Pricing Problems with Buyer Preselection

Vittorio Bilò
Univ. of Salento, Lecce, Italy
vittorio.bilo@unisalento.it

Michele Flammini
GSSI Institute, L'Aquila, Italy & Univ. of L'Aquila, L'Aquila, Italy
michele.flammini@univaq.it

Gianpiero Monaco
Univ. of L'Aquila, L'Aquila Italy
gianpiero.monaco@univaq.it

Luca Moscardelli
Univ. of Chieti-Pescara, Pescara, Italy
luca.moscardelli@unich.it

Abstract

We investigate the problem of preselecting a subset of buyers participating in a market so as to optimize the performance of stable outcomes. We consider four scenarios arising from the combination of two stability notions, item and bundle envy-freeness, with the two classical objective functions, i.e., the social welfare and the seller's revenue. When adopting the notion of item envy-freeness, we prove that, for both the two objective functions, the problem cannot be approximated within $n^{1-\varepsilon}$ for any $\varepsilon > 0$, and provide tight or nearly tight approximation algorithms. We also prove that maximizing the seller's revenue is NP-hard even for a single buyer, thus closing an open question. Under bundle envy-freeness, instead, we show how to transform in polynomial time any stable outcome for a market involving only a subset of buyers to a stable one for the whole market without worsening its performance, both for the social welfare and the seller's revenue. Finally, we consider multi-unit markets, where all items are of the same type and are assigned the same price. For this specific case, we show that buyer preselection can improve the performance of stable outcomes in all of the four considered scenarios, and we design corresponding approximation algorithms.

2012 ACM Subject Classification Theory of computation → Computational pricing and auctions

Keywords and phrases Pricing problems, Envy-freeness, Revenue maximization, Social Welfare maximization

Digital Object Identifier 10.4230/LIPIcs.MFCS.2018.47

Related Version A two-page extended abstract [8] of a preliminary version of this paper appeared in the proceedings of AAMAS 2018.

1 Introduction

Determining an efficient pricing strategy is a fundamental problem in many business activities, as it affects both the seller's revenue and the customers' or buyers' satisfaction. Usually, optimal prices are the result of a challenging counterbalancing process: selecting low prices, for instance, may be profitable for the seller when it attracts considerably more customers, but, at the same time, in case of limited supply, it may leave some buyers unsatisfied, thus generating discontent. In particular, this happens when a customer is negated the right to buy her preferred set of items, or even any item at all, despite the fact that she is willing to pay for the posted prices. In this case she is often called *loser*, as opposed to a customer receiving items, called *winner*. For such a reason, pricing problems are traditionally considered under the hypothesis of **envy-freeness** [25, 32], which prescribes that, once a pricing strategy has been established, items have to be allocated to buyers in such a way that no one would prefer a different set of items.

However, if from the one hand safeguarding the losers' interests shelters the seller from possible future losses due to their dissatisfaction, on the other hand, a result by [7] shows that, in certain markets, an intrinsic and unavoidable hurdle to the construction of a good quality envy-free solution may come from the presence of a set of "disturbing" customers, that is, a set of buyers such that at least one of them gets envious in any assignment of sufficiently high revenue. This observation naturally leads to the following intriguing question: *"What happens if envy-freeness is restricted to apply only to the set of winners? Can the seller raise enough more revenue (with certainty) today to compensate the (uncertain) future loss of potential customers?"*. Such a relaxed form of envy-freeness models indeed the situation in which the seller is allowed the freedom to discard any subset of buyers from the given instance, so as to get rid of envious losers in the assignment she would like to propose. In addition, one might consider such buyers "preselection" as a situation in which the seller advertises about the existence of the market only targeted buyers that, once involved, will all consider the allocation fair; the excluded ones then are simply unaware and won't feel any unfairness.

1.1 Our Contribution

Motivated by the above discussion, we introduce and investigate the *buyer preselection problem* in which, given a pricing problem P with n buyers and m items, we are interested in computing the best possible envy-free solution that can be achieved by removing any arbitrary subset of buyers from P. We consider four scenarios arising from the combination of two stability notions, called item and bundle envy-freeness, respectively, with the two classical objective functions, namely, the social welfare and the seller's revenue.

In an item envy-free allocation, given a pricing of the items, each buyer gets the subset maximizing her utility among all possible subsets that can be created from the set of available items; in a bundle envy-free allocation, no buyer gets a better utility by receiving the bundle allocated to any winner. Observe that these allocations are always guaranteed to exist, as it suffices to assign all items an arbitrarily high price, so that no winner is possible.

For item envy-free allocations and both objective functions, we show that the buyer preselection problem cannot be approximated within $n^{1-\epsilon}$ for every $\epsilon > 0$, unless $\mathsf{P} = \mathsf{ZPP}$. On the positive side, under the objective of social welfare, we design an n-approximation algorithm, while, for the case of revenue maximization, we give an $O(n \log m)$-approximation. In particular, these results are obtained as follows: all but one buyer are discarded from the given instance, so that we are left with a pricing problem with a single buyer. While such a problem is solvable in polynomial time under the objective of social welfare, for revenue maximization, it already exhibits challenging combinatorial structures and, to the best of our knowledge, has been considered before only in [5]. In this paper an $O(\log m)$-approximation is provided, but no lower bounds on the problem complexity are given. We show that the problem is NP-hard, thus solving the corresponding longstanding open problem raised by the authors.

We stress that efficient preselection can be profitable under two orthogonal directions: from the one hand, the removal of a subset of pathological envious buyers may increase the value of the optimal solution; from the other hand, even when this does not happen or it has only a modest impact, simplifying the combinatorial structure defined by the valuation functions of the winners may lead to the design of better approximation algorithms. In fact, for the two preselection problems obtained by considering bundle envy-free allocations, we show how to transform in polynomial time any allocation which is bundle envy-free only for the subset of the winners to a bundle envy-free allocation for all buyers without worsening its performance. Hence, although this transformation implies that, in this case, buyer preselection cannot improve the performance of stable outcomes, it can be used to map any bundle envy-free allocation for a subset of winners obtained through preselection back to a bundle envy-free allocation involving all buyers. Therefore, it can be used as an algorithmic tool for computing good stable outcomes when preselection is not allowed. In fact, it can be first exploited to simplify the combinatorics of the problem, and then for mapping back the computed solution to one encompassing all the buyers.

To this respect, consider for instance the case in which buyers can be partitioned in two (or more) sets A and B (each of which, for instance, containing only unit-demand buyers or only single minded ones or only multi-demand buyers with additive valuations), and such that (i) valuations of players in A and B, if considered separately, allow to compute an r_A-approximate (resp. r_B-approximate) solution of the problem restricted to set A (resp. B) by exploiting known or new simpler algorithms/techniques and (ii) considering the whole instance would make difficult the direct application of such algorithms/technique for solving P. By performing preselection we can easily obtain an $r = 2\max(r_A, r_B)$-approximation for BP(P) (selecting the best solution among the one only for A and the one only for B), and thus, by the transformation, also for the initial problem P.

Finally, we consider the multi-unit case, where all items are of the same type and are assigned the same price. We show how preselection can improve the revenue and the social welfare of both item and bundle envy-free solutions. In particular, for item envy-free allocations, we show a tight multiplicative factor of m for both objective functions. For bundle envy-free allocations, we show a lower multiplicative bound of 2 for both the revenue and the social welfare, and prove that it is tight for the objective of revenue maximization. We also provide tight results on the complexity of computing optimal solutions for the buyer preselection problem under envy-freeness.

Due to space limitations, some proofs are only sketched or omitted.

1.2 Related Work

The literature on envy-free pricing is so vast that it cannot be exhaustively covered here. For such a reason, we simply refer to the achievements which are mostly related to the model of [28] we consider in this paper.

For the social welfare maximization, the VCG mechanism [33, 19, 26] provides an optimal solution to the envy-free pricing problem. However, while this mechanism is efficiently computable in markets with unit-demand buyers, yet for single-minded ones its computation becomes NP-hard. Approximate solutions are still possible in this case thanks to the results of [4]. Also Walrasian Equilibria [34] provide an optimal solution to the problem [6]; however, they are guaranteed to exist only under very stringent hypothesis on the buyers' valuation functions [27].

