problem. An instance of this problem is a directed graph $G = (V, E)$, and the goal is to find a partition of $V$ into $(S, V \setminus S)$, where $S \subseteq V$, such that the number of edges that cross this cut, i.e., edges $(u, v)$ such that $u \in S$ but $v \notin S$, is maximized. This problem is not $(12/13 + \epsilon)$-approximable, for any constant $\epsilon > 0$, unless P=NP [10].

Given an instance $G = (V, E)$ of the unweighted Max Dicut problem, we create an instance of UdBp-Min by introducing one item for each node in $V$ and 48 customers for each edge in $E$. For the edge $(u, v) \in E$, we introduce customers with value 0 for all items in $V \setminus \{u, v\}$ and with the following values for $u$ and $v$:

| number of customers $c$ | 9 | 3 | 3 | 15 | 1 | 3 | 6 | 2 | 6 |
|---|---|---|---|---|---|---|---|---|---|
| $v_c(u)$ | 0 | 0 | 0 | 1 | 1 | 2 | 2 | 2 | 4 |
| $v_c(v)$ | 1 | 2 | 4 | 0 | 4 | 0 | 1 | 4 | 0 |

Given these customers, the only bundles for which there exist customers with non-zero valuation are singleton bundles and bundles $\{u, v\}$ for edges $(u, v) \in E$. Hence, we can focus on setting prices for these bundles. Furthermore, any pricing can be transformed into a pricing using only the prices 1, 2, 3, and 4 and achieving at least the same revenue as follows: If we have a pricing with prices below 1, we can first increase all these prices to 1 without decreasing the revenue, then we can increase all prices strictly between 1 and 2 to 2 without decreasing the revenue and so on.

If we have already set prices for the singleton bundles $\{u\}$ and $\{v\}$ and there is an edge $(u, v) \in E$, then the (not necessarily unique) price for the bundle $\{u, v\}$ maximizing the revenue is determined. A case analysis yields the following table showing an optimal price for $\{u, v\}$ and the total revenue obtained from all customers belonging to edge $(u, v)$ for the different choices of $p(\{u\})$ and $p(\{v\})$:

| $p(\{u\})$ | 1 | 1 | 1 | 1 | 2 | 2 | 2 | 2 | 3 | 3 | 3 | 3 | 4 | 4 | 4 | 4 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| $p(\{v\})$ | 1 | 2 | 3 | 4 | 1 | 2 | 3 | 4 | 1 | 2 | 3 | 4 | 1 | 2 | 3 | 4 |
| opt. price for $p(\{u, v\})$ | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 2 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| total revenue | 48 | 48 | 48 | 48 | 48 | 48 | 48 | 54 | 42 | 42 | 42 | 42 | 48 | 48 | 48 | 48 |

The following information from the previous table is crucial:

- For any pricing with $p(\{u\}) \neq 3$ and $p(\{v\}) \neq 3$, there is a choice for $p(\{u,v\})$ for which all customers belonging to edge $(u,v)$ yield a total revenue of 48.
- For any pricing with $p(\{u\}) \neq 2$ or $p(\{v\}) \neq 4$, there is no choice for $p(\{u,v\})$ for which they yield a larger revenue than 48.
- If $p(\{u\}) = 2$ and $p(\{v\}) = 4$, then there is a choice for $p(\{u,v\})$ for which all customers belonging to edge $(u,v)$ yield a total revenue of 54.

Based on this information, we can relate the maximum directed cut in the graph $G$ and the revenue-maximizing pricing for the instance of UdBp-Min that we have constructed. If the graph $G$ has a cut $(S, V \setminus S)$ crossed by $\ell$ edges, then there exists a pricing for the instance of UdBp-Min with a revenue of $48 \cdot (|E| - \ell) + 54 \cdot \ell$. For this, we assign a price of 2 to every set $\{u\}$ with $u \in S$ and a price of 4 to every set $\{u\}$ with $u \notin S$. The prices for the sets $\{u,v\}$ for edges $(u,v) \in E$ are chosen according to the previous table.

