On Social Envy-Freeness
in Multi-Unit Markets

Michele Flammini,[†‡] **Manuel Mauro,**[†] **Matteo Tonelli**[†]

[†]Gran Sasso Science Institute, L'Aquila, Italy
[‡]University of L'Aquila, L'Aquila, Italy
michele.flammini@univaq.it, manuel.mauro@gssi.it, matteo.tonelli@gssi.it

Abstract

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 a same item being sold and each buyer is assigned a set of identical items. We consider the four different cases arising by considering different buyers valuations, i.e., single-minded or general, and by adopting 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, we provide an optimal allocation algorithm for general valuations with item-pricing, under the assumption of social graphs of bounded treewidth. 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 relationships. Some of our results close hardness open questions or improve already known ones in the literature concerning the classical setting without sociality.

Introduction

The choice of the prices a firm has to set for goods or services put on sale is a non-trivial issue that has to be faced in setting up a business. One of the main aims of the sellers is to find pricing policies that maximize their revenue, while guaranteeing good levels of customer satisfaction. In this respect, the fair allocation of goods, resources or services is a crucial matter, especially when selling items available in a limited amount, like for instance in bandwidth allocation. In such a setting, one of the most used concepts to describe fairness is the so called *envy-freeness*. This notion was already introduced in the second half of the past century (Foley 1967; Varian 1974), and can be intended in different ways. In the field of combinatorial auctions, two basic versions have been mostly considered:

- An *envy-free* allocation is an allocation in which each buyer receives a bundle of goods among the ones that maximize her utility.

- A *pair envy-free* allocation is an allocation in which no buyer finds convenient to switch the bundle she receives with the one of any other buyer.

Several works in the literature addressed the problem of pricing and *envy-free* allocations (Guruswami et al. 2005; Briest 2008; Chen and Deng 2010; Hartline and Koltun 2005; Brânzei et al. 2016). Similarly, different papers considered the notion of *pair envy-freeness* (Colini-Baldeschi et al. 2014; Feldman et al. 2012; Fiat and Wingarten 2009; Monaco, Sankowski, and Zhang 2015).

As future market scenarios are predicted to become more and more decentralized and pervasive, new and more realistic constraints need to be introduced in our models. A major concern is the assumption of full information about the context possessed by the buyers, like the awareness of the existence of all the other buyers and, more important, of their allocations. In fact, when dealing with online, highly dynamic and distributed environments, such a global knowledge might be unfeasible. The issue of modeling the locality of mutual influences in game theory was already considered in graphical games (Kearns, Littman, and Singh 2001) and explicitly taken into account in (Bilò et al. 2010; 2011), 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, several works in fair allocations of goods in absence of pricing assumed an individuals view of the subjective well-being as based on a comparison with peers, that is restricting (pair) envy-freeness constraints to social neighbors (Abebe, Kleinberg, and Parkes 2017; Chevaleyre, Endriss, and Maudet 2007). 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 (Abebe, Kleinberg, and Parkes 2017).

Our contribution

We focus on the notion of sociality in the pricing problem. 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, edges mutual knowledge between the corresponding endpoints, and each buyer can only be envious of the bundles received by her neighbors. We can 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:

envy-free $\subseteq$ *pair envy-free* $\subseteq$ *social envy-free*.

Besides investigating the time complexity of determining social envy-free revenue maximizing outcomes and good approximate solutions, we investigate the increase of revenue due to the incomplete knowledge of buyers by providing proper bounds on the price of envy-freeness, 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 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. 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 put on sale are homogeneous, like for example commodity markets. We consider two different hypotheses 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. For what concerns pricing, we consider two kinds of non-discriminatory pricing policies: *item-pricing*, where a unique price p equal for the all identical items must be set, and *bundle-pricing*, where the seller is allowed to assign different non-proportional price for each bundle size.