For the revenue maximization, [28, 29, 18, 5, 10] design logarithmic approximation algorithms for various special cases of the problem. Relative hardness results have been given by [9, 11, 13, 12, 20]. Further variants have been considered by [14, 16, 22, 3, 7, 15].

[23] propose an interesting relaxation of the notion of Walrasian Equilibrium, called Combinatorial Walrasian Equilibrium (CWE), obtained by grouping items into bundles so as to induce a "reduced market" to which, then, applying the notion of Walrasian Equilibrium. They show the existence of a CWE yielding a 2-approximation of the optimal social welfare and that of a CWE yielding a logarithmic approximation of the optimal revenue.

Furthermore, [31] study the case of revenue maximization in markets with multi-unit items under both item and bundle envy-freeness when allowing both item and bundle pricing. Such setting has been extended in [24] where the authors consider a social graph of the buyers and envies can arise only between neighbors.

In our model, by preselecting buyers, we basically require that all the winners are envy-free. Settings in which, in a similar way, envy-freeness is not guaranteed for all buyers, but only for the winners, are studied in [1, 2, 17]. In particular, [17] considers "weak" Walrasian equilibirum, a relaxed version of Walrasian equilibirum in which the goal is that of maximizing the number of envy-free buyers, with the condition that all the winners must be envy-free. [1, 2] consider a relaxed version of envy-freeness: in their model, identical items have to be sold to buyers, with every buyer constituting a node of a given unweighted graph; adjacent winning buyers have to pay similar prices for the received item, while the losers cannot envy. This feature is exploited to achieve higher revenue with respect to the classical case in which there cannot be losers, even if it makes the computational problem harder.

2 Model and Definitions

2.1 Markets

A *market* is a tuple $\Gamma = (N, M, (v_i)_{i \in N})$, where N is a set of n buyers, M is a set of m items, and for every buyer $i \in N$, $v_i : 2^M \to \mathbb{R}_{\geq 0}$ is a *valuation function* expressing, given a set of items $X \subseteq M$, the amount of money that buyer i is willing to pay for X; we assume that $v_i(\emptyset) = 0$ for every buyer $i \in N$.

Depending on the definition of the valuation functions, different types of markets can be modeled. In the most general case, called *market with combinatorial valuations*, function v_i is completely arbitrary for every buyer $i \in N$. In a *market with unit-demand buyers*, $v_i(X) = 0$ for every $i \in N$ and $X \subseteq M$ with $|X| > 1$, that is, every buyer is only interested in singleton sets. In a *market with single-minded buyers*, for every $i \in N$, there exists a unique set of items $X \subseteq M$ such that $v_i(X) \neq 0$, that is, every buyer is only interested in a particular set of items. In a *market with additive valuations*, $v_i(X) = \sum_{j \in X} v_i(\{j\})$ for every $i \in N$ and $X \subseteq M$. We stress that, while the representation of a market with combinatorial valuations may require $\Omega(n2^m)$ bits, $\Theta(nm)$ bits suffice to represent the last three types of markets. Finally, in a *market with multi-unit items*, all the m items are of the same type and so, for every $i \in N$, the valuation function becomes of the form $v_i : \{0, 1, \ldots, m\} \to \mathbb{R}_{\geq 0}$, since it is only required to specify how much a buyer evaluates a set of k items, for every $k \in \{1, \ldots, m\}$ (clearly, $v_i(0) = 0$); thus, also in this case, the market can be represented with $\Theta(nm)$ bits.

2.2 Stable Outcomes

Fix a market Γ. A *price vector* is an m-tuple $\mathbf{p} = (p_1, \ldots, p_m)$ such that, for every $j \in M$, $p_j \geq 0$ is the price of item j.[1] We denote by $\mathbf{0}^m$ the price vector assigning price 0 to all items. Given a price vector $\mathbf{p}$ and a set of items $X \subseteq M$, $u_i(X, \mathbf{p}) = v_i(X) - \sum_{j \in X} p_j$ is

[1] For the case of markets with multi-unit items, it is only required to fix the price of a single item so that vector $\mathbf{p}$ collapses to a real number $p \geq 0$.

the *utility* of buyer i when buying X. The *demand set* of buyer i for the price vector $\mathbf{p}$ is the set $D_i(\mathbf{p}) = \text{argmax}_{X \subseteq M} u_i(X, \mathbf{p})$ of subsets of items maximizing i's utility according to the prices specified by $\mathbf{p}$. An *allocation vector* is an n-tuple $\mathbf{X} = (X_1, \ldots, X_n)$ such that $X_i \subseteq M$ is the set of items sold to buyer i. The allocation vector $\mathbf{X} = (X_1, \ldots, X_n)$ is *feasible* if $X_i \cap X_{i'} = \emptyset$ for each $i \neq i' \in N$. An *outcome* is a pair $(\mathbf{X}, \mathbf{p})$ such that $\mathbf{X}$ is feasible. Denote with $\mathsf{OUT}(\Gamma)$ the set of outcomes of Γ. An outcome $(\mathbf{X}, \mathbf{p})$ is *individually-rational* if $u_i \geq 0$ for every $i \in N$.

Denote as $M(\mathbf{X}) = \bigcup_{i \in N} X_i$ the set of items sold to some buyer according to a feasible allocation vector $\mathbf{X}$. Buyer i is a *winner* if $X_i \neq \emptyset$ and $W(\mathbf{X})$ denotes the set of all winners in $\mathbf{X}$. For an item $j \in M(\mathbf{X})$, denote with $b_{\mathbf{X}}(j)$ the buyer $i \in W(\mathbf{X})$ such that $j \in X_i$. When the allocation vector is clear from the context, we simply write $b(j)$.

The following concepts define two types of stable outcomes for Γ.

▶ **Definition 1.** An individually-rational outcome $(\mathbf{X}, \mathbf{p})$ is **item envy-free** if $u_i(X_i, \mathbf{p}) \geq u_i(T, \mathbf{p})$ for every buyer $i \in N$ and $T \subseteq M$, that is, $X_i \in D_i(\mathbf{p})$ for every $i \in N$.

▶ **Definition 2.** An individually-rational outcome $(\mathbf{X}, \mathbf{p})$ is **bundle envy-free** if $u_i(X_i, \mathbf{p}) \geq u_i(X_j, \mathbf{p})$ for every two buyers $i, j \in N$.

Denote with $\mathsf{IEF}(\Gamma)$ and $\mathsf{BEF}(\Gamma)$, the sets of item envy-free and bundle envy-free outcomes for Γ, respectively. Notice that $\mathsf{IEF}(\Gamma) \subseteq \mathsf{BEF}(\Gamma)$ and $\mathsf{IEF}(\Gamma) \neq \emptyset$ (and so also $\mathsf{BEF}(\Gamma) \neq \emptyset$), since the outcome $(\mathbf{X}, \mathbf{p})$ such that $X_i = \emptyset$ for every $i \in N$ and $p_j = \infty$ for every $j \in M$ is individually-rational and item envy-free.

2.3 Pricing Problems

Fix an outcome $(\mathbf{X}, \mathbf{p})$ for a market Γ. The revenue raised by $(\mathbf{X}, \mathbf{p})$ is $\mathsf{REV}(\mathbf{X}, \mathbf{p}) = \sum_{j \in M(\mathbf{X})} p_j$. The social welfare generated by $(\mathbf{X}, \mathbf{p})$ is $\mathsf{SW}(\mathbf{X}, \mathbf{p}) = \sum_{i \in N} u_i(X_i) + \sum_{j \in M(\mathbf{X})} p_j = \sum_{i \in W(\mathbf{X})} v_i(X_i)$. Note that the social welfare does not depend on the price vector $\mathbf{p}$ and that $\mathsf{SW}(\mathbf{X}, \mathbf{p}) \geq \mathsf{REV}(\mathbf{X}, \mathbf{p})$.

