GRAN SASSO SCIENCE INSTITUTE
SCHOOL OF ADVANCED STUDIES
Scuola Universitaria Superiore

PhD Thesis

Social Effects on Envy-Freeness and Discrimination of Prices in Multi-Unit Markets

PhD Program in Computer Science: XXXI cycle

Author:

Manuel Mauro
manuel.mauro@gssi.it

Advisor:

Prof. Michele Flammini
michele.flammini@gssi.it

October 2019

GSSI Gran Sasso Science Institute
Viale Francesco Crispi, 7 - 67100 L'Aquila - Italy

Declaration of Authorship

I undersigned, Manuel Mauro, hereby declare that this thesis, entitled "Social Effects on Envy-Freeness and Discrimination of Prices in Multi-Unit Markets", and the work herein presented is an original contribution obtained during my PhD program under the guidance of my supervisor Prof. Michele Flammini. In particular, I confirm that:

- Chapter 2 is based on the paper "On Social Envy-Freeness in Multi-Unit Markets", first appeared in the Thirty-Second AAAI Conference on Artificial Intelligence [56] and later extended and published in Artificial Intelligence, Elsevier [57].
- Chapter 3 is based on the paper "On Fair Price Discrimination in Multi-Unit Markets", first appeared in the International Joint Conferences on Artificial Intelligence [55] and for which an extended version has been accepted with major revision for publication in Artificial Intelligence, Elsevier [58].
- Chapter 4 contains novel results currently under review.

Signed:

Date:

Acknowledgements

Firstly, I would like to thank the Gran Sasso Science Institute for selecting me and providing me the opportunity to conduct my thesis within the Computer Science department. It has been a fulfilling journey.

I would like to express my sincere gratitude to my advisor Prof. Flammini for his constant support during my PhD study, for his patience, motivation, and immense passion. His guidance and humor were invaluable in all the time of research and writing of this thesis. I would like to thank him for the time he invested in me, for our many discussions, as well as for his kindness and encouragement. His knowledge on the subject has contributed enormously to the result.

I would also like to thank Prof. Markakis, Prof. Moscardelli, and all the anonymous reviewers who greatly improved this work: thank you very much for the time you devoted in order to contribute to my thesis, for sharing your knowledge and providing insightful comments.

Finally yet importantly, sincere thanks to my all my colleagues and especially to Matteo Tonelli, with whom I spent countless hours during these wonderful years. This thesis would not have been possible without his precious contribution.

Manuel Mauro

Abstract

Multi-unit markets model the real-world scenario in which a seller is willing to sell multiple copies of a single good to many buyers, like in the case of commodities, retailer goods, subscriptions, etc., but they are also a powerful abstraction for resource allocation problems, such as power supply in manufacture systems, cargo space in transportation industry, bandwidths in the radio spectrum, and many more.

The basic model is commonly enriched with a number of constraints in order to more closely describe settings typically arising in practice, such as a limited supply of items, customers' intelligent behavior (envy-freeness), a seller aiming at revenue or social welfare maximization, etc. The literature is rich of results for this specific and more general kinds of markets, but only recently researchers have started investigating relationships and knowledge shared among buyers. Seminal works have considered buyers as individuals of a population who transfer their knowledge and thus behave accordingly.

In the present thesis we extend or introduce concepts related to sociality among buyers in multi-unit markets and study problems of revenue and social welfare maximization from an algorithmic game theory perspective.

Sociality and price discrimination are the focus of our work, with resulting frameworks for a more accurate description of the real-world scenarios related to multi-unit markets and a picture of the computational complexity of the challenges arising from the interplay of the various requirements considered in the model.

The contribution of this work can be broadly divided into two different respects. One concerns a suitable relaxation of the envy-freeness notion induced by social relationships, obtained by restricting the corresponding constraints only to known peers, instead of the whole population of buyers. The related results are presented in Chapter 2. The second line concerns forms of price discrimination that buyers can consider as fair. More precisely, the prices proposed to neighbor buyers cannot be arbitrarily, but they should be reasonably close. The original results obtained in this setting are described in Chapter 3 and Chapter 4.

In both cases, we propose suitable frameworks for representing the arising scenarios, and provide optimal efficient algorithms, and hardness and approximation results concerning the sellers' revenue and the social welfare maximization.

List of Publications

This is the list of the author's publications arising from the thesis work, on which this doctoral dissertation is based:

1. Michele Flammini, Manuel Mauro, and Matteo Tonelli. On social envy-freeness in multi-unit markets. In *32nd AAAI Conference on Artificial Intelligence*. AAAI Press, Palo Alto, California, USA, 2018;
2. Michele Flammini, Manuel Mauro, and Matteo Tonelli. On fair price discrimination in multi-unit markets. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI, Stockholm, Sweden*. IJCAI Org, Vienna, Austria, 2018;
3. Michele Flammini, Manuel Mauro, and Matteo Tonelli. On social envy-freeness in multi-unit markets. *Artificial Intelligence*, 2018. Elsevier, Amsterdam, Netherlands;
4. Michele Flammini, Manuel Mauro, and Matteo Tonelli. On fair price discrimination in multi-unit markets. *Accepted with major revision for publication in Artificial Intelligence*, 2019. Elsevier, Amsterdam, Netherlands;
5. Michele Flammini, Manuel Mauro, Matteo Tonelli, and Cosimo Vinci. Inequity aversion pricing in multi-unit markets. In *Manuscript*, 2018.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Multi-unit markets model the real-world setting of a market in which a single seller is willing to sell a set M of m identical items to a set N of n buyers. Being the items equivalent, buyers value bundles (subsets) of items only according to their size. In this type of markets, items on sale could be commodities, homogeneous server grids, supercomputers in computer networks, power supply in manufacture systems, cargo space in transportation industry, bandwidths in the radio spectrum, and many other kinds of goods or resources. In many real world scenarios, maximizing the revenue of the seller is the natural objective, but when setting prices and allocating goods it is important to have a sound model for buyers' behavior. One of the most studied model is *envy-freeness*. Individuals described by this notion always maximize the difference between their valuation for the item they buy and the price the seller sets for it, that is, their utility.

In the present work we study solution concepts for multi-unit markets that allow, at the same time, to model characteristic features of the modern internet economy and greatly increase seller's revenue or social welfare. In order to attain this objectives we follow two different directions. In the first one, addressing the problem of modeling the distributed nature of internet markets, we study the effects of buyers' social relationship on envy-freeenss. In particular, we study what happens when buyers can share information only with the ones they are connected to and envy-freeness cannot extend beyond their neighborhood. In the second one, closely connected to the previous one, we empower the seller with modern tools like targeted advertisement and customized deals. The former tool allows the seller to select the subset of buyers she wants to sell her items to, while with the latter tool the seller is able to ask different prices for equivalent items. Even if this could seem unfair, we will see that the practice is widespread and well received by the buyers. These mechanisms are among those which characterize the modern internet economy and offer important gain in terms of revenue and even social welfare.

Pricing-based mechanisms have been applied to many problems of resources allocation in the context of multi-agent systems. Due to their expressiveness, these mechanisms have been proven to be able to model a huge variety of real-world scenarios. In this respect, multi-unit markets are one of the most basic forms, and they have been widely investigated in this context, because on the one hand they are still able to capture many realistic settings, and on the other hand, while already revealing the inherent intricacies of markets, they offer more chances of achieving computational tractability or better approximations. Tools and techniques developed in multi-unit markets are important as they can give insights and turn useful when addressing more general pricing-problems. As a consequence, not surprisingly the multi-unit model is central in combinatorial auctions and has been studied in a large body of literature [14, 19, 45, 47, 48, 52, 63, 78], with a strong focus on the determination of incentive-compatible pricing-based mechanisms. Unfortunately, many of the related auctions, e.g., the VCG mechanism [40, 67, 92] (for bidders with unlimited budgets) or the ascending auction [45, 84] (for bidders with budgets), sell identical products at unfairly different prices. Furthermore, the shortcomings of the VCG mechanism in terms of revenue for the seller gave rise to an important research direction that was initiated in [53, 64, 65], and resulted in a sequence of follow-up results for which we refer to Hartlines book [70] for a comprehensive discussion.

Envy-freeness is closely related to the fundamental economic solution concept of Walrasian equilibrium [93], which requires both the outcome of a market to be envy-free and the market itself to "clear" (i.e., each unsold item has to be priced zero) meaning that the whole demand from the buyers is fully satisfied and no item is left. The Walrasian equilibrium has proved to be an elegant tool and benchmark for the analysis of competitive markets in economics, capturing both the behavior of buyers and the efficiency of the whole market, but may not always exists and even when its existence is guaranteed the revenue it generates for the sellers is far from being optimal. The latter concern, the maximization of seller's revenue, is an overarching requirement in all the real-world scenarios we listed. Because of this reason, envy-free pricing [68] has been proposed as a relaxation of the classical Walrasian equilibrium, prone to yield an increased revenue. Thus, it is no longer required for the market to clear, but only the satisfaction of the buyers according to their behavioral properties. As a result, an envy-free pricing always exists and the focus moves on determining one that maximizes seller's revenue.

The notion of envy-freeness is also connected to another classical stability concept in multi-agent systems, that is, the Nash equilibrium. In fact, given a pricing of the items and up to some coordination in order to break ties, buyers naturally and autonomously converge to an envy-free allocation [49].

Several works in the literature have addressed the problem of pricing and *envy-free* allocations [22, 23, 31, 68, 71]. Similarly, different papers have considered the closely related notion of *pair envy-freeness* [41, 52, 54, 82].

As the internet economy makes today's markets global and more and more distributed, the concept of envy-freeness needs some refinement in order to model this reality. In fact, a major concern about the standard definition is the assumption of a single buyer's full knowledge about all the prices and allocations of items proposed to all the others. Moreover, many tools give to today's sellers the ability to precisely target special classes of buyers, offering great possibility in terms of revenue maximization and at the same time bypassing the envy-freeness requirements.

The same tools can be exploited in order to deploy various mechanisms of price discrimination. Even if charging customers different prices for the same item can be perceived as unfair at first, it is actually widespread. For instance, when buying tickets, customers who arrive earlier usually get a better price. In the opposite way, the price of electronic devices like computers and cameras decreases over time. We could define this form of price discrimination as *temporal* since buyers are charged different prices according to their time of purchase. Another possible form of price discrimination, that we could define as *spatial* price discrimination, is also consolidated in the sale of the same good at discordant prices at different sites, like in different countries or shops. A further type of discrimination occurs in users profiling, when premium customers such as frequent-fliers are charged lower prices with respect to the general public, or when some classes of buyers are recognized to be eligible for better treatment, such as age discounts or financial aids.

One of the first classifications of this kind of mechanisms is maybe due to A. Pigou in [87]. Introductions to the concept of price discrimination together with the main arising scenarios and the related motivation can be found in [26, 79, 86, 88, 91].

1.1 Related Work

The first work to study the problem of revenue maximization under the constraints of envy-freeness is due to Guruswami et al. [68]. The setting they studied is the very general model of combinatorial markets, in which a set of different items is on sale and for each item multiple copies are available. They were able to show that, in case of unit-demanded buyers (i.e., they have positive valuations only for sets consisting of a single item) or single-minded (i.e., each of them demands a fixed subset of items), computing an optimal envy-free pricing is APX-hard, even in case of unlimited supply. On the opposite side, they were able to devise an $O(\log n)$-approximation algorithm for unit-demand buyers with limited supply, and an $O(\log n + \log m)$-approximation for single-minded buyers with unlimited supply. Later, this last result has been extended to general valuation functions in [13]. Still for single-minded buyers with unlimited supply, [54] provided a polynomial-time algorithm in case prices can be assigned to subsets of items, plus other hardness results for specific variants of the problem. Studying special cases for the problem, Chen and Deng [31] were able to develop a polynomial-time algorithm for unit-demand buyers, in the case in which each buyer has a positive valuation for at most two items. A stronger lower bound to the computational tractability of such a problem has been obtained by Briest [23] who showed that, under the R3SAT assumption proposed in [50], the unit-demand envy-free pricing problem cannot be approximated within $O(log^{\epsilon} n)$ for some $\epsilon > 0$. A similar lower bound for envy-free pricing with unlimited supply has been derived by Demaine et al. [44]. An even stronger lower bound stands when buyers are single-minded and items have limited supply, that is the $\Omega(\sqrt{m})$ inapproximability given in [66]. On the other hand, Cheung and Swamy [36] showed a nearly-optimal $O(\sqrt{m} \cdot \log m_{max})$-approximation algorithm, where m_{max} is the maximum among all the supplies of the different items. Finally, Hartline and Koltun [71] designed near-linear and near-cubic time approximation schemes under the assumption that the number of distinct items for sale is constant.

Following suggestions from classical papers [60, 90], multi-unit markets with multiple identical items have been mostly investigated under the pair envy-freeness notion. In particular, [52] studied the case of budgeted buyers with additive valuations. Starting from the observation that incentive-compatibility requires price discrimination, they focused on outcomes that are pair envy-free rather than incentive-compatible, considering progressively stronger levels of price discrimination. In particular, they investigated item-pricing, i.e., no discrimination at all with all identical items being assigned the same price, and bundle-pricing, that is, a mildly fair form of price discrimination in which sets of "shrink-wrapped" items are sold with a different non-proportional price for each bundle size. After showing that the progressive discrimination levels allow a corresponding increase in revenues, they proved that the revenue maximization problem is NP-hard, and provided a polynomial-time 2-approximation algorithm for item-pricing. Such an approximation has been improved to an FPTAS in [41], who also proved that in case of heterogeneous goods it is impossible to approximate the maximum revenue in single-minded instances within $O(min(n, m)^{1/2\varepsilon})$ for any $\varepsilon > 0$, unless $P = NP$. Some improvements of these results were provided in [42].

[82] investigated envy-freeness in multi-unit markets. In particular, they considered buyers with arbitrary valuation functions and no budgets. They gave hardness and approximation results for the revenue maximization problem in several cases arising by assuming envy-freeness or pair envy-freeness, item- or bundle-pricing, and single-minded (with free disposal) or general valuation buyers. More precisely, for single-minded buyers they showed that in all the cases the revenue maximization problems is NP-hard. Moreover, they provided corresponding FPTASs, with the exception of the pair envy-freeness item-pricing case, for which they gave an $O(\log n)$ approximation. For general valuations, in case of item-pricing they gave a polynomial-time algorithm for envy-freeness and an $O(\log n)$-approximation one for pair envy-freeness. For bundle-pricing they gave an $\Omega(log^{\epsilon} n)$ approximability lower bound for some $\epsilon > 0$ under the R3SAT assumption, plus an $O(\log n \cdot \log m)$ approximation.

In [20], Branzei et al. investigated Walrasian envy-free pricing mechanisms in multi-unit markets. They focused on (monotone) pricing mechanisms, identifying a natural mechanism selecting the minimum Walrasian envy-free price, in which for two buyers the best response dynamic converges from any starting profile, and for which they conjecture convergence for any number of buyers. The authors considered the same setting in [21], where, assuming budgeted buyers with additive valuations, they studied the dynamics of Walrasian envy-free pricing mechanisms showing that, for any such a mechanism, the best response dynamic starting from truth-telling converges to a pure Nash equilibrium with a small loss in revenue and welfare.

Finally, suitable models for explicitly taking into account different forms of fair price discrimination in multi-unit markets have been proposed in [5, 7, 55], with corresponding hardness and approximation results on revenue maximization. To the same aim, mild forms of price discrimination have been investigated also in [52, 54].

1.2 Multi-Unit Markets

Multi-unit markets are a special case of combinatorial markets, in which a single seller is willing to sell a set M of m identical item to a set N of n buyers.

According to the buyers' valuations for the different bundles sizes, to the price scheme adopted by the seller, to the considered envy-freeness notion, and to other details specified below, different interesting scenarios arise.

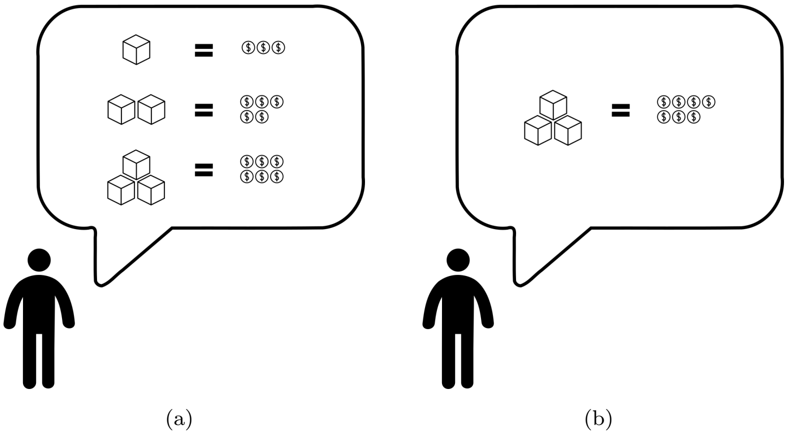

FIGURE 1.1: In Figure 1.1(a) we can see a buyer with general valuations, who expresses how much she is willing to pay for each bundle size. In Figure 1.1(b) instead we can see a single-minded buyer who is interested only in a specific bundle size and has null valuation for all the others.

Buyers' Valuations The most general way to describe buyer's valuations is by means of what are usually referred to as *general valuations*, that is, buyers who freely express an amount they are willing to pay for of every possible size for a bundle of items.

Quite on the opposite side of the spectrum of the possible valuation functions lay *single-minded* buyers, who are interested only in a specific bundle size, for which they are willing to pay a certain amount, and in no other bundle.

Figure 1.1 gives a graphical representation of an instance of both kind of buyers. Observe that both valuation functions have real-world use cases. For instance, a big computer manufacturer could be interested in buying the whole production of touch screens from smaller producers expressing different valuations according to the supply, while a car owner is only interest in purchasing four wheels from a tire repairer, no less because his car needs all of them, no more because disposing of the exceeding ones has a cost (sometimes getting rid of the extra items is free, in the present work we explore also this case).

Pricing Schemes The seller can adopt different strategies in giving a price to a bundle of items. Probably the most trivial one is to offer a price per item (item-pricing) and linearly derive the price of a bundle according to its size. However, often sellers offer some kind of discount for big bundles, thus pricing them in a nonlinear way (bundle-pricing). (See Figure 1.2.)

These two pricing schemes are non-discriminatory, in the sense that a certain bundle is offered at the same price to every buyer, but sometimes it is acceptable for two buyers to receive the same item or bundle at different prices, e.g., a ticket for a flight is sold at

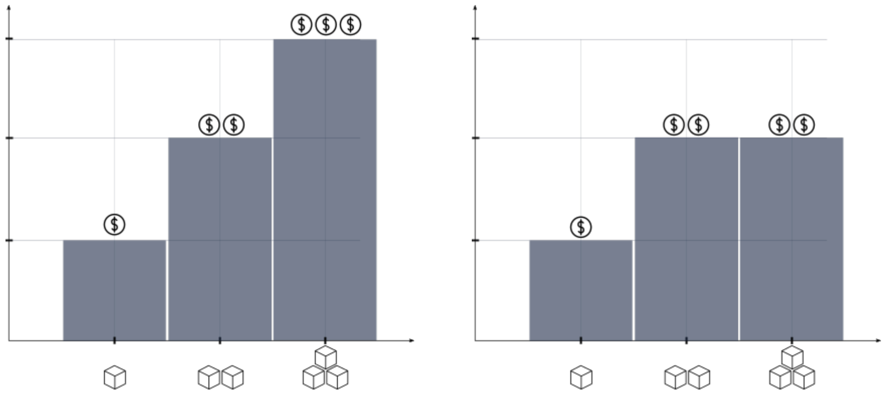

(a) (b)

Figure 1.2: In Figure 1.2(a) we can see the simple scheme in which the seller sets the same price for each single unit of good. In Figure 1.2(b) instead we can see a more general pricing scheme, usually referred to as bundle-pricing, in which multiple items can be packed together and given a special price. In this case, for instance, we have three items of the same good sold at the price of two of them, that is, a classical offer in supermarkets.

different prices to buyers buying it at different points in time, or special discounts are offered to premium customers.

Considering different pricing schemes is very important because, as we will see soon, they can greatly increase the seller's revenue or the social welfare. Moreover, they are widespread in practice.

Envy-freeness Buyers are individuals with their own objectives and their behavior is not coordinated by some external entity. In the case of (combinatorial) markets such behavior is usually modeled by the notion of envy-freeness, which comes in different flavors.

In the classical notion of envy-freeness, buyers look at their peers in order to assess their outcome. If a buyer gets a bundle which would make another buyer happier than she is with her own, the allocation has *envy*. In order for allocations to be stable and represent real outcomes, they must not contain this kind of situations. We say that they must be *(pair) envy-free.*

In another standard notion of envy-freeness, buyers always maximize their utility, e.g. in the case of buyers with general valuations, once the seller has set a price for the bundles on sale, a buyer will choose to buy (or will be happy if she receives) the bundle which maximizes the difference between her valuation and the price of that specific bundle (see Figure 1.3). In case of single-minded buyers, the choice is between the unique preferred

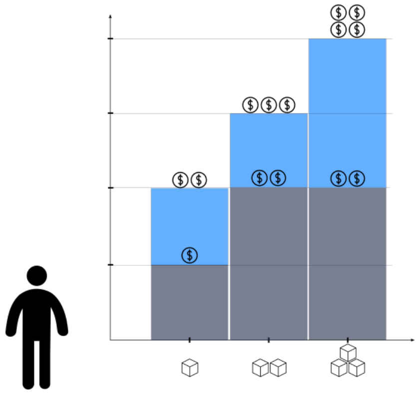

FIGURE 1.3: In Figure, given the seller (bundle-) pricing for the items on sale, the buyer evaluates which bundle is her preferred one by subtracting the price from her initial valuation. The light blue part of the column represent her utility for the specific bundle.

bundle or no items at all. Observe that this notion is a stronger requirement with respect to the previous one, as buyers can envy bundles not allocated to any other one.

Supply The amount of items on sale is an important aspect to take into account. In the case of identical items two different scenarios are of interests: *limited supply* and *unlimited supply*. The former is easily linked to any material good, while the latter is of interest for its connection to digital goods, like video streaming, software, e-books, etc.

Social Welfare and Revenue Maximization There are many ways to measure an outcome for a market, that is, the way bundles are assigned to buyers and the payment the seller asks for. Among the most well studied and of practical interest we have the *seller's revenue* and the *social welfare*. The first one measures, as the name clearly states, the amount of money the seller gets from the outcome and models all the real-world situations in which the seller is a selfish agent, e.g., private companies, individuals, etc. The second one tries to measure how well both buyers and seller perform in the market and it is useful when we want to model selfless sellers like government agencies or nonprofit ones.

Price Discrimination In many contexts selling the same item at different prices to different buyers is acceptable and sometimes desirable. This is for instance the case of tickets; in fact, as the event approaches, the same ticket can be sold at much greater price with respect to the previous months. Other examples are products sold in more than one country, for which bigger markets are usually allowed to get the item at smaller prices, or items for which the seller puts on sale both a normal and a special edition (only slightly different), thus discriminating between buyers who are willing to pay more

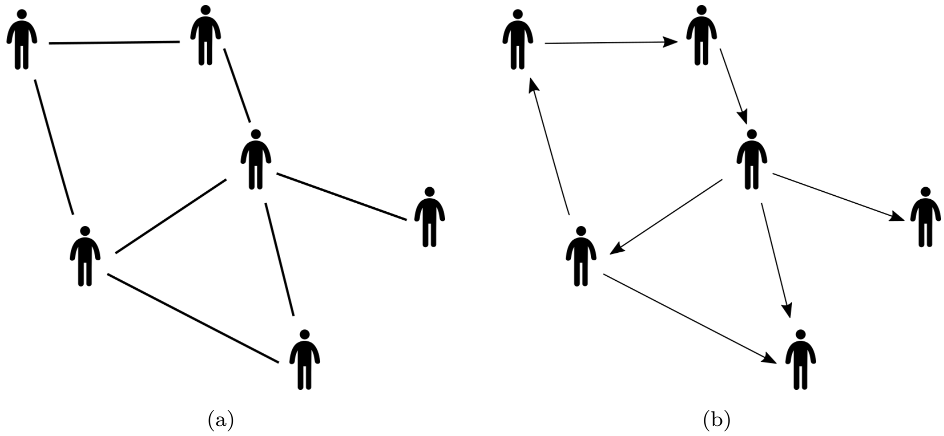

FIGURE 1.4: Buyers' information about the allocations and prices offered to other buyers can be modeled by the use of social networks. In Figure 1.4(a) a simple graph over the set of buyers models a symmetric relationship, like friendship. The directed graph in Figure 1.4(b) instead models a more complex asymmetric relationship, like knowledge.

for the small extras and buyers who don't. We will see that this pricing mechanisms can greatly increase both seller's revenue and social welfare, making them worth to be investigated.

Buyers Preselection Personalized marketing, also known as one-to-one marketing or individual marketing, is a marketing strategy by which companies leverage data analysis and digital technology to deliver individualized messages and product offerings to current or prospective customers. Advancements in data collection methods, analytics, digital electronics, and digital economics, have enabled marketers to deploy more effective real-time and prolonged customer experience personalization tactics.

Social Relationships among Buyers Due to the development of the electronic commerce, it is possible for a seller to have customers distributed in large environments or all around the world. This simple fact makes unrealistic the assumption of buyers' have complete knowledge about the prices and bundles allocated to others. Shortage of this knowledge can be exploited by the seller in order to increase her revenue or even social welfare. A possible way to model this limitation in buyers' knowledge is using a directed or undirected graph over the set of buyers, representing the fact that each single buyer has information only about the subset of her neighbor buyers (See Figure 1.4).

1.3 Dissertation Overview

The present work is organized in five chapters. The first chapter, the current one, contains an introduction to the topic, to the related work, and to the problems we will tackle in the coming chapters. Three chapters follow with all the results of our work. Finally, the last chapter gives a global summary of our results, some conclusive remarks, and outlines some worth investigating research directions.

In Chapter 2 we study what happens when buyers share information about the bundles of items with their social friends, while the information about the prices of bundles are public. Thus, envy-freeness constraint can be relaxed to deal only with adjacent buyers, and in addition only mild forms of price discrimination are considered, usually referred to as bundle-pricing, in which the seller is able to propose different non-linear prices for bundles of different sizes.

In Chapter 3 we start tackling the problem of price discrimination and develop a framework which defines when the corresponding pricing schemes can be considered fair by buyers. More in details, unlike Chapter 2, in the framework we develop here the social graph represents the knowledge each buyer has of the prices offered to her neighbors. Buyers in such a social network accept only prices that are lower or equal to those offered to their neighbors. Such a framework, while allowing price discrimination, limits its extent, thus inducing a notion that we define as *fair price discrimination*.

In Chapter 4, driven by the widespread practice of targeted advertisement, we empower the seller with another skill, that is, the ability of accepting only a subset of buyers to the market, while maintaining the same capability to offer a fair discriminatory pricing from Chapter 3. We also consider relaxed constraints of fair price discrimination, so as to contemplate also other recent related works in the scientific literature.

Chapter 2

Social Envy-Freeness

We consider a market setting in which buyers are individuals of a population, whose relationships are represented by an underlying social graph. Given buyers' valuations for the items being sold, an outcome consists of a pricing of the objects and an allocation of bundles to the buyers. An outcome is social envy-free if no buyer strictly prefers the bundles of her neighbors in the social graph. We focus on the revenue maximization problem in multi-unit markets, in which there are multiple copies of the same item being sold, and each buyer is assigned a subset of identical items. We consider four different cases that arise when considering two different buyers valuations, i.e., single-minded or general, and by adopting two different forms of pricing, that is, item- or bundle-pricing. For all the above cases we show the hardness of the revenue maximization problem and give corresponding approximation results. All our approximation bounds are optimal or nearly optimal. Moreover, under the assumption of social graphs of bounded treewidth, we provide an optimal allocation algorithm for general valuations with item-pricing. Finally, we determine optimal bounds on the corresponding price of envy-freeness, that is, on the worst-case ratio between the maximum revenue that can be achieved without envy-freeness constraints, and the one obtainable in case of social envy-freeness. Some of our results close hardness open questions or improve already known ones in the literature concerning the standard setting without sociality.

2.1 Introduction

In this chapter we consider a market setting in which buyers are individuals of a population, whose relationships are represented by an underlying social graph. A single seller must decide how to sell a set of identical items to a group of buyers who openly declare their valuations for bundles of items. In such a setting we define an outcome to be social envy-free if no buyer strictly prefers the bundles allocated to her neighbors in the social graph.

The issue of modeling the locality of mutual influences in game theory has already been considered in graphical games [74] and explicitly taken into account in [15, 16], where the authors introduced the existence of a social graph of the players, under the assumption that the payoff of each player is affected only by the strategies of the adjacent ones, representing somehow her neighborhood. Similarly, works in fair allocation of goods in absence of pricing assumed an individual's view of the subjective well-being based on a comparison with peers, that is, restricting (pair) envy-freeness constraints to social neighbors [1, 39]. In this setting, the price of envy-freeness has been defined as the worst-case ratio between the total utility reachable by any allocation of goods and the one that can be achieved satisfying envy-freeness constraints [1].

The price of envy-freeness has been defined in [25] in the context of fair allocation of divisible and indivisible goods without pricing, and suitably bounded under different assumptions. It was implicitly considered also in pricing problems in [52, 54] with respect to price discrimination. Such a notion was extended in [1] to the social case, that is, when (pair) envy-freeness constraints must be satisfied only between neighbors in a social graph.

Distributed mechanisms for allocating indivisible goods under the absence of central control were investigated in [37–39]. In such a setting, given an underlying social structure, agents can locally agree on deals to exchange some of the goods in their possession, again under the assumption of envy-freeness restricted only to social neighbors. The authors studied the convergence properties of such distributed mechanisms when used as fair division procedures.

Finally, [17, 18] considered the possibility of limiting the view of some buyers, not admitting them in the market. This might be seen as related to a social envy-freeness setting in which such buyers are isolated in the social graph.

2.1.1 Summary

We focus on the notion of sociality in the pricing problem and investigate its effects in multi-unit markets. Namely, we consider the *social envy-freeness* setting in which buyers are members of a social population, whose relationships are modeled by means of an undirected *social graph*. In such a graph, nodes represent buyers and edges represent mutual knowledge between the corresponding endpoints. Each buyer can be envious only of the bundles received by her neighbors. We notice that *social envy-freeness* is a relaxation of the notion of *pair envy-freeness*, which is in this respect a relaxation of *envy-freeness*. Thus, if we consider the spaces of pricings and allocations, according to these solution concepts we can highlight the following hierarchy:

$$\textit{envy-free} \subseteq \textit{pair envy-free} \subseteq \textit{social envy-free}.$$

We focus on multi-unit markets, that is, on the problem of pricing and allocating m identical items to n different buyers, so that each buyer is assigned a given bundle or subset of item copies. As already observed, this particular setting, in which the valuation of each buyer depends only on the number of goods she receives, is well suited to model all of those real-world scenarios in which the items on sale are homogeneous, like in commodity markets.

Besides investigating the time complexity of determining social envy-free outcomes which maximize revenue, and good approximate solutions, we investigate the increase of revenue due to the incomplete knowledge of buyers (with respect to pair envy-freeness) providing proper bounds on the price of envy-freeness. Such a quantity is defined as the worst-case ratio between the maximum revenue that can be achieved without envy-freeness constraints, that is, without social relationships between the buyers, and the one obtainable in case of social relationships that require the fulfillment of the respective envy-freeness constraints.

We consider two different assumptions on buyers' valuations: *single-minded*, in which each buyer has a strictly positive valuation only for bundles of a given fixed size, called preferred bundles, and *general valuations*, in which she has a different unrestricted valuation for each possible size. We remark that, since in combinatorial auctions the number of possible bundles is exponential in the number of items on sale, single-minded buyers have been largely investigated as a more tractable subproblem still of practical interest [10, 33, 34, 81, 83, 85]. When focusing on the multi-unit case, they correspond to the well-known knapsack auctions proposed in [3]. In fact, in the literature such valuations are usually called knapsack valuations, but in this thesis we will stick to referring to them as single-minded, since agents are desiring subsets of a single cardinality.

For what concerns pricing, we consider two kinds of assignment policies: the non-discriminatory *item-pricing*, where a unique price p equal for the all identical items must be set, and the mildly discriminating *bundle-pricing*, where the seller is allowed to assign a different (non-proportional) price to each bundle size.

For all the four different cases that arise, we show the following results on the revenue maximization problem (see also Tables 2.1-2.2):

i. *Single-minded, item-pricing*: We show that the problem is NP-hard for standard pair envy-freeness, and give a corresponding FPTAS, that is, an algorithm able to provide a solution whose revenue is at least a $(1+\varepsilon)$ approximation of the maximum achievable one for any fixed $\varepsilon > 0$, in time polynomial in $1/\varepsilon$ and in the size of the input instance. Similarly, we prove the strong NP-hardness of the problem for social envy-freeness and provide a corresponding PTAS, that is, again an algorithm able to return a $(1+\varepsilon)$ approximation of the maximum achievable revenue for any fixed $\varepsilon > 0$, but in time polynomial only in the size of the input instance and no longer in $1/\varepsilon$. Clearly, all such approximation bounds are optimal.

ii. *Single-minded, bundle-pricing*: We show that the problem is NP-hard both for pair and social envy-freeness, and give two corresponding FPTASs.

iii. *General valuations, item-pricing*: We provide a polylogarithmic lower bound on the achievable approximation ratio for pair envy-freeness (and thus also for social envy-freeness), while the $O(\log n)$-approximation algorithm provided in [82] for pair envy-freeness directly extends to social envy-freeness. In fact, such an algorithm actually provides a logarithmic fraction of the optimal revenue that can be obtained even ignoring envy-freeness constraints. As in [23, 82], our hardness result relies on a weaker conjecture with respect to $P \neq NP$, i.e., on the $R3SAT$-hardness of the problem (see [50]). Moreover, we give an optimal allocation algorithm for social graphs with bounded treewidth. We observe that, even if this class of graphs is not typical in social networks, it allows to treat in a unified setting interesting specific social scenarios, such as spatial proximity along a line (paths) and hierarchical relationships (trees).

iv. *General valuations, bundle-pricing*: In this case, while a polylogarithmic lower bound on the achievable approximation ratio was already known [82], we give an $O(\log n)$-approximation algorithm for social (and thus also pair) envy-freeness, improving upon the previous $O(\log n \cdot \log m)$ bound for pair envy-freeness given in [82]. When considering specific social topologies, we show that the problem is APX-hard even in the case of empty social graphs, i.e., with all nodes isolated, and ignoring the supply constraint. In fact, in this case the revenue maximization problem is equivalent to the so-called Max-Buying problem, that in [2] has been proven APX-hard and approximable within ratio 1.59. By the equivalence, such an approximation directly applies also to revenue maximization.

v. Finally, for all the above cases we provide optimal bounds on the price of envy-freeness (Table 2.3). We remark that, while the upper bounds concern every possible social graph topology, the lower bounds hold even in the very restricted case of paths or chains of nodes.

As already mentioned, some of our results close hardness open questions or improve already known approximation ones concerning the standard pair envy-freeness setting.

Before concluding the subsection, in order to better understand the nature of our results and their relationship with the previous literature, let us emphasize some important aspects.

	Single-minded		
	Standard	Social	Standard with Free Disposal
Item-pricing	**NP-hard**	**NP-hard (strong)**	NP-hard [1]
	FPTAS	**PTAS**	~~$O(\log n)$~~ **$(2+\varepsilon)$**
Bundle-pricing	**NP-hard**		NP-hard [1]
	FPTAS		FPTAS [1]

Table 2.1: Hardness and approximation results in the case of single-minded buyers (in bold those obtained in the present work).

	General Valuations		
	Standard	Social	Specific topologies
Item-pricing	**$\Omega(\log^{\epsilon} n)$**		**polytime - bounded treewidth**
	$O(\log n)$ [1]		
Bundle-pricing	$\Omega(\log^{\epsilon} n)$ [1]		**APX-hard - empty graph,unlim.supply**
	~~$O(\log n \log m)$~~ [1] **O(log n)**		**1.59 - empty graph, unlim.supply**

Table 2.2: Hardness and approximation results in the case of buyers with general valuations (in bold those obtained in the present work).

	Single-minded	General Valuations
Item-pricing	**2**	**Θ(log n)**
Bundle-pricing	**1**	**Θ(log n)**

Table 2.3: Price of envy-freeness bounds. The upper bounds hold with respect to any social graph, the lower bounds even for paths.

Free Disposal Differently from [82], where single-minded buyers under pair envy-freeness have also been considered, we do not assume the free disposal feature, in which buyers have the same valuation for all the bundles of size at least equal to their preferred one. Non-free disposal is realistic in cases in which the shrink-wrapped bundles exceed a given allowance of volume or capacity and cannot be partitioned, so that the utility of receiving a bigger bundle is completely lost. Another example is when the surplus items in the bundle must be disposed of, offsetting the utility that would be achieved with only the strictly needed ones. Finally, non-free disposal yields a direct correspondence of our model with the knapsack auctions proposed in [3], correspondence that otherwise is lost.

