

Pricing commodities, or How to sell when buyers have restricted valuations

Robert Krauthgamer¹, Aranyak Mehta², and Atri Rudra³

¹ Weizmann Institute, Rehovot, Israel and IBM Almaden, San Jose, CA.

`robert.krauthgamer@weizmann.ac.il`

² Google, Inc., Mountain View, CA. *

`aranyak@google.com`

³ University at Buffalo, The State University of New York, NY. *

`atri@cse.buffalo.edu`

Abstract. How should a seller price his goods in a market where each buyer prefers a single good among his desired goods, and will buy the cheapest such good, as long as it is within his budget? We provide efficient algorithms that compute near-optimal prices for this problem, focusing on a commodity market, where the range of buyer budgets is small. We also show that our technique (which is based on LP-rounding) easily extends to a different scenario, in which the buyers want to buy all the desired goods, as long as they are within budget.

1 Introduction

Pricing goods to maximize revenue is a critical yet difficult task in almost any market. We study the case of a monopolistic seller (only one seller in the market), a restricted scenario that is already quite challenging. One difficulty is to estimate the demand curves (amount of demand at different prices), but even complete knowledge of the demand curves is sufficient only in rather simple cases, e.g. if the monopolist sells only a single type of good, or if the various goods he sells cater to different sub-markets. In such cases, the revenue-maximizing prices can be determined for each good separately, directly from that good's demand curve.

But what if goods of different types are sold all in the same market? Now, *the seller's own goods could be competing against each other* for the attention of the same buyer. This is generally true of a seller who wants to tap into multiple market segments. For example, Dell sells many models of laptops with varying features catering to varying needs of its consumers, but then it must price the different models carefully so that they do not eat into each other's revenue. As an example on a smaller scale, consider the pricing of movie shows. Different shows are priced differently (for example, matinee vs. evening shows) to attract different audience sections. Again, pricing is critical— a very cheap matinee show might eat into the evening show revenue and decrease the overall revenue.

* Work done in part while at the IBM Almaden Research Center.

On the other hand, multiple goods might lead to higher prices by complementing each other. A very visible example is the marketing of Apple’s iPod and various accessories. The strategy there is not to sell the iPod in isolation but to offer various accessories. These accessories vary from items that are expensive (for example, a charger) to items that are inexpensive (for example, songs from iTunes). Pricing for revenue maximization becomes computationally complex precisely because of this interaction between different goods. Indeed, Aggarwal, Feder, Motwani and Zhu [1] and also Guruswami, Hartline, Karlin, Kempe, Kenyon and McSherry [2] studied the computational aspects of these pricing problems, showing that in various such settings, computing the optimal prices is NP-hard.

One setting, referred to as *unit-demand consumers* in [2], is where each buyer wants to buy one good out of his desired set, as follows: There are m buyers, each of whom has an arbitrary set of desirable goods and a spending budget. The (single) seller knows the buyers’ types (i.e. desired set and budget) and needs to set a price for each of the n goods. Once prices are set, every buyer buys the (single) cheapest good in his set, provided it is within his budget (breaking ties arbitrarily). Another setting, referred to as *single-minded consumers* in [2], differs from the first setting in that now each buyer only wants to buy the desired set as a bundle. That is, once prices are set, every buyer buys the entire set of his desired goods, provided its total cost is within his budget (if not, he buys nothing).

Throughout, we shall assume that the desired set of every buyer has size at most k . As we shall soon see, even the case of small k is nontrivial and interesting. In addition, we shall assume that the goods are available in *unlimited supply*, that is, the seller can sell any number of copies of the item without paying any marginal cost of production.

Several results are known about computing prices that maximize revenue in these two settings. In [1], it is shown that the problem of maximizing revenue in the unit-demand case is not only NP-hard, but APX-hard.⁴ For this problem, they also give an $O(\log m)$ approximation algorithm (which uses the best single price). In [2], similar results are shown independently, and it is shown in addition that for the single-minded setting, maximizing revenue is APX-hard and that there is $\log(nm)$ approximation algorithm (which again uses the best single price). Demaine, Feige, Hajiaghayi, and Salavatipour [3] show that the above results are more or less optimal in the general single minded bidder problem – under some complexity assumptions, there is a fixed $\delta > 0$ such that the problem cannot be approximated to within a factor of $\log^\delta n$. Balcan and Blum [4] present a 4-approximation algorithm for single-minded bidders and $k = 2$. Their algorithm extends to larger k , with $O(k)$ approximation. As was observed in [5], their arguments apply to the unit-demand case as well.

These results depict a rather grim landscape (at least computationally) for the problem of pricing to maximize revenue. However, many real-life instances

⁴ An optimization problem is APX-hard if there exists a constant $\rho > 1$ such that it is NP-hard to approximate the optimum within factor ρ .

are more specialized, and thus, a more practice-oriented approach is to identify restrictions, under which one can beat the aforementioned $O(\log n)$ factor, or better yet, obtain a very small constant-factor approximation. In particular, every percent of improvement counts in practice, requiring us to improve one (small) constant approximation factor to another.