Given a market Γ, let $\mathsf{sol}(\Gamma) \subseteq \mathsf{OUT}(\Gamma)$ denote any subset of outcomes for Γ and $\mathsf{obj} : \mathsf{OUT}(\Gamma) \to \mathbb{R}_{\geq 0}$ denote an objective function associating a non-negative value to every outcome for Γ. Let $opt(\Gamma, \mathsf{sol}, \mathsf{obj}) := \text{argmax}_{(\mathbf{X}, \mathbf{p}) \in \mathsf{sol}(\Gamma)} \{\mathsf{obj}(\mathbf{X}, \mathbf{p})\}$ denote the set of outcomes in $\mathsf{sol}(\Gamma)$ maximizing the objective function obj.

▶ **Definition 3.** The **pricing problem** $P = (\Gamma, \mathsf{sol}, \mathsf{obj})$ is an optimization problem which , given a market Γ, a set of outcomes $\mathsf{sol}(\Gamma)$ and an objective function obj, asks for an outcome $o^*(P) \in opt(P)$.

In this paper, we consider the cases in which $\mathsf{sol}(\Gamma) \in \{\mathsf{IEF}(\Gamma), \mathsf{BEF}(\Gamma)\}$ and $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$.

Oracles. Fix a pricing problem $P = (\Gamma, \mathsf{sol}, \mathsf{obj})$. As we have seen, when Γ is a market with combinatorial valuations, any algorithm for P needs to deal with an input of exponential size. In order to circumvent this problem and remain within the realm of polynomial time algorithms, it is usually assumed that functions v_is are not given as an input of the problem and are replaced by a polynomial time (with respect to n and m) *oracle* providing information about a buyer's valuation function. An oracle is usually assumed to answer two types of questions: a *value query* which, given a buyer $i \in N$ and a set of items X, returns the valuation $v_i(X)$, and a *demand query* which, given a price vector $\mathbf{p}$ and a buyer $i \in N$, returns any set in $D_i(\mathbf{p})$.

We remark that all algorithms of this paper exploiting oracle calls are polynomial also in the sense that they call the oracle a polynomial number of times. Therefore, when also the oracle is polynomially computable, the algorithm computation is fully polynomial.

2.4 The Buyer Preselection Problem

Given a market $\Gamma = (N, M, (v_i)_{i\in N})$ and a subset of buyers $N' \subseteq N$, the submarket of Γ induced by N' is the market $\Gamma(N') = (N', M, (v_i)_{i\in N'})$.

▶ **Definition 4.** The **buyer preselection problem** is an optimization problem $\mathsf{BP}(P)$ which, given a pricing problem $P = (\Gamma, \mathsf{sol}, \mathsf{obj})$ with $\Gamma = (N, M, (v_i)_{i\in N})$, asks for a pair $(N^*(\mathsf{BP}(P)), o^*(\mathsf{BP}(P)))$ such that $N^*(\mathsf{BP}(P)) \in \text{argmax}_{N'\subseteq N}\{opt(\Gamma(N'), \mathsf{sol}, \mathsf{obj})\}$ and $o^*(\mathsf{BP}(P)) \in opt(\Gamma(N^*(\mathsf{BP}(P))), \mathsf{sol}, \mathsf{obj})$, that is, $o^*(\mathsf{BP}(P))$ is the best outcome which can be realized in all possible submarkets of Γ.

Clearly, by definition, for every pricing problem P, $\mathsf{obj}(o^*(P)) \leq \mathsf{obj}(o^*(\mathsf{BP}(P)))$, that is, buyer preselection can only improve the quality of the optimal solution.

3 Results for Item Envy-Free Outcomes

In this section, we consider the buyer preselection problem $\mathsf{BP}(\Gamma, \mathsf{IEF}, \mathsf{obj})$ with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$. We start by providing a lower bound on its approximability attained by exploiting an approximation-preserving reduction from the maximum independent set problem.

▶ **Theorem 5.** *Let $P = (\Gamma, \mathsf{IEF}, \mathsf{obj})$ be a pricing problem with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$. For every $\epsilon > 0$, the buyer preselection problem $\mathsf{BP}(P)$ cannot be approximated within $n^{1-\epsilon}$, unless $\mathbf{P} = \mathbf{ZPP}$, even when Γ is a market with single-minded buyers.*

Proof. We prove the claim through an approximation-preserving reduction from the maximum independent set problem, in which, given an undirected graph, it is asked for a subset of nodes, no two of which are adjacent, of maximum cardinality. To this aim, consider an instance of the maximum independent set problem defined by a graph $G = (V, E)$ and denote with $\delta_i(G)$ the set of edges incident to node i. We create a market $\Gamma = (N, M, (v_i)_{i\in N})$ with single-minded buyers as follows: we set $N = V$, $M = E$ and, for every $i \in N$, we define the valuation function v_i in such a way that, for every $X \subseteq M$,

$$v_i(X) = \begin{cases} 1 & \text{if } X = \delta_i(G) \\ 0 & \text{otherwise.} \end{cases}$$

Fix any subset of buyers $N' \subseteq N$. Given an individually-rational outcome $o = (\mathbf{X}, \mathbf{p}) \in \mathsf{OUT}(\Gamma(N'))$, define $V_{\mathsf{REV}}(o) := \{i \in V : \sum_{j\in X_i} p_j > 0\}$ and $V_{\mathsf{SW}}(o) := \{i \in V : v_i(X_i) > 0\}$. By construction of the valuation functions, both $V_{\mathsf{REV}}(o)$ and $V_{\mathsf{SW}}(o)$ have to be independent sets for G. Let $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$. Since $\sum_{j\in X_i} p_j \leq 1$ and $v_i(X_i) \leq 1$ for every $i \in N'$, it follows that

$$\mathsf{obj}(o) \leq |V_{\mathsf{obj}}(o)|. \tag{1}$$

Let $V^* \subseteq V$ be a maximum independent set for G. It is easy to see that the outcome $(\mathbf{X}^*, \mathbf{p}^*)$ such that $X_i^* = \delta_i(G)$ for every $i \in V^*$ and

$$p_j^* = \begin{cases} 1/\delta_i(G) & \text{if } j \in \delta_i(G) \\ 0 & \text{otherwise.} \end{cases}$$

is an item envy-free outcome (actually it is also a Walrasian equilibrium because the market clears, i.e, every unsold item is assigned price zero) for market $\Gamma(N')$. This implies

$$\mathsf{obj}(o^*(\mathsf{BP}(P))) \geq \mathsf{obj}(\mathbf{X}^*, \mathbf{p}^*) = |V^*|. \tag{2}$$

Assume, for the sake of contradiction, that there exists an approximation algorithm for $\mathsf{BP}(P)$ returning an outcome o such that $n^{1-\epsilon}\mathsf{obj}(o) \geq \mathsf{obj}(o^*(\mathsf{BP}(P)))$ for some $\epsilon > 0$. Using (2), we get $n^{1-\epsilon}\mathsf{obj}(o) \geq |V^*|$ which combined with (1) implies that $n^{1-\epsilon}|V_{\mathsf{obj}}(o)| \geq |V^*|$: a contradiction to the inapproximability result for the maximum independent set problem. ◀

As a positive result, we show that, by building upon (approximation) algorithms for pricing problems defined on markets with a unique buyer, it is possible to obtain approximation algorithms for the buyer preselection problem. In particular, the following theorem can be proved by considering the buyer providing, when being alone in the market, the best possible outcome with respect to the considered objective function.