If the optimal pricing of the instance of UdBp-Min yields a revenue of $48 \cdot (|E| - \ell) + 54 \cdot \ell$ for some $\ell \in \mathbb{N}$, then there exists a cut $S$ in $G$ that is crossed by $\ell$ edges. To see this, we can assume w.l.o.g. that in the optimal pricing all singleton bundles have a price of either 2 or 4 because replacing every price of 1 or 3 by a price of 4 does not decrease the revenue. Then if $S$ consists of exactly those nodes whose corresponding singleton bundle has a price of 2, there are $\ell$ edges crossing the cut $(S, V \setminus S)$.

An optimal directed cut of any graph $G = (V, E)$ is crossed by at least $|E|/4$ edges. In order to see this, consider the undirected (multi)-graph $G'$ obtained from $G$ by removing the directions of the edges. In a maximum undirected cut $(S, V \setminus S)$ of $G'$ at least half of the edges have one endpoint in $S$ and one endpoint in $V \setminus S$. This means that at least a quarter of the edges go from $S$ to $V \setminus S$ or at least a quarter of the edges go from $V \setminus S$ to $S$.

Assume there was an algorithm achieving a $(\frac{428}{429} + \epsilon)$-approximation for UdBp-Min. Let $\ell^*$ denote the maximum number of edges crossing any cut in graph $G$, then $48 \cdot (|E| - \ell^*) + 54 \cdot \ell^* = 48 \cdot |E| + 6 \cdot \ell^*$ is the revenue of the optimal pricing in the instance of UdBp-Min described above. Hence, the algorithm computes a pricing with revenue $48 \cdot |E| + 6 \cdot \ell$ with

$$\frac{48 \cdot |E| + 6 \cdot \ell}{48 \cdot |E| + 6 \cdot \ell^*} \geq c = \frac{428}{429} + \epsilon .$$

From this, we derive

$$\ell \geq c \cdot \ell^* - 8 \cdot |E| \cdot (1 - c) \geq c \cdot \ell^* - 32 \cdot \ell^* \cdot (1 - c) = \ell^* \cdot (33c - 32) \geq \ell^* \cdot \left(\tfrac{12}{13} + \epsilon\right).$$

Hence, $\frac{\ell}{\ell^*} \geq \frac{12}{13} + \epsilon$, contradicting the hardness of the Max Dicut problem. □

## 4 Approximation Algorithm for Non-Uniform UdBp-Min

**Definition 7.** *For a customer $c$ let $\mathcal{V}_c = \{v \mid \exists i \in \mathcal{I} : v_c(i) = v\}$ denote the range of her valuation function. We say that a* UdBp-Min *instance $(\mathcal{I}, \mathcal{C})$ is* $\alpha$-coarse *for some $\alpha > 1$, if for every $c \in \mathcal{C}$ and all $v, v' \in \mathcal{V}_c$ with $v \neq v'$ it holds that either $v \geq \alpha v'$ or $v' \geq \alpha v$.*

**Algorithm 2:** Approximation Algorithm for $\alpha$-coarse UDBP-MIN.

**(1)** Given an $\alpha$-coarse instance $(\mathcal{I}, \mathcal{C})$, construct a uniform instance $(\mathcal{I}, \mathcal{C}')$ as follows: For every $c \in \mathcal{C}$ and every $v \in \mathcal{V}_c$, add a customer $c(v)$ with value $v$ and set of desired items $S_{c(v)} = L_c^v$.

**(2)** Compute an optimal solution $(\mathcal{B}, p)$ on this uniform instance.

**(3)** Return $\big(\mathcal{B}, (1-\alpha^{-1})p\big)$.

**Theorem 8.** *Algorithm 2 achieves approximation guarantee* $(1/4)(1-\alpha^{-1})^2$ *on* $\alpha$*-coarse instances of* UDBP-MIN.

Theorem 8 above is an immediate consequence of the following two lemmas.

**Lemma 9.** *Let* $R^*$ *denote the optimal revenue obtainable on* $\alpha$*-coarse instance* $(\mathcal{I}, \mathcal{C})$ *and* $R'$ *the maximum revenue obtainable on the uniform instance* $(\mathcal{I}, \mathcal{C}')$ *constructed by the algorithm. It holds that* $R' \geq R^*$.

