

Inequity Aversion Pricing in Multi-Unit Markets

Michele Flammini¹ and Manuel Mauro² and Matteo Tonelli³ and Cosimo Vinci⁴

Abstract. We build upon previous models for differential pricing in social networks and fair price discrimination in markets, 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 the inapproximability within any constant factor 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 seller’s revenue maximization with high probability, thus improving upon the general case.

1 INTRODUCTION

The study of differential pricing in economics can be traced back to the beginning of the 20th century [34, 36, 38, 39], and it is progressively attracting increasing interest. In fact, this feature is nowadays widely used in several settings, like in raising the price of the tickets as time approaches the related event, in assigning different prices for the same products in geographically apart markets, in offering discounts to premium customers, and so forth. Several works outlined various levels of price discrimination [37], while others tried to figure out proper forms that buyers can perceive as fair [3, 4, 27, 28]. Another related concept of customers’ satisfaction widely investigated in markets is envy-freeness. Namely, buyers, given the bundle of goods they receive and the assigned prices, should not prefer other bundles or bundles provided to other buyers. Seminal works in the field are [7, 13, 22, 30, 31], where authors describe logarithmic approximation algorithms for maximizing the seller’s revenue under various assumptions, while hardness of approximations complementing the previous results can be found in [12, 16, 17, 18, 24]. Many variants can be also found in the works of [5, 9, 19, 20, 21, 27]. None of these works however considered discriminatory pricing policies, except for a mild form called bundle pricing, in which non-proportional or lower average prices are assigned to bundles of bigger sizes. However, each bundle preserves a unique price for all the buyers.

1.1 Our contribution

In this paper we build upon previous works dealing with explicit forms of differential pricing [3, 4, 28]. Namely, like in [3, 4] we as-

sume 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 [28], we model the underlying scenario as a multi-unit market with relaxed fair price discrimination constraints. This type of markets has been largely investigated [11, 27, 29, 35], also for the purpose of determining good approximating incentive-compatible mechanisms [25, 33], as well as in other related auctions [2]. In fact, it is typical of many real-world situations where homogeneous items are on sale, like in commodity markets. We adopt one of the most investigated envy-freeness notions in this setting, proposed in [30], according to which each buyer must be assigned a bundle which maximizes her utility, i.e., the difference between her valuation for the item and its price.

We study four different cases arising by considering *i.* social welfare or seller’s revenue maximization, and *ii.* single-minded or general valuations. Our results comprise approximation algorithms and hardness of approximation results for all the above-mentioned cases, including the inapproximability within any constant factor on the achievable revenue under the unique game conjecture (see Table 1). It is worth mentioning that such a UGC-hardness result applies also to the model of [4], thus strengthening their APX-hardness.

Furthermore, we consider the case of specific topologies, focusing on the fundamental property that has been shown to characterize real world social networks, i.e., graphs where the node degrees follow a power-law distribution [15, 26]. In particular, we provide polynomial time algorithms able to achieve constant approximations on the maximum revenue with high probability, both for single-minded and general valuations (see Table 2). To the best of our knowledge, this is one of the first results on the social influence in markets showing that, assuming a paradigmatic social topology, better approximations (from logarithmic to constant) can be accomplished with respect to the general unrestricted case. On this respect, we further remark that the previous work in this setting considered topologies, like bounded treewidth, that while yielding efficient solutions, are more representative of situations in which social relationships are structured in a non-spontaneous way or according to external factors.

Due to space constraints, some of the less significant proofs are omitted and left to the full version of this paper.

	Single-Minded	General Valuations
Social Welfare	NP-hard (Thm. 4)	strong NP-hard (Thm. 10)
	FPTAS (Thm. 4)	2 (Thm. 9)
Revenue	$\omega(1)$ (Thm. 6)	$\omega(1)$ (Cor. 12)
	$O(\log n)$ (Cor. 5)	$O(\log n + \log m)$ (Cor. 11)

Table 1. Our hardness results (first and third row) and approximation results (second and fourth row).

¹ Gran Sasso Science Institute, Italy, email: michele.flammini@gssi.it

² Gran Sasso Science Institute, Italy, email: manuel.mauro@gssi.it

³ Gran Sasso Science Institute, Italy, email: matteo.tonelli@gssi.it

⁴ Gran Sasso Science Institute, Italy, email: cosimo.vinci@gssi.it

	Apx.	Prob.
Single-Minded (Unl. Supply) (Thm. 15)	$O(1)$	$1 - n^{-1}$
Single-Minded (Thm. 17)	$O(1)$	$1 - e^{-w}, \forall w > 0$
General Valuations (Thm. 17)	$O(1)$	$1 - e^{-w}, \forall w > 0$

Table 2. Approximation results for the problem of maximizing the revenue with high probability (n denotes the number of buyers).

1.2 Related work

Envy-freeness in multi-unit markets has been investigated in [27], where the authors studied budgeted buyers with additive valuations and considering progressively stronger levels of price discrimination. After showing that the different levels correspond to increasing maximum revenues, they proved that the revenue maximization problem is NP-hard, and provided a polynomial time 2-approximation algorithm for item-pricing without discrimination.

