

On Fair Price Discrimination in Multi-Unit Markets

Michele Flammini^{1,2}, Manuel Mauro¹, Matteo Tonelli¹

¹ Gran Sasso Science Institute, L'Aquila, Italy

² University of L'Aquila, L'Aquila, Italy

michele.flammini@univaq.it, manuel.mauro@gssi.it, matteo.tonelli@gssi.it

Abstract

Discriminatory pricing policies, even if at first glance can be perceived as unfair, are widespread. In fact, pricing differences for the same item among different national markets are common, or forms of discrimination based on the time of purchase, like in tickets' sales. In this work we propose a framework for capturing the setting of "fair" discriminatory pricing and study its application to multi-unit markets, in which many copies of the same item are on sale. Our model is able to incorporate the fundamental discrimination settings proposed in the literature, by expressing individual buyers constraints for assigning prices by means of a social relationship graph, modeling the information that each buyer can acquire about the prices assigned to the other buyers. After pointing out the positive effects of fair price discrimination, we investigate the computational complexity of maximizing the social welfare and the revenue in these markets, providing hardness and approximation results under various assumptions on the buyers valuations and on the social graph topology.

1 Introduction

Charging customers different prices for the same item can be perceived as unfair at first, but it is actually widespread. As an example, when buying flight tickets, customers who arrive earlier usually get a better price. In an opposite way, the price of electronic devices like computers and cameras decreases over time. Beside this form of *temporal* discrimination, *spatial* discrimination is also consolidated in the sale of the same good with varying prices at different sites, like in different countries or shops. The lack of knowledge of the prices offered at different locations is often modeled in the literature as a *cost of search*, that is the cost that buyers have to pay in order to get information about the different prices at the various sites. Exploiting this cost, a firm is able to discriminate on subsets of buyers on the base of their mutual distances. A further type of discrimination occurs in users profiling, when premium customers such as frequent-flyers are charged lower prices with respect to the general public, or

when some classes of buyers are recognized to be eligible for a better treatment, such as age discounts or financial aids.

Nice introductions on price discrimination together with the main arising scenarios and the related motivation can be found in [Varian, 1989; Philips, 1983; Stokey, 1979; Laffont *et al.*, 1998]. An interesting classification is due to A. Pigou in [Pigou, 1920], who detected the following three classes:

- *First-degree* or *perfect* price discrimination, where the seller is admitted to set a different price for each unit of good, so as to perfectly adapt to each buyer's willingness to buy the item.
- *Second-degree* or *non-linear* pricing, also known as bundle pricing, in which the seller is able to set a price based on the amount of items purchased by a single buyer, e.g. a smaller price per item for buyers who buy a bigger amount.
- *Third-degree* price discrimination, in which different buyers may be charged different linear prices, that is they are asked a fixed price per item which may be different from the one of the other buyers. Examples are student discounts or charging different prices in different days of the week.

The problem of maximizing the seller's revenue in combinatorial markets under the notion of envy-freeness has been widely investigated. Important works in the field are [Guruswami *et al.*, 2005; Hartline and Yan, 2011; Cheung and Swamy, 2008; Balcan *et al.*, 2008; Briest and Krysta, 2006], in which the authors designed logarithmic approximation algorithms under different assumptions. On the other hand, related hardness results were given by [Briest, 2008; Chalermsook *et al.*, 2012; 2013b; 2013a; Demaine *et al.*, 2008]. Further variants were considered by [Chen and Deng, 2010; Chen *et al.*, 2011; Feldman *et al.*, 2012; Anshelevich *et al.*, 2017; Bilbø *et al.*, 2017; Chen *et al.*, 2016].

All these works did not consider discriminatory policies, or are related only to the second-degree discrimination, which is a somewhat weaker form of discrimination, as each bundle must be offered at the same (even if not linear) price to all the buyers in the market. Therefore, in this paper we focus on the more proper forms of discrimination related to the first- and third-degree scenarios. In particular, we introduce a unifying framework for "fair" discrimination pricing policies, which takes into account individual constraints among the buyers induced by their social relationships or by the information they can acquire about the other buyers. Namely, an arc from buyer i to buyer k indicates that i knows k or has access to the same buying strategy without paying any search

cost, for instance buying at the same site. Then, in order for i to be satisfied or feel to be treated fairly, i can only be offered the same or a lower price for the same good with respect to k . Our framework is very general, as by modifying the structure of the graph is in turn able to incorporate in a unified setting all the fundamental “fair” price discrimination forms proposed in the literature, including the first- and third-degree ones. For instance, a directed chain of arcs can express temporal dependencies, undirected social graphs can model in a nice way spatial discrimination through the submarkets induced by their connected components, and directed trees of strongly connected components the hierarchical profiling of classes of users.

Under the above setting, we investigate the case of a single item on sale in a limited amount of copies, which is usually referred to as a multi-unit market and represents many real world scenarios (e.g. commodity markets) [Feldman *et al.*, 2012; Brânzei *et al.*, 2016; Monaco *et al.*, 2015]. More precisely, we have a single seller who is willing to sell m identical copies to n buyers, who are interested in buying one or more copies. Buyers have a valuation function which describes how much they are willing to spend for each given number of copies, and resorting on these valuations the seller must decide what is the price per copy and the allocation. The multi-unit scenario considered by the above papers allows only to assign the same price per copy to all buyers, or in case of bundle pricing different non-linear prices to bundles containing a certain number of identical items. In any case, the same bundle must be offered at the same price to all the buyers.