*Proof.* We have seen in Proposition 2 that the collection of bundles $\mathcal{B}^*$ sold in the revenue-maximizing solution of instance $(\mathcal{I}, \mathcal{C})$ consists only of level sets of the customers from $\mathcal{C}$. Now assume that we offer all bundles from $\mathcal{B}^*$ at the same prices to the uniform customers constructed by the algorithm. For each customer $c \in \mathcal{C}$ purchasing her level-$v$ set $L_c^v$ at price $p \leq v$, we have a uniform customer $c(v)$ with value $v$ for any item in $L_c^v$ by construction, both of which experience utility $v - p$ from buying bundle $L_c^v$. On the other hand, customer $c(v)$'s values for all items are no larger than those of customer $c$ and, consequently, she values no bundle higher than $c$. Since buying $L_c^v$ at price $p$ is the utility maximizing choice for $c$, so it is for $c(v)$ and it follows that we collect as much revenue from $c(v)$ in the uniform instance as we do from $c$ in the original instance. Summing over all $c \in \mathcal{C}$ yields the claim. ☐

**Lemma 10.** *Let* $(\mathcal{B}, p)$ *be an optimal solution to the uniform* UDBP-MIN *instance* $(\mathcal{I}, \mathcal{C}')$ *constructed by the algorithm resulting in revenue* $R'$*. Then solution* $(\mathcal{B}, (1-\alpha^{-1})p)$ *yields revenue at least* $(1/4)(1-\alpha^{-1})^2 R'$ *on the original* $\alpha$*-coarse instance* $(\mathcal{I}, \mathcal{C})$.

*Proof.* Let $(\mathcal{B}, p)$ be an optimal solution to the uniform UDBP-MIN instance $(\mathcal{I}, \mathcal{C}')$. By Proposition 3 we may w.l.o.g. assume that $\mathcal{B}$ is a subset of the desired sets of customers from $\mathcal{C}'$ and so it is also a subset of the level sets of the original non-uniform customers from $\mathcal{C}$. We can also assume w.l.o.g. that every customer who buys a set buys her desired set and no subset.

Let $\mathcal{C}'_+ = \{c(v) \in \mathcal{C}' \,|\, L_c^v \in \mathcal{B} \text{ and } v/2 \leq p(L_c^v) \leq v\}$ denote the set of customers who purchase their set of desired items at a price of at least half their value. Note, that it must be the case that customers in $\mathcal{C}'_+$ contribute total revenue of at least $R'/2$. It is easy to argue that if this was not the case, multiplying all prices by a factor of 2 would increase overall revenue, contradicting the optimality of $(\mathcal{B}, p)$.

Let us refer to the original set of non-uniform customers which have at least one corresponding customer in $\mathcal{C}'_+$ as $M$ and sort the customers in $\mathcal{C}'_+$ according to the non-uniform customer they represent and their values. For a customer $c \in M$, we define $\ell_c$ to be the number of corresponding customers in $\mathcal{C}'_+$. Formally, let us denote

$$\mathcal{C}'_+ = \bigcup_{c\in M} \bigcup_{i=1}^{\ell_c} \Big\{ c\big(v_i^c\big) \Big\},$$

where $v_1^c > v_2^c > \cdots > v_{\ell_c}^c$ for all $c \in M$. Now let $\mathcal{C}'_* = \bigcup_{c\in M} \{c(v_1^c)\}$ be the thinned out version of $\mathcal{C}'_+$ which only contains the uniform customer with highest value for each original customer $c \in M$. Let $R'_*$ be the total revenue collected from customers in $\mathcal{C}'_*$. It holds that