Monaco et al. [35] investigated multi-unit markets with buyers having no budgets and gave hardness and approximation results for the revenue maximization problem in several cases arising by assuming different notions of envy-freeness, item- or bundle-pricing, and single-minded or general valuation buyers.

Social influence in markets has been recently considered in [1, 23, 29], where the authors focused on a relaxed notion of social envy-freeness restricting envy only to social neighbors, and applied it to problems like cake cutting, distributed negotiation, and multi-unit markets. The authors of [3, 4, 9, 10] considered also the possibility of dismissing some of the buyers from the market, which translates into limiting the social awareness acting on the social graph topology.

Frameworks more related to the present paper have been studied in [3, 4, 28]. In particular, in [3] 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. In their setting, every node contributes a revenue which is equal to a given function of the assigned price. The authors gave an optimal revenue maximizing algorithm if buyers preselection is not allowed, while with preselection and arbitrary revenue functions they proved that the problems is hard to approximate within $n^{1-\epsilon}$ for any fixed $\epsilon > 0$, where n is the number of buyers, and NP-hard even for constant and single-valued revenue functions, this last family being the one considered in this paper. Finally, they provided polynomial time algorithms for trees and bounded treewidth graphs. However, unlike in our paper, they assumed that prices can belong to a given finite set.

Along this line of research, the authors of [4] proved that the revenue maximization problem remains NP-hard even if we restrict the revenue functions to be only single value (and not constant), and gave improved approximation results, especially when the number of distinct prices is small. Moreover, they showed the APX-hardness for the special case in which neighbor prices must be identical. Finally, they extended the model of [3] to allow the assignment of more than one item per node, which corresponds to multi-unit markets with single-minded buyers and unlimited supply. Various approximation algorithms are given also in this setting. Our inapproximability result applies even to their basic model.

A related model of fair price discrimination in multi-unit markets has been investigated in [28]. While it extends the above frameworks by allowing general valuations, and thus with the possibility of allocating bundles of different sizes to a single buyer, it does not allow the exclusion of buyers, and it does not consider any slackness in fair pricing. The authors investigated the computational complexity of

the social welfare and the revenue maximization problems, providing hardness and approximation results under various assumptions on the buyers' valuations and on the social graph topology.

As already mentioned, we borrow the model of fair price discrimination from [28] and extend it with the features of buyers' preselection and additive slackness in fairness constraints of [3, 4]. This allows to generalize in a unified framework all the above models in different respects. In particular, while the notion of fair pricing is of [3], generalizing [28], buyers with single-minded valuations are able to incorporate the extended model of [4]. Moreover, as in [28], our model is able to express the further constraint of a limited supply of items and general buyers' valuations, according to which buyers are not interested in a single bundle, but have different utilities for different sizes. We remark that, without limited supply and general valuations, the envy-freeness constraints are meaningless, and as it happens in some previous related papers, can be ignored.

We finally stress that preselection of buyers should not be intended in the negative sense of excluding some of them, but rather in the positive way that is able to capture common real world situations in which offers are not made to all customers, but only to prospective or promising ones. This normally happens in the new targeted marketing strategies, that deliver individualized messages and product offerings, thanks to profiling and data analysis techniques. Hence, exclusion should be more properly interpreted as inclusion of buyers receiving offers.

2 PRELIMINARIES

We define a multi-unit market μ as 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 homogeneous items, and each v_i is the valuation function of buyer $i \in N$. Items in M are considered identical by the buyers, so that for every $i \in N$, we can represent v_i as a vector $v_i = (v_i(1), \dots, v_i(m))$ expressing, for each natural number $j \in \{1, \dots, m\}$, the maximum amount $v_i(j) \in \mathbb{R}$ that 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 the *single-minded* case, in which every buyer i is interested only in bundles of a *preferred bundle size* m_i , and *general valuations*, i.e., the unrestricted case.

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 $\bar{p} = (p_1, \dots, p_n)$ 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 $N(i) = \{k \in N \mid (i, k) \in E\}$. 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 1. A price vector \bar{p} is *fair* w.r.t. 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 $\bar{X} = (X_1, \dots, X_n)$ such that $X_i \subseteq M$ is the set of items sold to buyer i , and we call a

pair (\bar{X}, \bar{p}) an *outcome*. (\bar{X}, \bar{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 (\bar{X}, \bar{p}) is *envy-free* if $X_i \in \arg \max_{X \subseteq M} u_i(X, p_i)$ 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, i.e., *fair outcomes*.

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

We study the (fair) *pricing problems* of determining fair outcomes that maximize two fundamental metrics: **i.** The *social welfare* $sw(\bar{X}, \bar{p}) = \sum_{i=1}^n v_i(|X_i|)$. **ii.** The *seller's revenue* $r(\bar{X}, \bar{p}) = \sum_{i=1}^n p_i \cdot |X_i|$.

We denote as $opt_{sw}(\mu, G)$ (resp. $opt_r(\mu, G)$) the maximum possible social welfare (resp. revenue) achievable by an outcome for μ fair under G .