▶ **Lemma 6.** *Given a buyer preselection problem* $\mathsf{BP}(P)$*, where* $P = (\Gamma, \mathsf{IEF}, \mathsf{obj})$ *with* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$*, if there exists a polynomial time algorithm* $\mathcal{A}$ *returning an outcome for the pricing problem defined on markets with a unique buyer whose objective value is at least an* α *fraction of the optimal social welfare, then* $\mathsf{BP}(P)$ *admits an* αn*-approximation algorithm.*

Proof. Assume that $o^*(\mathsf{BP}(P)) := (\mathbf{X}^*, \mathbf{p}^*)$. For every $i \in N^*(\mathsf{BP}(P))$, define

$$\mathsf{obj}_i(o^*(\mathsf{BP}(P))) = \begin{cases} \sum_{j \in X_i^*} p_j^* & \text{if } \mathsf{obj} = \mathsf{REV}, \\ v_i(X_i^*) & \text{if } \mathsf{obj} = \mathsf{SW}. \end{cases} \tag{3}$$

Consider the preselection algorithm which, given Γ, returns a buyer $i^* \in N$ such that $i^* \in \operatorname{argmax}_{i \in N} \{\max_{X \subseteq M}\{v_i(X)\}\}$. Let $o \in \mathsf{IEF}(\Gamma(\{i^*\}))$ be the outcome returned by $\mathcal{A}$ when executed on the pricing problem $P' = (\Gamma(\{i^*\}), \mathsf{IEF}, \mathsf{obj})$. We have

$$\begin{aligned} \mathsf{obj}(o^*(\mathsf{BP}(P))) &= \sum_{i \in N^*(\mathsf{BP}(P))} \mathsf{obj}_i(o^*(\mathsf{BP}(P))) \\ &\leq \sum_{i \in N^*(\mathsf{BP}(P))} v_i(X_i^*) \\ &\leq n v_{i^*}(X_{i^*}^*) \\ &\leq \alpha n \mathsf{obj}(o), \end{aligned}$$

where the first inequality comes from (3) and the fact that $\sum_{j \in X_i^*} p_j \leq v_i(X_i^*)$ because of individual rationality, the second inequality follows from the definition of i^*, and the third inequality comes from the hypothesis on algorithm $\mathcal{A}$. ◀

As a consequence of Theorem 5 and Lemma 6, we have that the buyer preselection problem $\mathsf{BP}(\Gamma, \mathsf{IEF}, \mathsf{obj})$, with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, admits a polynomial time algorithm providing the best possible approximation guarantee, whenever the pricing problem defined on markets with a unique buyer can be solved in polynomial time with respect to the social welfare objective function. For such a reason, in the following subsection, we focus on the solution of the latter problem.

3.1 Pricing Problems Defined on Markets with a Unique Buyer

Throughout this subsection, since there is only one buyer in the market, for the sake of simplicity we remove the pedis 1 from the notation.

As a warmup, we start by considering the simpler case in which $\mathsf{obj} = \mathsf{SW}$.

▶ **Claim 7.** *Let* Γ *be a market with a single buyer and combinatorial valuations. The pricing problem* $(\Gamma, \mathsf{IEF}, \mathsf{SW})$ *can be solved in polynomial time.*

In fact, observe that an outcome $(X^*, \mathbf{0}^m)$ such that $X^* \in \text{argmax}_{X \subseteq M} v(X)$ verifies $(X^*, \mathbf{0}^m) \in opt(P)$, and a set $X^* \in \text{argmax}_{X \subseteq M} v(X)$ can be obtained in polynomial time by using the price vector $\mathbf{0}^m$ as the input of an oracle demand query.

We show in the next theorem that the case of $\mathsf{obj} = \mathsf{REV}$ yields an **NP**-hard problem, thus solving a longstanding open problem left by [5] (where in Lemma 7 an approximation algorithm with no hardness result is provided). Theorem 8 can be proved by exploiting a polynomial reduction from 3SAT, in which a given boolean formula ϕ is transformed into a market with a unique buyer, whose items are the literals of ϕ.

▶ **Theorem 8.** *Let* Γ *be a market with a single buyer and combinatorial valuations. The pricing problem* $(\Gamma, \mathsf{IEF}, \mathsf{REV})$ *is* **NP***-hard.*

Proof. We prove the claim through a reduction from 3SAT. To this aim, given a boolean formula ϕ, let $\mathsf{V}(\phi)$ denote the set of its variables and $\mathsf{L}(\phi)$ the set of all possible literals on variables in $\mathsf{V}(\phi)$; moreover, denote $\nu = |\mathsf{V}(\phi)|$. Throughout this proof, we assume $\nu \geq 4$. An assignment for ϕ is a function $f : \mathsf{V}(\phi) \to \{0, 1\}$ assigning to each variable of ϕ a boolean value. Denote with $F(\phi)$ the set of all possible assignments for ϕ and with $\phi(f)$ the boolean value obtained by evaluating all literals occurring in ϕ according to f. ϕ is satisfiable if there exists an assignment $f \in F(\phi)$ such that $\phi(f) = 1$ and it is unsatisfiable if, for every assignment $f \in F(\phi)$, $\phi(f) = 0$. Given a set of literals $X \subseteq \mathsf{L}(\phi)$, with a little abuse of notation, we write $X \in F(\phi)$ whenever there exists an assignment $f \in F(\phi)$ such that X contains all and only those literals which are evaluated 0 according to f. A formula ϕ is an instance of 3SAT if ϕ is expressed in Conjunctive Normal Form and each clause is the disjunction of 3 literals, so that ϕ can be completely expressed by listing the set $C = \{c_1, \ldots, c_k\}$ of its clauses, where each clause c_i $(i = 1, \ldots, k)$ is a set of three literals.

Given an instance of 3SAT $\phi := C$, we construct a market $\Gamma = (\{1\}, M, v)$ with a unique buyer such that $M = \mathsf{L}(\phi)$ and the valuation function v is defined as follows:

$$v(X) = \begin{cases} 1 & \text{if } |X| = 1, \\ 3 + \epsilon & \text{if } X \in C, \\ \nu & \text{if } X \in F(\phi), \\ 0 & \text{otherwise,} \end{cases}$$

where $\epsilon > 0$ is arbitrarily small.

Clearly Γ can be constructed in polynomial time with respect to the representation of ϕ. However, in order to complete the reduction, we have to construct an oracle which can answer both demand and value queries in polynomial time. To this aim, observe that, given a set of literals X, checking whether X is a singleton set, or $X \in C$, or $X \in F(\phi)$ can be performed in polynomial time, so that value queries can be efficiently answered. In order to provide an efficient answer to a demand query, we first observe that the cardinalities of sets $\mathsf{L}(\phi)$ and C are polynomial in the representation of ϕ, and therefore, given a pricing vector, a set of items in $\mathsf{L}(\phi)$ and C yielding the highest utility can be efficiently computed by enumeration. Then, in order to compute a set $X \in F(\phi)$ of maximum utility, note that all the candidate sets have the same valuation so that, in order to return one with the highest utility, we simply need to choose, for each variable in $\mathsf{V}(\phi)$, the related literal having the lowest price. Hence, market Γ can be generated and managed in polynomial time.

Now, in order to complete the proof, we show that there exists an outcome $o \in \mathsf{OUT}(\Gamma)$ such that $\mathsf{REV}(o) = \nu$ if and only if ϕ is satisfiable.

Assume first that ϕ is satisfiable and let f be a satisfying assignment for ϕ. Let X be the set of literals which are evaluated 0 in f and let $\mathbf{p}$ be the pricing vector such that all literals in X are priced 1, while all literals in $\mathsf{L}(\phi) \setminus X$ are priced ∞. By definition, we have $u(X, \mathbf{p}) = 0$, so that $(X, \mathbf{p})$ is individually-rational. Moreover, $(X, \mathbf{p}) \in D(\mathbf{p})$. In fact, for every $X' \neq X$ such that $X' \in F(\phi)$, $u(X', \mathbf{p}) < 0$ since X' has to contain at least one item priced ∞; for every $X' \in C$, $u(X', \mathbf{p}) < 0$ since, as we have $\phi(f) = 1$, X' has to contain at least one item priced ∞; for every $X' \in \mathsf{L}(\phi)$, $u(X', \mathbf{p}) \leq 0$ by construction. Hence, $(X, \mathbf{p}) \in \mathsf{IEF}(\Gamma)$ and $\mathsf{REV}(X, \mathbf{p}) = \nu$.