1.1 Our Contribution

In this paper, we study the problem of maximizing the social welfare or the seller’s revenue by means of fair price discrimination in multi-unit markets according to the above framework. Namely, the price per copy of each buyer must be at most the lowest proposed to the buyers she knows directly. We investigate four different cases arising by considering *i.* social welfare and revenue maximization, and *ii.* single-minded and general valuations, that is when buyers are interested only in a specific amount of item copies, or do not have any restriction. Moreover, we consider different relevant social graph topologies.

We first point out that the gain from discrimination can be considerable. In fact, it can increase the seller’s revenue or the social welfare up to a multiplicative factor equal to the number of item copies. We then provide polynomial time algorithms or hardness and approximation results in all the above mentioned cases.

More precisely, for single minded buyers, we show that the problems of maximizing the social welfare and the revenue are strongly NP-Hard. We then provide a polynomial time 2-approximation algorithm for the social welfare maximization and an $O(\log n)$ -approximation one for the revenue. Finally, still for the revenue maximization, we give an improved $(2 + \epsilon)$ -approximation for the case of undirected social graphs.

For general valuations, we first provide a general construction able to reduce the corresponding instances into “nearly” equivalent single-minded ones. More precisely, the solutions

provided by our single minded approximation algorithms for the reduced instances can be transformed into solutions for the initial instances without performance losses. Therefore, we again achieve a 2-approximation for the social welfare maximization and an $O(\log n + \log m)$ one for the revenue, $n \cdot m$ being the number of buyers in the reduced instances. We remark that our reduction is also of independent interest, as it might be exploited in other settings for extending results from single-minded to general multi-unit instances.

Finally, we provide polynomial time optimal algorithms for all the realistic cases mentioned before, that is for undirected social graphs, graphs with a constant number of strongly connected components, and arborescences or directed trees of strongly connected components.

1.2 Related Work

Envy-freeness in multi-unit markets has already been investigated in [Feldman *et al.*, 2012], where the authors studied the case of budgeted buyers and gave a hardness result together with a constant approximation algorithm for the problem of finding an allocation of optimal revenue. Later on [Monaco *et al.*, 2015; Brânzei *et al.*, 2016] investigated the problem with slightly different assumptions, providing several hardness and approximation results. In particular, [Monaco *et al.*, 2015] considered revenue maximization in multi-unit markets under both the notion of *envy-freeness* and *pair envy-freeness*, both for item- and bundle-pricing. In this work we refer to the notion of *envy-freeness* given in [Guruswami *et al.*, 2005].

The social influence over the set of buyers of a market has been recently studied in [Abebe *et al.*, 2017; Chevaleyre *et al.*, 2007; Flammini *et al.*, 2018], where the authors generalized the concept of pair envy-freeness in the context of limited mutual knowledge of the buyers, restricting the pair envy-freeness constraints only to social neighbors. These works applied such a notion of social (or local) envy-freeness to the problems of cake cutting, distributed negotiation, and multi-unit markets. Notice that, despite the apparent similarity of our model with the one in [Flammini *et al.*, 2018], which considers multi-unit markets with social graphs, even if some of our techniques are inspired to such a paper, our results refer to a completely different setting, as sociality influences discrimination of prices and not relaxations of the envy-freeness constraints. Finally, [Bilò *et al.*, 2017; Bilò *et al.*, 2018] considered the possibility of limiting the view of some buyers, not admitting them in the market. This concept is somewhat related to sociality, as it translates in limiting the social awareness of buyers acting on the social graph topology. In all the above cases, social graphs have been exploited to induce relaxed notions of envy-freeness.

Forms of price discrimination induced by social relationships have been considered in [Alon *et al.*, 2013], where buyers who are neighbors in the social graph demand a difference of their proposed prices bounded by a parameter α . Several results on revenue maximization are provided under different assumptions, such as the possibility of excluding some buyers. Furthermore, [Amanatidis *et al.*, 2016] extended such results to a generalized model having a non empty overlap with the multi-unit markets with single-minded valuations herein considered. However, the two settings are different, since on

the one hand we consider additional supply and envy-freeness constraints, on the other hand they admit a certain tolerance in the discrimination due to the parameter α . Moreover, we also investigate the case of general valuations.

Finally, [Feldman *et al.*, 2012] mentioned the possibility of considering, besides bundling, other forms of “unfair” price discrimination, together with their potential positive effects on revenue maximization. Although the topic has been widely addressed in economics [Phlips, 1983; Stokey, 1979; Laffont *et al.*, 1998], we believe that fair discrimination policies induced by social relationships have not been sufficiently investigated in the context of envy-free pricing and are the focus of the present work.

2 Preliminaries

A multi-unit market η can be represented by a tuple $(N, M, (v_i)_{i \in N})$, where $N = \{1, \dots, n\}$ is a set of n buyers, M is a set of m identical items and for every buyer $i \in N$, $v_i = (v_i(1), \dots, v_i(m))$ is a valuation function or vector which expresses, given a subset of items $X \subseteq M$ of size j , the amount of money $v_i(j) \in \mathbb{R}$ that buyer i is willing to pay for X . We assume that $v_i(0) = 0$ and $v_i(j) \geq 0$ for every j , $1 \leq j \leq m$ and buyer $i \in N$.

We distinguish the following two different cases, according to the imposed restrictions on the valuation functions: *single-minded*, with buyers interested only in a certain amount of items and thus having positive valuation only for bundles of such a size, and *general valuations*, i.e., the unrestricted case. In the sequel, when dealing with single-minded valuations, we will denote by m_i the unique *preferred* bundle size for which buyer i has a strictly positive valuation.