$$\begin{aligned} R'_* &\geq \frac{1}{2} \sum_{c\in M} v_1^c = \frac{1}{2}\big(1-\alpha^{-1}\big) \sum_{c\in M} \left( \sum_{i=0}^{\infty} \alpha^{-i} \right) v_1^c \\ &= \frac{1}{2}\big(1-\alpha^{-1}\big) \sum_{c\in M} \sum_{i=0}^{\infty} (\alpha^{-i} v_1^c) \geq \frac{1}{2}\big(1-\alpha^{-1}\big) \sum_{c\in M} \sum_{i=1}^{\ell_c} v_i^c \geq \frac{1}{2}\big(1-\alpha^{-1}\big) \frac{1}{2} R', \end{aligned}$$

where we use the facts that $v_i^c \leq \alpha^{-i+1} v_1^c$ since instance $(\mathcal{I}, \mathcal{C})$ is $\alpha$-coarse and customer $c(v_i^c)$ cannot contribute more than $v_i^c$ to the overall revenue of at least $R'/2$ collected from customers in $\mathcal{C}'_+$.

Finally, let us fix a single uniform customer $c(v) \in \mathcal{C}'_*$ purchasing bundle $L_c^v$ at price $p$. We observe that it must be the case that all bundles $B \subset L_c^v$ must have a price of at least $p$, as otherwise purchasing $L_c^v$ could not be $c(v)$'s utility maximizing choice. Now consider the non-uniform customer $c$ corresponding to $c(v)$ and the effect of reducing all prices by a factor of $(1-\alpha^{-1})$. Customer $c$ has utility

$$v - \big(1-\alpha^{-1}\big)p \geq v - \big(1-\alpha^{-1}\big)v = \alpha^{-1} v$$

from purchasing bundle $L_c^v$ at price $(1-\alpha^{-1})p$. Her value for any bundle containing items from outside $L_c^v$ is at most $\alpha^{-1}v$ by $\alpha$-coarseness, so none of these bundles can yield higher utility even at price 0. Bundles strictly contained in $L_c^v$ could potentially yield higher utility, but by our argument above the price of any such bundle is at least $(1-\alpha^{-1})p$ after decreasing prices and we conclude that customer $c$ contributes at least as much revenue as $c(v)$ under the decreased prices.

It follows that when offered bundles $\mathcal{B}$ at prices $(1-\alpha^{-1})p$, customers $\mathcal{C}$ generate overall revenue of at least $(1-\alpha^{-1})R'_* \geq (1/4)(1-\alpha^{-1})^2 R'$, which completes the proof. □

Finally, we briefly mention that the reduction described above has interesting applications in other varaiants of unit-demand pricing, as well. By UDP-PL we refer to the item pricing problem as defined in Section 1.1 with an additional *price ladder constraint* [1] $\pi$, i.e., a predefined relative order of item prices $p_{\pi(1)} \leq \cdots \leq p_{\pi(n)}$. It is known that uniform UDP-PL can be solved in polynomial time via a dynamic programming approach and, by our reduction, we obtain a $(1/4)(1-\alpha^{-1})^2$-approximation for general $\alpha$-coarse UDP-PL. This stands in sharp contrast to UdP without price ladder contraint, which does not allow for constant approximation gurantees even on coarse instances [4].

## 5 Approximability of UdBp-Max

In this section we turn to UdBp-Max where customers are strictly optimistic and assign to every bundle the maximum value of any item contained in it. We will see that this model is fundamentally different from UdBp-Min as it turns out to be equivalent to the pure item pricing problem.

Let $(\mathcal{I}, \mathcal{C})$ be an instance of UdBp-Max and let $(\mathcal{B}, p)$ be an optimal solution to that instance. Then we can transform $(\mathcal{B}, p)$ into a solution $(\mathcal{B}', p')$ yielding the same revenue where $\mathcal{B}'$ consists of singleton sets only. For this, we just need to replace any non-singleton bundle in $\mathcal{B}$ that is bought by a subset $\{c_1, \ldots, c_\ell\} \subseteq \mathcal{C}$ of customers by a set of bundles $\{i_1\}, \ldots, \{i_\ell\}$ where $i_j$ denotes the item from $\mathcal{B}$ that customer $c_j$ values the most. All these new bundles are offered for price $p(\mathcal{B})$. For a customer $j$ bundle $\{i_j\}$ has the same value as bundle $\mathcal{B}$ and for every other customer it has at most the same value. Hence, customer $j$ will buy $\{i_j\}$ and the other customers in $\mathcal{C} \setminus \{c_1, \ldots, c_\ell\}$ are not affected by replacing $\mathcal{B}$.