[1]These results first appeared in (or directly follow from) [82]

Even if free disposal might appear a subtle assumption, as described in more detail in Subsection 2.3.3, it significantly changes the combinatorial structure of the problem, making our results incomparable. In particular, the NP-hardness for free disposal does not extend to non-free disposal and vice versa. Moreover, neither free disposal approximation algorithms directly yield to corresponding non-free disposal ones nor the other way around. As a matter of fact, we also provide a $(2 + \varepsilon)$-approximation algorithm for any $\varepsilon > 0$ for free disposal under pair envy-freeness and item-pricing, which is quite different in structure with respect to the non-free disposal algorithms. This improves upon the $O(\log n)$ approximation of [82]. For bundle-pricing an FPTAS was already given in [82], but it does not apply to non-free disposal.

Input Size It is important to remark that, while the size of the representation of the instances with general valuations is polynomial in m, as different valuations must be specified for different bundle sizes, in the single-minded case the dependence is logarithmic in m, as for each buyer it is sufficient to specify just the size of her preferred bundle, together with the corresponding valuation. As a consequence, in a bit counter-intuitive way, approximation algorithms for general valuations do not directly give corresponding ones for single-minded buyers, since their time complexity might be not polynomial in $\log m$. In a similar way, single-minded hardness results do not apply to general valuations, as the impossibility of being polynomial in $\log m$ does not imply the impossibility of being polynomial in m.

Pair Envy-Freeness versus Social Envy-Freeness A final remark concerns the relationships between these two notions of envy-freeness and the mutual implications of the corresponding results.

As pointed out in the introduction, pair envy-freeness is equivalent to social envy-freeness in complete social graphs. As a consequence, hardness results for pair envy-freeness directly imply corresponding ones for social envy-freeness by restriction on complete social graphs. Such hardness results might not be tight, though. An example of this will be given in the sequel for single-minded buyers and item-pricing, where the NP-hardness holds in the strong sense only for social envy-freeness.

On the other hand, any r-approximation algorithm A for social envy-freeness under arbitrary social graphs yields in a direct way an r-approximation algorithm for pair-envy freeness, simply running A on complete social graphs. Therefore, social envy-freeness approximation results can only be worse with respect to pair envy-freeness ones, and the only way to provide better bounds is that of restricting to specific social graph topologies. As an example, in the sequel we provide an optimal algorithm for bounded treewidth social graphs, general valuations and item-pricing, which is a hard problem in the unrestricted case.

As a final remark, like in previous papers, some of our r-approximation algorithms actually guarantee an r fraction of the revenue of the optimal outcomes that can be obtained by ignoring envy-freeness constraints, or equivalently of social envy-free outcomes on empty graphs, while at the same time satisfying pair envy-freeness, or social envy-freeness in complete graphs. Such algorithms clearly directly yield r-approximations for arbitrary social graphs.

2.2 Preliminaries

A multi-unit market μ can be represented by a tuple $(N, M, (v_i)_{i\in N})$, where $N = \{1, \ldots, n\}$ is a set of n buyers, M is a set of m identical items and for every buyer $i \in N$, $v_i = (v_i(1), \ldots, v_i(m))$ 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 would be willing to pay for buying 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 that bundle size, and *general valuations*, i.e., the unrestricted case. In the sequel, when dealing with single-minded valuations, we call *preferred* by a buyer i the unique bundle for which i has a strictly positive valuation, and denote by m_i its size. Moreover, we let n_j be the number of buyers with preferred bundles of size j and $J = \{j|n_j > 0\}$.

A *price vector* is an m-tuple $\overline{p} = (\overline{p}(1), \ldots, \overline{p}(m))$ such that, for every j, $1 \leq j \leq m$, $\overline{p}(j) \geq 0$ is the price of a bundle of size j. Given a price vector $\overline{p}$ and a set of items $X \subseteq M$, $u_i(X, \overline{p}) = v_i(|X|) - \overline{p}(|X|)$ is the utility of buyer i when buying X.

Since items in M are identical, we consider the following two different pricing schemes, called *item-pricing* and *bundle-pricing*, respectively. In the former, the seller must assign a single non-negative price $p \in \mathbb{R}$ to all the identical items, so that the price owed by each buyer for a bundle X is $\overline{p}(|X|) = |X| \cdot p$. In the latter, the seller has the freedom to give different (non-proportional) prices $\overline{p}(j) \in \mathbb{R}$ to bundles of size j. Therefore, the only constraint is that the prices owed by buyers receiving bundles of the same size must be coincident. In the following, in item-pricing we denote an outcome simply as $(\overline{X}, p)$, that is, by specifying the single price assigned to each of the identical items.

An *allocation vector* is an n-tuple $\overline{X} = (X_1, \ldots, X_n)$ such that $X_i \subseteq M$ is the set of items sold to buyer i. Notice that, in the case of multi-unit markets, it would also be possible to define an allocation as a tuple $\overline{X} = (x_1, \ldots, x_n)$ where $x_i \in \mathbb{N}$ is the number of units assigned to buyer i. Nonetheless, throughout the present work we will use the former definition because of its connection to the original setting of combinatorial markets.

A *feasible* outcome of market μ is a pair $(\overline{X}, \overline{p})$ satisfying the following conditions:

1. supply constraint: $\sum_{i=1}^{n} |X_i| \leq m$;
2. individual rationality: $u_i(X_i, \overline{p}) \geq 0$ for every $i \in N$.

We assume that buyers in N are individuals of a population, whose relationships are represented by an underlying undirected social graph $G = (N, E)$. In such a setting, given an outcome $(\overline{X}, \overline{p})$, each buyer $i \in N$ is aware only of the bundles assigned to the other buyers she knows, that is, belonging to the subset $N(i) = \{j \in N | \{i, j\} \in E\}$ of her neighbors in G.

Definition 2.1. A feasible outcome $(\overline{X}, \overline{p})$ for market μ is *social envy-free* or simply *stable* under G if $u_i(X_i, \overline{p}) \geq u_i(X_j, \overline{p})$ for every buyer $i \in N$ and $j \in N(i)$.

Thus, an outcome is stable if no buyer strictly prefers the bundles assigned to the buyers she knows. Notice that, if G is complete, the above definition corresponds to the standard notion of pair envy-freeness defined in the literature.

For a market μ and an outcome $(\overline{X}, \overline{p})$ associated with μ, the revenue raised by the seller is given by $r(\overline{X}, \overline{p}) = \sum_{i=1}^{n} \overline{p}(|X_i|)$. The *pricing problem* for μ consists in determining an outcome $(\overline{X}, \overline{p})$ of maximum revenue that is, stable under G.

Let $opt(\mu, G)$ be the maximum possible revenue achievable by a stable outcome for μ under G and $opt(\mu)$ be the highest possible revenue achievable by a feasible allocation for μ without considering envy-freeness constraints.

Definition 2.2. Given a set of market instances $\mathcal{M}$ and a family of social graphs $\mathcal{G}$, the *price of envy-freeness* $c(\mathcal{M}, \mathcal{G})$ of $\mathcal{M}$ and $\mathcal{G}$ is the worst-case ratio between the maximum revenue that can be achieved in the markets in $\mathcal{M}$ without considering envy-freeness constraints, and the one induced by the outcomes that are stable according to the social graphs in $\mathcal{G}$, that is, $c(\mathcal{M}, \mathcal{G}) = \sup_{\mu \in \mathcal{M}, G \in \mathcal{G}} \frac{opt(\mu)}{opt(\mu, G)}$.

In a sense, the price of envy-freeness can be seen as related to the classical price of anarchy notion, since it quantifies the loss of revenue that one has to pay to enforce stability according to buyers' social relationships. For the sake of brevity, in the following we call (SINGLE,ITEM)-pricing (resp. (GENERAL,ITEM)-, (SINGLE,BUNDLE)- and (GENERAL,BUNDLE)-pricing) the standard pricing problem restricted to the instances of multi-unit markets with single-minded valuations and item-pricing (resp. general valuations and item-pricing, single-minded valuations and bundle-pricing, and general valuations and bundle-pricing). Moreover, we will call such problems *social*, when considering a social graph of knowledge of the buyers. So for instance, in the social (SINGLE,ITEM)-pricing problem, we are given in input a single-minded multi-unit market μ and a social graph G and we want to determine a revenue-maximizing outcome with item-pricing for μ which is stable under G.

Clearly, since the standard problem corresponds to the restriction of the social problem to the instances in which G is fixed to be a complete graph, every hardness result concerning the standard problem extends to the social version, while every approximation algorithm for the social problem also applies to the standard version.

We will often reduce the pricing problem to a variant of the knapsack problem called MULTIPLE-CHOICE KNAPSACK. In such a problem, we are given t classes $\{O_1, \dots, O_t\}$ of objects to pack in a knapsack of capacity k. Each object $o_{j,h} \in O_j$ has a profit $z_{j,h}$ and a weight $w_{j,h}$, and we must pick exactly 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 [80], MULTIPLE-CHOICE KNAPSACK is NP-hard, but it admits an FPTAS. Notice that, by adding in each O_j a dummy object of profit 0 and weight 0, we can get an equivalent version of the problem in which we must pick *at most* one object per class. In the present work, unless explicitly stated, we will always refer to this alternative formulation.

2.3 Single-Minded Valuations

In this section we consider single-minded buyers in multi-unit markets.

2.3.1 Item-Pricing

Let us first focus on item-pricing. The following fact will be useful in the sequel.

Lemma 2.3. *In the case of single-minded buyers, the price* p^{opt} *of an optimal stable item-pricing outcome* $(\overline{X}^{opt}, \mathrm{p}^{opt})$*, that is, maximizing the seller's revenue, belongs to the set* $\mathbb{P} := \{\frac{v_i(m_i)}{m_i} | i \in N\}$*, where* m_i *is the size of the preferred bundle of buyer* i*.*

Proof. Assume that $(\overline{X}^{opt}, p^{opt})$ with $p^{opt} \notin \mathbb{P}$ is an optimal stable outcome and let $p = \frac{v_i(m_i)}{m_i} \in \mathbb{P}$ be the minimum ratio in $\mathbb{P}$ greater than p^{opt}. Consider then the outcome $(\overline{X}^{opt}, p)$.

Since the allocations of $(\overline{X}^{opt}, p)$ and $(\overline{X}^{opt}, p^{opt})$ are coincident, $(\overline{X}^{opt}, p)$ still satisfies the supply constraints. Furthermore, it also satisfies the individual rationality, since the price is increased only until the minimum $\frac{v_i(m_i)}{m_i}$ greater than the optimal price p^{opt}, and thus no utility of an allocated buyer can become negative. Finally $(\overline{X}^{opt}, p)$ is social envy-free, because in $(\overline{X}^{opt}, p^{opt})$ every buyer receives her favorite bundle among the allocated ones and this can change in $(\overline{X}^{opt}, p)$ only if some allocated buyer starts strictly preferring the empty bundle due to the rise of the price. However, as just mentioned before, this cannot be the case.

In conclusion, $(\overline{X}^{opt}, p)$ is a stable outcome with a strictly higher revenue than $(\overline{X}^{opt}, p^{opt})$ a contradiction. □

Having claimed the above fact, let us first focus on the traditional market scenario in which all the buyers can envy all the others, that is, with a complete social graph K_n. Under such an assumption the following result holds, which is similar to the one provided in [82] for a slightly different problem (single-minded instances with free disposal under envy-freeness).

Theorem 2.4. *The* (SINGLE,ITEM)*-pricing problem is NP-hard.*

Proof. We prove this via a reduction from the NP-hard problem SUBSET-SUM in which, given a subset of natural numbers $A \subseteq \mathbb{N}$ and $k \in \mathbb{N}$, we have to find a subset $S \subseteq A$ such that $\sum_{a_i \in S} a_i = k$. Given an instance $(A = \{a_1, \ldots, a_n\}, k)$ of SUBSET-SUM, we construct the following instance μ of the (SINGLE,ITEM)-pricing problem:

- $N = \{1, \ldots, n\}$;
- $M = \{1, \ldots, k\}$;
- $v_i(a_i) = a_i$ and $v_i(j) = 0$ for $j \neq a_i$.

In such a reduction, a subset $S \subseteq A$ can be associated with an outcome of μ in which buyer i gets a bundle of size a_i if $a_i \in S$, and vice versa every outcome has a corresponding subset S. By Lemma 2.3, any optimal outcome for μ must assign price 1 to all the items and thus is also stable, since it provides utility 0 to all the buyers. Therefore, there exists a subset $S \subseteq A$ of overall sum k if and only if the maximum revenue of a stable outcome for μ is k. □

On the other hand, the problem can be well approximated.

Algorithm 1: FPTAS for the (SINGLE,ITEM)-pricing problem

Input: An instance $\mu = (N, M, (v_i)_{i \in N})$ of the (SINGLE,ITEM)-pricing problem, accuracy parameter ε.
Output: An item-pricing outcome $(\overline{X}, p)$.
$S = \emptyset$;
for *each price* $p \in \mathbb{P} := \{\frac{v_i(m_i)}{m_i}, i \in N\}$ **do**
  N_j^+ = set of buyers having strictly positive utility with price p for a bundle of size j ;
  N_j^0 = set of buyers having utility 0 with price p for a bundle of size j;
  $n_j^+ = |N_j^+|$ and $n_j^0 = |N_j^0|$;
  Construct the following instance $K(\mu) = (\overline{O}, \overline{z}, \overline{w}, k)$ of MULTIPLE-CHOICE KNAPSACK:
    $k = m$;
    $O_j = \{o_{j,0}, \ldots, o_{j,n_j^0}\}$ for bundle j preferred by at least one buyer;
    $z_{j,h} = w_{j,h} = j(n_j^+ + h)$ for $h = 0, \ldots, n_j^0$;
  $T \longleftarrow FPTAS-$MULTIPLE-CHOICE KNAPSACK$(K(\mu), \varepsilon)$;
  for *each* $o_{j,h} \in T$;
  do
    Let $H \subseteq N_j^0$ be a set of h buyers having utility 0 with price p for bundle j;
    Assign a bundle X_i of size j to each buyer $i \in N_j^+ \cup H$;
  end
  $S = S \cup \{(\overline{X}, p)\}$;
end
return $(\overline{X}, p) = argmax_{(\overline{X}', p') \in S}\ r(\overline{X}', p')$.

Theorem 2.5. *The* (SINGLE,ITEM)*-pricing problem admits an FPTAS.*

Proof. Consider Algorithm 1 in the figure.

According to Lemma 2.3, for each candidate optimal price $p = \frac{v_i(m_i)}{m_i}$, $i \in N$, we construct a proper instance $K(\mu) = (\overline{O}, \overline{z}, \overline{w}, k)$ of MULTIPLE-CHOICE KNAPSACK as follows.

Let N_j^+ (resp. N_j^0) be the set of buyers having strictly positive utility (resp. utility equal to 0) with price p for the bundle of size j, and let $n_j^+ = |N_j^+|$ (resp. $n_j^0 = |N_j^0)|$) be its cardinality.

We let the size of the knapsack be $k = m$, and we associate a class $O_j = \{o_{j,0}, \ldots, o_{j,n_j^0}\}$ of $n_j^0 + 1$ objects to each bundle size j preferred by at least one buyer, with $z_{j,h} = w_{j,h} = j \cdot (n_j^+ + h)$ for each $o_{j,h} \in O_j$. In the algorithm, picking an object $o_{j,h} \in O_j$ corresponds to assigning a bundle of size j to all the buyers in N_j^+ and to h buyers in N_j^0. Notice in fact that, by the stability constraint, if a bundle of size j is assigned, then all buyers in N_j^+ must receive it, plus any subset of buyers in N_j^0.

Therefore, any feasible solution T for the instance $K(\mu)$ of MULTIPLE-CHOICE KNAPSACK is naturally associated by the algorithm to a stable outcome $(\overline{X}, p)$ for the initial instance μ of the pricing problem. In fact, since T doesn't exceed the knapsack capacity, that is, equal to the number of items on sale, $(\overline{X}, p)$ doesn't sell more then m items. Moreover, $(\overline{X}, p)$ is stable, since if a bundle of size j is assigned, then all buyers in N_j^+ who might be envious receive it. Finally, by construction $(\overline{X}, p)$ has revenue $r \cdot p$, where r is the profit of T. Thus, any $(1-\varepsilon)$-approximate solution T for $K(\mu)$ provides a $(1-\varepsilon)$-approximate solution $(\overline{X}, p)$ of the optimal stable outcome of μ assigning price p to the items in M. The theorem then follows by observing that, according to Lemma 2.3, the algorithm considers the price of an optimal outcome in at least one iteration, the FPTAS for MULTIPLE-CHOICE KNAPSACK returns in polynomial-time in $1/\varepsilon$ and the size of μ a $(1-\varepsilon)$-approximate solution for each fixed price, and finally Algorithm 1 outputs the outcome of maximum revenue obtained in all the iterations. □

Example: Consider the market μ with 6 single-minded buyers and supply m equal to 10. The valuation functions of the buyers are:

$$v_1(4) = 4 \quad v_2(3) = 4 \quad v_3(3) = 3$$
$$v_4(2) = 3 \quad v_5(2) = 2 \quad v_6(2) = 2$$

The set of the candidate optimal prices is $\mathbb{P} = \left\{\frac{3}{2}, \frac{4}{3}, 1\right\}$.

An example of construction of MULTIPLE-CHOICE KNAPSACK for $p = 1$, that is, $K(\mu) = (\overline{O}, \overline{z}, \overline{w}, 10)$, is given in the following table:

O_4		O_3		O_2	
$o_{4,0}$	$z_{4,0} = w_{4,0} = 0$	$o_{3,0}$	$z_{3,0} = w_{3,0} = 3$	$o_{2,0}$	$z_{2,0} = w_{2,0} = 2$
$o_{4,1}$	$z_{4,1} = w_{4,1} = 4$	$o_{3,1}$	$z_{3,1} = w_{3,1} = 6$	$o_{2,1}$	$z_{2,1} = w_{2,1} = 4$
				$o_{2,2}$	$z_{2,2} = w_{2,2} = 6$

The optimal allocation solution of $K(\mu)$ is $\{o_{3,1}, o_{2,1}\}$, corresponding to the following allocation $\overline{X}$:

$$|X_1| = 0 \quad |X_2| = 3 \quad |X_3| = 3$$
$$|X_4| = 2 \quad |X_5| = 2 \quad |X_6| = 0$$

□

For the social version of the problem the following stronger hardness result holds.

Theorem 2.6. *The social* (SINGLE,ITEM)*-pricing problem is strongly NP-hard.*

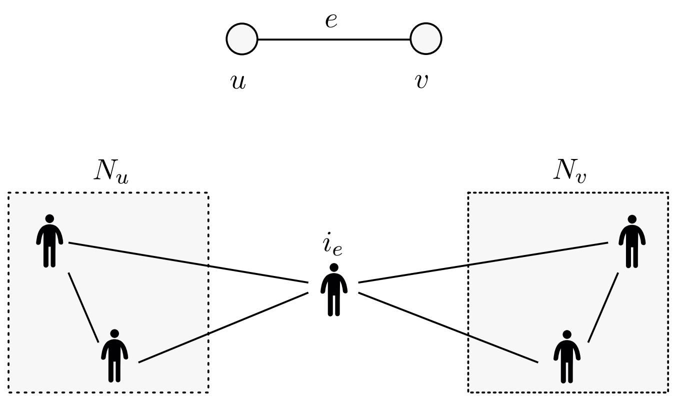

FIGURE 2.1: A simple example of a reduced instance for social (SINGLE,ITEM)-pricing problem.

Proof. In order to prove the claim, we provide a polynomial-time reduction from DENSEST K-SUBGRAPH. An instance of such a problem is given by an undirected graph $H = (V, F)$ and an integer k, and we want to find a subset $S \subseteq V$ of cardinality $|S| \leq k$ that maximizes the number of edges in the subgraph induced by S. Given an instance of DENSEST K-SUBGRAPH, we construct an instance (μ, G) of the social (SINGLE,ITEM)-pricing problem as follows.

In the market μ:

- there is a set N_u of $|F| + 1$ buyers to each $u \in V$, in such a way that $v_i(1) = 1 + \varepsilon$ for each $i \in N_u$, where $\varepsilon = \frac{k}{|V|+1}$;
- for each $e \in F$ there is a buyer i_e with $v_{i_e}(1) = 1$;
- there is a distinguished buyer w with $v_w(|V|(|F| + 1)) = |V|(|F| + 1)$;
- there are $m = (|V| + k)(|F| + 1) + |F|$ items.

In the social graph $G = (N, E)$:

- there is an edge $\{i, i'\} \in E$ for every pair of buyers i, i' such that $i, i' \in N_u$ for some $u \in V$;
- there is an edge $\{i, i_e\} \in E$ for every $i \in N_u$ and edge e incident to u in H.

We now show that H admits a subgraph of size k with h edges if and only if the reduced instance has an outcome $(\overline{X}, p)$ stable under G of revenue $r(\overline{X}, p) = (|V| + k)(|F| + 1)+$ h.

($\Rightarrow$) Assume that $S \subseteq V$ is such that $|S| = k$ and the graph induced by S has h edges. Consider the outcome $(\overline{X}, p)$ for the reduced instance where $p = 1$ and bundles are allocated to the buyers in $\{w\} \cup \bigcup_{u \in S} N_u \cup \{i_e \in N \mid e = \{u, z\},\ u, z \in S\}$, i.e., the buyer w, buyers corresponding to vertices in S, and buyers associated to edges between such vertices.

Notice that all the buyers i_e to whom no item is allocated have utility 0 for their preferred bundle and thus cannot be envious. Moreover, the remaining not allocated buyers, those belonging to the subsets N_u with $u \notin S$, do not have any neighbor in G receiving a bundle, and thus they cannot envy as well. Therefore, $(\overline{X}, p)$ is stable and its achieved revenue, which corresponds to the number of assigned items, is equal to $(|V| + k)(|F| + 1) + h$.
($\Leftarrow$) Assume on the contrary that there exists a stable outcome $(\overline{X}, p)$ for the reduced instance of revenue $(|V| + k)(|F| + 1) + h$.

In $(\overline{X}, p)$, price p cannot be strictly greater than $1 + \varepsilon$, otherwise we would have revenue 0. If $p > 1$ it is possible to assign bundles only to the buyers in the sets N_u, so that the achieved revenue can be at most $(1 + \varepsilon)[|V|(|F| + 1)] < (|V| + k)(|F| + 1) + h$, as $\varepsilon = \frac{k}{|V|+1}$. Therefore, it must be $p \leq 1$. In this case, in order to have $r(\overline{X}, p) = (|V| + k)(|F| + 1) + h$, we must assign to the distinguished buyer w her preferred bundle. By the envy-freeness constraints, if a buyer in a set N_u gets her bundle, then all the other buyers in N_u must get theirs. Moreover, by the supply constraint, items can be allocated to buyers in at most k of the subsets N_u, $u \in V$, and in order to achieve the desired revenue to at least k of them. So, in $(\overline{X}, p)$, preferred bundles are allocated to w, to all the buyers in exactly k of the subsets N_u with $u \in V$, and to h buyers i_e. Let us then consider the solution S for the instance of DENSEST K-SUBGRAPH such that $S = \{u \mid \forall i \in N_u,\ |X_i| = 1\}$. Since buyers of exactly k sets N_u are allocated, S is feasible. Thus, it remains to show that in the subgraph induced by S there are at least h edges. In $(\overline{X}, p)$ we allocate bundles to h buyers i_e. Moreover, as already remarked, if a buyer i_e with $e = \{u, z\}$ receives her preferred item, then also all the buyers in set N_u and N_z must have theirs, so that $u, z \in S$. In conclusion, each allocated buyer i_e corresponds to an edge e in the subgraph induced by S, which in turn has h edges. □

By the above theorem, an FPTAS for the problem does not exist, unless P=NP. On the other hand, we now show an optimal result concerning the approximability, that is, the existence of a PTAS.

Before providing the algorithm, let us describe the main ideas involved.

Proof idea. Given an instance (μ, G) of the pricing problem, a fixed price p, $j \leq m$ and $h \leq n$, assume we can efficiently solve the *restricted subproblem* of determining a feasible allocation of bundles of size j to h buyers stable under G, if it exists. Recalling that J is the set of at most n bundle sizes for which there exists at least one buyer with a preferred bundle of that size, under the above assumption we could even obtain an FPTAS, again relying on a proper instance $K(\mu, G, p)$ of MULTIPLE-CHOICE KNAPSACK:

- the knapsack capacity is fixed to $k = m$;
- for each bundle size $j \in J$ there is a class $O_j = \{o_{j,h}$: there exists an allocation of bundles of size j to h buyers stable under $G\}$;
- $z_{j,h} = w_{j,h} = j \cdot h$.

Clearly, such an instance of MULTIPLE-CHOICE KNAPSACK has size polynomial in the one of (μ, G), as there are at most n classes, each of at most n objects. Any solution S for the instance of value r can be associated to an outcome of revenue $r \cdot p$ simply by adding in the outcome, for each object $o_{j,h} \in S$, the stable allocation under G of bundles of size j to h buyers. Therefore, by Lemma 2.3, running the FPTAS for all the candidate optimal prices $p = \frac{v_i(m_i)}{m_i}$, $i \in N$, and selecting the best returned outcome in terms of revenue, we get a $(1 - \varepsilon)$-approximation of the optimum solution in time polynomial in the size of (μ, G) and in $1/\varepsilon$, i.e., an FPTAS for the social (SINGLE,ITEM)-pricing problem.

Unfortunately, according to Theorem 2.6, the basic assumption underlying the algorithm cannot be feasible, that is, the restricted subproblem cannot be solved in polynomial time (as it would imply an FPTAS for a strongly NP-hard problem). However, the following approximation result holds.

Lemma 2.7. *Given an instance (μ, G) of the social* (SINGLE,ITEM)*-pricing problem, a fixed price p, $j \leq m$ and $h \leq n$, the problem of finding a revenue-maximizing feasible allocation of bundles of size j to at most h buyers stable under G admits a PTAS.*

Proof. Let N_j^+ (resp. N_j^0) be the set of buyers with strictly positive utility (resp. with utility equal to 0) for a bundle of size j when the price of the items is fixed to p, and let G_j^+ be the subgraph of G induced by N_j^+.

Any optimal outcome $(\overline{X}, p)$ for μ stable under G must be such that, for each connected component of G_j^+, either all the buyers receive a bundle of size j or none of them.

For a fixed $\varepsilon > 0$, call *big* a maximal connected component of G_j^+ containing at least $\varepsilon \cdot h$ buyers, and *small* the remaining (maximal) connected components. Then, any optimal outcome $(\overline{X}, p)$ allocating bundles of size j to at most h buyers, can assign bundles to at most $\lfloor \frac{1}{\varepsilon} \rfloor$ big components. In addition, it can allocate to some small components, and finally to some buyers of utility 0 in N_j^0 that are adjacent in G only to all the allocated components in G_j^+.

Consider then the following procedure. For each of the at most $O(n^{\lfloor \frac{1}{\varepsilon} \rfloor})$ subsets of big components, first allocate bundles to all the buyers belonging to such components; proceed by allocating bundles to small components, until either we cannot assign another component without exceeding h buyers, or all the small components have been allocated. In the first case, since no small component fits into the limit of h allocated buyers, the overall number of allocated buyers is at least $(1-\varepsilon)h$. In the second case, continue allocating to buyers of utility 0 in N_j^0 until either allocating to h buyers in total, thus determining the optimal allocation of bundles of size j, or all the buyers of utility 0 have been allocated. In this last case, again we get the optimal allocation with respect to the initially selected big components, since all the compatible buyers have been allocated.

Therefore, such a procedure provides a $(1-\varepsilon)$-approximation of the best stable allocation of bundles of size j to at most h buyers, and runs in time polynomial in the input size (and exponential in $1/\varepsilon$). □

According to the above lemma, for each fixed price p, we can efficiently determine an instance $K_\varepsilon(\mu, G, p)$ of MULTIPLE-CHOICE KNAPSACK which suitably approximates $K(\mu, G, p)$:

- the knapsack capacity is fixed to $k = m$;
- for each bundle size $j \in J$ there is a class $O_j = \{o_{j,l}$: the above PTAS run on some h, $1 \leq h \leq n$, returns an allocation of bundles of size j to $l \leq h$ buyers stable under $G\}$;
- $z_{j,l} = w_{j,l} = j \cdot l$.

Clearly, if $T = \{o_{j_1,h_1}, \ldots, o_{j_t,h_t}\}$ is a feasible solution for $K(\mu, G, p)$, then there exists a feasible solution $T_\varepsilon = \{o_{j_1,l_1}, \ldots, o_{j_t,l_t}\}$ for $K_\varepsilon(\mu, G, p)$ such that $(1-\varepsilon)h_q \leq l_q \leq h_q$ for each q, $1 \leq q \leq t$. Therefore, if $opt(K)$ is the measure of the optimal solution of $K(\mu, G, p)$ and $opt(K_\varepsilon)$ the one of $K_\varepsilon(\mu, G, p)$, we have that $opt(K_\varepsilon) \geq (1-\varepsilon)opt(K)$. Notice moreover that T_ε is feasible also for $K(\mu, G, p)$ and again yields a corresponding outcome for μ stable under G of proportional revenue.

We are now ready to prove the following theorem.

Theorem 2.8. *The social* (SINGLE,ITEM)*-pricing problem admits a PTAS.*

Algorithm 2: PTAS for the social (SINGLE,ITEM)-pricing problem

Input: An instance (μ, G) of the social (SINGLE,ITEM)-pricing problem, accuracy parameter ε.
Output: An item-pricing outcome $(\overline{X}, p)$.
$S = \emptyset$;
for *each price* $p \in \mathbb{P} := \{\frac{v_i(m_i)}{m_i}, i \in N\}$ **do**
  Compute instance $K_{\varepsilon/2}(\mu, G, p)$ of MULTIPLE-CHOICE KNAPSACK;
  $T \longleftarrow FPTAS-$MULTIPLE-CHOICE KNAPSACK$(K_{\varepsilon/2}(\mu, G, p), \varepsilon/2)$;
  Let $(\overline{X}, p)$ be the outcome corresponding to T;
  for *each* $o_{j,l} \in T$ **do**
    $\overline{X}$ stably assigns a bundle of size j to l buyers (according to PTAS);
  end
  $S = S \cup \{(\overline{X}, p)\}$;
end
return $(\overline{X}, p) = argmax_{(\overline{X}', p') \in S}\ r(\overline{X}', p')$.

Proof. Consider Algorithm 2 in the figure. For all the candidate optimal prices $p \in \mathbb{P}$ established in Lemma 2.3, it constructs a corresponding instance $K_{\varepsilon/2}(\mu, G, p)$ exploiting the PTAS of Lemma 2.7 with accuracy parameter $\varepsilon/2$ and runs on $K_{\varepsilon/2}(\mu, G, p)$ the FPTAS of MULTIPLE-CHOICE KNAPSACK with accuracy $\varepsilon/2$. Among all the returned solutions, it selects the one yielding the outcome $(\overline{X}, p)$ of maximum revenue and provides such an outcome in output.

The complexity of the algorithm is polynomial in the input size (and exponential in $1/\varepsilon$), and recalling that $opt(\mu, G)$ is the revenue of an optimal stable outcome for (μ, G), the revenue of $(\overline{X}, p)$ is $r(\overline{X}, p) = p \cdot (1 - \varepsilon/2) opt(K_{\varepsilon/2}) \geq p \cdot (1 - \varepsilon/2)^2 opt(K) \geq (1 - \varepsilon) opt(\mu, G)$. □

2.3.2 Bundle-Pricing

We now focus on bundle-pricing. First of all, like in item-pricing, the following negative result holds.

Theorem 2.9. *The* (SINGLE,BUNDLE)*-pricing problem is NP-hard.*

Proof. We prove the claim by the same reduction of Theorem 2.4. In fact, again a subset $S \subseteq A$ can be associated with an outcome of μ in which buyer i gets a bundle of size a_i if $a_i \in S$, and vice versa every outcome has a corresponding subset S. Moreover, by the individual rationality constraint, in any optimal stable outcome the price of an allocated bundle of size j can be at most j, and it must be at least j, because otherwise we could get a strictly better stable outcome by raising all the bundles prices to become equal to their cardinalities. Therefore, there exists a subset $S \subseteq A$ of overall sum k if and only if the maximum revenue of a stable outcome for μ is k. □

However, an optimal approximation can be achieved.

Theorem 2.10. *The social* (SINGLE,BUNDLE)*-pricing admits an FPTAS.*

Algorithm 3: FPTAS for the social (SINGLE,BUNDLE)-pricing problem

---

Input: An instance (μ, G) of the social (SINGLE,BUNDLE)-pricing problem,
accuracy parameter ε.
Output: A bundle-pricing outcome $(\overline{X}, \overline{p})$.
for $j \in J$ **do**
  let $i_1, \ldots, i_{n_j}$ be an ordering of the n_j buyers with preferred bundles of size j
  such that $v_{i_1}(j) \geq \ldots \geq v_{i_{n_j}}(j)$;
end
Construct the following instance $K = (\overline{O}, \overline{z}, \overline{w}, k)$ of MULTIPLE-CHOICE KNAPSACK:
  $k = m$;
  $O_j = \{o_{j,i,h} : i \in N, h \in M\}$ for $j \in J$;
  $z_{j,i,h} = v_{j,h}(j) \cdot h$;
  $w_{j,i,h} = j \cdot h$;
$T \longleftarrow$ MULTIPLE-CHOICE KNAPSACK(K, ε);
for *each* $o_{j,i,h} \in T$ **do**
  Set $p(j) = v_{j,h}(j)$;
  Let $|X_i| = j$ for $i = i_{j,1}, \ldots, i_{j,h}$;
end
return $(\overline{X}, \overline{p})$.

Proof. Consider a fixed bundle size $j \leq m$ and let $i_{j,1}, \ldots, i_{j,n_j}$ be an ordering of the buyers with preferred bundles of size j such that $v_{i_{j,1}}(j) \geq \ldots \geq v_{i_{j,n_j}}(j)$.

Given any $h \leq n_j$, setting price $v_{i_{j,h}}(j)$ for bundles of size j allows the stable allocation of h bundles to the first h buyers in the ordering.

Since single-minded buyers can envy only buyers with the same preferred bundle size, the above allocation can be independently performed for each different bundle size with at least one buyer preferring it. Therefore, again the problem reduces to a proper instance of MULTIPLE-CHOICE KNAPSACK given in Algorithm 3, where the insertion of an object $o_{j,i,h}$ in the knapsack corresponds to the stable assignment of h bundles of size j to the first h buyers $i_{j,1}, \ldots, i_{j,h}$ in the ordering. □

Notice that Algorithm 3 is independent of the social graph G, as the returned allocation is stable under the complete graph (and thus under any graph). In fact, as we will show in Section 2.5, social envy-freeness in (SINGLE,BUNDLE)-pricing does not affect the revenue of the optimum stable outcome.

Notice also that, by restriction on complete social graphs, the following result is directly implied by the previous theorem.

Corollary 2.11. *The* (SINGLE,BUNDLE)*-pricing admits an FPTAS.*

2.3.3 Free Disposal

As already observed, so far we did not assume the free-disposal on buyers' valuation function. Under free-disposal each buyer i maintains the same valuation for all the bundles of at least her preferred size m_i, that is, $v_i(j) = v_i(m_i)$ for all $j \geq m_i$. While for the envy-freeness notion this assumption does not affect at all the space of all possible stable allocations, this is not true any longer under pair envy-freeness. In this last case, even if at first glance free disposal might appear a subtle assumption, it seems to influence the underlying combinatorial structure of the problem significantly, so that free disposal results do not extend in a straightforward way to non-free disposal and vice versa. In particular, this makes our contributions incomparable with the ones of [82]. In this paper, the authors provided an FPTAS for all the single-minded cases under free disposal, except for pair envy-freeness and item pricing, for which they gave a $O(\log n)$ approximation.

We now improve the latter result, resorting to a quite different algorithm with respect to the ones given in the previous subsections, achieving a $(2+\varepsilon)$ approximation, for any fixed $\varepsilon > 0$.

Before describing the algorithm in detail and proving its correctness, let us briefly discuss the main underlying ideas.

Proof idea. In order to provide the approximation algorithm, we first focus on the subproblem of finding a pair envy-free allocation for a given price p that maximizes the number of assigned items. This is accomplished by first ignoring the limited supply constraint, and thus possibly allocating a greater number of items by means of an optimal dynamic algorithm. Then, using Lemma 2.12 already given in [82], we can detect a price $p' \geq p$ which yields a pair envy-free allocation respecting the supply constraint, that assigns at least $m/2$ items, thus losing at most a multiplicative factor of 2 in approximation with respect to an optimal allocation.

The following lemma has been proved by Monaco et al. in [82].

Lemma 2.12. *Given a pair envy-free outcome* (X, p) *which does not satisfy the supply constraint, it is possible to find in polynomial time a feasible and pair envy-free outcome* (X', p') *such that* $r(X', p') \geq \frac{m}{2} \cdot p$*.*

The main technique used in [82] to prove the lemma is a careful selection of a subset of the bundles of the original allocation along with an increment of the original price.

A key issue with the non-free disposal setting is that, differently from Lemma 2.3, it is no longer possible to fix a polynomially bounded set $\mathbb{P}$ of candidate optimal prices, that is, containing at least one price p^{opt} yielding an optimal stable allocation $(\overline{X}^{opt}, p^{opt})$. However, by a sampling argument on all the prices that are powers of $(1+\varepsilon/2)$, we can construct a set of prices $\mathbb{P}_\varepsilon$ that is, still polynomially bounded in the input size and in $1/\varepsilon$, while containing at least one price p which is a $(1+\varepsilon/2)$ fraction of an optimal price p^{opt}, i.e., such that $p^{opt}/(1+\varepsilon/2) \leq p \leq p^{opt}$.

