Pricing Problems with Buyer Preselection

Vittorio Bilò　　　　　　　　　　　　　　　　　　　　VITTORIO.BILO@UNISALENTO.IT
University of Salento, Italy.

Michele Flammini　　　　　　　　　　　　　　　　　　　MICHELE.FLAMMINI@GSSI.IT
GSSI, L'Aquila, Italy.

Gianpiero Monaco　　　　　　　　　　　　　　　　　　GIANPIERO.MONACO@UNIVAQ.IT
University of L'Aquila, Italy.

Luca Moscardelli　　　　　　　　　　　　　　　　　　LUCA.MOSCARDELLI@UNICH.IT
University of Chieti-Pescara, Italy.

Abstract

We investigate the problem of preselecting a subset of buyers (also called agents) 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, namely market envy-freeness and agent envy-freeness, with the two state-of-the-art objective functions of social welfare and seller's revenue. When insisting on market envy-freeness, we prove that the problem cannot be approximated within $n^{1-\varepsilon}$ (with n being the number of buyers) for any $\varepsilon > 0$, under both objective functions; we also provide approximation algorithms with an approximation ratio tight up to subpolynomial multiplicative factors for social welfare and the seller's revenue. The negative result, in particular, holds even for markets with single-minded buyers. We also prove that maximizing the seller's revenue is NP-hard even for a single buyer, thus closing a previous open question. Under agent envy-freeness and for both objective functions, instead, we design a polynomial time algorithm transforming any stable outcome for a market involving any subset of buyers into a stable outcome for the whole market without worsening its performance. This result creates an interesting middle-ground situation where, if on the one hand buyer preselection cannot improve the performance of agent envy-free outcomes, on the other one it can be used as a tool for simplifying the combinatorial structure of the buyers' valuation functions in a given market. Finally, we consider the restricted case of multi-unit markets, where all items are of the same type and are assigned the same price. For these markets, we show that preselection may improve the performance of stable outcomes in all of the four considered scenarios, and design corresponding approximation algorithms.

1. Introduction

Pricing-based mechanisms (Nagle & Holden, 2002) provide a fundamental approach for conflict resolution in several competitive scenarios, such as the allocation of scarce resources in multi-agent systems or the sale of items with limited supplies in monopolistic markets. While satisfying suitable constraints for coping with the autonomous and self-interested agents' behavior, such as incentive-compatibility and/or fairness, these mechanisms also seek for the optimization of a given objective function.

We consider a market with perfect information in which there is *one seller* selling m items to n buyers (also called agents). In the most general case, every buyer has a valuation for every possible set of items, and her utility is given by the valuation she has for the set of bought items minus the sum of the prices assigned to them by the pricing mechanism. Usually, optimal prices are the result of a challenging counterbalancing process. For instance, if, on the one hand, selecting low prices may be profitable for the seller when it attracts considerably more customers, on the other hand, it may leave some buyers unsatisfied, thus generating discontent. In particular, this happens when, in case of limited supply, 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 them. In

this case she is often called *loser*, as opposed to a customer receiving items, called *winner*. This situation is practically illustrated by the following example.

Example 1. Consider a market with 2 identical items (in this case, the pricing mechanism has to decide a unique price p) and 2 single-minded buyers, i.e., being interested only at buying a unique combination of items in the sense that any other combination gives them zero value. Buyer 1 has a valuation of 3 for obtaining both items, while buyer 2 has a valuation of 2 for a single item.

Notice that, if $\frac{3}{2} < p \leq 2$, only buyer 2 is willing to buy, and one item remains unsold. So, the seller's revenue is at most 2. On the other hand, if the price is set to $p = \frac{3}{2}$, buyer 1 can purchase both items, leaving nothing to buyer 2. Thus, even though the seller gets a higher revenue of 3, buyer 2 is negated the right to buy an item, despite the fact that she is willing to pay for it: a sort of unfairness is present in the market.

For this reason, pricing problems are traditionally considered under the hypothesis of ***envy-freeness*** (Foley, 1967; Varian, 1974), which prescribes that, once a pricing strategy has been established, items have to be allocated to buyers in such a way that no one can get a higher utility by buying a different set of items.

Apart from fairness, envy-free pricing possesses additional appealing properties. From a theoretical point of view, the following structural parallelism can be claimed: envy-freeness can be seen, in the context of markets, as a notion corresponding to the one of equilibria in the context of game theory or multi-agent systems. In fact, it decentralizes the market in the sense that buyers, up to a proper tie-breaking rule, autonomously converge to the allocation dictated by a sort of equilibrium, since it corresponds to their preferred choices (Easley & Kleinberg, 2010). From an economic point of view, envy-free pricing, by safeguarding the losers' interests, shelters the seller from possible future losses due to their dissatisfaction. However, Bilò et al. (2017) show 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?"*. This 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 buyer "preselection" as an application of *personalized marketing* (Kim, Kim, Oh, & Ryu, 2010; Weng & Liu, 2004) in which the seller, knowing the buyer valuations for instance by performing market research, advertises about the existence of the supply only to targeted buyers that, once involved, will all consider the allocation fair; the excluded ones then are simply unaware and will not 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 PP with n buyers and m items, sold by a unique seller, we are interested in computing the best possible envy-free solution that can be achieved by removing a subset of buyers from PP. We consider four scenarios arising from the combination of two stability notions, called *market envy-freeness* and *agent envy-freeness*, with the two classical objective functions to be maximised, namely, the social welfare and the seller's revenue.

In a market envy-free allocation, given a pricing of the items, each buyer gets the subset maximizing her utility (given by the difference between the valuation she has for the subset and the prices assigned to the items belonging to it) among all possible subsets that can be created from the set of available items; in an agent envy-free allocation, no buyer gets a better utility by receiving the bundle allocated to any winner. Thus, a market envy-free allocation is also agent envy-free. Moreover, both types of allocations are always guaranteed to exist, as it suffices to assign all items an arbitrarily high price, so that no winner is possible.

It is worth remarking that market envy-freeness is the strongest considered notion of envy-freeness, as it looks to any possible subset of items in the whole market (and for this reason we have called it *market* envy-freeness), while agent envy-freeness only cares about the bundles allocated to other winning buyers, and is therefore called *agent* envy-freeness.

Throughout all the paper, we focus on polynomial time approximation algorithms, i.e., we require that the running time of the algorithms is polynomial in the size of the input instance. We first focus on the most general case, called *market with combinatorial valuations*, in which the buyers' valuations are completely arbitrary. For market 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 $\mathsf{P} = \mathsf{NP}$. On the positive side, under the objective of maximising social welfare, we design an n-approximation algorithm, while, for the case of revenue maximization, we give an $O(n \log m)$-approximation. These results are summarized in the first row of Table 1. In particular, they 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 this 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 by Balcan et al. (2008). There, an $O(\log m)$-approximation is provided, but no lower bound on the problem complexity is given. We show that the problem is NP-hard, thus solving the corresponding open question. We also conjecture (see the discussion at the beginning of Section 3.2 for all details) that the pricing problem with k buyers cannot be approximated up to a factor smaller than k unless $\mathsf{P} = \mathsf{NP}$. As a first evidence of this conjecture, we provide, in Theorem 11, a formal proof for the case of $k = 2$. These results are summarized in Table 2.

Market type	Objective function	Negative results	Positive results
n buyers, combinatorial valuations	Social welfare	no $n^{1-\epsilon}$-approximation, $\forall \epsilon > 0$ (Theorem 6)	n-approximation (Theorem 10)
	Revenue		$O(n \log m)$-approximation (Theorem 10)
n buyers, multi-unit market with identical items	Social welfare		polynomial time solvable (Theorem 21)
	Revenue		
n single-minded buyers, multi-unit market with identical items	Social welfare		**FPTAS** (Theorem 22)
	Revenue	NP-hardness (Theorem 23)	

Table 1: Computational results for the buyer preselection problem under market envy-freeness.

We stress that efficient preselection can be profitable under two orthogonal directions: on the one hand, the removal of a subset of envious buyers may increase the value of the optimal solution; on 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 agent envy-free allocations, we show in Theorems 16 and 17 how to transform in polynomial time any allocation which is agent envy-free only for the subset of the winners to an agent envy-free allocation for all buyers without worsening its performance, with respect to both objective functions. Hence, although this transformation implies that, for agent envy-freeness, buyer preselection cannot improve the performance of stable outcomes, it can be used to map any agent envy-free allocation for a subset of winners obtained through preselection back to an agent envy-free allocation involving all buyers. Therefore, it can be used as an algorithmic tool for computing good stable

Market type	Objective function	Negative results	Positive results
1 buyer, combinatorial valuations	Social welfare		polynomial time solvable (Claim 8)
	Revenue	NP-hardness (Theorem 9)	$O(\log m)$-approximation (Balcan et al., 2008)
2 buyers, combinatorial valuations	Social welfare	no $(2 - \epsilon)$-approximation, $\forall \epsilon > 0$ (Theorem 11)	
	Revenue		

Table 2: Computational results for pricing problems (without preselection) with few buyers under market envy-freeness.

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.

In this respect, consider for instance the case in which the set of buyers of a pricing problem PP can be partitioned in two (or more) subsets A and B, each containing only buyers of the same type, such as, unit-demand buyers or single-minded buyers or buyers with additive valuations. Assume also that the valuations of agents in A and B, when considered separately, allow to compute an r_A-approximate (resp. r_B-approximate) solution of the problem restricted to subset A (resp. B) by exploiting known or new simpler algorithms/techniques. This would give an easy r-approximation for pricing problem PP with $r = r_A + r_B$, by selecting the best of the two solutions obtained when considering the two subsets separately, whereas approximating PP directly would be much more challenging. In fact, consider an optimal solution of value opt and let opt_A and $opt_B = opt - opt_A$ be the contributions to opt due to agents in A and B, respectively. When selecting the best of the two solutions, we obtain a solution of value $v \geq \max\left\{\frac{opt_A}{r_A}, \frac{opt_B}{r_B}\right\} = \max\left\{\frac{opt_A}{r_A}, \frac{opt - opt_A}{r_B}\right\}$ that is minimized when opt_A is such that $\frac{opt_A}{r_A} = \frac{opt - opt_A}{r_B}$, which gives $opt_A = opt\frac{r_A}{r_A + r_B}$ and $v \geq \frac{opt}{r_A + r_B}$.

Envy-freeness	Objective function	Lower bound	Upper bound
market envy-freeness	Social welfare	$m - \epsilon, \forall \epsilon > 0$ (Claim 20)	m (Claim 19)
	Revenue		
agent envy-freeness	Social welfare	$(2 - \epsilon), \forall \epsilon > 0$ (Claim 25)	
	Revenue		2 (Theorem 24)

Table 3: How preselection can improve the objective functions of envy-free solutions in multi-unit markets with identical items. All lower bounds hold even in the case of single-minded buyers.

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 market and agent envy-free solutions. In particular, for market envy-free allocations, we show a tight multiplicative factor of m for both objective functions. For agent 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. These results are summarized in Table 3. We also provide essentially tight results on the

complexity of computing optimal solutions for the buyer preselection problem under market envy-freeness: these results are summarized in the last two rows of Table 1.

1.1.1 COMPARISON BETWEEN OUR AND PREVIOUS RESULTS

In this subsection, we make a comparison between our results and the ones contained in other two papers (Flammini, Mauro, & Tonelli, 2019; Guruswami, Hartline, Karlin, Kempe, Kenyon, & McSherry, 2005) that consider a similar model.

Guruswami et al. (2005) were the first to deal with the problem of computing a market envy-free pricing of maximum revenue[1]. They consider the unit-demand setting where every buyer would like to buy at most one item, and the single-minded one where every buyer is only interested in a particular set of items. For the unit-demand setting they prove, by a reduction from the vertex cover problem, that computing a market envy-free pricing of maximum revenue is APX-hard even for the special case of unlimited supply. Moreover, they provide an $O(\log n)$ approximation algorithm for the limited supply unit-demand setting by using a Walrasian Equilibrium (that is a market envy-free outcome with the additional requirement that the market clears) with reserved prices. For the single-minded setting they prove, by a reduction from the max cut problem, that computing a market envy-free pricing of maximum revenue is APX-hard even for the special case of unlimited supply. Moreover, they present a simple algorithm assigning the same price to every item and prove that it is an $O(\log n + \log m)$-approximation for the unlimited supply case.

Our paper differs from the one by Guruswami et al. (2005) in a variety of aspects. In particular, we consider the more general setting of combinatorial valuations and also focus of the notion of agent envy-freeness and on the objective function of social welfare. More importantly, we deal with the problem of preselecting buyers in order to get higher revenue or social welfare; this problem was not considered by Guruswami et al. We thus provide a new set of results, obtained by exploiting, in our proofs, different arguments and techniques. Specifically, we make use of reductions from the Maximum Independent Set problem, the Satisfiability problem and the Subset Sum problem to show our hardness results, whereas reductions from vertex cover and max cut problems were used by Guruswami et al. Moreover, the provided approximation or exact polynomial time algorithms are completely different from the ones given by Guruswami et al. Finally, we also exploit Walrasian Equilibria, but only when maximizing the social welfare under agent envy-freeness and in a different way with respect to the one considered by Guruswami et al.

Flammini et al. (2019) consider the problem of computing envy-free pricing of maximum revenue for the multi-unit setting where all the items are identical. In particular, they consider social envy-freeness where buyers are the nodes of a given graph and envy is possible only among neighbors in the graph. We notice that social envy-freeness is a relaxation of agent envy-freeness. They provide tight complexity results for single-minded and general buyers valuations. Moreover, they introduce the price of envy-freeness and provide tight bounds for it in the considered setting.

In our paper, as in the one by Flammini et al. (2019), we also consider the multi-unit setting. However, we get a different set of results. In particular, we show how preselection can improve both the revenue and the social welfare of both market and agent envy-free solutions. Moreover, we provide almost tight results on the complexity of computing optimal solutions for the buyer preselection problem under market envy-freeness. None of these problems were considered by Flammini et al. Finally, like other papers dealing with the multi-unit setting (Flammini et al., 2019; Monaco, Sankowski, & Zhang, 2015), in some cases we exploit polynomial reductions to the multiple-choice knapsack problem in order to prove our positive results in a constructive way.

---

1. Given that they only consider the notion of market envy-freeness, they call it simply *envy-freeness* without any further specification.

1.2 Related Work

The literature on envy-free pricing is so vast that it cannot be exhaustively covered here. For this reason, we simply refer to the achievements which are mostly related to the model introduced by Guruswami et al. (2005) we consider in this paper.