By the individual rationality constraint, for any feasible outcome (\bar{X}, \bar{p}) , it holds $sw(\bar{X}, \bar{p}) \geq r(\bar{X}, \bar{p})$, 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 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 $\bar{p} \in (\mathbb{R} \cup \{\perp\})^n$. The notion of fair pricing is then extended as follows:

Definition 3. A price vector \bar{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 are then computed considering only buyers not receiving bottom prices.

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

Proposition 1. 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.

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, neither hardness results for single-minded buyers directly extend to general valuations, nor approximation bounds for general valuations automatically transfer to single-minded instances.

3 SINGLE-MINDED

We first provide optimal results for the social welfare.

Theorem 4. (SINGLE,WELFARE)-pricing is NP-hard, but admits an FPTAS.

Proof (sketch). Both the hardness and the FPTAS are obtained by showing a direct correspondence between this problem and KNAPSACK. \square

Regarding revenue maximization, an approximation algorithm can be obtained with the same approach used in Theorem 5 of [28] for the same case without bottom prices.

Corollary 5. (SINGLE,REVENUE)-pricing admits a $\log n$ -approximation algorithm.

The approximation factor above presented is nearly optimal, as the following negative result holds.

Theorem 6. It is Unique-Game-hard to approximate (SINGLE,REVENUE)-pricing within any constant factor.

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 been shown to be Unique-Game-hard to approximate within a factor of $\Omega(\frac{\delta}{\log^2 \delta})$ [6] (see [32] for details on the Unique-Game conjecture and hardness). 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:

Definition 7. 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 H if:

- i.** (Coloring) $\forall S_i \in \mathcal{S}, \forall u, v \in S_i, (u, v) \notin F$.
- ii.** Each node $u \in V$ has at most one neighbor in each subset $S_i \in \mathcal{S}$.

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

Lemma 8. Any graph $H = (V, F)$ with node degrees at most δ admits a good partition \mathcal{S} which can be found in polynomial time and such that $|\mathcal{S}| \leq \delta^2 + 1$.

Proof of Lemma 8. Consider the graph $H^2 = (V, F^2)$, where:

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

By construction H^2 has degree at most δ^2 . By the Brooks' Theorem [14] 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. \square

We are ready to prove Theorem 6.

Proof of Theorem 6. Consider the following reduction from INDEPENDENT SET instances $H = (V, F)$ with bounded degree δ to instances (μ, G) of (SINGLE,REVENUE)-pricing:

- i.** Find a good partition $\mathcal{S} = \{S_1, \dots, S_D\}$ of H , 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 $\alpha(i, h) = \alpha(h, i) = 0$.

iv. For each edge $\{u, v\} \in F$, with $u \in S_d, v \in S_{d'}$ and $d < d'$, add to G all arcs (i, h) such that $i \in \mathcal{N}_u$ and $h \in \mathcal{N}_v$ with $\alpha(i, h) = 0$.

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

We are going to prove our claim by showing that (\Rightarrow) if H 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 H has an independent set of cardinality at least $\frac{k}{2}$.

(\Rightarrow) Let $I \subseteq V$ be an independent set of H having cardinality k . Consider the outcome (\bar{X}, \bar{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 $|X_i| = 0$ and $p_i = 2^{-d}$.

Notice that \bar{p} is fair under G . In fact, by construction of G , the set of the neighbors of buyer $i \in \mathcal{N}_v$ in G is a subset of $\mathcal{N}_v \cup \bigcup_{u \in \Delta(v)} \mathcal{N}_u$, where $\Delta(v)$ denotes the set of the neighbors of node v in H . 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, \bar{p} is fair. Furthermore, (\bar{X}, \bar{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, $r(\bar{X}, \bar{p}) = k$.

(\Leftarrow) Assume that the reduced instance admits an outcome (\bar{X}, \bar{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 $p_i \leq v_i(1)$, as this can only increase the revenue. Similarly, as \bar{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 (\bar{X}, \bar{p}) can be described by means of a vector $\bar{\pi} \in (\mathbb{R} \cup \{\perp\})^{|V|}$, where component π_v is equal to the price proposed to all buyers in \mathcal{N}_v .

After these preliminary remarks, let us construct a suitable subset of nodes $I \subseteq V$ 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 H , 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 (\bar{X}, \bar{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 cannot have more than 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 $|I| \geq \frac{k}{2}$.

We observe that the above reduction is polynomial if δ is constant

with respect to the input size. Finally, since INDEPENDENT SET with maximum degree δ is Unique-Game-hard to approximate within a factor of $\frac{t\delta}{\log^2 \delta}$ for some constant $t > 0$, and because of the above polynomial reduction, we have that, for any fixed $c \geq 1$, there exists a sufficiently large integer δ such that any c -approximation algorithm for (SINGLE, REVENUE)-pricing can be polynomially turned into a $\frac{t\delta}{\log^2 \delta}$ -approximation algorithm for INDEPENDENT SET with maximum degree δ , and this shows the inapproximability of (SINGLE, REVENUE)-pricing within any constant factor $c \geq 1$. \square

We remark that, since in the reduction every buyers has preferred size 1, the UCG-hardness closes an open question raised in [4], extending their APX-hardness to inapproximability results holding for any constant approximation factor.