We thus pay special attention to *commoditized markets*,⁵ where the range of buyers budgets is restricted to a “small” set \mathcal{B} . In one such restriction, $\mathcal{B} = \{1, C\}$ is a doubleton, representing a bimodal market in which buyers are divided into poor and rich. For example, buyers coming from different referring websites such as `lastminutedeals.com` and `hotels.com` might have significantly different budgets for booking a hotel room. As another example, a tourist might be willing to pay for a Broadway show a significantly higher amount than a local. Another motivation for studying such markets could be the low descriptive complexity for the different buyers’ budget types, or equivalently a low communication complexity to identify a buyer’s budget. In yet another restriction, $\mathcal{B} = [1, C]$ is a small interval, representing a market with little variation, say within 50%, in the valuation of different buyers, and clearly there are numerous examples for such markets. Note that in both cases, the buyers can be completely idiosyncratic regarding their desired goods, as only the budgets are restricted.

1.1 Results and Techniques

Unit-demand setting. In Section 2, we consider inputs with $\mathcal{B} = \{1, C\}$ (i.e., bimodal markets) and $k = 2$ (i.e. a desired set is a pair of goods). On the one hand, the APX-hardness results [1, 2] mentioned above are actually shown for such restricted instances (in fact, for $C = 2$). On the other hand, obtaining $2 - \frac{1}{C}$ approximation is rather easy — simply choose the best single price (same for all goods) among $\{1, C\}$ — obviously a very naive solution, but already better than the (more general) 4 approximation that can be derived from [4]. The challenge in this regime is to improve the approximation below 2, and indeed we present an algorithm achieving factor $\frac{3}{2} - \frac{1}{2C}$. Observe that even when C is not too large, this is a significant improvement (e.g. for $C = 2$, from 1.5 to 1.25). This approach easily extends to larger k , in which case the approximation we achieve is $2 - \frac{1}{k} - \frac{k-1}{kC}$.

Our algorithms are based on randomized rounding of a linear programming (LP) relaxation, a powerful paradigm that is often useful for discrete optimization (for example, see the survey of Srinivasan [6]). We “round” the prices suggested by the LP to prices in the discrete (“integral”) set $\{1, C\}$. The rationale behind this rounding is that an optimal pricing may always choose prices from the set $\{1, C\}$. However, it is interesting to note that the pricing problem does not require the prices to be “discrete”, and thus, the real reason behind our rounding procedure is the following: In contrast with a “standard” randomized

⁵ A commoditized (also commodified) market is one characterized by price-competition with little or no differentiation by brand.

rounding algorithm, where the probability (with which we round a variable upwards) depends linearly on the corresponding LP variable, we use a probability that is polynomial in the LP variable. The only other non-linear randomized LP rounding algorithms that we are aware of are the approximation algorithm of Goemans and Williamson [7] for MAX SAT, and that for finding the densest k -subgraph problem that is attributed to Goemans [8]. The crux is that at every basic feasible solution of our LP relaxation, all the prices are half-integral [9, Chap. 14] (modulo a normalization factor), and this fact greatly simplifies the choice of the polynomial—in fact, our rounding procedure raises the variables to a power. Interestingly, the value of the power is a function of C .

We further show that our algorithm can be derandomized, and that its approximation matches the LP’s integrality gap, and thus it is optimal with respect to this LP. In addition, we observe that in the case where budgets come from an interval $\mathcal{B} = [1, C]$, a simple algorithm achieves $1 + \ln C$ approximation by computing the best single price, and that this factor matches the integrality gap of an LP relaxation that extends the LP mentioned above for the case $\{1, C\}$.

Single-Minded setting. Recently, Khandekar, Könemann and Markakis (Private Communication) have studied the case of single-minded bidders with desired sets of size at most 2, and the same budget for all the buyers, and gave a $4/3$ approximation algorithm. Subsequently (but using independently derived techniques), we found out that our LP rounding approach mentioned above is easy to adapt to this setting as well, achieving $\frac{6+\sqrt{2}}{5+\sqrt{2}} \approx 1.15$ approximation. In Section 3 we briefly present this algorithm, and show a matching integrality gap. Again, this problem is known to be APX-hard because the results of [2] are actually shown for such restricted instances. Further, our algorithm obtains much better approximation than a $3/2$ approximation achievable by choosing the best single price in the set $\{1/2, 1\}$, which was already better than the (more general) 4 approximation of [4].

Online pricing. Finally, we consider in Section 4 inputs with $k = 2$ (and no restriction on the budget). Using a variation of the algorithm designed by Balcan and Blum [4], we design an algorithm that works even in an *online setting*, where goods arrive sequentially (together with the bids of all the buyers interested in that good), and the seller has to determine the price of a good immediately as it arrives. This model may correspond for instance to Comcast cable TV selling video on demand, where new offerings are announced (with prices) on a regular basis. Our algorithm achieves 4 approximation, compared to the best (offline) prices. We note that [4] also give an online pricing algorithm, but in their setting buyers arrive online, and the prices (of a fixed set of goods) need to be updated.

Truthful Mechanisms We assume throughout the paper that the seller knows the budget of each bidder. We may also be interested in settings where the seller does not know such information about the market. Balcan, Blum, Hartline and Mansour [10] show that every approximation algorithm for revenue maximization can be converted into a truth-revealing mechanism. They design a general

technique that loses only an additional factor of $1 + \epsilon$ in the approximation, if certain technical conditions (like sufficiently many bidders) are satisfied. Similarly to [4], we note that this technique is applicable in our setting, and thus converts our algorithms to truthful mechanisms, provided that the number of bidders is at least (roughly) Cn/ϵ^2 .