The final algorithm works by probing all possible prices $p \in \mathbb{P}_\varepsilon$ and returning the best-determined allocation.

Let us now proceed with the details. For the sake of simplicity, without loss of generality let us assume that the minimum buyers' valuation is 1. In fact, by a simple scaling argument, it is possible to obtain an equivalent instance satisfying such a condition.

Let v_{max} be the maximum buyers' valuation. Given $\varepsilon > 0$, we define:

$$\mathbb{P}_\varepsilon = \left\{(1+\varepsilon/2)^\ell \mid -\log_{(1+\varepsilon/2)} m \leq \ell \leq \log_{(1+\varepsilon/2)} v_{max}\right\}$$

Such a set $\mathbb{P}_\varepsilon$ has cardinality $O\left(\log_{(1+\varepsilon/2)} m + \log_{(1+\varepsilon/2)} v_{max}\right) = O\left(\frac{\log v_{max} + \log m}{\varepsilon}\right)$. Moreover, since the optimal price p^{opt} associated to an optimal stable allocation $(\overline{X}^{opt}, p^{opt})$ must clearly be such that $\frac{1}{m} \leq p^{opt} \leq v_{max}$, $\mathbb{P}_\varepsilon$ contains the above claimed price p such that $p^{opt}/(1+\varepsilon/2) \leq p \leq p^{opt}$.

Let us then define the following subproblem: given a market μ and a price p, find a pair envy-free outcome $(\overline{X}, p)$ maximizing the number of items sold, ignoring the limited supply constraint. We are now going to provide an optimal algorithm (Algorithm 5) for such a subproblem, which relies on dynamic programming. Before describing Algorithm 5 in detail, let us show how it can be exploited to provide the claimed approximation for revenue maximization problem.

Theorem 2.13. *The* (SINGLE,ITEM)*-pricing with free disposal admits a* $(2+\varepsilon)$*-approximation algorithm.*

Proof. Given an instance μ of (SINGLE,ITEM)-pricing with free disposal and the accuracy parameter ε, consider Algorithm 4. Let $(\overline{X}^{opt}, p^{opt})$ be an optimal outcome, and let $p = (1+\varepsilon/2)^{\lfloor \log_{(1+\varepsilon/2)} p^{opt} \rfloor}$ be the price in $\mathbb{P}_\varepsilon$ such that $p^{opt}/(1+\varepsilon/2) \leq p \leq p^{opt}$.

Algorithm 4: A $(2+\varepsilon)$-approximation algorithm for (SINGLE,ITEM)-pricing with free disposal

Input: An instance μ of (SINGLE,ITEM)-pricing with free disposal, accuracy parameter ε.
Output: An item-pricing pair envy free outcome $(\overline{X}, p)$.
$S \leftarrow \emptyset$;
Let $\mathbb{P}_\varepsilon = \left\{(1+\varepsilon/2)^\ell \mid -\log_{(1+\varepsilon/2)} m \leq \ell \leq \log_{(1+\varepsilon/2)} v_{max}\right\}$;
for *each* $p \in \mathbb{P}_\varepsilon$ **do**
  Let $(\overline{X}^p, p)$ be the outcome returned by Algorithm 5 on (μ, p) ;
  Let $(\overline{X}^{p'}, p')$ be feasible outcome extracted from $(\overline{X}^p, p)$ as in Lemma 2.12;
  $S \longleftarrow S \cup \{(\overline{X}^{p'}, p')\}$;
end
$(\overline{X}, p) \longleftarrow \arg\max_{(\overline{X}^{p'}, p') \in S} \{r(\overline{X}^{p'}, p')\}$;
return $(\overline{X}, \overline{p})$;

Notice that, since $p \leq p^{opt}$, the number of items allocated by Algorithm 5 for p is greater or equal to the corresponding one obtained by Algorithm 5 for p^{opt}. More precisely, consider the allocation $\overline{X}'$ where $|X'_i| = |X_i^{opt}|$ if $X_i^{opt} \neq \emptyset$, and $|X'_i|$ is equal to the bundle size that maximizes i's utility for price p among the ones already allocated in $\overline{X}^{opt}$. Clearly, $\overline{X}'$ is pair envy-free under p and allocates at least as many items as $\overline{X}^{opt}$. Thus, by the optimality of Algorithm 5, it returns an allocation $\overline{X}^p$ such that $\sum_{i=1}^n |X_i^p| \geq \sum_{i=1}^n |X'_i| \geq \sum_{i=1}^n |X_i^{opt}|$.

If $\sum_{i=1}^n |X_i^p| > m$, that is, $(\overline{X}^p, p)$ does not satisfy the limited supply constraint, $(\overline{X}^p, p)$ is then transformed using Lemma 2.12 into a pair envy-free outcome $(\overline{X}^{p'}, p')$ such that $m/2 \leq \sum_{i=1}^n |X_i^{p'}| \leq m$ and $p' \geq p$.

In conclusion, since $(\overline{X}^{p'}, p')$ allocates at least half of the items of $(\overline{X}^{opt}, p^{opt})$ at a price $p' \geq p \geq p^{opt}/(1+\varepsilon/2)$, $r(\overline{X}^{p'}, p') \geq \frac{r(\overline{X}^{opt}, p')}{2} \geq \frac{r(\overline{X}^{opt}, p^{opt})}{2(1+\varepsilon/2)} = \frac{r(\overline{X}^{opt}, p^{opt})}{2+\varepsilon}$. The theorem follows by observing that Algorithm 4 considers price p in at least one iteration, and thus returns an outcome of revenue at least $r(\overline{X}^{p'}, p')$. □

In order to complete the proof, it remains to describe the above claimed optimal dynamic programming algorithm and show its correctness.

To this aim, given a pair envy-free allocation $(\overline{X}, p)$ maximizing the number of items sold for a fixed market μ and price p, let us denote by $B_{\overline{X}}$ the set of all the strictly positive bundle sizes assigned in $\overline{X}$. By the pair envy-freeness constraint and the optimality of $(\overline{X}, p)$, each buyer having a non-negative utility for at least one size in $B_{\overline{X}}$, must get the bundle with size in $B_{\overline{X}}$ that maximizes her utility. Therefore, given $B_{\overline{X}}$, we can reconstruct $\overline{X}$ in polynomial time, and thus it is sufficient that Algorithm 5 determines $B_{\overline{X}}$.

Notice also that, given a fixed price p, each buyer i such that $m_i \leq \frac{v_i(m_i)}{p}$ has non-negative utility for all the bundle sizes in the interval $\left[m_i, \left\lfloor \frac{v_i(m_i)}{p} \right\rfloor\right]$. Since in item-pricing the cost of a bundle increases linearly in its cardinality, while by the single-minded free disposal assumption the valuation remains fixed, in such an interval the utility is strictly decreasing in the bundle size. Therefore, given $B_{\overline{X}}$, the bundle X_i assigned to i has cardinality equal to the minimum of size in $\left[m_i, \left\lfloor \frac{v_i(m_i)}{p} \right\rfloor\right] \cap B_{\overline{X}}$, and 0 if such a set is empty.

We recall that m in single-minded instances can be exponential in the input size. Hence, in order to efficiently compute an optimal set of bundle sizes $B_{\overline{X}}$, we show how to determine a superset B_p of $B_{\overline{X}}$ of polynomially bounded size.

Lemma 2.14. *Given market μ and price p, let $B_p = \left\{ \left\lfloor \frac{v_i(m_i)}{p} \right\rfloor \middle| i \in N \right\}$. Then, $B_{\overline{X}} \subseteq B_p$ for every possible outcome $(\overline{X}, p)$ maximizing the number of items sold in μ at price p.*

Proof. Assume by contradiction that a bundle of size $j \notin B_p$ is allocated in $(\overline{X}, p)$, and let N_j be set of buyers $i \in N$ such that $|X_i| = j$. Consider the outcome $(\overline{X}', p)$ in which we swap all the bundles of size j with ones of cardinality $\min_{i \in N_j} \left\lfloor \frac{v_i(m_i)}{p} \right\rfloor$. Notice that, since $\overline{X}$ is pair envy-free, no bundle size j' with $\min_{i \in N_j} m_i \leq j' < j$ is allocated by $\overline{X}$ and $\overline{X}'$. In fact, if this is not the case, as each buyer $i \in N_j$ prefers the bundle with least cardinality in $[m_i, j]$, N_j would contain an envious buyer. This implies that $(\overline{X}', p)$ is also pair envy-free, while allocating a larger number of items, thus contradicting the optimality of $(\overline{X}, p)$. □

Algorithm 5: A subroutine for the $(2 + \varepsilon)$-approximation algorithm for (SINGLE,ITEM)-pricing with free disposal

---

Input: A market μ together with a price p.
Output: An item-pricing pair envy-free outcome $(\overline{X}^p, p)$ that allocates the maximum number of items (possibly exceeding supply constraints).
Let $B_p = \left\{ \left\lfloor \frac{v_i(m_i)}{p} \right\rfloor \middle| i \in N \right\}$;
for $j = m, \ldots, 1$ *with* $j \in B_p$ **do**
  Compute $x(j)$, the number of items sold by a maximal allocation, given that the minimum size for an allocated bundle is j:
  $x(j) = \max \left\{ \max_{k \in B_p, k > j} \left\{ x(k) + \Delta_k^j \right\}, j \cdot |N_j| \right\}$;
  Keep track of the corresponding allocations;
end
return $(\overline{X}^j, p)$ associated to $\max_j x(j)$;

We are now ready to present Algorithm 5. For every integer $j \geq 1$, let N_j be the set of buyers having non-negative utility for bundles of size j at price p. Let $x(j)$ be the number of items sold by an optimal solution $(\overline{X}^j, p)$ of the subproblem, conditioned to the fact that j is the minimum bundle size in $(\overline{X}^j, p)$. We can have one of the following two cases:

i. $(\overline{X}^j, p)$ allocates only bundles of size j, achieving revenue $r(j) = r(\overline{X}^j, p) = j \cdot |N_j|$;

ii. $(\overline{X}^j, p)$ can be determined by an optimal outcome $(\overline{X}^k, p)$ with minimum bundle size k for a given $k > j$, just assigning bundles of size j to the buyers in N_j, and the remaining ones as in $(\overline{X}^k, p)$.

Notice that in case ***ii.*** the already allocated buyers of N_j in $(\overline{X}^k, p)$, that is, belonging to $N_j \cap N_k$, in $(\overline{X}^j, p)$ will see their bundles of size k swapped with ones of size j. Moreover, all the buyers with a bundle greater than k in $(\overline{X}^k, p)$ have preferred size greater than k, and thus will not envy buyers with bundles of smaller size. Therefore, $(\overline{X}^j, p)$ preserves the pair envy-freeness, as each buyer still receives the minimum size bundle, among the allocated, providing her a non-negative utility. If we set $\Delta_k^j = j \cdot |N_j| - k \cdot |N_j \cap N_k|$, the amount of items sold in this case is $x(j) = x(k) + \Delta_k^j$.

Therefore, for each $j \in B_p$, it is possible to compute $x(j)$ in decreasing order of j using the following explicit formula:

$$x(j) = \max\left\{ \max_{k \in B_p, k > j} \left\{ x(k) + \Delta_k^j \right\}, j \cdot |N_j| \right\} \tag{2.1}$$

Finally, the number of items sold by an optimal solution $(\overline{X}, p)$ of the subproblem can be computed by choosing the maximum among all the $x(j)$ values with $j \in B_p$. The set of bundle $B_{\overline{X}}$ associated to $(\overline{X}, p)$ can be determined by simply keeping track of the partial solutions at each step of the dynamic programming. Clearly, being B_p of polynomially bounded size, Algorithm 5 runs in polynomial time. This concludes the correctness proof of the $(2 + \varepsilon)$-approximation algorithm for (SINGLE,BUNDLE)-pricing with free disposal.

2.4 General Valuations

In this section we consider buyers with general valuations in multi-unit markets.

2.4.1 Item-pricing

Again, we first focus on item-pricing. For such a case, the tractability of the problem in the standard setting, that is, in the case of pair envy-freeness, was an open problem raised in [82].

For this setting we provide a hardness of approximation result that relies on a conjecture on the hardness of deciding whether a random 3-SAT formula is satisfiable [50]. More precisely, given random 3-SAT formula with n variables and $m = \Delta n$ clauses, the conjecture states that for every $\varepsilon > 0$, and for a large constant Δ that does not depend on n, there is no polynomial time algorithm that is able to identify formulae with $(1-\varepsilon)m$ satisfiable clauses. A problem is said to be $R3SAT$-hard if the existence of a polynomial time algorithm solving it would falsify such a conjecture. In our proof we provide a polynomial time reduction from the following $R3SAT$-hard problem.

Definition 2.15. In MES, the problem of finding a maximum expanding sequence [23], we are given a universe set U and an ordered collection of some of its subsets $\mathcal{C} = (S_1, S_2, ..., S_c)$. An expanding sequence $\phi = (\phi(1) < ... < \phi(\ell))$ of length $|\phi| = \ell$ is a selection of sets $S_{\phi(1)}, ..., S_{\phi(\ell)}$, such that for each y, $1 \leq y \leq \ell$, $S_{\phi(y)} \nsubseteq \bigcup_{l=1}^{y-1} S_{\phi(l)}$. The problem requires to find the expanding sequence of maximum length.

Definition 2.16. An instance of MES is said to be κ-separable if the sequence of the subsets $\mathcal{C}$ can be partitioned in the order into κ subsequences or *classes* $\mathcal{C}_1, ..., \mathcal{C}_\kappa$, where each subsequence does not contain intersecting sets. In [23] it has been shown that there exists an $\epsilon > 0$ such that MES is $R3SAT$-hard to approximate within $O\left(f(c)^\epsilon\right)$, when restricted to $f(c)$-separable instances, where $f : \mathbb{N} \to \mathbb{R}^+$ satisfies:

i. f is non-decreasing;

ii. $f(a) \leq a$ and $f(a^b) \leq f(a)^b$ for all $b \geq 1$, $a \in \mathbb{N}$.

For the sake of brevity, from now on, we denote $f(c)$ as κ.

Before providing our reduction, let us informally discuss the underlying ideas.

Proof idea. In our reduction we encode elements of the universe U as bundles of a market, and the ordered collection of sets $\mathcal{C} = (S_1, S_2, ..., S_c)$ as buyers with particular valuation functions.

More precisely, in the reduced instance buyers associated to a set S_i have a constant valuation over a set of bundle sizes which encodes the elements of S_i. Since their valuation is constant over this set, no matter the chosen price, buyers have a higher utility for smaller bundles. This means that a buyer would envy another buyer if the latter manages to get a smaller bundle for which she has a positive valuation.

Another important detail is that we group many identical (with respect to their valuation) buyers together. The bigger the group, the smaller the constant valuation of its buyers, attaining a kind of constant purchasing power among the groups, given by the product of their size and their valuation. Observe that, in order to sell items to a group at its exact valuation, we must avoid allocating a bundle encoding elements of the set S_i associated to the group to another group with a lower valuation.

In other words, since bundles encode elements of U and buyers' envy behaves as just described, in order to sell bundles to groups at their exact valuation, for each group we must find a bundle representing an element that has not been covered by the buyers with higher valuation, to whom we have already allocated other bundles.

The strategy just described is only one of the possible strategies an algorithm can exploit in order to find a stable outcome for our reduced market. However, as we show in the next theorem, it is able to determine a solution which has revenue within a constant factor of the optimal one, and at the same time can be easily reconstructed starting from any optimal outcome.

As you may have noticed, so far we did not specify anything concerning the optimal price for the reduced market. In fact, the market has been suitably constructed in such a way that the revenue of an optimal outcome is almost independent on the chosen price. More precisely, the choice of a price can affect the achieved revenue only by a constant factor with respect to the optimal one, thus shifting the computational complexity of the revenue maximization problem for these particular instances to that of finding a stable allocation of the buyers.

The hardness follows by finally observing that the reduction is approximation preserving, since the revenue of an optimal solution of the reduced instance is within a constant multiplicative factor of the optimal one of MES.

After this brief description of the proof idea we are ready to prove the following hardness theorem.

Theorem 2.17. *Approximating* (GENERAL,ITEM)*-pricing within a factor of* $O(\log^{\epsilon} n)$*, for some* $\epsilon > 0$*, is R3SAT-hard.*

Proof. In order to prove the claim, we provide an approximation preserving polynomial time reduction from MES [23].

Consider the following reduction. Given a κ-separable instance for MES $S_1, ..., S_c \subseteq U$ with corresponding classes $\mathcal{C}_1, ..., \mathcal{C}_\kappa$, we construct an instance μ of (GENERAL,ITEM)-pricing as follows:

- for each $o \in U$ and $1 \leq k \leq \kappa$, let

$$B_o^k = \left\{ 2^{\kappa-h}|U| + o \;\middle|\; h \in \mathbb{N},\ k \leq h \leq \kappa \right\}$$

.

- we associate to each $S_y \in \mathcal{C}_k$ a set I_y of 2^k buyers such that each $i \in I_y$ has valuation function

$$v_i(j) = \begin{cases} 2^{\kappa-k}|U| + |U| & \text{if } j \in \bigcup_{o \in S_y} B_o^k \\ 0 & \text{otherwise} \end{cases}$$

A clarifying picture for the case of a 4-separable instance of MES is given in Figure 2.2.

We prove that the claim holds even in case of unlimited supply, or analogously by setting the total number of items in μ equal to $c \cdot 2^\kappa \cdot 2^\kappa \cdot |U|$.

In the following for the sake of simplicity we will say that a set I_y associated to a given $S_y \in \mathcal{C}_k$ is of class k.

We can immediately observe that each buyer belonging to a given set I_y of class k is only interested in bundles of cardinality in $\bigcup_{o \in S_y} B_o^k$ and, regardless of the price, she always prefers the smallest one. Therefore, in any stable outcome, either all the buyers in I_y do not receive any bundle, or they all receive bundles of the same fixed size j. In this case we say that I_y *supports* size j.

Notice also that for any $q \neq o$ we have $B_q^k \cap B_o^h = \emptyset$. Therefore, there is no bundle size with a strictly positive valuation for both the buyers in two sets I_y and $I_{y'}$ when S_y and $S_{y'}$ are disjoint. In particular, by definition of κ-separability, this implies that I_y and $I_{y'}$ cannot support the same bundle size if S_y and $S_{y'}$ belong to the same class $\mathcal{C}_k$. In other words, every bundle size can be supported by at most one set I_y per class.

In order to prove that the reduction is approximation-preserving, we rely on the following lemmata.

Lemma 2.18. *If the reduced instance μ admits a stable outcome $(\overline{X}, p)$ with revenue r and $p \neq 1$, then it also admits a stable outcome $(\overline{X}', 1)$ with revenue $\frac{r}{4}$.*

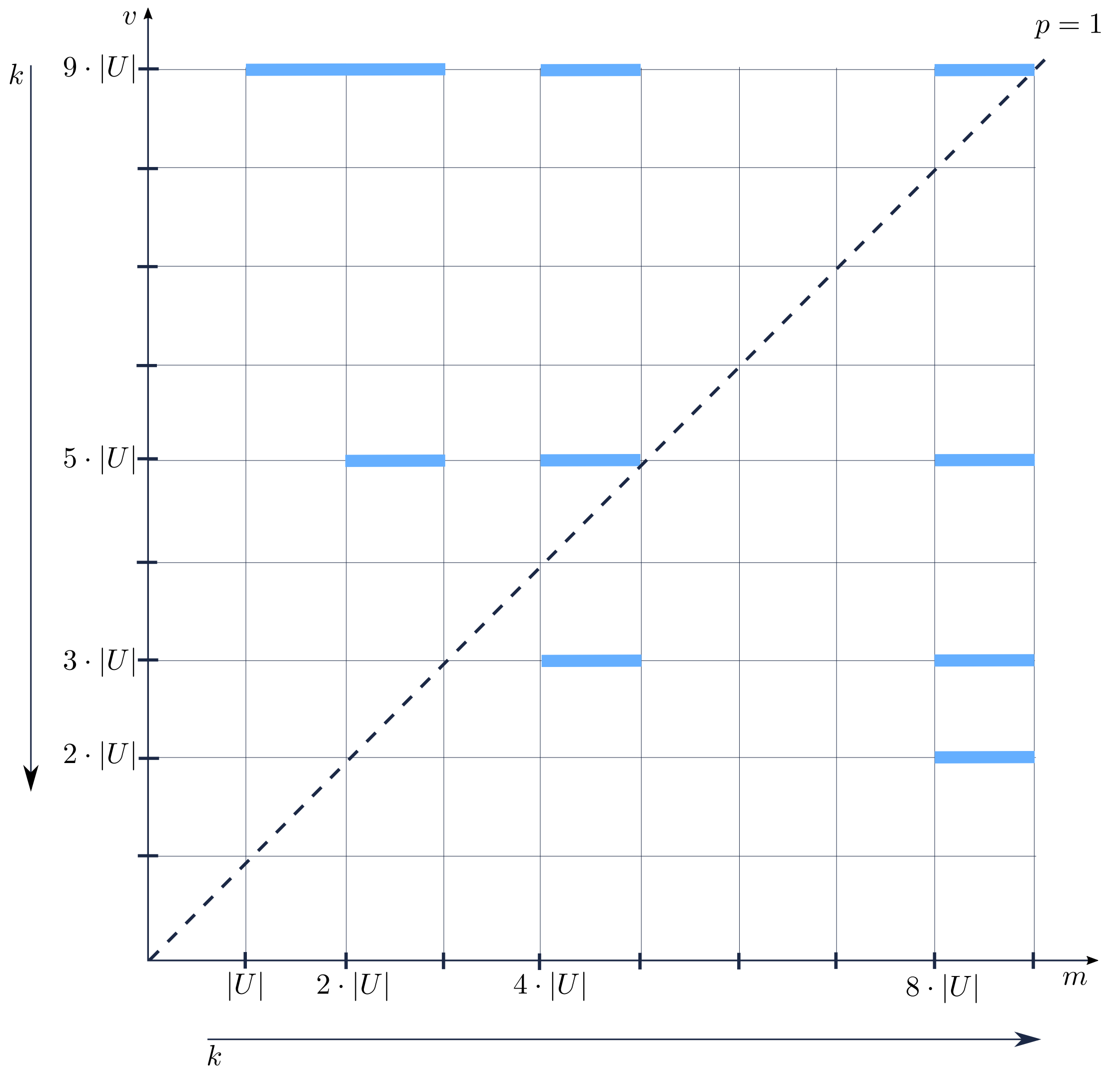

FIGURE 2.2: A market obtained from a 4-separable instance of MES; buyers are interested in the highlighted bundles.

Proof. If $p < 1$, then setting the price equal to 1 does not change the preferred bundle of any buyer, as by construction she has a non-negative utility for a given bundle with price p if and only if the same holds with price 1. Therefore outcome $(\overline{X}, 1)$ is stable and guarantees revenue strictly greater than r.

Assume then $p > 1$. In this case we can express p as $\alpha 2^{\pi}$, with $1 \leq \alpha < 2$ and $\pi \in \mathbb{N}$. Let I_y be a set of buyers of class k supporting a given size $2^{\kappa-h}|U| + o$ with $o \in S_y$ and $k \leq h \leq \kappa$.

We first show that $2^{\kappa-h+\pi}|U| + o$ is a size for which buyers in I_y have a strictly positive valuation, that is, $\kappa - h + \pi \leq \kappa - k$.

Since in $(\overline{X}, p)$ the utilities of all the buyers are non-negative, $\left(2^{\kappa-h}|U| + o\right)\alpha 2^{\pi} \leq 2^{\kappa-k}|U| + |U|$, so that

$$2^{\kappa-h} \leq \frac{2^{\kappa-k-\pi}}{\alpha} + \frac{1}{\alpha 2^{\pi}} - \frac{o}{|U|} \leq 2^{\kappa-k-\pi} + \frac{1}{\alpha 2^{\pi}}. \tag{2.2}$$

Then, if $\pi = 0$, since $k \leq h \leq \kappa$, both $2^{\kappa-k-\pi}$ and $2^{\kappa-h}$ are positive integers, while term $\frac{1}{\alpha 2^{\pi}}$ is strictly smaller than 1. Therefore, inequality (2.2) implies $2^{\kappa-h} \leq 2^{\kappa-k-\pi}$, that is, $\kappa - h + \pi \leq \kappa - k$. If $\pi \geq 1$, since $2^{\kappa-h}$ is a positive integer and $\frac{1}{\alpha 2^{\pi}} < \frac{1}{2}$, if (2.2) holds $2^{\kappa-k-\pi} > \frac{1}{2}$, that is, $\kappa - k - \pi \geq 0$. Therefore, again (2.2) implies $2^{\kappa-h} \leq 2^{\kappa-k-\pi}$ and consequently $\kappa - h + \pi \leq \kappa - k$.

Consider now the outcome $(\overline{X}', 1)$ for the reduced instance constructed as follows:

i. all the buyers in a set I_y of class k supporting size $2^{\kappa-h}|U| + o$ in $(\overline{X}, p)$ receive a bundle of size $2^{\kappa-h+\pi}|U| + o$;

ii. all the remaining buyers either get their preferred bundle among the ones allocated in the previous step, or none.

We now prove that $(\overline{X}', 1)$ is stable. First of all notice that, since $\kappa - h + \pi \leq \kappa - k$, all the buyers have non-negative utility for the bundles they receive. Assume then by contradiction that there exists a set I_y of buyers that are envious in $(\overline{X}', 1)$. By construction, they got their bundles in step ***i.***, and since step ***ii.*** does not produce new bundles sizes, they must envy the buyers in a set of buyers $I_{y'}$ that got their bundles in step ***i.***. Let $2^{\kappa-h}|U| + o$ and $2^{\kappa-h'}|U| + o'$ be the bundle sizes supported in $(\overline{X}', 1)$ by I_y and $I_{y'}$, respectively. Since under any price all buyers prefer the bundle of minimum cardinality among the ones providing them a non-negative utility, recalling that $(\overline{X}, 1)$ is stable,

$$2^{\kappa-h'}|U| + o' \geq 2^{\kappa-h}|U| + o \tag{2.3}$$

On the other hand, since buyers in I_y envy the ones in $I_{y'}$, we have $o' \in I_y$ and

$$2^{\kappa-h'+\pi}|U| + o' < 2^{\kappa-h+\pi}|U| + o \tag{2.4}$$

Combining inequalities (2.4) and (2.3) we obtain $2^{\pi}\left(2^{\kappa-h'} - 2^{\kappa-h}\right) < 2^{\kappa-h'} - 2^{\kappa-h}$, that implies $h < h'$. Since both h and h' are integers, we have $h \leq h' - 1$. Therefore, combining with inequality (2.3), we obtain $2^{\kappa-h'}|U| + o' \geq 2^{\kappa-h}|U| + o \geq 2^{\kappa-h'+1}|U| + o$, i.e., $o' \geq 2^{\kappa-h'}|U| + o$. Then, since $\kappa - h' \geq 0$, it must be $o' > |U| + o'$: a contradiction.

Therefore, $(\overline{X}', 1)$ is stable and it only remains to show that it guarantees a revenue $r(\overline{X}', 1)$ which is at least one fourth the revenue $r(\overline{X}, \alpha 2^{\pi})$. Let $\mathcal{I}$ be the family of the sets of buyers that receive at least one item in $\overline{X}$ and let $2^{\kappa-h_y+\pi}|U|+o$ be the size supported by any $I_y \in \mathcal{I}$. Since by construction all the sets $I_y \in \mathcal{I}$ receive bundles also in $(\overline{X}',1)$,

$$\frac{r(\overline{X}',1)}{r(\overline{X},\alpha 2^{\pi})} \geq \frac{\sum_{I_y\in\mathcal{I}} 2^{\kappa-h_y+\pi}|U|+o}{\sum_{I_y\in\mathcal{I}} \left(2^{\kappa-h_y}|U|+o\right)\alpha 2^{\pi}} \geq$$

$$\geq \frac{1}{\alpha}\cdot\frac{\sum_{I_y\in\mathcal{I}} 2^{\kappa-h_y+\pi}|U|}{\sum_{I_y\in\mathcal{I}} 2^{\kappa-h_y+\pi}|U| + \sum_{I_y\in\mathcal{I}} o2^{\pi}} \geq \frac{1}{2\alpha}\cdot\frac{\sum_{I_y\in\mathcal{I}} 2^{\kappa-h_y+\pi}|U|}{\sum_{I_y\in\mathcal{I}} 2^{\kappa-h_y+\pi}|U|} \geq \frac{1}{4}$$

□

Lemma 2.19. *If μ admits a stable outcome $(\overline{X},1)$ with revenue r, then it also admits a stable outcome $(\overline{X}',1)$ with revenue at least r that satisfies the following property $\mathcal{P}$: "for each bundle of size $2^{\kappa-k}|U|+o$ allocated in $\overline{X}'$, there is a subset I_y of class k that supports $2^{\kappa-k}|U|+o$".*

Proof. If $(\overline{X},1)$ satisfies property $\mathcal{P}$ the claim holds by setting $\overline{X}'=\overline{X}$, otherwise let B be the set of the sizes of the bundles allocated in $(\overline{X},1)$ and let $j=2^{\kappa-k}|U|+o$ be a size in B whose (unique) set of highest class supporting j has class $h<k$. Consider then the outcome $(\overline{X}',1)$ in which each buyer gets her preferred bundle among the ones with sizes in $B' = B \cup \{2^{\kappa-h}|U|+o\} \setminus \{2^{\kappa-k}|U|+o\}$.

By construction $(\overline{X}',1)$ is stable. For what concerns its achieved revenue we observe that, since we just increased one size in B', recalling that in $(\overline{X}',1)$ each buyer prefers the smallest size in B' for which she has a non-negative utility, we have that all the buyers not receiving a bundle of size $2^{\kappa-k}|U|+o$ in $(\overline{X},1)$ still get a bundle of the same size in $(\overline{X}',1)$. All the sets of buyers I_y receiving $2^{\kappa-k}|U|+o$ in $(\overline{X},1)$, being of class at most h, have non-negative utility for size $2^{\kappa-h}|U|+o$. Thus, in $(\overline{X}',1)$ they either support $2^{\kappa-h}|U|+o$ or a different size $2^{\kappa-k'}|U|+o'$ such that $2^{\kappa-k'}|U|+o' < 2^{\kappa-h}|U|+o$ and $2^{\kappa-k'}|U|+o' > 2^{\kappa-k}|U|+o$.

In conclusion, $(\overline{X}',1)$ has at least the same revenue as that of $(\overline{X},1)$. Therefore, since the maximum bundle size and the number of different allocated sizes are limited, by iterating this process a finite number of times we must finally reach the claimed outcome. □

We are now ready to prove our main claim. To this aim, it is sufficient to show that (⇒) if the MES instance admits an expanding sequence of length ℓ, then μ admits a solution with revenue at least $\ell 2^{\kappa}|U|$ and (⇐) if μ admits a solution with revenue $\ell 2^{\kappa}|U|$, then the corresponding MES instance admits an expanding sequence of length $\frac{\ell}{16}$.

(⇒) Suppose that the MES instance admits an expanding sequence $S_{\phi(1)}, ..., S_{\phi(\ell)}$ of length ℓ. Let N_y be the set of the elements newly covered by $S_{\phi(y)}$ in the sequence. Consider the following set of bundle sizes B. For each $S_{\phi(y)} \in \mathcal{C}_k$ in the expanding sequence, put in B integer $2^{\kappa-k}|U| + o$, for an arbitrarily chosen $o \in N_y$.

Let then $(\overline{X}, p)$ be an outcome where $p = 1$ and $\overline{X}$ gives to each buyer her preferred bundle with size in B. Clearly, in such a solution no buyer can be envious, so that $(\overline{X}, p)$ is stable. Then, as $p = 1$, it remains to prove that at least $\ell 2^{\kappa}|U|$ items are sold.

Since $p = 1$, every buyer in $I_{\phi(y)}$ with $S_{\phi(y)} \in \mathcal{C}_k$ has non-negative utility for all (and only) the bundles sizes in $\bigcup_{o \in S_{\phi(y)}} B_o^k$, so that her preferred assigned bundle is the one with least cardinality in $B \cap \bigcup_{o \in S_{\phi(y)}} B_o^k$. We now prove that such a bundle has size at least $2^{\kappa-k}|U|$. By construction, we know that B contains size $2^{\kappa-k}|U| + o$ for some o in $S_{\phi(y)}$. Hence, it is enough to show that B does not contain any other size $2^{\kappa-h}|U| + o'$ with $h > k$ and $o' \in S_{\phi(y)}$. Assume by contradiction that B contains such an integer. Then, since $k < h$, it must be that o' is a newly covered element by a subset $S_{\phi(y')} \in \mathcal{C}_h$ in the expanding sequence with $\phi(y) < \phi(y')$: an absurd, since o' has been previously covered by $S_{\phi(y)}$.

In conclusion, we have that, for any $S_{\phi(y)} \in \mathcal{C}_k$ in the expanding sequence, at least $|I_{\phi(y)}| = 2^k$ buyers receive a bundle of size at least $2^{\kappa-k}|U|$. Therefore, in $(\overline{X}, p)$ globally at least $\ell 2^{\kappa}|U|$ items are sold.

(⇐) Suppose that the reduced instance admits a stable outcome $(\overline{X}, p)$ of revenue $\ell 2^{\kappa}|U|$. Then, by Lemma 2.18 and Lemma 2.19, there exists a stable outcome $(\overline{X}', 1)$ of revenue $r(\overline{X}', 1)$ at least $\frac{\ell}{4} 2^{\kappa}|U|$, where any allocated bundle of size $j = 2^{\kappa-k}|U| + o$ is supported by a subset I_y of class k. By the definition of the valuation functions of the buyers and by the above observations on the reduction, size j can be supported by at most one set $I_{y'}$ for each class h, $1 \leq h \leq k$. Since each such $I_{y'}$ of class h has cardinality 2^h, this implies that at least half of all the players supporting j are contained in I_y, so that at least half of the revenue contributed to $r(\overline{X}', 1)$ by the bundles of size j is due to I_y. Let us call *maximal* any such a subset I_y, that is, supporting size $2^{\kappa-k}|U| + o$ for some $o \in U$, and let $\mathcal{I}^{max}$ be the family of all the maximal sets I_y. We now to prove that $|\mathcal{I}^{max}| \geq \frac{\ell}{16}$.

Denoting by $2^{\kappa-k_y}|U| + o_y$ the size supported by any given $I_y \in \mathcal{I}^{max}$ of class k_y, we then have

$$\frac{\ell}{8} 2^{\kappa}|U| \leq \frac{r(\overline{X}', 1)}{2} \leq \sum_{I_y \in \mathcal{I}^{max}} 2^{k_y} \left(2^{\kappa-k_y}|U| + o_y\right) \leq \sum_{I_y \in \mathcal{I}^{max}} \left(2^{\kappa}|U| + 2^{k_y} o_y\right) \leq$$

$$\leq 2 \cdot 2^{\kappa}|U||\mathcal{I}^{max}|$$

and this implies that $|\mathcal{I}^{max}| \geq \frac{\ell}{16}$.

It remains to prove that the sequence induced by $\mathcal{I}^{max}$ is an expanding sequence. In fact, given any $I_y \in \mathcal{I}^{max}$ supporting size $2^{\kappa-k_y}|U| + o_y$, we have that element o_y is newly covered by S_y in the sequence. If not, it means that o_y belongs also to some $S_{y'} \in \mathcal{C}_{k_{y'}}$ in the sequence, with $k_{y'} < k_y$. However, this is not possible, as otherwise buyers in I'_y would have preferred bundle $2^{\kappa-k_y}|U|+o_y$, thus contradicting the maximality of $I_{y'}$ and consequently the fact that $S_{y'}$ belongs to the sequence.

In order to have a polynomial time reduction we choose $f(c) = \log(c)$, which completes the proof. □

For what concerns the determination of approximated solutions, we notice that the $O(\log n)$-approximation algorithm of [82] for the standard pair envy-freeness problem, given in input any market μ, returns an outcome whose revenue is at least a $\log n$ fraction of the optimal revenue that can be achieved without considering any envy-freeness constraint. Moreover, such an outcome guarantees that no buyer envies any other buyer, and thus is stable with respect to any social graph. Therefore, such an algorithm directly corresponds to a $O(\log n)$-approximation algorithm also for the social (GENERAL,ITEM)-pricing problem.

Let now focus our attention to specific classes of graphs. The following theorem provides a better bound for a restricted class of social graphs.

Theorem 2.20. *The social* (GENERAL,ITEM)*-pricing problem restricted to social graphs of bounded treewidth admits an optimal polynomial time algorithm.*

Proof. Let us first provide a simplified construction for the case in which the social graph G is a tree. Arbitrarily fixing a root r in G, we can exploit the tree structure of the graph to derive a recursive construction of an optimal stable outcome for a given price. More precisely, once we fix a bundle size j with $0 \leq j \leq m$ and a supply bound b with $0 \leq b \leq m$, we can compute the revenue of an optimal *restricted* outcome for a subtree T of G that assigns a bundle of size j to the root i of T and globally at most b items to T. In order to properly define the recursion, we allow the value $j = 0$, that is, we assume that not assigning any item to a buyer corresponds to the assignment of a bundle of size 0.