Since items in M are identical, we consider the pricing scheme called *item-pricing*, in which the seller must assign a single non-negative price per item $p_i \in \mathbb{R}$ to each buyer i . The amount owed by i for a bundle of items X is thus $p_i \cdot |X|$, so that her utility for receiving X is $u_i(X, p_i) = v_i(|X|) - p_i \cdot |X|$. In the sequel we will denote by $\bar{p} = (p_1, \dots, p_n)$ the *price vector* or n -tuple such that for every i , $1 \leq i \leq n$, $p_i \geq 0$ is the price per item assigned to buyer i .

We assume that buyers in N are individuals of a population, whose relationships are represented by an underlying directed social graph $G = (N, E)$. Each buyer $i \in N$ is aware only of the prices proposed to the other buyers she knows directly, that is belonging to the subset $N(i) = \{k \in N \mid (i, k) \in E\}$ of her neighbors in G . Under such an assumption, it is possible to define the following notion of fair price discrimination.

Definition 1. A price vector \bar{p} is *fair* with respect to the social graph $G = (N, E)$ if $p_i \leq p_k$ for every $(i, k) \in E$.

An *allocation vector* is an n -tuple $\bar{X} = (X_1, \dots, X_n)$ such that $X_i \subseteq M$ is the set of items sold to buyer i .

A *feasible* outcome for market η is a pair (\bar{X}, \bar{p}) satisfying the following supply constraint: $\sum_{i=1}^n |X_i| \leq m$.

We say that a feasible outcome (\bar{X}, \bar{p}) for market η is *envy-free* if $X_i \in \operatorname{argmax}_{X \subseteq M} u_i(X, p_i)$ for every buyer $i \in N$. Notice that, for every $i \in N$, since $v_i(0) = 0$, envy-freeness implies the classical assumption of individual rationality of the buyers, that is $u_i(X_i, p_i) \geq 0$.

Definition 2. A feasible outcome (\bar{X}, \bar{p}) is *fair* under the social graph G if it is envy-free and its price vector is fair with respect to G .

Given an outcome (\bar{X}, \bar{p}) , the achieved social welfare is $sw(\bar{X}, \bar{p}) = \sum_{i=1}^n v_i(|X_i|)$ and the revenue raised by the seller is $r(\bar{X}, \bar{p}) = \sum_{i=1}^n p_i \cdot |X_i|$. The (fair) *pricing problem* consists in determining a feasible outcome (\bar{X}, \bar{p}) for market η , which is fair under G and maximize social welfare or revenue. In the following we will denote by $opt_{sw}(\eta, G)$ (resp. $opt_r(\eta, G)$) the maximum possible social welfare (resp. revenue) achievable by an outcome for η fair under G , and $opt_{sw}(\eta)$ (resp. $opt_r(\eta)$) be the highest possible one achievable without price discrimination, or analogously by an outcome fair under the complete social graph. Clearly, by the individual rationality constraint, in any feasible outcome (\bar{X}, \bar{p}) , $sw(\bar{X}, \bar{p}) \geq r(\bar{X}, \bar{p})$, so that also $opt_{sw}(\eta) \geq opt_r(\eta)$ and $opt_{sw}(\eta, G) \geq opt_r(\eta, G)$.

For the sake of brevity, we call (SINGLE,WELFARE)-pricing (resp. (GENERAL,WELFARE)-, (SINGLE,REVENUE)- and (GENERAL,REVENUE)-pricing) the pricing problem restricted to the instances of multi-unit markets with single-minded valuations and social welfare maximization (resp. general valuations and social welfare maximization, single-minded and revenue maximization, and general valuations and revenue maximization).

We will often reduce the pricing problem to a variant of the MULTIPLE-CHOICE KNAPSACK problem in which we are given ℓ classes $\{O_1, \dots, O_\ell\}$ of objects to pack in a knapsack of capacity k . Each object $o_{h,j} \in O_h$ has a profit $z_{h,j}$ and a size $w_{h,j}$, and we must pick at most one object from each class so as to maximize the sum of the profits of the selected objects without exceeding the knapsack capacity k . As shown in [Lawler, 1979], the problem is NP-hard, but it admits an FPTAS. MULTIPLE-CHOICE KNAPSACK has already been exploited by [Monaco *et al.*, 2015; Brânzei *et al.*, 2016; Flammini *et al.*, 2018] in order to provide optimal algorithms or efficient approximations in different scenarios concerning multi-unit markets. Some of our algorithms are inspired to such constructions.

Before concluding the section, let us remark that discrimination of prices is a particularly worth investigating issue, as it can increase the achievable social welfare and revenue up to an m multiplicative factor. This holds even for the very simple case of two single-minded buyers, as shown by the following example: buyer 1 with preferred size 1 and valuation $v_1(1) = 1 + \varepsilon$ for small ε , and buyer 2 with preferred size m and valuation $v_2(m) = m$. Without discrimination the only possibility is selling a bundle of size 1 to buyer 1 at price at most $1 + \varepsilon$, yielding social welfare and revenue at most $1 + \varepsilon$. With price discrimination it is possible to assign a bundle of size m to buyer 2 at price 1 per item, and no item to buyer 1 asking her price at least $1 + \varepsilon$. This provides social welfare and revenue m . Notice that this is strict, because it is always possible to achieve an m fraction of the optimum with discrimination, just assigning a bundle only to the buyer having the highest possible valuation per item, that is maximizing the ratio $v_i(m_i)/m_i$.