For the social welfare maximization without money transfers, Bouveret and Lang (2008) and Karp et al. (2014) consider the problem of assigning items to agents in an agent envy-free way. When money transfers are allowed, the VCG mechanism (Clarke, 1971; Groves, 1973; Vickrey, 1961) provides an optimal solution to the agent 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 by Archer et al. (2003). Also Walrasian Equilibria, which are market envy-free outcomes with the additional requirement that the market clears (Walras, 1954), provide an optimal solution to the problem (Bikhchandani & Mamer, 1997); however, they are guaranteed to exist only under very stringent hypothesis on the buyers' valuation functions (Gul & Stacchetti, 1999).

For the revenue maximization, several studies design logarithmic approximation algorithms for various special cases of market and agent envy-free problems (Balcan et al., 2008; Briest & Krysta, 2006; Cheung & Swamy, 2008; Guruswami et al., 2005; Hartline & Yan, 2011). Related hardness results are alsto studied in the literature (Briest, 2008; Chalermsook, Chuzhoy, Kannan, & Khanna, 2012; Chalermsook, Laekhanukit, & Nanongkai, 2013a, 2013b; Demaine, Feige, Hajiaghayi, & Salavatipour, 2008). Further variants such as large markets (Anshelevich, Kar, & Sekar, 2015), sharp demands where each buyer asks for a fixed quantity of items (Bilò et al., 2017; Chen & Deng, 2010; Chen, Deng, Goldberg, & Zhang, 2016), market with metric substitutability among the items (Chen, Ghosh, & Vassilvitskii, 2011), and buyers with budgets (Colini-Baldeschi, Leonardi, Sankowski, & Zhang, 2014; Feldman, Fiat, Leonardi, & Sankowski, 2012) are also considered.

Feldman et al. (2016) 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, Monaco et al. (2015) study the case of revenue maximization in multi-unit markets under both market and agent envy-freeness when allowing both item and bundle pricing.

In our model, by preselecting buyers, we basically require that all the winners are envy-free (we consider both market and agent envy-freeness). Settings in which, in a similar way, envy-freeness is not guaranteed for all buyers, but only for the winners, are studied in the literature (Alon, Mansour, & Tennenholtz, 2013; Amanatidis, Markakis, & Sornat, 2016; Chen & Rudra, 2008; Elbassioni, Fouz, & Swamy, 2010; Nyman, Su, & Zerbib, 2020; Segal-Halevi & Nitzan, 2019; Segal-Halevi & Suksompong, 2019). In particular, Chen and Rudra (2008) consider "weak" Walrasian Equilibrium, a relaxed version of Walrasian Equilibrium in which the goal is that of maximizing the number of market envy-free buyers, with the condition that all the winners must be market envy-free. Alon et al. (2013) and Amanatidis et al. (2016) consider a relaxed version of market 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. Elbassioni et al. (2010) consider an approximated variant of agent envy-freeness that has to be guaranteed for the winners, in the case of bundle pricing (differently than the one of item pricing addressed in this paper). Nyman et al. (2020) study how to guarantee fairness for a portion of agents when cutting multiple cakes. Segal-Halevi and Nitzan (2019) deal with the division of a cake among groups, with the constraint of

guaranteeing fairness to a given portion of each group. Analogously, Segal-Halevi and Suksompong (2019), study a similar notion of group fairness for the case of indivisible goods.

Moreover, other settings in which related fairness notions are dealt with in a similar way (i.e., by guaranteeing the considered fairness notion only for a portion of agents) are also studied in the literature (Segal-Halevi, 2018; Ortega, 2018; Searns, 2020). In particular, Segal-Halevi considers the problem of redividing an heterogeneous resource in a way that balances fairness with ownership rights, guaranteeing ownership retention to at least for a portion of agents. Ortega deals with the monotonicity of assignments when merging two-sided matching markets, and succeeds to guarantee it for one half of agents. Searns considers the notion of maximin share guarantee when allocating indivisible items, by showing that this fairness notion can be obtained for a fraction of agents.

Finally, some papers deal with multi-unit markets and consider a relaxed definition of market envy-freeness, in which, given a social graph, every agent (corresponding to a node) can only see the prices of items assigned to her neighbors (Flammini, Mauro, & Tonelli, 2018; Flammini et al., 2019). Similarly, several works in fair allocation of goods in absence of pricing assume an individual's view of the subjective well-being as based on a comparison with peers, that is restricting agent envy-freeness constraints to social neighbors (Abebe, Kleinberg, & Parkes, 2017; Chevaleyre, Endriss, & Maudet, 2007b). Moreover, distributed mechanisms for allocating indivisible goods under the absence of central control are investigated in the literature (Chevaleyre, Endriss, Estivie, & Maudet, 2007a; Chevaleyre et al., 2007b; Chevaleyre, Endriss, & Maudet, 2017); in this 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 agent envy-freeness restricted only to social neighbors.

2. Model and Definitions

In this section, we formally define the *buyer preselection problem*. In order to do that, we first have to define *markets*, *stability* notions and *pricing problems*.

2.1 Markets

A *market* with perfect information is a tuple $\Gamma = (N, M, (v_i)_{i \in N})$. We assume there is a seller who owns a set M of m items and wants to sell these items to a set N of n buyers (also called agents). 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 $i \in N$. The seller sets a price for each item. If a buyer buys a subset of items, her utility is the difference between the valuation and the purchase price, i.e., the amount of money she "saves" compared to her valuation. We will formally define the buyers' utility in the next subsection. We assume that a buyer is indifferent between buying one of two or more different subsets of items giving her the same non-negative utility. In particular, a buyer is indifferent between buying a subset of items giving her zero utility and buying no items at all.

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 $X \subseteq M$. A valuation function v is *monotone*, if $v(X) \geq v(X')$ for each $X' \subseteq X$. Finally, in a *multi-unit market with identical 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, $v_i(0) = 0$).

In this paper we mostly focus on markets with combinatorial valuations (Sections 3 and 4) and on multi-unit markets with identical items (Section 5).

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 $j \in M$, $p_j \in \mathbb{R}_{\geq 0} \cup \{\infty\}$ is the price of item j.[2] 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 the *utility* of buyer i when buying X. The *demand set* of buyer i for the price vector $\mathbf{p}$ is the family $D_i(\mathbf{p}) = \mathrm{argmax}_{X \subseteq M} u_i(X, \mathbf{p})$ 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. Given any $i \in N$, a set of items X and a price vector $\mathbf{p}$, the pair $(X, \mathbf{p})$ is *individually-rational for buyer i* if $u_i(X, \mathbf{p}) \geq 0$; given any allocation vector $\mathbf{X} = (X_1, \ldots, X_n)$, an outcome $(\mathbf{X}, \mathbf{p})$ is *individually-rational* if, for every $i \in N$, it holds that the pair $(X_i, \mathbf{p})$ is individually rational for buyer i.

Given an outcome $(\mathbf{X}, \mathbf{p})$, denote by $M(\mathbf{X}) = \bigcup_{i \in N} X_i$ the set of items sold to some buyer according to a feasible allocation vector $\mathbf{X}$. Moreover, buyer i is a *winner* if $X_i \neq \emptyset$; $W(\mathbf{X})$ denotes the set of all winners in $\mathbf{X}$. For an item $j \in M(\mathbf{X})$, denote by $b_{\mathbf{X}}(j)$ 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 $b(j)$.

We say that a buyer i envies a set of items T in outcome $(\mathbf{X}, \mathbf{p})$ if $u_i(X_i, \mathbf{p}) < u_i(T, \mathbf{p})$. The following concepts define two types of stable outcomes for Γ.

Definition 1. *An individually-rational outcome $(\mathbf{X}, \mathbf{p})$ is **market 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 $i \in N$, i.e., no buyer in $(\mathbf{X}, \mathbf{p})$ envies any subset of items.*

Definition 2. *An individually-rational outcome $(\mathbf{X}, \mathbf{p})$ is **agent envy-free** if $u_i(X_i, \mathbf{p}) \geq u_i(X_j, \mathbf{p})$ for every two buyers $i, j \in N$, i.e., no buyer in $(\mathbf{X}, \mathbf{p})$ envies any subset of items being assigned to other buyers.*

Item set Buyer	{1}	{2}	{1,2}
1	1	0	5
2	4	0	0

Table 4: The buyers' valuations of example 2.

Example 2. Consider a market $\Gamma = (N, M, (v_i)_{i \in N})$ with $n = 2$ buyers and $m = 2$ items. The buyers' valuations are summarized in Table 4.

Consider the outcome $(\mathbf{X}, \mathbf{p}) = ((\{1,2\}, \emptyset), (2,3))$ assigning both items to buyer 1 for a total price of 5. First of all, notice that, given the price vector $\mathbf{p}$, buyer 1 is obtaining the most profitable item set. This implies that $(\mathbf{X}, \mathbf{p})$ is agent envy-free, because buyer 2 cannot increase her utility by purchasing the set of items sold to buyer 1. However, $(\mathbf{X}, \mathbf{p})$ is not market envy-free because buyer 2 envies the item set $\{1\}$ that would provide her a utility equal to $v_2(\{1\}) - p_1 = 4 - 2 = 2$.

Consider now the outcome $(\mathbf{X}, \mathbf{p}') = ((\{1,2\}, \emptyset), (4,1))$ with the same allocation vector and a different price vector. Again, on the one hand, given the price vector $\mathbf{p}'$, buyer 1 is obtaining the most profitable item

---

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

set. On the other hand, under price vector $\mathbf{p}'$, buyer 2 cannot obtain a positive utility by choosing a different item set. Therefore, $(\mathbf{X}, \mathbf{p}')$ is market envy-free (and clearly also agent envy-free).

Denote by $\mathsf{MEF}(\Gamma)$ and $\mathsf{AEF}(\Gamma)$ the sets of market envy-free and agent envy-free outcomes for Γ, respectively. Notice that $\mathsf{MEF}(\Gamma) \subseteq \mathsf{AEF}(\Gamma)$ and $\mathsf{MEF}(\Gamma) \neq \emptyset$ (and so also $\mathsf{AEF}(\Gamma) \neq \emptyset$), 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 market envy-free.

In our proofs we sometimes need also the following fundamental notion of stability.

Definition 3. *A* **Walrasian Equilibrium** *is a market envy-free outcome with the additional requirement that the market clears, i.e., every unsold item is assigned price zero.*

We denote by $\mathsf{WE}(\Gamma)$ the set of Walrasian Equilibria for Γ. Notice that $\mathsf{WE}(\Gamma) \subseteq \mathsf{MEF}(\Gamma) \subseteq \mathsf{AEF}(\Gamma)$.

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 $\mathsf{SW}(\mathbf{X}, \mathbf{p}) = \sum_{i \in N} u_i(X_i) + \sum_{j \in M(\mathbf{X})} p_j = \sum_{i \in W(\mathbf{X})} v_i(X_i)$, i.e., it is given by the sum of the utilities of the seller and of all buyers. Note that, since the utility of the seller, i.e., the revenue $\sum_{j \in M(\mathbf{X})} p_j$, is cancelled out by the total price paid by the buyers, social welfare does not depend on the price vector $\mathbf{p}$; moreover, under the assumption of individual rationality, it holds that $\mathsf{SW}(\mathbf{X}, \mathbf{p}) \geq \mathsf{REV}(\mathbf{X}, \mathbf{p})$.

In the following, we formally define the *pricing problem*. Roughly speaking, it is an optimization problem that, given in input a market (i.e., a set of items and a set of buyers with their valuations), an envy-free notion (i.e., market envy-freeness or agent envy-freeness) and an objective function (i.e., social welfare or revenue), returns an allocation of items to buyers and a pricing vector for the items (i.e., an outcome) maximizing the objective function and satisfying the given envy-free notion.

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 $\mathrm{opt}(\Gamma, \mathsf{sol}, \mathsf{obj}) := \mathrm{argmax}_{(\mathbf{X}, \mathbf{p}) \in \mathsf{sol}(\Gamma)} \{\mathsf{obj}(\mathbf{X}, \mathbf{p})\}$ denote the set of outcomes in $\mathsf{sol}(\Gamma)$ maximizing the objective function obj.

Definition 4. *The* **pricing problem** $\mathsf{PP} = (\Gamma, \mathsf{sol}, \mathsf{obj})$ *is an optimization problem which, given a market Γ, an envy-freeness notion* sol *and an objective function* obj*, asks for an outcome* $o^*(\mathsf{PP}) \in \mathrm{opt}(\mathsf{PP})$.

In this paper, we consider the cases in which $\mathsf{sol} \in \{\mathsf{MEF}, \mathsf{AEF}\}$ and $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$.

2.4 Oracles

Fix a pricing problem $\mathsf{PP} = (\Gamma, \mathsf{sol}, \mathsf{obj})$. As we have seen, when Γ is a market with combinatorial valuations, any algorithm for PP 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 in this paper exploiting oracles are polynomial also in the sense that they call the oracle a polynomial number of times. Therefore, when the oracle is polynomially computable, the algorithm computation is polynomial as well.

2.5 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 $\Gamma(N') = (N', M, (v_i)_{i \in N'})$.

We are now ready to define the *buyer preselection problem*. Roughly speaking, it is defined upon a pricing problem $\mathsf{PP} = (\Gamma, \mathsf{sol}, \mathsf{obj})$, with $\Gamma = (N, M, (v_i)_{i \in N})$, and returns a pair (N^*, o^*). In particular, it aims at selecting a subset $N^* \subseteq N$ of buyers to admit to the market so that submarket $\Gamma(N^*)$ induced by N^* is such that the returned outcome o^* maximizes, among all possible choices of the buyers to admit to the market, the objective function (i.e., social welfare or revenue) while satisfying the given envy-free notion.

Definition 5. *The* **buyer preselection problem** *is an optimization problem* $\mathsf{BPP}(\mathsf{PP})$ *which, given a pricing problem* $\mathsf{PP} = (\Gamma, \mathsf{sol}, \mathsf{obj})$ *with* $\Gamma = (N, M, (v_i)_{i \in N})$, *asks for a pair* $(N^*(\mathsf{BPP}(\mathsf{PP})), o^*(\mathsf{BPP}(\mathsf{PP})))$ *such that* $N^*(\mathsf{BPP}(\mathsf{PP})) \in argmax_{N' \subseteq N}\{\mathrm{opt}(\Gamma(N'), \mathsf{sol}, \mathsf{obj})\}$ *and* $o^*(\mathsf{BPP}(\mathsf{PP})) \in \mathrm{opt}(\Gamma(N^*(\mathsf{BPP}(\mathsf{PP}))), \mathsf{sol}, \mathsf{obj})$, *that is,* $o^*(\mathsf{BPP}(\mathsf{PP}))$ *is the best outcome which can be realized in all possible submarkets of* Γ.