Let $M_p^i(j,b)$ be the number of items sold in the above optimal outcome, which in turns has revenue $p \cdot M_p^i(j,b)$. We use symbol $\bot$ to denote infeasibility, that is, the fact that under such restrictions a stable outcome for T does not exist. If i is a leaf, we have

$$M_p^i(j,b) = \begin{cases} j & \text{if } v_i(j) - jp \geq 0 \ \ \textit{and} \ \ j \leq b \\ \bot & \text{otherwise} \end{cases} \tag{2.5}$$

If i is not a leaf, $M_p^i(j,b)$ can be recursively constructed by optimally combining the optimal restricted outcomes for its subtrees in T that do not make i and her children envious and globally satisfy the supply constraint. Such a problem can be formulated as an instance of MULTIPLE CHOICE KNAPSACK, according to the original version in which we must pick exactly one object per class. In such an instance, an object $o_p^{i'}(j',b')$ with utility $M_p^{i'}(j',b')$ and weight b' represents an optimal restricted outcome assigning a bundle of size j' to child i' of i and at most b' items at the subtree rooted at i', if such an outcome exists. Then, we can associate to node i' the class $O_p^i(j,b,i')$ of all the objects $o_p^{i'}(j',b')$ not creating envies between i and i' and not exceeding budget b together with the bundle of size j of node i. Namely, $O_p^i(j,b,i')$ contains the objects $o_p^{i'}(j',b')$ such that $v_i(j) - jp \geq v_i(j') - j'p$, $v_{i'}(j') - j'p \geq v_{i'}(j) - jp$, $b' \leq b - j$ and $M_p^{i'}(j',b') \neq \bot$.

The knapsack capacity is set to b. The built instance of MULTIPLE CHOICE KNAPSACK has all values polynomially bounded in the size of the instance and thus admits an exact polynomial-time algorithm (via dynamic programming). If it has a feasible solution, we let $\mathrm{OPT}_p^i(j,b)$ be its value, otherwise we set $\mathrm{OPT}_p^i(j,b) = \bot$.

We can then compute $M_p^i(j,b)$ for an intermediate node i as:

$$M_p^i(j,b) = \begin{cases} \bot & \text{if } v_i(j) \leq jp \text{ or } \mathrm{OPT}_p^i(j,b) = \bot \\ j + \mathrm{OPT}_p^i(j,b) & \text{otherwise} \end{cases} \tag{2.6}$$

Recalling that r is the arbitrarily chosen root, the maximum number of allocated items, given price p, can be then expressed as:

$$M_p^* = \max_{j \leq m} M_p^r(j,m) \tag{2.7}$$

Exploiting (2.5), (2.6) and (2.7), we can compute in polynomial time M_p^*. More precisely, for each node i, bundle j and $b \leq m$, we can compute each $M_p^i(j,b)$, according to a post-order visit on G. This ensures that, when $M_p^i(j,b)$ is computed, all the values $M_p^{i'}(j',b')$ associated to the children of i are already known. Keeping track of the bundles chosen at every step, we can finally return the allocation which realizes M_p^*.

In order to complete the algorithm, we recall that since the number of candidate prices that can lead to an optimal solution is polynomially bounded in the size of the instance (see Lemma 2 in [82]), we can repeat this procedure for all such prices, returning the best computed solution in terms of revenue, thus finally obtaining an exact polynomial-time algorithm for tree social graphs.

The approach mentioned above can be extended to social graphs of bounded treewidth as follows.

Definition 2.21. A tree decomposition of a graph $G = (V, E)$ is a pair $(\mathcal{D}, T)$ where $\mathcal{D} = \{D_1, ..., D_l\}$ is a family of node subsets of G, and T is a set of edges between the components in $\mathcal{D}$ such that:

i. $\bigcup_{D_i \in D} D_i = V$;

ii. for each $\{u, w\} \in E$ there is at least a $D_i \in \mathcal{D}$ such that $u, w \in D_i$;

iii. for each $u \in D_i \cap D_l$, we have $u \in D_h$ for all sets D_h along the unique path from D_i to D_l in T;

iv. the edges in T do not form any cycle.

The treewidth of a graph G is the minimum number k for which there exists a tree decomposition $(\mathcal{D}, T)$ of G with $|D_i| \leq k + 1$ for all $D_i \in \mathcal{D}$. It is well-known that if k is a constant such a decomposition can be found in polynomial time.

Given a price p, we can exploit such a decomposition in order to design a polynomial-time procedure that maximizes the number of items sold or analogously the achievable revenue under price p for social graphs of bounded (constant) treewidth. Extending the algorithm for trees, we can describe an allocation of bundles to buyers of a certain $D_i = \{u_1, \ldots, u_t\}$ with a t-tuple of sizes $\overline{j}$ or simply $j = (j(u_1), \ldots, j(u_t))$. Since each D_i has cardinality $t \leq k + 1$, the number of all possible allocations for each component is $O(m^{k+1})$, that is, polynomially bounded.

Now, let us arbitrarily fix a root r in $(\mathcal{D}, T)$, and define $\mathcal{C}_i$ as the set of child nodes of D_i in T and C_i as $\bigcup_{D_h \in \mathcal{C}_i} D_h$.

Let then $F_p^{D_i}$ be the set of $|D_i|$-tuples j such that:

- for each $u \in D_i$, $v_u(j(u)) - j(u)p \geq 0$;
- for each $u, w \in D_i$ such that $\{u, w\} \in E$, if u and w receive bundles of sizes $j(u)$ and $j(w)$, respectively, neither u or w are envious, i.e., $v_u(j(u)) - j(u)p \geq v_u(j(w)) - j(w)p$ and $v_w(j(w)) - j(w)p \geq v_w(j(u)) - j(u)p$.

Given any $D_h \in \mathcal{C}_i$, let $F_p^{D_i,D_h}(j)$ be set of $|D_h|$-tuples s such that:

- for each $u \in D_h$, $v_u(s(u)) - s(u)p \geq 0$;
- for each $u, w \in D_h$ such that $\{u, w\} \in E$, if u and w receive bundles of sizes $s(u)$ and $s(w)$, respectively, neither u nor w are envious;
- for each $u \in D_i$ and $w \in D_h$ such that $\{u, w\} \in E$, if u and w receive bundles of sizes $j(u)$ and $s(w)$, respectively, neither u or w are envious;
- for each $u \in D_i \cap D_h$, $s(u) = j(u)$.

Therefore, if for any $D_i \in D$ and $D_h \in \mathcal{C}_i$ all buyers $u \in D_i$ receive bundles according to tuple j and all buyers $w \in D_h$ bundles according to a tuple $s \in F_p^{D_i,D_h}(j)$, the corresponding allocation is stable. Note that property ***iii.*** of Definition 2.21 ensures that such allocation is consistent (i.e., exactly one bundle is allocated to each buyer). We can then compute the allocation with the maximum number M_p^* of sold items for a given price p, through the following formula:

- if D_i is a leaf:

$$M_p^{D_i}(j, b) = \begin{cases} \bot & \text{if } j \notin F_p^{D_i} \text{ or } \sum_{u \in D_i} j(u) > b \\ \sum_{u \in D_i} j(u) & \text{otherwise} \end{cases} \tag{2.8}$$

- if D_i is not a leaf, again we rely on a suitable instance of MULTIPLE CHOICE KNAPSACK for combining the optimal restricted outcomes of the subtrees of D_i in T as follows: for each $D_h \in \mathcal{C}_i$, consider the class of objects $O_p^{D_i}(j, b, D_h)$ defined as:

$$\left\{ o_p^{D_h}(s, b') \,\middle|\, s \in F_p^{D_i D_h}(t),\ b' \leq b - \sum_{u \in D_i - D_h} j(u),\ M_p^{D_h}(s, b') \neq \bot \right\}$$

  where each $o_p^{D_h}(s, b')$ has utility $M_p^{D_h}(s, b')$ and weight b', and represents an optimal outcome for the subtree rooted at D_h assigning the tuple of bundles corresponding to s to the nodes in D_h and globally at most b' items to all the nodes in the subtree of D_h. The capacity of the knapsack is set to $b - \sum_{u \in D_i - C_i} j(u)$. Denoting the optimal value of this instance by $\mathrm{OPT}_p^{D_i}(j, b)$ if a feasible solution exists, and setting $\mathrm{OPT}_p^{D_i}(j, b) = \bot$ otherwise, we then have:

$$M_p^{D_i}(j, b) = \begin{cases} \bot \text{ if } j \notin F_p^{D_i} \text{ or } \mathrm{OPT}_p^{D_i}(j, b) = \bot \\ \sum_{u \in D_i - C(i)} j(u) + \mathrm{OPT}_p^{D_i}(j, b) \text{ otherwise} \end{cases}$$

- compute M_p^* using:

$$M_p^* = \max_{j \in F_p^{D_r}} \left\{ M_p^{D_r}(j, m) \right\}$$

Since the number of possible tuples for each given D_i is polynomially bounded, as in the previous case we can find a stable outcome (X, p) that maximizes the number of items sold at price p through a post order visit on (D, T). Recalling that the number of possible optimal prices is polynomially bounded, we can then devise an exact polynomial-time algorithm iterating the described procedure for each such a price and returning the best computed solution in terms of revenue. □

2.4.2 Bundle-Pricing

Let us now consider bundle-pricing. For such a case, an $\Omega(\log^\epsilon n)$ inapproximability result for some $\epsilon > 0$ in the standard setting has been proven in [82], together with a $O(\log n \log m)$-approximation algorithm. However, we are able to prove the following nearly optimal approximation bound, also holding for the more general social instances.

Theorem 2.22. *There exists a $\frac{\log n}{1-e^{-1}}$-approximation algorithm for the social* (GENERAL,BUNDLE)*-pricing problem.*

Algorithm 6: A $\log n$-approximation algorithm for social (GENERAL,BUNDLE)-pricing

---

Input: An instance (μ, G) of the social (GENERAL,BUNDLE)-pricing problem.
Output: A bundle-pricing outcome $(\overline{X}, \overline{p})$.
Let $B = \{b_{lj} : l \in N, j \in M\}$;
Construct the following instance $(K_{N \cup B}, \overline{z}, \overline{w}, k)$ of MAX-KNAPSACK-MATCHING:
  $k = m$;
  $K_{N \cup B}$ = complete bipartite graph on $N \cup B$;
  $z(\{i, b_{lj}\}) = v_i(j)$ for $l \in N$;
  $w(\{i, b_{lj}\}) = j$ for $l \in N$;
Let $(\overline{X}, \overline{p})$ be the stable outcome returned by the $\log n$-approximation algorithm for the UNIT-DEMAND pricing problem on the matching returned by the previous step;
return $(\overline{X}, \overline{p})$;

Proof. Given an instance (μ, G) with $\mu = (N, M, (v_i)_{i \in N})$ of the social (GENERAL,BUNDLE)-pricing problem, consider the instance η of a market with unit-demand buyers constructed as follows. The set of buyers in η is N, i.e. the same of μ, and we include, for each bundle size j in μ, a set of n items $B_j = \{b_{j,1} \ldots, b_{j,n}\}$, under the understanding that assigning item $b_{j,l}$ for some $l \in \{1, \ldots, n\}$, to buyer i in η corresponds to allocating a bundle of size j to i in μ. The n copies in each B_j guarantee that the possibility of assigning a bundle of size j to every buyer is taken into account in η. Every buyer i has valuation $v_i(j)$ for each item $b_{j,l} \in B_j$.

We can represent η by means of a complete bipartite graph $K_{N,B}$ with node set $N \cup B$, $B = B_1 \cup \ldots \cup B_m$, and all possible edges between N and B, where each edge $\{i, b_{j,l}\}$ has weight $w(\{i, b_{j,l}\}) = j$ and value $z(\{i, b_{j,l}\}) = v_i(j)$. Any matching in $K_{N,B}$ then corresponds to an allocation of items in η and of bundles in μ to the buyers in N.

By a little abuse of notation, given two subsets $N' \subseteq N$ and $B' \subseteq B$, let us denote by $z(N', B')$ the maximum value of a matching between N' and B'. A stable outcome for η can be obtained by exploiting an algorithm presented in [68] (see Theorem 3.5), that guarantees a revenue which is at least $z(N, B)/(2 \cdot \log n)$, that is at least equal to the value of any possible allocation of buyers to item in η that can be obtained even without considering envy-freeness constraints. By construction, such an outcome corresponds to an outcome for μ which is stable under G, as it prevents envies between every possible pair of buyers. However, unfortunately such an outcome might not be feasible, because it completely ignores the supply constraints, that is, it might assign more than m items.

In order to solve this problem, we now show how to suitably preselect a subset $B' \subseteq B$ of items having the property that the overall weight $\sum_{b_{j,l} \in B'} w(\{i, b_{j,l}\}) \leq m$ and the value of the maximum matching between N and B' is close to optimality. Namely, if $B^* \subseteq B$ is an optimal selection of items, that is such that $\sum_{b_{j,l} \in B*} w(\{i, b_{j,l}\}) \leq m$ and $z(N, B^*)$ is maximized, B' has the property that $z(N, B') \geq (1 - e^{-1}) z(N, B^*)$. Then, by applying the algorithm of [68] to the submarket of η containing only the items in B', we get a stable outcome for η that this time corresponds to a feasible stable outcome for μ under G and has revenue at least $z(N, B')/\log n \geq (1 - e^{-1}) z(N, B^*)/\log n \geq (1 - e^{-1}) opt(\mu, G)/\log n$, thus proving the claim.

To this aim, it is possible to show that $z(N, B')$ (with fixed argument N) is a non-decreasing submodular set function with respect to B', that is, it satisfies the properties that $z(N, B') \leq z(N, B'')$ if $B' \subseteq B''$ and $Z(N, B') + Z(N, B'') \geq Z(N, B' \cup B'') + Z(N, B' \cap B'')$ for every $B', B'' \subseteq B$. We can then see the problem of determining a subset B' yielding a matching of maximum value while not exceeding overall weight m as an instance of the problem of maximizing a non-decreasing submodular set function subject to a knapsack constraint, for which a $(1 - e^{-1})$-approximation algorithm was provided in [89]. □

Example: We give a simple example of MAX-KNAPSACK-MATCHING construction determined in the previous algorithm, in order to clarify the various involved details.

Consider a market μ with $|N| = 2$ buyers and $m = 3$ items, and assume the two buyers to be connected by an edge in the social graph. The valuations are:

$$v_1(j) = \begin{cases} 1 & j = 1 \\ 2 & j = 2 \\ 3 & j = 3 \end{cases} \qquad v_2(j) = \begin{cases} 4 & j = 1 \\ 4 & j = 2 \\ 4 & j = 3 \end{cases}$$

The instance $(K_{N \cup B}, \overline{z}, \overline{w}, k)$ of MAX-KNAPSACK-MATCHING constructed in the algorithm is such that the complete bipartite graph on $N \cup B$ has 6 edges with the following weights and values:

$$w(\{i, b_{l,j}\}) = \begin{cases} 1 & j = 1 \\ 2 & j = 2 \\ 3 & j = 3 \end{cases}$$

for $i = 1, 2$ and $l = 1, 2$;

$$z(\{1, b_{l,j}\}) = \begin{cases} 1 & j = 1 \\ 2 & j = 2 \\ 3 & j = 3 \end{cases} \qquad z(\{2, b_{l,j}\}) = \begin{cases} 4 & j = 1 \\ 4 & j = 2 \\ 4 & j = 3 \end{cases}$$

for $l = 1, 2$.

The maximum knapsack matchings for this example are given by the edges $\{1, b_{1,2}\}$, $\{2, b_{2,1}\}$ or $\{2, b_{1,2}\}$, $\{1, b_{2,1}\}$, and both of them have a total value of 6. However, due to its approximation, the algorithm during the execution might consider a different one. □

Before concluding the section, let us turn our attention to specific social graph topologies. Differently from the item-pricing case, for which we could provide an optimal polynomial time algorithm for bounded treewidth graphs, determining revenue-maximizing outcomes in this context appears to be a rather more difficult task. In fact, as we are going to show below, even if we ignore both the pair envy-freeness (i.e. we restrict to empty social graphs) and the limited supply constraints, the problem cannot be approximated below a given constant, as it is equivalent to the well studied Max-Buying problem. Such a problem has been proved to be APX-hard in [2] and approximable within a ratio of 1.59. Thus the following result holds.

Corollary 2.23. *The social* (GENERAL,BUNDLE)*-pricing restricted on instances having empty social graphs and unlimited supply of items is APX-hard and approximable within a ratio of* 1.59*.*

Proof. In the Max-Buying problem we are given a set $A = \{1, \ldots, m\}$ of products and n data samples $(R_1, B_1), \ldots, (R_n, B_n)$, each associated to one of n different consumers. R_i for $i = 1, \ldots, n$ is a *rank* or ordering of the products according to consumer i's preference. Let $b_{i,j}$ be the budget that i has for buying product j, and let $B_i = \{b_{i,j}, | 1 \leq j \leq m\}$ be her sequence of budgets. Consumers are assumed to be consistent, i.e., the order of their budgets for the various products obeys their preference ordering. Furthermore, given a pricing $\overline{p} = \{p_1, \ldots, p_m\}$ of the products, they will buy the most expensive one they can afford. Let p^i the price of such a product for consumer i, and $p^i = 0$ if i cannot afford product. The goal is to find a price vector $\overline{p}$, that maximizes the revenue defined as $\sum_{i=1}^{n} p^i$.

The equivalence of Max-Buying and the social (GENERAL,BUNDLE)-pricing problem on empty graphs and unlimited supply is pretty direct. In fact, products correspond to bundle sizes, consumers to buyers, and budgets to valuations. Moreover, without pair envy-freeness and supply constraints, the revenue maximization objective is equivalent to that of assigning to each consumer the most expensive product she can afford, exactly as in Max-Buying.

The claim thus follows by the results in [2]. □

Example: Consider an instance of max buying with two products $\mathbf{a}, \mathbf{b}$ and three customers $\mathbf{1}, \mathbf{2}, \mathbf{3}$. The customers' budgets for the products are given in the following table:

	1	**2**	**3**
a	7	4	8
b	10	12	6

In the corresponding instance of (GENERAL,BUNDLE)-pricing with unlimited supply and empty social graph, we have three buyers with the following valuation functions:

$$v_1(1) = 7 \quad v_2(1) = 4 \quad v_3(1) = 8$$
$$v_1(2) = 10 \quad v_2(2) = 12 \quad v_3(2) = 6$$

where products $\mathbf{a}$ and $\mathbf{b}$ are respectively represented by the bundle sizes 1 and 2. Setting $p_{\mathbf{a}} = 7$ and $p_{\mathbf{b}} = 12$, we have that costumers 1 and 3 purchase product $\mathbf{a}$, while customer 2 purchases product $\mathbf{b}$. The corresponding price vector for (GENERAL,BUNDLE)-pricing is $\overline{p} = (7, 12)$, inducing an optimal allocation $\overline{X}$ with $|X_1| = |X_3| = 1$ and $|X_2| = 2$. □

2.5 Price of Envy-Freeness

In this section we focus on the price of envy-freeness, that is, on the worst-case ratio between the maximum revenue that can be achieved without pair envy-freeness constraints, and the one obtainable in case of social relationships. In particular, we provide tight and asymptotically tight bounds for all the considered cases.

For single-minded buyers the following two theorems hold.

Theorem 2.24. *The price of envy-freeness for* (SINGLE,ITEM)*-pricing is* 2*.*

Proof. The following counterexample shows a lower bound of 2 to the price of envy-freeness for this case: we have 3 buyers and an odd number $m = 2k - 1$ of items. The buyers valuations for their preferred bundles are $v_1(k) = k(1 + \varepsilon)$, $v_2(k) = k(1 + \varepsilon)$ and $v_3(k-1) = k-1$, with ε suitably small, and the social graph G has a single edge between 1 and 2. The outcome setting price 1 per item and assigning a bundle of size k to buyer 1 and of size $k - 1$ to buyer 3 achieves revenue m, but it is not stable, as buyer 2 envies 1. According to Lemma 2.3, the best stable outcome can be obtained by setting price 1 or $1 + \varepsilon$ per item. In the first case the only possibility is that of allocating a single bundle of size $k - 1$ to buyer 3, while in the second case a single bundle of size k either to 1 or 2, that both have utility 0 for such a bundle and thus cannot envy each other. In both cases the revenue is at most $k(1+\varepsilon)$ and the claim follows for k growing to infinity and ε tending to 0.

For what concerns the upper bound, consider a feasible outcome $(\overline{X}, p)$ maximizing the revenue without considering the envy-freeness constraints. If $(\overline{X}, p)$ is stable the upper bound trivially holds. Let then $(\overline{X}', p)$ be the outcome which extends $(\overline{X}, p)$ by allocating to each envying buyer her preferred bundle. Clearly by the optimality of $(\overline{X}, p)$, $(\overline{X}', p)$ cannot satisfy the supply constraint, that is, it allocates more than m items. Let $p' \geq p$ be the highest possible price such that the outcome $(\overline{X}'', p')$ assigning to each buyer with non-negative utility her preferred bundle does not satisfy the supply constraint. Consider all the allocated buyers in $(\overline{X}'', p')$ having utility 0 for the bundles they receive. By the assumption on p', if all such buyers have bundles of size less than $m/2$, we can remove a proper subset of them so as to get an outcome selling at least $m/2$ items at price p'. On the contrary, if there exists one such buyer i with a bundle of size $j \geq m/2$, then we can pick the outcome which assigns a single bundle of size j to the buyer i' that evaluates it most, at a price $p'' = v_{i'}(j)/j \geq v_i(j)/j = p' \geq p$. In both cases, the arising outcome sells at least $m/2$ items at price at least p per item, thus achieving at least half of the revenue of $(\overline{X}, p)$. For what concerns the envy-freeness constraints, the former outcome starts from an unfeasible envy-free outcome and deallocates only buyers with null utility, thus preserving envy-freeness. The latter allocates only one bundle of a given size to the buyer who values it the most, at an overall price equal to her exact valuation, thus no other buyer can be envious. □

Notice that the above lower bound of 2 on the price of envy-freeness for item-pricing holds even for one of the simplest topologies, that is, for the family $\mathcal{P}$ of the path networks, i.e., $c(\mathcal{M}, \mathcal{P}) \geq 2$. On the other hand, the upper bound $c(\mathcal{M}, \mathcal{F}) \leq 2$ holds with respect to any possible graph family $\mathcal{F}$.

In the case of bundle-pricing instead we have the following result.

Theorem 2.25. *The price of envy-freeness for* (SINGLE,BUNDLE)*-pricing is* 1*.*

Proof. Let us now consider (SINGLE,BUNDLE)-pricing and again let $(\overline{X}, \overline{p})$ be a feasible outcome maximizing the revenue without considering the envy-freeness constraints. Let n_j be the number of buyers to whom a bundle of size j is allocated in $(\overline{X}, \overline{p})$, and let $i_1, \ldots, i_{n_j}$ be such buyers, listed in non-increasing order of valuation for bundle j, that is, in such a way that $v_{i_1}(j) \geq v_{i_2}(j) \geq \ldots \geq v_{i_{n_j}}(j)$. Clearly, the price at which a bundle of size j is sold must be equal to $v_{i_{n_j}}(j)$, as a lower price might be increased to $v_{i_{n_j}}(j)$ yielding a feasible outcome of higher revenue, thus contradicting the optimality of $(\overline{X}, \overline{p})$, while a higher price would cause buyer i_{n_j} to have a negative utility, thus violating the individual rationality constraint. Without loss of generality we can assume that all the other buyers have valuations at most equal to $v_{i_{n_j}}(j)$ for a bundle of size j, otherwise we can insert them one by one in the list of allocated buyers for bundle j and take out the one with lowest valuation without decreasing the revenue. Therefore, all the not allocated buyers with preferred bundle of size j cannot have a strictly positive utility for such a bundle. Thus, by the single-minded assumption, $(\overline{X}, \overline{p})$ is also stable, hence the claim. □

Let us turn our attention to general valuations. The following asymptotically tight bounds on the price of envy-freeness hold.

Theorem 2.26. *The price of envy-freeness for* (GENERAL,ITEM)*-pricing is* $\Theta(\log n)$*.*

Proof. The $O(\log n)$ bound can be determined constructively by observing that the $O(\log n)$-approximating algorithm from [82] for item-pricing returns outcomes stable under any G whose revenue is at least a logarithmic fraction of the maximum revenue that can be achieved by ignoring envy-freeness constraints.

For what concerns the corresponding lower bounds, consider first the following counterexample for item-pricing: a market μ with n buyers, n^2 items and valuations

$$v_i(j) = \begin{cases} \frac{n}{i} & \text{for } j = 1, \ldots, n \\ 0 & \text{otherwise} \end{cases}$$

Observe first that the outcome $(\overline{X}, p)$ for the empty social graph with $|X_i| = \lfloor \frac{n}{i} \rfloor$ and $p = 1$ gives revenue $\Theta(n \log n)$.

Consider the social network G consisting of a simple path over all the buyers listed in the order, that is, with edge set $\{\{i, i+1\} | 1 \leq i < n\}$. Since buyers have higher utility for smaller bundles (regardless of the price), the allocation just described is not stable because each buyer $i < n$ envies buyer $i + 1$. In order to characterize the optimal outcomes stable under G, observe first that any allocation in which a buyer $i < n$ does not receive any bundle and buyer $i + 1$ receives one cannot be stable, as i would envy $i + 1$. Thus, the only stable outcomes are the ones assigning only bundles to a prefix of buyers. Notice also that in such a prefix all the assigned bundles must have identical size, otherwise at least one buyer would envy a neighbor with a smaller bundle. Summarizing, an optimal outcome $(\overline{X}, p)$ assigns bundles of a given size to the prefix of all the buyers having non-negative utilities for the fixed price p. More precisely, it sells a bundle of a given size j to each buyer $i \in N$ for which $\frac{n}{i} \geq j \cdot p$, that is, to $O(\frac{n}{j \cdot p})$ buyers, and no bundle of a different size. This can provide revenue at most $O(\frac{n}{j \cdot p} \cdot j \cdot p) = O(n)$, hence the claim. □

Theorem 2.27. *The price of envy-freeness* (GENERAL,BUNDLE)*-pricing is* $\Theta(\log n)$.

Proof. Similarly to the previous theorem, the $O(\log n)$ bound can be determined constructively by observing that Algorithm 6 for bundle-pricing returns outcomes stable under any G whose revenue is at least a logarithmic fraction of the maximum revenue that can be achieved by ignoring envy-freeness constraints.

As far as the corresponding lower bound is concerned, consider the following counterexample (see also Figure 2.3): a market μ with n buyers, n^2 items and valuations

$$v_i(j) = \begin{cases} \frac{1}{i} & \text{for } j = 1, \ldots, n \\ 0 & \text{otherwise.} \end{cases}$$

In the case of an empty social graph, the outcome $(\overline{X}, \overline{p})$ with $|X_i| = i$ and

$$\overline{p}(j) = \begin{cases} \frac{1}{j} & \text{for } j = 1, \ldots, n \\ \infty & \text{otherwise.} \end{cases}$$

yields revenue $\Theta(\log n)$.

Consider again the social network G consisting of a simple path with edge set $\{\{i, i+1\} | 1 \leq i < n\}$. The outcome just described is not stable, as similarly to item-pricing each buyer $i < n$ envies $i+1$.

Since buyers have the same valuation over all the bundles, only pricing schemes which do not discriminate on the bundle size are stable for μ. Thus the set of optimal prices reduces to those of the form $\frac{1}{i}$ and a corresponding optimal stable outcome can sell to at most i buyers, providing revenue $O(1)$ to the seller. □

We remark that, as for single-minded instances, the lower bounds on the price of envy-freeness hold even for the family $\mathcal{P}$ of the path networks, that is, $C(\mathcal{M}, \mathcal{P}) = \Omega(\log n)$, while the upper bound $C(\mathcal{M}, \mathcal{F}) = O(\log n)$ with respect to any possible graph family $\mathcal{F}$.

2.6 Conclusions and Future Work

We considered a framework that explicitly takes into account social relationships in the notion of envy-freeness. We focused on multi-unit markets, and besides closing open questions or improving known approximation results in the standard case, we gave positive and negative results concerning the determination of revenue maximizing socially stable solutions and proper bounds on the price of envy-freeness, that is, on the revenue loss due to enforcing stability according to social relationships.

A major open question concerns general valuations. In fact, similarly to the item-pricing case, for bundle-pricing it would be interesting to characterize the approximability of the problem for specific social topologies, like bounded treewidth graphs. As already remarked, the problem here seems considerably more difficult with respect to the item-pricing case, as it is APX-hard even in the case of empty social graphs, i.e., with all nodes isolated, and ignoring the limited supply.

Moreover, both for item- and bundle-pricing, it would be worthwhile investigating other relevant restricted families of social graphs, such as bounded-degree ones or characteristic topologies in this setting. It would also be worth considering some further restricted pricing policies, like bundle prices inducing a price per item which is decreasing in the bundle size.

Another nice research direction is that of considering other market scenarios, like unit-demand markets, or providing refined results for other suitable restricted families of

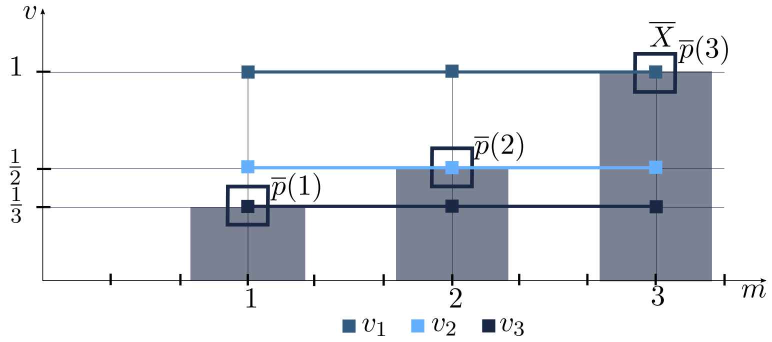

(a)

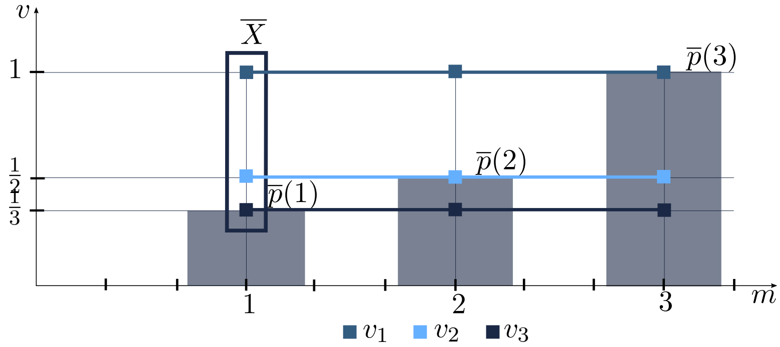

(b)

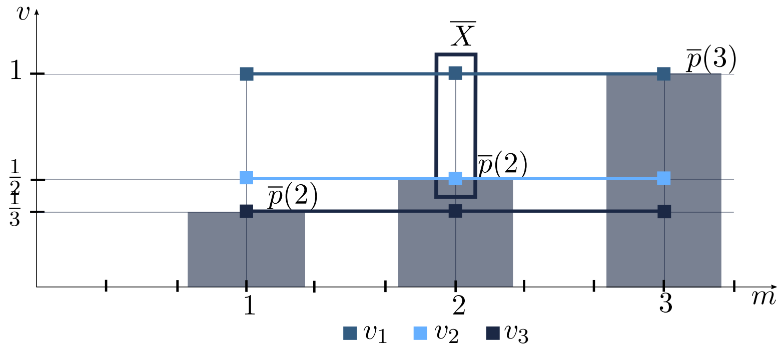

(c)

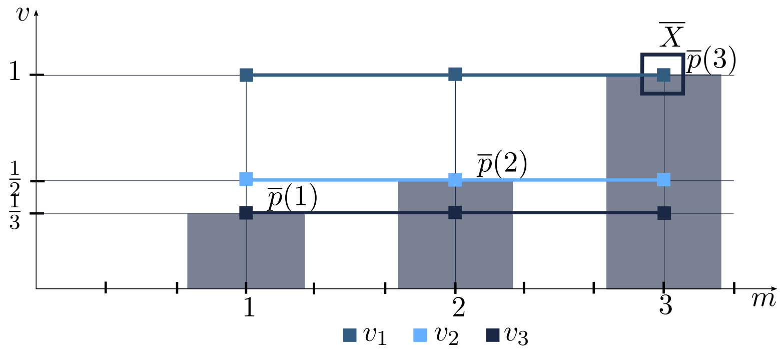

(d)

FIGURE 2.3: An instance of the lower-bound for (GENERAL,BUNDLE)-pricing in the case of a market with three buyers. The outcome in 2.3(a) gives revenue $1 + \frac{1}{2} + \frac{1}{3}$ but it is not pair envy-free, since buyer 2 envies buyer 3 and buyer 1 envies both buyer 2 and 3. In 2.3(b) three buyers receive bundles of size 1 at price $\frac{1}{3}$, in 2.3(c) two buyers receive bundles of size 2 at price $\frac{1}{2}$, and finally in 2.3(d) buyer 1 receives a bundle of size 3 at price 1. Each of these outcomes are pair envy-free but give only revenue 1.

valuation functions to which our results seem not to extend directly, such as subadditive and submodular ones, additive valuations with budgets, additive valuations with cut-off thresholds, that is, linear up to a certain size and then constant, and so forth.

It would be nice also to investigate other forms of social influence, or to determine how modifications of the social graph, for instance adding or removing edges, can affect the achievable revenue.

Finally, even if in the envy-freeness setting the market clearing assumption has been traditionally avoided mainly for the purpose of maximizing the seller's revenue, it would be worthwhile providing similar results for other global objectives, such as the social welfare.

Chapter 3

Fair Price Discrimination

Discriminatory pricing policies, even if often 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 "fair" price discrimination policies that can be tolerated by customers, 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 polynomial time algorithm, hardness, and approximation results under various assumptions on the buyers' valuations and on the social graph topology.

3.1 Introduction

Charging customers different prices for the same item can be perceived as unfair at first, but it is actually widespread. For instance, when buying flight tickets, customers who arrive earlier usually get a better price. In the opposite way, the price of electronic devices like computers and cameras decreases over time. Besides this form of *temporal* discrimination, *spatial* discrimination is also common 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 increase its revenue by discriminating against subsets of buyers on the basis of their mutual distances. A further type of discrimination occurs in user profiling, when premium customers such as frequent-fliers are charged lower prices with respect to the general public, or when some classes of buyers are recognized to be eligible for better treatment, such as age discounts or financial aids.

Introductions on price discrimination together with the main arising scenarios and the related motivation can be found in [26, 79, 86, 88, 91]. An interesting classification is due to A. Pigou in [87], who detected the following three classes:

i. *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.

ii. *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 larger amount.

iii. *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.

In this chapter, we focus on price discrimination policies applied to homogeneous items, or units/copies of the same good. This setting is the basic scenario among combinatorial auctions to which price discrimination can be applied and is able to model several realistic situations. In fact, items could be commodities in markets, goods like electronic devices, power supply in manufacture systems, slots for ads in web search pages, bandwidths in the radio spectrum, and so forth. As a consequence, not surprisingly the multi-unit model is central in combinatorial auctions and has been studied in a large body of literature [14, 19, 45, 47, 48, 52, 63], with a strong focus on the determination of incentive-compatible pricing-based mechanisms. Unfortunately, many of the related auctions, e.g., the VCG mechanism [40, 67, 92] (for bidders with unlimited budgets) or the ascending auction [45, 84] (for bidders with budgets), sell identical products at different prices.

However, as pointed out by the economic literature [8, 12], this is problematic and in some cases even forbidden by the international commerce law [62]. Therefore, Feldman et al. [52] further emphasized the importance of the approach in the scientific community of turning the attention to revenue maximizing auctions that are "envy-free", rather than incentive-compatible. In particular, in one of the most widely used settings, an *envy-free* allocation is an assignment in which each buyer receives a bundle of goods among the ones that maximize her utility. In this chapter we will not consider pair-envy freeness.

In this chapter, 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 modeling price discrimination policies that can be perceived as "fair", 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 it 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 (SCC) the hierarchical profiling of classes of users, where different premium classes may be offered a different treatment or different discounts.

Forms of price discrimination induced by social relationships have been considered in [6], 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 have been provided under different assumptions. Furthermore, [7] 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. In particular, the overlap holds only when $\alpha = 0$, the graph is undirected, and there is an unlimited supply of items. In all the other cases ($\alpha > 0$ or directed graphs or limited supply) the models are incomparable. Moreover, we also investigate the case of general valuations.

3.1.1 Summary

In this chapter, we study the problem of maximizing the social welfare and the seller's revenue in multi-unit markets by means of fair price discrimination policies according to the above framework. Namely, the price per item of each buyer must be at most the lowest one proposed to the buyers she knows directly.

We investigate four different cases arising by considering:

i. social welfare and revenue maximization;

ii. single-minded and general valuations. Respectively when buyers are interested only in a specific amount of items, or in any number of them but with a possibly different valuation.

Moreover, we consider different relevant social graph topologies (see also Figure 3.1).

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 items. We then provide polynomial time algorithms or hardness and approximation results in all the cases mentioned above.