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 [28]. Given a buyer $i \in N$, let $S_i = \{m_i^1, \dots, m_i^\ell\}$ be the bundle sizes that are in the demand set of i for at least one positive price, that is, among the preferred ones for such a 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, \dots, i_ℓ , 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 of [28] 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 has not such a property. Hence, we devise an ad-hoc procedure that 2-approximates the social welfare, while allocating prefixes of marginal buyers.

Theorem 9. (GENERAL, WELFARE)-pricing admits a 2-approximation algorithm.

Proof. Given an instance (μ, G) of (GENERAL, WELFARE)-pricing, let (μ', G') be the associated output of the reduction of [28]. Consider the following algorithm for maximizing $opt_{sw}(\mu', G')$:

- 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:
 - (\bar{X}', \bar{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;
 - (\bar{X}'', \bar{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 nonexistent all the marginal buyers are allocated in (\bar{X}', \bar{p}') , which 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 not items to all the other buyers.

- Return $\arg \max\{sw(\bar{X}', \bar{p}'), sw(\bar{X}'', \bar{p}'')\}$.

Notice that both (\bar{X}', \bar{p}') and (\bar{X}'', \bar{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 $p_{i_j} \leq \frac{v_{i_j}(m_{i_j})}{m_{i_j}}$ we have $|X'_{i_j}| = |X''_{i_j}| = m_{i_j}$, both \bar{X}' and \bar{X}'' are also envy-free.

Notice also that $sw(\bar{X}', \bar{p}') + sw(\bar{X}'', \bar{p}'')$ 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 (\bar{X}', \bar{p}') and (\bar{X}'', \bar{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. \square

Considering general valuations worsens the complexity of finding an outcome that maximizes social welfare. In fact, we are able to prove the following theorem.

Theorem 10. (GENERAL,WELFARE)-pricing is strongly NP-hard.

In the same fashion of Corollary 5, it is possible to exploit the 2-approximation for the social welfare in order to obtain the following result, in which the approximation derives from the number of buyers in the reduced instance $(n \cdot m)$.

Corollary 11. (GENERAL,REVENUE)-pricing admits a $2(\log n + \log m)$ -approximation algorithm.

The hardness result provided in Theorem 6 directly extends to general valuations, as in the provided reduction m is polynomially bounded in the size of the instance.

Corollary 12. (GENERAL,REVENUE)-pricing is Unique-Game-hard to approximate within any constant factor.

This last negative bound extends also to the setting proposed in [4]. In fact, the same reduction can be used setting the slackness between neighbouring equal to 0.

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. Due to space constraints, many details of algorithms and proofs are only sketched.

Let $\bar{d} = (d_1, d_2, \dots, d_n)$ be a non-decreasing sequence or vector of n strictly positive integers, whose sum is even. We assume that \bar{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 \bar{d} can be suitably upper bounded as $\sum_{h=k+1}^n n(h) = O\left(\frac{n}{k^{\gamma-1}}\right)$.

Let $\mathcal{G}_{n,\bar{d}}$ be the class of graphs with node set $N = \{1, 2, \dots, n\}$, in which the sequence of node degrees listed in non-decreasing order coincides with \bar{d} . We assume that the social graph G is randomly drawn in $\mathcal{G}_{n,\bar{d}}$ 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 $d_{\pi(i)}$.

Before providing constant approximations for power-law graphs, let us give the following key lemma, which will be useful in the sequel.

Lemma 13. Given any family of graphs \mathcal{G} and a fixed integer $k > 0$, if a $(k+1)$ -coloring for any graph in \mathcal{G} exists and can be determined in polynomial time, then: (i) (SINGLE,REVENUE)-pricing and (GENERAL,REVENUE)-pricing restricted to social graphs in \mathcal{G} admit a $(k+1+\epsilon)$ -approximation algorithm for any $\epsilon > 0$; (ii) (SINGLE,REVENUE)-pricing with unlimited supply restricted to social graphs in \mathcal{G} admits a $(k+1)$ -approximation algorithm.

Proof. We first show part (i). Consider the single-minded case. 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 $(1+\epsilon/(k+1))$ -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 for knapsack on the equivalent knapsack instance with capacity m containing an object o_i for every buyer $i \in N$ with profit $z_i = v_i(m_i)$ 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 $v_i(m_i)/m_i$ 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 $(k+1+\epsilon)$ -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 $r \geq opt_r(\mu, G)/(k+1)$ 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 .

When considering general valuations, we consider the above algorithm, but we compute the $(1+\epsilon/(k+1))$ -approximation for the submarket restricted to each set of buyers N_i by running the FPTAS on the equivalent multi-choice knapsack instance defined as follows: 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 . As in the previous case, the returned solution can be directly translated to an outcome of the original problem with the same revenue, and the above algorithm provides again a $(k+1+\epsilon)$ -approximation.