Secondly, assume that there exists an outcome $(X, \mathbf{p}) \in \mathsf{IEF}(\Gamma)$ such that $\mathsf{REV}(X, \mathbf{p}) = \nu$. By construction of the valuation function, this is possible only if $X \in F(\phi)$ and $\sum_{j \in X} p_j = \nu$, so that $u(X, \mathbf{p}) = 0$. However, since $v(\{\ell\}) = 1$ for every $\ell \in \mathsf{L}(\phi)$, $(X, \mathbf{p}) \in \mathsf{IEF}(\Gamma)$ implies that it must also be $p_j \geq 1$ for each $j \in X$. Hence, we can conclude that $p_j = 1$ for each $j \in X$. Assume, for the sake of contradiction, that ϕ is unsatisfiable. This implies that the assignment induced by all literals in $\mathsf{L}(\phi) \setminus X$ cannot satisfy ϕ, that is, there exists a clause $c_j \in C$ such that $c_j \subseteq X$. This implies $u(c_j, \mathbf{p}) = 3 + \epsilon - 3 = \epsilon > 0$ thus contradicting $(X, \mathbf{p}) \in \mathsf{IEF}(\Gamma)$. Hence, ϕ has to be satisfiable. ◀

On the positive side, [5] derive an $O(\log m)$-approximation algorithm for this problem. An interesting feature of this algorithm is that its performance guarantee holds also with respect to the maximum social welfare which is an upper bound to the maximum revenue, thus allowing the application of Lemma 6.

Hence, by combining Lemma 6 with Claim 7 and the result of [5], we obtain the following upper bounds.

▶ **Theorem 9.** *Let $P = (\Gamma, \mathsf{IEF}, \mathsf{obj})$ be a pricing problem with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$. The buyer preselection problem $\mathsf{BP}(P)$ admits an n-approximation when $\mathsf{obj} = \mathsf{SW}$, and an $O(n \log m)$-approximation when $\mathsf{obj} = \mathsf{REV}$.*

It is worth noticing that, given the proof of Lemma 6, the preselection claimed in Theorem 9 is someway "oblivious", i.e., it can be obtained by exploiting the minimum possible number of oracle queries, that is only a single oracle (demand) query for each buyer.

In light of the lower bound given in Theorem 5, the upper bounds given in Theorem 9 are asymptotically tight both for $\mathsf{obj} = \mathsf{SW}$ (unless $\mathbf{P} = \mathbf{ZPP}$) and for $\mathsf{obj} = \mathsf{REV}$ when $m = o(n)$ (unless $\mathbf{P} = \mathbf{ZPP}$).

4 Results for Bundle Envy-Free Outcomes

In this section, we consider the buyer preselection problem $\mathsf{BP}(\Gamma, \mathsf{BEF}, \mathsf{obj})$ with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$. Since we deal with bundle envy-free outcomes, we suppose that the price of any unsold item is infinite. Formally, given an outcome (N, o), where $o = (\mathbf{X}, \mathbf{p})$, for the buyer preselection problem $\mathsf{BP}(\Gamma, \mathsf{BEF}, \mathsf{obj})$ with $\Gamma = (N, M, (v_i)_{i \in N})$, we have that, for any $j \in M \setminus M(\mathbf{X})$, $p_j = \infty$.

We start by considering the buyer preselection problem $\mathsf{BP}(\Gamma, \mathsf{BEF}, \mathsf{REV})$.

The following theorem shows how it is possible to transform a bundle envy-free solution $(\bar{N}, \bar{o})$ with preselection into anther one (without preselection) having a non-smaller revenue.

▶ **Theorem 10.** *Given any solution* $(\bar{N}, \bar{o})$ *for the buyer preselection problem* $\mathsf{BP}(\Gamma, \mathsf{BEF}, \mathsf{REV})$, *with* $\Gamma = (N, M, (v_i)_{i \in N})$, *it is possible to compute in polynomial time an outcome* $o \in \mathsf{BEF}(\Gamma)$ *for problem* P *such that* $\mathsf{REV}(o) \geq \mathsf{REV}(\bar{o})$.

Theorem 10 can be proved by exploiting the notion and properties of maximal solutions, i.e, solutions in which no item price can be increased without changing the allocation (notice that an optimal solution is maximal), and by providing a constructive algorithm working on maximal solutions and allocating bundle of items to the players in $N \setminus \bar{N}$ that envy other winners. In particular, we first show that in a maximal solution $(\bar{N}, \bar{o})$ with preselection (i) for any buyer with positive utility, there exists another sold bundle providing her the same utility and (ii) there always exists a winner buyer having utility equal to 0. Given these properties, it is possible to exchange among the buyers the assigned bundles so that an excluded envious buyer can be assigned her preferred sold bundle $\bar{X}_j$ and can be therefore added to the solution without generating envy. Buyer j, getting bundle $\bar{X}_j$ in $(\bar{N}, \bar{o})$, gets another one providing her the same utility (by property (i) such a bundle always exists). This process can be iterated until a winner buyer with utility equal to 0 is reached and removed from the set of winners (it is possible to show that in this process a buyer is never considered twice and therefore by property (ii) it always terminates).

We now consider the buyer preselection problem $\mathsf{BP}(\Gamma, \mathsf{BEF}, \mathsf{SW})$. Analogously to Theorem 10 holding for the revenue maximization case, next theorem shows that a bundle envy-free outcome for the buyer preselection problem can be efficiently transformed in a bundle envy-free outcome, for the corresponding pricing problem (without preselection), having at least the same social welfare.

▶ **Theorem 11.** *Given any solution* $(\bar{N}, \bar{o})$ *for the buyer preselection problem* $\mathsf{BP}(\Gamma, \mathsf{BEF}, \mathsf{SW})$, *with* $\Gamma = (N, M, (v_i)_{i \in N})$, *it is possible to compute in polynomial time an outcome* $o \in \mathsf{BEF}(\Gamma)$ *for problem* P *such that* $\mathsf{SW}(o) \geq \mathsf{SW}(\bar{o})$.

Theorem 11 can be proved by considering a new market Γ' with unit-demand buyers, in which the set of buyers is N and there is an item for every bundle sold in $(\bar{N}, \bar{o})$. In fact, given that markets with unit-demand buyers always admit a Walrasian Equilibrium computable in polynomial time and that, by the well known *First Welfare Theorem*, Walrasian equilibria maximize social welfare over all possible outcomes, it can be easily obtained, by suitably setting the price of the items in Γ as a function of those in Γ', a bundle envy-free solution for Γ with no excluded buyer and having the same social welfare of $\bar{o}$.

As a consequence of Theorems 10 and 11, we obtain the following corollary.

▶ **Corollary 12.** *For* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, *given an* α*-approximate solution* (N, o) *(with* $\alpha \geq 1$, *notice that when* $\alpha = 1$ *the corollary holds for optimal solutions) for the buyer preselection problem* $\mathsf{BP}(P)$ *with* $P = (\Gamma, \mathsf{BEF}, \mathsf{obj})$ *and* $\Gamma = (N, M, (v_i)_{i \in N})$, *it is possible to compute in polynomial time an outcome* $o' \in \mathsf{BEF}(\Gamma)$ *approximating the optimal solution of* P *by a factor equal to* α.

On the one hand, Corollary 12 tells us that, for $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, any inapproximability result holding for a pricing problem $P = (\Gamma, \mathsf{BEF}, \mathsf{obj})$ directly extends to the buyer preselection problem $\mathsf{BP}(P)$; on the other hand, it tells us that any α-approximation algorithm for $\mathsf{BP}(P)$ is also an α-approximation algorithm for P. Thus, as discussed in the Our Contribution subsection, preselection can be exploited as an algorithmic framework for designing approximation algorithms for the normal market scenario (without preselection).

5 The multi-unit case

In this section we study the multi-unit case, in which all the m items are of the same type. Recall that, for every $i \in N$, the valuation function becomes of the form $v_i : \{1, \ldots, m\} \to \mathbb{R}_{\geq 0}$ and that the market can be represented with $\Theta(nm)$ bits. Furthermore, in this case an allocation vector X can be specified by the number of items assigned to each buyer, i.e. $X = (x_1, \ldots, x_n)$, with $x_i \in \{0, \ldots, m\}$ for any $i = 1, \ldots, n$, and $\sum_{i=1}^{n} x_i \leq m$. Finally, we set the item pricing to p (i.e., all the multi-items have the same price) and the total price for selling x items is px.