For all the four different arising cases we show the hardness of the revenue maximization problem and determine corresponding approximation results (see Table 1). All our approximation bounds for single-minded valuations are optimal. For general valuations, in case of 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 (Monaco, Sankowski, and Zhang 2015) for pair envy-freeness directly extends to social envy-freeness. As in (Briest 2008; Monaco, Sankowski, and Zhang 2015), our hardness result relies on a weaker conjecture with respect to $P \neq NP$, i.e., on the $R3SAT$-hardness of the problem (see (Feige 2002)). Moreover, we give an optimal allocation algorithm for bounded-treewidth social graphs. For general valuations under bundle-pricing, while a polylogarithmic lower bound on the achievable approximation ratio was already known (Monaco, Sankowski, and Zhang 2015), we give an $O(\log n)$-approximation algorithm for social (and thus pair) envy-freeness, thus improving upon

the previous $O(\log n \cdot \log m)$ bound for pair envy-freeness given in (Monaco, Sankowski, and Zhang 2015). Finally, for all the cases we provide optimal bounds on the price of envy-freeness (Table 2). Some of our results close hardness open questions or improve already known approximation ones concerning the classical pair envy-freeness setting.

	Single-minded		General Valuations	
	Classical	Social	Classical	Social
Item-pricing	**NP-hard**	**NP-hard (strong)**	$\Omega(\log^\epsilon n)$	
	FPTAS	**PTAS**	$O(\log n)$	
Bundle-pricing	**NP-hard**		$\Omega(\log^\epsilon n)$	
	FPTAS		~~$O(\log n \log m)$~~ **O(log n)**	

Table 1: General hardness and approximation results.

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

Table 2: The price of envy-freeness bounds.

Related Works

There is an extensive literature concerning envy-free pricing. For the revenue maximization, (Guruswami et al. 2005; Hartline and Yan 2011; Cheung and Swamy 2008; Balcan, Blum, and Mansour 2008; Briest and Krysta 2006) designed logarithmic approximation algorithms for various special cases of the problem. Related hardness results were given by (Briest 2008; Chalermsook et al. 2012; Chalermsook, Laekhanukit, and Nanongkai 2013b; 2013a; Demaine et al. 2008). Further variants were considered by (Chen and Deng 2010; Chen, Ghosh, and Vassilvitskii 2011; Feldman et al. 2012; Anshelevich, Kar, and Sekar 2015; Bilò, Flammini, and Monaco 2017; Chen et al. 2016).

In (Monaco, Sankowski, and Zhang 2015) authors gave hardness and approximation results on revenue maximization in markets with multi-unit items under both notion of *envy-freeness* and *pair envy-freeness* when allowing both item- and bundle-pricing. We remark that for single-minded buyers they admitted the free disposal feature, in which buyers have the same valuation for all the bundles of size at least equal to their preferred one. However, our results for single-minded buyers and pair envy-freeness are related to the canonical version in which free disposal is not admitted.

The price of envy-freeness was defined in (Caragiannis et al. 2009) in the context of fair allocation of divisible and indivisible goods without pricing and suitably bounded under different assumptions. Such a notion was extended in (Abebe, Kleinberg, and Parkes 2017) to the social case, that is when envy-freeness constraints must be satisfied only between neighbors in the social graph.

Distributed mechanisms for allocating indivisible goods under the absence of central control where investigated in (Chevaleyre, Endriss, and Maudet 2007; Chevaleyre et al. 2007; Chevaleyre, Endriss, and Maudet 2017). 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 for such distributed mechanisms when used as fair division procedures.

Finally, (Bilò, Flammini, and Monaco 2017) considered the possibility of limiting the view of some buyers, not admitting them in the market. This might be seen as a related to a social envy-freeness setting in which such buyers are isolated in the social graph.

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.
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 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 classical notion of pair envy-freeness defined in the literature.

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

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. 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 $c(\mathcal{M}, \mathcal{G}) = \sup_{\mu \in \mathcal{M}, G \in \mathcal{G}} \frac{opt(\mu)}{opt(\mu, G)}$.

For the sake of brevity, in the following we call (SINGLE,ITEM)-pricing (resp. (GENERAL,ITEM)-, (SINGLE,BUNDLE)- and (GENERAL,BUNDLE)-pricing) the classical 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 classical problem corresponds to the restriction to complete social graphs, every hardness result concerning the classical problem extends to the social version, while every approximation algorithm for the social problem also applies to the classical 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 $\{O_1, \ldots, O_t\}$ 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 at most one object from each class so as to maximize the sum of the profits of the selected objects without exceeding the knapsack capacity k. As shown in (Lawler 1979), MULTIPLE-CHOICE KNAPSACK is NP-hard, but it admits an FPTAS.