For what concerns part (ii), consider again the above method. In the case of unlimited supply, we observe that, for any fixed color i , the optimal solution for the submarket restricted to N_i can be computed in polynomial time, i.e., it suffices assigning price $v_i(m_i)/m_i$ per item to each buyer in N_i . Thus, a $(k+1)$ -approximation can be determined in a similar fashion. \square

As a direct consequence of the above lemma, constant approximation algorithms can be obtained for graphs with maximum degree bounded by a constant k (thanks to the well-known greedy algorithm returning a $k+1$ coloring), 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 low 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: once drawn $G \in \mathcal{G}_{n,\bar{d}}$ according to the above random process, consider the subset $N' \subseteq N$ of buyers of degree at most k , and then run one of the algorithms of Lemma 13 (according to the specific pricing problem) for bounded degree graphs on the instance (μ, G') , where G' is the subgraph induced by N' .

Lemma 14. *POWER-LAW executed on randomly drawn social graphs in $\mathcal{G}_{n,\bar{d}}$ has constant expected approximation ratio for (SINGLE,REVENUE)-pricing and (GENERAL,REVENUE)-pricing.*

Proof. We prove that the expected revenue of the above algorithm is $\Omega(\text{opt}_r(\mu, G))$.

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 $|N'| \geq n/2$. For a given optimal outcome, let N_{opt} be the set of buyers receiving a bundle and let m_i be the bundle assigned to each buyer $i \in N_{\text{opt}}$. 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 in N_{opt} having degree at most k in G , that is, $S = \sum_{i \in N_{\text{opt}}} v_i(m_i) X_i$. Then, since POWER-LAW exploiting a $(k+1)$ -coloring of the buyers in N' returns a solution of revenue at least $S/(k+1+\epsilon)$ (by Lemma 13) and k is constant, it is sufficient to asymptotically bound the expected value $\mathbb{E}[S]$ of S .

To this aim, by the linearity of expectation, we have that

$$\begin{aligned} \mathbb{E}[S] &= \mathbb{E}\left[\sum_{i \in N_{\text{opt}}} v_i(m_i) X_i\right] \\ &= \sum_{i \in N_{\text{opt}}} v_i(m_i) \cdot \mathbb{E}[X_i] \\ &= \sum_{i \in N_{\text{opt}}} v_i(m_i) \cdot \mathbb{P}(X_i = 1) \\ &\geq \sum_{i \in N_{\text{opt}}} v_i(m_i)/2 \\ &= \text{opt}_r(\mu, G)/2 \end{aligned}$$

thus proving the claim. \square

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 and single-minded valuations, it is possible to obtain a bound with high probability by pre-processing 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, called POWER-LAW-2:

- once drawn $G \in \mathcal{G}_{n,\bar{d}}$, order the buyers non-increasingly with

respect to their preferred valuations, and let P be the prefix of the first $l = 200 \ln n$ 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 $(X_{1,P}, \bar{p}_{1,P})$ be the resulting outcome;
- run POWER-LAW on the submarket containing only buyers in $N \setminus P$ and their induced subgraph $G_{N \setminus P}$, and assign bottom prices to all the remaining buyers; let $(X_{2,P}, \bar{p}_{2,P})$ be the corresponding outcome;
- return the best of the two outcomes.

Notice that $(X_{1,P}, \bar{p}_{1,P})$ 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, $(X_{1,P}, \bar{p}_{1,P})$ can be computed in such a way by probing all the possible subsets of P , whose number is polynomially bounded.⁵

We are now able to prove the following theorem.

Theorem 15. *POWER-LAW-2, when run on randomly drawn social graphs in $\mathcal{G}_{n,\bar{d}}$ and in case of unlimited supply, has probability at least $1 - 1/n$ of returning an outcome with revenue at least a constant fraction of the optimal one for (SINGLE,REVENUE)-pricing.*

Proof. We show that the revenue of the algorithm is $\Omega(\text{opt}_r(\mu, G))$ with probability $1 - 1/n$.

If the overall revenue contributed to $\text{opt}_r(\mu, G)$ by the prefix P of the first $l = 200 \ln n$ buyers is higher than the one contributed by the remaining buyers, then the algorithm returns a solution of revenue at least $\text{opt}_r(\mu, G)/2$.

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, $S = \sum_{i \in N \setminus P} v_i(m_i) X_i$.

Since the algorithm returns a solution of revenue at least $S/(k+1)$ (by Lemma 13) and k is constant, we equivalently show that $S = \Omega(\text{opt}_r(\mu, G))$ with high probability. To this aim, we show that $\mathbb{P}(S \leq \text{opt}_r(\mu, G)/8) \leq 1/n$, that is $\mathbb{P}(S > \text{opt}_r(\mu, G)/8) \geq 1 - 1/n$. Observe that $\mathbb{E}[S] = \mathbb{E}[\sum_{i \in N \setminus P} v_i(m_i) X_i] = \sum_{i \in N \setminus P} v_i(m_i) \cdot \mathbb{E}[X_i] = \sum_{i \in N \setminus P} v_i(m_i) \cdot \mathbb{P}(X_i = 1) \geq \sum_{i \in N \setminus P} v_i(m_i)/2 \geq \sum_{i \in N} v_i(m_i)/4 \geq \sum_{i \in N} \text{opt}_r(\mu, G)/4$. Thus, since $\mathbb{P}(S \leq \text{opt}_r(\mu, G)/8) \leq \mathbb{P}(S \leq \mathbb{E}[S]/2) = \mathbb{P}(\mathbb{E}[S] - S \geq \mathbb{E}[S]/2)$, we can equivalently show that $\mathbb{P}(\mathbb{E}[S] - S \geq \mathbb{E}[S]/2) \leq 1/n$. Observe that S can be defined as $S = \sum_{i=1}^{|N'|} Y_i$, where each $Y_i \geq 0$ is a random sample without replacement from