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

Coming back to Example 1, it is worth noticing that, without considering preselection, all market envy-free outcomes cannot sell (both) items to the first buyer: in this case, the price should be at most $\frac{3}{2}$ and the second buyer would be envious. Therefore, market envy-free outcomes can sell an item only to the second buyer at price $\frac{3}{2} < p \leq 2$. In particular, an outcome maximizing social welfare and revenue is $o = ((\emptyset, \{1\}), 2)$, i.e., the one selling a copy of the item only to the second buyer at price 2: it holds that $\mathsf{REV}(o) = \mathsf{SW}(o) = 2$. If we exploit buyer preselection, we can exclude from the market the second buyer, and we can sell both items to the first buyer at price $\frac{3}{2}$, thus obtaining outcome $o^* = ((\{1, 2\}), \frac{3}{2})$: in such a way, we obtain $\mathsf{REV}(o^*) = \mathsf{SW}(o^*) = 3$, thus increasing the performance of envy-free solutions. In Section 5 it is shown that preselection can be profitable, in terms of social welfare and revenue, also when considering agent envy-free outcomes.

3. Results for Market Envy-Free Outcomes

In this section, we consider the buyer preselection problems $\mathsf{BPP}(\Gamma, \mathsf{MEF}, \mathsf{SW})$ and $\mathsf{BPP}(\Gamma, \mathsf{MEF}, \mathsf{REV})$ in their most general form, that is, when Γ is a market with combinatorial valuations. We start by providing a severe lower bound on their best-possible approximation guarantee. We achieve this result by designing an approximation-preserving reduction from the Maximum Independent Set problem, which cannot be approximated within a factor of $n^{1-\epsilon}$ unless $\mathsf{P} = \mathsf{NP}$ (Zuckerman, 2007). Interestingly, this lower bound holds even under the more stringent assumption of single-minded buyers.

Theorem 6. *Let* $\mathsf{PP} = (\Gamma, \mathsf{MEF}, \mathsf{obj})$ *be a pricing problem with* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$ *and* Γ *being a market with combinatorial valuations. For every* $\epsilon > 0$, *the buyer preselection problem* $\mathsf{BPP}(\mathsf{PP})$ *cannot be approximated within* $n^{1-\epsilon}$, *unless* $\mathsf{P} = \mathsf{NP}$, *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 a market 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 $\Gamma(V^*)$. This implies

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

Assume, for the sake of contradiction, that there exists an approximation algorithm for $\mathsf{BPP}(\mathsf{PP})$ returning an outcome o such that $n^{1-\epsilon}\mathsf{obj}(o) \geq \mathsf{obj}(o^*(\mathsf{BPP}(\mathsf{PP})))$ 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. $\qquad \square$

As a positive result, we show that, by building upon (approximation) algorithms for pricing problems defined on markets with a constant number of buyers, it is possible to obtain approximation algorithms for the buyer preselection problem. The key ingredient is provided by the following lemma.

Lemma 7. *Given a buyer preselection problem* $\mathsf{BPP}(\mathsf{PP})$*, where* $\mathsf{PP} = (\Gamma, \mathsf{MEF}, \mathsf{obj})$ *with* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$*, and a positive integer* $k \leq n$*, if there exists a polynomial time algorithm* $\mathcal{A}_k$ *running in time* $T(k)$ *and returning an outcome for the pricing problem defined on markets with* k *buyers whose objective value is at least an* α_k *fraction of the optimal social welfare, then* $\mathsf{BPP}(\mathsf{PP})$ *admits an* $\left(\lceil \frac{n}{k} \rceil \alpha_k\right)$*-approximation algorithm running in time* $O(n^k)T(k)$*.*

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

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

so that $\mathsf{obj}_i(o^*(\mathsf{BPP}(\mathsf{PP})))$ is the contribution of buyer i to the objective value in the optimal solution of $\mathsf{BPP}(\mathsf{PP})$. Let

$$N_k^* = \begin{cases} N^*(\mathsf{BPP}(\mathsf{PP})) & \text{if } |N^*(\mathsf{BPP}(\mathsf{PP}))| \leq k, \\ \mathsf{argmax}_{N \subseteq N^*(\mathsf{BPP}(\mathsf{PP})):|N| \leq k} \sum_{i \in N} v_i(X_i^*) & \text{otherwise,} \end{cases}$$

be a subset of $N^*(\mathsf{BPP}(\mathsf{PP}))$ having cardinality $\min\{k, |N^*(\mathsf{BPP}(\mathsf{PP}))|\}$ and maximizing the overall valuation of the allocated bundles. Roughly speaking, given an optimal outcome $o^*(\mathsf{BPP}(\mathsf{PP}))$ of the buyer preselction problem, N_k^* just maximizes the sum of valuations of allocated bundles among all subsets of at most k (preselected) buyers, independently from the considered objective function obj. For $k \geq |N^*(\mathsf{BPP}(\mathsf{PP}))|$, we get

$$\sum_{i \in N^*(\mathsf{BPP}(\mathsf{PP}))} v_i(X_i^*) \leq \left\lceil \frac{|N^*(\mathsf{BPP}(\mathsf{PP}))|}{k} \right\rceil \sum_{i \in N_k^*} v_i(X_i^*), \tag{4}$$

by definition. When $k < |N^*(\mathsf{BPP(PP)})|$, set $N^*(\mathsf{BPP(PP)})$ can be partitioned into $\left\lceil \frac{|N^*(\mathsf{BPP(PP)})|}{k} \right\rceil$ subsets each having cardinality k, except for at most one set which we denote by S. Clearly, any set S' of cardinality k satisfies $\sum_{i \in S'} v_i(X_i^*) \leq \sum_{i \in N_k^*} v_i(X_i^*)$ by the definition of N_k^*. Moreover, also $\sum_{i \in S} v_i(X_i^*) \leq \sum_{i \in N_k^*} v_i(X_i^*)$ holds as S can be extended to a set of cardinality k by adding $k - |S|$ buyers from $N_k^* \setminus S$ without worsening the overall valuation of the allocated bundles. Thus, we conclude that (4) holds for any value of k.

Consider algorithm Best which, given Γ and k, calls $\mathcal{A}_k$ on all $O(n^k)$ possible subsets of k buyers and returns the solution, let us call it o_{Best}, with the largest objective value. We show that the approximation guarantee of o_{Best} is $\lceil \frac{n}{k} \rceil \alpha_k$. Let $o \in \mathsf{MEF}(\Gamma(N_k^*))$ be the outcome returned by $\mathcal{A}_k$ when executed on the pricing problem $(\Gamma(N_k^*), \mathsf{MEF}, \mathsf{obj})$. We have

$$
\begin{aligned}
\mathsf{obj}(o^*(\mathsf{BPP(PP)})) &= \sum_{i \in N^*(\mathsf{BPP(PP)})} \mathsf{obj}_i(o^*(\mathsf{BPP(PP)})) \\
&\leq \sum_{i \in N^*(\mathsf{BPP(PP)})} v_i(X_i^*) \\
&\leq \left\lceil \frac{n}{k} \right\rceil \sum_{i \in N_k^*} v_i(X_i^*) \\
&\leq \left\lceil \frac{n}{k} \right\rceil \alpha_k \mathsf{obj}(o) \\
&\leq \left\lceil \frac{n}{k} \right\rceil \alpha_k \mathsf{obj}(o_{\mathsf{Best}}),
\end{aligned}
$$

where the first inequality comes from (3) and the fact that, because of individual rationality, $\sum_{j \in X_i^*} p_j^* \leq v_i(X_i^*)$; the second inequality follows from (4) together with $|N^*(\mathsf{BPP(PP)})| \leq n$; the third inequality comes from the approximation guarantee of algorithm $\mathcal{A}_k$; and the last inequality holds as Best returns the outcome with highest objective value. $\qquad\square$

As a consequence of Theorem 6 and Lemma 7, the buyer preselection problem $\mathsf{BPP}(\Gamma, \mathsf{MEF}, \mathsf{obj})$, with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, admits a polynomial time algorithm asymptotically achieving the best possible approximation guarantee (up to subpolynomial factors), whenever there is a constant value k for which the pricing problem defined on markets with k buyers admits a polynomial time algorithm computing a market envy-free outcome providing a k-approximation of the optimal social welfare. Clearly, the value of k for which better approximation guarantees are ideally expected is $k = 1$. For this reason, in the following subsection, we focus on the solution of pricing problems with a unique buyer.

3.1 Pricing Problems Defined on Markets with a Unique Buyer

Throughout this subsection, we assume that there is only one buyer in the market. Therefore, for the sake of simplicity, we remove the subscript 1 from the notation.

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

Claim 8. *Let Γ be a market with a single buyer and combinatorial valuations. The pricing problem $\mathsf{PP} = (\Gamma, \mathsf{MEF}, \mathsf{SW})$ can be solved in polynomial time.*

In fact, observe that an outcome $(X^*, \mathbf{0}^m)$ such that $X^* \in \mathrm{argmax}_{X \subseteq M} v(X)$ is such that $(X^*, \mathbf{0}^m) \in \mathrm{opt}(\mathsf{PP})$, and a set $X^* \in \mathrm{argmax}_{X \subseteq M} v(X)$ 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 $\mathsf{obj} = \mathsf{REV}$ is NP-hard. Before our contribution, only an approximation algorithm was known (Balcan et al., 2008, Lemma 7). Towards this end, we exploit a

polynomial time reduction from 3SAT, in which a given boolean formula ϕ is transformed into a market with a unique buyer and monotone combinatorial valuations, whose items are the literals of ϕ.

Theorem 9. *Let Γ be a market with a single buyer and combinatorial monotone valuations. The pricing problem* $\mathsf{PP} = (\Gamma, \mathsf{MEF}, \mathsf{REV})$ *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 literals on variables in $\mathsf{V}(\phi)$, i.e., either a variable or its negation. Denote $\nu = |\mathsf{V}(\phi)|$, so that $|\mathsf{L}(\phi)| = 2\nu$. Throughout this proof, we assume $\nu \geq 4$. We shall denote by v_i the ith variable in $\mathsf{V}(\phi)$ and by x_i and $\tilde{x}_i$ the two literals related to v_i. An assignment for ϕ is a function $f : \mathsf{V}(\phi) \to \{0, 1\}$ associating each variable of ϕ with 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 the following rule: if $f(v_i) = 1$, then $x_i = 1$ and $\tilde{x}_i = 0$, whereas, if $f(v_i) = 0$, then $x_i = 0$ and $\tilde{x}_i = 1$. An assignment f can be also defined by listing the ν literals which are evaluated 1 according to f. ϕ is satisfiable if there exists an assignment $f \in F(\phi)$ such that $\phi(f) = 1$ and is not satisfiable if, for every assignment $f \in F(\phi)$, $\phi(f) = 0$. Given a set of literals $X \subseteq \mathsf{L}(\phi)$, we say that X *represents* $\mathsf{V}(\phi)$ if X contains at least one literal for each variable in $\mathsf{V}(\phi)$. Observe that any set X representing $\mathsf{V}(\phi)$ needs have at least ν elements. 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 $C = \{c_1, \ldots, c_\gamma\}$ of its clauses, where each clause is a set of three literals.

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

$$
v(X) = \begin{cases}
\nu & \text{if } X \text{ represents } \mathsf{V}(\phi), \\
3 + \epsilon & \text{if } X \text{ does not represent } \mathsf{V}(\phi) \text{ and } X \text{ contains a clause in } C, \\
0 & \text{if } X = \emptyset, \\
1 & \text{otherwise,}
\end{cases}
$$

where $\epsilon > 0$ is arbitrarily small. Observe that, by the definitions of ϵ and ν, the valuation of a set X either stays the same or increases when adding new items, so that v is monotone.

For example, the instance $C = \{\{x_1, x_2, \tilde{x}_3\}, \{\tilde{x}_1, x_2, \tilde{x}_4\}, \{\tilde{x}_2, x_3, \tilde{x}_4\}, \{x_1, x_3, x_4\}\}$ defines the boolean formula

$$
\phi := (x_1 \vee x_2 \vee \tilde{x}_3) \wedge (\tilde{x}_1 \vee x_2 \vee \tilde{x}_4) \wedge (\tilde{x}_2 \vee x_3 \vee \tilde{x}_4) \wedge (x_1 \vee x_3 \vee x_4),
$$

which has four variables $\mathsf{V}(\phi) = \{v_1, v_2, v_3, v_4\}$ and eight literals $\mathsf{L}(\phi) = \{x_1, \tilde{x}_1, x_2, \tilde{x}_2, x_3, \tilde{x}_3, x_4, \tilde{x}_4\}$. The assignment $f_1 = \{x_1, x_2, x_3, x_4\}$ satisfies ϕ as $f_1 \cap c_j \neq \emptyset$ for each $j \in [\gamma]$, while the assignment $f_2 = \{\tilde{x}_1, \tilde{x}_2, x_3, x_4\}$ does not satisfy ϕ, as it does not intersect c_1. The set of items $X_1 = \{x_1, x_2, \tilde{x}_2, x_3, x_4\}$ represents $\mathsf{V}(\phi)$ as $X_1 \cap \{x_i, \tilde{x}_i\} \neq \emptyset$ for each $i \in [\nu]$, while the set of items $X_2 = \{x_1, x_2, \tilde{x}_2, x_3, \tilde{x}_3\}$ does not represent $\mathsf{V}(\phi)$, as $X_2 \cap \{x_4, \tilde{x}_4\} = \emptyset$. Moreover, X_2 contains clause c_1. Thus, $v(X_1) = \nu$ and $v(X_2) = 3 + \epsilon$. Finally, $X_3 = \{x_1, x_2, \tilde{x}_2, x_3\}$ neither represents $\mathsf{V}(\phi)$ nor contains a clause in C, so that $v(X_3) = 1$.

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 first that both ν and γ are polynomial in the representation of ϕ. Hence, given a set of literals X, determining whether it contains a clause in C, or whether it represents $\mathsf{V}(\phi)$ can be performed in polynomial time, so that value queries can be efficiently answered. In order to provide an efficient response to a demand query, we first observe that, by the definition of v, the answer to a demand query is a set X satisfying exactly one of the following four properties: (1) X represents $\mathsf{V}(\phi)$, (2) X contains a clause in C but does not represent $\mathsf{V}(\phi)$, (3) X is the empty set; (4) X is non-empty and neither

represents $V(\phi)$ nor contains a clause in C. Moreover, as item prices are non-negative and the valuation of all sets satisfying the same property is the same, a set X_i^* maximizing the buyer's utility among all the ones satisfying properties (1), (2) and (4) can be obtained by selecting the ν items obtained by choosing, for each variable in $V(\phi)$, the related literal having the lowest price (doable in $O(\nu)$ time), the three items of smallest total price among all the triplets belonging to a clause in C (doable in $O(\gamma)$ time) and the item of smallest price (doable in $O(\nu)$ time), respectively. It follows that a response to the demand query can be done by selecting the set providing the higher buyer's utility among the empty set (giving a utility equal to 0) and the sets maximizing the utility among all the ones satisfying properties (1), (2) and (4), computed as described above. Thus, market Γ can be generated and managed in polynomial time.

