Buying Cheap is Expensive: Hardness of Non-Parametric Multi-Product Pricing

Patrick Briest* Piotr Krysta†

Abstract

We investigate non-parametric unit-demand pricing problems, in which we want to find revenue maximizing prices for products $\mathcal{P}$ based on a set of consumer profiles $\mathcal{C}$. A consumer profile consists of a number of non-zero budgets for different products and possibly an additional product ranking. Once prices are fixed, each consumer chooses to buy one of the products she can afford based on some predefined selection rule. We distinguish between the min-buying, max-buying, and rank-buying models.

For the min-buying model we show that it is not approximable within $\mathcal{O}(\log^{\varepsilon}|\mathcal{C}|)$ for some constant $\varepsilon > 0$, unless NP $\subseteq$ DTIME$(n^{\mathcal{O}(\log\log n)})$, thereby closing the gap between the known algorithmic results and previous lower bounds. We also prove inapproximability within $\mathcal{O}(\ell^{\varepsilon})$, ℓ being an upper bound on the number of non-zero budgets per consumer, and $\mathcal{O}(|\mathcal{P}|^{\varepsilon})$ under slightly stronger assumptions and provide matching upper bounds. Surprisingly, these hardness results hold even if a price ladder constraint, i.e., a predefined order on the prices of all products, is given.

For the max-buying model a PTAS exists if a price ladder is given. We give a matching lower bound by proving strong NP-hardness. Assuming limited product supply, we analyze a generic local search algorithm and prove that it is 2-approximate. Finally, we discuss implications for the rank-buying model.

1 Introduction

In recent years many technical improvements have made increasingly convenient access to all kinds of Internet services available to a broad public, making the Internet the world's largest market place. But to both consumers and companies the Internet offers possibilities far beyond the capabilities of traditional markets. Many websites that compare the prices offered by different companies for a certain product help customers make optimal buying decisions. For companies, running such websites might be just as profitable, because they help gathering data about the preferences of a huge number of potential customers and boosting sales by optimizing product profiles and applying intelligent pricing schemes tailored for the specific market.

Aiming at the latter objective Glynn, Rusmevichientong and Van Roy [22] defined the *non-parametric multi-product pricing problem*. Consumers are characterized by their budgets for different products and a selection rule describing how a consumer selects a product among those she can afford once prices are fixed. Since consumers will buy exactly one product if they can afford it, they are usually referred to as *unit-demand*. Glynn et al. propose three different selection rules. In the *rank-buying* model each consumer has a ranking of all the products she is interested in. When prices are fixed she buys the highest ranked product with a price below her respective budget. In the *min-buying* and *max-buying* models a consumer buys the product with lowest or highest price not exceeding her budget, respectively. The objective of the problem is to compute prices of the products and (possibly) a corresponding allocation of the products to consumers that maximize total revenue. Rusmevichientong et al. [22, 34] show that the min-buying model, where each consumer has the same budget for all products she desires, allows a polynomial time algorithm, assuming a *price ladder constraint*, i.e., a predefined total order on the prices of all products. Such a constraint is often implied by the set of products in question. Aggarwal, Feder, Motwani and Zhu [1] give approximation algorithms for all three models: a PTAS for both rank-buying and max-buying with price ladder, a 1.59-approximation for max-buying without price ladder, and a logarithmic approximation for any of the above models, assuming unlimited supply of the products. In the limited supply case a 4-approximation is derived for max-buying with price ladder. There are many practical situations in which it is desirable to be able to handle limited supply, as well. Besides the obvious point that it might not be possible to increase production capacity beyond a certain limit, even artificially limiting product supply can sometimes be rewarding.

Further results about the limited supply case were given by Guruswami, Hartline, Karlin, Kempe, Kenyon and McSh-

*Department of Computer Science, University of Dortmund, Otto-Hahn-Str. 14, 44221 Dortmund, Germany. E-mail: patrick.briest@cs.uni-dortmund.de. The author is supported by DFG grant Kr 2332/1-2 within Emmy Noether program.

†Department of Computer Science, University of Dortmund, Otto-Hahn-Str. 14, 44221 Dortmund, Germany. E-mail: piotr.krysta@cs.uni-dortmund.de. The author is supported by DFG grant Kr 2332/1-2 within Emmy Noether program.

erry [24], who introduced another selection rule inspired by the notion of truthfulness in auction design. In the *envy-free pricing problem* a consumer buys the product that maximizes her personal utility, i.e., the difference between the product's price and her respective budget. A set of prices together with a corresponding allocation of the products is envy-free, if every consumer indeed receives the product maximizing her utility. Guruswami et al. present an algorithm with logarithmic approximation ratio for this problem.

So far, there have been large gaps between the lower and upper bounds on the approximation ratio for almost all of these pricing problems, the only exceptions being the max-buying model without price ladder, for which a constant approximation and APX-hardness are known [1], and the *single-minded* envy-free model with unlimited supply [15, 24], discussed in Section 1.1. To our knowledge there are no non-constant inapproximability results known for any of the above unit-demand pricing problems.

As our main contribution we resolve the question of approximability of most of the above unit-demand pricing models, putting emphasis on the *hardness of approximation*. In particular we prove near-tight hardness results for the min-buying and max-buying models, and some versions of the rank-buying model (including the most general). Many of our hardness results show the first non-constant, logarithmic, and even polynomial inapproximability for those problems. We also give algorithmic results, which close the gap in approximability of some of those models. Finally, the problem is studied from a game theoretic standpoint. Namely, we study the multi-player game obtained by assuming that the price of every product is determined by a distinct agent trying to maximize her personal revenue, and present a bound on the *price of anarchy* (cf. [26, 32]) in this game.

1.1 Related Work Following the introduction of *algorithmic mechanism design* [31] as a major field of interest for computer science, a lot of research has been done on problems motivated by economical questions. While in *combinatorial auction design* the main goal is to motivate agents to participate truthfully in the protocol, the optimized objective is twofold. On one side social welfare is to be maximized. Various such results have been obtained for the case of *single-minded* agents [4, 13, 27, 30]. On the other side an auctioneer is clearly interested in auctions that generate high revenue. Goldberg et al. [23] and Fiat et al. [19] first investigated whether and how these two objectives can be combined. While originally only randomized revenue maximizing protocols were known, meanwhile a first deterministic protocol was designed by Aggarwal et al. [2].

While truthfulness of auctions can be assured when agents are single-minded and, thus, of a severely restricted kind, the situation is much more complicated for more general types of agents [8, 28]. Only recently it has been shown by Lavi and Swamy [29] and by Dobzinski, Nisan and Schapira [17] that in fact randomization can help to overcome this difficulty in many practically relevant cases. Another focus in general auction design is on algorithms for winner determination. Here, the incentives of single players are left aside and the goal is to find an allocation of the products that guarantees high overall social welfare. For recent results on so-called *submodular bidders* see [16] and references therein.

Besides the unit-demand pricing problem, a closely related line of research is multi-product pricing with single-minded consumers. Such consumers are interested in buying a single set of products rather than a single product out of a set of alternatives. Guruswami et al. [24] derive a logarithmic approximation for this problem with unlimited supply. Recently, Demaine et al. [15] have shown logarithmic hardness of approximation for this model. To our knowledge this is the only non-constant inapproximability result known for any of the discussed pricing problems. Further results on this so-called *single-minded unlimited supply pricing problem* are also found in [7, 11, 12, 25], where interest is paid mainly to various types of restricted problem instances.

1.2 Preliminaries Throughout the rest of the paper the setting will be as follows. Given a set of products $\mathcal{P}$ and consumer samples $\mathcal{C}$ with budgets $b(c, e)$ for all $c \in \mathcal{C}$, $e \in \mathcal{P}$ we want to assign prices $p(e)$ to the products that maximize the revenue from the resulting sale. This sale depends on how consumers decide whether and which product to buy once prices have been fixed. We differentiate between the min-buying, max-buying, and rank-buying models.

DEFINITION 1. **(Unit-Demand Pricing – UDP)** *We are given products $\mathcal{P}$ and consumer samples $\mathcal{C}$ consisting of budgets $b(c, e) \in \mathbb{R}_0^+$ for all $c \in \mathcal{C}$, $e \in \mathcal{P}$ and rankings $r_c : \mathcal{P} \to \{1, \ldots, |\mathcal{P}|\}$. For a price assignment $p : \mathcal{P} \to \mathbb{R}_0^+$ we let $\mathcal{A}(p) = \{c \in \mathcal{C} \mid \exists e \in \mathcal{P} : p(e) \le b(c, e)\}$ refer to the set of consumers that can afford to buy any product under p. In the no price ladder scenario (NPL) we want to find prices p that maximize*

- $\sum_{c \in \mathcal{A}(p)} \min\{p(e) \mid p(e) \le b(c, e)\}$.
 (UDP-MIN-NPL)
- $\sum_{c \in \mathcal{A}(p)} \max\{p(e) \mid p(e) \le b(c, e)\}$.
 (UDP-MAX-NPL)
- $\sum_{c \in \mathcal{A}(p)} p\big(\operatorname{argmin}\{r_c(e) \mid e : p(e) \le b(c, e)\}\big)$.
 (UDP-RANK-NPL)

Given a price ladder constraint (PL), $p(e_1) \le \cdots \le p(e_{|\mathcal{P}|})$, UDP-{MIN,MAX,RANK}-PL *asks for a price assignment p satisfying this constraint.*

The above definition assumes that all products are available in unlimited supply and, thus, any number of consumers requesting to buy some product can be satisfied. Let UDP-MAX-{PL,NPL} be defined as above and assume that of any product e there are only s_e many copies available. In this *limited supply case* we want to find not only a price assignment p, but also a feasible allocation $a : \mathcal{C} \to \mathcal{P}$ of the products, where $a(c)$ is the product given to consumer c. Allocation a is feasible if no more than s_e copies of product e are allocated and each consumer receives the most expensive product she can afford which is not sold out, i.e., if a consumer receives a product that is not the most expensive she can afford, then it must be the case that all affordable products with a higher price are sold out. Given prices p, finding the optimal allocation reduces to solving an instance of maximum weighted b-matching in a bipartite graph, where $b = (s_e)_{e\in\mathcal{P}}$ for vertices corresponding to products and $b = (1)_{c\in\mathcal{C}}$ for consumer vertices, which can be done in polynomial time [14].

Finally, we define a restricted version of the rank-buying model. UDP-RANK-{PL,NPL} *with consistent budgets* requires that for every consumer $c \in \mathcal{C}$, we have that $b(c, e) \geq b(c, f)$ whenever $r_c(e) < r_c(f)$ for all products $e, f \in \mathcal{P}$.