	Single-minded	
	General Graphs	Special Graphs
Social Welfare	**NP-hard (strong)**	**FPTAS - constant number of SCC**
	PTAS	
Revenue	**NP-hard (strong)**	$(2+\varepsilon)$ **approximation - undirected graphs, FPTAS - constant number of SCC**
	log n	

TABLE 3.1: Hardness and approximation results for single-minded valuations pricing problems.

	General Valuations	
	General Graphs	Special Graphs
Social Welfare	**NP-hard (strong)**	**polytime - undirected graphs, constant number of SCC, arborescences of SCC**
	PTAS	
Revenue	**NP-hard (strong)**	
	log n + log m	

TABLE 3.2: Polynomial time, hardness, and approximation results for general valuations pricing problems. Notice that, because of different input sizes, there is no direct way to translate results between single-minded and general valuations.

In more detail, for all the four different cases that arise, and for the specific relevant classes of graphs, we show the following results (see also Tables 3.1, 3.2):

i. *Single-minded, social welfare*: We show that the maximization problem is strongly NP-hard and provide a corresponding PTAS, that is, an algorithm able to return

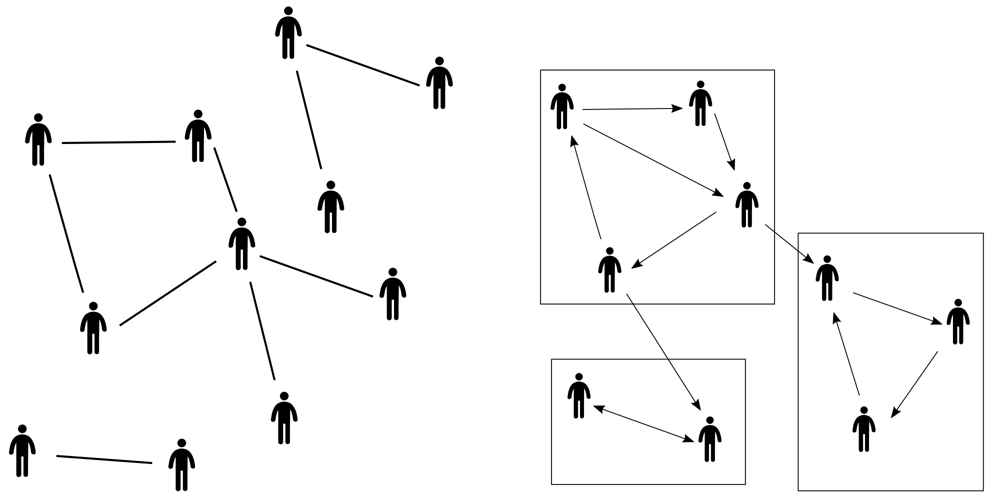

(a) (b)

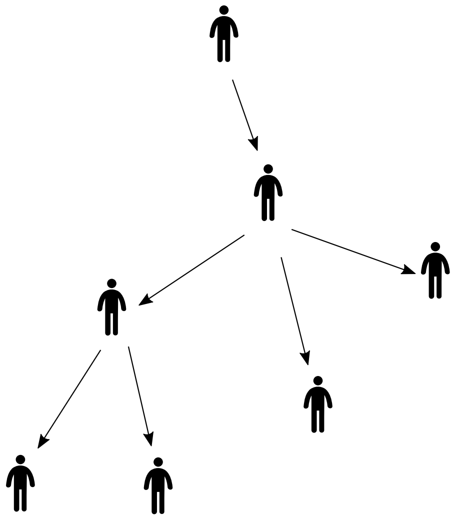

(c)

FIGURE 3.1: In Figure 3.1(a) we depicted an undirected social graph with three connected components. We will see later that according to our definition of fair discriminatory pricing buyers belonging to the same connected component demand the same price from the seller. In Figure 3.1(b) a directed graph of which we highlighted the strongly-connected components (SCC). We will give some improved approximation results when the number of SCCs is bounded by a constant. Finally,Figure 3.1(c) depicts an arborescence, for which we will give a formal definition later. We give improved approximation results also for this class of directed graphs.

an outcome with social welfare at least equal to a $(1+\varepsilon)$ fraction of the maximum achievable one, for any fixed $\varepsilon > 0$, in time polynomial only in the size of the input instance (but not in $1/\varepsilon$); clearly, the strong NP-hardness implies that such an approximation bound is optimal.

ii. *Single-minded, revenue*: We prove that the strong NP-hardness also holds in this case, and give a $O(\log n)$-approximation algorithm.

iii. *General valuations, social welfare*: We first provide a general reduction able to transform instances with general valuations 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. As a consequence, we give a PTAS also for social welfare maximization over buyers with general valuations, and show this approximation bound is tight, as the strong NP-hardness for single-minded instances also extends to this case.

iv. *General valuations, revenue*: By means of the same reduction introduced in ***iii.***, we provide a $O(\log n + \log m)$-approximation algorithm, where the bound comes from the logarithmic approximation for the single-minded case, by observing that $n \cdot m$ is the number of buyers in the reduced single-minded instances; again, we show that the strong NP-hardness for single-minded instances also extends to this case.

v. *Specific topologies*: For single-minded instances, while for an arbitrary graph we developed a logarithmic approximation for revenue maximization, we are able to provide an improved $(2+\varepsilon)$ approximation for the case of undirected social graphs, plus an FPTAS for social graphs having a constant number of strongly connected components. For general valuations, we provide polynomial time algorithms both for social welfare and revenue maximization for the following social graph topologies: undirected graphs, graphs with a constant number of strongly-connected components, and arborescences of strongly connected components. The above topologies are quite relevant, as they allow to model different types of price discrimination that are typical in this setting.

As a final observation, we remark that our reduction from general to single-minded markets is of independent interest, as it might be exploited in other settings in order to extend results from single-minded to general multi-unit instances.

3.2 Preliminaries

Here we will introduce some specific definition for the current setting while bearing in mind the more general definitions from Chapter 2. A multi-unit market μ 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 $i \in N$.

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 $\overline{p} = (p_1, \dots, p_n)$ 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 $N(i) = \{k \in N | (i, k) \in E\}$ of her neighbors in G. Under such an assumption, it is possible to define the following notion of fair price discrimination.

Definition 3.1. A price vector $\overline{p}$ is *fair* with respect to the social graph $G = (N, E)$ if $p_i \leq p_k$ for every $(i, k) \in E$.

We remark that in the framework just introduced, a price must be proposed to each buyer and, since for a fair price vector $\overline{p}$ if both (i, j) and (j, k) are in E then also $p_i \leq p_k$ must hold, it becomes indifferent to consider either G or its transitive closure. Previous works in the literature also addressed the case in which such a transitive property can be arbitrarily broken by allowing the seller to propose a special price $\perp$ to buyers in such a way that they are excluded from the market, that is, they do not mind the prices proposed to their neighbors.

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

A *feasible* outcome for market μ is a pair $(\overline{X}, \overline{p})$ satisfying the following supply constraint:

$$\sum_{i=1}^{n} |X_i| \leq m$$

We say that a feasible outcome $(\overline{X}, \overline{p})$ for market μ is *envy-free* if for every buyer $i \in N$:

$$X_i \in argmax_{X \subseteq M} u_i(X, p_i)$$

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 3.2. A feasible outcome $(\overline{X}, \overline{p})$ is *fair* for μ under the social graph G, or simply fair, if it is envy-free and its price vector is fair with respect to G.

Given an outcome $(\overline{X}, \overline{p})$, the achieved social welfare is $sw(\overline{X}, \overline{p}) = \sum_{i=1}^{n} v_i(|X_i|)$ and the revenue raised by the seller is $r(\overline{X}, \overline{p}) = \sum_{i=1}^{n} p_i \cdot |X_i|$. The (fair) *pricing problem* consists in determining a feasible outcome $(\overline{X}, \overline{p})$ for market μ, which is fair under G and maximizes the social welfare or the revenue.

In the following we will denote by $opt_{sw}(\mu, G)$ (resp. $opt_r(\mu, G)$) the maximum possible social welfare (resp. revenue) achievable by an outcome for μ fair under G, and by $opt_{sw}(\mu)$ (resp. $opt_r(\mu)$) 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 fair outcome $(\overline{X}, \overline{p})$, $sw(\overline{X}, \overline{p}) \geq r(\overline{X}, \overline{p})$, so that also $opt_{sw}(\mu) \geq opt_r(\mu)$ and $opt_{sw}(\mu, G) \geq opt_r(\mu, G)$.

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).

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 a 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 proposing her a price of at least $1 + \varepsilon$ (if 1 doesn't know 2). 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 remind 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.

3.3 Single-Minded Valuations

In this section we consider single-minded buyers, interested only in a specific amount of items. We consider the social welfare and revenue maximization cases separately.

3.3.1 Social Welfare Maximization

Before providing our maximization algorithms, let us show the following hardness results.

Theorem 3.3. *The* (SINGLE,WELFARE)*-pricing problem is strongly NP-hard.*

Proof. We prove the claim by providing a polynomial-time reduction from the DENSEST K-SUBGRAPH problem in which, given an undirected graph $H = (V, F)$ and an integer k, the aim is 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, k) of DENSEST K-SUBGRAPH, the reduced instance has:

- 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 $m = k(2|F| + 1) + |F|$;
- preferred size $2|F| + 1$ with valuation $v_{i_v}(2|F| + 1) = 2|F| + 1 + \varepsilon$ for each $v \in V$ and ε suitably small, and preferred size 1 with $v_{i_e}(1) = 1$ 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. □

The above hardness result is tight. In fact, we are able to devise a corresponding PTAS, which is inspired to the classical one for the knapsack problem. Before presenting our approximation scheme, it is useful first to provide some preliminary results, followed by a simpler 2-approximation result.

To this aim, notice first that in the single-minded case, the envy-freeness constraint translates in dictating that each buyer $i \in N$ with strictly positive utility for her unique preferred bundle must be allocated the corresponding items. Consider then the *precedence* set $N^{>}(i)$ of i containing all the buyers k such that: $v_k(m_k)/m_k > v_i(m_i)/m_i$ and for whom there exists a direct path from k to i in G. Since 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, then k has strictly positive utility for hers. Therefore, by the above remark on the envy-freeness constraint, if buyer i gets her preferred bundle, then also every $k \in N^{>}(i)$ must get hers. This implies that if $\sum_{k\in\{i\}\cup N^{>}(i)} m_k > m$, i cannot be allocated in any fair outcome, otherwise it would not be possible to respect the supply constraint.

Generalizing the above argument, given a subset of buyers $B \subseteq N$, and defined the precedence set of B as $B^{>} = \bigcup_{i\in B} N^{>}(i)$, if $\sum_{k\in B\cup B^{>}} m_k > m$, there cannot exist any fair outcome allocating bundles to all the buyers in B.

In a similar way, it is possible to define the set $N^{\leq}(i)$ of all the buyers k such that $v_k(m_k)/m_k \leq v_i(m_i)/m_i$ and there is a direct path from k to i in G. Then, even if by the fairness condition the price proposed to all the buyers in $N^{\leq}(i)$ must be at most the one of i, it is possible to allocate a bundle to i and not to buyers in $N^{\leq}(i)$ by setting price $p_i = v_i(m_i)/m_i$ for i and $p_k = p_i$ for every $k \in N^{\leq}(i)$. In fact, in this way no $k \in N^{\leq}(i)$ would have a strictly positive utility. Let $B^{\leq} = \bigcup_{i\in B} N^{\leq}(i)$ for any $B \subseteq N$.

Keeping the previous observations in mind, consider the procedure described in Algorithm 7. Given (μ, G) and a subset of buyers B, the procedure returns a minimal fair outcome assigning bundles to buyers in B, that is, an outcome that allocates bundles only to the set of buyers $B \cup B^{>}$ and describes a fair pricing on all the buyers.

Notice that in the procedure, in order to fulfill the fairness condition on prices, some buyer $k \in B$ might receive a price during an iteration of the while loop in which another buyer $i \in B$ is considered, since it might be the case that k belongs to $N^{>}(i)$.

The following lemma concerns the outcome returned by Algorithm 7.

Lemma 3.4. *Algorithm 7 returns in polynomial time a fair outcome allocating B and a minimal subset of buyers in $N \setminus B$, if such an outcome exists, otherwise it returns a negative answer.*

Algorithm 7: Minimal-outcome assigning bundles to a specific subset of buyers.

Input: An instance (μ, G) of the (SINGLE,WELFARE)-pricing problem, a subset of buyers $B \subseteq N$.
Output: A minimal fair outcome $(\overline{X}, \overline{p})$ allocating buyers in B.
if $\sum_{i \in B \cup B^{>}} m_i > m$ **then**
  return negative answer;
end
Assign a bundle of size m_i to every buyer $i \in B \cup B^{>}$;
Set $p_i =$ *undefined* for every buyer $i \in N$;
while *there exists a buyer in B with undefined price* **do**
  Pick a buyer $i \in B$ with *undefined* price having minimum value $v_i(m_i)/m_i$;
  Assign price $p_k = v_i(m_i)/m_i$ to every buyer k with *undefined* price in $\{i\} \cup N^{\leq}(i) \cup N^{>}(i)$;
end
Let $p_i = \infty$ for all the buyers i in $N \setminus (B \cup B^{\leq} \cup B^{>})$ with *undefined* prices;
return the resulting outcome $(\overline{X}, \overline{p})$;

Proof. The algorithm runs in polynomial time since it considers each buyer in B at most once in order to propose her a price and changes it at most $O(|B|)$ times before returning an outcome.

If $\sum_{i=i \in B \cup B^{>}} m_i > m$, the claimed outcome does not exist and the procedure returns a negative answer. So, let us assume that $\sum_{i=i \in B \cup B^{>}} m_i \leq m$.

In this case, the minimality of the outcome comes from the fact that it allocates only the strictly needed buyers for a fair outcome allocating buyers in B, that is, buyers in $B \cup B^{>}$.

The envy-freeness constraint holds by observing that no buyer who is not receiving items has strictly positive utility for her preferred bundle. In fact, only the buyers being proposed a finite price and not receiving any bundle could be envious. But each such a buyer k belongs to $B^{\leq}$ and thus is assigned a price equal to $v_i(m_i)/m_i \geq v_k(m_k)/m_k$, where i is the buyer with minimum $v_i(m_i)/m_i$ value in B she can reach via a directed path. Thus k does not have strictly positive utility for her preferred bundle.

Finally, the fairness condition on prices comes directly from the construction, as prices are explicitly assigned in a fashion that fulfills such a requirement. In particular, the price p_i assigned to every buyer i is the minimum one of a buyer in B she can reach via a directed path, infinite if such a path does not exist. □

We are now ready to provide our 2-approximation algorithm.

Lemma 3.5. *There exists a* 2*-approximation algorithm for (*SINGLE,WELFARE*)-pricing.*

Proof. The algorithm first executes a *pre-processing phase* exploiting Algorithm 7 in order to identify the set N' of all the buyers i that can possibly be allocated in a fair outcome, that is, having precedence sets $N^>(i)$ satisfying the supply constraint $\sum_{k \in \{i\} \cup N^>(i)} m_k \leq m$.

By the above considerations, optimal outcomes do not allocate bundles to any $i \in N \setminus N'$. Consider then the ordering of the buyers in the algorithm, in which buyers in N' are listed before the ones in $N \setminus N'$ in non increasing order with respect to their ratios $v_i(m_i)/m_i$, that is, in such a way that $v_k(m_k)/m_k \geq v_i(m_i)/m_i$ if $k < i$.

Notice that, for every $i \in N'$, we have $N^>(i) \subseteq N'$. More precisely, every buyer $k \in N^>(i)$ precede i in the order, that is, are such that $k < i$.

Let $k \in N'$ be the first buyer in the order such that Algorithm 7 on $(\mu, G, \{1, \ldots, k\})$ returns a negative answer, or $k = |N'|$ if such a buyer does not exist. Notice that, for every $i \in N'$, we have $N^>(i) \subseteq N'$. In particular, all the buyers $h \in N^>(i)$ precede i in the order, that is, are such that $h < i$. Then, in an equivalent way, we can define k as the first buyer such that $\sum_{i=1}^{k} m_i > m$, or $k = |N'|$ if $\sum_{i=1}^{|N'|} m_i \leq m$. It is immediate to see that $\sum_{i=1}^{k} v_i(m_i)$ upper bounds the value $opt_{sw}(\mu, G)$, as it is the social welfare achieved by the optimal fair solution, in the case of $\sum_{i=1}^{k} m_i \leq m$, or by the optimal fair outcome for a bigger amount of items $\sum_{i=1}^{k} m_i > m$.

Consider then the following two outcomes:

i. $(\overline{X}, \overline{p})$ from Algorithm 7 on $(\mu, G, \{1, \ldots, k-1\})$, that allocates only buyers $i < k$ and the remaining strictly needed ones;

ii. $(\overline{X}', \overline{p}')$ from Algorithm 7 on $(\mu, G\{k\})$, that allocates buyer k and the remaining strictly needed ones (existing by the pre-processing phase).

By Lemma 3.4 and the pre-processing phase, both the outcomes are fair. Thus, it remains to show that returning the best among the two of them provides the claimed 2-approximation. In fact, the achieved social welfare is equal to $\max\{sw(\overline{X}, \overline{p}), sw(\overline{X}', \overline{p}')\} \geq \frac{sw(\overline{X}, \overline{p}) + sw(\overline{X}', \overline{p}')}{2} \geq \frac{opt_{sw}(\mu, G)}{2}$. ☐

Algorithm 8 can be exploited in order to obtain a PTAS for the (SINGLE,WELFARE)-pricing problem as follows. For a given threshold $\delta > 0$, in Algorithm 8 if no buyer has valuation greater than δ, then the first of the two considered outcomes, that is, the one allocating only buyers $i < k$, has social welfare at least equal to $opt_{sw}(\mu, G) - \delta$. On the other end, if there are *large* buyers, that is, with valuation greater than δ, in every optimal outcome at most $opt_{sw}(\mu, G)/\delta$ of them can receive their preferred bundle.

Algorithm 8: 2-approximation algorithm for the (SINGLE,WELFARE)-pricing problem.

Input: An instance (μ, G) of the (SINGLE,WELFARE)-pricing problem.
Output: An outcome $(\overline{X}, \overline{p})$ fair with respect to G.
Let N' be the subset buyers $i \in N$ such that Algorithm 7 on $(\mu, G, \{i\})$ returns a fair outcome;
Sort all the buyers in N in such a way that the ones in N' come before the ones in $N \setminus N'$ in non increasing order with respect to their values $v_i(m_i)/m_i$;
Let k be the first buyer in N' such that Algorithm 7 on $(\mu, G, \{1, \ldots, k\})$ returns a negative answer, and $k = |N'|$ if such a buyer does not exist;
return the best solution among:
- $(\overline{X}, \overline{p})$ from Algorithm 7 on $(\mu, G, \{1, \ldots, k-1\})$;
- $(\overline{X}', \overline{p}')$ from Algorithm 7 on $(\mu, G, \{k\})$;

Starting from this observation, denoting by L the set of large buyers, the refined algorithm exhaustively considers all the possible subsets $B \subseteq L$ of size at most $opt_{sw}(\mu, G)/\delta$, checks by means of Algorithm 7 whether there exist fair outcomes allocating bundles to the buyers in B, and in this case as in Algorithm 8 completes by assigning bundles to the non-large buyers considered in non increasing order with respect to their values $v_i(m_i)/m_i$, while not violating the supply constraint. Returning the best determined outcome for all the considered sets B, the achieved social welfare is at least $opt_{sw}(\mu, G) - \delta$, and the running time is polynomial in the input size and in $2^{opt_{sw}(\mu,G)/\delta}$. Then, in order to get the claimed PTAS, it is sufficient to set $\delta = \varepsilon \cdot opt_{sw}(\mu, G)$.

Unfortunately, the value $opt_{sw}(\mu, G)$ is not known, but it is possible to exploit Algorithm 8 to the aim of computing a suitable approximation s such that $opt_{sw}(\mu, G)/2 \leq s \leq opt_{sw}(\mu, G)$. Setting $\delta = \varepsilon \cdot s$ finally gives us the desired approximation and time complexity.

We are now ready to claim the following theorem.

Theorem 3.6. *There exists a PTAS for (*SINGLE,WELFARE*)-pricing.*

Proof. By Lemma 3.5, the social welfare s of the outcome returned by the first execution of Algorithm 8 is such that $opt_{sw}(\mu, G)/2 \leq s \leq opt_{sw}(\mu, G)$.

Hence, for the threshold $\delta = \varepsilon \cdot s$, in any fair outcome the number of large buyers receiving bundles can be at most $\frac{opt_{sw}(\mu,G)}{\delta} = \frac{opt_{sw}(\mu,G)}{\varepsilon \cdot s} \leq \frac{2}{\varepsilon}$.

The algorithm exhaustively considers all the possible subsets B of at most $\frac{2}{\varepsilon}$ large buyers, and preliminarily verifies whether B can be allocated in a fair outcome by means of Algorithm 7, which translates into checking if $\sum_{k \in B \cup B^>} m_k \leq m$.

Consider the specific iteration in which a set B of large buyers allocated in an optimal outcome is examined by the algorithm. By the assumption, no large buyer can be in the precedence set $(B \cup N^*)^>$, where N^* contains the non-large buyers allocated in the optimal outcome. Let N' be the set of the non large buyers $i \in N \setminus (B \cup B^>)$ that can possibly be allocated with B, that is, such that:

i. Algorithm 7 on $(\mu, G, B \cup \{i\})$ returns a fair outcome;

ii. $N^>(i) \cap L = \emptyset$.

Consider the ordering of the algorithm in which all the buyers in N' come before the remaining ones in $N \setminus N'$ in non increasing order with respect to their values $v_i(m_i)/m_i$, and let k be the buyer determined by the algorithm, that is, the first in the order such that Algorithm 7 on $(\mu, G, B \cup \{1, \ldots, k\})$ returns a negative answer, or $k = |N'|$ if such a buyer does not exist. As in the 2-approximation algorithm, for every $i \in N'$, we have $N^>(i) \subseteq N'$. In particular, all the buyers $h \in N^>(i)$ precede i in the order, that is, $h < i$. Therefore, denoted by b the number of items that must be allocated to buyers in $B \cup B^>$, this translates in defining k as the first buyer in the order such that $b + \sum_{i=1}^{k} m_i > m$, or $k = |N'|$ if $b + \sum_{i=1}^{|N'|} m_i \leq m$.

Consider then $(\overline{X}_B, \overline{p}_B)$ returned from Algorithm 7 on $(\mu, G, B \cup \{1, \ldots, k-1\})$. We observe that $(\overline{X}_B, \overline{p}_B)$, with the additional allocation of buyer k, would be the optimal fair outcome (if $b + \sum_{i=1}^{|N'|} m_i \leq m$) or the optimal fair outcome for a higher number of supplied items $b + \sum_{i=1}^{k} m_i > m$. Therefore, since k is not a large buyer, the social welfare of $(\overline{X}_B, \overline{p}_B)$, and thus also of the outcome returned by the algorithm, is at least

$$sw(\overline{X}_B, \overline{p}_B) \geq opt_{sw}(\mu, G) - \delta \geq opt_{sw}(\mu, G) - \varepsilon \cdot s \geq (1 - \varepsilon) opt_{sw}(\mu, G).$$

The claim then follows by observing that the algorithm runs in time polynomial in the size of (μ, G) and $2^{1/\varepsilon}$. □

3.3.2 Revenue Maximization

As observed in the proof of Theorem 3.3, the same hardness result holds for the case of revenue.

Corollary 3.7. *The* (SINGLE,REVENUE)*-pricing problem is strongly NP-hard.*

A suitable approximation can also be found for revenue maximization.

Theorem 3.8. *For any $\varepsilon > 0$, there exists a $\frac{\log n}{1-\varepsilon}$-approximation algorithm for* (SINGLE,REVENUE)*-pricing.*

Algorithm 9: PTAS for the (SINGLE,WELFARE)-pricing problem.

Input: An instance (μ, G) of the (SINGLE,WELFARE)-pricing problem, accuracy parameter $\varepsilon > 0$.

Output: An outcome $(\overline{X}, \overline{p})$ fair with respect to G.

Run Algorithm 8 with input (μ, G) and let s be the social welfare of the returned outcome;

Set $\delta = \varepsilon \cdot s$ and let $L \subseteq N$ be the subset of large buyers in N, i.e., with valuation at least δ;

for *every* $B \subseteq L$ *with* $|B| \leq 2/\varepsilon$ **do**

  if *Algorithm 7 on* (μ, G, B) *returns a fair outcome* **then**

    Let N' be the subset of all the non-large buyers $i \in N$ such that Algorithm 7 on $(\mu, G, B \cup \{i\})$ returns a fair outcome and $N^{>}(i) \cap L = \emptyset$;

    Sort all the buyers in N in such a way that the ones in N' come before the ones in $N \setminus N'$ in non increasing order with respect to their values $v_i(m_i)/m_i$;

    Let k be the first buyer of N' such that Algorithm 7 on $(\mu, G, B \cup \{1, \ldots, k\})$ returns a negative answer, and $k = |N'|$ if such a buyer does not exist;

    Let $(\overline{X}_B, \overline{p}_B)$ be the outcome from Algorithm 7 on $(\mu, G, B \cup \{1, \ldots, k-1\})$;

  end

end

return the best solution $(\overline{X}_B, \overline{p}_B)$ among the ones determined for the various sets B;

Proof. Consider the outcome $(\overline{X}, \overline{p})$ returned by the PTAS for social welfare maximization, and without loss of generality let $1, \ldots, k$ be the allocated buyers, listed in non increasing order of values $v_i(m_i)/m_i$. Let h be the index that maximizes $\frac{v_h(m_h)}{m_h} \sum_{i=1}^{h} m_i$. Then, allocating items to buyers $1, \ldots, h$ at price $p'_i = max\{p_i, v_h(m_h)/m_h\}$ for every $i \in N$, we obtain an outcome of revenue at least $\frac{v_h(m_h)}{m_h} \sum_{i=1}^{h} m_i \geq \frac{\sum_{i=1}^{k} v_i(m_i)}{\log n} \geq \frac{(1-\epsilon) opt_{sw}(\mu, G)}{\log n} \geq \frac{(1-\epsilon) opt_r(\mu, G)}{\log n}$.

Clearly such an outcome is fair, as raising prices of an envy-free solution cannot introduce envious buyers, and the obtained prices remain fair with respect to G, as buyers with $p_i > v_h(m_h)/m_h$ do not have directed paths towards the buyers k with $p_k \leq v_h(m_h)/m_h$, and all the remaining buyers have the same price $v_h(m_h)/m_h$. □

3.4 General Valuations

In this section we consider general (unrestricted) valuations, that is, buyers able to express a different valuation for each possible bundle size.

We first observe that the hardness reduction of Theorem 3.3 for the single-minded case is polynomial also for general valuations, as the number of items m is polynomially bounded in the size of the starting DENSEST K-SUBGRAPH instance. More precisely, a reduction with m not polynomial, while being a polynomial time reduction for the single-minded buyers, as in this case the size of the instance depends on $\log m$, would not be polynomial in the general case, as it would generate a reduced instance of exponential size. On the contrary, by the polynomiality of m in the reduction, the negative results of Theorem 3.3 and Corollary 3.7 directly extend also to general valuations.

Corollary 3.9. *The* (GENERAL,WELFARE)- *and* (GENERAL,REVENUE)*-pricing problems are strongly NP-hard.*

We now provide a reduction from general to single-minded instances, that allows to suitably extend the approximation results of the previous section to (GENERAL,WELFARE)- and (GENERAL,REVENUE)-pricing problems. To this aim, we first need to analyze in a careful way the relationship between prices, valuations, and cardinalities of bundles that can appear in the demand set of a buyer $i \in N$ with general valuation, that is, for which i might have maximum utility at some price. For such a class of bundles we are able to prove the following lemma.

Lemma 3.10. *Let $S_i = \{m_i^1, \ldots, m_i^\ell\}$, with $m_i^1 < \ldots < m_i^\ell$, be the set of the bundle sizes that are preferred by i for at least one positive price p. Then the following properties hold:*

i. *For all $1 \leq j < \ell$,*

$$v_i\left(m_i^j\right) < v_i\left(m_i^{j+1}\right)$$

ii. *For all $1 < j < \ell$,*

$$\frac{v_i\left(m_i^j\right) - v_i\left(m_i^{j-1}\right)}{m_i^j - m_i^{j-1}} \geq \frac{v_i\left(m_i^{j+1}\right) - v_i\left(m_i^j\right)}{m_i^{j+1} - m_i^j}$$

iii. *(Boundary condition)*

$$\frac{v_i\left(m_i^1\right)}{m_i^1} \geq \frac{v_i\left(m_i^2\right) - v_i\left(m_i^1\right)}{m_i^2 - m_i^1}$$

Proof. For a bundle of size m_i^{j+1} to appear in the demand set of i, there must exist one positive price p such that $v_i(x) - p \cdot x$ is maximized for $x = m_i^{j+1}$. In particular, it must be $v_i\left(m_i^{j+1}\right) - p \cdot m_i^{j+1} \geq v_i\left(m_i^j\right) - p \cdot m_i^j$, that implies $p \leq \frac{v_i\left(m_i^{j+1}\right) - v_i\left(m_i^j\right)}{m_i^{j+1} - m_i^j}$. Since $m_i^{j+1} > m_i^j$, in order for p to be positive, property ***i.*** must hold.

As also m_i^j is in the demand set, there must also exist one positive price p' such that $v_i(x) - p' \cdot x$ is maximized for $x = m_i^j$. Similarly as above, p' must be such that

$p' \leq \frac{v_i(m_i^j)-v_i(m_i^{j-1})}{m_i^j - m_i^{j-1}}$. Moreover, since $v_i\left(m_i^j\right) - p' \cdot m_i^j \geq v_i\left(m_i^{j+1}\right) - p' \cdot m_i^{j+1}$ then it holds $p' \geq \frac{v_i(m_i^{j+1})-v_i(m_i^j)}{m_i^{j+1}-m_i^j}$, and therefore:

$$\frac{v_i\left(m_i^j\right) - v_i\left(m_i^{j-1}\right)}{m_i^j - m_i^{j-1}} \geq p' \geq \frac{v_i\left(m_i^{j+1}\right) - v_i\left(m_i^j\right)}{m_i^{j+1} - m_i^j},$$

which implies property ***ii.***

Finally, consider bundle size m_i^1 and the positive price p'' at which m_i^1 is preferred. As previously discussed, $p'' \geq \frac{v_i(m_i^2)-v_i(m_i^1)}{m_i^2-m_i^1}$. Furthermore, in order for i to have a non-negative utility for m_i^1, it must also be $p'' \leq \frac{v_i(m_i^1)}{m_i^1}$, which proves property ***iii.*** □

In order to provide our reduction from general valuations instances to single-minded ones, let us define for the sake of brevity the following values v_i^j:

$$v_i^j = \begin{cases} \frac{v_i(m_i^j)}{m_i^j} & \text{if } j = 1 \\ \frac{v_i(m_i^j)-v_i(m_i^{j-1})}{m_i^j-m_i^{j-1}} & \text{if } 1 < j \leq \ell \end{cases}$$

Notice that the values v_i^j are decreasing in j. Moreover, in the proof of Lemma 3.10, we indirectly observe that a given bundle of size m_i^j is in the demand set of buyer i if and only if $p_i \in \left[v_i^{j+1}, v_i^j\right]$. In particular, v_i^j is the maximum price at which bundle size m_i^j is in the demand set of buyer i.

These observations allow us to represent a buyer with a general valuation by means of a chain of ℓ single-minded buyers (see Figure 3.2). More precisely, given a market μ in which buyers have general valuations, consider the market with single-minded buyers μ' and graph G' built as follows:

- For each buyer i in μ, consider the set $S_i = \{m_i^1, \ldots, m_i^\ell\}$, with $m_i^1 < \ldots < m_i^\ell$, of the bundle sizes that are in i's demand set for some positive price p_i. To each $m_i^j \in S_i$ associate a single-minded *marginal* buyer i_j in μ', with valuation v_{i_j} for her preferred bundle defined as follows:

$$\left\{ \begin{array}{ll} v_{i_j}\left(m_i^j\right) = v_i\left(m_i^j\right) & \text{if } j = 1 \\ v_{i_j}\left(m_i^j - m_i^{j-1}\right) = v_i\left(m_i^j\right) - v_i\left(m_i^{j-1}\right) & \text{if } 1 < j \leq \ell \end{array} \right\}$$

- For each buyer i of μ, add to G' arcs (i_j, i_{j+1}) for each $1 \leq j < \ell$.

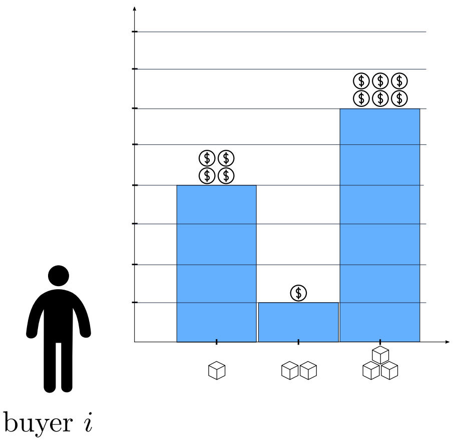

(a)

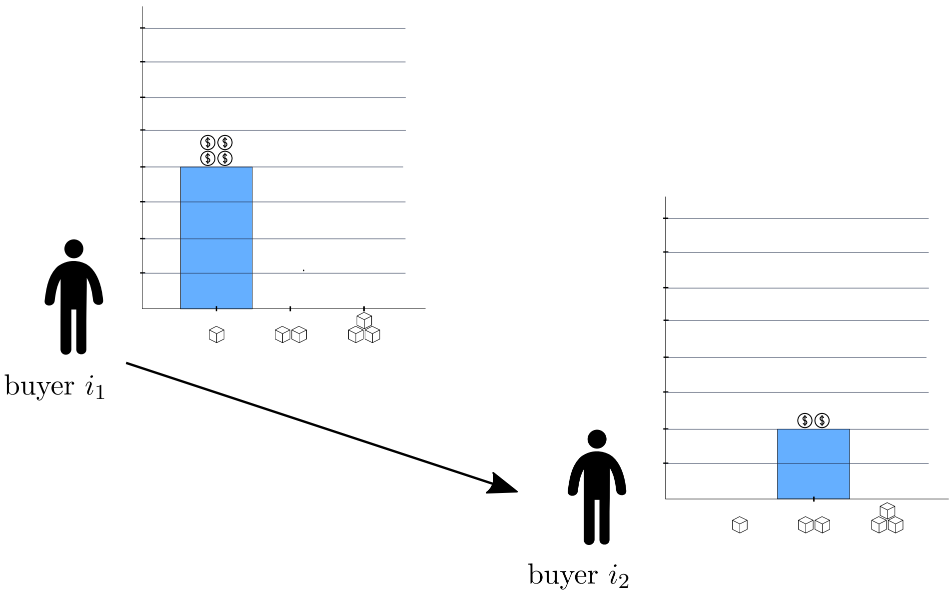

(b)

FIGURE 3.2: A simple example of the reduction from instances with buyers with general valuations to single-minded ones. In Figure 3.2(a) we depicted the original buyer i with general valuations. Notice that the bundle of size 2 is never in the demand set of i, no matter the price offered by the seller. In Figure 3.2(b) we have the two resulting single-minded buyers on the social network defined by the single arc between them. The former buyer is associated to $m_i^1 = 1$ while the latter to the bundle $m_i^2 = 3$. More precisely, buyer i_1 in Figure 3.2(b) has single-minded valuation $v_i(1) = v_{i_1}(m_i^1) = v_i(m_i^1) = 4$ while buyer i_2 has single-minded valuation equal to $v_{i_2}(2) = v_{i_2}(m_i^2 - m_i^1) = v_i(m_i^2) - v_i(m_i^1) = v_i(3) - v_i(1) = 2$

- For each pair of buyers i, l in market μ such that (i, l) belongs to G, add to G' an arc between every pair of marginal buyers associated to i and l.

Intuitively, in such a reduction, the allocation of a bundle of size m_i^j in the demand set of buyer i at price p_i corresponds to the allocation of bundles of respective preferred sizes $m_i^1, m_i^2 - m_i^2, \ldots, m_i^j - m_i^{j-1}$ to the prefix of the marginal buyers $i_1, \ldots, i_j$ associated to i in μ', still all at price p_i. Notice that, by the above arguments, since the bundle of size

m_i^j is in the demand set of i, it must be $p_i \in \left[v_i^{j+1}, v_i^j\right]$, and by Lemma 3.10 each buyer of μ' has non-negative valuation for her preferred bundle. Notice also that the overall amount of allocated items to the marginal buyers is m_i^j, the sum of their valuations is $v_i(m_i^j)$, and the sum of their utilities is $u_i(m_i^j)$. In other words, any solution for the reduced instance can be transformed into an equivalent one for the general case that maintains the same number of allocated items and the same social welfare and revenue.

However, the two instances are not completely equivalent in the case of ties, that is, if at a given price some marginal buyers of the same chain have utility exactly equal to 0 for their preferred bundles, it might be possible to allocate marginal buyers that do not form a prefix. While this could be useful to better fulfill the supply constraint, it would compromise the correspondence between the solutions of the two instances. However, such a correspondence can be maintained if we restrict to *conservative solutions*, that is, fair outcomes that allocate only (possibly empty) prefixes of buyers in the reduced instance. Clearly, such solutions have a one-to-one correspondence with the solutions of the original instance with general valuations.