With respect to the example outlined above, consider the pricing vector $\mathbf{p} = (2, 3, 1, 4, 2, 1, 0, 2)$ for the set of items $M = L(\phi) = \{x_1, \tilde{x}_1, x_2, \tilde{x}_2, x_3, \tilde{x}_3, x_4, \tilde{x}_4\}$. The item yielding the highest utility is x_4 for which $u(\{x_4\}, \mathbf{p}) = 1 - 0 = 1$. The set of items containing a clause and not representing $V(\phi)$ yielding the highest utility is c_4 for which $u(c_4, \mathbf{p}) = 3 + \epsilon - 4 = -1 + \epsilon$. The set of items representing $V(\phi)$ yielding the highest utility is $X = \{x_1, x_2, \tilde{x}_3, x_4\}$ for which $u(X, \mathbf{p}) = 4 - 4 = 0$. Finally, the empty set yields a utility equal to 0. Thus, the demand query answers by returning bundle $\{x_4\}$.

Now, in order to complete the proof, we show that there exists a market envy-free outcome $o \in MEF(\Gamma)$ such that $REV(o) = \nu$ 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 $L(\phi) \setminus X$ are priced ∞. By definition, X represents $V(\phi)$, thus, we have $u(X, \mathbf{p}) = 0$, so that $((X), \mathbf{p})$ is individually-rational. Moreover, $X \in D(\mathbf{p})$. In fact, for every $X' \neq X$ such that X' represents $V(\phi)$, $u(X', \mathbf{p}) < 0$ since X' has to contain at least one item priced ∞; for every X' such that X' does not represents $V(\phi)$ and X' contains a clause 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 non-empty X' such that X' neither represents $V(\phi)$ nor contains a clause in C, we have $v(X') = 1$, so that $u(X', \mathbf{p}) \leq 0$. Hence, $(X, \mathbf{p}) \in MEF(\Gamma)$ and $REV(X, \mathbf{p}) = \nu$.

Secondly, assume that there exists an outcome $(X, \mathbf{p}) \in MEF(\Gamma)$ such that $REV(X, \mathbf{p}) = \nu$. By construction of the valuation function, this is possible only if X represents $V(\phi)$ and $\sum_{j \in X} p_j = \nu$, so that $u(X, \mathbf{p}) = 0$. However, since $v(\{\ell\}) = 1$ for every $\ell \in L(\phi)$, $(X, \mathbf{p}) \in MEF(\Gamma)$ 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 not satisfiable. This implies that the assignment induced by all literals in $L(\phi) \setminus X$ 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 $(X, \mathbf{p}) \in MEF(\Gamma)$. Hence, ϕ has to be satisfiable. $\qquad\square$

On the positive side, Balcan et al. (2008) 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 7.

Hence, by combining Lemma 7 with Claim 8 and the result by Balcan et al. (2008), we obtain the following upper bounds.

Theorem 10. *Let Γ be a market with combinatorial valuations and let* $PP = (\Gamma, MEF, obj)$ *be a pricing problem with* $obj \in \{REV, SW\}$. *The buyer preselection problem* $BPP(PP)$ *admits an n-approximation when* $obj = SW$, *and an $O(n \log m)$-approximation when* $obj = REV$.

It is worth noticing that, given the proof of Lemma 7 when $k = 1$, the preselection claimed in Theorem 10 is somehow "oblivious", as it can be obtained by exploiting the minimum possible number of oracle queries. In fact, to determine the buyer having the highest valuation for a single bundle, a single demand query is sufficient for each buyer (it suffices asking for a bundle in the demand set when all prices are set equal to zero).

In light of the lower bound given in Theorem 6, the upper bounds given in Theorem 10 are (unless $P = NP$) asymptotically tight up to subpolynomial multiplicative factors (notice that even for $obj = SW$, the lower bound of $n^{1-\epsilon}$ for any $\epsilon > 0$ may be a subpolynomial multiplicative factor far from the upper bound of n).

3.2 Pricing Problems Defined on Markets with Few Buyers

From the inapproximability result of Theorem 6 it follows that, given any $\epsilon > 0$, for any $n \geq n(\epsilon)$, there exists a preselection problem with n buyers which cannot be approximated up to $n^{1-\epsilon}$ unless $P = NP$. Conversely, by Lemma 7 it follows that an $\left\lceil \frac{n}{k} \right\rceil \alpha_k$ approximation is possible if the pricing problem (without preselection) on markets with k buyers, k being a constant, admits an α_k-approximation (satisfying some additional assumptions).

It is worth noticing that a stronger negative result on the complexity of the preselection problem, stating that, given any $\epsilon > 0$, for any $n \geq n(\epsilon)$, there exists a preselection problem with n buyers which cannot be approximated up to $n - \epsilon$ (unless $P = NP$), would imply that the pricing problems (without preselection) with k buyers cannot be approximated up to a factor smaller than k (unless $P = NP$), even for every constant value of k. In fact, if the considered stronger result about the preselection problem held, it would follow that, given any $\epsilon > 0$, for any $n \geq n(\epsilon)$ such that k divides n and any integer $k \geq 1$, $\left\lceil \frac{n}{k} \right\rceil \alpha_k > n - \epsilon$, which implies $\alpha_k > k$.

Although these results cannot be formally claimed, we conjecture they may be true. As a first evidence of our conjectures, we provide, in the following, a formal proof for the case of $k = 2$, leaving the extension to any value of k as an open problem. The proof of this result builds on the ideas exploited in the proof of Theorem 9 with some additional technicalities and extends to even agent envy-freeness.

Theorem 11. *Let Γ be a market with two buyers and combinatorial monotone valuations. The pricing problem $PP = (\Gamma, sol, obj)$, with $sol \in \{MEF, AEF\}$ and $obj \in \{REV, SW\}$, cannot be approximated up to a factor $\alpha < 2$ unless $P = NP$.*

Proof. Again, we prove the claim through a reduction from 3SAT (see the proof of Theorem 9 for basic definitions and notation).

Given an instance $\phi := C$ of 3SAT, we construct a market $\Gamma = ([2], M, v_1, v_2)$ with two buyers as follows. The set of items M is obtained by creating γ different copies of each literal, one for each of the γ clauses, and then associating a different item to every copy; so $M = \{x_i^j : i \in [\nu], j \in [\gamma]\} \cup \{\tilde{x}_i^j : i \in [\nu], j \in [\gamma]\}$ and $|M| = 2\nu\gamma$, where x_i^j and $\tilde{x}_i^j$ are the jth copies of the two literals associated with variable $v_i \in V(\phi)$.

For example, if $C = \{c_1, c_2\} = \{\{x_1, x_2, x_3\}, \{\tilde{x}_1, x_2, \tilde{x}_3\}\}$, we have $V(\phi) = \{v_1, v_2, v_3\}$, $L(\phi) = \{x_1, \tilde{x}_1, x_2, \tilde{x}_2, x_3, \tilde{x}_3\}$ and $M = \{x_1^1, x_1^2, \tilde{x}_1^1, \tilde{x}_1^2, x_2^1, x_2^2, \tilde{x}_2^1, \tilde{x}_2^2, x_3^1, x_3^2, \tilde{x}_3^1, \tilde{x}_3^2\}$.

Given a set of items $X \subseteq M$, we say that X *represents* $V(\phi)$ if, for each variable $v_i \in V(\phi)$, X contains all items associated to the γ copies of one of the two literals associated to v_i, that is, for each $i \in [\nu]$, either $\{x_i^1, \ldots x_i^\gamma\} \subseteq X$ or $\{\tilde{x}_i^1, \ldots \tilde{x}_i^\gamma\} \subseteq X$. Thus, any set X representing $V(\phi)$ needs have at least $\nu\gamma$ items. Similarly, we say that X *represents* C if, for each clause $c_j \in C$, X contains the jth copy of at least one of the three literals contained in c_j. For instance, clause $c_1 = \{x_1, x_2, x_3\}$ is represented by $X = \{x_1^1\}$, but is not represented by $X = \{x_1^2, \tilde{x}_2^1\}$. As, by construction, each clause has its dedicated copies of the literals, no item can represent more than one clause. Hence, any set X representing C needs have at least γ items. The valuation functions of the two buyers are defined as follows:

$$v_1(X) = \begin{cases} 1 & \text{if } X \text{ represents } V(\phi), \\ 0 & \text{otherwise.} \end{cases}$$

$$v_2(X) = \begin{cases} 1 & \text{if } X \text{ represents } C \text{ and } |X| \geq \nu\gamma, \\ 0 & \text{otherwise.} \end{cases}$$

Observe that, for both buyers, the valuation of a set X either stays the same or increases when adding new items, so that both v_1 and v_2 are monotone.

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 implement an oracle which can answer demand and value queries in polynomial time. To this aim, observe first that both ν and γ are polynomial in the representation of ϕ. Hence, given a set of items X, determining whether it represents either $\mathsf{V}(\phi)$ or C can be performed in polynomial time, so that value queries can be efficiently answered for both buyers. In order to provide an efficient answer to a demand query for a given price vector $\mathbf{p}$, we consider the two buyers separately.

For buyer 1, observe that, by the definition of v_1, the answer to a demand query is either a set X representing $\mathsf{V}(\phi)$ or the empty set. Moreover, as item prices are non-negative and the valuation of all sets representing $\mathsf{V}(\phi)$ is the same under v_1, a set X^* maximizing the buyer's utility among all the ones representing $\mathsf{V}(\phi)$ can be obtained by selecting, for each variable in $\mathsf{V}(\phi)$, the literal whose copies amount for the lowest total price. Formally, for each $i \in [\nu]$, we choose $\{x_i^1, \ldots, x_i^\gamma\}$ if $\sum_{j \in [\gamma]} p(x_i^j) \leq \sum_{j \in [\gamma]} p(\tilde{x}_i^j)$ and $\{\tilde{x}_i^1, \ldots, \tilde{x}_i^\gamma\}$ otherwise. This choice can be performed in $O(\nu\gamma)$ time.

For buyer 2, observe that, by the definition of v_2, the answer to a demand query is either a set X representing C and such that $|X| \geq \nu\gamma$ or the empty set. Moreover, as item prices are non-negative and the valuation of all sets representing representing C and having cardinality at least $\nu\gamma$ is the same under v_2, a set X^* maximizing the buyer's utility among all the ones representing C and such that $|X^*| \geq \nu\gamma$ can be obtained by selecting, for each clause $c_j \in C$ and each literal in c_j, the cheapest among their jth copies and then the $\nu\gamma - \gamma$ cheapest items among the remaining ones. Formally, for each $j \in [\gamma]$, if $c_j = \{\ell_1, \ell_2, \ell_3\}$, any item belonging to $\arg\min\{p(\ell_i^j) : i \in [3]\}$ is chosen. Then, the cheapest items among the leftovers are added. The entire selection can be performed in $O(\nu\gamma)$ time, after having ordered the price vector.

Thus, market Γ can be generated and managed in polynomial time.

Now, in order to complete the proof, we show that there exists a market envy-free outcome $o \in \mathsf{MEF}(\Gamma)$ such that $\mathsf{SW}(o) = \mathsf{REV}(o) = 2$ if ϕ is satisfiable and that, whenever ϕ is not satisfiable, for any outcome $o \in \mathsf{OUT}(\Gamma)$, we have $\mathsf{SW}(o) \leq 1$ and $\mathsf{REV}(o) \leq 1$. As market envy-free outcomes are also agent envy-free ones, it follows that any α-approximation algorithm for the pricing problem $(\Gamma, \mathsf{sol}, \mathsf{obj})$, with $\mathsf{sol} \in \{\mathsf{MEF}, \mathsf{AEF}\}$, $\mathsf{obj} \in \{\mathsf{SW}, \mathsf{REV}\}$ and $\alpha < 2$, could be used to solve 3SAT, thus implying $\mathsf{P} = \mathsf{NP}$.

Assume first that ϕ is satisfiable and let f be a satisfying assignment for ϕ. Let L' be the set of literals which are false in f and let X denote the set of items corresponding to all the copies of literals in L', so that $|X| = \nu\gamma$. Let $\mathbf{p}$ be the pricing vector in which all items are priced $\frac{1}{\nu\gamma}$. By definition, X represents $\mathsf{V}(\phi)$, thus, we have $u_1(X, \mathbf{p}) = 0$, so that $(X, \mathbf{p})$ is individually rational for buyer 1. Moreover, as every set of items representing $\mathsf{V}(\phi)$ needs have at least $\nu\gamma$ items, none of them can give buyer 1 a utility larger than zero. Thus, $X \in D_1(\mathbf{p})$. As f is a satisfying assignment, by definition of L', it follows that $\mathsf{L}(\phi) \setminus \mathsf{L}'$ is a set of literals satisfying ϕ. This implies that, for any clause $c_j \in C$, there exists a literal $\ell \in \mathsf{L}(\phi) \setminus \mathsf{L}'$ such that $\ell \in c_j$. Thus, the jth copy of ℓ can represent clause c_j. As this property holds for each $c_j \in C$, set $M \setminus X$ represents C and satisfies $|M \setminus X| = \nu\gamma$. It follows that $u_2(M \setminus X, \mathbf{p}) = 0$, so that $(M \setminus X, \mathbf{p})$ is individually rational for buyer 2, and therefore outcome $o = ((X, M \setminus X), \mathbf{p})$ is individually-rational. Moreover, as every set of items yielding a non-zero valuation to buyer 2 needs have at least $\nu\gamma$ items, none of them can give buyer 2 a utility larger than zero. Thus, $M \setminus X \in D_2(\mathbf{p})$. Hence, we conclude that $o \in \mathsf{MEF}(\Gamma)$ and $\mathsf{SW}(o) = \mathsf{REV}(o) = 2$.