1.3 Contributions *Min-Buying:* Let us first focus on the UDP-MIN-{PL,NPL} problem with unlimited supply. The best known algorithm for this problem, which simply computes the optimum solution assigning the same price to every product, has an approximation factor of $\mathcal{O}(\log |\mathcal{C}|)$ (see Aggarwal et al. [1]). Surprisingly, it turns out that this simple algorithm is (close to) best possible, as we prove that there is no $\mathcal{O}(\log^{\varepsilon} |\mathcal{C}|)$-approximation algorithm for some absolute $\varepsilon > 0$, assuming NP $\not\subseteq$ DTIME($n^{\mathcal{O}(\log\log n)}$). In fact, an approximability threshold of Δ^{ε} for the independent set problem in graphs of degree at most Δ yields the same constant ε in our reduction. As so far no algorithms with approximation guarantee essentially below Δ are known for independent set, this suggests that UDP-MIN-{PL,NPL} does not allow approximation ratios essentially better than $\log |\mathcal{C}|$. We emphasize that this inapproximability result holds even in the presence of a price ladder constraint. This stands in a sharp contrast with the restricted version of UDP-MIN-PL, in which we assume that each consumer has the same budget for all the goods she desires [34] [1]. Surprisingly, after very few natural maximization problems with logarithmic approximation threshold have been known for quite some time (see [18] for one of the first examples), UDP is already the second problem from the field of product pricing (see [15]) for which such a threshold can be shown.

Applying a number of small modifications our reduction also yields almost tight hardness results when the approximation ratio is expressed in terms of ℓ, i.e., the maximum number of positive budgets of any consumer, and $|\mathcal{P}|$, the number of products. We prove that the problem is hard to approximate within $\mathcal{O}(\ell^{\varepsilon})$ and within $\mathcal{O}(|\mathcal{P}|^{\varepsilon})$ for some $\varepsilon > 0$, unless NP $\subseteq$ DTIME($2^{\mathcal{O}(n^{\delta})}$) for all $\delta > 0$ and present matching upper bounds. Specifically, there is a trivial $\mathcal{O}(|\mathcal{P}|)$-approximation and an approach of Balcan and Blum [7] implies an $\mathcal{O}(\ell)$-approximation for UDP-MIN-NPL.

Techniques: We use the classical method of graph products [9] to amplify the inapproximability threshold of the maximum independent set problem in bounded degree graphs. We first slightly extend the derandomized version of that construction due to Alon et al. [3] and parameterize the maximum degree of the constructed graph product in the number of its vertices. We then encode independence in such graphs by classes of geometrically increasing budgets in our pricing problem, where vertices correspond to products. The difficulty here is that independence needs to be enforced in a somewhat asymmetric way, i.e., based on a vertex coloring of the given graph, we can define collections of consumers that encode independence of a vertex from adjacent vertices with colors of smaller index, but we cannot do this in the opposite direction.

Max-Buying: Let us now switch our interest to the max-buying model. From the economical viewpoint this version finds less motivations, but as we will see our motivation to study this problem comes from its connection to the economically well motivated rank-buying model, and from the fact that the max-buying problem turns out to be tractable as compared to the min-buying problems.

The best previous algorithms for the max-buying problem were given by Aggarwal, Feder, Motwani and Zhu [1]. For UDP-MAX-NPL they present a 1.59-approximation based on a linear programming relaxation and randomized rounding and prove that the problem is NP-hard to approximate within 16/15. For UDP-MAX-PL they present a PTAS based on a rather involved dynamic programming approach. However, they left open the question whether this is the best possible algorithmic result that can be obtained in the presence of a price ladder constraint. We answer this question in the affirmative by proving strong NP-hardness of UDP-MAX-PL.

We then consider the effect of having to deal with limited product supply. The only known result for this situation is a 4-approximation for limited-supply UDP-MAX-PL [1]. We first have a closer look at the relation between the maximum supply and the problem's complexity. We show that UDP-MAX-NPL can be solved in polynomial time for unit-supply but becomes APX-hard already with maximum supply of only 2. On the algorithmic side, we analyze the performance of a generic local search algorithm and prove

[1] A polynomial time algorithm follows basically by observing that in the presence of a price ladder each consumer who is able to buy any product buys the product with smallest price according to the price ladder. This reduces the number of products to be considered for each consumer to one. Then one uses dynamic programming.

Variation	Previous [Lower] Upper	New Lower {Assumption}	New Upper
UDP-MIN-{PL,NPL}	[APX-hard], {UDP-MIN-NPL} $\mathcal{O}(\log\|\mathcal{C}\|)$	$\Omega(\ell^{\varepsilon})$ {NP $\not\subseteq$ DTIME($2^{\mathcal{O}(n^{\delta})}$)} $\Omega(\|\mathcal{P}\|^{\varepsilon})$ {NP $\not\subseteq$ DTIME($2^{\mathcal{O}(n^{\delta})}$)} $\Omega(\log^{\varepsilon}\|\mathcal{C}\|)$ {NP $\not\subseteq$ DTIME($n^{\mathcal{O}(\log\log n)}$)}	$\mathcal{O}(\ell)$ {NPL only} $\mathcal{O}(\|\mathcal{P}\|)$
UDP-MAX-NPL	[16/15], 1.59 (LP-based)	–	2 (combinatorial)
UDP-MAX-NPL {Limited supply}	[–], –	APX-hard {supply ≥ 2} in P {supply ≤ 1}	2
UDP-MAX-PL {Limited supply}	[–], 4	strongly NP-hard	–
UDP-MAX-PL	[–], PTAS	strongly NP-hard	–

Figure 1: Results apply to unlimited supply, unless stated otherwise. Hardness results with ε and complexity assumptions with δ are assumed to hold for some ε, $\delta > 0$. Previous upper and lower bounds are from [1].

that it yields a 2-approximation for limited or unlimited supply UDP-MAX-NPL. This complements our APX-hardness result for this problem, and in fact it is the first algorithm for the limited-supply case without price ladder with provable approximation guarantee. For unlimited supply UDP-MAX-NPL our ratio does not match the best known result, which gives a 1.59-approximation [1]. However, the previous algorithm is based on a rather problem specific LP-formulation and rounding techniques. Local search, on the other hand, appears to be a quite natural approach to a wide range of pricing problems. Seen in this light, we provide first evidence that this approach might indeed be promising also for more practical problems. We also show that the analysis can be extended to bound the price of anarchy in the related pricing game, in which each product is owned by an individual agent setting its price.

Rank-Buying: We point out that many of the discussed results can be transferred to the rank-buying model. All hardness results for UDP-MIN-{PL,NPL} hold for UDP-RANK-{PL,NPL} if we allow non-consistent budgets. Additionally, all known algorithmic results apply here, as well. If we require consistent budgets, we obtain strong NP-hardness for UDP-RANK-PL as in the max-buying case, complementing the existing PTAS.

Outline: The rest of the paper is organized as follows. Section 2 presents inapproximability results for UDP-MIN-{PL,NPL} and derives an $O(\ell)$-approximation. Section 3 starts with proving strong NP-hardness of UDP-MAX-PL. The remainder of the section shows APX-hardness of sparse limited-supply versions of this problem and analyzes a generic local search algorithm and the related unit-demand pricing game. Section 4 states which of the results apply to the rank-buying model. Section 5 concludes.

2 Min-Buying

We start by considering the min-buying model. Aggarwal et al. [1] give an algorithm with approximation guarantee $\mathcal{O}(\log|\mathcal{C}|)$ and prove that no algorithm that assigns only a constant number of different prices can beat this bound by more than a constant factor. We show that under very reasonable complexity theoretic assumptions no polynomial time algorithm can obtain approximation guarantee $\mathcal{O}(\log^{\varepsilon}|\mathcal{C}|)$, for some $\varepsilon > 0$ and, thus, the very simple algorithm of [1] turns out to be (close to) best possible. Interestingly, these results hold even in the price ladder scenario. We also show that a slightly stronger assumption leads to another strong inapproximability result in terms of the number ℓ of non-zero budgets per consumer or the total number of products $|\mathcal{P}|$. Finally, we supply the matching upper bound for one of the hardness results by deriving an $\mathcal{O}(\ell)$-approximation and a (trivial) $\mathcal{O}(|\mathcal{P}|)$-approximation for the no price ladder case.

2.1 Hardness of Approximation The hardness results of this section are based on a reduction of the independent set problem (IS). In order to obtain, e.g., logarithmic hardness as in Theorem 2.1 or in terms of parameter ℓ as in Theorem 2.2, it is important to have classes of restricted IS with different asymptotic inapproximability. Proposition 2.1, which extends a result from [3], states that we can obtain these restricted classes by considering graphs with maximum degree bounded in terms of their number of vertices. For a given graph G let $\alpha(G)$ denote the size of its maximum independent set.

PROPOSITION 2.1. *For arbitrary non-decreasing functions $f : \mathbb{N} \longrightarrow \mathbb{R}_+$ with $f(n) \leq n$ and $f(n^c) \leq f(n)^c$ for all $c \geq 1$, $n \in \mathbb{N}$, let $\mathcal{G}_f$ be the family of graphs $G = (V, E)$,*

$|V| = n$, *with maximum degree* $\Delta = \mathcal{O}(f(n))$. *There exists a constant* $\varepsilon > 0$, *such that it is NP-hard to approximate* $\alpha(G)$ *within* $\mathcal{O}(f(n)^{\varepsilon})$ *for* $G \in \mathcal{G}_f$.

Theorem 2.1 is based on a reduction of these restricted classes of IS and shows an approximability threshold for UDP-MIN-{PL,NPL}, which is (semi-) logarithmic in the number of consumers. Theorems 2.2 and 2.3 demonstrate the flexibility of the reduction and state corresponding thresholds expressed in terms of the number ℓ of non-zero budgets per consumer and the number of products $|\mathcal{P}|$, respectively.

THEOREM 2.1. UDP-MIN-{PL,NPL} *with unlimited supply is not approximable within* $\mathcal{O}(\log^{\varepsilon} |\mathcal{C}|)$ *for some* $\varepsilon > 0$, *unless NP* $\subseteq$ *DTIME*($n^{\mathcal{O}(\log\log n)}$).

Proof. Consider the family $\mathcal{G}$ of graphs $G = (V, E)$, $|V| = n$, with degree bounded by $\mathcal{O}(\log n)$. By Proposition 2.1 it is NP-hard to approximate $\alpha(G)$ for $G \in \mathcal{G}$ within $\mathcal{O}(\log^{\varepsilon} n)$. Towards a contradiction, we assume that there is a polynomial time algorithm with approximation guarantee $\mathcal{O}(\log^{\varepsilon-\delta} |\mathcal{C}|)$ for UDP-MIN-PL for some small $\delta > 0$. For a given graph $G = (V, E)$ from $\mathcal{G}$ let Δ denote its maximum degree. Clearly, we can compute a $(\Delta + 1)$-coloring of the vertices of G, which we denote by $V = V_0 \cup \ldots \cup V_\Delta$. For ease of notation let $V_i = \{v_{i,j} \,|\, j = 0, \ldots, |V_i| - 1\}$. Furthermore, by $\mathcal{V}(v_{i,j}) = \{v_{k,\ell} \,|\, \{v_{i,j}, v_{k,\ell}\} \in E$ and $k < i\}$ we refer to the vertices that are adjacent to $v_{i,j}$ and belong to a color class with index smaller than i. We proceed by defining an UDP-MIN-PL instance.