1.2 Related Work

The notion of revenue-maximizing pricing of goods in unlimited supply was introduced by Goldberg, Hartline, Karlin, Saks and Wright [11]. In their setting, the goods were “independent” and hence the optimization problem was trivial, and they focused on designing *truthful* mechanisms to maximize revenue. There have been numerous followup work, and we only mention here results that are directly related to our work.

Guruswami, Hartline, Karlin, Kempe, Kenyon and McSherry [2] considered the problem of revenue maximization in a variety of settings, including both unit-demand and single-minded bidders, and also envy-free pricing of goods in limited supply. As mentioned earlier, they showed logarithmic upper bounds and APX-hardness for both types of bidders. The results for the unit-demand case were also obtained independently by Aggarwal, Feder, Motwani and Zhu [1]. For single-minded bidders, a polylogarithmic hardness result, which complements the result above, was obtained by Demaine, Feige, Hajiaghayi, and Salavatipour [3]. The problem of the single-minded bidder case, where the size of the demand sets was upper bounded by k , was considered by Briest and Krysta [12] who gave an $O(k^2)$ approximation for the problem, and was improved by Blum and Balcan [4] to $O(k)$. For the special case of $k = 2$, they obtain a 4 approximation algorithm.

Another paper that is less directly related but was also a starting point for our work is the work of Bansal, Cheng, Cherniavsky, Rudra, Scheiber, Sviridenko [13], which studies a problem of pricing over time, that was proposed in [2]. A special case of their problem gives another interpretation for the unit-demand setting: The seller is selling just one type of good (in unlimited supply), and does so over a period of n days, and can set a different price on each day. Each of the m buyers has a subset of size k of the n days, which represent the days on which she can purchase the item, and will choose to buy a copy of the good at the cheapest price she sees over the k days. The seller’s aim is to maximize revenue.

1.3 Problem Definitions

Our pricing problems involve one seller and m buyers. The seller has a collection V of n goods (also called items). Each $j \in V$ is a digital good, i.e., the seller has 0 marginal cost of production, or equivalently, the number of copies of j is at least the number of buyers m . Once the seller sets the prices of the goods, each buyer will buy a collection of goods, based on his own utility function. The seller’s problem is to determine a price p_j of each good $j \in V$ so as to maximize

revenue. Depending on the utility functions of the buyers, we have the following variations of the pricing problem. The first variation is our main focus, but we will also show how the techniques we develop also work for the second variation.

1. *Unit-demand bidders*: We let $UD_k(\mathcal{B})$ denote the problem of item pricing for unit-demand bidders with sets of size at most k , and bids from the set \mathcal{B} , as follows. Buyer i has a budget of $u_i \in \mathcal{B}$ and a subset S_i of desirable goods, with $|S_i| \leq k$. He is interested in buying *exactly one* good from S_i , and given prices on the goods, he will buy the cheapest good in S_i , provided that its price is at most u_i . For a price vector $\mathbf{p} = (p_1, \dots, p_m)$, let $\pi_i(\mathbf{p})$ be the revenue that the seller obtains from buyer i if the prices are set to \mathbf{p} . Thus

$$\pi_i(\mathbf{p}) = \begin{cases} \min\{p_j : j \in S_i\} & \text{if } \min\{p_j : j \in S_i\} \leq u_i \\ 0 & \text{otherwise} \end{cases}$$

Thus the seller's problem is: Find \mathbf{p} so as to maximize $\sum_{i=1}^n \pi_i(\mathbf{p})$. We are interested in the following special cases of this problem (defined by different values of k and \mathcal{B}): (1) $UD_k([1, C])$ for $C > 1$ and (2) $UD_k(\{1, C\})$ for $C > 1$.

2. *Single-minded bidders*: We let $SM_k(\mathcal{B})$ denote the problem of item pricing for single minded bidders, who have sets of size k and bids from the set \mathcal{B} , as follows. Buyer i has a budget of $u_i \in \mathcal{B}$ and a subset S_i of desirable goods with $|S_i| \leq k$. He is interested in buying *all* the goods in the set S_i . For a price vector \mathbf{p} , let π_i be the revenue that the seller obtains from buyer i , if the prices are set to \mathbf{p} . Thus

$$\pi_i(\mathbf{p}) = \begin{cases} \sum_{j \in S_i} p_j & \text{if } \sum_{j \in S_i} p_j \leq u_i \\ 0 & \text{otherwise} \end{cases}$$

Again, the seller's problem is: Find \mathbf{p} so as to maximize $\sum_{i=1}^n \pi_i(\mathbf{p})$. We will show how our techniques for $UD_2(\{1, C\})$ extend to give an algorithm for the case $SM_2(\{1\})$.

The case of $k = 2$: Pricing on a graph

Following [4], for $k = 2$, $UD_2(\mathcal{B})$ becomes a problem of pricing the vertices of a graph, with the buyers' desired sets corresponding to the edges of the graph. This will be our main focus in our techniques and analysis. We study two settings of budget ranges: $\mathcal{B} = \{1, C\}$ and $\mathcal{B} = [1, C]$, for $C > 1$.