Before concluding the section let us remark 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.

Single-minded valuations

In this section we consider single-minded multi-unit markets. Let us first focus on item-pricing. The following fact will be useful in the sequel.

Lemma 3. *The price p^{opt} of an optimal stable item-pricing outcome $(\overline{X}^{opt}, p^{opt})$, that is maximizing the seller's revenue, belongs to the set $\mathbb{P} := \{\frac{v_i(m_i)}{m_i} | i \in N\}$.*

The following positive result concerns the classical case.

Theorem 4. *The (SINGLE,ITEM)-pricing problem is NP-hard but admits an FPTAS.*

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

Theorem 5. *The social (SINGLE,ITEM)-pricing problem is strongly NP-hard.*

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 we associate 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 $v_w(|V|(|F|+1)) = |V|(|F|+1)$; finally, 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 $u \in V$, and an edge $\{i, i_e\} \in E$ for every $i \in N_u$ and edge $e \in F$ incident to u in H.

It is possible to show that H admits a subgraph of size k with h edges if and only if μ has an outcome stable under G of revenue $(|V| + k)(|F| + 1) + h$. $\qquad\square$

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 involved ideas. 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 the 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 resorting on a proper instance $K(\mu, G, p)$ of MULTIPLE-CHOICE KNAPSACK, where 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 $G\}$, and $z_{j,h} = w_{j,h} = j \cdot h$.

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 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 5, 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 strong NP-hard problem). However, the following approximation result holds.

Lemma 6. *Given an instance (μ, G) of the (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.*

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

Proof. Consider the algorithm that, for all the candidate optimal prices $p \in \mathbb{P}$ established in Lemma 3, constructs a corresponding instance $K_{\varepsilon/2}(\mu, G, p)$ exploiting the PTAS of Lemma 6 with accuracy parameter $\varepsilon/2$, and then 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 $r(\overline{X}, p) = p \cdot (1 - \varepsilon/2)opt(K_{\varepsilon/2}) \geq (1 - \varepsilon/2)^2 opt(K) \geq (1 - \varepsilon)opt(\mu, G)$. $\qquad\square$

We now focus on bundle-pricing. The same reduction of Theorem 4 shows the following negative result.

Theorem 8. *The* (SINGLE,BUNDLE)-*pricing problem is NP-hard.*

However, an optimal approximation can be achieved.

Theorem 9. *The social* (SINGLE,BUNDLE)-*pricing admits an FPTAS.*

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 preferred player. Therefore, again the problem reduces to a proper instance of MULTIPLE-CHOICE KNAPSACK 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. $\square$

Notice that the algorithm of Theorem 9 is independent of the social graph G, as the returned allocation is stable under any graph. In fact, as shown in the following theorem, sociality in (SINGLE,BUNDLE)-pricing does not affect the quality of the optimal stable outcomes.

Theorem 10. *The price of envy-freeness is* 2 *for* (SINGLE,ITEM)-*pricing and* 1 *for* (SINGLE,BUNDLE)-*pricing.*

General valuations

In this section we consider multi-unit markets with general valuations. Again, we first focus on item-pricing. For such a case, the tractability in the classical setting was an open problem raised in (Monaco, Sankowski, and Zhang 2015). However, we are now able to answer this question by showing the following approximation hardness result.

Theorem 11. *Approximating* (GENERAL,ITEM)-*pricing within* $O(\log^\varepsilon n)$, *for some* $\varepsilon > 0$, *is R3SAT-hard.*

Proof. In order to prove the claim, we give an approximation preserving polynomial time reduction from MES (Maximum expanding sequence) (Briest 2008). In such a problem, we have a universe set U and an ordered collection of its subsets $\mathcal{C} = (S_1, S_2, \ldots, S_m)$. An expanding sequence $\phi = (\phi(1) < \ldots < \phi(\ell))$ of length $|\phi| = \ell$ is a selection of sets $S_{\phi(1)}, \ldots, S_{\phi(\ell)}$, such that for each y, $1 \leq y \leq \ell$, $S_{\phi(y)} \nsubseteq \bigcup_{l=1}^{y-1} S_{\phi(l)}$. We want to find the expanding sequence of maximum length.