Products / Price Ladder Constraint: For every $v_{i,j} \in V$ we have a product $e_{i,j}$. The price ladder is defined as

$$p(e_{0,0}) \leq p(e_{0,1}) \leq \cdots \leq p(e_{0,|V_0|-1}) \leq p(e_{1,0}) \leq \cdots .$$

Let $\mu = 4(\Delta + 1)$ and $\gamma = \mu^{-\Delta-1}/n$. For every product $e_{i,j}$ we define a corresponding threshold $p_{i,j} = \mu^{i-\Delta} + j\gamma$. Observe that thresholds are arranged according to the price ladder constraint and differ from each other by at least γ.

Consumers: For every $v_{i,j} \in V$ define a collection $\mathcal{C}_{i,j} = \{c^t_{i,j} \,|\, t = 0, \ldots, \mu^{\Delta-i} - 1\}$ of identical consumers with budgets $b(c^t_{i,j}, e_{i,j}) = p_{i,j}$ and $b(c^t_{i,j}, e_{k,\ell}) = p_{k,\ell}$ for all k,ℓ with $v_{k,\ell} \in \mathcal{V}(v_{i,j})$. In analogy to the coloring of G we denote consumers as $\mathcal{C} = \mathcal{C}_0 \cup \ldots \cup \mathcal{C}_\Delta$, where $\mathcal{C}_i = \bigcup_j \mathcal{C}_{i,j}$. Note, that all budgets are consistent with the thresholds we just defined. The complete construction is illustrated in Figure 2.

Soundness: Let opt_{UDP} denote the revenue made by an optimal price assignment on the above instance. We first argue that this defines an upper bound on the size of a maximum independent set in G, i.e., $opt_{UDP} \geq \alpha(G)$. Given an independent set V' of G, we can define a price assignment p as follows. If $v_{i,j} \in V'$ set $p(e_{i,j}) = p_{i,j}$, else set $p(e_{i,j}) = p_{i,j} + \gamma$. Since the $p_{i,j}$'s differ by at least γ this assignment is clearly in accordance with the price ladder constraint.

Now consider $v_{i,j} \in V'$ and the corresponding consumers $\mathcal{C}_{i,j}$. Since $v_{k,\ell} \notin V'$ for all $v_{k,\ell} \in \mathcal{V}(v_{i,j})$, each consumer $c^t_{i,j}$ can afford to buy product $e_{i,j}$ at its threshold price $p_{i,j}$, while the prices of all products $e_{k,\ell}$ are above their thresholds and, thus, exceed the consumers' respective budgets. Hence, $e_{i,j}$ is indeed the product with smallest price that any $c^t_{i,j}$ can afford. It follows that the overall revenue from consumers $\mathcal{C}_{i,j}$ is at least $|\mathcal{C}_{i,j}| \cdot p_{i,j} = \mu^{\Delta-i}\left(\mu^{i-\Delta} + j\gamma\right) \geq 1$. Thus, price assignment p results in revenue of at least $|V'|$ and we may conclude that $opt_{UDP} \geq \alpha(G)$.

Completeness: Assume now that our approximation algorithm returns a price assignment p. By $r(\mathcal{C})$ we refer to the overall revenue of this price assignment, $r(\mathcal{C}_{i,j})$ and $r(c^t_{i,j})$ denote the revenue made by sales to consumers in $\mathcal{C}_{i,j}$ and to $c^t_{i,j}$ alone, respectively. First observe that w.l.o.g. the price of each product $e_{i,j}$ is either $p_{i,j}$ or $p_{i,j} + \gamma$. To see this, note, that as long as this is not the case there is always a price that we can increase up to $p_{i,j}$ or decrease down to $p_{i,j} + \gamma$ without decreasing the overall revenue. Define $\mathcal{C}^+ = \left\{c^t_{i,j} \,|\, r(c^t_{i,j}) = p_{i,j}\right\}$ as the set of consumers buying at maximum possible price and $\mathcal{C}^- = \mathcal{C} \backslash \mathcal{C}^+$. Clearly $\mathcal{C}_{i,j} \subseteq \mathcal{C}^+$ or $\mathcal{C}_{i,j} \subseteq \mathcal{C}^-$ for all i and j, since all $c^t_{i,j}$'s budgets are identical. We want to show that a large portion of the solution's revenue is due to consumers in $\mathcal{C}^+$. Note, that a consumer $c^t_{i,j} \in \mathcal{C}^-$ buys at price at most $p_{i-1,|V_{i-1}|-1}$. Thus, we have:

$$\begin{aligned} r(\mathcal{C}^-) &= \sum_{\mathcal{C}_{i,j} \subseteq \mathcal{C}^-} r(\mathcal{C}_{i,j}) \leq \sum_{\mathcal{C}_{i,j} \subseteq \mathcal{C}^-} |\mathcal{C}_{i,j}| \cdot p_{i-1,|V_{i-1}|-1} \\ &\leq \sum_{\mathcal{C}_{i,j} \subseteq \mathcal{C}^-} \mu^{\Delta-i}(\mu^{i-1-\Delta} + n\gamma) \leq \sum_{\mathcal{C}_{i,j} \subseteq \mathcal{C}^-} 2\mu^{-1} \\ &= \sum_{\mathcal{C}_{i,j} \subseteq \mathcal{C}^-} \frac{1}{2(\Delta+1)} \leq \frac{n}{2(\Delta+1)} \end{aligned}$$

On the other hand, it clearly holds that $\alpha(G) \geq n/(\Delta + 1)$ and, in fact, it is straightforward to construct a price assignment resulting in revenue $n/(\Delta + 1)$. It follows that we may assume w.l.o.g. that $r(\mathcal{C}) \geq n/(\Delta + 1)$ and, thus, $r(\mathcal{C}^+) = r(\mathcal{C}) - r(\mathcal{C}^-) \geq (1/2)r(\mathcal{C})$. Define $V' = \{v_{i,j} \,|\, \mathcal{C}_{i,j} \subseteq \mathcal{C}^+\}$. Let $v_{i,j} \in V'$ and consider the corresponding consumers $\mathcal{C}_{i,j} \subseteq \mathcal{C}^+$. The revenue made by sales to consumers in $\mathcal{C}_{i,j}$ is $|\mathcal{C}_{i,j}| \cdot p_{i,j} = \mu^{\Delta-i}\left(\mu^{i-\Delta} + j\gamma\right) \leq 2$. We conclude that $|V'| = |\{v_{i,j} \,|\, \mathcal{C}_{i,j} \subseteq \mathcal{C}^+\}| \geq (1/2)r(\mathcal{C}^+)$. Finally, observe that V' is indeed a feasible independent set in G. To see this, consider $v_{i,j} \in V'$ and let $v_{k,\ell}$ be an adjacent vertex. If $k < i$, the fact that consumers $\mathcal{C}_{i,j}$ buy $e_{i,j}$ at price $p_{i,j}$ implies that the price of $e_{k,\ell}$ is strictly larger than its threshold $p_{k,\ell}$. It follows that $\mathcal{C}_{k,\ell} \nsubseteq \mathcal{C}^+$ and, thus, $v_{k,\ell} \notin V'$. If $k > i$, consumers $\mathcal{C}_{k,\ell}$ can afford to buy product $e_{i,j}$ at price $p_{i,j}$ and again $\mathcal{C}_{k,\ell} \nsubseteq \mathcal{C}^+$.

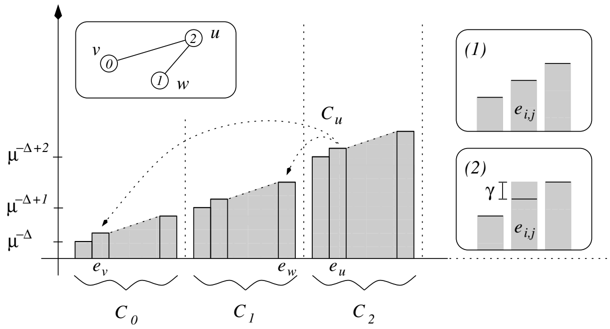

Figure 2: Products are arranged in blocks according to the $(\Delta+1)$-coloring of G, thresholds in block i are roughly $\mu^{-\Delta+i}$. The additional offset γ allows setting prices according to the price ladder. Consumers $\mathcal{C}_u$ belonging to vertex u (u stands for some $v_{i,j}$ here) have non-zero budgets for e_u and products in blocks with lower numbers corresponding to adjacent vertices. Cases on the right illustrate how price $p(e_{i,j})$ is set to indicate that $v_{i,j} \in V'$ (1), or $v_{i,j} \notin V'$ (2).

Remember that $|\mathcal{C}|$ denotes the number of consumers in our instance and note that $\log |\mathcal{C}| \leq \log n\mu^{\Delta} = \mathcal{O}(\log^{1+\gamma} n)$ for any $\gamma > 0$. Applying our initial assumption that $r(\mathcal{C})$ is an $\mathcal{O}(\log^{\varepsilon-\delta} |\mathcal{C}|)$-approximation to opt_{UDP} we finally obtain

$$|V'| \geq \frac{1}{2} r(\mathcal{C}^+) \geq \frac{1}{4} r(\mathcal{C}) \geq \frac{1}{\mathcal{O}(\log^{\varepsilon-\delta} |\mathcal{C}|)} opt_{UDP} \geq \frac{1}{\mathcal{O}(\log^{\varepsilon} n)} \alpha(G).$$

By Proposition 2.1 finding such an independent set is NP-hard. The size of our UDP-MIN-PL instance is roughly $n \cdot (\log n)^{\log n} = n^{\mathcal{O}(\log\log n)}$ and the running time of our approximation algorithm will be polynomial in this expression. Finally, observe that the proof will go through, as well, if we do not impose a price ladder constraint. □

By slightly changing the reduction in the proof above, we can obtain a similar inapproximability result in terms of the number ℓ of non-zero budgets of a single consumer or the total number of products, respectively. Starting from a graph $G \in \mathcal{G}_f$ with $f(n) = \mathcal{O}(n^{\delta})$ for some $\delta \geq 0$, we again define groups of products according to an $\mathcal{O}(f(n))$-coloring of G. Thresholds will be roughly powers of n. The resulting UDP-MIN-{PL,NPL} instance has size $n^{f(n)} = 2^{\mathcal{O}(n^{\gamma})}$ for some $\gamma > 0$ and is not approximable within $f(n)^{\varepsilon}$ by assumption. This yields the following result under a slightly stronger complexity theoretic assumption.