Let us also finally stress that in multi-unit markets, while

the size of the representation of instances with general valuations is polynomial in m , as different valuations must be specified for different bundle sizes, in single-minded instances the dependence is logarithmic in m , as for each buyer it is sufficient to specify the size of her unique preferred bundle, together with the corresponding valuation.

3 Single-Minded Valuations

We first show the following hardness results.

Theorem 3. *The (SINGLE,WELFARE)- and (SINGLE,REVENUE)-pricing problems are strongly NP-hard.*

Proof. We show the claim by reduction from the DENSEST K-SUBGRAPH problem in which, given an undirected graph $H = (V, F)$ and an integer k , we are interested in determining a subset $S \subseteq V$ with $|S| \leq k$ that maximizes the number of edges in the subgraph induced by S . Given an instance $(H = (V, F), k)$ of DENSEST K-SUBGRAPH, the reduced instance is constructed as follows: a buyer i_v for each $v \in V$ and a buyer i_e for each $e \in F$; arcs (i_u, i_e) and (i_v, i_e) in the social graph G for every $e = \{u, v\} \in F$; number of items $m = k(2|F| + 1) + |F|$; finally, non-null valuations are $v_{i_v}(2|F| + 1) = 2|F| + 1 + \varepsilon$ for each $v \in V$, with ε suitably small, and $v_{i_e}(1) = 1$ for every $e \in F$.

By construction, any fair outcome for the reduced instance that maximizes the social welfare (or revenue) must allocate k buyers i_v with $v \in V$, plus the highest possible number of buyers i_e with $e \in F$. Moreover, since in any fair outcome if any buyer i_e with $e = \{u, v\} \in F$ is allocated then also i_u and i_v must receive their preferred bundles, it is immediate to see that there exists a fair outcome of social welfare (resp. revenue) $k(2|F| + 1) + l$ if and only if there exists $S \subseteq V$ with $|S| \leq k$ inducing a subgraph of l edges in H , hence the claim. \square

However, the problem can be suitably approximated.

Theorem 4. *There exists a 2-approximation algorithm for (SINGLE,WELFARE)-pricing.*

Proof. In the claimed algorithm we first execute the following *pre-processing phase*. Let G' be the transitive closure of G and $N^>(i)$ be the set of buyers k such that $v_k(m_k)/m_k > v_i(m_i)/m_i$ and there exists an arc (k, i) in G' . Then, since there is a direct path from k to i in G and thus in any fair price vector the price assigned to k must be at most the one of i , if i has non-negative utility for her preferred bundle size m_i , then k has strictly positive utility for hers. Therefore, since in any envy-free allocation all the buyers with strictly positive utility must be allocated the corresponding items, if buyer i gets her preferred bundle, then also every $k \in N^>(i)$ must. This implies that, if $m_i + \sum_{k \in N^>(i)} m_k > m$, i can not be allocated in any fair outcome. Therefore, all such buyers are discarded in the pre-processing phase, and this can be clearly accomplished in polynomial time. The residual market η' thus preserves the optimal solutions of η and satisfies $m_i + \sum_{k \in N^>(i)} m_k \leq m$ for all the remaining buyers i .

Assume without loss of generality that buyers of η' are listed in non increasing order with respect to the ratios

$v_i(m_i)/m_i$, that is such that $v_i(m_i)/m_i \geq v_k(m_k)/m_k$ if $i < k$. Let k be the first buyer in the order such that $\sum_{i=1}^k m_i \geq m$. It is immediate to see that $\sum_{i=1}^k v_i(m_i)$ upper bounds the value $opt_{sw}(\eta', G)$ of the optimal solution for η' . Consider then the following two fair outcomes:

- i. (\bar{X}, \bar{p}) , where bundles X_i of size m_i are allocated only to buyers $i < k$ at price $p_i = v_{k-1}(m_{k-1})/m_{k-1}$ per item, while $p_i = \infty$ for $i \geq k$;
- ii. (\bar{X}', \bar{p}') , where bundles X_i of size m_i are allocated only to buyers in $\{k\} \cup N^>(k)$ at price $p_i = v_k(m_k)/m_k$ per item, while $p_i = \infty$ for all the remaining buyers.

Both solutions can be computed in polynomial time and by construction are fair under G also for the initial market η . It remains to show that returning the best of the two outcomes provides the claimed 2-approximation. In fact, the achieved social welfare is equal to $\max\{sw(\bar{X}, \bar{p}), sw(\bar{X}', \bar{p}')\} \geq \frac{sw(\bar{X}, \bar{p}) + sw(\bar{X}', \bar{p}')}{2} \geq \frac{opt_{sw}(\eta', G)}{2} = \frac{opt_{sw}(\eta, G)}{2}$. \square

A suitable approximation can be found also for the revenue maximization.

Theorem 5. *There exists a $O(\log n)$ -approximation algorithm for (SINGLE,REVENUE)-pricing.*

Proof. Consider the two outcomes (\bar{X}, \bar{p}) and (\bar{X}', \bar{p}') for the instance (η, G) of (SINGLE,REVENUE)-pricing determined by the algorithm of Theorem 4. If $sw(\bar{X}, \bar{p}) \geq sw(\bar{X}', \bar{p}')$, let h be the index that maximizes $\frac{v_h(m_h)}{m_h} \sum_{i=1}^h |X_i|$. Then, allocating items to buyers $1, \dots, h$ at price $v_h(m_h)/m_h$, we obtain a fair outcome for (η, G) of revenue $\frac{v_h(m_h)}{m_h} \sum_{i=1}^{h-1} |X_i| \geq \frac{\sum_{i=1}^{k-1} v_i(m_i)}{\log n}$.