Given a graph $G = (V, E)$ (possibly with self loops and multiple edges), along with edge weights $c_{ij} \in \mathcal{B}$ for every edge $(i, j) \in E$, the goal is to set prices p_i on every vertex i so as to maximize the total revenue. The revenue from an edge $(i, j) \in E$ becomes:

$$\pi_{ij} = \begin{cases} \min(p_i, p_j) & \text{if } \min(p_i, p_j) \leq c_{ij} \\ 0 & \text{otherwise.} \end{cases}$$

The case of $SM_2(\mathcal{B})$, studied in [4], is defined as a pricing problem on a graph analogously.

$\max \sum_{(i,j) \in E} \pi_{ij}$	subject to:
$\forall (i, j) \in E, c_{ij} = C$	$\pi_{ij} \leq 1 + p_i$ (1)
$\forall (i, j) \in E, c_{ij} = C$	$\pi_{ij} \leq 1 + p_j$ (2)
$\forall (i, j) \in E, c_{ij} = 1$	$\pi_{ij} \leq 1$ (3)
$\forall (i, j) \in E, i \neq j, c_{ij} = 1$	$\pi_{ij} \leq 2 - \frac{p_i + p_j}{C - 1}$ (4)
$\forall (i, i) \in E, c_{ii} = 1$	$\pi_{ii} \leq 1 - \frac{p_i}{C - 1}$ (5)
$\forall i \in V$	$0 \leq p_i \leq C - 1$ (6)
$\forall (i, j) \in E$	$\pi_{ij} \geq 0$ (7)

Fig. 1. LP relaxation for the unit-demand setting

2 Unit-Demand Buyers in Commoditized Markets

In this section, we look at pricing for unit demand bidders with restricted valuations. We start with valuations restricted to the set $\{1, C\}$ for some $C > 1$. In other words, we are interested in pricing schemes for the $UD_k(\{1, C\})$ model. Our main result is a pricing scheme that generates a revenue within a factor $\frac{(2k-1)C-k+1}{kC}$ of the optimal revenue (Theorem 2). For ease of exposition, we present the proofs for the $k = 2$ case, and mention the result for the general case in Section 2.5.

Our pricing scheme rounds an LP relaxation for the problem. Theorem 1 shows that our rounding algorithm (for the case $k = 2$) has an approximation factor of $\frac{3C-1}{2C}$. Our rounding procedure is tight (optimal) as we show in Section 2.4 that the integrality gap of our LP relaxation is at least $\frac{3C-1}{2C}$.

It is not difficult to verify that the LP in Figure 1 is a relaxation for our pricing problem $UD_2(\{1, C\})$.

2.1 On the optimal LP solutions

We first observe that an optimal basic feasible solution to the LP relaxation is half integral, in the sense that all the p_i variables are from the set $\{0, \frac{C-1}{2}, C-1\}$. Note that the actual “price” set for vertex i is $p_i + 1$.

Proposition 1. *Every optimal basic feasible solution to the LP in Figure 1 is half integral. That is, every extremal optimal assignment to the variables $\{p_i^*\}_{i \in V}$ satisfies the following: $p_i^* \in \{0, \frac{C-1}{2}, C-1\}$.*

Proof. Let $\{p_i^*\}_{i \in V}$ be the assignments to the $p_{(\cdot)}$ variables in some optimal assignment such that the values are not half integral. We will show that such an assignment is not an extremal optimal solution. In particular, we will exhibit two optimal assignments $p_{(\cdot)}^-$ and $p_{(\cdot)}^+$ such that for all $i \in V$, $p_i^* = \frac{1}{2}(p_i^- + p_i^+)$.

Define the following two subsets of vertices: $V^+ = \{i \mid \frac{C-1}{2} < p_i^* < C-1\}$ and $V^- = \{i \mid 0 < p_i^* < \frac{C-1}{2}\}$. Note that by the assumption on p^* , $V^- \cup V^+ \neq \emptyset$. Let $\epsilon > 0$ be a small enough number (to be defined later). We define the two related “price” assignments.

$$p_i^+ = \begin{cases} p_i^* + \epsilon & \text{if } i \in V^+ \\ p_i^* - \epsilon & \text{if } i \in V^- \\ p_i^* & \text{otherwise} \end{cases} \quad p_i^- = \begin{cases} p_i^* - \epsilon & \text{if } i \in V^+ \\ p_i^* + \epsilon & \text{if } i \in V^- \\ p_i^* & \text{otherwise} \end{cases}$$

Obviously, for all $i \in V$, $p_i^* = \frac{1}{2}(p_i^+ + p_i^-)$. To complete the proof, we will show that both $p_{(\cdot)}^+$ and $p_{(\cdot)}^-$ are optimal assignments. Let the “revenue” variables corresponding to $p_{(\cdot)}^*$, $p_{(\cdot)}^+$ and $p_{(\cdot)}^-$ be denoted by $\pi_{(\cdot)}^*$, $\pi_{(\cdot)}^+$ and $\pi_{(\cdot)}^-$. For the sake of contradiction assume w.l.o.g. that $\sum_{(i,j) \in E} \pi_{ij}^- < \sum_{(i,j) \in E} \pi_{ij}^*$. We aim to show that $\sum_{(i,j) \in E} \pi_{ij}^+ > \sum_{(i,j) \in E} \pi_{ij}^*$, which will contradict the optimality of $p_{(\cdot)}^*$.