⁵ Observe that prefix l has been set equal to $c \cdot \ln n$, where c is a constant. We set c equal to 200 (that is a high constant) since our aim is just showing that the returned approximation factor is constant, regardless of its precise value. Anyway the analysis of the algorithm can be strengthened in such a way that the lower c , the lower the (constant) approximation factor.

population $\mathcal{Y} := (v_i(m_i) : i \in N \setminus P)$ ⁶. By resorting to Bernstein's inequality (refer to Proposition 1.4 of [8]), we get

$$\mathbb{P}\left(\mathbb{E}[S] - S \geq \frac{\mathbb{E}[S]}{2}\right) \leq \exp\left(-\frac{(\mathbb{E}[S]/2)^2/|N'|}{2\sigma^2 + 2b\mathbb{E}[S]/(3|N'|)}\right) \quad (1)$$

where σ^2 is the variance of population \mathcal{Y} and b is the maximum preferred valuation of each buyer in $N \setminus P$. Observe that $\sigma^2 \leq \frac{1}{|N \setminus P|} \sum_{i \in N \setminus P} v_i(m_i)^2$. Since the l buyers in P have the l -th highest preferred valuations, we have that $b \leq \sum_{i \in N} v_i(m_i)/l$. Furthermore, since quantity $F := \frac{1}{|N \setminus P|} \sum_{i \in N \setminus P} v_i(m_i)^2$ is proportional to the sum of the squares of the buyers' preferred valuations, we have that F is maximum when l buyers have the maximum possible preferred valuation (i.e., $\sum_{i \in N} v_i(m_i)/l$) and the remaining ones have null preferred valuation. Thus, $\sigma^2 \leq \left(\frac{l}{|N \setminus P|}\right) \left(\frac{\sum_{i \in N} v_i(m_i)}{l}\right)^2 = \frac{(\sum_{i \in N} v_i(m_i))^2}{|N \setminus P|l} \leq \frac{(\sum_{i \in N} v_i(m_i))^2}{|N'|l}$. By incorporating the previous bounds on b and σ^2 in (1), and since $\sum_{i \in N} v_i(m_i)/4 \leq \mathbb{E}[S] \leq \sum_{i \in N} v_i(m_i)$, we get:

$$\begin{aligned} \mathbb{P}\left(\mathbb{E}[S] - S \geq \frac{\mathbb{E}[S]}{2}\right) &\leq \exp\left(-\frac{(\mathbb{E}[S]/2)^2/|N'|}{2\sigma^2 + 2b\mathbb{E}[S]/(3|N'|)}\right) \\ &\leq \exp\left(-\frac{(\sum_{i \in N} v_i(m_i)/8)^2}{2|N'| \sigma^2 + 2b(\sum_{i \in N} v_i(m_i))/3}\right) \\ &\leq \exp\left(-\frac{(\sum_{i \in N} v_i(m_i)/8)^2}{2(\sum_{i \in N} v_i(m_i))^2/l + 2(\sum_{i \in N} v_i(m_i))^2/(3l)}\right) \\ &\leq e^{-l/200} = e^{-\ln(n)} = 1/n, \end{aligned}$$

and this shows the claim. \square

Unfortunately, the argument in the previous theorem does not 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, called POWER-LAW-3:

- for a fixed constant parameter w , once randomly drawn $G \in \mathcal{G}_{n,\vec{d}}$, for all the possible subsets $P \subseteq N$ of $l = 200w$ buyers:
 - find the optimal outcome $(X_{1,P}, \bar{p}_{1,P})$ of the submarket induced by the buyers in P , assigning bottom prices to to all the remaining buyers;
 - run POWER-LAW on the submarket containing only the buyers in $N \setminus P$ and their induced subgraph $G_{N \setminus P}$, assign bottom prices to to all the remaining buyers, and let $(X_{2,P}, \bar{p}_{2,P})$ be the resulting outcome;
- return the best among all the outcomes $(X_{1,P}, \bar{p}_{1,P})$ and $(X_{2,P}, \bar{p}_{2,P})$, over all the subsets P of l buyers.

Notice that, since w is constant, in the initial phase the number of all the considered subsets of $l = 200w$ buyers is 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 w . A proof similar to the one of Theorem 15 restricted to the buyers allocated in an optimal outcome shows the following theorem.

⁶ A population is a multiset of values.

Theorem 16. *POWER-LAW-3, when run on randomly drawn social graphs in $\mathcal{G}_{n,\vec{d}}$, returns a constant approximation for (SINGLE,REVENUE)-pricing with prob. at least $1 - e^{-w}$.*

Proof (sketch). We show that the revenue of the algorithm is $\Omega(\text{opt}_r(\mu, G))$ with probability $1 - e^{-w}$.