Analogously, if $sw(\bar{X}', \bar{p}') > sw(\bar{X}, \bar{p})$, allocating items to all the buyers in $\{k\} \cup N^>(k)$ at price $v_k(m_k)/m_k$, we obtain a fair outcome of revenue at least $v_k(m_k)$.

The claim then follows by recalling that $\sum_{i=1}^k v_i(m_i) \geq opt_{sw}(\eta, G) \geq opt_r(\eta, G)$. \square

A refined improved approximation bound for revenue maximization can be achieved when focusing on undirected social graphs. We remark that under this assumption the problem remains computationally difficult. This comes directly by restriction to the case of complete social graphs, which corresponds to the classical problem without discrimination, shown NP-hard in [Monaco *et al.*, 2015].

Theorem 6. *The (SINGLE,REVENUE)-pricing problem restricted on undirected social graphs admits a $(2 + \varepsilon)$ -approximation.*

Proof. Let us first informally discuss the basic intuition behind the algorithm. Given an instance (η, G) of (SINGLE,REVENUE)-pricing, if G is undirected, then any fair pricing under G must assign the same price to all of the buyers belonging to a same connected component. Thus, if N_1, \dots, N_ℓ are the nodes corresponding to the maximal connected components of G , the problem translates in determining how many items must be allocated to each N_h . Let η_h be the *submarket* corresponding to the set of buyers in N_h . If we are able to solve the classical revenue maximization problem in multi-unit markets for every η_h and for every amount of

items or supply assigned to η_h , then it is possible to pack the items among the components in the best possible way using MULTIPLE-CHOICE KNAPSACK as follows. We construct an instance $K(\eta, G)$ having a class O_h for each η_h . Every O_h contains m objects $o_{h,j}$ for $1 \leq j \leq m$, where each $o_{h,j}$ represents a revenue maximizing envy-free outcome for η_h with supply j . In particular, each $o_{h,j}$ has profit $z_{i,j}$ equal to the revenue of such an outcome and size $w_{h,j} = j$. The knapsack capacity is set to m .

By construction, an optimal solution S^* for $K(\eta, G)$ directly induces a revenue maximizing fair outcome (\bar{X}, \bar{p}) for (η, G) . Unfortunately, since (η, G) has size depending on $\log m$ and $K(\eta, G)$ proportional to m , this construction is not polynomial. We solve this problem by sampling the supply values, losing only a constant factor of the optimal revenue. Namely, in $K(\eta, G)$ we do not insert in each class O_h the objects $o_{h,j}$ for all the possible sizes $j = 1, \dots, m$, but only for the at most $\lceil \log_{\frac{3}{2}} m \rceil + 2$ sizes $j = (\frac{3}{2})^0, \dots, (\frac{3}{2})^{\lceil \log_{\frac{3}{2}} m \rceil}, m$. The knapsack capacity is set to $\frac{3}{2}m$.

In order to measure the loss in revenue due to this sampling, consider a fair optimal outcome (\bar{X}, \bar{p}) for (η, G) . If j_h is the number of objects that (\bar{X}, \bar{p}) assigns to a given submarket η_h and k is the integer such that $(\frac{3}{2})^{k-1} < j_h \leq \min\{(\frac{3}{2})^k, m\}$, then in the sampled instance $K(\eta, G)$ there is an object $o_{h,j}$ in class O_h with size $j = \min\{(\frac{3}{2})^k, m\} < \frac{3}{2}j_h$. Since such an object represents an optimal outcome for submarket η_h with supply j , its profit $z_{i,j}$ is at least the one that (\bar{X}, \bar{p}) collects on the subset of buyers N_h of η_h . Therefore, all such objects correspond to a solution S for $K(\eta, G)$ of profit at least $r(\bar{X}, \bar{p})$ and total size at most $\frac{3}{2}m$. Hence, the optimal solution S^* for $K(\eta, G)$ achieves at least such a revenue. In order to extract an outcome allocating at most m items, we can split S^* in at most two sets S_1^* and S_2^* , each containing a subset of objects of overall size at most m , as follows. If S^* contains an object $o_{h,j}$ of size $j \geq \frac{m}{2}$, then $S_1^* = \{o_{h,j}\}$ and $S_2^* = S^* \setminus S_1^*$. If all objects in S^* have size less than $\frac{m}{2}$, then we start inserting objects in S_1^* in any order till not exceeding total size m , and the remaining ones in S_2^* . In this way, again S_1^* thus not exceed size m , and since its size is greater than $\frac{m}{2}$, S_2^* has size less than m . Consider then the set S_i^* of maximum profit between S_1^* and S_2^* . Allocating items only to the connected components corresponding to the objects $o_{h,j} \in S_i^*$ according to the maximum revenue solutions for the related submarkets η_h with supply j , and assigning an infinite price to all the buyers of the other connected components, we can finally collect at least one half of the revenue of S^* and thus of the optimal fair allocation (\bar{X}, \bar{p}) for (η, G) .