We first set $\epsilon = \frac{1}{4} \min\{\epsilon_1, \epsilon_2, \epsilon_3\}$ where:

$$\begin{aligned} \epsilon_1 &= \min\{|p_i^* - p_j^*| : p_i^* \neq p_j^*, (i, j) \in E \text{ and } c_{ij} = C\}, \\ \epsilon_2 &= \min\{C - 1 - p_i^* : (i, j) \in E \text{ and } c_{ij} = C\}, \\ \epsilon_3 &= \min\left\{\left|1 - \frac{p_i^* + p_j^*}{C-1}\right| : p_i^* + p_j^* \neq C-1, (i, j) \in E, \right. \\ &\quad \left. \text{and } c_{ij} = 1\right\}. \end{aligned}$$

Note that for $(i, j) \in E$ such that $i, j \notin V^+ \cup V^-$, $\pi_{ij}^* = \pi_{ij}^+ = \pi_{ij}^-$. Let us now consider an edge $(i, j) \in E$ with at least one end point in $V^+ \cup V^-$. To finish the proof, we will show that

$$\pi_{ij}^+ - \pi_{ij}^* = -(\pi_{ij}^- - \pi_{ij}^*). \quad (8)$$

First assume that $i \neq j$ and $c_{ij} = C$. In this case, $\pi_{ij}^* = 1 + \min(p_i^*, p_j^*)$, $\pi_{ij}^+ = 1 + \min(p_i^+, p_j^+)$ and $\pi_{ij}^- = 1 + \min(p_i^-, p_j^-)$. If $p_i^* = p_j^*$, then by the definitions of ϵ , $p_{(\cdot)}^+$ and $p_{(\cdot)}^-$, (8) is satisfied. If $p_i^* \neq p_j^*$, then by the definition of ϵ if p_i^* (p_j^*) is the minimum price for (i, j) , then so are p_i^+ (p_j^+) and p_i^- (p_j^-). Again by the definitions of $p_{(\cdot)}^+$ and $p_{(\cdot)}^-$, (8) is satisfied.

Now let us consider the case when $i \neq j$ and $c_{ij} = 1$. We now consider two subcases. First if $p_i^* + p_j^* \neq C-1$, then by the definitions of ϵ , $p_{(\cdot)}^+$ and $p_{(\cdot)}^-$, (8) is satisfied. Now if $p_i^* + p_j^* = C-1$ then either $p_i^* \in V^+$ and $p_j^* \in V^-$ or $p_i^* \in V^-$ and $p_j^* \in V^+$. In all these cases, $\pi_{ij}^* = \pi_{ij}^+ = \pi_{ij}^-$, which in particular implies that (8) is satisfied.

Similarly, one can show that for self loops, (8) is also satisfied.

2.2 A rounding algorithm

Consider the following randomized algorithm, where $\tau > 0$ is a parameter (to be chosen later).

Algorithm **Algo**(τ):

1. Solve the LP in Figure 1 and obtain an optimal basic feasible solution with prices variables $\{p_i\}_{i \in V}$.
2. For every $i \in V$, independently assign a price of C with probability $\left(\frac{p_i}{C-1}\right)^\tau$ and a price of 1 with probability $1 - \left(\frac{p_i}{C-1}\right)^\tau$.

We now analyze the performance of the rounding algorithm above.

Theorem 1. *For every $C > 1$, there is $\tau > 0$ such that **Algo**(τ) is a $(3C - 1)/(2C)$ approximation for the pricing problem with unit-demand bidders, $k = 2$, and budgets from $\mathcal{B} = \{1, C\}$. That is, the expected revenue of **Algo**(τ) is at least $\frac{2C}{3C-1}$ fraction of the optimum for $UD_2(\{1, C\})$.*

Proof. Set $\tau = \frac{1}{2} \log\left(\frac{3C-1}{C-1}\right)$. For notational convenience, we will denote **Algo**(τ) by **Algo**. Let the optimal (extremal) solution of the LP assign prices p_i^* to every vertex i and obtains a revenue of π_{ij}^* from every edge (i, j) . We will show that for every edge $(i, j) \in E$, the expected revenue of **Algo** from that edge is at least $\frac{2C}{3C-1} \cdot \pi_{ij}^*$. The result follows from the linearity of expectation. For the rest of the proof, it will be convenient to define, for every $i \in V$, $q_i = \frac{p_i^*}{C-1}$. By Proposition 1, we have $q_i \in \{0, \frac{1}{2}, 1\}$.

Let us first consider the case when $i \neq j$. We have two subcases.