Now assume that ϕ is not satisfiable. We claim that, for any set of items X representing $\mathsf{V}(\phi)$, there cannot exist a set $Y \subseteq M \setminus X$ representing C. Assume the contrary, by way of contradiction. That is, there exists a set of items $Y \subseteq M \setminus X$ such that, for each $c_j \in C$, Y contains the jth copy of a literal occurring in c_j. This implies that there exists a set of literals $\mathsf{L}' \in \mathsf{L}(\phi)$ such that $\mathsf{L}' \cap c_j \neq \emptyset$ for each $j \in [\gamma]$. As X represents $\mathsf{V}(\phi)$ and $X \cap Y = \emptyset$, it follows that, for any variable $v_i \in \mathsf{V}(\phi)$, L' cannot contain both x_i and $\tilde{x}_i$. Thus, by setting to true every literal in L', we obtain a satisfying assignment for ϕ, contradicting the

assumption that ϕ is not satisfiable. We conclude that, when ϕ is not satisfiable, any outcome $o \in \mathsf{OUT}(\Gamma)$, no matter whether o meets some envy-freeness property or not, cannot assign to both buyers a set of items of value 1. Thus, $\mathsf{SW}(o) \leq 1$ and $\mathsf{REV}(o) \leq 1$. $\qquad\square$

4. Results for Agent Envy-Free Outcomes

In this section, we consider the buyer preselection problem $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{obj})$ with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$ and Γ being a market with combinatorial valuations. Since we deal with agent envy-free outcomes, we suppose that the price of any unsold item is infinite. Formally, given a solution (N, o), where outcome $o = (\mathbf{X}, \mathbf{p})$, for the buyer preselection problem $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{obj})$ 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 $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$. We are going to show how it is possible to transform an agent envy-free solution $(\bar{N}, \bar{o})$ with preselection into anther one (without preselection) having a non-smaller revenue.

In order to do that, we first need an additional definition and some lemmas.

Definition 12. *A solution* $(\bar{N}, \bar{o})$, *where* $\bar{o} = (\bar{\mathbf{X}}, \bar{\mathbf{p}})$, *for the buyer preselection problem* $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$ *is called* **maximal** *if there is no solution* $(\bar{N}, \bar{o}')$ *for the same buyer preselection problem* $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$ *with the same subset of preselected agents* $\bar{N}$ *where outcome* $\bar{o}' = (\bar{\mathbf{X}}, \bar{\mathbf{p}}')$ *is such that* $\bar{p}'_j \geq \bar{p}_j$, *for any* $j \in M$, *and there is at least an item* $k \in M(\bar{\mathbf{X}})$ *such that* $\bar{p}'_k > \bar{p}_k$.

Informally, a solution $(\bar{N}, \bar{o})$ for the buyer preselection problem $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$ (i.e., where AEF holds for agents in $\bar{N}$) is maximal if it is not possible to obtain another one with the same allocation vector, in which we increase the price of at least one sold item while keeping the prices of the remaining items unchanged.

The following claim shows that any non-maximal solution can be efficiently transformed into another one (with the same allocation vector) being maximal.

Claim 13. *Given a solution* $(\bar{N}, \bar{o})$, *where* $\bar{o} = (\bar{\mathbf{X}}, \bar{\mathbf{p}})$, *for the buyer preselection problem* $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$, *it is always possible to compute in polynomial time a solution* $(\bar{N}, \bar{o}')$ *for the buyer preselection problem* $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$ *with the same allocation vector, i.e., with* $\bar{o}' = (\bar{\mathbf{X}}, \bar{\mathbf{p}}')$, *such that* $\mathsf{REV}(\bar{o}') \geq \mathsf{REV}(\bar{o})$.

Proof. Given an allocation vector $\bar{\mathbf{X}} = (\bar{X}_1, \ldots, \bar{X}_{|\bar{N}|})$, consider the following linear program.

$$
\begin{aligned}
\max \quad & \sum_{i \in \bar{N}} \sum_{j \in \bar{X}_i} \bar{p}_j \\
\text{subject to} \quad & v_i(\bar{X}_i) - \sum_{j \in \bar{X}_i} \bar{p}_j \geq v_i(\bar{X}_z) - \sum_{j \in \bar{X}_z} \bar{p}_j, \quad i \in \bar{N}, z \in \bar{N} \\
& v_i(\bar{X}_i) - \sum_{j \in \bar{X}_i} \bar{p}_j \geq 0, \quad\quad\quad\quad\quad\; i \in \bar{N} \\
& \bar{p}_j \geq 0 \quad\quad\quad\quad\quad\quad\quad\quad\quad\quad j = 1, \ldots, m
\end{aligned}
$$

Notice that, since we start from an agent envy-free solution, the linear program always admits a feasible solution. It is easy to see that, by solving the linear program, we get a pricing vector for the allocation $\bar{\mathbf{X}}$ such that the obtained solution is agent envy-free and the revenue is maximized. Clearly, this solution is also maximal. $\qquad\square$

The following lemmas show some interesting properties of maximal solutions that will be useful in the proof of Theorem 16.

Lemma 14. *Given a maximal solution $(\bar{N}, \bar{o})$, where $\bar{o} = (\bar{\mathbf{X}}, \bar{\mathbf{p}})$, for the buyer preselection problem* $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$ *with $\Gamma = (N, M, (v_i)_{i \in N})$, there exists a buyer $i \in W(\bar{\mathbf{X}})$ such that $u_i(\bar{X}_i, \bar{\mathbf{p}}) = 0$.*

Proof. Assume, for the sake of contradiction, that for any $i \in W(\bar{\mathbf{X}})$ it holds that $u_i(\bar{X}_i, \bar{\mathbf{p}}) > 0$. Let $t \in W(\bar{\mathbf{X}})$ be the winner with the minimum (greater than zero) utility. Then, for any $i \in W(\bar{\mathbf{X}})$, let us consider any item $j_i \in \bar{X}_i$ (i.e., j_i is any item sold to buyer i). By increasing the prices of all j_i, for any $i \in W(\bar{\mathbf{X}})$, by the quantity $u_t(\bar{X}_t, \bar{\mathbf{p}})$, we obtain another solution (with the same allocation vector $\bar{\mathbf{X}}$) for $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$. In particular, we notice that the new solution is still agent envy-free since, by raising the price of each bundle by the same quantity, we do not create new envy. Moreover, each buyer has still greater or equal than zero utility. This is a contradiction to the fact that solution $(\bar{N}, \bar{o})$ is maximal. $\qquad\square$

Lemma 15. *Given a maximal solution $(\bar{N}, \bar{o})$, where $\bar{o} = (\bar{\mathbf{X}}, \bar{\mathbf{p}})$, for the buyer preselection problem* $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$, *with $\Gamma = (N, M, (v_i)_{i \in N})$, for any buyer $i \in W(\bar{\mathbf{X}})$ such that $u_i(\bar{X}_i, \bar{\mathbf{p}}) > 0$, there exists $t_i \in W(\bar{\mathbf{X}})$, where $i \neq t_i$, such that $u_i(\bar{X}_i, \bar{\mathbf{p}}) = u_i(\bar{X}_{t_i}, \bar{\mathbf{p}})$.*

Proof. Clearly, since $(\bar{N}, \bar{o})$ is a solution for the $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$, then for any $i, i' \in W(\bar{\mathbf{X}})$ it holds that $u_i(\bar{X}_i, \bar{\mathbf{p}}) \geq u_i(\bar{X}_{i'}, \bar{\mathbf{p}})$. That is, among all the sold bundles, $\bar{X}_i$ is the one that gives maximum utility to buyer i.

Assume, for the sake of contradiction, that there exists $i \in W(\bar{\mathbf{X}})$ such that for any $i' \in W(\bar{\mathbf{X}}) \setminus \{i\}$ it holds that $u_i(\bar{X}_i, \bar{\mathbf{p}}) > u_i(\bar{X}_{i'}, \bar{\mathbf{p}})$. Let $l \in W(\bar{\mathbf{X}}) \setminus \{i\}$ such that $u_i(\bar{X}_l, \bar{\mathbf{p}}) \geq u_i(\bar{X}_v, \bar{\mathbf{p}})$, for any $v \in W(\bar{\mathbf{X}}) \setminus \{i\}$, i.e., $\bar{X}_l$ is the bundle giving maximum utility to buyer i if we exclude bundle $\bar{X}_i$. By increasing the price of an item $j \in \bar{X}_i$ by the quantity $u_i(\bar{X}_i, \bar{\mathbf{p}}) - \max\{u_i(\bar{X}_l, \bar{\mathbf{p}}); 0\}$ we obtain another solution (with the same allocation vector $\bar{\mathbf{X}}$ and higher revenue) for $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{REV})$. In fact, since we started from an agent envy-free outcome, from one hand we have that buyer i still gets the bundle for which she has the maximum utility (that is still greater or equal than zero). On the other hand, for each other buyer but i, it still holds that she does not envy buyer i (as well as none of the other buyers) since we have increased the price of the bundle assigned to i. This is a contradiction to the fact that solution $(\bar{N}, \bar{o})$ is maximal. $\qquad\square$

We are now ready to prove the following theorem.

Theorem 16. *Let $\Gamma = (N, M, (v_i)_{i \in N})$ be a market with combinatorial valuations and* $\mathsf{PP} = (\Gamma, \mathsf{AEF}, \mathsf{REV})$ *be a pricing problem. Given any solution $(\bar{N}, \bar{o})$ for the buyer preselection problem* $\mathsf{BPP}(\mathsf{PP})$, *it is possible to compute in polynomial time an outcome $o \in \mathsf{AEF}(\Gamma)$ for problem* PP *such that $\mathsf{REV}(o) \geq \mathsf{REV}(\bar{o})$.*

Proof. Without loss of generality, we can assume that solution $(\bar{N}, \bar{o})$ is maximal, because, otherwise, it can be transformed by Claim 13 in a maximal one with at least the same revenue.

Let $\bar{o} = (\bar{\mathbf{X}}, \bar{\mathbf{p}})$. We show how to constructively get in polynomial time the outcome $o \in \mathsf{AEF}(\Gamma)$ starting from $(\bar{N}, \bar{o})$. For any buyer $i \in N \setminus \bar{N}$, let U_i be the set of bundles belonging to $\bar{\mathbf{X}}$ giving maximum and positive (i.e., greater than zero) utility to i. Moreover, for any $i \in \bar{N}$, let $E_i \subseteq \bar{\mathbf{X}} \setminus \{\bar{X}_i\}$ be the set of bundles such that, for any $\bar{X}_j \in E_i$, $u_i(\bar{X}_i, \bar{\mathbf{p}}) = u_i(\bar{X}_j, \bar{\mathbf{p}})$, (i.e., E_i is the set of sold bundles giving to buyer i the same utility as bundle $\bar{X}_i$). By Lemma 15 we know that $E_i \neq \emptyset$ whenever $u_i(\bar{X}_i, \bar{\mathbf{p}}) > 0$.

If, for every buyer $i \in N \setminus \bar{N}$, it holds that $U_i = \emptyset$, it implies that $\bar{o} \in \mathsf{AEF}(\Gamma)$ and the theorem holds. Otherwise, let us define the following directed graph $G = (N, A)$. We first notice that given an allocation vector $\mathbf{X} = (X_1, X_2, \ldots)$, we have that the bundle X_i, which is the set of items sold to buyer i, univocally identifies the buyer i. That is, in the following we will use X_i also for identifying buyer i. For instance, E_i identifies both the set of bundles and the set of buyers getting bundles in E_i. The set of nodes of the graph G is the set of buyers N. Moreover, for any $i \in N \setminus \bar{N}$ and for any $j \in U_i$, there is a directed edge $(i, j) \in A$. Finally, for any $i \in \bar{N}$ and for any $j \in E_i$, there is a directed edge $(i, j) \in A$.

Let us consider any buyer $i \in N \setminus \bar{N}$ such that $U_i \neq \emptyset$. Then we argue that, in the graph G there is a directed path $a = \langle i, i', i'', \ldots, z, j \rangle$ from node i to another $j \in \bar{N}$ such that $u_j(\bar{X}_j, \bar{\mathbf{p}}) = 0$. Notice that such path can be efficiently computed by performing a breadth first search starting from node i. Suppose, for the sake of contradiction, that such path does not exist. It implies that the sub-tree rooted in the node i contains all winners, excluding i, having utility strictly greater than zero. Let us denote with $T_i \subseteq W(\bar{\mathbf{X}})$ the set of winners belonging to such sub-tree rooted in the node i. Notice that for any $j \in T_i$, it holds that $E_j \subseteq T_i$. By exploiting arguments similar to the ones exploited in the proofs of lemmas 14 and 15, it is possible to show that $(\bar{N}, \bar{o})$ is not a maximal solution: a contradiction. Indeed, for any $j \in T_i$, let $l_j \in W(\bar{\mathbf{X}}) \setminus E_j$ such that $u_j(\bar{X}_{l_j}, \bar{\mathbf{p}}) \geq u_j(\bar{X}_i, \bar{\mathbf{p}})$, for any $i \in W(\bar{\mathbf{X}}) \setminus E_j$, that is, $\bar{X}_{l_j}$ is the bundle giving maximum utility to buyer j among all the bundles assigned to buyers in $W(\bar{\mathbf{X}}) \setminus E_j$. Notice that $u_j(\bar{X}_j, \bar{\mathbf{p}}) > u_j(\bar{X}_{l_j}, \bar{\mathbf{p}})$. Let $\delta = \min_{j \in T_i}(u_j(\bar{X}_j, \bar{\mathbf{p}}) - u_j(\bar{X}_{l_j}, \bar{\mathbf{p}}))$. For any $j \in T_i$, let us consider any item $i_j \in \bar{X}_j$ (i.e., i_j is an item sold to buyer j). By increasing by δ the prices of all i_j, for any $j \in T_i$, we obtain another solution $(\bar{N}, \bar{o}')$ for the buyer preselection problem BPP(PP) with the same allocation vector $\bar{\mathbf{X}}$ where we have increased the price of at least one sold item while keeping the prices of the other ones unchanged. This is in contradiction with the fact that the solution $(\bar{N}, \bar{o})$ is maximal.

Therefore, given the path a, we can obtain a new allocation vector where buyer i becomes a winner and buyer j gets no item. Recall that $u_j(\bar{X}_j, \bar{\mathbf{p}}) = 0$ and therefore node j would not envy any bundle. In particular, node (i.e., buyer) i gets the bundle $\bar{X}_{i'}$, node i' gets the bundle $\bar{X}_{i''}$ and so on, and finally node j gets no item. That is, the new allocation vector $\bar{\mathbf{X}}'$ is as follows: any buyer not in the path a gets the same allocation as in $\bar{\mathbf{X}}$, and moreover $\bar{X}'_i = \bar{X}_{i'}, \bar{X}'_{i'} = \bar{X}_{i''}, \ldots, \bar{X}'_z = \bar{X}_j, \bar{X}'_j = \emptyset$.