A final observation concerns the fact that we are not really able to optimally solve $K(\eta, G)$, nor the subproblems for determining the profits of the objects $o_{h,j}$. However, in both cases we can use respectively the FPTAS of MULTIPLE-CHOICE KNAPSACK and the one in [Monaco *et al.*, 2015] for the classical revenue maximization problem in multi-unit markets. These details only account for an additional ε loss in the final approximation. \square

4 General Valuations

In this section we drop any restriction on the buyers' valuations. We first observe that the hardness reduction of Theorem 3 for the single-minded case is polynomial also for general valuations, as m is polynomially bounded in the size of the starting DENSEST K-SUBGRAPH instance. Thus, the same negative result directly extends also to general valuations.

Theorem 7. *The (GENERAL, REVENUE)- and (GENERAL, WELFARE)-pricing problems are strongly NP-hard.*

A nice reduction from the general to the single-minded case allows to suitably extend the approximation results of the previous section to general valuations.

Theorem 8. *There exists a 2-approximation algorithm for (GENERAL, WELFARE)-pricing.*

Proof. In order to prove the claim, we show how to reduce any instance (η, G) of (GENERAL, WELFARE)-pricing into an instance (η', G') of (SINGLE, WELFARE)-pricing with the following two properties:

- i. $opt_{sw}(\eta', G') \geq opt_{sw}(\eta, G)$;
- ii. exploiting the algorithm of Theorem 4 it is possible to obtain a fair outcome (\bar{X}', \bar{p}') for (η', G') which can be directly transformed into a fair outcome (\bar{X}, \bar{p}) for (η, G) with $sw(\bar{X}, \bar{p}) \geq sw(\bar{X}', \bar{p}')$.

This clearly proves the theorem.

The reduction is obtained as follows. Consider a buyer i of η , and let $S_i = \{m_i^1, \dots, m_i^\ell\}$, with $m_i^1 \leq \dots \leq m_i^\ell$, be the set of the bundle sizes that are preferred by i for some positive price p_i , that is maximizing her utility for such a price. By the definition of S_i , if $j < k$ then $v_i(m_i^j) \leq v_i(m_i^k)$, otherwise size m_i^k would be never preferred. Moreover, if bundle size m_i^j is preferred by i for a given price p_i , comparing the utility of i for sizes m_i^j, m_i^{j-1} and sizes m_i^j, m_i^{j+1} , we obtain $(v_i(m_i^j) - v_i(m_i^{j-1})) / (m_i^j - m_i^{j-1}) \geq p_i \geq (v_i(m_i^{j+1}) - v_i(m_i^j)) / (m_i^{j+1} - m_i^j)$. For the sake of brevity, in the sequel we denote the quantity $(v_i(m_i^j) - v_i(m_i^{j-1})) / (m_i^j - m_i^{j-1})$ simply as v_i^j .

In η' we associate ℓ distinct single-minded marginal buyers i_1, \dots, i_ℓ to buyer i as follows: i_1 has preferred size m_i^1 , that is valuation function $v_{i_1}(m_i^1) = v_i(m_i^1)$, 0 otherwise; furthermore, for each $j > 1$, i_j has valuation function $v_{i_j}(m_i^j - m_i^{j-1}) = v_i(m_i^j) - v_i(m_i^{j-1})$, 0 otherwise.

For what concerns the social graph G' , for each buyer i of η, G' contains arcs (i_j, i_k) for each $j > k$. Moreover, for each pair of buyers i, l such that (i, l) belongs to G , G' contains an arc (i_j, l_k) if $v_i^j > v_l^k$.

Intuitively, in such a reduction, the allocation of a bundle of size m_i^j to buyer i corresponds to the allocation of bundles of preferred size to the prefix of the marginal buyers i_1, \dots, i_j in η' .

Let us first prove that $opt_{sw}(\eta', G') \geq opt_{sw}(\eta, G)$. To this aim, we show how to transform a fair outcome (\bar{X}, \bar{p}) for (η, G) into a fair one (\bar{X}', \bar{p}') for (η', G') such that $sw(\bar{X}', \bar{p}') = sw(\bar{X}, \bar{p})$. If $|X_i| = m_i^j$ in (\bar{X}, \bar{p}) , then in (\bar{X}', \bar{p}') we allocate to all the marginal buyers i_k with

$k \leq j$ their preferred bundles at price p_i , and no bundle to all the remaining buyers (asking them an infinite price). In this way, for each buyer i of η , $|X_i| = \sum_{k=1}^j |X'_{i_k}|$, and $v_i(|X_i|) = \sum_{k=1}^j v_{i_k}(|X'_{i_k}|)$. Notice that, since the m_i^j is the preferred bundle of buyer i in η , we must have that $v_i^k \geq p_i$ for every $k \leq j$. Thus, $(\overline{X'}, \overline{p'})$ is feasible. Moreover, since we propose the same price p_i to all buyers i_j , if \overline{p} is a fair pricing under G , also $\overline{p'}$ is fair under G' . In conclusion, $(\overline{X'}, \overline{p'})$ is a fair outcome for (η', G') with $sw(\overline{X'}, \overline{p'}) = sw(\overline{X}, \overline{p})$.

Consider now the solution $(\overline{X'}, \overline{p'})$ for η' obtained by the algorithm of Theorem 4, where the buyers are ordered in non increasing order of their values v_i^j , and in case of ties for two buyers i_j, i_k with $j < k$, that is if $v_i^j = v_i^k$, putting i_j before i_k in the ordering (recall that by the above observations it must be $v_i^j \geq v_i^k$). Since by the ordering and by the construction of the algorithm $(\overline{X'}, \overline{p'})$ can allocate only a prefix of marginal buyers i_1, \dots, i_j in η' associated to each i of η , in $(\overline{X'}, \overline{p'})$ we can assign a bundle of size m_i^j to buyer i at the price proposed to i_j in $(\overline{X'}, \overline{p'})$. It is easy to see that $(\overline{X}, \overline{p})$ is a fair outcome for (η, G) with $sw(\overline{X}, \overline{p}) = sw(\overline{X'}, \overline{p'})$. \square

We are also able to devise a suitable approximation algorithm for the revenue maximization problem.