This implies that UdBp-Max and the pure item pricing problem are essentially the same problem. Hence, known results for the latter problem apply also to UdBp-Max. In particular, the revenue-maximizing item pricing can be approximated within factors that are logarithmic in the number of customers or linear in the number of distinct items [14] and, under appropriate complexity theoretic assumptions, no essential improvement is possible [4, 11].

## 6 Conclusions

We have introduced an extension of the unit-demand pricing problem in which bundles may be offered. This problem is interesting because it models the sales model of retailers like hotwire.com. We have seen that different assumptions about the customers' beliefs yield very different conclusions. While for the case of pessimistic customers we presented novel algorithmic results, the case of optimistic customers boils down to the pure item pricing problem.

There are many interesting questions open. One question arising directly from our results is whether there exists a constant factor approximation algorithm for general non-uniform instances of UdBp-Min without the additional assumption of $\alpha$-coarseness. It would also be very interesting to extend our model to different beliefs of customers. One could, e.g., study a model in which customers believe that an item is chosen uniformly at random from the set they buy.


## Acknowledgements

We thank Bobby Kleinberg for several insightful discussions.

## References


1. Gagan Aggarwal, Tomás Feder, Rajeev Motwani, and An Zhu. Algorithms for Multi-Product Pricing. In *Proc. of the 31st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 253–266, 2004.
2. Maria-Florina Balcan and Avrim Blum. Approximation Algorithms and Online Mechanisms for Item Pricing. In *Proc. of the 7th ACM Conference on Electronic Commerce (EC)*, 2006.
3. Maria-Florina Balcan, Avrim Blum, Jason D. Hartline, and Yishay Mansour. Mechanism Design via Machine Learning. In *Proc. of the 46th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 605–614, 2005.
4. Patrick Briest. Uniform Budgets and the Envy-Free Pricing Problem. In *Proc. of the 35th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 808–819, 2008.
5. Patrick Briest, Shuchi Chawla, Robert Kleinberg, and Matthew Weinberg. Pricing Randomized Allocations. In *Proc. of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010.
6. Patrick Briest and Piotr Krysta. Buying Cheap is Expensive: Hardness of Non-Parametric Multi-Product Pricing. In *Proc. of the 18th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007.
7. Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic Pricing via Virtual Valuations. In *Proc. of the 8th ACM Conference on Electronic Commerce (EC)*, pages 243–251, 2007.
8. Ning Chen, Arpita Ghosh, and Sergei Vassilvitskii. Optimal Envy-Free Pricing with Metric Substitutability. In *Proc. of the 9th ACM Conference on Electronic Commerce (EC)*, pages 60–69, 2008.
9. Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The Strong Perfect Graph Theorem. *Annals of Mathematics*, 164:51–229, 2006.
10. Pierluigi Crescenzi, Riccardo Silvestri, and Luca Trevisan. On Weighted vs Unweighted Versions of Combinatorial Optimization Problems. *Inf. Comput.*, 167(1):10–26, 2001.
11. Erik Demaine, Uriel Feige, MohammadTaghi HajiAghayi, and Mohammad Salavatipour. Combination Can Be Hard: Approximability of the Unique Coverage Problem. In *Proc. of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006.
12. Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael Saks, and Andrew Wright. Competitive Auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
13. Martin Grötschel, László Lovász, and Alexander Schrijver. The Ellipsoid Method and its Consequences in Combinatorial Optimization. *Combinatorica*, 1(2):169–197, 1981.
14. Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On Profit-Maximizing Envy-Free Pricing. In *Proc. of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1164–1173, 2005.
15. Roger Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
16. John Riley and Richard Zeckhauser. Optimal Selling Strategies: When to Haggle, When to Hold Firm. *Quarterly J. Economics*, 98(2):267–289, 1983.
17. John Thanassoulis. Haggling Over Substitutes. *J. Economic Theory*, 117:217–245, 2004.