In this way we get a new solution $(\bar{N}' = \bar{N} \cup \{i\}, \bar{o}')$, where $\bar{o}' = (\bar{\mathbf{X}}', \bar{\mathbf{p}})$, for the buyer preselection problem BPP(Γ, AEF, REV) with the same revenue as $(\bar{N}, \bar{o})$, and where node i is not excluded anymore. However, notice that the obtained solution could be not maximal anymore. Indeed, we need to make it maximal and such that buyer j would be still not envious. We can get it by using a linear program similar to the one defined in the proof of Claim 13, with some further constraints. In particular, we want that the total price of any bundle is at least as it is in solution $(\bar{N}, \bar{o})$. Notice that in the new allocation vector the set of sold bundles is the same, because we are just changing the allocation of bundles to buyers. Given the allocation $\bar{\mathbf{X}}'$, let us consider the following linear program.

$$
\begin{aligned}
\max \quad & \sum_{i \in \bar{N}'} \sum_{j \in \bar{X}'_i} \bar{p}'_j \\
\text{subject to} \quad & v_i(\bar{X}'_i) - \sum_{j \in \bar{X}'_i} \bar{p}'_j \geq v_i(\bar{X}'_z) - \sum_{j \in \bar{X}'_z} \bar{p}'_j, \quad i \in \bar{N}', z \in \bar{N}' \\
& \sum_{j \in \bar{X}'_i} \bar{p}'_j \geq \sum_{j \in \bar{X}'_i} \bar{p}_j, \quad i \in \bar{N}' \\
& \bar{p}'_j \geq 0 \quad j = 1, \ldots, m
\end{aligned}
$$

Notice that the above defined linear program admits at least a feasible solution, being $(\bar{N}', \bar{o}')$, where $\bar{o}' = (\bar{\mathbf{X}}', \bar{\mathbf{p}})$.

It is easy to see that, by solving the linear program, we get a maximal solution $(\bar{N}', \bar{o}'')$, where $\bar{o}'' = (\bar{\mathbf{X}}', \bar{\mathbf{p}}')$, for the buyer preselection problem BPP(Γ, AEF, REV). Indeed, remind that a solution for the buyer preselection problem is such that AEF holds among preselected buyers. If the solution to this linear programming is not maximal it means that we can get a new outcome for the buyer preselection problem with the same set of preselected buyers $\bar{N}'$ with the same allocation where we can increase the price of at least one sold item while keeping the prices of the remaining items unchanged. However, this new outcome satisfies all the constraints of the linear programming and this is a contradiction with the fact that in the linear programming we maximize the sum of the prices of sold items. Moreover, notice that in the obtained maximal solution, buyer j has still maximum utility at most zero, that is, buyer j is not envious and we do not need to exclude her from the market.

At this point, for each buyer i not in the market, we compute again (as before) the set U_i, and for any buyer i in the market the set E_i. We then repeat the same procedure described above. Notice that the sets U_i and E_i, and the edges in the graph G can be in general different at any step.

Finally, since at each step we include a new buyer in the market (without excluding any other buyer) and the prices of the bundles never decrease, we are able to get in polynomial time outcome $o \in \mathsf{AEF}(\Gamma)$ such that $\mathsf{REV}(o) \geq \mathsf{REV}(\bar{o})$, as desired. □

Item set Buyer	{1}	{2}	{3}	{4}	{5}	{6}	{7}	{8}
1	2	3	4	4	4	1	3	5
2	3	2	3	3	4	1	4	4
3	4	4	1	3	3	3	5	5
4	2	3	3	2	4	1	3	3
5	3	2	0	3	5	0	7	4

Table 5: The buyers' valuations of Example 3.

We now present an illustrating example showing how Theorem 16 works for a specific instance of the problem. In particular, we consider a market with additive valuations to make things simpler.

Example 3. Let $\Gamma = (N, M, (v_i)_{i \in N})$ be a market with additive valuations with $n = 5$ buyers and $m = 8$ items, and $\mathsf{PP} = (\Gamma, \mathsf{AEF}, \mathsf{REV})$ a pricing problem. The buyers' valuations for each item are reported in Table 5. Let us consider the solution $(\bar{N}, \bar{o})$ for the buyer preselection problem $\mathsf{BPP(PP)}$ where $\bar{N} = \{1, 2, 3\}$ and $\bar{o} = (\bar{\mathbf{X}}, \bar{\mathbf{p}}) = ((\bar{X}_1, \bar{X}_2, \bar{X}_3), (\bar{p}_1, \bar{p}_2, \bar{p}_3, \bar{p}_4, \bar{p}_5, \bar{p}_6, \bar{p}_7, \bar{p}_8)) = ((\{1, 2, 3\}, \{4, 5, 6\}, \{7, 8\}), (2, 3, 3, 3, 4, 1, 5, 4))$. Notice that $u_1(\bar{X}_1, \bar{\mathbf{p}}) = 1$, $u_2(\bar{X}_2, \bar{\mathbf{p}}) = 0$, and $u_3(\bar{X}_3, \bar{\mathbf{p}}) = 1$. It is easy to check that $(\bar{N}, \bar{o})$ is a maximal solution (see Definition 12) for the buyer preselection problem $\mathsf{BPP(PP)}$. In fact, if we increase the price of any item assigned to buyer 2 (i.e., $\bar{p}_4$, $\bar{p}_5$ and $\bar{p}_6$) her utility becomes negative. Moreover, if we increase the price of any item assigned to buyer 1 (i.e., $\bar{p}_1$, $\bar{p}_2$ and $\bar{p}_3$) her utility becomes smaller than one and the resulting outcome would not be agent envy-free since $u_1(\bar{X}_2, \bar{\mathbf{p}}) = 1$. Finally, if we increase the price of any item assigned to buyer 3 (i.e., $\bar{p}_7$ and $\bar{p}_8$) her utility becomes smaller than one and the resulting outcome would not be agent envy-free since $u_3(\bar{X}_2, \bar{\mathbf{p}}) = 1$.

We now show how, by exploiting Theorem 16, starting from the solution $(\bar{N}, \bar{o})$, we can compute an outcome $o \in \mathsf{AEF}(\Gamma)$ for problem PP such that $\mathsf{REV}(o) \geq \mathsf{REV}(\bar{o})$.

In the first iteration of the theorem we have that $U_4 = \emptyset$ since $u_4(\bar{X}_1, \bar{\mathbf{p}}) = 0$, $u_4(\bar{X}_2, \bar{\mathbf{p}}) = -1$ and $u_4(\bar{X}_3, \bar{\mathbf{p}}) = -2$. Moreover, $U_5 = \{\bar{X}_3\}$ since $u_5(\bar{X}_1, \bar{\mathbf{p}}) = -3$, $u_5(\bar{X}_2, \bar{\mathbf{p}}) = 0$ and $u_5(\bar{X}_3, \bar{\mathbf{p}}) = 2$. Finally, it is easy to check that $E_1 = \{\bar{X}_2\}$, $E_2 = \emptyset$, and $E_3 = \{\bar{X}_1, \bar{X}_2\}$. Therefore, we get the directed graph $G = (N, E)$ where $E = \{(1, 2), (3, 1), (3, 2), (5, 3)\}$. A graphical representation of graph G is reported in Figure 1.
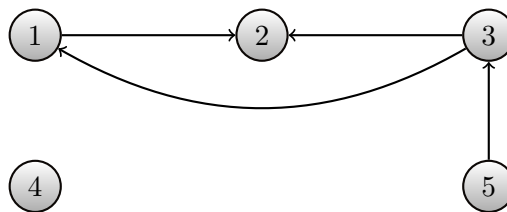

Figure 1: The graph G of Example 3.

Let us consider the directed path $a = \langle 5, 3, 2 \rangle$ in G (notice that we could also consider the path $a = \langle 5, 3, 1, 2 \rangle$). We get the new solution $(\bar{N}', \bar{o}')$ for the buyer preselection problem BPP(PP) where $\bar{N}' = \{1, 2, 3, 5\}$ and $\bar{o}' = (\bar{\mathbf{X}}', \bar{\mathbf{p}}') = ((\bar{X_1}', \bar{X_2}', \bar{X_3}', \bar{X_5}'), (\bar{p}_1', \bar{p}_2', \bar{p}_3', \bar{p}_4', \bar{p}_5', \bar{p}_6', \bar{p}_7', \bar{p}_8')) = ((\{1, 2, 3\}, \emptyset, \{4, 5, 6\}, \{7, 8\}), (2, 3, 3, 3, 4, 1, 5, 4))$. Notice that in the new solution buyer 5 is a winner and buyer 2 gets no item, and that we do not get a smaller revenue with respect to $(\bar{N}, \bar{o})$. We further notice that $(\bar{N}', \bar{o}')$ is not maximal. Therefore, we apply the linear program defined in the proof of Theorem 16 to $(\bar{N}', \bar{o}')$ and get, for instance, the following maximal solution $(\bar{N}'', \bar{o}'')$ where $\bar{N}'' = \bar{N}'$ and $\bar{o}'' = (\bar{\mathbf{X}}'', \bar{\mathbf{p}}'') = ((\bar{X_1}'', \bar{X_2}'', \bar{X_3}'', \bar{X_5}''), (\bar{p}_1'', \bar{p}_2'', \bar{p}_3'', \bar{p}_4'', \bar{p}_5'', \bar{p}_6'', \bar{p}_7'', \bar{p}_8'')) = ((\{1, 2, 3\}, \emptyset, \{4, 5, 6\}, \{7, 8\}), (2, 3, 4, 3, 4, 2, 7, 4))$. We notice that by solving the linear program we do not get a different allocation of items to buyers with respect to $(\bar{N}', \bar{o}')$, and that item prices do not decrease.

In the second iteration considered in the proof of the theorem we still have that $U_4 = \emptyset$.

We conclude that (N, o) where $o = (\mathbf{X}, \mathbf{p}) = ((X_1, X_2, X_3, X_4, X_5), (p_1, p_2, p_3, p_4, p_5, p_6, p_7, p_8)) = ((\{1, 2, 3\}, \emptyset, \{4, 5, 6\}, \emptyset, \{7, 8\}), (2, 3, 4, 3, 4, 2, 7, 4))$ is such that $o \in \mathsf{AEF}(\Gamma)$ and $\mathsf{REV}(o) \geq \mathsf{REV}(\bar{o})$ (in particular, in this example we have that $\mathsf{REV}(o) > \mathsf{REV}(\bar{o})$).

We now consider the buyer preselection problem $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{SW})$. Analogously to Theorem 16 holding for the revenue maximization case, next theorem shows that an agent envy-free outcome for the buyer preselection problem can be efficiently transformed in an agent envy-free outcome, for the corresponding pricing problem (without preselection), having at least the same social welfare.

Theorem 17. *Let $\Gamma = (N, M, (v_i)_{i \in N})$ be a market with combinatorial valuations and $\mathsf{PP} = (\Gamma, \mathsf{AEF}, \mathsf{SW})$ a pricing problem. Given any solution $(\bar{N}, \bar{o})$ for the buyer preselection problem $\mathsf{BPP}(\mathsf{PP})$, it is possible to compute in polynomial time an outcome $o \in \mathsf{AEF}(\Gamma)$ for problem PP such that $\mathsf{SW}(o) \geq \mathsf{SW}(\bar{o})$.*

Proof. We first give an intuitive explanation of what we do in this proof. Given any solution $(\bar{N}, \bar{o})$ for the buyer preselection problem $\mathsf{BPP}(\mathsf{PP})$ where the market is $\Gamma = (N, M, (v_i)_{i \in N})$, we create a new market $\Gamma' = (N, M', (v_i')_{i \in N})$ with unit-demand buyers and where the number of items is $|\bar{N}|$. More precisely, there is an item in M' for each sold bundle in $(\bar{N}, \bar{o})$. Moreover, for each buyer $i \in N$, the valuation $v_i'(\{j\})$ that i has for item j in market Γ' is equal to the valuation $v_i(\bar{X}_j)$ that i has for the corresponding sold bundle in $(\bar{N}, \bar{o})$ for market Γ (we assume, without loss of generality, that each buyer in $\bar{N}$ gets a non-empty bundle in the outcome $\bar{o}$). We then consider a Walrasian equilibrium for market Γ' which, we recall, is a market envy-free outcome with the additional requirement that the market clears and which has the important property of maximizing the social welfare. We finally obtain, by using the Walrasian equilibrium, an agent envy-free outcome for the pricing problem (without preselection) for market Γ with social welfare greater or equal than $\mathsf{SW}(\bar{o})$. We now proceed with the formal proof.

Without loss of generality, let us assume that the agents orderings in $\bar{N}$ and N are such that the i-th agent of $\bar{N}$ is also the i-th agent of N (for $i = 1, \ldots, |\bar{N}|$).

Let $\bar{o} = (\bar{\mathbf{X}}, \bar{\mathbf{p}})$ be the considered outcome for problem $\mathsf{BPP}(\Gamma, \mathsf{AEF}, \mathsf{SW})$.

Consider a new market $\Gamma' = (N, M', (v_i')_{i \in N})$ with unit-demand buyers, where $M' = \{1, \ldots, m'\}$ with $m' = |\bar{N}|$ and $(v_i')_{i \in N}$ are such that $v_i'(\{j\}) = v_i(\bar{X}_j)$. Consider the allocation vector $\mathbf{X}' = (X_1', \ldots, X_n')$ for market Γ', in which $X_i' = \{i\}$ for any $i = 1, \ldots, m'$, and $X_i' = \emptyset$ for any $i = m' + 1, \ldots, n$. It is easy to check that $\mathsf{SW}(\mathbf{X}', \mathbf{p}) = \mathsf{SW}(\bar{o})$ for any price vector $\mathbf{p}$.

Recall that, given any market Γ, $\mathsf{WE}(\Gamma)$ denotes the set of Walrasian Equilibria for Γ. By the well known *First Welfare Theorem* (Nisan, Roughgarden, Tardos, & Vazirani, 2007, Theorem 11.13), whenever they exist, Walrasian equilibria maximize social welfare over all possible outcomes. It is also known (Kelso & Crawford, 1982; Leme & Wong, 2017) that a market with gross substitutes valuations always admits a Walrasian Equilibrium that can be computed in polynomial time and with a polynomial number of calls to the value oracle. Moreover, a market with unit-demand buyers (as Γ') is a special case of a market with gross substitutes valuations in which the value oracle trivially requires polynomial time to be implemented, and

therefore always admits a Walrasian Equilibrium that can be computed in polynomial time. Thus, we can compute in polynomial time an outcome $o'' = (\mathbf{X}'', \mathbf{p}'')$ for market Γ' such that (i) $\mathsf{SW}(o'') \geq \mathsf{SW}(\mathbf{X}', \mathbf{p}) = \mathsf{SW}(\bar{o})$ (for any price vector $\mathbf{p}$) and (ii) $o'' \in \mathsf{WE}(\Gamma')$.