Theorem 9. *There exists a $O(\log n + \log m)$ -approximation algorithm for (GENERAL, REVENUE)-pricing.*

Proof. (Sketch) Given an instance (η, G) of (GENERAL, REVENUE)-pricing, using the same arguments of Theorem 8, we can obtain a fair outcome $(\overline{X'}, \overline{p'})$ for the reduced single-minded instance (η', G') with $sw(\overline{X'}, \overline{p'}) \geq opt_{sw}(\eta, G)/2 \geq opt_r(\eta, G)/2$, allocating only prefixes of buyers i_1, \dots, i_j in η' .

As in Theorem 5, we can extract from $(\overline{X'}, \overline{p'})$ an outcome $(\overline{X''}, \overline{p''})$ of revenue $r(\overline{X''}, \overline{p''}) \geq sw(\overline{X'}, \overline{p'})/\log n'$, where n' is the number of buyers in η' , still allocating only prefixes of buyers i_1, \dots, i_j in η' .

Again resorting to the arguments of Theorem 8, we can transform $(\overline{X''}, \overline{p''})$ into a fair outcome for (η, G) with the same revenue. The theorem then follows by observing that in the reduction $n' \leq n \cdot m$. \square

We now focus our attention to specific cases allowing optimal polynomial time algorithms. To this aim observe that, since buyers in a same strongly connected components must be offered identical prices, a convenient way to represent G is through the graph of its strongly connected components. Namely, each node represents a component and there is an arc between two nodes if the same happens between the corresponding components. Such a graph is clearly a DAG, and properly restricting its topology we can find interesting polynomial sub-cases. In particular, we are able to prove the following theorem, that includes all the interesting realistic cases mentioned in the introduction.

Theorem 10. *There exists a polynomial time optimal algorithm for (GENERAL, WELFARE)- and (GENERAL, REVENUE)-pricing under any of the following restrictions on the social graph G : i. G is undirected, ii.*

G contains a constant number of strongly connected components, iii. the graph of the strongly connected components of G is an arborescence.

Proof. (Sketch) For undirected social graphs, a polynomial time algorithm can be obtained by the same initial construction proposed in the proof of Theorem 6 by observing that, since the size of the input instance (η, G) now is polynomial in m , sampling is no longer necessary. Moreover, both the knapsack instance $K(\eta, G)$ and the determination of the optimal market outcomes corresponding to the objects $o_{h,j}$ can be solved in polynomial time.

Let us now consider the remaining two cases. Let N_1, \dots, N_ℓ be the sets of buyers corresponding to the strongly connected components of G , and η_1, \dots, η_ℓ be the corresponding submarkets. Since all the buyers in a same component must be offered the same price, as shown in [Monaco *et al.*, 2015], for each η_k there exists a set $\mathbb{P}_{N_k}^{opt}$ of polynomial size containing all the candidate optimal prices.

For a constant number k of strongly connected components, it is possible to exhaustively search among a polynomial number of candidate optimal price vectors and return the best determined outcome.

For arborescences of strongly connected components, the basic idea behind the algorithm is that, along a path of the arborescence, some submarkets can be forced by their neighbors to set some prices which they would not set on their own. However, these prices must be optimal for some other submarket on the path. We can then use dynamic programming to find an optimal solution by combing the optimal outcome of each single subtree of submarkets, in such a way that, for respecting the fairness, non-increasing prices are proposed along each path, and the global supply constraint is satisfied. This holds both for social welfare and revenue maximization. Details will appear in the full version of the paper. \square

5 Conclusions and Future Work

We proposed a framework able to capture many realistic scenarios in which a discriminatory pricing is acceptable by buyers. We considered both social welfare and revenue maximization, under different hypotheses on the buyers valuations and the social graph topology, providing polynomial time optimal algorithms, hardness results and suitable approximations. Our framework is very general and we believe will capture future research attention.

Beside improving our approximation results, it would be worth considering other types of markets with non identical items, unit-demand buyers, single-minded valuations on non homogeneous items, and so forth. It would be also interesting to consider different notions of envy-freeness, such as *pair envy-freeness*, *social envy-freeness*, *proportionality*, etc.

Finally, the issue of influencing the social graph topology so as to increase the achievable social welfare or revenue appears another challenging interesting research direction.