THEOREM 2.2. UDP-MIN-{PL,NPL} *with unlimited supply and with at most ℓ non-zero budgets per consumer is not approximable within $\mathcal{O}(\ell^{\varepsilon})$ for some constant $\varepsilon > 0$, unless NP $\subseteq$ DTIME($2^{\mathcal{O}(n^{\delta})}$) for all $\delta > 0$. Especially,* UDP-MIN-{PL,NPL} *with unlimited supply is not approximable within $\mathcal{O}(|\mathcal{P}|^{\varepsilon})$ under the same assumption.*

Finally, we want to point out that going to a weaker assumption than the one in Theorem 2.1 our reduction still shows that no constant factor approximation is possible. Applying the reduction to graphs of constant degree Δ the reduction yields a UDP-MIN-{PL,NPL} instance of polynomial size $\mathcal{O}(n^{\Delta})$, which is not approximable within Δ^{ε} for some $\varepsilon > 0$ by [3].

THEOREM 2.3. UDP-MIN-{PL,NPL} *with unlimited supply does not allow any constant approximation ratio, unless NP $\subseteq$ P.*

2.2 An $\mathcal{O}(\ell)$-Approximation We first observe that there is a trivial $\mathcal{O}(|\mathcal{P}|)$-approximation algorithm for both UDP-MIN-{PL,NPL} and UDP-RANK-{PL,NPL} with unlimited supply. We just sell a single product for the best price to all potential consumers. Let us then consider unlimited supply versions of UDP-{MIN,RANK}-NPL with a maximum number ℓ of non-zero budgets per consumer. In [7], Balcan and Blum present an $\mathcal{O}(\ell)$-approximation for the single-minded unlimited supply pricing problem. We briefly sketch the main idea of this algorithm and its application to UDP-MIN-NPL. The algorithm is based on a random partition $\mathcal{P} = \mathcal{Q} \cup \mathcal{R}$ of the products, where each product is placed in $\mathcal{Q}$ with probability $1/\ell$ and in $\mathcal{R}$ with probability $1 - 1/\ell$. The key observation is that every consumer has exactly one non-zero budget for a specific product $e \in \mathcal{Q}$ with probability at least $1/(\mathrm{e}\ell)$ and, thus, we obtain a solvable problem instance (with each consumer interested in exactly one product) that carries a $1/(\mathrm{e}\ell)$-fraction of the optimal revenue in expectation. Finally, [7] also shows how to derandomize the above algorithm. Thus, we obtain the following result.

THEOREM 2.4. UDP-{MIN,RANK}-NPL *with unlimited supply and at most ℓ non-zero budgets per consumer can be approximated in polynomial time within $\mathcal{O}(\ell)$.*

3 Max-Buying

We now turn to the max-buying model. We first consider the practically relevant case of UDP-MAX-PL. Aggarwal et al. [1] point out that given a price ladder constraint, rank-buying with consistent budgets reduces to UDP-MAX-PL and give a PTAS for this problem. We present a matching hardness result and settle the question of this model's computational complexity. We then investigate the effect of having to deal with limited product supply. This question has been addressed in [1] assuming that a price ladder is given. We show that apparently the problem does not get more complex if this assumption is removed.

3.1 Unlimited Supply and Price Ladder

The following theorem states that UDP-MAX-PL is strongly NP-hard. This resolves a previously open problem from [1].

THEOREM 3.1. UDP-MAX-PL *with unlimited supply is strongly NP-hard, even if each consumer has at most* 2 *non-zero budgets.*

Sketch of Proof. We show that MAX-2SAT $\leq_P$ UDP-MAX-PL (see [21] for NP-hardness of MAX-2SAT). We are given a collection of disjunctive clauses $c_1, \ldots, c_m$ of length at most 2 over variables $x_1, \ldots, x_n$ and some positive integer $s \in \mathbb{N}$ and want to decide whether there is a truth assignment $t : \{x_1, \ldots, x_n\} \to \{0, 1\}$ that simultaneously satisfies s of the clauses. Note, that w.l.o.g. we may assume that $n \leq m$.

Variable gadgets: For every variable x_i we construct a gadget $\mathcal{V}_i$ consisting of 2 products e_i, f_i and the following collection of consumers:

- α_i^j, $j = 1, \ldots, 4m$, with budgets $b(\alpha_i^j, e_i) = 1 + \frac{2i-2}{2m^2}$ and $b(\alpha_i^j, f_i) = 1 + \frac{2i-1}{2m^2}$.
- β_i^j, $j = 1, \ldots, 4m^3$, with budgets $b(\beta_i^j, e_i) = 1 + \frac{2i-1}{2m^2}$.
- γ_i^j, $j = 1, \ldots, 4m^3 + 4m$, with budgets $b(\gamma_i^j, f_i) = 1 + \frac{2i}{2m^2}$.

Budgets that are not explicitly stated are assumed to be 0. By $r(\mathcal{V}_i)$ we refer to the revenue made from sales to the above consumers. We want to calculate the value of $r(\mathcal{V}_i)$ depending on prices $p(e_i)$ and $p(f_i)$. It is w.l.o.g. to assume that prices are chosen from the set of distinct budget values of the above consumers. We are especially interested in the following two cases. If $p(e_i) = 1 + (2i-2)/(2m^2)$, $p(f_i) = 1 + (2i)/(2m^2)$ we say that $\mathcal{V}_i$ is *in state* 1. If $p(e_i) = p(f_i) = 1 + (2i-1)/(2m^2)$ we say that $\mathcal{V}_i$ is *in state* 0. In our interpretation variable gadgets in state 0 correspond to variables that are assigned the boolean value 0, variable gadgets in state 1 to variables that are assigned 1.

Let now $r_i^* = (4m^3 + 4m)(2 + (4i-2)/(2m^2))$. A simple calculation yields that $r(\mathcal{V}_i) = r_i^*$ if $\mathcal{V}_i$ is in either state 0 or state 1 and $r(\mathcal{V}_i) \leq r_i^* - 2m$ else. This will ensure that prices are always set in correspondence with a feasible truth assignment.

Clause gadgets: For every clause c_j we define a single consumer δ_j with budgets $b(\delta_j, e_i) = 1 + (2i-2)/(2m^2)$, if clause c_j contains literal x_i and $b(\delta_j, f_i) = 1 + (2i-1)/(2m^2)$, if c_j contains $\overline{x}_i$. Again, budgets that are not explicitly stated are set to 0.

We finally impose a price ladder constraint that requires that $p(e_1) \leq p(f_1) \leq p(e_2) \leq p(f_2) \leq \cdots$ and let $r^* = \sum_i r_i^*$. For the constructed UDP-MAX-PL instance we now ask whether there exists a price assignment p that results in overall revenue of at least $r^* + s$ for $s \in \mathbb{N}$ as in the MAX-2SAT instance. The idea of the construction is depicted in Figure 3.

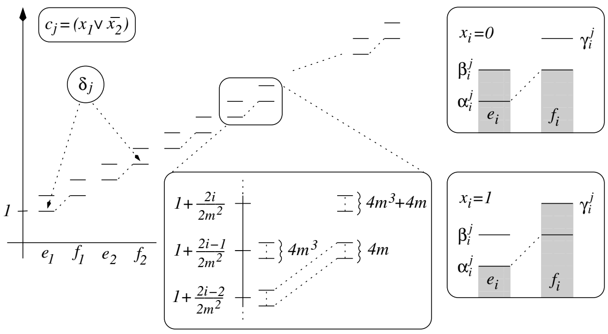

Figure 3: Consumers α_i^j, β_i^j, γ_i^j ensure that prices of e_i and f_i are always in state 0 ($x_i = 0$) or in state 1 ($x_i = 1$), both of which are consistent with the price ladder constraint. For each clause c_j we have a single consumer δ_j with non-zero budgets for the products corresponding to the literals of c_j.

Soundness: Let t be a truth assignment satisfying s of the clauses. If $t(x_i) = 0$ we set variable gadget $\mathcal{V}_i$ to state 0, if $t(x_i) = 1$ to state 1. Clearly, our price assignment is in accordance with the price ladder constraint. Obviously we obtain revenue r_i^* from each gadget $\mathcal{V}_i$. Furthermore, for each satisfied clause c_j it is clear that consumer δ_j can afford to buy one of the products corresponding to satisfied literals at a price strictly larger than 1. Thus, overall revenue is at least $r^* + s$.

Completeness: Let p be a price assignment resulting in revenue at least $r^* + s$. We construct a truth assignment t that satisfies s of the clauses. First observe that w.l.o.g. each variable gadget $\mathcal{V}_i$ is in either state 0 or state 1. This follows directly from the fact that setting prices differently would reduce profit from consumers α_i^j, β_i^j, γ_i^j by at least $2m$. On the other hand, there are at most m consumers of type δ_j and, thus, other consumers can generate profit of at most $m \cdot (1 + (2n)/(2m^2)) \leq m + 1$ by buying products e_i and f_i, respectively.

We can then define the obvious truth assignment t by $t(x_i) = 0$ if $\mathcal{V}_i$ is in state 0, $t(x_i) = 1$ if $\mathcal{V}_i$ is in state 1. For every consumer δ_j that can afford to buy a product under price assignment p the corresponding clause c_j is satisfied by t. Since revenue made by sales to consumers of type α_i^j, β_i^j and γ_i^j is precisely r^*, the number of consumers of type δ_j buying some product must be at least

$$\left\lceil s \cdot \left(1 + \frac{2m-1}{2m^2}\right)^{-1} \right\rceil \geq \left\lceil s \cdot \left(1 + \frac{1}{m}\right)^{-1} \right\rceil \geq s,$$

where we use the fact that $s \leq m$ and consumers δ_j buy at a price of at most $1 + (2m-1)/(2m^2)$. $\square$

3.2 Limited Supply We want to investigate the effect of assuming limited product supply and focus here on the no price ladder case. We first point out at what maximum supply the problem becomes hard to approximate and then present an approximation algorithm based on a local search approach. In the case of *unit-supply* we assume that there is exactly one copy of each product available, thus, $s_e = 1$ for all $e \in \mathcal{P}$. Since prices can be chosen w.l.o.g. from the set of distinct budget values, a given allocation now implies the optimal pricing and the problem reduces to finding an optimal weighted bipartite matching.

THEOREM 3.2. UDP-MAX-{PL,NPL} *with unit-supply can be solved in polynomial time.*

Surprisingly, the situation changes drastically already if maximum supply is increased to 2, which is enough to make the problem APX-hard. Note, that APX-hardness of general UDP-MAX-NPL has already been shown in [1]. However, our new reduction yields the precise maximum supply at which the problem becomes hard. The proof of Theorem 3.3 is omitted due to space limitations.