Let $Y = \bigcup_{j \in M(\mathbf{X}'')} \bar{X}_j$ be the set of all items of market Γ corresponding to items of market Γ' sold in allocation $\mathbf{X}''$, and, for any $i = 1, \ldots, n$ such that $X_i'' \neq \emptyset$, let k_i be the unique element of X_i''. Now, consider the outcome $o = (\mathbf{X}, \mathbf{p})$ for market Γ, where $\mathbf{X} = (X_1, \ldots, X_n)$ is such that, for any $i = 1, \ldots, n$,

- if $X_i'' = \emptyset$ then set $X_i = \emptyset$ as well;

- otherwise, set $X_i = \bar{X}_{k_i}$ and, for any $j \in \bar{X}_{k_i}$, set $p_j = \dfrac{p_{k_i}''}{|\bar{X}_{k_i}|}$.

Finally, for any $j \notin Y$, set $p_j = 0$.

Clearly, by construction, for any $i, j = 1, \ldots, n$ such that $X_j \neq \emptyset$, it holds that $v_i(X_j) = v_i'(\{k_j\})$ and $u_i(X_j, \mathbf{p}) = u_i'(\{k_j\}, \mathbf{p}'')$. Therefore, $\mathsf{SW}(o) = \mathsf{SW}(o'')$ (hence we obtain $\mathsf{SW}(o) \geq \mathsf{SW}(\bar{o})$) and, given that $o'' \in \mathsf{WE}(\Gamma')$, it follows that $o \in \mathsf{AEF}(\Gamma)$. $\qquad\square$

As a consequence of Theorems 16 and 17, we obtain the following corollary.

Corollary 18. *Let $\Gamma = (N, M, (v_i)_{i \in N})$ be a market with combinatorial valuations. For* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$*, given an α-approximate solution (N, o) (with $\alpha \geq 1$) for the buyer preselection problem $\mathsf{BPP}(\mathsf{PP})$ with $\mathsf{PP} = (\Gamma, \mathsf{AEF}, \mathsf{obj})$, it is possible to compute in polynomial time an outcome $o' \in \mathsf{AEF}(\Gamma)$ approximating the optimal solution of PP by a factor equal to α.*

On the one hand, Corollary 18 tells us that, for $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, any inapproximability result holding for a pricing problem $\mathsf{PP} = (\Gamma, \mathsf{AEF}, \mathsf{obj})$ directly extends to the buyer preselection problem $\mathsf{BPP}(\mathsf{PP})$; on the other hand, it tells us that any α-approximation algorithm for $\mathsf{BPP}(\mathsf{PP})$ is also an α-approximation algorithm for PP. Thus, as discussed in Section 1.1, 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 $v_i : \{0, 1, \ldots, m\} \to \mathbb{R}_{\geq 0}$. Notice that, in this case, the market requires $\Omega(nm)$ bits to be represented, because every buyer has a valuation for every possible number of items. 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 price to p (i.e., all the items have the same price) and the total price for selling x items is px.

A particular situation of multi-unit market arises when one assumes *single-minded* buyers; in this case, for every agent $i = 1, \ldots, n$, the valuation function v_i can be completely defined by specifying how much every agent valuates the set containing k_i items (being the only set she is interested in). Notice that, in this particular case, the market requires $\Omega(n \log m)$ bits to be represented, because, for every buyer i, it is necessary to represent integer k_i.

We first focus on the case of *market envy-free* solutions. The following claim says that preselection cannot improve the revenue and social welfare of envy-free solutions by a factor greater than m.

Claim 19. *Given a multi-unit market Γ with identical items and a pricing problem $\mathsf{PP} = (\Gamma, \mathsf{MEF}, \mathsf{obj})$ with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, it holds that $\mathsf{obj}(o^*(\mathsf{BPP}(\mathsf{PP}))) \leq m \cdot \mathsf{obj}(o^*(\mathsf{PP}))$.*

Proof. We now show that a market envy-free solution (without preselection) which is a $1/m$-approximation with respect to both the optimal social welfare and the optimal revenue (with preselection) can be obtained by selling $\bar{j}$ items to the buyer $\bar{i}$ such that $(\bar{i}, \bar{j}) \in \text{argmax}_{i=1,\dots,n;j=1,\dots,m} \frac{v_i(j)}{j}$ at price $\frac{v_{\bar{i}}(\bar{j})}{\bar{j}}$, i.e., we are considering outcome (X, p) such that $x_{\bar{i}} = \bar{j}$ and $x_i = 0$ for all $i \neq \bar{i}$, and $p = \frac{v_{\bar{i}}(\bar{j})}{\bar{j}}$.

First of all, notice that the utility of all buyers in outcome (X, p) is 0: buyer $\bar{i}$ has utility $v_{\bar{i}}(\bar{j}) - \bar{j}p = v_{\bar{i}}(\bar{j}) - \bar{j}\frac{v_{\bar{i}}(\bar{j})}{\bar{j}} = 0$ and all other buyers clearly have utility equal to 0 as they receive no item. The market envy-freeness follows because, for every buyer $i = 1, \dots, n$ (including buyer $\bar{i}$) and any $j = 1, \dots, m$, it holds that $v_i(j) - jp = v_i(j) - j\frac{v_{\bar{i}}(\bar{j})}{\bar{j}} \leq 0$ (by the choice of $\bar{i}$ and $\bar{j}$).

Finally, the approximation ratio follows by the fact that (X, p) is selling at least an item at price $p = \frac{v_{\bar{i}}(\bar{j})}{\bar{j}}$ and in no solution (even with preselection, selling at most all m items) there can exist an item sold at a price $p' > p$, because the corresponding buyer i buying j items would have utility $v_i(j) - jp' < 0$ for any $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$ (again by the choice of $\bar{i}$ and $\bar{j}$). $\qquad\square$

We notice that the above bound is tight: next claim shows that preselection can improve both the revenue and the social welfare of a multi-unit market with identical items of a factor at least $m - \epsilon$ (for any $\epsilon > 0$), in the case of single-minded buyers.

Claim 20. *For any $\epsilon > 0$, there exists a multi-unit market Γ with single-minded buyers and identical items, and a pricing problem $\mathsf{PP} = (\Gamma, \mathsf{MEF}, \mathsf{obj})$ with $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, such that $\mathsf{obj}(o^*(\mathsf{BPP}(\mathsf{PP}))) \geq (m - \epsilon)\mathsf{obj}(o^*(\mathsf{PP}))$.*

This claim follows from the following example.

Example 4. Consider a market with only two single-minded buyers. Buyer 1 has a valuation of $1 + \epsilon'$, for an $\epsilon' > 0$, for receiving any (one) item. Buyer 2 has a valuation of m for receiving all the m items.

Notice that, in any market envy-free outcome for the setting without preselection, buyer 2 receives no item. In fact, if buyer 2 receives m items, the item price p must be at most 1. Thus, buyer 1 would be envious since she would get no item (i.e., there are not 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 equal to $1 + \epsilon'$. However, if we exclude buyer 1 and consider the submarket with only buyer 2, we can sell m items at item price 1, thus obtaining a revenue and social welfare of m. We get that, for any $0 < \epsilon < m$, $\frac{m}{1+\epsilon'} \geq m - \epsilon$ when $0 < \epsilon' \leq \frac{\epsilon}{m-\epsilon}$.

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

Both Theorem 21, holding for the multi-unit case, and Theorem 22, holding for the special case of single-minded buyers, exploit reductions to the 0-1 Multiple-Choice Knapsack Problem, that is defined as follows.

The *0-1 Multiple-Choice Knapsack Problem (0-1 MCKP)* is a generalization of the classical Knapsack problem introduced by Lawler (1979). In this problem, we are given σ elements, partitioned into α classes $C_1, C_2, \dots, C_\alpha$, to be packed in some knapsack of capacity c. Notice that $\sigma = \sum_{i=1}^{\alpha} |C_i|$. For every $i = 1, \dots, \alpha$, each element $e \in C_i$ 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. It is known that this problem can be optimally solved in pseudo-polynomial time, i.e., in time $O(\sigma c)$ (Martello & Toth, 1990, Section 2.12).

Theorem 21. *Given a pricing problem $\mathsf{PP} = (\Gamma, \mathsf{MEF}, \mathsf{obj})$, where $\Gamma = (N, M, (v_i)_{i \in N})$ is a multi-unit market with identical items and $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, the buyer preselection problem $\mathsf{BPP}(\mathsf{PP})$ can be*

optimally solved in time $O(nm^2)$ (i.e., in polynomial time in the size of the instance, given that the instance requires $\Omega(nm)$ space for its representation).

Proof. Recall that solving the buyer preselection problem BPP(PP) corresponds to finding a subset N^* of agents to admit to the market and an optimal outcome $o^* = (X^*, p^*)$, where, since we are considering a multi-unit market, 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 BPP(PP).

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

$$\mathcal{P} = \{y \,|\, y > 0 \land \exists_{i \in N} \exists_{k,k' \in 0,\ldots,m, k \neq k'} \; v_i(k) - ky = v_i(k') - k'y\}.$$

$\mathcal{P}$ can be computed in $O(nm^2)$ time, i.e., in polynomial time in the size of the instance: in fact, for all buyers $i \in N$ and for all pairs (k, k') of integers belonging to $\{0, 1, \ldots, m\}$ such that $k \neq k'$, $\mathcal{P}$ contains the solution y of equality $v_i(k) - yk = v_i(k') - yk'$.

Coming back to Example 1, Table 6 shows that, in this case, $\mathcal{P} = \left\{\frac{3}{2}, 2, 3\right\}$.

i	k	k'	equation	solution
1	0	1	$0 = v_1(1) - y$	$y = v_1(1) = 0$
1	0	2	$0 = v_1(2) - 2y$	$y = \frac{v_1(2)}{2} = \frac{3}{2}$
1	1	2	$v_1(1) - y = v_1(2) - 2y$	$y = v_1(2) - v_1(1) = 3$
2	0	1	$0 = v_2(1) - y$	$y = v_2(1) = 2$
2	0	2	$0 = v_2(2) - 2y$	$y = \frac{v_2(2)}{2} = 0$
2	1	2	$v_2(1) - y = v_2(2) - 2y$	$y = v_2(2) - v_2(1) = -2$

Table 6: The values in set $\mathcal{P}$ relative to the market of Example 1.

Now, let (N^*, o^*), with $o^* = (X^*, p^*)$, be an optimal solution for BPP(PP) with the highest possible item price p^* and assume by contradiction that $p^* \notin \mathcal{P}$; 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 satisfy, for any $i = 1, \ldots, n$, $v_i(\bar{x}_i) - p^*\bar{x}_i \geq 0$ and y satisfying equality $v_i(\bar{x}_i) - y\bar{x}_i = 0$ belongs to $\mathcal{P}$. Since p^* induces a market envy-free solution, it has to satisfy, 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 the above constraints are not satisfied in a strict manner, i.e., it holds 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 satisfy all envy-free constraints (with some inequalities 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 agents to admit to the market and an allocation X^*. In fact, the optimal solution of BPP(PP) is given by the best solution among the ones obtained for all the candidate prices belonging to $\mathcal{P}$.

We now provide a polynomial reduction from our problem to 0-1 MCKP. Given an instance I of BPP(PP) 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 positive utility: let $\bar{U}_i = (\arg\max_{k=0,\ldots,m} (v_i(k) - kp^*)) \setminus \{0\}$ be the set containing these values. Notice that if a buyer i has negative utility for all possible bundles of $k = 1, \ldots, m$ items, then $\arg\max_{k=0,\ldots,m} (v_i(k) - kp^*) = \{0\}$ and $\bar{U}_i = \emptyset$. It is easy to check that, in

every market envy-free solution, allocation X must satisfy $x_i \in \bar{U}_i$ for every agent $i = 1, \ldots, n$ receiving at least an item. For every $i = 1, \ldots, n$, consider set $\bar{U}_i$: for every $k \in \bar{U}_i$, we add to class C_i an element e such that $\gamma_e = k$ and

- $\beta_e = kp^*$ if obj $=$ REV;

- $\beta_e = v_i(k)$ if obj $=$ SW.

Notice that the total number of elements σ is at most nm. If we consider again the market of Example 1, we have $\alpha = 2$ classes and capacity $c = 2$. Moreover, if we set $p^* = \frac{3}{2}$ and consider obj $=$ SW, $\bar{U}_1 = \{2\}$ and $\bar{U}_2 = \{1\}$: we therefore add to class C_1 an element e_1 such that $\gamma_{e_1} = 2$ and $\beta_{e_1} = v_1(2) = 3$, and to class C_2 an element e_2 such that $\gamma_{e_2} = 1$ and $\beta_{e_2} = v_2(1) = 2$.

Given a solution E 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 agents to admit to the market and an allocation $\bar{X}$, as follows: $\bar{N}$ contains the agents 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 $i \in \bar{N}$, $\bar{x}_i = \gamma_e$ if there exists an element $e \in C_i \cap E$. Let $\bar{o} = (\bar{X}, p^*)$; clearly, obj$(\bar{o}) = \beta$. Furthermore, by the way $\bar{U}_i$ is defined, outcome $\bar{o}$ is market envy-free. Coming back to the example considered above, if we consider solution E containing only element e_1 (of total profit 3), then we admit to the market only buyer 1 and allocate 2 items to her, thus obtaining a market envy-free outcome of social welfare 3.

Conversely, given a subset $\bar{N}$ of preselected agents and an outcome $o = (X, p^*)$ for I, it is possible to obtain a solution for I' with the profit being equal to obj(o) 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 market envy-free). Clearly, the total profit $\beta = \sum_{e \in E} \beta_e = $ obj(o).

The claim follows by noticing that, since $\sigma \leq nm$ and $c = m$, the pseudo-polynomial algorithm by Martello and Toth (1990) is in fact polynomial with respect to the size of the instance of problem BPP(PP). $\square$

For the special case of single-minded buyers, by exploiting the same ideas used for proving Theorem 21, the following theorem provides an **FPTAS**, i.e., a fully polynomial time approximation scheme. More precisely, an **FPTAS** for a problem is an approximation scheme providing, for any $\epsilon > 0$, a $(1 + \epsilon)$-approximation algorithm whose time complexity is polynomial both in the input size and in $1/\epsilon$; we call ϵ the **FPTAS** parameter.