References

[Abebe *et al.*, 2017] R. Abebe, J. M. Kleinberg, and D. C. Parkes. Fair division via social comparison. In *16th Conf.*

- on *Autonomous Agents and MultiAgent Systems, AAMAS*, pages 281–289, 2017.
- [Alon *et al.*, 2013] N. Alon, Y. Mansour, and M. Tennenholtz. Differential pricing with inequity aversion in social networks. In *14th ACM Conf. on Electronic Commerce, EC*, pages 9–24, 2013.
- [Amanatidis *et al.*, 2016] G. Amanatidis, E. Markakis, and K. Sornat. Inequity Aversion Pricing over Social Networks: Approximation Algorithms and Hardness Results. In *41st Int. Symp. on Mathematical Foundations of Computer Science MFCS*, volume 58, pages 9:1–9:13, 2016.
- [Anshelevich *et al.*, 2017] E. Anshelevich, K. Kar, and S. Sekar. Envy-free pricing in large markets: Approximating revenue and welfare. *ACM Transactions on Economics and Computation*, 5(3):16, 2017.
- [Balcan *et al.*, 2008] M. Balcan, A. Blum, and Y. Mansour. Item pricing for revenue maximization. In *9th ACM Conf. on Electronic Commerce EC*, pages 50–59, 2008.
- [Bilò *et al.*, 2017] V. Bilò, M. Flammini, and G. Monaco. Approximating the revenue maximization problem with sharp demands. *Theor. Comp. Science*, 662:9–30, 2017.
- [Bilò *et al.*, 2018] V. Bilò, M. Flammini, and L. Moscardelli. On the impact of buyers preselection in pricing problems. In *17th Conf. on Autonomous Agents and MultiAgent Systems, AAMAS*, 2018.
- [Brânzei *et al.*, 2016] S. Brânzei, A. Filos-Ratsikas, P. B. Miltersen, and Y. Zeng. Envy-free pricing in multi-unit markets. *arXiv:1602.08719*, 2016.
- [Briest and Krysta, 2006] P. Briest and P. Krysta. Single-minded unlimited supply pricing on sparse instances. In *17th ACM-SIAM Symp. on Discr. Alg., SODA*, pages 1093–1102, 2006.
- [Briest, 2008] P. Briest. Uniform budgets and the envy-free pricing problem. In *Automata, Languages and Programming, 35th Int. Colloq., ICALP*, pages 808–819, 2008.
- [Chalermsook *et al.*, 2012] P. Chalermsook, J. Chuzhoy, S. Kannan, and S. Khanna. Improved hardness results for profit maximization pricing problems with unlimited supply. In *Approx., Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 73–84, 2012.
- [Chalermsook *et al.*, 2013a] P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Graph products revisited: Tight approximation hardness of induced matching, poset dimension and more. In *24th ACM-SIAM Symp. on Discr. Alg., SODA*, pages 1557–1576, 2013.
- [Chalermsook *et al.*, 2013b] P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *54th IEEE Symp. on Foundations of Computer Science, FOCS*, pages 370–379, 2013.
- [Chen and Deng, 2010] N. Chen and X. Deng. Envy-free pricing in multi-item markets. In *Automata, Languages and Programming, 37th Int. Colloq., ICALP*, pages 418–429, 2010.
- [Chen *et al.*, 2011] N. Chen, A. Ghosh, and S. Vassilvitskii. Optimal envy-free pricing with metric substitutability. *SIAM J. on Computing*, 40(3):623–645, 2011.
- [Chen *et al.*, 2016] N. Chen, X. Deng, P. W. Goldberg, and J. Zhang. On revenue maximization with sharp multi-unit demands. *J. of Combinatorial Optimization*, 31(3):1174–1205, 2016.
- [Cheung and Swamy, 2008] M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *49th IEEE Symp. on Foundations of Computer Science, FOCS*, pages 35–44, 2008.
- [Chevaleyre *et al.*, 2007] Y. Chevaleyre, U. Endriss, S. Estivie, N. Maudet, et al. Reaching envy-free states in distributed negotiation settings. In *20th Int. Joint Conf. on Artificial Intell., IJCAI*, volume 7, pages 1239–1244, 2007.
- [Demaine *et al.*, 2008] E. D. Demaine, U. Feige, M. T. Hajiaghayi, and M. R. Salavatipour. Combination can be hard: Approximability of the unique coverage problem. *SIAM J. on Computing*, 38(4):1464–1483, 2008.
- [Feldman *et al.*, 2012] M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *ACM Conf. on Electronic Commerce, EC*, pages 532–549, 2012.
- [Flammini *et al.*, 2018] M. Flammini, M. Mauro, and M. Tonelli. On social envy-freeness in multi-unit markets. In *32nd AAAI Conference on Artificial Intelligence*, 2018.
- [Guruswami *et al.*, 2005] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *16th ACM-SIAM Symp. on Discr. Alg., SODA*, pages 1164–1173, 2005.
- [Hartline and Yan, 2011] J. D. Hartline and Q. Yan. Envy, truth, and profit. In *12th ACM Conf. on Electronic Commerce EC*, pages 243–252, 2011.
- [Laffont *et al.*, 1998] Jean-Jacques Laffont, Patrick Rey, and Jean Tirole. Network competition: Ii. price discrimination. *The RAND J. of Economics*, 29(1):38–56, 1998.
- [Lawler, 1979] E. L. Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 4(4):339–356, 1979.
- [Monaco *et al.*, 2015] G. Monaco, P. Sankowski, and Q. Zhang. Revenue maximization envy-free pricing for homogeneous resources. In *24th Int. Joint Conf. on Artificial Intell., IJCAI*, pages 90–96, 2015.
- [Phlips, 1983] L. Phlips. *The Economics of Price Discrimination*. Cambridge University Press, 1983.
- [Pigou, 1920] A. Pigou. *The Economics of Welfare*. Palgrave Macmillan UK, 1920.
- [Stokey, 1979] N. L. Stokey. Intertemporal price discrimination. *Quarterly J. of Economics*, 93(3):355–371, 1979.
- [Varian, 1989] H. R. Varian. Chapter 10 price discrimination. volume 1 of *Handbook of Industrial Organization*, pages 597 – 654. Elsevier, 1989.