An instance of MES is said κ-separable if the sequence of the subsets $\mathcal{C}$ can be partitioned in the order into κ subsequences or *classes* $\mathcal{C}_1, \ldots, \mathcal{C}_\kappa$, where each subsequence does not contain intersecting sets. In (Briest 2008) it has been shown that there exists an $\varepsilon > 0$ such that MES is $R3SAT$-hard to approximate within $O(f(m)^\varepsilon)$, when restricted to $f(m)$-separable instances.

Consider the following reduction. Given a κ-separable instance of MES with subsets $S_1, \ldots, S_m \subseteq U$ of corresponding classes $\mathcal{C}_1, \ldots, \mathcal{C}_\kappa$, we construct an instance μ of (GENERAL,ITEM)-pricing as follows: for each $o \in U$ and $\mathcal{C}_k$, $1 \leq k \leq \kappa$, let $B_o^k = \{ 2^{\kappa-h}|U| + o \mid h \in \mathbb{N}, \ k \leq h \leq \kappa \}$; 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 $v_i(j) = 2^{\kappa-k}|U| + |U|$ if $j \in \bigcup_{o \in S_y} B_o^k$, while $v_i(j) = 0$ otherwise.

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 $m \cdot 2^\kappa \cdot 2^\kappa \cdot |U|$.

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 not any bundle size with 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 resort on the following lemmata.

Lemma 12. *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* $\frac{r}{4}$.

Lemma 13. *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* $2^{\kappa-k}|U| + o$ ".

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

$(\Rightarrow)$ Suppose that the MES instance admits an expanding sequence $S_{\phi(1)}, \ldots, S_{\phi(\ell)}$ 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.

($\Leftarrow$) Suppose that the reduced instance admits a stable outcome $(\overline{X}, p)$ of revenue $\ell 2^{\kappa}|U|$. Then, by Lemma 12 and Lemma 13, 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, $1 \le h \le k$. 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 $|\mathcal{I}^{max}| \ge \frac{\ell}{16}$.

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

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

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

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 $f(m) = \log(m)$, which completes the proof. $\qquad\square$

For what concerns the determination of approximated solutions, we notice that the $O(\log n)$-approximation algorithm of (Monaco, Sankowski, and Zhang 2015) for the classical 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 14. *The social* (GENERAL,ITEM)-*pricing problem restricted to social graphs of bounded tree-width admits an optimal polynomial time algorithm.*

Proof. We 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 fixed a bundle size j with $0 \le j \le m$ and a supply bound b with $0 \le b \le m$, 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 $p \cdot M_p^i(j, b)$, and let us use symbol $\perp$ to denote unfeasibility, that is the fact that under such restrictions a stable outcome for T does not exist.

Then, if i is a leaf, $M_p^i(j, b) = j$ if $v_i(j) - jp \ge 0$ and $j \le b$, while $M_p^i(j, b) = \perp$ otherwise.

If i is not a leaf, $M_p^i(j, b)$ can be recursively constructed optimally combining the optimal restricted outcomes for its subtrees in T that do not make i and its children envious and globally satisfy the supply constraint. Such a problem can be formulated as an instance of MULTIPLE CHOICE KNAPSACK, in which 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 $v_i(j) - jp \ge v_i(j') - j'p$, $v_{i'}(j') - j'p \ge v_{i'}(j) - jp$, $b' \le b - j$ and $M_p^{i'}(j', b') \ne \perp$.

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 $\text{OPT}_p^i(j, b)$ be its value, otherwise we set $\text{OPT}_p^i(j, b) = \perp$.

We can then compute $M_p^i(j, b)$ for an intermediate node i as $M_p^i(j, b) = \perp$ if $v_i(j) \le jp$ or $\text{OPT}_p^i(j, b) = \perp$, while

$M_p^i(j, b) = j + \text{OPT}_p^i(j, b)$ otherwise.

The maximum number of allocated items, given price p, can then be expressed as $M_p^* = \max_{j \le m} M_p^r(j, m)$

Exploiting the above formulas, we can compute in polynomial time M_p^*. More precisely, for each node i, bundle j and $b \le m$, 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 (Monaco, Sankowski, and Zhang 2015), 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 extension to bounded treewidth graphs will be given in the full version of the paper. □

Let us now consider bundle-pricing. For such a case, an inapproximability $\Omega(\log^\epsilon n)$ result for some $\epsilon > 0$ in the classical setting was proven in (Monaco, Sankowski, and Zhang 2015), together with a $O(\log n \log m)$-approximation algorithm. However, we are able to prove the following nearly optimal approximation bound, holding also for the more general social instances.