Case 1a: $c_{ij} = 1$. In this case $\pi_{ij}^* \leq \min(1, 2 - q_i - q_j)$, while **Algo** obtains an expected revenue of $0 \cdot (q_i^\tau q_j^\tau) + 1 \cdot (1 - q_i^\tau q_j^\tau) = 1 - (q_i q_j)^\tau$. When $q_i = q_j = 1$ then both the LP and **Algo** obtain a revenue of 0. When $q_i + q_j = \frac{3}{2}$ then the ratio of the revenue obtained by **Algo** to π_{ij}^* (which is $1/2$) is $2(1 - \frac{1}{2}^\tau) > 1 - \frac{1}{2^{2\tau}}$.⁶ Finally, when $q_i + q_j \leq 1$, then $\pi_{ij}^* = 1$, while **Algo** obtains the least revenue when $q_i = q_j = \frac{1}{2}$, which implies a ratio of at least $1 - \frac{1}{2^{2\tau}} = \frac{2C}{3C-1}$ in all the possibilities.

Case 1b: $c_{ij} = C$. In this case $\pi_{ij}^* \leq 1 + (C - 1) \min(q_i, q_j)$, while **Algo** obtains an expected revenue of $C \cdot (q_i^\tau q_j^\tau) + 1 \cdot (1 - q_i^\tau q_j^\tau) = 1 + (C - 1)(q_i q_j)^\tau$. W.l.o.g. assume that $q_j \geq q_i$. Thus, the ratio of the revenue obtained by **Algo** and π_{ij}^* is at least:

$$\begin{aligned} \min_{q_i, q_j \in \{0, \frac{1}{2}, 1\}, q_j \geq q_i} \frac{1 + (C - 1)(q_i q_j)^\tau}{1 + (C - 1) \min(q_i, q_j)} &\geq \min_{q_i \in \{0, \frac{1}{2}, 1\}} \frac{1 + (C - 1)q_i^{2\tau}}{1 + (C - 1)q_i} \\ &= \frac{1 + \frac{C-1}{2^{2\tau}}}{1 + \frac{C-1}{2}} = \frac{2C}{3C-1}. \end{aligned} \quad (9)$$

⁶ To see why this is true set $a = 2^{-\tau}$ and note that we have to show that $2 - 2a > 1 - a^2$, which is true for $a > 1$. The latter inequality is true as $\tau > 0$.

We now consider the case $i = j$. Again we have two sub cases.

Case 2a: $c_{ii} = 1$. In this case $\pi_{ii}^* \leq \min(1, 1 - q_i) = 1 - q_i$, while **Algo** gets a revenue of $0 \cdot q_i^\tau + 1 \cdot (1 - q_i^\tau) = 1 - q_i^\tau$. Thus, the ratio of the revenue of **Algo** to π_{ii}^* is at least

$$\min_{q_i \in \{0, \frac{1}{2}, 1\}} \frac{1 - q_i^\tau}{1 - q_i} = \min(1, 2 - 2^{1-\tau}) \geq 1 - \frac{1}{2^{2\tau}} = \frac{2C}{3C - 1}.$$

Case 2b: $c_{ii} = C$. In this case $\pi_{ii} \leq 1 + (C - 1)q_i$. The expected revenue for **Algo** is $1 \cdot (1 - q_i^\tau) + C \cdot q_i^\tau = 1 + (C - 1)q_i^\tau \geq 1 + (C - 1)q_i^{2\tau}$. Thus, from (9), the ratio is at least $\frac{2C}{3C - 1}$.

Thus, in all cases for every edge $(i, j) \in E$, **Algo** obtains an expected revenue of at least $\frac{2C}{3C - 1} \cdot \pi_{ij}^*$, as desired.

2.3 Derandomization

Algorithm **Algo**(τ) can be derandomized in a straightforward way using standard techniques. In particular, observe that the analysis of the randomized rounding step only required pairwise independence among the random choices. One can use a small family of pairwise independent random variables (see the survey [14] for such constructions) and exhaustively try all the possibilities in this space.

Alternatively, one can employ the method of conditional expectation [15, 6], since the expected revenue after randomized rounding is an easy formula to calculate (given the probabilities).

2.4 A tight integrality gap

Next, we show that Theorem 1 is the best one can hope from any algorithm that rounds the LP. Formally, we prove the following.

Proposition 2. *There exist an instance of $UD_2(\{1, C\})$ for which the the integrality gap of the LP in Figure 1 is at least $\frac{3C - 1}{2C}$.*

Proof. Consider the graph with two vertices and C parallel edges— one of which has a cost of C and the rest have a cost of 1. (This assumes that C is integral; if however C is not integral, we need to choose an appropriate number of cost 1 edges and cost C edges such that their ratio is C .) The optimal revenue is C . However, the LP can set a price of $\frac{C+1}{2}$ on both the vertices to get a revenue of $\frac{C+1}{2}$ from the cost C edge and a revenue of 1 from each of the cost 1 edges. Thus, the integrality gap is at least

$$\frac{1 \cdot (\frac{C+1}{2}) + (C - 1) \cdot 1}{C} = \frac{3C - 1}{2C}.$$

2.5 The general case

The results presented for $k = 2$ in the previous sections can be suitably modified to work for the general case. The LP relaxation for general k is the natural one. For example, the constraint (4), the sum $p_i + p_j$ will be replaced by $\sum_{j=1}^k p_{i_j}$ for the hyperedge $(p_{i_1}, p_{i_2}, \dots, p_{i_k})$. The “half-integrality” gap result will now say that the prices are in the set $\{0, (1 - 1/k)(C - 1), C - 1\}$. Finally, we can prove the following counterparts of Theorem 1 and Proposition 2 by straightforward generalizations of their proofs