Let $l = 200w$. For a given optimal outcome, let N_{opt} be the set of buyers receiving a bundle, and let m_i be the bundle assigned to each buyer $i \in N_{\text{opt}}$. If $|N_{\text{opt}}| \leq l$ then the algorithm returns an optimal outcome. If $|N_{\text{opt}}| > l$, order the buyers $i \in N_{\text{opt}}$ non-increasingly with respect to valuation $v_i(m_i)$, and consider the prefix P of the first l buyers of N_{opt} .

If the overall revenue contributed to $\text{opt}_r(\mu, G)$ by the prefix P is higher than the one contributed by the remaining buyers, then the algorithm returns a solution of revenue at least $\text{opt}_r(\mu, G)/2$, since P is one of the subsets analyzed by the algorithm.

On the other hand, if such a contribution is lower, consider the subset $N_{\text{opt}} \setminus P$ of the remaining buyers not in the prefix. If $|N_{\text{opt}} \setminus P| \leq l$, then the algorithm has determined the optimal solution of some set P' of l buyers containing $N_{\text{opt}} \setminus P$, and thus the optimal solution restricted to buyers in P' guarantees a revenue of at least $\text{opt}_r(\mu, G)/2$.

If $|N_{\text{opt}} \setminus P| > l$, let k be the constant selected by the algorithm, i.e. such that the set $N' \subseteq N_{\text{opt}} \setminus P$ of the buyers of degree at most k has cardinality $|N'| \geq (|N_{\text{opt}}| - l)/2$. 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 $N_{\text{opt}} \setminus P$ of degree at most k in G , that is, $S = \sum_{i \in N_{\text{opt}} \setminus P} v_i(m_i) X_i$.

Since the algorithm returns a solution of revenue at least $S/(k + 1 + \epsilon)$ (by Lemma 13) and k is constant, we equivalently show that $S = \Omega(\text{opt}_r(\mu, G))$ with high probability. To this aim, we show that $\mathbb{P}(S \leq \text{opt}_r(\mu, G)/8) \leq e^{-w}$, that is $\mathbb{P}(S > \text{opt}_r(\mu, G)/8) \geq 1 - e^{-w}$. Then, by exploiting arguments similar to the ones in Theorem 15, we obtain the claim. \square

By exploiting the same proof as in Theorem 16, it is possible to show the following theorem for general valuations.

Theorem 17. *POWER-LAW-3, when run on randomly drawn social graphs in $\mathcal{G}_{n,\vec{d}}$, returns a constant approximation for (GENERAL,REVENUE)-pricing with prob. at least $1 - e^{-w}$.*

6 CONCLUSIONS

It would be interesting to close the logarithmic 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 the so-called pair- and social envy-freeness. Moreover, we adopted the basic item-pricing policy, but what about other relaxed forms of pricing, like bundle-pricing? 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.

Acknowledgements

This work was partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR ‘‘Algorithms, Games, and Digital Markets’’.