Let us then call *conservative* an algorithm for determining a fair outcome for single-minded instances if when executed on a reduced instance it returns only conservative solutions. Let us also say that a conservative algorithm for social welfare (resp. revenue) maximization has conservative approximation ratio ρ if it returns a fair solution whose social welfare (resp. revenue) is at least an ρ fraction of the best conservative fair outcome.

Theorem 3.11. *If there exists an algorithm with conservative ratio ρ for* (SINGLE, WELFARE)*-pricing (resp.* (SINGLE, REVENUE)*-pricing), then there exists also a ρ-approximation algorithm for* (GENERAL, WELFARE)*-pricing (resp.* (GENERAL, REVENUE)*-pricing).*

Proof. Let us first consider the social welfare maximization case.

Given any instance (μ, G) of (GENERAL,WELFARE)-pricing, let (μ', G') be the reduced instance of (SINGLE,WELFARE)-pricing constructed according to the above reduction. Let us denote by $opt^c_{sw}(\mu', G')$ the value of an optimal conservative fair outcome for (μ', G'). We prove the claim by showing that (μ', G') satisfies the following two properties:

i. $opt^c_{sw}(\mu', G') \geq opt_{sw}(\mu, G)$;

ii. a fair conservative outcome $(\overline{X}', \overline{p}')$ for (μ', G') can be directly transformed into a fair outcome $(\overline{X}, \overline{p})$ for (μ, G) with $sw(\overline{X}, \overline{p}) \geq sw(\overline{X}', \overline{p}')$.

This clearly proves the theorem.

Let us first prove that $opt^c_{sw}(\mu', G') \geq opt_{sw}(\mu, G)$. To this aim, we show how to transform a fair outcome $(\overline{X}, \overline{p})$ for (μ, G) into a fair conservative one $(\overline{X}', \overline{p}')$ for (μ', G') such that $sw(\overline{X}', \overline{p}') = sw(\overline{X}, \overline{p})$.

Let $m_i^1 < \ldots < m_i^\ell$ be the set of the bundle sizes that are preferred by a given buyer i of μ for at least one positive price p. If $|X_i| = m_i^j$ in $(\overline{X}, \overline{p})$, then in $(\overline{X}', \overline{p}')$ 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 μ, $|X_i| = m_i^j = m_i^1 + \sum_{k=2}^{j}(m_i^k - m_i^{k-1}) = \sum_{k=1}^{j} |X'_{i_k}|$, and $v_i(|X_i|) = \sum_{k=1}^{j} v_{i_k}(|X'_{i_k}|)$. Thus, $(\overline{X}', \overline{p}')$ is feasible, as it allocates the same number of items, and has the same social welfare of $(\overline{X}, \overline{p})$. It remains to show that $(\overline{X}', \overline{p}')$ is also fair. To this aim, we observe that, since m_i^j is the size of a preferred bundle of buyer i in μ, and since by Lemma 3.10 we have $v_i^k \geq v_i^j$ for $k < j$, we must have that $v_i^k \geq p_i$ for every $k \leq j$. With the same argument, we can say that $p_i \geq v_i^k$ for all $k \geq j$, which means that no allocated buyer in $(\overline{X}', \overline{p}')$ has negative utility, nor any non-allocated buyer is envious. 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 conservative outcome for (μ', G') with $sw(\overline{X}', \overline{p}') = sw(\overline{X}, \overline{p})$.

Consider now a fair conservative outcome $(\overline{X}', \overline{p}')$ for μ', let i be a buyer of μ for which such an outcome allocates a non-empty prefix of marginal buyers $i_1, \ldots, i_j$, and let p'_{i_j} be the price proposed to i_j. Notice that, by the fairness constraint on prices $p'_{i_k} \leq p'_{i_j}$ for every $k < i$ and without loss of generality, we can assume $p'_{i_k} = p'_{i_j}$ for $k < j$. By Lemma 3.10, this does not provide a negative utility to the marginal buyers i_k with $k < i$, it maintains the fairness of the pricing p', as it respects the constraints of the last prices p'_{i_j} of the allocated prefixes, and finally it does not decrease the social welfare or the revenue. Then, (μ', G') directly translate to a fair outcome $(\overline{X}, \overline{p})$ for (μ, G) with $sw(\overline{X}, \overline{p}) \geq sw(\overline{X}', \overline{p}')$, simply by assigning in $(\overline{X}, \overline{p})$ a bundle of size m_i^j to buyer i at price $p_i = p'_{i_j}$.

The proof for the case of revenue maximization uses the same identical arguments, simply considering the revenue in place of the social welfare. □

Unfortunately, the algorithms presented in the previous section for single-minded instances are not conservative. However, we are able to equip them with some tailored refinements in order to infuse such a property, thus obtaining corresponding approximation algorithms by virtue of Theorem 3.11.

In particular, the definition of the precedence set of each marginal buyer i_j associated to a buyer i of μ is modified in such a way that $N^{>}(i_j)$ also contains all the previous buyers in her chain, that is, every i_k with $k < i$. Notice that, by Lemma 3.10, for such buyers it must be $v_i^k \geq v_i^j$, but as the inequality is not strict, this containment was not holding if $v_i^k = v_i^j$. As a consequence of the modified definition of the precedence sets, the algorithms for single-minded instances in the previous section always return conservative outcomes, as for every allocated buyer i_j also the ones in $N^{>}(i_j)$ are allocated, and thus all the prefix $i_1, \ldots, i_j$. However, a further refinement for preserving the approximation in the algorithms is necessary. In particular, for the PTAS (and for Algorithm 8) this is needed for imposing the following crucial property: in the sorting of the buyers in the set N', when a buyer is considered in the algorithm, all the ones in her precedence set occur before her in the order. This is accomplished by properly breaking ties when two buyers i_k, i_j with $k < j$ of a same buyer i of μ are such that $v_i^j = v_i^k$. In this case, it suffices to put i_k before i_j in the ordering. The same expedient aimed at breaking ties is adopted in the revenue maximization algorithm of Theorem 3.8.

With such tailored refinements in mind, we can prove the following corollary.

Corollary 3.12. *There exists a PTAS for* (GENERAL, WELFARE)*-pricing and a* $\frac{\log n + \log m}{1-\varepsilon}$*-approximation algorithm for* (GENERAL, REVENUE)*-pricing, for any fixed* $\varepsilon > 0$.

Proof. A proof, identical to the one presented in Theorem 3.6, shows that the PTAS for single-minded instances with the above refinements has conservative ratio $\frac{1}{1-\varepsilon}$. Analogously, being $n' = n \cdot m$ the number of buyers in the reduced instance (μ', G'), the conservative ratio of the revenue maximization algorithm of Theorem 3.8 is $\frac{\log n'}{1-\varepsilon} = \frac{\log n + \log m}{1-\varepsilon}$. The claim then follows by Theorem 3.11. □

3.5 Special Classes of Social Graphs

We now focus on specific social graph topologies and present optimal polynomial time algorithms or improved approximations.

Notice that, while in the single-minded case the social welfare can be arbitrarily approximated, only a logarithmic approximation is achieved for the revenue. However, a better approximation can be accomplished if the social graph is undirected. We remark that under this assumption the problem remains computationally difficult. This comes directly by restriction on complete social graphs, which correspond to the classical problem without price discrimination, shown to be NP-hard in [82].

On the other hand, building on the results in [82] and on techniques developed in [20, 56], we are able to provide a constant approximation algorithm (see Algorithm 10).

Algorithm 10: $(2+\varepsilon)$-approximation algorithm for (SINGLE,REVENUE)-pricing for undirected social graphs

Input: An instance (μ, G) of the (SINGLE,REVENUE)-pricing problem with G undirected, accuracy parameter $\varepsilon > 0$.
Output: An outcome $(\overline{X}, \overline{p})$ fair with respect to G.
Let $N_1, \ldots, N_l$ be the connected components of G;
Compute instance $K_{\varepsilon/4}(\mu, G)$ of MULTIPLE-CHOICE KNAPSACK;
$S \longleftarrow FPTAS-$MULTIPLE-CHOICE KNAPSACK$(K_{\varepsilon/4}(\mu, G), \varepsilon/4)$;
Extract outcome $(\overline{X}, \overline{p})$ as suggested by the solution S for $K_{\varepsilon/4}(\mu, G)$;
return $(\overline{X}, \overline{p})$;

Theorem 3.13. *The* (SINGLE,REVENUE)*-pricing problem, restricted to undirected social graphs, admits a* $(2+\varepsilon)$*-approximation.*

Proof. Let us first informally discuss the basic intuition behind the algorithm. Given an instance (μ, G) of (SINGLE,REVENUE)-pricing, if G is undirected, then any pricing fair with respect to G must assign the same price to all of the buyers belonging to the same connected component. Thus, if $N_1, \ldots, N_\ell$ are the subsets of buyers corresponding to the connected components of G, the problem translates to determining how many items must be allocated to each N_h. Let μ_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 μ_h and for every amount of items or supply assigned to μ_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 $K(\mu, G)$ having a class O_h for each μ_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 μ_h with supply j. In particular, each $o_{h,j}$ has profit $z_{h,j}$ 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 $K(\mu, G)$ induces a revenue-maximizing fair outcome $(\overline{X}, \overline{p})$ for (μ, G). Unfortunately, since (μ, G) has size depending on $\log m$ and $K(\mu, G)$ 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 $K(\mu, G)$ we do not insert in each class O_h the objects $o_{h,j}$ for all the possible sizes $j = 1, \ldots, m$, but only for the at most $\lfloor \log_{\frac{3}{2}} m \rfloor + 2$ sizes $j = (\frac{3}{2})^0, \ldots, (\frac{3}{2})^{\lfloor \log_{\frac{3}{2}} m \rfloor}, m$. The knapsack capacity is set to $\frac{3}{2}m$.

In order to measure the loss in revenue due to this sampling, consider an optimal fair outcome $(\overline{X}, \overline{p})$ for (μ, G). If j_h is the number of objects that $(\overline{X}, \overline{p})$ assigns to a given submarket μ_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 $K(\mu, G)$ 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 μ_h with supply j, its profit $z_{h,j}$ is at least the one that $(\overline{X}, \overline{p})$ collects on the subset of buyers N_h of μ_h. Therefore, all such objects correspond to a solution S for $K(\mu, G)$ of profit at least $r(\overline{X}, \overline{p})$ and total size at most $\frac{3}{2}m$. Hence, the optimal solution S^* for $K(\mu, G)$ achieves at least such a revenue. In order to extract an outcome allocating at most m items, we can split S^* into 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 sizes 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^* does 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 μ_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 $(\overline{X}, \overline{p})$ for (μ, G).

Unfortunately, we are not really able to optimally solve $K(\mu, G)$, nor the subproblems for determining the profits of the objects $o_{h,j}$. However, in both cases we can use respectively the FPTAS for MULTIPLE-CHOICE KNAPSACK from [80] and the one from [82] for the classical revenue maximization problem in multi-unit markets, losing only an additional ε factor in the final approximation.

More precisely, we construct an instance $K_{\varepsilon/4}(\mu, G) = (\overline{O}, \overline{z}, \overline{w}, k)$ of MULTIPLE-CHOICE KNAPSACK as follows:

- $k = \frac{3}{2}m$
- $O_h = \{o_{h,j} : j = \frac{3}{2}^0, \ldots, \frac{3}{2}^{\lfloor \log_{\frac{3}{2}} m \rfloor}, m\}$ for $h = 1, \ldots l$
- z_{hj} = revenue of the envy-free (non-discriminatory) outcome returned by the FPTAS for the submarket N_h with j items on sale and accuracy parameter $\varepsilon/4$
- $w_{h,j} = j$

Then, since every $o_{h,j}$ represents an envy-free outcome for μ_h with supply j which is a $(1 + \varepsilon/4)$-approximation of the optimal one, denoted as $S^*_{\varepsilon/4}$ the optimal solution for $K_{\varepsilon/4}(\mu, G)$ and by S^* the one for $K(\mu, G)$, we have that $\sum_{o_{h,j} \in S^*_{\varepsilon/4}} z_{h,j} \geq \frac{\sum_{o_{h,j} \in S^*} z_{h,j}}{1+\varepsilon/4}$.

Therefore, by running the FPTAS with accuracy parameter $\varepsilon/4$ on $K_{\varepsilon/4}(\mu, G)$, we finally obtain a solution S such that $\sum_{o_{h,j} \in S} z_{h,j} \geq \frac{\sum_{o_{h,j} \in S^*_{\varepsilon/4}} z_{h,j}}{1+\varepsilon/4} \geq \frac{\sum_{o_{h,j} \in S^*} z_{h,j}}{(1+\varepsilon/4)^2} \geq \frac{\sum_{o_{h,j} \in S^*} z_{h,j}}{1+\varepsilon/2}$.

By the same argument above, we can finally split S in two sets S_1 and S_2, each containing a subset of objects of overall size at most m, in such a way that at least one of them will have total profit at least $\frac{1}{2} \frac{\sum_{o_{h,j} \in S^*} z_{h,j}}{1+\varepsilon/2} = \frac{\sum_{o_{h,j} \in S^*} z_{h,j}}{2+\varepsilon}$. Hence, similarly as above, allocating items only to the connected components corresponding to the selected objects $o_{h,j}$ (according to the solutions determined by the FPTAS for the related submarkets μ_h with supply j), and assigning an infinite price to all the buyers of the other connected components, we can finally determine a fair outcome for (μ, G) collecting a revenue which is at least a $2 + \varepsilon$ fraction of the optimal one. □

Because of the transitive nature of the constraints on a fair-pricing, in the case of directed social graphs, the following remark holds.

Remark 3.14. In the case of directed social graphs, the seller must offer identical prices to buyers in the same strongly-connected component.

Given the previous remark, a convenient way to represent G is through the graph of its strongly-connected components. More precisely, let $N_1, \ldots, N_\ell$ be the subsets of buyers corresponding to the strongly-connected components of G, and μ_h be the *submarket* reduced to the set of buyers in N_h. Then, in such a graph each node represents a component N_h, and there is an arc between two nodes if the same happens between the corresponding components. Clearly, the graph of the strongly-connected components is a directed acyclic graph (DAG), and properly restricting its topology we can find meaningful more tractable sub-cases.

To this aim, let us first generalize a result from [82], which states that, given any multi-unit market μ, the set of the optimal prices has polynomial size. In particular, let S_i be the set of bundle sizes that are preferred by i for at least one positive price p. Then, the following lemma holds.

Lemma 3.15. *There exist fair outcomes for (μ, G) maximizing the social welfare or the revenue in which all the prices proposed to the buyers belong to the set $\mathbb{P}^{opt} = \{\frac{v_i(j)}{j} | i \in N, j \in S_i\}$.*

Proof. Assume that $(\overline{X}^{opt}, \overline{p})$ is a fair outcome for (μ, G) maximizing the social welfare or the revenue, and assume that there exists a buyer i whose proposed price p_i does not belong to $\mathbb{P}^{opt}$. Let $p^\top \in \mathbb{P}^{opt}$ be the maximum ratio in $\mathbb{P}^{opt}$ smaller than p_i, and let $p^\perp \in \mathbb{P}^{opt}$ be the minimum ratio in $\mathbb{P}^{opt}$ greater than p_i. Consider then the outcome

$(\overline{X}^{opt}, \overline{p'})$ where the price proposed to any buyer i is $p'_i = p^\top$ if $p^\perp < p_i < p^\top$, $p'_i = p_i$ otherwise.

Since the two outcomes $(\overline{X}^{opt}, \overline{p})$ and $(\overline{X}^{opt}, \overline{p'})$ have the same allocation, $(\overline{X}^{opt}, \overline{p'})$ still satisfies the supply constraint. Furthermore, it also satisfies the individual rationality, since the price is increased to $p^\top \in \mathbb{P}^{opt}$ only for the buyers with previous price p_i between $p^\perp$ and $p^\top$. In fact, each such a buyer either already had negative utility, and thus was not allocated, or her new utility does not become negative. Finally, $(\overline{X}^{opt}, \overline{p'})$ is envy-free, because each buyer with strictly positive utility continues to receive her preferred bundle, and it is fair, as the modification of the prices respects the fairness constraints.

In conclusion, $(\overline{X}^{opt}, \overline{p'})$ is a fair outcome with social welfare and revenue at least equal to the ones $(\overline{X}^{opt}, \overline{p})$. Thus, by iterating the above procedure, we can finally obtain a fair optimal outcome respecting the conditions of the claim. □

Notice that $\mathbb{P}^{opt}$ has always polynomially bounded size. In fact, $|\mathbb{P}^{opt}| \leq n$ for single-minded instances and $|\mathbb{P}^{opt}| \leq n \cdot m$ for general ones.

Let us now consider special directed topologies allowing better results. A first relevant case is when G has a constant number of strongly-connected components. In fact, like undirected graphs, this might represent geographically apart submarkets (the strongly-connected components), with some limited possibility of interaction, like the case of national markets, in which there are a few countries, each containing millions or billions of individuals.

Under this assumption, the following positive results hold.

Theorem 3.16. *The* (SINGLE,WELFARE)- *and* (SINGLE,REVENUE)*-pricing problems restricted to social graphs with a constant number of strongly-connected components admit an FPTAS.*

Proof. Given an instance (μ, G) of the (SINGLE,WELFARE)- or (SINGLE,REVENUE)-pricing problem, let $N_1, \ldots, N_\ell$ be the subsets of buyers corresponding to the strongly-connected components of G, and μ_h be the submarket corresponding to the set of buyers in N_h, for $h = 1, \ldots, \ell$.

Similarly to the undirected case, any fair pricing with respect to G must assign the same price to all of the buyers belonging to the same submarket or strongly connected component. Then, since ℓ is constant, by Lemma 3.15 it is possible exhaustively search all the ℓ-tuples of prices in $\mathbb{P}^{opt}$. For every such a tuple $\overline{p} = (p_1, \ldots, p_\ell)$, where each p_h is the price offered to all the buyers in N_h, we first check whether $\overline{p}$ respects the fairness constraints, that is, if $p_h \leq p_k$ whenever there is a directed path from component h to component k in G. If this holds, considering that items must be allocated to all the buyers with strictly positive utility, we further check if a feasible allocation for $\overline{p}$ exists, that is, if less than m items must be allocated to such buyers. If yes, we focus on the subset N' of the remaining buyers with utility exactly equal to 0 for their preferred bundles. In fact, these are the only left ones that may possibly receive other items. Denoted as m' the residual supply, that is, the number of items that remain to be assigned, we formulate the problem of allocating bundles to buyers in N' so as to maximize the welfare or the seller's revenue as an equivalent knapsack instance $K(\mu, \overline{p})$, in which there is an object o_i with profit and weight m_i for every $i \in N'$. The knapsack capacity is set to m'. We can then run the FPTAS for the problem in order to obtain a $(1+\varepsilon)$-approximation for such a price. Considering the best solution obtained for all the price vectors $\overline{p}$, and allocating bundles to buyers accordingly, yields the claimed $(1+\varepsilon)$-approximation for the pricing problem in time polynomial in the input size and in $1/\varepsilon$, hence the claim. □

Let us now focus on general valuations. In this case, the fact that the input size now is polynomial in m instead of $\log m$, gives us a chance to achieve better tractability results. In particular, the following theorem holds.

Theorem 3.17. *There exists a polynomial time algorithms for* (GENERAL,WELFARE)-*pricing problem (resp* (GENERAL,REVENUE)-*pricing problem) if the social graph is undirected or contains a constant number of strongly-connected components.*

Proof. Let us first focus on revenue maximization. A polynomial time algorithm can be obtained by the same initial construction proposed in the proof of Theorem 3.13, by observing that, since the size of the input instance (μ, G) now is polynomial in m, sampling is no longer necessary. Moreover, for each submarket μ_h corresponding to a connected component of G, the determination of the optimal market outcomes associated to the objects $o_{h,j}$ can be performed in polynomial time by means of the algorithm shown in [82]. Finally, since the initial knapsack instance $K(\mu, G)$ has capacity bounded by m, an optimal solution can be determined in polynomial time by exploiting the pseudo-polynomial time algorithm for MULTIPLE-CHOICE KNAPSACK described in [80].

For what concerns the social welfare, we observe that the algorithm of [82] can be slightly modified for efficiently determining envy-free (non-discriminatory) outcomes maximizing the social welfare in multi-unit markets. We can then use exactly the same approach described above, by setting in the knapsack instance $K(\mu, G)$ any $z_{h,j}$ equal to the social welfare of an optimal envy-free outcome with supply j for the submarket μ_h, finally obtaining a polynomial time algorithm also for the social welfare maximization problem.

For what concerns the case in which G has a constant number of strongly-connected components, exactly the same proof of Theorem 3.16 applies by observing that, since the size of the instance now is polynomial in m, the knapsack instances $K(\mu, \overline{p})$ can now be solved in polynomial time. □

Another interesting case for general valuations that turns out to be tractable is given by the arborescences of strongly-connected components, which are relevant in the hierarchical profiling of users, where each strongly-connected component corresponds to buyers of a same class, and to different classes of premium buyers are offered different special discounts. Moreover, arborescences properly contain directed paths, an abstraction for temporal price discrimination, which is another meaningful and widespread kind of price discrimination.

Definition 3.18. An *arborescence* is a directed graph in which for a node r, called the root, and any other node v, there exists a unique directed path which connects r to v.

The following positive result holds.

Theorem 3.19. *There exists a polynomial time optimal algorithm for* (GENERAL,WELFARE)- *and* (GENERAL,REVENUE)*-pricing if the graph of the strongly-connected components of G is an arborescence.*

Proof. Let T be the arborescence of the strongly-connected components $N_1, \ldots, N_\ell$ of G, in which each node h corresponds to a strongly-connected component N_h, and let r be the root of T. In order to provide an optimal fair outcome for the input instance (μ, G) of (GENERAL,REVENUE)-pricing, we build partial solutions starting from the leaves of T and reaching the root r through a post-order visit of T.

We use a dynamic programming approach, where the quantity $M_h(p, j, b)$ represents the value of an optimal allocation which assigns at most b items at prices greater or equal to p to buyers in the subgraph of T rooted at h, and at most j items at price exactly p to the buyers in N_h. Clearly, the quantity $M_r^* = \max_{p,j} M_r(p, j, m)$ gives the solution to the problem. Recalling that by Lemma 3.15 it is possible to restrict to the polynomial size set of prices $\mathbb{P}^{opt}$, let us now discuss how it is possible to compute the above quantities and the corresponding outcomes in polynomial time.

First of all, if h is a leaf, $M_h(p,j,b)$ for every $p \in \mathbb{P}^{opt}$, $j = 1,\dots,m$ and $b = j,\dots,m$ can be computed as follows:

$$M_h(p,j,b) = \begin{cases} OPT_h(p,j) & \text{if there exists an envy-free allocation which sells} \\ & \text{at most } j \text{ items at price } p \text{ to buyers in } N_h, \\ \perp & \text{otherwise,} \end{cases}$$

where $OPT_h(p,j)$ is the optimal revenue of an envy-free outcome which sells at most j items in the submarket of the buyers in N_h at price p. Such a revenue can be computed by solving the following equivalent instance of the MULTIPLE-CHOICE-KNAPSACK problem:

- j is the knapsack's capacity;
- $O_i = \{o^{p,j'}$: a bundle of size j' is in the demand set of buyer $i \in N_h$ a price $p\}$, with
  - $z_{p,j'} = p \cdot j'$ for each $o^{p,j'} \in O_i$,
  - $w_{p,j'} = j'$ for each $o^{p,j'} \in O_i$.

If h is not a leaf, $M_h(p,j,b) = \perp$ if there exists no envy-free outcome which sells at most j items at price p to buyers in N_h, otherwise we can rely on the fact that we have already solved all the subproblems associated to the children of h and we can formulate the problem of optimally joining these solutions as the following instance of MULTIPLE-CHOICE-KNAPSACK:

- b is the knapsack's capacity;
- $O_i = \{o^{p,j'}$: a bundle of size j' is in the demand set of buyer $i \in N_h$ a price $p\}$, with
  - $z_{p,j'} = p \cdot j'$ for each $o^{p,j'} \in O_i$
  - $w_{p,j'} = j'$ for each $o^{p,j'} \in O_i$
- $O_k = \{o^{p',j',b'} : M_k(p',j',b') \neq \perp,\ p' \in \mathbb{P}^{opt},\ p' \geq p, j' = 1,\dots,b,\ b' = j',\dots,b\}$, for each child k of h in T, with
  - $z_{p',j',b'} = M_h(p',j',b')$ for each $o^{p',j',b'} \in O_k$
  - $w_{p',j',b'} = b'$ for each $o^{p',j',b'} \in O_k$

As we said, the maximum revenue for the market can be expressed as

$$M_r^* = \max_{p,j} M_r(p,j,m)$$

for $p \in \mathbb{P}^{opt}$ and $j = 1, \ldots, m$. Such a quantity can be computed in polynomial time by observing that all the above MULTIPLE-CHOICE-KNAPSACK instances have polynomially bounded capacity and thus can be solved in polynomial time by the pseudo-polynomial algorithm of [80]. Moreover, since we can restrict on prices in $\mathbb{P}^{opt}$, the dynamic programming algorithm requires to solve only a polynomial number of such instances. Thus, to prove the claim, it is sufficient to observe that a fair outcome with the corresponding revenue can be determined by just keeping track of the solutions yielding the various quantities. We remark that an object $o^{p,j} \in O_i$ in a knapsack solution represents the allocation of a bundle of size j at a price p to the corresponding buyer $i \in N_h$, while an object $o^{p,j,b} \in O_k$ the adoption of a solution of revenue $M_k(p, j, b)$ for the buyers in the subtree rooted at k in T.

For what concerns the social welfare maximization, it suffices to modify, in the above instances of MULTIPLE-CHOICE-KNAPSACK, the profit of the objects in the classes O_i associated to the buyers in $i \in N_h$ as $z_{p,j'} = v_i(j')$ for each $o^{p,j'} \in O_i$. Apart from this change, the proof is identical. □

3.6 Conclusions and Future Work

We proposed a framework able to capture many realistic scenarios in which discriminatory pricing is acceptable for buyers. We considered both social welfare and revenue maximization, under different hypotheses on buyers' valuations and social graph topologies, providing polynomial time optimal algorithms, hardness results, and suitable approximations. We also provided a useful reduction from general to single-minded instances that is of independent interest and can be possibly used also in other settings concerning multi-unit markets.

Besides 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 also be interesting to consider other relevant specific social graph topologies, and to find better approximation results for arborescences of strongly-connected components in the case of single-minded buyers. Moreover, what about different notions of envy-freeness, such as *pair envy-freeness*, *social envy-freeness*, *proportionality*, and so forth.

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

As a last observation, we would like to remark that our framework for fair price discrimination is very general and we believe that it will capture future research attention.

Chapter 4

Inequity Aversion

We build upon the framework for fair price discrimination in markets presented in the previous chapter and other models proposed in the literature concerning differential pricing in social networks, considering a setting in which multiple units of a single product must be sold to selected buyers so as to maximize the seller's revenue or the social welfare, while limiting the differences of the prices offered to social neighbors. We first consider the case of general social graph topologies, and provide optimal or nearly-optimal hardness and approximation results for the related optimization problems under various meaningful assumptions, including a polylogarithmic lower bound on the achievable revenue under the unique game conjecture. Then, we focus on topologies that are typical of social networks. Namely, we consider graphs where the node degrees follow a power-law distribution, and show that it is possible to obtain constant or good approximations for the sellers revenue maximization with high probability, thus improving upon the general case.

4.1 Introduction

A powerful mechanism often used in combination with price discrimination is the finely-grained selection of buyers, for instance, exploiting targeted advertisement mechanisms like those offered by Google or Facebook. This mechanism is able to greatly improve seller's revenue but also social welfare as we will prove in this chapter. Companies with the technology enabling them to deploy this kind of targeted sales, still have the need to detect the best subset of buyers to whom directing their special offers. In the present chapter, relying on a framework for price discrimination similar to the one described in the previous chapter, we will study the problem of selecting the optimal set of buyers to target in order to maximize seller's revenue or social welfare.

Frameworks more closely related to the present chapter have been studied in [6, 7, 55]. In particular, in [6] the authors defined a form of price discrimination constrained by a social graph, precisely requiring a difference bounded by an additive constant in the prices offered to two social neighbors. Several results on revenue maximization have been provided, also in case of buyers preselection. Furthermore, [7] extended the previous model to a setting allowing the assignment of more than one item per node, which corresponds to multi-unit markets with single-minded buyers and unlimited supply.

In this chapter we study market scenarios in which the seller is able to propose prices in $\mathbb{R} \cup \{\perp\}$, meaning the exclusion of buyers receiving the latter price. Notice that such a variant has been considered also in [6] and [7]. In their model, the possibility of setting bottom prices increases the maximum achievable revenue and at the same time makes the problem intractable and hard to approximate within a factor of $O(n^{1-\varepsilon})$, for any $\varepsilon > 0$. For our setting, we show an $O(\log n)$-approximating algorithm, plus a hardness of approximation which is optimal up to a polylogarithmic factor. Namely, we show that the problem is Unique-Game-Hard to approximate within a factor $\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$.

The model of fair price discrimination investigated in Chapter 3, while extending the above frameworks by allowing general valuations, does not allow the exclusion of buyers by means of special ($\perp$) prices. Furthermore, it does not consider any slackness in fair pricing.

As already mentioned, we borrow the model of fair price discrimination from Chapter 3 and extend it with the features of buyers preselection and additive slackness in fairness constraints considered in [6, 7].

4.1.1 Summary

In this chapter we build upon previous works dealing with explicit forms of differential pricing [6, 7] and on the model introduced in Chapter 3. Namely, like in [6, 7] we assume that buyers are members of a social population and that the difference of the prices offered to social neighbors must be suitably bounded. Like in Chapter 3, we model the underlying scenario as a multi-unit market with relaxed fair price discrimination constraints. As already mentioned, we borrow the model of fair price discrimination from Chapter 3 and extend it with the features of buyers preselection and additive slackness in fairness constraints considered in [6, 7].

Our results comprise approximation algorithms and hardness of approximation theorems for all the above-mentioned cases (see Table 4.1). Furthermore, we consider specific topologies that are typical of social networks, i.e., graphs where the node degrees follow a power-law distribution, and we show constant approximations with high probability for the revenue on such a class (see Table 4.2).

As in [6, 7], we assume that buyers belong to a social network, in which there is an arc from buyer i to buyer k if i knows k. Then, in order for an assignment of prices to be perceived as fair by i, the price per item proposed to i cannot be higher than the one offered to k, plus a suitable slackness additive factor α_{ik}.

This framework, differently from [6, 7], is able to incorporate also the classical envy-free markets setting. In particular, assuming general valuations, it can express the fact that each buyer might consider multiple options, that is, bundles of different sizes. Moreover, it can model the case of limited supply, meaning that no more than m items can be allocated to buyers. [6, 7] in fact considered only single-minded buyers, that is, interested in purchasing only a given number of items, with unlimited supply.

Moreover, the framework generalizes the fair price discrimination model of Chapter 3 in two respects. First of all, it allows the additive slackness in the fair pricing constraints, which in Chapter 3 was set to 0 for every pair of buyers. Moreover, as in [6, 7], it allows the assignment of a distinguished $\perp$ price to buyers, representing their exclusion from the market. This corresponds to the relevant feature considered in the literature of preselecting buyers, for instance with targeted advertisement or with the typical mechanism of presales used by e-commerce sites.

We first show the effectiveness of price discrimination and buyers preselection, by quantifying the substantial gain in the achievable social welfare and revenue.

Then, we study four different cases arising by considering:

i. social welfare or seller's revenue maximization;

ii. single-minded or general valuations.

Where single-minded buyers are interested only in purchasing a certain amount of items, while buyers with general valuations are free to express the price that they are willing to pay for each possible bundle size. Our results comprise approximation algorithms and hardness of approximation theorems for all the above-mentioned cases.

More precisely, we prove the following results:

i. *Single-minded, social welfare*: we show the NP-hardness of the problem, together with a fully polynomial-time approximation scheme (FPTAS).

ii. *Single-minded, revenue*: We prove that, under the unique game conjecture, it is hard to approximate the maximum revenue within a factor of $\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$, and provide a $O(\log n)$-approximation.

iii. *General valuations, social welfare*: We prove the strong NP-hardness of the problem and give a 2-approximation.

iv. *General valuations, revenue*: We show that the inapproximability within a factor of $\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$ stands also here and describe a $O(\log n + \log m)$-approximation algorithm.

Furthermore, we consider specific topologies that are typical of social networks, i.e., graphs where the node degrees follow a power-law distribution, and we show the following results for the revenue maximization problem:

i. *Single-minded with unlimited supply*: We give a polynomial-time algorithm which returns a constant approximation of the maximum revenue with probability $1 - n^{-1}$.

ii. *Single-minded*: We give a scheme that, given a parameter l, returns a constant approximation of the maximum revenue with probability $1 - e^{-l}$, $\forall l > 0$. The running time is polynomial in the size of the instance and exponential in the size of the parameter l.

iii. *General valuations*: Also in this case we are able to devise a scheme which satisfies the properties of the above-mentioned one.

	Single-Minded	General Valuations
Social Welfare	NP-hard	NP-hard (strong)
	FPTAS	2
Revenue	$\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$	$\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$
	$O(\log n)$	$O(\log n + \log m)$

TABLE 4.1: Hardness and approximation results.

	Approximation	Probability
Single-Minded with Unlimited Supply	$O(1)$	$1 - n^{-1}$
Single-Minded	$O(1)$	$1 - e^{-l}$, $\forall l > 0$
General Valuations	$O(1)$	$1 - e^{-l}$, $\forall l > 0$

TABLE 4.2: Probabilistic approximation results.

4.2 Preliminaries

We define a multi-unit market as a tuple $(N, M, (v_i)_{i \in N})$, where $N = \{1, \ldots, n\}$ is a set of n buyers, M is a set of m identical items and for every buyer $i \in N$, $v_i = (v_i(1), \ldots, v_i(m))$ is a valuation function (or vector) expressing, for each natural number j, the maximum amount $v_i(j) \in \mathbb{R}$ that buyer i is willing to pay for a subset of items $X \subseteq M$ of size j. We assume $v_i(0) = 0$ and $v_i(j) \geq 0$ for every $i \in N$ and j, $1 \leq j \leq m$. We consider both buyers *single-minded* and with *general valuations.*

We adopt a classical pricing scheme, which is natural in case of identical items, usually referred as *item-pricing.* In such a scheme, the seller assigns a single non-negative price per item $p_i \in \mathbb{R}$ to each buyer i. Thus, buyer i owes $p_i \cdot |X|$ for a bundle of items X and her utility for receiving X is given by $u_i(X, p_i) = v_i(|X|) - p_i \cdot |X|$. We denote by $\overline{p} = (p_1, \ldots, p_n)$ the vector of all the prices assigned to the buyers in the market.

We assume buyers to be individuals of a population and we represent this by means of a directed social graph $G = (N, E)$. Such a graph captures the notion of buyers' awareness of the prices proposed to other buyers, more precisely a buyer i is only aware of the prices that the seller proposes to her neighbors $N(i) = \{k \in N | (i, k) \in E\}$. As in previous models for fair price discrimination, we assume that arcs of G are weighted according to a given *slackness* function α specifying, for each arc $(i, k) \in E$, a slackness factor $\alpha(i, k) \geq 0$. Starting from G, it is possible to define the following concept of fair price discrimination.

Definition 4.1. A price vector $\overline{p}$ is *fair* with respect to the social graph $G = (N, E)$ if $p_i \leq p_k + \alpha(i, k)$ for every $(i, k) \in E$.

We define an *allocation vector* as an n-tuple $\overline{X} = (X_1, \ldots, X_n)$ such that $X_i \subseteq M$ is the set of items sold to buyer i, and we call a pair $(\overline{X}, \overline{p})$ an *outcome.* $(\overline{X}, \overline{p})$ is a *feasible* outcome for market μ if it satisfies the supply constraint $\sum_{i=1}^{n} |X_i| \leq m$. Moreover, a feasible outcome $(\overline{X}, \overline{p})$ is *envy-free* if $X_i \in argmax_{X \subseteq M} u_i(X, p_i)$ 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, $u_i(X_i, p_i) \geq 0$.

We are now ready to define the solutions to our markets, that is, *fair outcomes.*

Definition 4.2. A feasible outcome $(\overline{X}, \overline{p})$ is *fair* under the social graph G if it is envy-free and its price vector is fair with respect to G.

We study the (fair) *pricing problems* of determining fair outcomes that maximize two fundamental metrics:

i. *social welfare*: $sw(\overline{X}, \overline{p}) = \sum_{i=1}^{n} v_i(|X_i|)$

ii. *seller's revenue*: $r(\overline{X}, \overline{p}) = \sum_{i=1}^{n} p_i \cdot |X_i|$

We use the notation $opt_{sw}(\mu, G)$ (resp. $opt_r(\mu, G)$) for the maximum possible social welfare (resp. revenue) achievable by an outcome for μ fair under G, and $opt_{sw}(\mu)$ (resp. $opt_r(\mu)$) for the highest possible one achievable without price discrimination (or analogously by an outcome fair under the complete social graph).