A particular situation of multi-unit market, moreover, arises when one assumes *single-minded* buyers; in this case, since, for every player $i = 1, \ldots, n$, valuation function v_i can be completely defined by specifying how much every player valuates the set containing k_i items (being the only set she is interested in), $\Theta(n \log m)$ bits suffice to represent the market.

We first focus on the case of *item envy-free* solutions. Next claim shows that preselection can improve both the revenue and the social welfare of a market with multi-unit items, even in the case of single-minded buyers.

▶ **Claim 13.** *For any $\epsilon > 0$, there exists a market Γ with single-minded buyers and multi-unit items, and a pricing problem $P = (\Gamma, \mathsf{IEF}(\Gamma), \mathsf{obj})$ with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, such that $\mathsf{obj}(o^*(\mathsf{BP}(P))) \geq (m - \epsilon)\mathsf{obj}(o^*(P))$.*

In fact, consider a market with only two single-minded buyers. Buyer 1 valuates $1 + \epsilon'$, for a small $\epsilon' > 0$, for receiving any (one) item. Buyer 2 valuates m for receiving all the m items. Notice that, in any item envy-free outcome for the setting without preselection, buyer 2 receives no item. In fact, if buyer 2 receives m items, the item pricing p must be at most 1. Thus, buyer 1 would be envious since she would get no item (i.e., there are no enough items). Therefore, for the market without preselection, the only feasible solutions are selling one item to buyer 1 at price at least 1 and at most $1 + \epsilon'$. Thus the optimal revenue and the optimal social welfare are at most $1 + \epsilon'$. However, if we consider the submarket with only buyer 2, i.e., we exclude buyer 1, we can sell m items at item pricing 1, thus obtaining a revenue and social welfare of m.

We notice that the above bound is tight. It is easy to see that preselection cannot improve the revenue and social welfare of envy-free solutions by a factor greater than m. In fact, it is sufficient to sell items to the buyer i such that $(i, j) = \text{argmax}_{i=1,\ldots,n; j=1,\ldots,m} \frac{v_i(j)}{j}$.

Now we focus on the computation of optimal or approximate solution for the preselection problem in the case of item envy-free outcomes, both for the social welfare and the revenue objective functions.

▶ **Theorem 14.** *Given a pricing problem $P = (\Gamma, \mathsf{IEF}(\Gamma), \mathsf{obj})$, where $\Gamma = (N, M, (v_i)_{i \in N})$ is a market with multi-unit items and $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, the buyer preselection problem $\mathsf{BP}(P)$ can be optimally solved in polynomial time.*

Proof. Recall that solving the buyer preselection problem $\mathsf{BP}(P)$ corresponds to find a subset N^* of players to admit to the market and an optimal outcome $o^* = (X^*, p^*)$, in which in the considered case of multi-unit items p^* is just the price of a single item.

In order to prove the claim, we first show that it is possible to compute in polynomial time a set $\mathcal{P}$ containing the optimal price p^* for problem $\mathsf{BP}(P)$.

Consider the set $\mathcal{P}$ defined as follows:

$$\begin{aligned} \mathcal{P} \quad = \quad & \{y | v_i(k) - yk = v_i(k') - yk', i \in N, \\ & k, k' \in \{0, 1, \ldots, m\}, k \neq k'\}. \end{aligned}$$

Roughly speaking, for all buyers $i \in N$ and for all couples (k, k') of integers belonging to $\{0, 1, \ldots, m\}$ such that $k \neq k'$, $\mathcal{P}$ contains the solution y of equality $v_i(k) - yk = v_i(k') - yk'$. Clearly, $\mathcal{P}$ can be computed in $O(nm^2)$ time, i.e., in polynomial time in the size of the instance. Now, assume by contradiction that the optimal solution (N^*, o^*) for $\mathsf{BP}(P)$ with the highest possible item price assign to an item price $p^* \notin \mathcal{P}$, and let $p' \in \mathcal{P}$ the smallest element of $\mathcal{P}$ such that $p' > p^*$. Notice that this element $p' \in \mathcal{P}$ has to exist, because p^* must verify, for any $i = 1, \ldots, n$, $v_i(\bar{x}_i) - p^*\bar{x}_i \geq 0$ and y verifying equality $v_i(\bar{x}_i) - y\bar{x}_i = 0$ belongs to $\mathcal{P}$. Since p^* induces an envy-free solution, it has to verify, for every $i = 1, \ldots, n$, constraint $v_i(\bar{x}_i) - p^*\bar{x}_i \geq v_i(j) - p^*j$ for any $j = 0, \ldots, m$; since, $p^* \notin \mathcal{P}$, given how $\mathcal{P}$ is defined, it follows that all above constraints are not verified in a strict manner, i.e. it is verified that, for every $i = 1, \ldots, n$, $v_i(\bar{x}_i) - p^*\bar{x}_i > v_i(j) - p^*j$ for any $j = 0, \ldots, m$. Therefore, p' still continues to verify all envy-free constraints (with some constraints possibly become strict). It follows that we have found a new envy-free outcome o' such that $\mathsf{SW}(o') = \mathsf{SW}(o^*)$ and $\mathsf{REV}(o') > \mathsf{REV}(o^*)$: a contradiction to the fact that o^* was the optimal outcome with the highest possible item price.

Therefore, since the number of values that can be assigned to p^* in order to obtain an optimal outcome is polynomial in the size of the instance, it remains to show that, given a fixed price p^*, it is possible to optimally compute in polynomial time a subset N^* of players to admit to the market and an allocation X^*. In fact, the optimal solution of $\mathsf{BP}(P)$ is given by the best solution among the ones obtained for all the candidate prices belonging to $\mathcal{P}$.

The *0-1 Multiple-Choice Knapsack Problem (0-1 MCKP)* is a generalization of the classical Knapsack problem introduced in [30]. In this problem, we are given α classes $C_1, C_2, \ldots, C_\alpha$ of elements to pack in some knapsack of capacity c. For every $i = 1, \ldots, n$, each element $e \in C_i$ (let $class(e) = i$ be the index of the class to which e belongs) has a profit β_e and a volume γ_e, and the problem is to choose a set E containing *at most* one element from each class such that the profit sum is maximized without the volume sum exceeding capacity c. In [21] it is shown that it can be optimally solved in pseudo-polynomial time, i.e., in time $O(c)$.

We now provide a polynomial reduction from our problem to 0-1 MCKP. Given an instance I of $\mathsf{BP}(P)$ and fixed a price p^*, we construct an instance I' of 0-1 MCKP as follows: We have $\alpha = n$ classes (one class for each buyer) and the capacity $c = m$. Consider, for every buyer $i = 1, \ldots, n$, the number of items providing her with the highest possible utility: let $\bar{U}_i = \arg\max_{k=1,\ldots,m} v_i(k) - kp^*$ be the set containing these values. It is easy to check that, in every item envy-free solution, allocation X must satisfy $x_i \in \bar{U}_i$ for every player $i = 1, \ldots, n$. For every $i = 1, \ldots, n$, consider set $\bar{U}_i$: we add to class C_i an element e for every $k \in \bar{U}_i$ such that $\gamma_e = k$ and

- $\beta_e = k$ if $\mathsf{obj} = \mathsf{REV}$;
- $\beta_e = v_i(k)$ if $\mathsf{obj} = \mathsf{SW}$.

Given a solution for I' with total profit $\beta = \sum_{e \in E} \beta_e$, it is possible to obtain a solution for I with fixed price p^*, i.e., a subset $\bar{N}$ of players to admit to the market and an allocation $\bar{X}$, as follows: $\bar{N}$ contains the players associated to classes containing an element belonging to E, i.e., $\bar{N} = \{i | C_i \cap E \neq \emptyset\}$, and the allocation vector $\bar{X} = (\bar{x}_1, \ldots, \bar{x}_n)$ is such that, for every $i = 1, \ldots, n$, $x_i = \gamma_e$ if there exists an element $e \in C_i \cap E$. Clearly, $\mathsf{obj}((\bar{X}, p^*)) = \beta$. Furthermore, by the way $\bar{U}_i$ is defined, outcome $(\bar{X}, p^*)$ is item envy-free.