Theorem 2. *For every $C > 1$, there an algorithm that is a $\frac{(2k-1)C-k+1}{kC}$ approximation for the pricing problem with unit-demand bidders with demand size at most k and budgets from $\mathcal{B} = \{1, C\}$.*

Proposition 3. *There exist an instance of $UD_k(\{1, C\})$ for which the the integrality gap of the LP used above is at least $\frac{(2k-1)C-k+1}{kC}$.*

2.6 Budget Range $[1, C]$

Another interesting restriction on the range of buyer’s budgets is to an interval $\mathcal{B} = [1, C]$, which clearly generalizes the previous doubleton case $\{1, C\}$. For this case, denoted $UD_2([1, C])$, we obtain the following approximation. The proof proceeds by considering the best single price (same price for all the goods) in the range $[1, C]$, and is deferred to the full version.

Proposition 4. *For every $C > 1$, there is a polynomial time $1 + \ln C$ approximation algorithm for the unit-demand pricing problem with $k = 2$ and budgets from $\mathcal{B} = [1, C]$.*

One can try a natural extension of our LP-relaxation technique for $\{1, C\}$ to this more general case $[1, C]$. However, it turns out that the resulting LP has integrality gap $1 + \ln C$, and thus cannot offer improved approximation.

3 Single-Minded Buyers in Commoditized Markets

We now consider the pricing problem for single minded bidders when all the bidders have the same budget, which can be assumed w.l.o.g. to be 1. That is, we are interested in pricing schemes for the $SM_2(\{1\})$ model. We extend our techniques from Section 2 to get a pricing algorithm with an approximation factor of $\frac{6+\sqrt{2}}{5+\sqrt{2}} \approx 1.156$ (Theorem 3). As in the case of single-minded bidders, our rounding procedure is tight (optimal), as we show that this LP relaxation has a matching integrality gap.

It is not difficult to verify the LP in Figure 2 is a relaxation for our problem $SM_2(\{1\})$.

As in the $UD_2(\{1, C\})$, we first observe that an optimal basic feasible solution to the LP relaxation is half integral.

$\max \sum_{(i,j) \in E} \pi_{ij}$	subject to:	
$\forall (i,i) \in E$	$\pi_{ii} \leq p_i$	(10)
$\forall (i,j) \in E, i \neq j$	$\pi_{ij} \leq p_i + p_j$	(11)
$\forall (i,j) \in E, i \neq j$	$\pi_{ij} \leq 2 - p_i - p_j$	(12)
$\forall i \in V$	$0 \leq p_i \leq 1$	(13)
$\forall (i,j) \in E$	$\pi_{ij} \geq 0$	(14)

Fig. 2. LP relaxation for the single-minded bidders setting

Proposition 5. *Every optimal basic feasible solution to the LP in Figure 2 is half integral. That is, every extremal optimal assignment to the variables $\{p_i^*\}_{i \in V}$ satisfies the following: $p_i^* \in \{0, \frac{1}{2}, 1\}$.*

The proof is very similar to that of Proposition 1 and is omitted (the only possible “tight” edge can be for an edge $(i, j) \in E$, such that $i \neq j$ and $p_i^* + p_j^* = 1$).

We next analyze the following randomized algorithm.

- Algorithm **Algo_{SM}**:
1. Solve the LP in Figure 2 to obtain an optimal basic feasible solution with price variables $\{p_i\}_{i \in V}$.
 2. Fix prices according to the three schemes below and pick the one that generates the maximum revenue.
 - (a) Assign a price p_i to vertex i .
 - (b) If $p_i \neq 1$, assign a price of p_i to vertex i , else assign a price of $1/2$.
 - (c) If $p_i \neq 1/2$, assign a price of p_i to vertex i , else assign a price of 0 with probability $1/\sqrt{2}$ and a price of 1 with probability $1 - 1/\sqrt{2}$.

Theorem 3. *Algo_{SM} achieves $\frac{6+\sqrt{2}}{5+\sqrt{2}}$ approximation for the pricing problem with single-minded bidders, desired sets of size at most 2, and unit budgets. That is, expected revenue of Algo_{SM} is at least $\frac{5+\sqrt{2}}{6+\sqrt{2}}$ fraction of the optimum for $SM_2(\{1\})$.*

The proof is deferred to the full version. The rounding procedure above is tight (optimal), as the following proposition shows (proof in the full version).