By the individual rationality constraint, for any feasible outcome $(\overline{X}, \overline{p})$, it holds

$$sw(\overline{X}, \overline{p}) \geq r(\overline{X}, \overline{p})$$

so that also $opt_{sw}(\mu) \geq opt_r(\mu)$ and $opt_{sw}(\mu, G) \geq opt_r(\mu, G)$.

As in previous models of fair price discrimination, we consider the additional option of preselecting subsets of buyers admitted to the market. In fact, this feature allows the seller to break transitivity chains of price dependencies in the social graph, considerably increasing the maximum revenue achievable in some cases. Formally, we model this by introducing a distinguished bottom price $\perp$ for buyers to be excluded, yielding a corresponding price vector $\overline{p} \in (\mathbb{R} \cup \{\perp\})^n$. The notion of fair pricing is then extended as follows:

Definition 4.3. A price vector $\overline{p}$ is *fair* with respect to the social graph $G = (N, E)$ and the slackness function α if $p_i \leq p_k + \alpha(i, k)$ for every $(i, k) \in E$ such that $p_i \neq \perp$ and $p_k \neq \perp$.

The social welfare and the seller's revenue metrics are then computed considering only buyers not receiving bottom prices.

The advantage of price-discrimination in terms of achievable social welfare and revenue has already been explored in Chapter 3, with an increase that can reach a multiplicative factor of m. 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 a 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 equal to 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 $v_i(m_i)/m_i$.

We now show similar results witnessing the effectiveness of bottom prices with respect to multi-unit markets where only fair-price discrimination is allowed.

The following preliminary result shows the effectiveness of allowing the exclusion of buyers in terms of achievable performance of fair outcomes.

Proposition 4.4. *Let μ and G be respectively a market and its corresponding social graph. Allowing bottom prices in μ can increase the optimal social welfare and revenue of fair outcomes by a multiplicative factor equal to m, and such a bound is tight.*

Proof. Consider the following instance (μ, G) with two single-minded buyers: buyer 1 with $v_1(1) = 1 + \varepsilon$ and buyer 2 with $v_2(m) = m$; graph G contains only the arc $(1, 2)$. If we do not allow bottom prices, the best possible outcome is the one that sells only one item to buyer 1 at a price $1 + \varepsilon$, obtaining social welfare and revenue $1 + \varepsilon$. On the other hand, if bottom prices are allowed, we can set $p_1 = \perp$ and $p_2 = 1$, allocating m items to buyer 2 and achieving social welfare and revenue m.

In order to prove the this bound is tight, consider an optimal outcome $(\overline{X}_\perp, \overline{p}_\perp)$ for an instance (μ, G) using bottom prices. Let v_{max} be equal to $\max_{i \in N, j \leq m}\{\frac{v_i(j)}{j}\}$. Clearly, $v_{max} \cdot m$ is an upper bound to $sw(\overline{X}_\perp, \overline{p}_\perp)$, and then also to $r(\overline{X}_\perp, \overline{p}_\perp)$. Now, let i_{max} and j_{max} be respectively a buyer and a bundle size such that $v_{i_{max}}(j_{max}) = v_{max}$. Consider the outcome $(\overline{X}, \overline{p})$ where $|X_{i_{max}}| = j_{max}$, $X_i = \emptyset$ for all the other buyers, and $p_1 = \ldots = p_n = v_{max}$. Since the utility of all the buyers are not strictly positive, X is envy free. Moreover, as all the proposed prices are coincident, $\overline{p}$ is fair under G. The claim then follows by observing that $\overline{p}$ does not have any entry equal to $\perp$ and $sw(\overline{X}, \overline{p}) = r(\overline{X}, \overline{p}) = v_{max}$. □

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).

Let us finally stress that in multi-unit markets, while the size of the representation of an instance with general valuations is polynomial in m, as different valuations must be specified for every different bundle size, 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. Thus, in a quite counter-intuitive way, hardness results for single-minded buyers do not directly extend to general valuations, and vice versa approximation bounds for general valuations do not automatically transfer to single-minded instances.

4.3 Single-Minded Valuations

We first provide optimal results for the social welfare.

Theorem 4.5. (SINGLE,SOCIAL)*-pricing is NP-hard, but admits an FPTAS (see also Algorithm 11).*

Proof. In order to get the claimed FPTAS, we transform (μ, G) into an equivalent instance $\mathcal{K} = (O, \overline{w}, \overline{z}, k)$ of KNAPSACK as follows:

- O contains an object o_i for each buyer i with profit $z_i = v_i(m_i)$ and weight $w_i = m_i$;
- the knapsack capacity k is set equal to m.

We now prove that ($\Rightarrow$) for each solution $O^* \subseteq O$ of $\mathcal{K}$ with profit z^* there exists a fair outcome $(\overline{X}, \overline{p})$ with $sw(\overline{X}, \overline{p}) = z^*$, and ($\Leftarrow$) vice versa.

($\Rightarrow$): Given a solution $O^* \subseteq O$ for $\mathcal{K}$ with profit z^*, consider the outcome $(\overline{X}, \overline{p})$ defined as follows:

$$|X_i| = \begin{cases} m_i & \text{if } o_i \in O^* \\ 0 & \text{otherwise} \end{cases}$$

$$p_i = \begin{cases} \min_{o_k \in O^*} v_k(m_k)/m_k & \text{if } o_i \in O^* \\ \perp & \text{otherwise} \end{cases}$$

By construction, $\sum_{i \in N} |X_i| = \sum_{o_i \in O^*} w_i \leq k = m$ and $\sum_{i \in N} v_i(X_i) = \sum_{o_i \in O^*} z_i$, that is, $(\overline{X}, \overline{p})$ satisfies the supply constrains and has social welfare $sw(\overline{X}, \overline{p}) = z^*$. Furthermore, since each buyer with price different from $\perp$ gets her preferred bundle, $(\overline{X}, \overline{p})$ is also envy-free. Finally, all buyers are either proposed price $\perp$ or the unique price $\min_{o_k \in O^*} v_k(m_k)/m_k$, and thus $\overline{p}$ is fair under any social graph.

($\Leftarrow$): Given a fair outcome $(\overline{X}, \overline{p})$ with $sw(\overline{X}, \overline{p}) = z^*$, consider the solution $O^* \subseteq O$ for $\mathcal{K}$, where $o_i \in O^*$ if $|X_i| = m_i$. Then O^* is feasible as $\sum_{o_i \in O^*} w_i = \sum_{i \in N} |X_i| \leq m = k$ and has total profit equal to $\sum_{o_i \in O^*} z_i = \sum_{i \in N \, s.t. \, |X_i| = m_i} v_i(m_i) = sw(\overline{X}, \overline{p})$.

Having shown the above equivalence, it is possible to provide the claimed FPTAS for (SINGLE,SOCIAL)-pricing as follows: we transform in polynomial time the input instance (μ, G) into the equivalent one of KNAPSACK, we run the well known FPTAS for such a problem, and finally transform in polynomial time the output to an outcome of the initial problem as described in $\Rightarrow$) to obtain a $(1 + \varepsilon)$-approximation of the optimal social welfare, hence the claim. □

Algorithm 11: FPTAS for the (SINGLE,SOCIAL)-pricing problem

Input: An instance (μ, G) of the (SINGLE,SOCIAL)-pricing problem, accuracy parameter ε.
Output: A fair outcome $(\overline{X}, \overline{p})$.
Construct the following instance $K(\mu) = (O, \overline{z}, \overline{w}, k)$ of KNAPSACK:
  $O = \{o_1, \ldots, o_n\}$ for each buyer $i \in N$;
  $z_i = v_i(m_i)$ for $i \in N$;
  $w_i = m_i$ for $i \in N$;
  $k = m$;
$O^* \longleftarrow$ FPTAS-KNAPSACK;
$|X_i| = m_i$ if $o_i \in O^*$, 0 otherwise;
$p_i = \min_{o_k \in O^*} v_k(m_k)/m_k$ if $o_i \in O^*$, $\perp$ otherwise;
return $(\overline{X}, \overline{p})$;

Regarding revenue maximization, an approximation algorithm can be obtained, with the same approach used in Chapter 3, for the same case without bottom prices.

Corollary 4.6. (SINGLE,REVENUE)*-pricing admits a* $(\frac{\log n}{1-\varepsilon})$*-approximation algorithm (see Algorithm 12).*

Proof. Consider the output $(\overline{X}, \overline{p})$ of the algorithm described in Theorem 4.5, and let $N_{\overline{X}} \subseteq N$ be the set of buyers allocated in $(\overline{X}, \overline{p})$. Without loss of generality assume that buyers in $N_{\overline{X}}$ are ordered non-increasingly with respect to the ratios $v_i(m_i)/m_i$, and let h be the index maximizing $\frac{v_h(m_h)}{m_h} \sum_{i=1}^{h} m_i$. Consider the price vector $\overline{p}'$ with $p'_i = \frac{v_h(m_h)}{m_h}$ for all $i \in N_{\overline{X}}$ such that $i \leq h$, and $p'_i = \perp$ otherwise. Let $(\overline{X}', \overline{p}')$ be the outcome allocating bundles only to buyers with price different from $\perp$. Then, $(\overline{X}', \overline{p}')$ is a fair outcome for (μ, G) of revenue $\frac{v_h(m_h)}{m_h} \sum_{i=1}^{h-1} |X_i| \geq \frac{\sum_{i=1}^{k-1} v_i(m_i)}{\log n}$. Furthermore $\overline{p}$ is fair under G, since the same price is proposed to all buyers i such that $p'_i \neq \perp$. ☐

Algorithm 12: A $(\frac{\log n}{1-\varepsilon})$-approximation algorithm for (SINGLE,REVENUE)-pricing.

Input: An instance (μ, G) of the (SINGLE,REVENUE)-pricing problem.
Output: A fair outcome $(\overline{X'}, \overline{p'})$.
$(\overline{X}, \overline{p}) \longleftarrow$ outcome from Algorithm 11 on input (μ, G);
Let $N_{\overline{X}} \subseteq N$ be the set of buyers allocated in $(\overline{X}, \overline{p})$, non-increasingly ordered with respect to the ratios $v_i(m_i)/m_i$;
Let $h = argmax \frac{v_h(m_h)}{m_h} \sum_{i=1}^{h} m_i$;
$|X_i| = m_i$ for all $i \in N_{\overline{X}}$ such that $i \leq h$, 0 otherwise;
$p'_i = \frac{v_h(m_h)}{m_h}$ for all $i \in N_{\overline{X}}$ such that $i \leq h$, $\perp$ otherwise;
return $(\overline{X'}, \overline{p'})$;

The approximation factor above presented is nearly optimal, as we are also able to prove a strong negative result, but first we need to introduce an important conjecture.

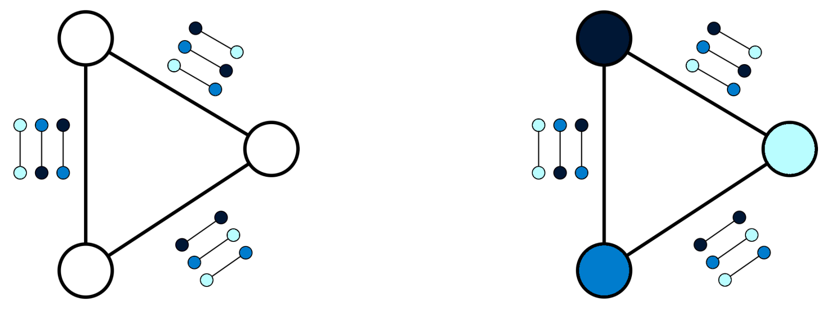

(a) (b)

FIGURE 4.1: In Figure 4.1(a) we depicted an instance of UNIQUE LABEL COVER that admits a solution, as shown in Figure 4.1(b).

Unique game conjecture A formulation of the conjecture is based on the problem UNIQUE LABEL COVER, which is in turn strictly connected to the classical problem of coloring the vertices of a graph. Formally, given in input a graph $G = (V, E)$ and k colors, the objective is to assign colors to the vertices of such a graph. Differently from the classical definition, which requires adjacent vertices to have different colors, the constraints are expressed by the mean of a family of permutations $\pi_e : [k] \to [k]$ for $e \in E$, meaning that if e is the edge $\{u, v\}$ and u has color i then v is required to get color $\pi_e(i)$.

It is easy to see that if the constraints admit a coloring for the graph, it is possible to find such a coloring by checking all the available colors for a fixed vertex and deriving the full coloring of the graph (see Figure 4.1), but even if such a coloring does not exist, we can try to maximize the number of constraints satisfied (see Figure 4.2).

Given this brief introduction we are finally able to state the conjecture. In fact it states that for every value of $\varepsilon > 0$, no matter how small, there is a number of colors k (possibly very large) for which, looking at graphs with constraints on k colors, it is NP-hard to tell the difference between those for which at least $(1 - \varepsilon)$ fraction of the constraints can be satisfied and those for which at most an ε fraction of the constraints can be satisfied.

The above conjecture has been first proposed in [75] and since then has given birth to an important body of hardness-of-approximation results in the literature [51, 69, 73, 76, 77].

Theorem 4.7. *It is Unique-Game-hard to approximate* (SINGLE,REVENUE)*-pricing within a factor of* $o(\frac{\sqrt{\log n}}{\log^2 \log n})$

In order to prove our claim, we resort on known hardness results on the INDEPENDENT SET problem on graphs with maximum degree bounded by δ. Such a problem has be shown to be Unique-Game-hard to approximate within a factor of $o(\frac{\delta}{\log^2 \delta})$ [11]. We first exploit the bound on the node degrees of the input graph in order to find a *good partition* of the nodes. More precisely:

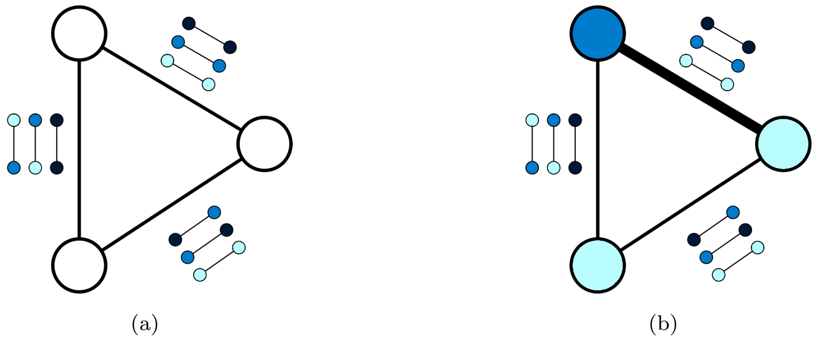

FIGURE 4.2: In Figure 4.2(a) we depicted an instance of UNIQUE LABEL COVER that does not admit a solution but for which it is possible to satisfy $\frac{2}{3}$ of the constraints as shown in Figure 4.2(b).

Definition 4.8. Let $H = (V, F)$ be a graph, and let $\mathcal{S} = \{S_1, \dots, S_\kappa\}$ be a partition of V. We say that $\mathcal{S}$ is a *good partition* for G if:

i. (*coloring*) $\forall S_i \in \mathcal{S}$, $\forall u, v \in S_i$, $(u, v) \notin F$;

ii. each node has at most one neighbor in each subset $\mathcal{S}$.

When the maximum degree of H is bounded by δ, the following result holds:

Lemma 4.9. *Any graph $H = (V, F)$ with node degrees at most δ admits a good partition $\mathcal{S}$ with $|\mathcal{S}| \leq \delta^2 + 1$, and such a partition can be found in polynomial time.*

Proof. Consider the graph $H^2 = (V, F^2)$, where:

$$F^2 = F \cup \{\{u, v\} \in V^2 | k \in V, \{u, k\}, \{k, v\} \in E\}.$$

By construction H^2 has degree at most δ^2. By the Brooks' Theorem H^2 can be colored using at most $\delta^2 + 1$ colors, and such a coloring can be found in polynomial time. Let $\mathcal{S}$ be the partition of V induced by such a coloring. Since $F \subseteq F^2$, property ***i.*** holds. Suppose then by contradiction that property ***ii.*** does not hold for $\mathcal{S}$, that is, there exists $w \in V$ that has two neighbors v, u belonging to $S_i \in \mathcal{S}$. Since v and u share a neighbor in H, $\{v, u\} \in F^2$, but this implies that $\mathcal{S}$ is not induced by a coloring of H^2: a contradiction. □

We are ready to prove Theorem 4.7.

Proof. Consider the following reduction from INDEPENDENT SET instances $H = (V, F)$ with bounded degree δ to instances (μ, G) of (SINGLE,REVENUE)-pricing:

i. let $\mathcal{S} = \{S_1, \dots S_D\}$ be a good partition of G, with $D \leq \delta^2 + 1$;

ii. for each node $u \in S_d$ add a set $\mathcal{N}_u$ of 2^d single-minded buyers with valuation 2^{-d} only for bundles of size 1;

iii. for each $v \in V$ and pair of buyers $i, h \in \mathcal{N}_v$ add in the social graph G arcs (i, h), (h, i) with $\alpha(i, h) = \alpha(h, j) = 0$;

iv. for each edge $\{u, v\} \in F$, with $u \in S_d$, $v \in S_{d'}$ and $d < d'$, add in G all arcs (i, h) with $i \in \mathcal{N}_u$ and $h \in \mathcal{N}_v$;

v. consider unlimited supply (or equivalently set it to $|V| \cdot 2^D$).

We are going to prove our claim by showing that ($\Rightarrow$) if G admits an independent set of cardinality k, then the reduced instance admits revenue at least k; and ($\Leftarrow$) if the reduced instance admits revenue k then G has an independent set of cardinality at least $\frac{k}{2}$.

($\Rightarrow$) Let $I \subseteq V$ be an independent set of G with cardinality k. Consider the outcome $(\overline{X}, \overline{p})$ for the reduced instance in which $|X_i| = 1$ and $p_i = 2^{-d}$ if $i \in \mathcal{N}_v$, $v \in S_d$ and $v \in I$, otherwise $|X_i| =$ and $p_i = 2^{-d}$.

Notice that $\overline{p}$ is fair under G. In fact, by construction, denoted as $\Delta(v)$ the set of the neighbors of node v in H, the set of the neighbors of buyer $i \in \mathcal{N}_v$ in G is a subset of $\mathcal{N}_v \cup \bigcup_{u \in \Delta(v)} \mathcal{N}_u$. Then, if $v \notin I$, $p_i = \perp$ and no fairness constraints on p_i must hold. If instead $v \in I$, all neighbors of i in $\mathcal{N}_v$ get the same price, and, since I is independent, $p_k = \perp$ for all buyers k in $\bigcup_{u \in \Delta(v)} \mathcal{N}_u$. Therefore, $\overline{p}$ is fair. Furthermore, $(\overline{X}, \overline{p})$ is envy-free, since a price equal to their valuation is proposed to all buyers receiving a bundle of cardinality 1, while all the ones not receiving any item get price $\perp$. Finally observe that for each $v \in I$ with $v \in S_d$, there are 2^d buyers buying at price 2^{-d}, ensuring revenue 1 for each node in I. Therefore, $r(\overline{X}, \overline{p}) = k$.

($\Leftarrow$) Assume that the reduced instance admits an outcome $(\overline{X}, \overline{p})$ with revenue k. Since we are under the hypothesis of unlimited supply, without loss of generality we can assume that $|X_i| = 1$ for each buyer i such that $p_i \leq v_i(1)$, since this can only increase the revenue. Similarly, as $\overline{p}$ is fair, if a price $p_i \neq \perp$ is proposed to a buyer $i \in \mathcal{N}_v$, then the price proposed to all the other buyers in $\mathcal{N}_v$ must be either p_i or $\perp$. Thus, we can assume that $p_k = p_i$ and a bundle is assigned to all the buyers in $k \in \mathcal{N}_v$. Under these assumptions $(\overline{X}, \overline{p})$ can be described by means of a vector $\overline{\pi} \in (\mathbb{R} \cup \{\perp\})^{|V|}$, where component π_v is equal to the price proposed to all buyers in $\mathcal{N}_v$.

After these preliminary remarks, consider the subset of nodes I built as follows:

i. consider all the subsets $S_d \in \mathcal{S}$ in an inverse order with respect to their index d;

ii. for each $v \in S_d$ such that $\pi_v \neq \perp$, add v to I, set $\pi_v = \perp$ and set $\pi_u = \perp$ for all $u \in \Delta(v)$.

Since each time that we add a node to I we set $\pi_u = \perp$ for all its neighbors, it is not possible to add to I two adjacent nodes in G, and thus I is independent. We now prove that $|I| \geq \frac{k}{2}$. Let $\rho = \sum_{v \in V} \pi_v |\mathcal{N}_v|$, considering π_v as 0 if $\pi_v = \perp$. Clearly, before starting running the building procedure described above, $\rho = k$, as it coincides with the revenue of outcome $(\overline{X}, \overline{p})$. Moreover, after I is built, $\rho = 0$, since all $\pi_v = \perp$. Furthermore, after adding node $v \in S_d$ to I, ρ decreases by exactly $\sum_{u \in \{v\} \cup \Delta(v)} \pi_u |\mathcal{N}_u|$. By the fairness constraints and since we are considering subset of nodes in a decreasing order, in this step $\pi_u \leq \pi_v$ for all $u \in \Delta(v)$, and as $v_i(1) = 2^{-d}$ for all the buyers $i \in \mathcal{N}_v$, we have that:

$$\sum_{u \in \{v\} \cup \Delta(v)} \pi_u |\mathcal{N}_u| \leq 2^{-d} \sum_{u \in \{v\} \cup \Delta(v)} |\mathcal{N}_u| \leq 2^{-d} \sum_{d'=1}^{d} 2^{d'} \leq 2,$$

where the second inequality derives from the fact that a node v can't have more then one neighbor in the same subset S_d. We then have that ρ decreases by at most 2 each time a node is added to I, and thus $|I| \geq \frac{k}{2}$.

To conclude the proof, we remark that in order to have a polynomial number of buyers and consequently a polynomial time reduction it is sufficient to choose $\delta = \sqrt{\log n}$. □

4.4 General Valuations

In order to achieve good approximations for general valuations, we resort on the following reduction to the single-minded case provided in Chapter 3. Given a buyer $i \in N$, let $S_i = \{m_i^1, ..., m_i^\ell\}$ be the bundle sizes that are in the demand set of i for at least one positive price, listed in non-decreasing order. Let $m_{i_1} = m_i^1$ and $m_{i_j} = m_i^j - m_i^{j-1}$ for $2 \leq j \leq \ell$. The reduction transforms buyer i into ℓ single-minded *marginal buyers* $i_1, \ldots, i_\ell$, where i_j has preferred bundle size m_{i_j} and valuation $v_{i_j}(m_{i_j}) = v_i(m_i^j) - v_i(m_i^{j-1})$. The reduced social graph $G' = (N', E')$ is such that $(i_j, i'_{j'}) \in E'$ if and only if $(i, i') \in E$.

The authors have shown that the ratios $\frac{v_{i_j}(m_{i_j})}{m_{i_j}}$ are non-increasing in j. Moreover, if an approximation algorithm for single-minded buyers applied on a reduced instance allocates bundles only to prefixes of marginal buyers, then its solution can be transformed back into an outcome for the initial problem preserving the same approximation ratio.

Unfortunately, the FPTAS given in the previous section for single-minded instances does not have such a property. Hence, we devise an ad-hoc procedure that 2-approximates the social welfare, while allocating prefixes of marginal buyers.

Theorem 4.10. (GENERAL,SOCIAL)*-pricing admits a* 2*-approximation algorithm (see Algorithm 13).*

Proof. Let us first recap the reduction provided in Chapter 3. Given a buyer $i \in N$, let $S_i = \{m_i^1, ..., m_i^\ell\}$ be the bundle sizes that are in the demand set of i for at least one positive price, listed in non-decreasing order. Let $m_{i_1} = m_i^1$ and $m_{i_j} = m_i^j - m_i^{j-1}$ for $2 \leq j \leq \ell$. The reduction transforms buyer i into ℓ single-minded *marginal buyers* $i_1, \ldots, i_\ell$, where i_j has preferred bundle size m_{i_j} and valuation $v_{i_j}(m_{i_j}) = v_i(m_i^j) - v_i(m_i^{j-1})$. The reduced social graph $G' = (N', E')$ is such that $(i_j, i'_{j'}) \in E'$ if and only if $(i, i') \in E$.

The authors have shown that the ratios $\frac{v_{i_j}(m_{i_j})}{m_{i_j}}$ are non-increasing in j. Moreover, if an approximation algorithm for single-minded buyers applied on a reduced instance allocates bundles only to prefixes of marginal buyers, then its solution can be transformed back into an outcome for the initial problem preserving the same approximation ratio.

Unfortunately, the FPTAS given in the previous section for single-minded instances does not have such a property. Hence, we devise an ad-hoc procedure that 2-approximates the social welfare, while allocating prefixes of marginal buyers.

Given an instance (μ, G) of (GENERAL,SOCIAL)-pricing, let (μ', G') be the associated output of the reduction of Chapter 3. Consider the following algorithm for the maximization of $opt_{sw}(\mu', G)$:

- Sort all marginal buyers by the ratios $\frac{v_{i_j}(m_{i_j})}{m_{i_j}}$ in non-increasing order, where in case of ties the marginal buyers i_j of a same buyer i are listed in order of j; let $\pi(i_j)$ be the position in the order of each i_j.
- Compute the following two outcomes:
  - $(\overline{X}', \overline{p}')$: Let i'_h be the last marginal buyer in the order such that

$$\sum_{i_j | \pi(i_j) \leq \pi(i'_h)} m_{i_j} \leq m$$

    . Set $|X'_{i_j}| = m_{i_j}$ and $p'_{i_j} = \frac{v_{i'_h}(m_{i'_h})}{m_{i'_h}}$ if $\pi(i_j) \leq \pi(i'_h)$, and $|X'_{i_j}| = 0$ and $p_{i_j} = \perp$ otherwise;
  - $(\overline{X}'', \overline{p}'')$: Let i''_l be the marginal buyer following i'_h in the order, that is, such that $\pi(i''_l) = \pi(i'_h) + 1$ (if not existing all the marginal buyers are allocated

in $(\overline{X}', \overline{p}')$, that is in turn an optimal solution). For all i_j with $j \leq l$, set $|X''_{i_j}| = m_{i_j}$ and $p''_{i_j} = 0$, while set price equal to $\perp$ and give no items to all the other buyers.

- Return argmax$\{sw(\overline{X}', \overline{p}'), sw(\overline{X}'', \overline{p}'')\}$

Notice that both $(\overline{X}', \overline{p}')$ and $(\overline{X}'', \overline{p}'')$ are fair under G', as in both a unique price different from bottom is proposed. Furthermore, since for each buyer i_j with a non bottom price $p_{i_j} \leq \frac{v_{i_j}(m_{i_j})}{m_{i_j}}$ and $|X'_{i_j}| = |X''_{i_j}| = m_{i_j}$, both $\overline{X}'$ and $\overline{X}''$ are also envy-free.

Notice also that $sw(\overline{X}', \overline{p}') + sw(\overline{X}'', \overline{p}'')$ is an upper bound on $opt_{sw}(\mu', G')$, as the union of their allocated buyers corresponds to an optimal outcome for supply bigger than m. Thus choosing the best of the two solutions ensures approximation ratio equal to 2. The claim then follows by observing that both $(\overline{X}', \overline{p}')$ and $(\overline{X}'', \overline{p}'')$ allocate only prefixes of marginal buyers, hence by the properties of the reduction they can be turned back into corresponding outcomes for (μ, G) with the same approximation ratio. ☐

Algorithm 13: A 2-approximation algorithm for (GENERAL,SOCIAL)-pricing

---

Input: An instance (μ, G) of the (GENERAL,SOCIAL)-pricing problem.
Output: A fair outcome $(\overline{X}, \overline{p})$.
Sort all marginal buyers by the ratios $\frac{v_{i_j}(m_{i_j})}{m_{i_j}}$ in non-increasing order, where in case of ties the marginal buyers i_j of a same buyer i are listed in order of j;
Let $\pi(i_j)$ be the position in the order of each i_j;
Let i'_h be the last marginal buyer in the order such that $\sum_{i_j | \pi(i_j) \leq \pi(i'_h)} m_{i_j} \leq m$;
$|X'_{i_j}| = m_{i_j}$ if $\pi(i_j) \leq \pi(i'_h)$, $|X'_{i_j}| = 0$ otherwise;
$p'_{i'_j} = \frac{v_{i'_h}(m_{i'_h})}{m_{i'_h}}$ if $\pi(i'_j) \leq \pi(i'_h)$, $p_{i_j} = \perp$ otherwise;
Let i''_l be the marginal buyer following i'_h in the order, that is, such that $\pi(i''_l) = \pi(i'_h) + 1$ (if not existing all the marginal buyers are allocated in $(\overline{X}', \overline{p}')$, that is in turn an optimal solution);
$|X''_{i_j}| = m_{i_j}$ for all i_j with $j \leq l$, 0 otherwise;
$p''_{i_j} = 0$ for all i_j with $j \leq l$, $\perp$ otherwise;
return $(\overline{X}, \overline{p}) = \text{argmax}\{sw(\overline{X}', \overline{p}'), sw(\overline{X}'', \overline{p}'')\}$;

Considering general valuations worsens the complexity of finding an outcome that maximizes social welfare.

Theorem 4.11. (GENERAL,SOCIAL)*-pricing is strongly NP-hard.*

Proof. In a similar way as in [57], in order to prove this result we provide a polynomial-time reduction from DENSEST K-SUBGRAPH. In such a problem, 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. Let $\varphi = (|F|+1)(|V|+1)$, $\psi = (|F|+1)$, and $\varepsilon = \frac{1}{k\psi+1}$. Given an instance $(H = (V, F), k)$ of DENSEST K-SUBGRAPH, consider the reduced instance (μ, G) built as follows:

- To each node $\mathbf{v} \in V$ associate one buyer $i_{\mathbf{v}}$ with valuation function

$$v_{i_{\mathbf{v}}}(j) = \begin{cases} 2\varphi & \text{if } j = \varphi \\ 2\varphi + (1+\varepsilon)\psi & \text{if } j = \varphi + \psi \\ 0 & \text{otherwise} \end{cases}$$

- To each node $\mathbf{e} \in F$ associate one buyer $i_{\mathbf{e}}$ with valuation function

$$v_{i_{\mathbf{e}}}(j) = \begin{cases} 2\varphi & \text{if } j = \varphi \\ 2\varphi + 1 & \text{if } j = \varphi + 1 \\ 0 & \text{otherwise} \end{cases}$$

- Set $m = (|V| + |F|)\varphi + k\psi + |F|$.
- The social graph $G = (I, E)$ is built as follows: if $\mathbf{e} = \{\mathbf{u}, \mathbf{v}\} \in F$, then $(i_{\mathbf{v}}, i_{\mathbf{e}})$, $(i_{\mathbf{u}}, i_{\mathbf{e}}) \in E$, setting $\alpha_{\mathbf{ue}} = \alpha_{\mathbf{ve}} = 0$.

Let us analyze the contribution of each buyer to the social welfare. Clearly, if a price is not proposed to a generic buyer, her contribution is 0. If instead a price $p_i \leq 2$ is proposed, buyer i contributes at least 2φ, having a non-negative utility for a bundle of size at least φ. For a generic buyer $i_{\mathbf{u}}$, if $p_{i_{\mathbf{u}}} \leq 1+\varepsilon$, it possible to allocate to her ψ more items, increasing in this way her contribution to the social welfare of exactly $(1+\varepsilon)\psi$. For a generic buyer $i_{\mathbf{e}}$ instead, if $p_{i_{\mathbf{e}}}$ is dropped to 1, it is possible to allocate to her 1 more item, increasing her contribution to the social welfare of exactly 1.

Since $\varphi > (1+\varepsilon)(|V|)\psi + |F|$, an outcome that maximizes the social welfare must propose a price $p_i \leq 2$ to all the buyers, allocating at least φ items to all of them. Since $\psi > |F|$, with the remaining supply $kD + |F|$, we can possibly allocate ψ more items to at most k buyers $i_{\mathbf{v}}$, and one more item to each of the $|F|$ buyers $i_{\mathbf{e}}$. However, by the construction of G, setting $p_{i_{\mathbf{e}}} = 1$ for some $\mathbf{e} = (\mathbf{u}, \mathbf{v})$ forces us to propose price 1 also to buyers $i_{\mathbf{u}}, i_{\mathbf{v}}$. Then, in order for such buyers not to be envious, we must allocate ψ more items to each of them. We conclude by observing that the solution that maximizes the social welfare in (μ, G) allocates φ items to all buyers, ψ more items to each buyer $i_{\mathbf{u}}$ belonging to subset S, with $|S| \leq k$, and one more item to all the buyers $i_{(\mathbf{u},(\mathbf{v}))}$ such that $i_{\mathbf{u}}, i_{\mathbf{v}} \in S$. The social-welfare of an optimal solution is then equal to $2(|V| + |F|)\varphi + (1+\varepsilon)k\psi + h$, where h is the cardinality of the DENSEST K-SUBGRAPH of H. Finding an optimal S is thus equivalent to finding a DENSEST K-SUBGRAPH in (V, F). □

In the same fashion of Corollary 4.6, it is possible to exploit the 2-approximation for the social welfare in order to obtain the following result.

Corollary 4.12. (GENERAL,REVENUE)*-pricing admits a* $2(\log n + \log m)$*-approximation algorithm.*

The hardness result provided in Theorem 4.7 directly extends to general valuations, as in the provided reduction m is polynomially bounded in the size of the instance.

Corollary 4.13. (GENERAL,REVENUE)*-pricing is Unique-Game-hard to approximate within a factor of* $o(\frac{\sqrt{\log n}}{\log^2 \log n})$.

4.5 Social Networks

We now focus on graph topologies that are typical of social networks. Namely, we assume that node degrees in G respect a power law distribution. This class of graphs, also called scale-free, has been largely investigated in the literature as the paradigmatic model of the web graph and other common graphs arising from social relationships. While in the previous sections good approximations bound have been already obtained for the social welfare without any restriction on the structure of the network, we here provide better results for the revenue maximization.

Let $\overline{d} = (d_1, d_2, \ldots, d_n)$ be a non-decreasing sequence or vector of n strictly positive integers, whose sum is even. We assume that $\overline{d}$ respects a power law distribution. Namely, for any fixed integer $k > 0$, the number $n(k)$ of integers d_i with $d_i = k$ is proportional to $k^{-\gamma}$, where typically $2 < \gamma < 3$. In other words, $c \cdot n \cdot k^{-\gamma} \leq n(k) \leq c' \cdot n \cdot k^{-\gamma}$, for three given constants c, c' and γ such that $c < c'$. As it can be easily checked, the number of integers d_i with $d_i > k$ in $\overline{d}$ can be suitably upper bounded as $\sum_{h=k+1}^{n} n(h) = O\left(\frac{n}{k^{\gamma-1}}\right)$.

Let $\mathcal{G}_{n,\mathbf{d}}$ be the class of graphs with node set $N = \{1, 2, \ldots, n\}$, in which the sequence of node degrees listed in non-decreasing order coincides with $\overline{d}$. We assume that the social graph G is randomly drawn in $\mathcal{G}_{n,\mathbf{d}}$ uniformly selecting a permutation of buyers π in such a way that buyer i is associated to position $\pi(i)$ of the degree sequence, with corresponding degree $d_{\pi(i)}$.

Let us first focus on the single-minded case. Before providing nice approximations for power-law graphs, let us give the following key lemma, which will be useful in the sequel.

Lemma 4.14. *Given any family of graphs $\mathcal{G}$ and a fixed integer $k > 0$, if a k-coloring for any graph in $\mathcal{G}$ exists and can be determined in polynomial time, then* (SINGLE,REVENUE)*-pricing restricted to social graphs in $\mathcal{G}$ admits a $(k+\varepsilon)$-approximation algorithm, for any $\varepsilon > 0$ (see Algorithm 14).*

Proof. Once colored the nodes of the graph, consider the subset of buyers N_i with a fixed color i. Since N_i forms an independent set, a $(1 + \varepsilon/k)$-approximation for the submarket containing only the buyers in N_i can be easily determined by completely ignoring the fair price discrimination constraints and running the FPTAS of knapsack on the equivalent knapsack instance with capacity m containing an object o_i for every buyer $i \in N$ with profit $z_i = v(m_i)$ and weight m_i. In fact, the returned solution can be directly translated to an outcome of the original problem with the same revenue, by assigning a preferred bundle of size m_i at price $v(m_i)/m_i$ per item to every buyer i corresponding to a selected object, and discarding the remaining buyers by means of bottom prices.

Starting from the above-collected outcomes, a $(k+\varepsilon)$-approximation can be determined simply by returning the best of them, say associated to a given color i, completed by assigning bottom prices to all the buyers not in N_i.

In fact, at least one set N_i contributes $r \geq opt_r(\mu, G)/k$ to the optimal revenue of a fair outcome for the initial instance (μ, G), and the optimal solution for the submarket restricted to N_i has revenue at least r. □

Algorithm 14: A $(k+\varepsilon)$-approximation algorithm for (SINGLE,REVENUE)-pricing problem, $\varepsilon > 0$, for a special family of graphs $\mathcal{G}$.