REFERENCES

- [1] R. Abebe, J. M. Kleinberg, and D. C. Parkes, 'Fair division via social comparison', in *16th Conf. on Autonomous Agents and MultiAgent Systems, AAMAS*, pp. 281–289, (2017).
- [2] G. Aggarwal and J. D. Hartline, 'Knapsack auctions', in *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pp. 1083–1092, (2006).
- [3] N. Alon, Y. Mansour, and M. Tennenholtz, 'Differential pricing with inequity aversion in social networks', in *14th ACM Conf. on Electronic Commerce, EC*, pp. 9–24, (2013).
- [4] G. Amanatidis, E. Markakis, and K. Sornat, 'Inequity Aversion Pricing over Social Networks: Approximation Algorithms and Hardness Results', in *41st Int. Symp. on Mathematical Foundations of Computer Science MFCS*, volume 58, pp. 9:1–9:13, (2016).
- [5] E. Anshelevich, K. Kar, and S. Sekar, 'Envy-free pricing in large markets: Approximating revenue and welfare', *ACM Transactions on Economics and Computation*, **5**(3), 16, (2017).
- [6] P. Austrin, S. Khot, and M. Safra, 'Inapproximability of vertex cover and independent set in bounded degree graphs', in *24th Conference on Computational Complexity, CCC*, pp. 74–80. IEEE, (2009).
- [7] M. Balcan, A. Blum, and Y. Mansour, 'Item pricing for revenue maximization', in *9th ACM Conf. on Electronic Commerce, EC*, pp. 50–59, (2008).
- [8] Rémi Bardenet and Odalric-Ambrym Maillard, 'Concentration inequalities for sampling without replacement', *Bernoulli*, **21**(3), 1361–1385, (2015).
- [9] V. Bilò, M. Flammini, and G. Monaco, 'Approximating the revenue maximization problem with sharp demands', *Theor. Comp. Science*, **662**, 9–30, (2017).
- [10] V. Bilò, M. Flammini, and L. Moscardelli, 'On the impact of buyers preselection in pricing problems', in *17th Conf. on Autonomous Agents and MultiAgent Systems, AAMAS*, (2018).
- [11] 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*, pp. 80:1–80:14, (2017).
- [12] P. Briest, 'Uniform budgets and the envy-free pricing problem', in *Automata, Languages and Programming, 35th Int. Colloq., ICALP*, pp. 808–819, (2008).
- [13] P. Briest and P. Krysta, 'Single-minded unlimited supply pricing on sparse instances', in *17th ACM-SIAM Symp. on Discr. Alg., SODA*, pp. 1093–1102, (2006).
- [14] R. L. Brooks, 'On colouring the nodes of a network', *Mathematical Proceedings of the Cambridge Philosophical Society*, **37**(2), 194–197, (1941).
- [15] S. Chakrabarti, *Mining the Web: Discovering Knowledge from Hypertext Data*, Morgan-Kaufman, 2002.
- [16] P. Chalermsook, J. Chuzhoy, S. Kannan, and S. Khanna, 'Improved hardness results for profit maximization pricing problems with unlimited supply', in *Approx., Randomization, and Combinatorial Optimization. Algorithms and Techniques*, 73–84, (2012).
- [17] P. Chalermsook, B. Laekhanukit, and D. Nanongkai, 'Graph products revisited: Tight approximation hardness of induced matching, poset dimension and more', in *24th ACM-SIAM Symp. on Discr. Alg., SODA*, pp. 1557–1576, (2013).
- [18] P. Chalermsook, B. Laekhanukit, and D. Nanongkai, 'Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses', in *54th IEEE Symp. on Foundations of Computer Science, FOCS*, pp. 370–379, (2013).
- [19] N. Chen and X. Deng, 'Envy-free pricing in multi-item markets', in *Automata, Languages and Programming, 37th Int. Colloq., ICALP*, pp. 418–429, (2010).
- [20] N. Chen, X. Deng, P. W. Goldberg, and J. Zhang, 'On revenue maximization with sharp multi-unit demands', *J. of Combinatorial Optimization*, **31**(3), 1174–1205, (2016).
- [21] N. Chen, A. Ghosh, and S. Vassilvitskii, 'Optimal envy-free pricing with metric substitutability', *SIAM J. on Computing*, **40**(3), 623–645, (2011).
- [22] M. Cheung and C. Swamy, 'Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply', in *49th IEEE Symp. on Foundations of Computer Science, FOCS*, pp. 35–44, (2008).
- [23] Y. Chevaleyre, U. Endriss, S. Estivie, N. Maudet, et al., 'Reaching envy-free states in distributed negotiation settings', in *20th Int. Joint Conf. on Artificial Intell., IJCAI*, volume 7, pp. 1239–1244, (2007).
- [24] E. D. Demaine, U. Feige, M. T. Hajiaghayi, and M. R. Salavatipour, 'Combination can be hard: Approximability of the unique coverage problem', *SIAM J. on Computing*, **38**(4), 1464–1483, (2008).
- [25] Shahar Dobzinski and Noam Nisan, 'Mechanisms for multi-unit auctions', *J. Artif. Intell. Res.*, **37**, 85–98, (2010).
- [26] David A. Easley and Jon M. Kleinberg, *Networks, Crowds, and Markets - Reasoning About a Highly Connected World*, Cambridge University Press, 2010.
- [27] M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski, 'Revenue maximizing envy-free multi-unit auctions with budgets', in *ACM Conf. on Electronic Commerce, EC*, pp. 532–549, (2012).
- [28] M. Flammini, M. Mauro, and M. Tonelli, 'On fair price discrimination in multi-unit markets', in *27th Int. Joint Conf. on Artificial Intell., IJCAI*, pp. 247–253, (2018).
- [29] M. Flammini, M. Mauro, and M. Tonelli, 'On social envy-freeness in multi-unit markets', *Artificial Intelligence*, **269**, 1–26, (2019).
- [30] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry, 'On profit-maximizing envy-free pricing', in *16th ACM-SIAM Symp. on Discr. Alg., SODA*, pp. 1164–1173, (2005).
- [31] J. D. Hartline and Q. Yan, 'Envy, truth, and profit', in *12th ACM Conf. on Electronic Commerce, EC*, pp. 243–252, (2011).
- [32] S. Khot, 'On the power of unique 2-prover 1-round games', in *34th ACM Symp. on Theory of Computing, STOC*, pp. 767–775, (2002).
- [33] 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).
- [34] Jean-Jacques Laffont, Patrick Rey, and Jean Tirole, 'Network competition: II. price discrimination', *The RAND J. of Economics*, **29**(1), 38–56, (1998).
- [35] G. Monaco, P. Sankowski, and Q. Zhang, 'Revenue maximization envy-free pricing for homogeneous resources', in *24th Int. Joint Conf. on Artificial Intell., IJCAI*, pp. 90–96, (2015).
- [36] L. Philips, *The Economics of Price Discrimination*, Cambridge University Press, 1983.
- [37] A. Pigou, *The Economics of Welfare*, Palgrave Macmillan UK, 1920.
- [38] N. L. Stokey, 'Intertemporal price discrimination', *Quarterly J. of Economics*, **93**(3), 355–371, (1979).
- [39] H. R. Varian, 'Chapter 10 price discrimination', volume 1 of *Handbook of Industrial Organization*, 597 – 654, Elsevier, (1989).