THEOREM 3.3. UDP-MAX-NPL *with limited supply* 2 *or larger is APX-hard.*

To obtain a matching algorithmic result, we will analyze the approximation guarantee of a generic local search approach to UDP-MAX-NPL. We let $r(p, a)$ refer to the overall revenue generated by price assignment p and corresponding allocation a. Unless stated otherwise we assume that a is chosen optimally. We start by briefly describing algorithm LOCALSEARCH. For a given price assignment p let $[p \,|\, p(e) = p']$ refer to the price assignment obtained by changing the price of e to p'.

> LOCALSEARCH: Initialize p arbitrarily and compute the optimal allocation a. While there exist product e and price $p' \neq p(e)$ such that
>
> $$r(p, a) < r([p \,|\, p(e) = p'], a'),$$
>
> where a' is the optimal allocation given prices $[p \,|\, p(e) = p']$, set $p(e) = p'$.

We next show that the total revenue generated by a locally optimal solution lies within a factor of 2 off the globally optimal solution's value.

THEOREM 3.4. *Let p be the price assignment returned by algorithm* LOCALSEARCH*, p^* an optimal price assignment and a, a^* the respective allocations. Then $r(p^*, a^*)/r(p, a) \leq 2$ and, thus, algorithm* LOCALSEARCH *achieves approximation ratio* 2 *for* UDP-MAX-NPL *with limited or unlimited supply. Furthermore, this bound is tight.*

Proof. Consider prices p and allocation a. We then define $C_e = (a^*)^{-1}(e)$, $L_e = \{c \in C_e \,|\, p(a(c)) < p^*(e)\}$ and $r_e = p(e)|a^{-1}(e)|$, i.e., C_e refers to the set of consumers buying e in an optimal solution, L_e is the subset of these consumers that buy at a price below $p^*(e)$ in the solution returned by LOCALSEARCH. Furthermore, we let $\Delta_e = \sum_{c \in L_e} (p^*(e) - p(a(c)))$ refer to the loss of revenue compared to the optimal solution incurred by consumers in C_e. Changing price $p(e)$ to $p^*(e)$ (or leaving it as it is in case it should happen to be just $p^*(e)$) defines price assignment $p' = [p \,|\, p(e) = p^*(e)]$ and corresponding allocation a'. Since we do not know what a' should look like we define an alternative allocation a'' as follows. First, we set $a''(c) = \varnothing$ for all consumers c with $a(c) = e$. We then set $a''(c) = e$ for all $c \in L_e$. For all other consumers we do not change allocation a and let $a''(c) = a(c)$. First observe that allocation a'' does not allocate more copies of any item than there are available, since $|L_e| \leq |C_e| \leq s_e$ and no product besides e can be sold to more consumers than in a. It immediately follows that $r(p', a') \geq r(p', a'')$. We observe that

$$\begin{aligned}
& r(p', a') - r(p, a) \geq r(p', a'') - r(p, a) \\
= & \sum_{c \notin L_e \cup a^{-1}(e)} p(a(c)) + \sum_{c \in L_e} p^*(e) - \sum_{c \in \mathcal{C}} p(a(c)) \\
\geq & \sum_{c \in \mathcal{C}} p(a(c)) + \sum_{c \in L_e} (p^*(e) - p(a(c))) \\
& - \sum_{c \in a^{-1}(e)} p(a(c)) - \sum_{c \in \mathcal{C}} p(a(c)) = \Delta_e - r_e.
\end{aligned}$$

By the fact that $r(p, a)$ cannot be improved by changing a single price $p(e)$ we have that $r(p', a') - r(p, a) \leq 0$ and, thus, $r_e \geq \Delta_e$. (If price $p(e)$ did not have to be changed because it was already $p^*(e)$ the same inequality follows from the optimality of allocation a.) We let now $r_e^* = p^*(e)|C_e|$ denote the revenue made by product e in the optimal solution. We can then write that

$$\begin{aligned}
2 \cdot r(p, a) &= \sum_{e \in \mathcal{P}} r_e + \sum_{c \in \mathcal{C}} p(a(c)) \\
&\geq \sum_{e \in \mathcal{P}} \big(r_e + \sum_{c \in C_e} p(a(c))\big) \\
&\geq \sum_{e \in \mathcal{P}} (r_e + r_e^* - \Delta_e) \geq \sum_{e \in \mathcal{P}} r_e^* = r(p^*, a^*).
\end{aligned}$$

This completes the first part of the proof. We note that there are quite simple examples proving that the above analysis is tight, which are omitted here due to space limitations. $\square$

So far, we have argued that algorithm LOCALSEARCH terminates with a solution that is a 2-approximation with respect to the optimal revenue. We have not, however, argued about the algorithm's running time. In order to obtain polynomial running time, only a small change needs to be applied. Instead of choosing any improving step, we need to find in each iteration the new price that will give maximum increase in revenue. This yields the following theorem.

THEOREM 3.5. UDP-MAX-NPL *with limited or unlimited supply and integral budgets can be approximated in polynomial time within a factor of* 2.

Proof. Assume that we choose in each iteration the new price that will give maximum increase in revenue. Let r be the revenue of the current solution, r^* the revenue of an optimal solution and assume that $r^* - 2r \geq \phi$. Using the same notation as in the proof of Theorem 3.4 we have

$$\sum_{e \in \mathcal{P}} (\Delta_e - r_e) = \sum_{e \in \mathcal{P}} (r_e + \Delta_e - 2r_e) = r^* - 2\sum_{e \in \mathcal{P}} r_e \geq \phi$$

and, thus, there must exist a product e, such that $r_e \leq \Delta_e - \phi/n$, where n denotes the number of products in the instance. It follows that revenue increases by at least ϕ/n in each iteration and, thus, after k iterations it must be true that $\phi \leq r^*(1-(2/n))^k$, since in the first iteration it holds that $\phi \leq r^*$. We assume that all budgets are integral. It follows that the overall revenue increases by at least 1 in each iteration. Now let $k^* = n \cdot \lceil \ln r^* \rceil + 1$. After k^* iterations we have that

$$\phi \leq r^* \left(1 - \frac{2}{n}\right)^{n \cdot \ln r^*} - 1 \leq r^* \cdot \mathrm{e}^{-\ln r^*} - 1 = 0,$$

and, thus, we can terminate the algorithm after k^* iterations with an approximation guarantee of 2. Note, that we do not need to know the value of r^*. For (weakly) polynomial running time it is sufficient to upper bound r^* by the sum of consumers' maximum budgets. □

3.3 The Price of Anarchy We want to point out that the analysis of algorithm LOCALSEARCH can be extended to bound the price of anarchy (the worst case ratio between the revenue of an optimal solution and any Nash equilibrium, see, e.g., [26]) in the pricing game we obtain if we let an individual player fix the price of each product. Since it can be shown that pure *Nash equilibria* do not generally exist, we will have to work here with the concept of mixed equilibria.

As we have argued before, we can restrict the set of potential prices to the set of distinct budget values given by the consumer samples. As a consequence, the strategy space of each player is of finite size and, thus, Nash's theorem [20] guarantees the existence of mixed equilibria.

Interestingly, the price of anarchy turns out to be 2, so in order to obtain good revenue in the max-buying scenario not even a global objective seems to be necessary. In the mixed strategy scenario every player $j \in \mathcal{P}$ defines a probability distribution P_j over a set of possible prices for her product. Let random variable R_j refer to the revenue of player j. A set of strategies $P^{eq} = (P_1^{eq}, \ldots, P_n^{eq})$ are at *Nash equilibrium*, if for every player j we have that $\mathrm{E}_{P^{eq}}[\,R_j\,] \geq \mathrm{E}_{(P_{-j}^{eq}, P_j')}[\,R_j\,]$ for all $P_j' \neq P_j^{eq}$, i.e., if no player can increase her expected revenue by changing her current strategy P_j^{eq}. (P_{-j}^{eq}, P_j') stands for the vector of strategies P^{eq} with strategy P_j^{eq} replaced by P_j'.

THEOREM 3.6. *The price of anarchy in the unit-demand max-buying pricing game is* 2.

We point out that the situation is quite different in the min-buying or rank-buying models, where it is easy to show that the price of anarchy is unbounded.

4 Rank-Buying

We finally turn to the rank-buying model and briefly describe which of the results presented in the previous sections apply here. Remember that UDP-RANK-{PL,NPL} *with consistent budgets* requires that for every consumer $c \in \mathcal{C}$, we have that $r_c(e) < r_c(f)$ implies $b(c,e) \geq b(c,f)$ for all products $e, f \in \mathcal{P}$. Given a price ladder constraint, UDP-RANK-PL with consistent budgets reduces to UDP-MAX-PL and, thus, the PTAS from [1] can be applied. It is straightforward to modify the proof of Theorem 3.1 in order to fit the rank-buying model.

THEOREM 4.1. UDP-RANK-PL *with unlimited supply and consistent budgets is strongly NP-hard, even if each consumer has at most* 2 *non-zero budgets.*

If we do not require consistent budgets, the problem immediately becomes a lot more intractable. Given a price ladder constraint, it is now straightforward to reduce any UDP-MIN-PL instance to a corresponding instance of UDP-RANK-PL by simply defining appropriate rankings. It is also straightforward to argue that the proof of Theorem 2.1 works for rank-buying without price ladder, as well, which implies similar hardness for UDP-RANK-NPL.

THEOREM 4.2. UDP-RANK-{PL,NPL} *with unlimited supply (allowing non-consistent budgets) is not approximable within* $\mathcal{O}(\log^{\varepsilon}|\mathcal{C}|)$ *for some* $\varepsilon > 0$*, unless NP* $\subseteq$ *DTIME(*$n^{\mathcal{O}(\log\log n)}$*). Allowing at most* ℓ *non-zero budgets per consumer it is not approximable within* $\mathcal{O}(\ell^{\varepsilon})$ *for some* $\varepsilon > 0$*, unless NP* $\subseteq$ *DTIME(*$2^{\mathcal{O}(n^{\delta})}$*) for all* $\delta > 0$*. Especially, it is not approximable within* $\mathcal{O}(|\mathcal{P}|^{\varepsilon})$ *under the same assumption. Assuming NP* $\nsubseteq$ *P,* UDP-RANK-{PL,NPL} *is not approximable within any constant factor.*

5 Conclusions and Open Problems

We have shown (near)-tight inapproximability and hardness results for a number of variations of the unit-demand pricing problem. Nevertheless, some interesting cases have still not been settled. Both UDP-MIN-PL and UDP-RANK-PL in the general case have turned out to allow no approximation guarantees essentially beyond the known logarithmic ratio. On the other hand, both problems become solvable exactly in polynomial time, if we require that each consumer c_i's budgets are either 0 or $v_i > 0$ (the uniform budget case). It is an interesting open question if this problem variation allows any constant approximation ratio in the no price ladder scenario. (APX-hardness follows from [24].) Also the complexity of *envy-free* (or *max-gain*) pricing as considered in [24] remains unresolved. It would be very interesting to obtain non-constant lower bounds for this problem, as well. We have presented a 2-approximation for UDP-MAX-NPL with limited supply. The best known approximation ratio for UDP-MAX-PL with limited supply, on the other hand, is 4 [1] and no lower bounds besides strong NP-hardness as shown in our paper are known. It would be very interesting to see whether a PTAS for the limited supply case is possible.