Theorem 15. *There exists a $\frac{\log n}{1 - e^{-1}}$-approximation algorithm for the social* (GENERAL,BUNDLE)*-pricing problem.*

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 $b_{j,l}$ 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 the algorithm presented in (Guruswami et al. 2005), which guarantees a revenue which is at least $z(N, B)/\log n$, 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 $\sum_{b_{j,l} \in B'} w(\{i, b_{j,l}\}) \le m$ 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 $\sum_{b_{j,l} \in B*} w(\{i, b_{j,l}\}) \le m$ and $z(N, B^*)$ is maximized, B' has the property that $z(N, B') \ge (1 - e^{-1})z(N, B^*)$. Then, by applying the algorithm of (Guruswami et al. 2005) 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 $z(N, B')/\log n \ge (1 - e^{-1})z(N, B^*)/\log n \ge (1 - e^{-1})opt(\mu, G)/\log n$, 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', i.e., that $z(N, B') \le z(N, B'')$ if $B' \subseteq B''$ and $Z(N, B') + Z(N, B'') \ge Z(N, B' \cup B'') + Z(N, B' \cap B'')$ 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 (Sviridenko 2004). □

We conclude the section providing optimal asymptotic bounds for the price of envy-freeness for general valuations.

Theorem 16. *The price of envy-freeness is $\Theta(\log n)$ both for* (GENERAL,ITEM)*-pricing and* (GENERAL,BUNDLE)*-pricing.*

Conclusions and Future Work

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. Moreover, both for item- and bundle-pricing, it would be worth investigating other relevant restricted families of social graphs, such as bounded-degree ones. Another nice research direction is that of considering other market scenarios, like unit-demand markets, or obtained by properly restricting the classes of valuation functions. Finally, it would be nice to consider other forms of social influence, or determining how limiting buyers visibility might increase the achievable revenue.

References

Abebe, R.; Kleinberg, J. M.; and Parkes, D. C. 2017. Fair division via social comparison. In *16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 281–289.

Anshelevich, E.; Kar, K.; and Sekar, S. 2015. Envy-free pricing in large markets: Approximating revenue and welfare. In *42nd International Colloquium on Automata, Languages and Programming (ICALP)*, 52–64.

Balcan, M.; Blum, A.; and Mansour, Y. 2008. Item pricing for revenue maximization. In *9th ACM Conference on Electronic Commerce (EC)*, 50–59.

Bilò, V.; Fanelli, A.; Flammini, M.; and Moscardelli, L. 2010. When ignorance helps: Graphical multicast cost sharing games. *Theoretical Computer Science* 411(3):660–671.

Bilò, V.; Fanelli, A.; Flammini, M.; and Moscardelli, L. 2011. Graphical congestion games. *Algorithmica* 61(2):274–297.

Bilò, V.; Flammini, M.; and Monaco, G. 2017. Approximating the revenue maximization problem with sharp demands. *Theoretical Computer Science* 662:9–30.

Brânzei, S.; Filos-Ratsikas, A.; Miltersen, P. B.; and Zeng, Y. 2016. Envy-free pricing in multi-unit markets. *CoRR* abs/1602.08719.

Briest, P., and Krysta, P. 2006. Single-minded unlimited supply pricing on sparse instances. In *17th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1093–1102.

Briest, P. 2008. Uniform budgets and the envy-free pricing problem. In *35th International Colloquium on Automata, Languages and Programming (ICALP)*, 808–819.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2009. The efficiency of fair division. In *5th International Workshop on Internet and Network Economics (WINE)*, 475–482.

Chalermsook, P.; Chuzhoy, J.; Kannan, S.; and Khanna, S. 2012. Improved hardness results for profit maximization pricing problems with unlimited supply. In *15th International Workshop on Approximation, Randomization, and Combinatorial Optimization (APPROX)*, 73–84.

Chalermsook, P.; Laekhanukit, B.; and Nanongkai, D. 2013a. Graph products revisited: Tight approximation hardness of induced matching, poset dimension and more. In *24th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1557–1576.