Conversely, given a subset $\bar{N}$ of preselected players and an outcome $o = (X, p^*)$ for I, it is possible to obtain a solution for I' as follows: for every $i \in \bar{N}$, add to E element $e \in C_i$ such that $\gamma_e = x_i$ (by recalling the definition of $\bar{U}_i$, it holds that this element e belongs to C_i because outcome o is item envy-free). Clearly, the total profit $\beta = \sum_{e \in E} \beta_e = \mathsf{obj}((X, p^*))$.

The claim follows by noticing that, since $c = m$, the pseudo-polynomial algorithm of [21] is in fact polynomial with respect to the size of the instance of problem $\mathsf{BP}(P)$. ◀

For the special case of single-minded buyers in which an instance of the buyer preselection problem can be represented by $\Theta(n \log m)$ bits, by exploiting the same ideas used for proving Theorem 14, the following theorem provides an **FPTAS**.

▶ **Theorem 15.** *Given a pricing problem* $P = (\Gamma, \mathsf{IEF}(\Gamma), \mathsf{obj})$*, where* $\Gamma = (N, M, (v_i)_{i \in N})$ *is a market with single-minded buyers and multi-unit items, for* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$*, the buyer preselection problem* $\mathsf{BP}(P)$ *admits a* fully polynomial approximation scheme.

We complement the result of Theorem 15 by showing a tight lower bound to the problem of computing the maximum revenue to the case of item envy-free with single-minded buyers and multi-unit items. The following claim can be proved by exploiting a reduction from the Subset Sum problem.

▶ **Claim 16.** *Given a pricing problem* $P = (\Gamma, \mathsf{IEF}(\Gamma), \mathsf{REV})$*, where* $\Gamma = (N, M, (v_i)_{i \in N})$ *is a market with single-minded buyers and multi-unit items, the buyer preselection problem* $\mathsf{BP}(P)$ *is* **NP***-Hard.*

Proof. We use a reduction from the Subset Sum problem, that is defined as follows: given a set of integers and an integer s, does any non-empty subset sum to s? Let n be the number of elements $\{a_1, a_2, \ldots, a_n\}$ in the given instance of the Subset sum problem, and let s be the required sum. We assume that $\sum_{i=1}^{n} a_i > s$, as otherwise the problem is trivial. For all a_i in the input of the Subset sum problem, we create a corresponding buyer i with the following valuation. Buyer i has valuation a_i for receiving a_i items, and zero otherwise. Moreover, we set the number of items $m = s$. In such case, if there is a solution for the Subset sum problem, then by setting the item pricing $p = 1$ and selling to corresponding buyers gives us a feasible and envy-free outcome in which the revenue equals to m. It is easy to see that m is an upper bound to the maximum revenue. On the other hand, if the revenue of the optimal outcome to the buyer preselection problem is equal to m, we can obtain the solution to the subset sum problem. Notice that we can obtain a revenue of m only if we sell exactly m items at item pricing $p = 1$. In fact, at item pricing $p > 1$, we sell no item, and at item pricing $p < 1$, we do not get an optimal solution. ◀

We now focus on the case of *bundle envy-free* solutions. We first notice that the results of Section 4 claiming that bundle envy-free solutions do not improve the quality of outcomes (with respect neither to social welfare nor to revenue maximization) do not hold for the multi-unit case, because in this case it is not possible to change the price of some item without influencing the other ones. We start by showing that preselection can improve both the revenue and the social welfare of a market with multi-unit items, even in the case of single-minded buyers.

▶ **Claim 17.** *For any* $\epsilon > 0$*, there exists a market* Γ *with single-minded buyers and multi-unit items, and a pricing problem* $P = (\Gamma, \mathsf{BEF}(\Gamma), \mathsf{obj})$ *with* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$*, such that* $\mathsf{obj}(o^*(\mathsf{BP}(P))) \geq (2 - \epsilon)\mathsf{obj}(o^*(P))$*.*

In fact, consider a market with $x + 1$ single-minded buyers. Buyer 1 valuates x for receiving x items. Buyer i, for any $i = 2, 3, \ldots, x + 1$, valuates $1 + \epsilon'$, for a small $\epsilon' > 0$, for receiving any (one) item. Finally, the number of items is $2x - 1$. Notice that, in any bundle envy-free outcome for the setting without preselection, if buyers 1 receives x items, it implies that the item pricing p must be at most 1. It further implies that, in such outcome, no buyer i, for any $i = 2, 3, \ldots, x + 1$, can get items. The reason is that such buyers have positive utility for receiving one item, but the number of items is not sufficient to satisfy all of them. Therefore, the only chance is selling no bundle of one item. Thus, on one hand, without preselection, the best revenue and social welfare is $x(1+\epsilon')$. It can be obtained by selling one item to buyers i, for any $i = 2, 3, \dots, x+1$, at price $1+\epsilon'$. On the other hand, if we consider the submarket with only buyers $1, 2, \dots, x$ (i.e., we exclude one buyer that valuates $1+\epsilon'$ for receiving any (one) item), we can sell $2x-1$ items at item pricing 1, that is x items to buyer 1, and one item to buyers i, for any $i = 2, \dots, x$, thus obtaining a revenue and social welfare of $2x-1$.

We now show that, for the market with multi-unit items and general valuations, preselection can improve the revenue by a multiplicative factor of at most 2, thus closing in a tight way the previous bound.

▶ **Theorem 18.** *Given a solution $(\bar{N}, \bar{o})$ for the buyer preselection problem* $\mathsf{BP}(\Gamma, \mathsf{BEF}, \mathsf{REV})$*, with $\Gamma = (N, M, (v_i)_{i \in N})$ being a market with multi-unit items, it is possible to compute in polynomial time an outcome $o \in \mathsf{BEF}(\Gamma)$ for problem P such that $\mathsf{REV}(\bar{o}) \leq 2\mathsf{REV}(o)$.*

Proof. Let $\bar{o} = (\bar{X}, \bar{p})$. In the following we will show how to compute in polynomial time an outcome $o \in \mathsf{BEF}(\Gamma)$ with $o = (X, p)$ such that (i) $p \geq \bar{p}$ and (ii) $|M(X)| \geq \frac{|M(\bar{X})|}{2}$, i.e., outcome o sells at least one half of the items sold by outcome $\bar{o}$ at a price at least equal to $\bar{p}$. Clearly, this directly implies that $\mathsf{REV}(o) \geq \frac{\mathsf{REV}(\bar{o})}{2}$.

For every $j \in \{1, \dots, m\}$, let a_j be the number of items that could be sold to some buyers in bundles of cardinality j at price $\bar{p}$ (we require that these buyers obtain a non-negative utility for a bundle of j items). More formally, for every $j = 1, \dots, m$, let $B_j(\bar{X}) = \{i | v_i(j) - j\bar{p} \geq 0\}$ be the subset of players obtaining a non-negative utility for a bundle with j items; then, $a_j = j|B_j(\bar{X})|$.

We divide the proof in two disjoint cases.