References

[1] G. Aggarwal, T. Feder, R. Motwani, and A. Zhu. Algorithms for Multi-Product Pricing. In *Proc. of 31st ICALP*, 2004.

[2] G. Aggarwal, A. Fiat, A.V. Goldberg, J.D. Hartline, N. Immorlica, and M. Sudan. Derandomization of Auctions. In *Proc. of 37th STOC*, 2005.

[3] N. Alon, U. Feige, A. Wigderson, and D. Zuckerman. Derandomized Graph Products. *Computational Complexity*, **5**(1): 60–75, 1995.

[4] A. Archer, C.H. Papadimitriou, K. Talwar, and É. Tardos. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. In *Proc. of 14th SODA*, 2003.

[5] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof Verification and Hardness of Approximation Problems. *J. ACM*, **45**(3):501–555, 1998.

[6] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation.* Springer, 1999.

[7] N. Balcan and A. Blum. Approximation Algorithms and Online Mechanisms for Item Pricing. In *Proc. of 7th EC*, 2006.

[8] Y. Bartal, R. Gonen, and N. Nisan. Incentive Compatible Multi-Unit Combinatorial Auctions. In *Proc. of 9th TARK*, 2003.

[9] P. Berman and G. Schnitger. On the Complexity of Approximating the Independent Set Problem. *Information and Computation*, **96**(1): 77–94, 1992.

[10] A. Blum. *Algorithms for Approximate Graph Coloring.* Ph.D. thesis, MIT Laboratory for Computer Science MIT/LCS/TR-506, May 1991.

[11] M. Bouhtou, A. Grigoriev, S. van Hoesel, A. van der Kraaij, and M. Uetz. Pricing Network Edges to Cross a River. In *Proc. of 2nd WAOA*, 2004.

[12] P. Briest and P. Krysta. Single-Minded Unlimited-Supply Pricing on Sparse Instances. In *Proc. of 17th SODA*, 2006.

[13] P. Briest, P. Krysta, and B. Vöcking. Approximation Techniques for Utilitarian Mechanism Design. In *Proc. of 37th STOC*, 2005.

[14] W.J. Cook, W.H. Cunningham, W.R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization.* John Wiley & Sons, 1998.

[15] E.D. Demaine, U. Feige, M.T. Hajiaghayi, and M.R. Salavatipour. Combination Can Be Hard: Approximability of the Unique Coverage Problem. In *Proc. of 17th SODA*, 2006.

[16] S. Dobzinski, and M. Schapira. An Improved Approximation Algorithm for Combinatorial Auctions with Submodular Bidders. In *Proc. of 17th SODA*, 2006.

[17] S. Dobzinski, N. Nisan, and M. Schapira. Truthful Randomized Mechanisms for Combinatorial Auctions. In *Proc. of 38th STOC*, 2006.

[18] U. Feige, M. Halldorsson, G. Kortsarz, and A. Srinivasan. Approximating the Domatic Number. *Siam Journal on Computing*, **32**(1): 172–195, 2002.

[19] A. Fiat, A.V. Goldberg, J.D. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. of 34th STOC*, 2002.

[20] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, 1991.

[21] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness.* Freeman, 1979.

[22] P. Glynn, P. Rusmevichientong, and B. Van Roy. A Non-Parametric Approach to Multi-Product Pricing. *Operations Research*, **54**(1):82–98, 2006.

[23] A.V. Goldberg, J.D. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. of 12th SODA*, 2001.

[24] V. Guruswami, J.D. Hartline, A.R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On Profit-Maximizing Envy-Free Pricing. In *Proc. of 16th SODA*, 2005.

[25] J. Hartline and V. Koltun. Near-Optimal Pricing in Near-Linear Time. In *Proc. of WADS*, 2005.

[26] E. Koutsoupias and C. H. Papadimitriou. Worst-Case Equilibria. In *Proc. of 16th STACS*, 1999.

[27] P. Krysta. Greedy Approximation via Duality for Packing, Combinatorial Auctions and Routing. In *Proc. of 30th MFCS*, 2005.

[28] R. Lavi, A. Mu'alem, and N. Nisan. Towards a Characterization of Truthful Combinatorial Auctions. In *Proc. of 44th FOCS*, 2003.

[29] R. Lavi and C. Swamy. Truthful and Near-Optimal Mechanism Design via Linear Programming. In *Proc. of 46th FOCS*, 2005.

[30] D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. of 1st EC*, 1999.

[31] N. Nisan and A. Ronen. Algorithmic Mechanism Design. In *Proc. of 31st STOC*, 1999.

[32] C. H. Papadimitriou. Algorithms, Games, and the Internet. In *Proc. of 33rd STOC*, 2001.

[33] C.H. Papadimitriou and M. Yannakakis. Optimization, Approximation and Complexity Classes. *Journal of Computer and System Sciences*, **43**, 425–440, 1991.

[34] P. Rusmevichientong. *A Non-Parametric Approach to Multi-Product Pricing: Theory and Application.* Ph.D. thesis, 2003.

A Independent sets and graph products

For a graph $G = (V, E)$, $|V| = n$, let $\alpha(G)$ refer to the size of a maximum independent set in G. Let $\mathcal{G}_a$ and $\mathcal{G}_b$ be two families of graphs with maximum degree bounded by 3 and $\alpha(G) \le an$ for $G \in \mathcal{G}_a$, $\alpha(G) \ge bn$ for $G \in \mathcal{G}_b$. As a direct consequence of the PCP theorem [6, 5, 33] one obtains:

PROPOSITION A.1. *There exist constants $0 < a < b < 1$, such that given $G \in \mathcal{G}_a \cup \mathcal{G}_b$ it is NP-hard to decide whether $G \in \mathcal{G}_a$ or $G \in \mathcal{G}_b$.*

The following is a standard concept that allows amplification of the above hardness.

DEFINITION 2. ([9, 3]) *Let $G = (V, E)$ be a graph and $k \in \mathbb{N}$. The k-fold* graph product *$G^k = (V^k, E_k)$ of G is defined by $V^k = V \times \cdots \times V$ and $\{(u_1, \ldots, u_k), (v_1, \ldots, v_k)\} \in E_k$ if and only if $\{u_1, \ldots, u_k, v_1, \ldots, v_k\}$ is not an independent set in G.*

Berman and Schnitger [9] and Blum [10] consider so-called *randomized graph products*, which are obtained as the subgraph induced by a random sample of the vertices of G^k. Alon et al. [3] show how this construction can be derandomized by replacing the sampling procedure of [9]. Given graph $G = (V, E)$, we construct a non-bipartite d-regular Ramanujan graph H, which has the same vertices as G and constant degree d that depends only on a and b. Vertices of the *derandomized graph product DG^k* are obtained by choosing a vertex of H uniformly at random and taking a random walk of length $k - 1$ starting at this vertex. For $k = O(\log n)$ the number nd^{k-1} of such random walks is polynomial and, thus, DG^k can be constructed deterministically in polynomial time. The edges of DG^k are defined as before. Now let dA be the (symmetric) adjacency matrix of H, where $\lambda_0 \ge \lambda_1 \ge \cdots \ge \lambda_{n-1}$ are eigenvalues of matrix A, and let $\lambda = \max\{\lambda_1, |\lambda_{n-1}|\}$. The following is a slightly simplified version of Theorem 1 of [3], which gives an upper and lower bound on the size of the maximum independent set in DG^k.

THEOREM A.1. ([3]) *For any given graph G and arbitrary $k \in \mathbb{N}$ it holds that*

$$\begin{aligned} & \alpha(G)d^{k-1}\left(\frac{\alpha(G)}{n} - \lambda\right)^{k-1} \\ \le \quad & \alpha(DG^k) \le \alpha(G)d^{k-1}\left(\frac{\alpha(G)}{n} + \lambda\right)^{k-1}. \end{aligned}$$

We now state a slightly extended version of Theorem 3 of [3]. We include the proof just to show that we can express the maximum degree of DG^k in terms of the number of its vertices.

Theorem 2.1 *For arbitrary non-decreasing functions $f : \mathbb{N} \longrightarrow \mathbb{R}_+$ with $f(n) \le n$ and $f(n^c) \le f(n)^c$ for all $c \ge 1$, $n \in \mathbb{N}$, let $\mathcal{G}_f$ be the family of graphs $G = (V, E)$, $|V| = n$, with maximum degree $\Delta = \mathcal{O}(f(n))$. There exists a constant $\varepsilon > 0$, such that it is NP-hard to approximate $\alpha(G)$ within $\mathcal{O}(f(n)^\varepsilon)$ for $G \in \mathcal{G}_f$.*

Proof. Let $\mathcal{G}_a$ and $\mathcal{G}_b$ be defined as above and let $G \in \mathcal{G}_a \cup \mathcal{G}_b$, $G = (V, E)$, $|V| = n$. Choosing $0 < a < b < 1$ appropriately it is NP-hard to decide whether $G \in \mathcal{G}_a$ or $G \in \mathcal{G}_b$ by Proposition A.1. We now consider the k-fold derandomized graph product $DG^k = (DV, DE)$.

By its construction we have that $|DV| = nd^{k-1}$. Let $(v_1, \ldots, v_k) \in DV$ and assume that there are indices i and j, such that $\{v_i, v_j\} \in E$. In this case it follows that $\{(v_1, \ldots, v_k), (w_1, \ldots, w_k)\} \in DE$ for all $(w_1, \ldots, w_k)$. Thus, DG^k contains a number of vertices of degree $nd^{k-1} - 1$. We define the modified graph $\widetilde{DG}^k = (\widetilde{DV}, \widetilde{DE})$ by removing all these vertices from DG^k. We observe that $\alpha(\widetilde{DG}^k) = \alpha(DG^k)$. By Theorem A.1 an independent set of size bn in G results in an independent set of size at least $bnd^{k-1}(b - \lambda)^{k-1}$ in DG^k. If less than this number of vertices are contained in $\widetilde{DE}^k$, it follows that $G \in \mathcal{G}_a$. Thus, w.l.o.g. we may assume that

$$bnd^{k-1}(b - \lambda)^{k-1} \le |\widetilde{DV}| \le nd^{k-1}.$$

In $\widetilde{DG}^k$ an edge $\{(v_1, \ldots, v_k), (w_1, \ldots, w_k)\}$ exists only if there are indices i and j, such that $\{v_i, w_j\} \in E$. We fix $(v_1, \ldots, v_k)$ and count the maximum number of adjacent vertices. There are k^2 possibilities to select i and j. Fixing indices fixes v_i as well and, by the fact that G has maximum degree 3, there are at most 3 possible choices for w_j. Finally, there remain d^{k-1} possibilities to choose the random walk generating $(w_1, \ldots, w_k)$. Thus, $\widetilde{DG}^k$ has maximum degree $\Delta \le 3k^2 d^{k-1}$.