Input: A market μ and a graph G belonging to a family of graphs $\mathcal{G}$ such that there exists $k > 0$ for which a k-coloring for any graph in $\mathcal{G}$ exists and can be determined in polynomial time.
Output: A fair outcome $(\overline{X}, \overline{p})$.
k-color the graph G;
Let N_i be the subset of buyers with a fixed color i;
Construct the following KNAPSACK instance $K(N_i) = (O, \overline{z}, \overline{w}, k)$:
  $O = \{o_i, i \in N_i\}$;
  $z_i = v_i(m_i)$ for $i \in N$;
  $w_i = m_i$ for $i \in N$;
  $k = m$;
Let $(\overline{X}^i, \overline{p}^i)$ be the outcome which assigns to all the buyers in the KNAPSACK solution their preferred bundle at their valuation, while giving $\perp$ prices to all the remaining buyers;
return $(\overline{X}, \overline{p}) = \text{argmax } r(\overline{X}^i, \overline{p}^i)$;

As a direct consequence of the above lemma, constant approximation algorithms can be obtained for graphs with maximum degree bounded by a constant (thanks to the well-known greedy coloring algorithm), for planar graphs, bipartite graphs and for many other classes of graphs.

Unfortunately, power law graphs do not have a constant bounded degree, thus not allowing a direct application of the above lemma. However, their average degree is constant. More precisely, for any choice of the constant parameters c, c' and γ, there exists a "small" constant integer k such that the number of nodes with degree greater than k is at most $n/2$.

Starting from the above observation, let us consider the following algorithm, called POWER-LAW (see also Algorithm 15): once drawn $G \in \mathcal{G}_{n,\mathbf{d}}$ according to the above random process, consider the subset $N' \subseteq N$ of buyers of degree at most k, and then run the above algorithm for bounded degree graphs on the instance (μ, G'), where G' is the subgraph induced by N'.

Lemma 4.15. *POWER-LAW executed on randomly drawn social graphs in $\mathcal{G}_{n,\mathbf{d}}$ has constant expected approximation ratio for* (SINGLE,REVENUE)*-pricing.*

Proof. We prove that the expected revenue of the above algorithm is $\Omega(opt_r(\mu, G))$.

Let k be the constant selected by the algorithm, i.e. such that the set N' of the buyers of degree at most k has cardinality $|N'| \geq n/2$. Let X_i be the random variable equal to 1 if buyer i has degree at most k in G, $X_i = 0$ otherwise, and let S be the random variable corresponding to the sum of the preferred valuations of the buyers of degree at most k in G, that is, $S = \sum_{i \in N} v_i(m_i)X_i$. Then, since POWER-LAW exploiting a $(k+1)$-coloring returns a solution of revenue at least $S/(k+1+\varepsilon)$ and k is constant, it is sufficient to asymptotically bound the expected value $E(S)$ of S.

To this aim, by the linearity of expectation, we have that $E(S) = E(\sum_{i \in N} v(m_i)X_i) = \sum_{i \in N} v(m_i) \cdot E(X_i) = \sum_{i \in N} v(m_i) \cdot Prob(X_i = 1) \geq \sum_{i \in N} v(m_i)/2 \geq opt_r(\mu, G)/2$, thus proving the claim.

The proof for the case of limited supply holds just repeating the above argument restricting on the subset of the buyers that are allocated in an optimal fair outcome for μ and G. □

Ideally, we would like to prove that the outcome returned by POWER-LAW has constant approximation not only in expectation, but also with high probability. Unfortunately, this is not guaranteed in general, as it can be easily checked in case a single buyer has a very high valuation for her preferred bundle, while all the others have negligible valuations. In this case, the probability that the returned solution has a constant approximation can be bounded only by 1/2.

However, in case of unlimited supply, it is possible to obtain a bound with high probability by preprocessing the buyers with the highest valuations, so as to reduce the variance of the random variable S, when restricted to the remaining buyers with lower valuations. Namely, consider the following algorithm: once drawn $G \in \mathcal{G}_{n,\mathbf{d}}$, order the buyers non-increasingly with respect to their preferred valuations, and let P be the prefix of the first $l = 8 \ln n$ buyers; determine the optimal solution for the submarket containing only buyers in P and their induced subgraph G_P, and complete it by assigning bottom prices to all the remaining buyers; let $(X_1, \overline{p1})$ be the resulting outcome; run POWER-LAW and let $(X_2, \overline{p2})$ be the corresponding outcome; return the best of the two outcomes.

Notice that $(X_1, \overline{p1})$ can be easily computed in polynomial time. In fact, given the subset P^* of P of the buyers allocated in an optimal outcome for P, the prices yielding the maximum revenue for P^* can be determined as follows. Order the buyers non-decreasingly with respect to the ratios $v_i(m_i)/m_i$. For each buyer i considered in such an order, set p_i to be the maximum possible value such that $p_i \leq v_i(m_i)/m_i$ and $p_i \leq p_k + \alpha(i,k)$ for every buyer $k \in P^*$ with $k < i$ and $(i,k) \in E$. In other words, p_i is set to the maximum possible value compatible with the individual rationality of i and the fairness constraints for the pricing. Thus, $(X_1, \overline{p1})$ can be computed in such a way by probing all the possible subsets of P, whose number is polynomially bounded.

We are now able to prove the following theorem.

Algorithm 15: POWER-LAW procedure.

Input: A market μ and a graph $G \in \mathcal{G}_{n,\mathbf{d}}$ drawn according to the above random process.
Output: A fair outcome $(\overline{X}, \overline{p})$.
Let k be such that the number of nodes with degree greater than k is at most $n/2$;
Let $N' \subseteq N$ be the set of buyers of degree at most k;
Run Algorithm 14 for bounded degree graphs on the instance (μ, G'), where G' is the subgraph induced by N';
return $(\overline{X}, \overline{p})$ given by Algorithm 14;

Theorem 4.16. *The above algorithm, run on randomly drawn social graphs in $\mathcal{G}_{n,\mathbf{d}}$, in case of unlimited supply returns constant approximation ratio for* (SINGLE,REVENUE)*-pricing with probability at least $1 - 1/n$.*

Proof. Again we show that the revenue of the algorithm is $\Omega(opt_r(\mu, G))$ with probability $1 - 1/n$.

If the contribute to $opt_r(\mu, G)$ due to the prefix P of the first $l = 8 \ln n$ buyers is higher with respect to the one of the remaining ones, then the algorithm returns a solution of revenue at least $opt_r(\mu, G)/2$.

On the other hand, if such a contribution is lower, consider the subset $N \setminus P$ of the remaining buyers not in the prefix. Let k be the constant selected by the algorithm, i.e. such that the set $N' \subseteq N \setminus P$ of the buyers of degree at most k has cardinality $|N'| \geq (n-l)/2$. Let X_i be the random variable equal to 1 if buyer i has degree at most k in G, $X_i = 0$ otherwise, and let S be the random variable corresponding to the sum of the preferred valuations of the buyers of $N \setminus P$ of degree at most k in G, that is, $S = \sum_{i \in N \setminus P} v_i(m_i) X_i$. Again, the algorithm returns a solution of revenue at least $S/(k+1+\varepsilon)$ and k is constant, so that the expected value of S is $E(S) = E(\sum_{i \in N \setminus P} v(m_i) X_i) = \sum_{i \in N \setminus P} v(m_i) \cdot E(X_i) = \sum_{i \in N \setminus P} v(m_i) \cdot Prob(X_i = 1) \geq \sum_{i \in N \setminus P} v(m_i)/2 \geq \sum_{i \in N} v(m_i)/4 \geq \sum_{i \in N} opt_r(\mu, G)/4$.

We now show that $S = \Omega(opt_r(\mu, G))$ with high probability. To this aim, let us first observe that all buyers in $N \setminus P$, not being in the prefix P, have preferred valuations at most $(\sum_{i \in N} v(m_i))/l$. Moreover, it is possible to show that the variance of S is maximum when the overall sum of all the valuations of the buyers in $N \setminus P$ is compacted in a set N'' of l buyers, that is, the preferred valuations of the buyers in N'' are all equal to $(\sum_{i \in N} v(m_i))/l$, and all the other buyers not in N'' have null valuations. Therefore, it is sufficient to show that the probability of S being at least half of its expected value is high in this specific case. Under such an assumption, such a probability corresponds to the one that at least $l/4$ buyers of N'' in the random extraction of G receive degree at most k. It is possible to check that the associated random variable S' follows a hypergeometric distribution with expectation $l/2$. Then, by the tail bounds of such a distribution, the probability of less than $l/4$ successes for S' is at most $e^{-l/8} = 1/n$, thus proving the claim. □

Unfortunately, the argument in the previous theorem doesn't work for limited supply, as the values of the outcomes can significantly differ from the sum of the buyers' preferred valuations. However, we can make the probability of having a constant approximation arbitrarily high at the expense of the running time by means of the following algorithm (see also Algorithm 16): for a fixed constant parameter l, once randomly drawn $G \in \mathcal{G}_{n,\mathbf{d}}$, find optimal outcomes for all the possible subsets of at most $8l$ buyers and let $(X_1, \overline{p1})$ be the best resulting outcome (completed with bottom prices for the other buyers); run POWER-LAW and let $(X_2, \overline{p2})$ be the corresponding outcome; return the best of the two outcomes.

Algorithm 16: Constant (probabilistic) approximation algorithm for (SINGLE,REVENUE) with limited supply.

Input: An instance (μ, G) of the (SINGLE,REVENUE)-pricing problem and parameter l.
Output: A fair outcome $(\overline{X}, \overline{p})$.
for *a fixed constant parameter l, once randomly drawn $G \in \mathcal{G}_{n,\mathbf{d}}$* **do**
  Find optimal outcomes for all the possible subsets of at most $8l$ buyers;
end
Let $(X_1, \overline{p1})$ be the best resulting outcome (completed with bottom prices for the other buyers);
Let $(X_2, \overline{p2})$ be the outcome from POWER-LAW procedure;
return $(\overline{X}, \overline{p}) = \text{argmax } \{r(X_1, \overline{p1}), r(X_2, \overline{p2})\}$;

Notice that, since l is constant, in the initial phase the number of considered sets of $8l$ buyers are polynomial, and for each of them an optimal outcome can be obtained in polynomial time by an exhaustive search. Therefore, the algorithm has running time polynomial in the input size, but exponential in the parameter l. A proof similar to the one of Theorem 4.16 restricted to the buyers allocated in an optimal outcome shows the following theorem.

Theorem 4.17. *The above algorithm run on randomly drawn social graphs in $\mathcal{G}_{n,\mathbf{d}}$ returns a constant approximation ratio for* (SINGLE,REVENUE)*-pricing with probability at least $1 - e^{-l}$.*

We only briefly discuss the case of general valuations. The basic POWER-LAW algorithm is modified simply by formulating the problem of finding an optimal outcome for the subset of buyers associated to each color as an equivalent instance of the multiple-choice-knapsack problem, in which objects are partitioned in classes and at most one object per class can be selected. In particular, we associate to every buyer i a class containing m objects, each corresponding to a bundle size j and having profit $v_i(j)$ and weight j; the knapsack capacity is set to m. Then, we run the FPTAS for such a problem and transform the solution in an outcome with the same revenue, similarly to the single-minded case. Again, in order to achieve a constant approximation ratio arbitrarily high probability, the algorithm is combined with a preliminary exhaustive search of the best outcomes for all the possible subsets of $8l$ buyers, where l is a constant parameter.

By similar arguments to the ones of Theorem 4.17, it is possible to show the following theorem.

Theorem 4.18. *The above algorithm run on randomly drawn social graphs in $\mathcal{G}_{n,\mathbf{d}}$ returns a constant approximation ratio for* (GENERAL,REVENUE)*-pricing with probability at least $1 - e^{-l}$.*

4.6 Conclusions and Future Work

It would be nice to close the polylogarithmic gaps on the approximability of the maximum revenue on general social topologies. Moreover, it would be worth providing better probabilistic bounds for some of the approximation algorithms on power law graphs, or even a good approximation in the worst case.

It would be also interesting to consider other classical notions of envy-freeness, such as pair- and social envy-freeness. Moreover, we adopted the basic item-pricing policy, but also other relaxed forms of pricing, like bundle-pricing, might be considered. As in the previous related papers, it would be nice to investigate also the case of a limited set of allowable prices. Finally, it would be interesting to consider more general markets and other relevant social graph topologies.

Chapter 5

Conclusions

As mentioned in the abstract of this work, the final result of our effort is a set of new tools for a more accurate description of the real-world problems related to multi-unit markets, and a picture of the computational complexity of the problems arising from the interplay of various requirements considered in the model. Before suggesting some future directions for our work we would like to recap our results.

Social envy-freeness We considered a generalization of the standard concept of pair envy-freeness which takes into account buyers' limited knowledge and we explored the four cases arising from the problem of maximizing revenue under the assumptions of: ***i.*** *single-minded* buyers or with *general valuations*; ***ii.*** *item-pricing* or *bundle-pricing* schemes.

We proved the results showcased in Tables 5.1-5.2. More in detail, for the case of single-minded buyers we showed that in the case of item-pricing the problem is NP-hard for standard pair envy-freeness, and gave a corresponding FPTAS. Similarly, we proved the strong NP-hardness of the problem for social envy-freeness and provided a corresponding PTAS. While in the case of bundle-pricing we showed that the problem is NP-hard both for pair and social envy-freeness, and gave two corresponding FPTASs.

In the case of buyers with general valuations, at first, considering item-pricing schemes we provided a polylogarithmic lower bound on the achievable approximation ratio for pair envy-freeness (and thus also for social envy-freeness), while the $O(\log n)$-approximation algorithm provided in [82] for pair envy-freeness directly extends to social envy-freeness. Moreover, we gave an optimal allocation algorithm for social graphs with bounded treewidth. Later, considering bundle-pricing we observed that in [82] the authors had given a polylogarithmic lower bound on the achievable approximation ratio, and we gave an $O(\log n)$-approximation algorithm for social (and thus also pair) envy-freeness, improving upon their previous $O(\log n \cdot \log m)$ bound for pair envy-freeness given in [82]. Considering specific social topologies, we showed that the problem is APX-hard even in the case of empty social graphs, i.e., with all nodes isolated, and ignoring the supply constraint.

Finally, for all the above cases we provided optimal bounds on the price of envy-freeness (Table 5.3).

	Single-minded		
	Standard	Social	Standard with Free Disposal
Item-pricing	**NP-hard**	**NP-hard (strong)**	NP-hard [1]
	FPTAS	**PTAS**	~~$O(\log n)$~~ $\mathbf{(2+\varepsilon)}$
Bundle-pricing	**NP-hard**		NP-hard [1]
	FPTAS		FPTAS [1]

TABLE 5.1: Hardness and approximation results in the case of single-minded buyers.

	General Valuations		
	Standard	Social	Specific topologies
Item-pricing	$\mathbf{\Omega(\log^{\epsilon} n)}$		**polytime - bounded treewidth**
	$O(\log n)$ [1]		
Bundle-pricing	$\Omega(\log^{\epsilon} n)$ [1]		**APX-hard - empty graph,unlim.supply**
	~~$O(\log n \log m)$~~ [1] **O(log n)**		**1.59 - empty graph, unlim.supply**

TABLE 5.2: Hardness and approximation results in the case of buyers with general valuations.

	Single-minded	General Valuations
Item-pricing	**2**	$\mathbf{\Theta(\log n)}$
Bundle-pricing	**1**	$\mathbf{\Theta(\log n)}$

TABLE 5.3: Price of envy-freeness bounds. The upper bounds hold with respect to any social graph, the lower bounds even for paths.

Fair price discrimination We developed a framework which takes into account buyers' limited knowledge and defines a notion of fair price discrimination. In this setting we explored the four cases arising from maximizing social welfare or revenue in markets with *single-minded* buyers or buyers with *general valuations*.

We pointed out that the gain from price 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 items on sale.

[1]These results first appeared in (or directly follow from) [82]

We proved the results exhibited in Tables 5.4-5.5. I more detail, for single-minded buyers we showed that the social welfare maximization problem is strongly NP-hard and provided a corresponding PTAS. While in the case of revenue maximization, with the strong NP-hardness holding also in this case, we gave a $O(\log n)$-approximation algorithm.

In the case of buyers with general valuations our main result is a general reduction able to transform this kind of instances into "nearly" equivalent single-minded ones. As a consequence, we gave a PTAS also for social welfare maximization over buyers with general valuations, and showed that this approximation bound is tight, as the strong NP-hardness for single-minded instances also extends to this case. Furthermore we provided an $O(\log n + \log m)$-approximation algorithm for the problem of maximizing revenue.

Finally we considered special topologies for the social graph. For single-minded instances we provided an improved $(2 + \varepsilon)$ approximation for the case of undirected social graphs and an FPTAS for social graphs having a constant number of strongly connected components. While for general valuations, we provided polynomial time algorithms both for social welfare and revenue maximization for the following social graph topologies: undirected graphs, graphs with a constant number of strongly-connected components, and arborescences of strongly connected components.

<table>
<tr><td rowspan="2"></td><td colspan="2">Single-minded</td></tr>
<tr><td>General Graphs</td><td>Special Graphs</td></tr>
<tr><td rowspan="2">Social Welfare</td><td>NP-hard (strong)</td><td rowspan="4">$(2+\varepsilon)$ approximation - undirected graphs,
FPTAS - constant number of SCC</td></tr>
<tr><td>PTAS</td></tr>
<tr><td rowspan="2">Revenue</td><td>NP-hard (strong)</td></tr>
<tr><td>log n</td></tr>
</table>

Table 5.4: Hardness and approximation results for single-minded valuations pricing problems.

<table>
<tr><td rowspan="2"></td><td colspan="2">General Valuations</td></tr>
<tr><td>General Graphs</td><td>Special Graphs</td></tr>
<tr><td rowspan="2">Social Welfare</td><td>NP-hard (strong)</td><td rowspan="4">polytime - undirected graphs,
constant number of SCC,
arborescences of SCC</td></tr>
<tr><td>PTAS</td></tr>
<tr><td rowspan="2">Revenue</td><td>NP-hard (strong)</td></tr>
<tr><td>log n + log m</td></tr>
</table>

Table 5.5: Polynomial time, hardness, and approximation results for general valuations pricing problems.

Inequity aversion We considered a slight generalization of the framework for fair price discrimination and introduced the concept of buyers preselection, a powerful mechanism that we can find in online advertisement industry, and which gives great benefits. In fact, we showed that employing this mechanism gives a substantial gain in the achievable social welfare and revenue. In the context of this framework we studied the four different cases arising by considering: ***i.*** social welfare or seller's revenue maximization; ***ii.*** single-minded or general valuations.

More precisely, for single-minded buyers we showed the NP-hardness of the problem of social welfare maximization, together with an FPTAS. While for the problem of revenue maximization we proved that, under the unique game conjecture, it is hard to approximate the maximum revenue within a factor of $\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$, and provided a $O(\log n)$-approximation.

In the case of buyers with general valuations, we proved the social welfare maximization to be strong NP-hardness and gave a 2-approximation. While we showed that the problem of revenue maximization is impossible to approximate within a factor of $\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$ also in this case and described a $O(\log n + \log m)$-approximation algorithm (see Table 5.6).

Probably our most interesting contributions to this setting are the results on specific social graph topologies. In fact we considered a class of graphs that is typical of social networks, i.e., graphs where the node degrees follow a power-law distribution, and we showed three probabilistic approximation algorithm for the revenue maximization problem. More precisely, we gave a polynomial-time algorithm which returns a constant approximation of the maximum revenue with probability $1 - n^{-1}$ in the case of single-minded buyers and unlimited supply. Both for single minded and for general valuations,we described a scheme that, given a parameter l, returns a constant approximation of the maximum revenue with probability $1 - e^{-l}$, $\forall l > 0$. The running time is polynomial in the size of the instance and exponential in the size of the parameter l (see Table 5.7).

	Single-Minded	General Valuations
Social Welfare	NP-hard	NP-hard (strong)
	FPTAS	2
Revenue	$\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$	$\Omega(\frac{\sqrt{\log n}}{\log^2 \log n})$
	$O(\log n)$	$O(\log n + \log m)$

TABLE 5.6: Hardness and approximation results.

	Approximation	Probability
Single-Minded with Unlimited Supply	$O(1)$	$1 - n^{-1}$
Single-Minded	$O(1)$	$1 - e^{-l}$, $\forall l > 0$
General Valuations	$O(1)$	$1 - e^{-l}$, $\forall l > 0$

TABLE 5.7: Probabilistic approximation results.

Future Work

In the present work we focused on the modeling of social behavior and the exploration of price discrimination techniques in the context of multi-unit markets. We focused on these two aspects, because on the one hand they are typical of current real-world electronic markets (and resource allocation problems), and on the other hand they offer possibilities of huge gains in terms of objective functions maximization (such as revenue and social welfare).

We carefully selected different valuation functions for buyers, different notions of buyer's intelligent behavior (envy-freeness), many different notions of pricing schemes, and many other characteristics of a market in order to model a wide set of real-world scenarios. This allowed us to draw a well detailed picture of how these different assumptions interact and how they modify the complexity of the resulting pricing problems.

We outlined future promising research directions and relevant open questions originating from our work at the end of the previous chapters. Besides them, we remark that in general there are many other valuation functions, such as subadditive and submodular ones, additive valuations with budgets, additive valuations with cut-off thresholds, and many other that are worth investigating. As well as other forms of envy-freeness or "stability", like *proportionality*, or other objective global functions that certainly remain of research interest.

Not all of our results are tight, for sure it's worth investing in reducing several approximation gaps, and for many scenarios that we considered it would be very interesting to find special classes of graphs that represent real-world social networks and for which at the same time pricing problems can be better solved.

As a guideline to future research, we strongly believe that price discrimination, buyers preselection, and models for the limited knowledge of buyers are novel and powerful tools for the most important emerging real-world applications of combinatorial markets.

Bibliography

[1] Rediet Abebe, Jon M. Kleinberg, and David C. Parkes. Fair division via social comparison. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS, São Paulo, Brazil*, pages 281–289, 2017.

[2] Gagan Aggarwal, Tomás Feder, Rajeev Motwani, and An Zhu. Algorithms for multi-product pricing. In *Automata, Languages and Programming: 31st International Colloquium, ICALP, Turku, Finland, July 12-16. Proceedings*, pages 72–83. Springer, 2004.

[3] Gagan Aggarwal and Jason D Hartline. Knapsack auctions. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, Miami, Florida, USA, January 22-26*, pages 1083–1092. Society for Industrial and Applied Mathematics, 2006.

[4] William Aiello, Fan Chung Graham, and Linyuan Lu. A random graph model for power law graphs. *Experimental Mathematics*, 10(1):53–66, 2001.

[5] Noga Alon, Yishay Mansour, and Moshe Tenneholtz. Differential pricing with inequity aversion in social networks. In *ACM Conference on Electronic Commerce, EC, Philadelphia, PA, USA, June 16-20*, pages 9–24. ACM, 2013.

[6] Noga Alon, Yishay Mansour, and Moshe Tennenholtz. Differential pricing with inequity aversion in social networks. In *Proceedings of the fourteenth ACM Conference on Electronic Commerce, EC 2013, Philadelphia, PA, USA, June 16-20, 2013*, pages 9–24, 2013.

[7] Georgios Amanatidis, Evangelos Markakis, and Krzysztof Sornat. Inequity aversion pricing over social networks: Approximation algorithms and hardness results. In *41st International Symposium on Mathematical Foundations of Computer Science, MFCS 2016, August 22-26, 2016 - Kraków, Poland*, pages 9:1–9:13, 2016.

[8] Eric T Anderson and Duncan I Simester. Price stickiness and customer antagonism. *The Quarterly Journal of Economics*, 125(2):729–765, 2010.

[9] Elliot Anshelevich, Koushik Kar, and Shreyas Sekar. Envy-free pricing in large markets: Approximating revenue and welfare. *ACM Transactions on Economics and Computation*, 5(3):16, 2017.

[10] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2):129–150, 2004.

[11] Per Austrin, Subhash Khot, and Muli Safra. Inapproximability of vertex cover and independent set in bounded degree graphs. *Theory of Computing*, 7(1):27–43, 2011.

[12] Benjamin Avi-Itzhak, Hanoch Levy, and David Raz. A resource allocation queueing fairness measure: Properties and bounds. *Queueing Systems*, 56(2):65–71, 2007.

[13] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings 9th ACM Conference on Electronic Commerce EC, Chicago, IL, USA*, pages 50–59, 2008.

[14] Yair Bartal, Rica Gonen, and Noam Nisan. Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge, TARK, Bloomington, Indiana, USA, June 20-22*, pages 72–87, 2003.

[15] Vittorio Bilò, Angelo Fanelli, Michele Flammini, and Luca Moscardelli. When ignorance helps: Graphical multicast cost sharing games. *Theoretical Computer Science*, 411(3):660–671, 2010.

[16] Vittorio Bilò, Angelo Fanelli, Michele Flammini, and Luca Moscardelli. Graphical congestion games. *Algorithmica*, 61(2):274–297, 2011.

[17] Vittorio Bilò, Michele Flammini, and Gianpiero Monaco. Approximating the revenue maximization problem with sharp demands. *Theoretical Computer Science*, 662:9–30, 2017.

[18] Vittorio Bilò, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. Pricing problems with buyer preselection. In *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK*, pages 47:1–47:16, 2018.

[19] Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings 6th ACM Conference on Electronic Commerce, EC, Vancouver, BC, Canada, June 5-8*, pages 44–51, 2005.

[20] S. Brânzei, A. Filos-Ratsikas, P. B. Miltersen, and Y. Zeng. Envy-free pricing in multi-unit markets. *arXiv:1602.08719*, 2016.

[21] Simina Brânzei and Aris Filos-Ratsikas. Walrasian dynamics in multi-unit markets. *CoRR*, abs/1712.08910, 2017.

[22] Simina Brânzei, Aris Filos-Ratsikas, Peter Bro Miltersen, and Yulong Zeng. Walrasian pricing in multi-unit auctions. In *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS, August 21-25, Aalborg, Denmark*, pages 80:1–80:14, 2017.

[23] Patrick Briest. Uniform budgets and the envy-free pricing problem. In *Automata, Languages and Programming, 35th International Colloquium, ICALP, Reykjavik, Iceland, Proceedings, Part I: Tack A: Algorithms, Automata, Complexity, and Games*, pages 808–819, 2008.

[24] Patrick Briest and Piotr Krysta. Single-minded unlimited supply pricing on sparse instances. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, Miami, Florida, USA*, pages 1093–1102, 2006.

[25] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012.

[26] Dennis W Carlton and Jeffrey M Perloff. *Modern industrial organization.* Pearson Higher Ed, 2015.

[27] D.W. Carlton and J.M. Perloff. *Modern Industrial Organization.* The Addison-Wesley series in economics. Pearson/Addison Wesley, 2005.

[28] Parinya Chalermsook, Julia Chuzhoy, Sampath Kannan, and Sanjeev Khanna. Improved hardness results for profit maximization pricing problems with unlimited supply. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques - 15th International Workshop, APPROX, and 16th International Workshop, RANDOM, Cambridge, MA, USA. Proceedings*, pages 73–84, 2012.

[29] Parinya Chalermsook, Bundit Laekhanukit, and Danupon Nanongkai. Graph products revisited: Tight approximation hardness of induced matching, poset dimension and more. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, New Orleans, Louisiana, USA*, pages 1557–1576, 2013.

[30] Parinya Chalermsook, Bundit Laekhanukit, and Danupon Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time)

approximation hardnesses. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS, Berkeley, CA, USA*, pages 370–379, 2013.

[31] Ning Chen and Xiaotie Deng. Envy-free pricing in multi-item markets. In *Automata, Languages and Programming, 37th International Colloquium, ICALP, Bordeaux, France, Proceedings, Part II*, pages 418–429, 2010.

[32] Ning Chen, Xiaotie Deng, Paul W Goldberg, and Jinshan Zhang. On revenue maximization with sharp multi-unit demands. *Journal of Combinatorial Optimization*, 31(3):1174–1205, 2016.

[33] Ning Chen, Xiaotie Deng, and Xiaoming Sun. On complexity of single-minded auction. *Journal of Computer and System Sciences*, 69(4):675–687, 2004.

[34] Ning Chen, Xiaotie Deng, and Hong Zhu. Combinatorial auction across independent markets. In *Proceedings 4th ACM Conference on Electronic Commerce, EC, San Diego, California, USA, June 9-12*, pages 206–207. ACM, 2003.

[35] Ning Chen, Arpita Ghosh, and Sergei Vassilvitskii. Optimal envy-free pricing with metric substitutability. *SIAM Journal on Computing*, 40(3):623–645, 2011.

[36] Maurice Cheung and Chaitanya Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS, Philadelphia, PA, USA*, pages 35–44, 2008.

[37] Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, Nicolas Maudet, et al. Reaching envy-free states in distributed negotiation settings. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI, Hyderabad, India*, volume 7, pages 1239–1244, 2007.

[38] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Distributed fair allocation of indivisible goods. *Artificial Intelligence*, 242:1–22, 2017.

[39] Yann Chevaleyre, Ulrich Endriss, and Nicolas Maudet. Allocating goods on a graph to eliminate envy. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, Vancouver, British Columbia, Canada*, pages 700–705, 2007.

[40] Edward H Clarke. Multipart pricing of public goods. *Public choice*, 11(1):17–33, 1971.

[41] Riccardo Colini-Baldeschi, Stefano Leonardi, Piotr Sankowski, and Qiang Zhang. Revenue maximizing envy-free fixed-price auctions with budgets. In *Web and Internet Economics - 10th International Conference, WINE, Beijing, China. Proceedings*, pages 233–246, 2014.

[42] Riccardo Colini-Baldeschi, Stefano Leonardi, and Qiang Zhang. Revenue maximizing envy-free pricing in matching markets with budgets. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*, pages 207–220, 2016.

[43] Riccardo Colini-Baldeschi, Stefano Leonardi, and Qiang Zhang. Revenue maximizing envy-free pricing in matching markets with budgets. In *Web and Internet Economics - 12th International Conference, WINE, Montreal, Canada, December 11-14, Proceedings*, pages 207–220. Springer, 2016.

[44] Erik D Demaine, Uriel Feige, MohammadTaghi Hajiaghayi, and Mohammad R Salavatipour. Combination can be hard: Approximability of the unique coverage problem. *SIAM Journal on Computing*, 38(4):1464–1483, 2008.

[45] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

[46] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. *J. Artif. Intell. Res.*, 37:85–98, 2010.

[47] Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. *CoRR*, abs/1401.3834, 2014.

[48] Shahar Dobzinski and Noam Nisan. Multi-unit auctions: Beyond roberts. *J. Economic Theory*, 156:14–44, 2015.

[49] David Easley and Jon Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world.* Cambridge University Press, 2010.

[50] Uriel Feige. Relations between average case complexity and approximation complexity. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, STOC, Montréal, Québec, Canada*, pages 534–543, 2002.

[51] Uriel Feige and Michel X. Goemans. Aproximating the value of two prover proof systems, with applications to MAX 2sat and MAX DICUT. In *Third Israel Symposium on Theory of Computing and Systems, ISTCS 1995, Tel Aviv, Israel, January 4-6, 1995, Proceedings*, pages 182–189, 1995.

[52] Michal Feldman, Amos Fiat, Stefano Leonardi, and Piotr Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *ACM Conference on Electronic Commerce, EC, Valencia, Spain*, pages 532–549, 2012.

[53] Amos Fiat, Andrew V. Goldberg, Jason D. Hartline, and Anna R. Karlin. Competitive generalized auctions. In *Proceedings on 34th Annual ACM Symposium on*

Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada, pages 72–81, 2002.

[54] Amos Fiat and Amiram Wingarten. Envy, multi envy, and revenue maximization. In *Internet and Network Economics, 5th International Workshop, WINE, Rome, Italy. Proceedings*, pages 498–504, 2009.

[55] Michele Flammini, Manuel Mauro, and Matteo Tonelli. On fair price discrimination in multi-unit markets. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI, Stockholm, Sweden.* IJCAI Org, Vienna, Austria, 2018.

[56] Michele Flammini, Manuel Mauro, and Matteo Tonelli. On social envy-freeness in multi-unit markets. In *32nd AAAI Conference on Artificial Intelligence.* AAAI Press, Palo Alto, California, USA, 2018.

[57] Michele Flammini, Manuel Mauro, and Matteo Tonelli. On social envy-freeness in multi-unit markets. *Artificial Intelligence*, 2018. Elsevier, Amsterdam, Netherlands.

[58] Michele Flammini, Manuel Mauro, and Matteo Tonelli. On fair price discrimination in multi-unit markets. *Accepted with major revision for publication in Artificial Intelligence*, 2019. Elsevier, Amsterdam, Netherlands.

[59] Michele Flammini, Manuel Mauro, Matteo Tonelli, and Cosimo Vinci. Inequity aversion pricing in multi-unit markets. In *Manuscript*, 2018.

[60] Duncan K Foley. Resource allocation in the public sector. *Yale Economic Essays*, 7:73–76, 1967.

[61] Alan M. Frieze, Michael Krivelevich, and Clifford D. Smyth. On the chromatic number of random graphs with a fixed degree sequence. *Combinatorics, Probability & Computing*, 16(5):733–746, 2007.

[62] Damien Geradin and Nicolas Petit. Price discrimination under ec competition law. *The pros and cons of price discrimination*, pages 21–63, 2005.

[63] Gagan Goel, Vahab S. Mirrokni, and Renato Paes Leme. Clinching auction with online supply. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, New Orleans, Louisiana, USA, January 6-8*, pages 605–619, 2013.

[64] Andrew V. Goldberg and Jason D. Hartline. Competitive auctions for multiple digital goods. In *Algorithms - ESA 2001, 9th Annual European Symposium, Aarhus, Denmark, August 28-31, 2001, Proceedings*, pages 416–427, 2001.

[65] Andrew V. Goldberg and Jason D. Hartline. Competitiveness via consensus. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA*, pages 215–222, 2003.

[66] Alexander Grigoriev, Joyce van Loon, René Sitters, and Marc Uetz. How to sell a graph: Guidelines for graph retailers. In *Graph-Theoretic Concepts in Computer Science, 32nd International Workshop, WG 2006, Bergen, Norway, June 22-24, 2006, Revised Papers*, pages 125–136, 2006.

[67] Theodore Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

[68] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, Vancouver, British Columbia, Canada*, pages 1164–1173, 2005.

[69] Venkatesan Guruswami, Prasad Raghavendra, Rishi Saket, and Yi Wu. Bypassing UGC from some optimal geometric inapproximability results. *ACM Trans. Algorithms*, 12(1):6:1–6:25, 2016.

[70] Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 122, 2013.

[71] Jason D. Hartline and Vladlen Koltun. Near-optimal pricing in near-linear time. In *Algorithms and Data Structures, 9th International Workshop, WADS, Waterloo, Canada, Proceedings*, pages 422–431, 2005.

[72] Jason D. Hartline and Qiqi Yan. Envy, truth, and profit. In *Proceedings 12th ACM Conference on Electronic Commerce EC, San Jose, CA, USA*, pages 243–252, 2011.

[73] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.

[74] Michael J. Kearns, Michael L. Littman, and Satinder P. Singh. Graphical models for game theory. In *UAI: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence, University of Washington, Seattle, Washington, USA*, pages 253–260, 2001.

[75] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 767–775, 2002.

[76] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable csps? *SIAM J. Comput.*, 37(1):319–357, 2007.

[77] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-epsilon. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.

[78] Piotr Krysta, Orestis Telelis, and Carmine Ventre. Mechanisms for multi-unit combinatorial auctions with a few distinct goods. *J. Artif. Intell. Res.*, 53:721–744, 2015.

[79] Jean-Jacques Laffont, Patrick Rey, and Jean Tirole. Network competition: Ii. price discrimination. *The RAND J. of Economics*, 29(1):38–56, 1998.

[80] Eugene L Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 4(4):339–356, 1979.

[81] Daniel Lehmann, Liadan Ita Oćallaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, 2002.

[82] Gianpiero Monaco, Piotr Sankowski, and Qiang Zhang. Revenue maximization envy-free pricing for homogeneous resources. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI, Buenos Aires, Argentina*, pages 90–96, 2015.

[83] Ahuva Mu'Alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.

[84] Noam Nisan, Jason Bayer, Deepak Chandra, Tal Franji, Robert Gardner, Yossi Matias, Neil Rhodes, Misha Seltzer, Danny Tom, Hal R. Varian, and Dan Zigmond. Google's auction for TV ads. In *Automata, Languages and Programming, 36th International Colloquium, ICALP, Rhodes, Greece, July 5-12, Proceedings, Part II*, pages 309–327, 2009.

[85] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory.* Cambridge University Press, 2007.

[86] Louis Phlips. *The Economics of Price Discrimination.* Cambridge University Press, 1983.

[87] Arthur Cecil Pigou. *The Economics of Welfare.* Palgrave Macmillan UK, 1920.

[88] Nancy Laura Stokey. Intertemporal price discrimination. *Quarterly J. of Economics*, 93(3):355–371, 1979.

[89] Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.

[90] Hal Ronald Varian. Equity, envy, and efficiency. *Journal of economic theory*, 9(1):63–91, 1974.

[91] Hal Ronald Varian. Chapter 10 price discrimination. In *Price Discrimination*, volume 1 of *Handbook of Industrial Organization*, pages 597 – 654. Elsevier, 1989.

[92] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

[93] Léon Walras, W. Jaffé, American Economic Association, and Royal Economic Society (Great Britain). *Elements of Pure Economics: Or the Theory of Social Wealth.* American Economic Association. American Economic Association and the Royal Economic Association and the Royal Economic Society, 1954.