- If there exists $j \in \{1, \dots, m\}$ such that $a_j \geq \frac{|M(\bar{X})|}{2}$, outcome $o = (X, p)$ is such that $x_k = j$ for every $k \in B_j(\bar{X})$, and $x_k = 0$ otherwise. By setting $p = \bar{p}$, by the definition of a_j, we know that a_j items could be sold without generating envy.
  If $a_j \leq m$, we are done.
  If $j \geq \frac{|M(\bar{X})|}{2}$, we can increase the price p so that only buyer i, with i such that $v_i(j) = \max_{k=1}^{n} v_k(j)$, is assigned the bundle (notice that in this way no other buyer is envious).
  It remains to deal with the subcase in which $a_j > m$ and $j < \frac{|M(\bar{X})|}{2}$: We increment price p until the number of buyers x with positive utility is such that $xj \leq m$. We then assign bundles of j items to all buyers with positive utility and to as many buyers with zero utility as possible. Notice that (i) since $j < \frac{|M(\bar{X})|}{2}$ implies that $m - j > \frac{|M(\bar{X})|}{2}$, this process leads to obtain at least $\frac{|M(\bar{X})|}{2}$ assigned items and (ii) again no buyer is envious.
- If for all $j \in \{1, \dots, m\}$ it holds that $a_j < \frac{|M(\bar{X})|}{2}$, outcome $o = (X, \bar{p})$ (with the same price of outcome $\bar{o}$) is computed as follows.
  For any $i = 1, \dots, n$ and $k = 1, \dots, m$, let $u_i^{*k} = \max_{t=1}^{k} v_i(t) - t\bar{p}$ be the maximum possible utility of buyer i for bundles of at most k items and $b_i^{*k} = \min\{j | v_i(j) - j\bar{p} = u_i^{*k}\}$ the minimum size of a bundle of maximum utility for buyer i. Moreover, for any $k = 1, \dots, m$ and $j = 1, \dots, k$, let $B_{k,j} = \{i | b_i^{*k} = j\}$ be the set of buyers having j items as their best bundle of maximum utility, among all bundles made up to k items, and resolving ties by selecting the bundle of minimum size.
  Clearly, $B_{k,j}$ can be computed in time $O(poly(n, m))$. For $k = 1, \dots, m$, consider allocation $X^k = (x_1^k, \dots, x_n^k)$ such that, for every $i = 1, \dots, n$, $x_i^j = j$ if $i \in B_{k,j}$ and $x_i^j = 0$ otherwise. By the definition of $B_{k,j}$ it follows that allocation X^k is envy-free. Moreover, it can be easily verified that $|M(X^1)| = a_1$ and, for any $j = 2, \dots, m$, $|M(X^j)| - |M(X^{j-1})| \leq a_j$.

If $|M(X^m)| \leq m$, the claim trivially follows by setting $X = X^m$ because we are allocating at least all items allocated in $\bar{o}$.
Otherwise, let k' be the minimum value of $k = 2, \ldots, m$ such that $|M(X^{k'})| > m$: the claim follows by setting $X = X^{k'-1}$. In fact, since $|M(X^{k'})| - |M(X^{k'-1})| \leq a_{k'} \leq \frac{|M(\bar{X})|}{2}$, it follows that $|M(X^{k'-1})| \geq m - \frac{|M(\bar{X})|}{2} \geq \frac{|M(\bar{X})|}{2}$. ◀

It is worth noticing that, on the one hand, preselection can be exploited as an algorithmic framework for designing good approximation algorithms (loosing only a moltiplicative factor of 2) for the normal market scenario without preselection; on the other hand, since the optimal revenue with preselection is at most twice the one without, an α-approximation algorithm for the normal market without preselection, is a 2α-approximation one for market with preselection.

6 Final remarks and Future work

Many results holding for the item envy-free outcomes and social welfare objective function extend to the notion of Walrasian equilibria, that are item envy-free outcomes with the additional requirement that the market clears, i.e., every unsold item is assigned price zero. In particular, the inapproximability result of Theorem 5 and the n-approximation algorithm for the buyer preselection problem of Theorem 9 directly extend to Walrasian equilibria. Notice also that for the remaining uncovered cases, that is when the goal is that of optimizing the seller's revenue, there is no reason for requiring market clearance, a condition clearly limiting the power of setting prices so as to maximize the revenue.

The main left open problems are: for markets with a unique buyer, closing the gap between the **NP**-hardness and the logarithmic approximation for the case of revenue maximization and item envy-free solutions; for the multi-unit case with bundle envy-free outcomes, determining an upper bound to the social welfare improvement achievable by preselection and setting the complexity of computing optimal solutions, for both the revenue and the social welfare cases.

References

1 Noga Alon, Yishay Mansour, and Moshe Tennenholtz. Differential pricing with inequity aversion in social networks. In *Proc. of EC*, pages 9–24, 2013.

2 Georgios Amanatidis, Evangelos Markakis, and Krzysztof Sornat. Inequity aversion pricing over social networks: Approximation algorithms and hardness results. In *Proc. of MFCS*, pages 9:1–9:13, 2016.

3 E. Anshelevich, K. Kar, and S. Sekar. Envy-free pricing in large markets: Approximating revenue and welfare. In *Proc. of ICALP*, pages 52–64. Springer, 2015.

4 A. Archer, C. H. Papadimitriou, K. Talwar, and É. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2):129–150, 2003.

5 M. F. Balcan, A. Blum, and Y. Mansour. Item pricing for revenue maximization. In *Proc. of EC*, pages 50–59, 2008.

6 S. Bikhchandani and J. W. Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2):386–413, 1997.

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

8 V. Bilò, M. Flammini, G. Monaco, and L. Moscardelli. On the impact of buyers preselection in pricing problems. In *Proc. of AAMAS*, 2018.

9 P. Briest. Uniform budgets and the envy-free pricing problem. In *Proc. of ICALP*, pages 808–819. Springer, 2008.
10 P. Briest and P. Krysta. Single-minded unlimited supply pricing on sparse instances. In *Proc. of SODA*, pages 1093–1102. ACM Press, 2006.
11 P. Chalermsook, J. Chuzhoy, S. Kannan, and S. Khanna. Improved hardness results for profit maximization pricing problems with unlimited supply. In *Proc. of APPROX*, pages 73–84. Springer, 2012.
12 P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Graph products revisited: Tight approximation hardness of induced matching, poset dimension and more. In *Proc. of SODA*, pages 1557–1576. ACM Press, 2013.
13 P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *Proc. of FOCS*, pages 370–379. IEEE Computer Society, 2013.
14 N. Chen and X. Deng. Envy-free pricing in multi-item markets. In *Proc. of ICALP*, pages 418–429. Springer, 2010.
15 N. Chen, X. Deng, P. W. Goldberg, and J. Zhang. On revenue maximization with sharp multi-unit demands. *Journal of Combinatorial Optimization*, 31(3):1174–1205, 2016.
16 N. Chen, A. Ghosh, and S. Vassilvitskii. Optimal envy-free pricing with metric substitutability. *SIAM Journal on Computing*, 40(3):623–645, 2011.
17 Ning Chen and Atri Rudra. Walrasian equilibrium: Hardness, approximations and tractable instances. *Algorithmica*, 52(1):44–64, 2008.
18 M. Cheung and C. Swamy. Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *Proc. of FOCS*, pages 35–44. IEEE Computer Society, 2008.
19 E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
20 E. D. Demaine, U. Feige, M. Hajiaghayi, and M. R. Salavatipour. Combination can be hard: Approximability of the unique coverage problem. *SIAM Journal on Computing*, 38(4):1464–1483, 2008.
21 K. Dudzinski and S. Walukiewicz. Exact methods for the knapsack problem and its generalizations. *European Journal of Operational Research*, 28(1):3–21, 1987.
22 M. Feldman, A. Fiat, S. Leonardi, and P. Sankowski. Revenue maximizing envy-free multi-unit auctions with budgets. In *Proc. of EC*, pages 532–549. ACM Press, 2012.
23 M. Feldman, N. Gravin, and B. Lucier. Combinatorial walrasian equilibrium. *SIAM Journal on Computing*, 45(1):29–48, 2016.
24 M. Flammini, M. Mauro, and M. Tonelli. On social envy-freeness in multi-unit markets. In *Proc. of AAAI*, 2018.
25 D. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7:45–98, 1967.
26 T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
27 F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87:95–124, 1999.
28 V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proc. of SODA*, pages 1164–1173. ACM Press, 2005.
29 J. Hartline and Q. Yan. Envy, truth, and profit. In *Proc. of EC*, pages 243–252. ACM Press, 2011.
30 E. L. Lawler. Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 44(4):339–356, 1979.
31 G. Monaco, P. Sankowski, and Q. Zhang. Revenue maximization envy-free pricing for homogeneous resources. In *Proc. of IJCAI*, pages 90–96, 2015.
32 H. R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.
33 W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
34 L. Walras. *Elements of Pure Economics*. Allen and Unwin, 1954.