Proposition 6. *There exist an instance of $SM_2(\{1\})$ for which the the integrality gap of the LP in Figure 2 is at least $\frac{6+\sqrt{2}}{5+\sqrt{2}}$.*

4 An online 4-approximation

In this section, we consider the following online version of the $UD_2(\cdot)$ and $SM_2(\cdot)$ problems. Buyers are assumed to be “in the system” at the beginning and the goods arrive in an online fashion. When a good arrives, any buyer who is interested submits a bid and the seller has to price this good before the next good arrives. We assume that every buyer is interested in at most two items and the seller knows the identity of each buyer. Further, the seller knows about the exact set of elements the buyer is interested in only after the buyer had bid for *both* the items he is interested in. The price that a buyer pays follows the same rules as in $UD_2(\cdot)$ and $SM_2(\cdot)$ models respectively. In the graph abstraction of the $UD_2(\cdot)$ and $SM_2(\cdot)$, the online model has the following interpretation. At every step, a vertex in the underlying graph arrives. Once a vertex appears, all the edges incident on it (along with the edge weights) are revealed to the seller. The only way the seller knows about the other end point of an edge is if that vertex had arrived earlier. Under these constraints, the seller has to price every vertex as it arrives, so as to make as much revenue as possible. For the rest of the section, we will only talk about the $UD_2(\cdot)$ model. The discussion holds equally well for the $SM_2(\cdot)$ model (just replace the prices of ∞ by 0).

The algorithm in [4] works in this model *if* when a vertex arrives, the seller has the full information about the edge. That is, if the other end point is in the “future” then the seller also knows about this other end point. We now restate the algorithm in [4] that works in this scenario. Initially with probability 1/2 decide on “left” or “right”. For the ease of exposition, assume that the algorithm chose left. When a vertex (say i) arrives, with probability 1/2 tag it as a left vertex or a right vertex (unless it is already assigned a tag). If i is a right vertex then assign it a price ∞ . Otherwise look at the set of neighbors of i (recall that the seller knows everything about the edge incident on i). If some neighbor j has not arrived yet then assign j one of the tags with equal probability. Let $N'(i)$ denote the set of neighbors of i that are tagged right. Now consider all edges between i and $N'(i)$ and set the price of i to be the best fixed price given that the vertices in $N'(i)$ are priced at infinity. By the analysis in [4], this algorithm is 4-competitive.

We now consider the more general model, where the seller has no information about the vertices that are yet to arrive. For this model, we consider the following refinement of the algorithm in [4]. For any vertex i , let p_i^* denote the best fixed price for vertex i , given that all of its neighbors are priced at ∞ . Recall that in our online model, once a vertex arrives, the seller knows the weights of all the incident edges. Thus, the seller can calculate the price p_i^* . Given this, the online algorithm is very simple.

Algorithm: When each vertex i arrives,

- Compute its best fixed price p_i^* .
- With probability 1/2 set its price $p_i = \infty$ and with probability 1/2 set $p_i = p_i^*$.

We have the following performance guarantee (proof in the full version).

Theorem 4. *For the online $UD_2(\cdot)$ model, the algorithm above is 4-competitive in the expected sense.*

References

1. Aggarwal, G., Feder, T., Motwani, R., Zhu, A.: Algorithms for multi-product pricing. In: 31st International Colloquium on Automata, Languages and Programming (ICALP). Volume 3142 of Lecture Notes in Computer Science., Springer (2004) 72–83
2. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proceedings of the 16th annual ACM-SIAM symposium on Discrete algorithms, Society for Industrial and Applied Mathematics (2005) 1164–1173
3. Demaine, E.D., Feige, U., Hajiaghayi, M.T., Salavatipour, M.R.: Combination can be hard: approximability of the unique coverage problem. In: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA). (2006) 162–171
4. Balcan, M.F., Blum, A.: Approximation algorithms and online mechanisms for item pricing. In: Proceedings of the 7th ACM conference on Electronic commerce, ACM Press (2006) 29–35
5. Briest, P., Krysta, P.: Buying cheap is expensive: Hardness of non-parametric multi-product pricing. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA). (2007)
6. Srinivasan, A.: Approximation algorithms via randomized rounding: a survey. In Karonski, M., Promel, H.J., eds.: Lectures on Approximation and Randomized Algorithms. Series in Advanced Topics in Mathematics. Polish Scientific Publishers PWN (1999) 9–71
7. Goemans, M.X., Williamson, D.P.: New $3/4$ -approximation algorithms for the maximum satisfiability problem. *SIAM Journal on Discrete Mathematics* **7** (1994) 656–666
8. Goemans, M.X.: Mathematical programming and approximation algorithms. Lecture at Udine School, Udine, Italy. (1996)
9. Vazirani, V.V.: Approximation Algorithms. Springer Verlag, New York, NY (2001)
10. Balcan, M.F., Blum, A., Hartline, J.D., Mansour, Y.: Mechanism design via machine learning. In: 46th Annual IEEE Symposium on Foundations of Computer Science, IEEE Computer Society (2005) 605–614
11. Goldberg, A., Hartline, J., Karlin, A., Saks, M., Wright, A.: Competitive Auctions. *Games and Economic Behavior* **55**(2) (2006) 242–269
12. Briest, P., Krysta, P.: Single-minded unlimited supply pricing on sparse instances. In: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA). (2006) 1093–1102
13. Bansal, N., Cheng, N., Cherniavsky, N., Rudra, A., Scheiber, B., Sviridenko, M.: Pricing to impatient bidders. In: Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA). (2007) To Appear.
14. Luby, M., Wigderson, A.: Pairwise independence and derandomization. *Found. Trends Theor. Comput. Sci.* **1**(4) (2006) 237–301
15. Alon, N., Spencer, J.H.: The probabilistic method. John Wiley & Sons, Inc., New York (1992)