For d-regular Ramanujan graphs it is known that $\lambda \approx 2\sqrt{d-1}/d$. By choosing the constant degree $d \ge 2$ of H sufficiently large we have that

$$\lambda < \frac{2}{\sqrt{d}} \le \frac{1}{3}(b - a).$$

By Theorem A.1 the gap between the cases that $G \in \mathcal{G}_a$ and $G \in \mathcal{G}_b$ is then amplified to

$$\frac{bnd^{k-1}(b-\lambda)^{k-1}}{and^{k-1}(a+\lambda)^{k-1}} \ge \left(\frac{b-\lambda}{a+\lambda}\right)^k > (1+\lambda)^k.$$

Using the fact that $d \approx 4/\lambda^2$ and choosing a constant γ, such that $(4/\lambda^2)^\gamma \approx (1 + \lambda)$, we obtain that $(1 + \lambda)^k \ge d^{\gamma k}$.

Given $G \in \mathcal{G}_a \cup \mathcal{G}_b$, $G = (V, E)$ and $|V| = n$, we choose (for the rest of this proof $\log$ is to the base of d)

$$k = c \log f(n)^{\delta} \quad \text{with} \quad c = (\log \frac{3}{2} b^{-1})^{-1}$$

for some $\delta \in (0, 1)$ such that $c\delta < 1$ (note that $\frac{3}{2}b^{-1} < d$ and $c > 1$). Thus, the number of vertices N of $\widetilde{DG}^k$ is lower bounded by

$$bnd^{k-1}(b - \lambda)^{k-1} \geq bn \frac{d^{k-1}}{(\frac{3}{2}b^{-1})^{k-1}} = \Omega(n),$$

where we use the fact that $\lambda \leq b/3$ and $\frac{3}{2}b^{-1} < d$. The maximum degree Δ of $\widetilde{DG}^k$ is upper bounded by $3(c \cdot \log f(n))^2 f(n)^{c\delta}$. Using that $\log^2 f(n) = o(f(n)^{1-c\delta})$ and the fact that f is non-decreasing we get that $\Delta = O(f(N))$. The gap between the cases $G \in \mathcal{G}_a$ and $G \in \mathcal{G}_b$ is amplified to

$$d^{\gamma k} = f(n)^{c\gamma\delta} \geq f(N)^{c\gamma\delta/2},$$

where we use that $n \geq \sqrt{N}$ and $f(\sqrt{N}) \geq \sqrt{f(N)}$ by our assumption. Choosing $\varepsilon = c\gamma\delta/2$ yields the claim. ☐

B The Missing Proofs

Theorem 3.3 UDP-MAX-NPL *with limited supply* 2 *or larger is APX-hard.*

Proof. We show an approximation preserving reduction from MAXCUT. It is known that MAXCUT is APX-hard even for graphs with maximum degree 3 (see, e.g., [6]). Let $G = (V, E)$ have such bounded degree. We transform G into an UDP-MAX-NPL instance as follows. For each vertex $v \in V$ we define 6 products and 6 consumers, both indexed by $v(0), \ldots, v(5)$, supply $s_{v(0)} = s_{v(2)} = s_{v(4)} = 2$, $s_{v(1)} = s_{v(3)} = s_{v(5)} = 1$ and budget values $b(c_{v(i)}, e_{v(i)}) = b(c_{v(i)}, e_{v(i+1)}) = 1$ for $i \in \{0, 2, 4\}$, $b(c_{v(i)}, e_{v(i)}) = b(c_{v(i)}, e_{v(i+1 \mod 6)}) = 2$ for $i \in \{1, 3, 5\}$. Budgets that are not specified are assumed to be 0. Each edge $e = \{v, w\} \in E$ can now be associated with unique products $e_{v(i)}$ and $e_{w(j)}$, where $i, j \in \{0, 2, 4\}$ and every product is associated with at most one edge. For edge e we define 2 consumers $c_{e(0)}$ and $c_{e(1)}$ with budgets $b(c_{e(0)}, e_{v(i)}) = b(c_{e(0)}, e_{w(j)}) = 1$, $b(c_{e(1)}, e_{v(i)}) = b(c_{e(1)}, e_{w(j)}) = 2$. This construction is depicted in Figure 4.

We start by stating some facts about the solution that an approximation algorithm for our pricing problem might return on this instance. First, we observe that we can w.l.o.g. assume that all prices in this solution are from $\{1, 2\}$, since prices above 2 cannot result in any revenue and prices below 2 can always be increased up to the next budget value. The second important observation is that for all vertices v from G we can w.l.o.g. assume that products $e_{v(0)}, e_{v(2)}, e_{v(4)}$ are assigned the same price, i.e., $p(e_{v(0)}) = p(e_{v(2)}) = p(e_{v(4)})$ for all $v \in V$. We show how any solution in which this is not the case can easily be turned into a solution of no smaller value, such that our assumption holds. For reasons of symmetry it is sufficient to consider the case that product $e_{v(0)}$ has been assigned the wrong price.

First, assume that $p(e_{v(0)}) = 2$, $p(e_{v(2)}) = p(e_{v(4)}) = 1$. In this situation, if $p(e_{v(1)}) = 2$, consumer $c_{v(0)}$ currently cannot afford to buy any product, resulting in revenue 0 from this consumer. If $p(e_{v(1)}) = 1$, then consumer $c_{v(1)}$ currently buys at price at most 1. In both cases, the revenue generated by consumers $c_{v(0)}, \ldots, c_{v(5)}$ is at most 8. By setting $p(e_{v(0)}) = p(e_{v(2)}) = p(e_{v(4)}) = 1$, $p(e_{v(1)}) = p(e_{v(3)}) = p(e_{v(5)}) = 2$ and $a(c_v(i)) = e_{v(i)}$ for all i this revenue increases to 9. On the other hand, if product $e_{v(0)}$ is associated with some edge e, only 1 consumer from $\{c_{e(0)}, c_{e(1)}\}$ can afford product $e_{v(0)}$ at price 2 and, thus, might be buying it. Revenue from this consumer decreases by no more than 1. Hence, we have transformed our solution without decreasing the overall revenue.

For the second case, let $p(e_{v(0)}) = 1$, $p(e_{v(2)}) = p(e_{v(4)}) = 2$. If $p(e_{v(5)}) = 2$, consumer $c_{v(4)}$ cannot afford any product. If $p(e_{v(5)}) = 1$, consumer $c_{v(5)}$ buys at price at most 1. Again setting $p(e_{v(0)}) = p(e_{v(2)}) = p(e_{v(4)}) = 2$, $p(e_{v(1)}) = p(e_{v(3)}) = p(e_{v(5)}) = 1$ and $a(c_v(i)) = e_{v(i+1 \mod 6)}$ makes overall revenue from consumers $c_{v(0)}, \ldots, c_{v(5)}$ increase by 1. On consumers $\{c_{e(0)}, c_{e(1)}\}$ revenue decreases by at most 1, because consumer $c_{e(1)}$ can still buy a product at price 2 after $p(e_{v(0)})$ is changed. This gives the above claim.

We now argue how any small constant factor approximation on the constructed problem instance yields a corresponding approximation for the MAXCUT problem. As we have seen we obtain solutions with prices in $\{1, 2\}$, $p(e_{v(0)}) = p(e_{v(2)}) = p(e_{v(4)})$, $p(e_{v(1)}) = p(e_{v(3)}) = p(e_{v(5)})$ and a corresponding allocation a for all $v \in V$. Thus, overall revenue from consumers $c_{v(0)}, \ldots, c_{v(5)}$ is exactly 9 for all $v \in V$. For consumers $\{c_{e(0)}, c_{e(1)}\}$ belonging to some edge $e = \{v, w\}$ it is simple to find the optimal allocation given prices $p(e_{v(i)})$, $p(e_{w(j)})$ of the corresponding products. If $p(e_{v(i)}) = p(e_{w(j)}) = 1$ then we can set $a(c_{e(0)}) = e_{v(i)}$, $a(c_{e(1)}) = e_{w(j)}$. If $p(e_{v(i)}) = p(e_{w(j)}) = 2$ then we let $a(c_{e(0)}) = \varnothing$, $a(c_{e(1)}) = e_{v(i)}$. If $p(e_{v(i)}) = 1$, $p(e_{w(j)}) = 2$ we define $a(c_{e(0)}) = e_{v(i)}$, $a(c_{e(1)}) = e_{w(j)}$. Thus, total revenue from consumers $c_{e(0)}$ and $c_{e(1)}$ is 2 if $p(e_{v(i)}) = p(e_{w(j)})$ and 3 if $p(e_{v(i)}) \neq p(e_{w(j)})$. We can then write the value of any such solution to UDP-MAX-NPL as $9n + 2m + c$, where $n = |V|$, $m = |E|$ and c is the number of edges $\{v, w\}$ such that $p(e_{v(0)}) \neq p(e_{w(0)})$. Given this solution we can immediately define a cut (S, T) of size c in G by setting $S = \{v \mid p(e_{v(0)}) = 1\}$, $T = V \backslash S$.

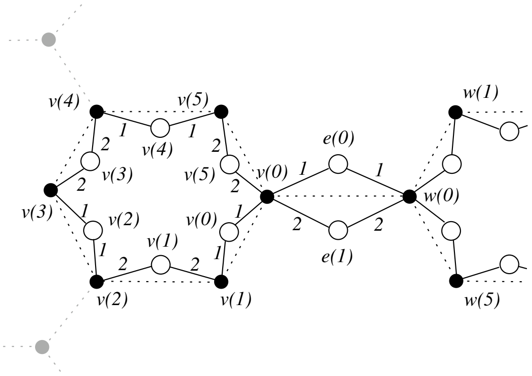

Figure 4: Construction from the proof of Theorem 3.3. Consumers are depicted as circles, products as points. Edges between consumers and products are labelled with the respective non-zero budgets.