Chalermsook, P.; Laekhanukit, B.; and Nanongkai, D. 2013b. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *54th IEEE Symposium on Foundations of Computer Science (FOCS)*, 370–379.

Chen, N., and Deng, X. 2010. Envy-free pricing in multi-item markets. In *37th International Colloquium on Automata, Languages and Programming (ICALP)*, 418–429.

Chen, N.; Deng, X.; Goldberg, P. W.; and Zhang, J. 2016. On revenue maximization with sharp multi-unit demands. *Journal of Combinatorial Optimization* 31(3):1174–1205.

Chen, N.; Ghosh, A.; and Vassilvitskii, S. 2011. Optimal envy-free pricing with metric substitutability. *SIAM Journal on Computing* 40(3):623–645.

Cheung, M., and Swamy, C. 2008. Approximation algorithms for single-minded envy-free profit-maximization

problems with limited supply. In *49th IEEE Symposium on Foundations of Computer Science (FOCS)*, 35–44.

Chevaleyre, Y.; Endriss, U.; Estivie, S.; and Maudet, N. 2007. Reaching envy-free states in distributed negotiation settings. In *20th International Joint Conference on Artificial Intelligence (IJCAI)*, 1239–1244.

Chevaleyre, Y.; Endriss, U.; and Maudet, N. 2007. Allocating goods on a graph to eliminate envy. In *22nd AAAI Conference on Artificial Intelligence*, 700–705.

Chevaleyre, Y.; Endriss, U.; and Maudet, N. 2017. Distributed fair allocation of indivisible goods. *Artificial Intelligence* 242:1–22.

Colini-Baldeschi, R.; Leonardi, S.; Sankowski, P.; and Zhang, Q. 2014. Revenue maximizing envy-free fixed-price auctions with budgets. In *10th International Conference on Web and Internet Economics (WINE)*, 233–246.

Demaine, E. D.; Feige, U.; Hajiaghayi, M.; and Salavatipour, M. R. 2008. Combination can be hard: Approximability of the unique coverage problem. *SIAM Journal on Computing* 38(4):1464–1483.

Feige, U. 2002. Relations between average case complexity and approximation complexity. In *34th ACM Symposium on Theory of Computing (STOC)*, 534–543.

Feldman, M.; Fiat, A.; Leonardi, S.; and Sankowski, P. 2012. Revenue maximizing envy-free multi-unit auctions with budgets. In *12th ACM Conference on Electronic Commerce (EC)*, 532–549.

Fiat, A., and Wingarten, A. 2009. Envy, multi envy, and revenue maximization. In *5th International Workshop on Internet and Network Economics (WINE)*, 498–504.

Foley, D. 1967. Resource allocation and the public sector. *Yale Economic Essays* 7(1):45–98.

Guruswami, V.; Hartline, J. D.; Karlin, A. R.; Kempe, D.; Kenyon, C.; and McSherry, F. 2005. On profit-maximizing envy-free pricing. In *16th ACM-SIAM Symposium on Discrete Algorithm (SODA)*, 1164–1173.

Hartline, J. D., and Koltun, V. 2005. Near-optimal pricing in near-linear time. In *9th International Workshop on Algorithm and Data Structures (WADS)*, 422–431.

Hartline, J. D., and Yan, Q. 2011. Envy, truth, and profit. In *Conference on Electronic Commerce (EC)*, 243–252. ACM.

Kearns, M. J.; Littman, M. L.; and Singh, S. P. 2001. Graphical models for game theory. In *17th Conference in Uncertainty in Artificial Intelligence (UAI)*, 253–260.

Lawler, E. L. 1979. Fast approximation algorithms for knapsack problems. *Math. Oper. Res.* 4(4):339–356.

Monaco, G.; Sankowski, P.; and Zhang, Q. 2015. Revenue maximization envy-free pricing for homogeneous resources. In *24th International Joint Conference on Artificial Intelligence (IJCAI)*, 90–96.

Sviridenko, M. 2004. A note on maximizing a submodular set function subject to a knapsack constraint. *Operation Research Letters* 32(1):41–43.

Varian, H. R. 1974. Equity, envy, and efficiency. *Journal of economic theory* 9(1):63–91.