Hence, the optimal solution on our pricing instance has value $9n + 2m + c^*$, where c^* is the size of a maximum cut in G. Assume now that we can obtain a $(1 - \varepsilon)$-approximation to the pricing problem. By $n \leq m$ (assuming G is not a tree) and $c^* \geq m/2$ we have

$$(1 - \varepsilon) \leq \frac{9n + 2m + c}{9n + 2m + c^*} \leq \frac{22c^* + c}{23c^*}$$

and, thus, $c/c^* \geq (1 - 23\varepsilon)$. Choosing ε appropriately small yields any arbitrarily small constant approximation ratio for MAXCUT. □

LEMMA B.1. *The analysis of algorithm* LOCALSEARCH *in Theorem 3.4 is tight.*

Proof. Consider a problem instance with 2 products indexed by $\mathcal{P} = \{1, 2\}$ and $k + 1$ consumers indexed by $\mathcal{C} = \{1, \ldots, k + 1\}$. Customers' budgets are $b(1, 1) = k$, $b(1, 2) = k - \varepsilon$, $b(2, 1) = 0$, $b(2, 2) = \varepsilon$ and $b(i, 1) = 1$, $b(i, 2) = 0$ for $i = 3, \ldots, k+1$. We assume that products are available in unlimited supply. It is straightforward to verify that prices $p(1) = k$, $p(2) = \varepsilon$ are locally optimal and result in revenue $k + \varepsilon$. Prices $p(1) = 1$, $p(2) = k - \varepsilon$, however, result in overall revenue of $2k - 1 - \varepsilon$. Choosing k and ε appropriately shows that a pure local search approach cannot give any approximation ratio better than 2. □

Before we present formal proofs of the results about the unit-demand pricing game let us introduce some notation to describe *mixed strategies* in more detail. Let $\mathcal{P} = \{1, \ldots, n\}$ be a set of players. Each player j needs to assign a price p_j to her product e_j, such as to maximize her revenue from sales to consumers $\mathcal{C}$. Allowing mixed strategies, every player defines a probability distribution P_j over a set of possible prices for her product e_j. For ease of notation we let $P = (P_1, \ldots, P_n)$, $P_{-j} = (P_1, \ldots, P_{j-1}, P_{j+1}, \ldots, P_n)$ and $(P_{-j}, P_j) = P$. Observe that we can w.l.o.g. allow only the budget values as possible prices and, thus, P_j is a discrete distribution. Since every set of fixed prices defines an optimal allocation, the distributions P_j define a probability distribution also over the set of allocations. We define R_j to be the random variable that describes the revenue of player j. A set of strategies $P^{eq} = (P_1^{eq}, \ldots, P_n^{eq})$ are at *Nash equilibrium*, if for every player j we have that

$$\mathrm{E}_{P^{eq}}\left[\, R_j \,\right] \geq \mathrm{E}_{(P_{-j}^{eq}, P_j')}\left[\, R_j \,\right] \quad \forall\, P_j' \neq P_j^{eq},$$

i.e., if no player can increase her expected revenue by changing her current strategy P_j^{eq}. Let prices $p_1^*, \ldots, p_n^*$ and allocation a^* be an optimal (i.e., revenue maximizing) solution to UDP-MAX-NPL. We let $C_j = (a^*)^{-1}(e_j)$ refer to the set of consumers that buy product e_j in this solution and define $L_j = \{c_i \in C_j \,|\, p_{a(c_i)} < p_j^*\}$, $H_j = C_j \backslash L_j$. For the remainder of this section it will be convenient to refer to players, their products and consumers only by their indices.

LEMMA B.2. *Consider a set of prices $p_1, \ldots, p_n$ with (optimal) allocation a and let $|L_j| = t$. If price p_j is changed to p_j^* and we recompute the optimal allocation b we have that $|b^{-1}(j)| \geq t$.*

Proof. Throughout this proof, set L_j is defined with respect to prices $p_1, \ldots, p_n$ and allocation a. Let us assume now that $|b^{-1}(j)| < t$. Clearly, there can be no consumer $i \in C_j$ with $p_{b(i)} < p_j^*$, since allocation b is chosen optimally and there are available copies of product j left unsold. It follows that there must exist a consumer $i_0 \in L_j$ with $b(i_0) \neq j$ and $p_{b(i_0)} \geq p_j^*$. Under this assumption we will show that allocation b is not optimal. The following chain of conclusions follows solely from the optimality of a. Since $p_{a(i_0)} < p_{b(i_0)}$ it must be the case that product $b(i_0)$ is sold out under allocation a, i.e., $|a^{-1}(b(i_0))| = s_{b(i_0)}$. Then there must be some consumer i_1 with $b(i_1) \neq a(i_1) = b(i_0)$. For this consumer it must be true that either $p_{b(i_1)} \leq p_{a(i_0)}$ (including the case that $b(i_1) = \varnothing$) or product $b(i_1)$ is sold out under a. Otherwise, modifying a by setting $a(i_0) = b(i_0)$ and $a(i_1) = b(i_1)$ would result in a feasible allocation with strictly higher revenue. By repeatedly applying this argument we obtain a chain $i_0, i_1, \ldots, i_s$ of consumers with $b(i_k) = a(i_{k+1})$ and $p_{b(i_s)} \leq p_{a(i_0)}$ (or $b(i_s) = \varnothing$). We can assume that $b(i_k) \neq j$ for all k. To see this, note, that otherwise we could for every consumer $i_0 \in L_j$ with $b(i_0) \neq j$ find a distinct consumer i_k with $b(i_k) = j$, which would in turn imply that $|b^{-1}(j)| \geq t$. The above argument is also depicted in Figure 5. We can define a feasible allocation c by going backwards along the constructed chain of consumers and allocating to each consumer the product she received under allocation a except for consumer i_0, who

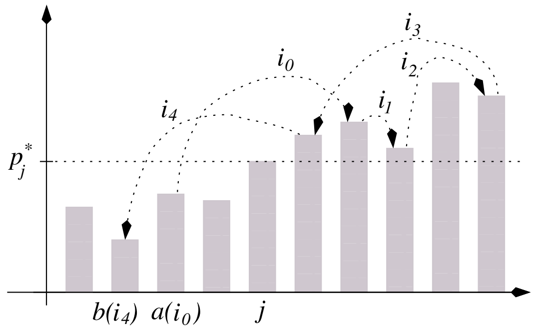

Figure 5: A chain of consumers switching to new products as in the proof of Lemma B.2, where $i_0 \in L_j$.

will now receive product j. Formally, we let $c(i_k) = a(i_k)$ for $k = 1, \ldots, s$, $c(i_0) = j$ and $c(i) = b(i)$ for all remaining consumers. We observe that

$$\begin{aligned} \sum_{k=0}^{s} p_{c(i_k)} &= p_j^* + \sum_{k=1}^{s} p_{c(i_k)} = p_j^* + \sum_{k=1}^{s} p_{a(i_k)} \\ &= p_j^* + \sum_{k=0}^{s-1} p_{b(i_k)} > \sum_{k=0}^{s} p_{b(i_k)}, \end{aligned}$$

where the last inequality follows from $p_{b(i_s)} \le p_{a(i_0)} < p_j^*$, since $i_0 \in L_j$. This contradicts the optimality of allocation b and, thus, finishes the proof. □

Theorem 3.6 *The price of anarchy in the unit-demand max-buying pricing game is* 2.

Proof. Let strategies $P^{eq} = (P_1^{eq}, \ldots, P_n^{eq})$ define a Nash equilibrium. We want to lower bound the expected revenue of agent j. We define a (deterministic) strategy P_j^* for agent j by $\Pr(p_j = p_j^*) = 1$ and let $P^* = (P_{-j}^{eq}, P_j^*)$ denote the modified set of strategies. By the definition of Nash equilibria we have that

$$\mathrm{E}_{(P_{-j}^{eq}, P_j^*)}\big[\, R_j \,\big] \le \mathrm{E}_{P^{eq}}\big[\, R_j \,\big].$$

By Lemma B.2 we can lower bound the expected revenue of agent j playing strategy P_j^* by

$$\mathrm{E}_{(P_{-j}^{eq}, P_j^*)}\big[\, R_j \,\big] \ge \sum_{t=0}^{|C_j|} t \cdot p_j^* \cdot \Pr_{P^{eq}}\big(|L_j| = t\big).$$

We can then write that

$$\begin{aligned} & \mathrm{E}_{P^{eq}}\big[\, R_j \,\big] + \mathrm{E}_{P^{eq}}\Big[\, \sum_{i \in C_j} p_{a(i)} \,\Big] \\ \ge\ & \mathrm{E}_{(P_{-j}^{eq}, P_j^*)}\big[\, R_j \,\big] + \mathrm{E}_{P^{eq}}\Big[\, \sum_{i \in C_j} p_{a(i)} \,\Big] \\ \ge\ & \sum_{t=0}^{|C_j|} t \cdot p_j^* \cdot \Pr_{P^{eq}}\big(|L_j| = t\big) \\ & \quad + \sum_{t=0}^{|C_j|} t \cdot p_j^* \cdot \Pr_{P^{eq}}\big(|H_j| = t\big) \\ =\ & \sum_{t=0}^{|C_j|} \Pr_{P^{eq}}\big(|L_j| = t\big) \cdot p_j^* \cdot |C_j| = p_j^* \cdot |C_j|, \end{aligned}$$

where we use the fact that

$$\Pr_{P^{eq}}\big(|H_j| = t\big) = \Pr_{P^{eq}}\big(|L_j| = |C_j| - t\big).$$

Let R denote the expected revenue of the equilibrium state, R_{opt} the revenue generated by the optimal solution. By using linearity of expectation we have that

$$\begin{aligned} 2 \cdot \mathrm{E}_{P^{eq}}[R] &= \sum_{j \in \mathcal{P}} \mathrm{E}_{P^{eq}}\big[\, R_j \,\big] + \mathrm{E}_{P^{eq}}\Big[\sum_{i \in \mathcal{C}} p_{a(i)}\Big] \\ &= \sum_{j \in \mathcal{P}} \Big(\mathrm{E}_{P^{eq}}\big[\, R_j \,\big] + \mathrm{E}_{P^{eq}}\Big[\, \sum_{i \in C_j} p_{a(i)} \,\Big] \Big) \\ &\ge \sum_{j \in \mathcal{P}} p_j^* \cdot |C_j| = R_{opt}. \end{aligned}$$

This gives the desired upper bound on the price of anarchy. We now give a simple corresponding lower bound. Consider a problem instance with 2 products $\mathcal{P} = \{1, 2\}$ each of which is available only once, i.e., $s_1 = s_2 = 1$, and 2 consumers $\mathcal{C} = \{1, 2\}$ with budgets $b(1,1) = \varepsilon$, $b(1,2) = 1$, $b(2,1) = 1$ and $b(2,2) = 1 + \varepsilon$. It is easy to see that the optimal solution generates revenue 2, while the pure strategies $p_1 = \varepsilon$ and $p_2 = 1 + \varepsilon$ define a Nash equilibrium which results in overall revenue $1 + 2\varepsilon$. Thus, the above bound is tight. □