Theorem 22. *Given a pricing problem* PP $= (\Gamma, \text{MEF}, \text{obj})$, *where* $\Gamma = (N, M, (v_i)_{i \in N})$ *is a multi-unit market with single-minded buyers and identical items, for* obj $\in \{\text{REV}, \text{SW}\}$, *the buyer preselection problem* BPP(PP) *admits a* fully polynomial approximation scheme *generating algorithms with execution time* $O\left(\frac{n^2}{\epsilon}\right)$, *i.e., polynomial both in the instance size (given that the instance requires* $\Omega(n \log m)$ *space for its representation) and in* $1/\epsilon$, *where* ϵ *is the* **FPTAS** *parameter.*

Proof. Since we are considering a multi-unit market with identical items, the proof is analogous to the one of Theorem 21: We will first show that it is possible to compute in polynomial time a set $\mathcal{P}$ containing the optimal price p^* for problem BPP(PP), and then we will show that, given a fixed price p^*, it is possible to optimally compute in polynomial time a subset N^* of agents to admit to the market and an allocation X^*.

The critical point here is that in this case an instance of the buyer preselection problem requires only $\Omega(n \log m)$ bits.

First of all, we have to check that the computation of the possible optimal prices is still polynomial. In order to compute, analogously to the proof of Theorem 21 (notice that if buyer i is single-minded, she does

not envy bundles with size different from k_i), set $\mathcal{P}$ containing the optimal price p^* for problem BPP(PP), we note that in this case $\mathcal{P}$ can be defined as follows:

$$\mathcal{P} = \{y | y > 0 \land \exists_{i \in N} \ v_i(k_i) - y k_i = 0\}.$$

Therefore, it is possible to compute set $\mathcal{P}$ in time $O(n)$.

Furthermore, since in this case, for every $i = 1, \dots, n$, $|\bar{U}_i| \leq 1$ (where $\bar{U}_i$ is defined as in the proof of Theorem 21), also the reduction from our problem to 0-1 MCKP is still polynomial and $\sigma \leq n$.

Finally, notice that in this case we cannot exploit the pseudo-polynomial algorithm by Martello and Toth (1990) having complexity $O(\sigma c) = O(nm)$, because m cannot be guaranteed to be polynomial in the size of the instance being $\Omega(n \log m)$. The claimed **FPTAS** result follows from the **FPTAS** provided by Lawler (1979) for 0-1 MCKP, running in time $O\left(\alpha \log \alpha + \frac{\alpha \sigma}{\epsilon}\right)$, that in our reduction becomes $O\left(\frac{n^2}{\epsilon}\right)$. □

We complement the result of Theorem 22 by showing the NP-hardness of the problem of computing the maximum revenue for the case of market envy-free in multi-unit markets with single-minded buyers and identical items. The following theorem is proved by exploiting a reduction from the Subset Sum problem.

Theorem 23. *Given a pricing problem* PP $= (\Gamma, \mathsf{MEF}, \mathsf{REV})$*, where* $\Gamma = (N, M, (v_i)_{i \in N})$ *is a multi-unit market with single-minded buyers and identical items, the buyer preselection problem* BPP(PP) *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 $\{a_1, a_2, \dots, a_n\}$ 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 $m = s$.

If there is a solution for the given instance of the Subset Sum problem, then by setting the item price $p = 1$ 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. In fact, every buyer has a value per item equal to 1, and therefore, by individual rationality, no buyer can be assigned items priced more than 1.

On the other hand, if the revenue of the optimal outcome to the buyer preselection problem is equal to m, we can obtain a solution to the Subset Sum problem. Notice that we can obtain a revenue of m only if we sell exactly m items at item price $p = 1$. In fact, at item price $p > 1$, we sell no item, and at item price $p < 1$, we do not get an optimal solution. □

We now focus on the case of *agent envy-free* solutions. We first notice that the results of Section 4 claiming that buyer preselection in agent envy-free solutions does not improve the quality of outcomes (with respect to neither social welfare nor revenue maximization) do not hold for the multi-unit case, because in this case it is not possible to change the price of some items without influencing the other ones; actually, as proved in Claim 25, preselection can improve, with respect to both social welfare and revenue maximization, the quality of outcomes.

We start by showing that, for multi-unit markets with identical items and general valuations, preselection can improve the revenue by a multiplicative factor of at most 2.

Theorem 24. *Given a solution* $(\bar{N}, \bar{o})$ *for the buyer preselection problem* BPP(PP)*, with* PP $=$ $(\Gamma, \mathsf{AEF}, \mathsf{REV})$ *and* $\Gamma = (N, M, (v_i)_{i \in N})$ *being a multi-unit market with identical items, it is possible to compute in polynomial time an outcome* $o \in \mathsf{AEF}(\Gamma)$ *for problem* PP *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 $o \in \text{AEF}(\Gamma)$ with $o = (X, p)$ such that (i) $p \geq \bar{p}$ and (ii) $|M(X)| \geq \frac{|M(\bar{X})|}{2}$: this directly implies that $\text{REV}(o) \geq \frac{\text{REV}(\bar{o})}{2}$, because outcome o sells at least one half of the items sold by outcome $\bar{o}$ at a price at least equal to $\bar{p}$. See Example 5 and the subsequent discussion for a clarifying example showing how this proof can be exploited in order to obtain an outcome $o \in \text{AEF}(\Gamma)$ with $\text{REV}(o) \geq \frac{\text{REV}(\bar{o})}{2}$.

For every $j = 1, \ldots, m$, 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 $j = 1, \ldots, m$, let $B_j(\bar{p}) = \{i | v_i(j) - j\bar{p} \geq 0\}$ be the subset of agents obtaining a non-negative utility for a bundle with j items; then, $a_j = j|B_j(\bar{p})|$.

We divide the proof in two disjoint cases.

- If there exists $j \in \{1, \ldots, m\}$ such that $a_j \geq \frac{|M(\bar{X})|}{2}$, outcome $o = (X, p)$ is such that $x_k = j$ for every $k \in B_j(\bar{p})$, 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, thus selling at least $m - j$ items. 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 $j \in \{1, \ldots, m\}$ it holds that $a_j < \frac{|M(\bar{X})|}{2}$, outcome $o = (X, \bar{p})$ (with the same price as in outcome $\bar{o}$) is computed as follows.

  For any $i = 1, \ldots, n$ and $k = 1, \ldots, m$, let $u_i^{*k} = \max_{t=0}^k \{v_i(t) - t\bar{p}\}$ be the maximum possible utility of buyer i for bundles of at most k items and $b_i^{*k} = \max\{j | 0 \leq j \leq k \wedge v_i(j) - j\bar{p} = u_i^{*k}\}$ the maximum size of a bundle for buyer i that achieves utility u_i^{*k}. Moreover, for any $k = 1, \ldots, m$ and $j = 0, \ldots, k$, 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 maximum size. Notice that, for any $k = 1, \ldots, m$, sets $B_{k,0}, \ldots, B_{k,k}$ partition the set N of buyers.

  Notice that $B_{k,j}$ can be computed in time *poly*(n, m). In fact, for any $i \in \{1, \ldots, n\}$, u_i^{*k} and b_i^{*k} can be computed in time being linear in $k \leq m$ (by enumerating all the k candidates that can realize the requested maximum) and $B_{k,j}$ can be computed in time proportional to nk (by computing b_i^{*k} for $i = 1, \ldots, n$, i.e., for all the n buyers in the market).

  For $k = 1, \ldots, m$, consider allocation $X^k = (x_1^k, \ldots, x_n^k)$ such that, for every $i = 1, \ldots, n$, $x_i^k = j$ if $i \in B_{k,j}$. 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 $j = 2, \ldots, m$, $|M(X^j)| - |M(X^{j-1})| \leq a_j$.

  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 $|M(X^{k'-1})| \geq m - \frac{|M(\bar{X})|}{2} \geq \frac{|M(\bar{X})|}{2}$.

$\square$

We now show that there exist multi-unit markets with identical items in which preselection can improve both the revenue and the social welfare in the case of single-minded buyers, thus closing in a tight way the previous bound.

Claim 25. *For any $\epsilon > 0$, there exists a multi-unit market Γ with single-minded buyers and identical items, and a pricing problem* $\mathsf{PP} = (\Gamma, \mathsf{AEF}(\Gamma), \mathsf{obj})$ *with* $\mathsf{obj} \in \{\mathsf{REV}, \mathsf{SW}\}$, *such that* $\mathsf{obj}(o^*(\mathsf{BPP}(\mathsf{PP}))) \geq (2 - \epsilon)\mathsf{obj}(o^*(\mathsf{PP}))$.

This claim follows from the following example.

Example 5. Consider a market with $x + 1$ single-minded buyers. Buyer 1 has a valuation of x for receiving x items. Buyer i, for any $i = 2, 3, \ldots, x + 1$, has a valuation of $1 + \epsilon'$, for a small $\epsilon' > 0$, for receiving any (one) item. Finally, the number of items is $2x - 1$.

Notice that, in any agent envy-free outcome for the setting without preselection, if buyer 1 receives x items, then the item price p must be at most 1. Moreover, 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 $i = 2, 3, \ldots, x+1$, at price $1+\epsilon'$. On the other hand, if we consider the submarket with only buyers $1, 2, \ldots, x$ (i.e., we exclude one buyer that has a valuation of $1+\epsilon'$ for receiving any (one) item), we can sell $2x - 1$ items at item price 1, that is x items to buyer 1, and one item to buyers i, for any $i = 2, \ldots, x$, thus obtaining a revenue and social welfare of $2x - 1$. We get that, for any $x > \frac{1}{\epsilon}$ and $0 < \epsilon < 2$, $\frac{2x-1}{x(1+\epsilon')} \geq 2 - \epsilon$ when $0 < \epsilon' \leq \frac{x\epsilon-1}{x(2-\epsilon)}$. Notice that $\frac{x\epsilon-1}{x(2-\epsilon)} > 0$.

If we consider the latter solution $(\bar{N}, \bar{o})$ of the buyer preselection problem with $\bar{N} = \{1, \ldots, x\}$ and $\bar{o} = (\bar{X}, \bar{p})$ such that $\bar{p} = 1$ and x items are sold to the first buyer, while other $x - 1$ items to the remaining $x - 1$ buyers (one item for each of them), we can exploit the proof of Theorem 24 as follows. First of all, notice that $\mathsf{REV}(\bar{o}) = |M(\bar{X})| = 2x - 1$ and, since $B_1(1) = \{2, 3, \ldots, x + 1\}$, $B_x(1) = \{1\}$ (and $B_j(1) = \emptyset$ for all $j \notin \{1, x\}$), it holds that $a_1 = a_x = x$. Therefore, we are in the first case considered in the proof, given that $a_1 = a_x \geq \frac{|M(\bar{X})|}{2}$. In fact, an agent envy-free solution o with $\mathsf{REV}(o) = x$ can be obtained by selling, at price $p = \bar{p} = 1$, either only a bundle of x items to buyer 1, or x bundles of one item to buyers $2, 3, \ldots, x + 1$.

It is worth noticing that, on the one hand, preselection can be exploited as an algorithmic framework for designing good approximation algorithms (losing only a multiplicative 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 market envy-free outcomes and the objective function of social welfare extend to the notion of Walrasian Equilibrium. In particular, the inapproximability result of Theorem 6 (as already noticed in its proof) and the n-approximation algorithm for the buyer preselection problem of Theorem 10 (given that the outcome considered in Claim 8 has all prices set to 0) 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 problems which are left open are: for markets with a unique buyer, closing the gap between NP-hardness and the logarithmic approximation for the case of revenue maximization and market envy-free

solutions; for markets with few buyers, proving (or disproving) the conjecture stated at the beginning of Section 3.2; for the multi-unit case with agent envy-free outcomes, determining an upper bound to the social welfare improvement achievable by preselection and settling the complexity of computing optimal solutions, for both objective functions.

Acknowledgements

References

Abebe, R., Kleinberg, J. M., & Parkes, D. C. (2017). Fair division via social comparison. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 281–289. ACM Press.

Alon, N., Mansour, Y., & Tennenholtz, M. (2013). Differential pricing with inequity aversion in social networks. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pp. 9–24. ACM Press.

Amanatidis, G., Markakis, E., & Sornat, K. (2016). Inequity aversion pricing over social networks: Approximation algorithms and hardness results. In *Proceedings of the 41st International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pp. 9:1–9:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

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

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

Balcan, M. F., Blum, A., & Mansour, Y. (2008). Item pricing for revenue maximization. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pp. 50–59. ACM Press.

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

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

Bilò, V., Flammini, M., Monaco, G., & Moscardelli, L. (2018a). On the impact of buyers preselection in pricing problems. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1871–1873. ACM Press.

Bilò, V., Flammini, M., Monaco, G., & Moscardelli, L. (2018b). Pricing problems with buyer preselection. In *Proceedings of the 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pp. 47:1–47:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Bouveret, S., & Lang, J. (2008). Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *J. Artif. Intell. Res.*, *32*, 525–564.

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

Briest, P., & Krysta, P. (2006). Single-minded unlimited supply pricing on sparse instances. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1093–1102. ACM Press.

Chalermsook, P., Chuzhoy, J., Kannan, S., & Khanna, S. (2012). Improved hardness results for profit maximization pricing problems with unlimited supply. In *Proceedings of the 15th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pp. 73–84. Springer.

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

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

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

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

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

Chen, N., & Rudra, A. (2008). Walrasian Equilibrium: Hardness, approximations and tractable instances. *Algorithmica*, *52*(1), 44–64.

Cheung, M., & Swamy, C. (2008). Approximation algorithms for single-minded envy-free profit-maximization problems with limited supply. In *Proceedings of the 49th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 35–44. IEEE Computer Society.

Chevaleyre, Y., Endriss, U., Estivie, S., & Maudet, N. (2007a). Reaching envy-free states in distributed negotiation settings. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1239–1244. ijcai.org.

Chevaleyre, Y., Endriss, U., & Maudet, N. (2007b). Allocating goods on a graph to eliminate envy. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI)*, pp. 700–705. AAAI Press.

Chevaleyre, Y., Endriss, U., & Maudet, N. (2017). Distributed fair allocation of indivisible goods. *Artificial Intelligence*, *242*, 1–22.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, *11*, 17–33.

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

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

Easley, D., & Kleinberg, J. (2010). *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press.

Elbassioni, K. M., Fouz, M., & Swamy, C. (2010). Approximation algorithms for non-single-minded profit-maximization problems with limited supply. In *Proceedings of the 6th International Workshop on Internet and Network Economics (WINE)*, pp. 462–472. Springer.

Feldman, M., Fiat, A., Leonardi, S., & Sankowski, P. (2012). Revenue maximizing envy-free multi-unit auctions with budgets. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pp. 532–549. ACM Press.

Feldman, M., Gravin, N., & Lucier, B. (2016). Combinatorial Walrasian Equilibrium. *SIAM Journal on Computing*, *45*(1), 29–48.

Flammini, M., Mauro, M., & Tonelli, M. (2018). On fair price discrimination in multi-unit markets. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 247–253. ijcai.org.

Flammini, M., Mauro, M., & Tonelli, M. (2019). On social envy-freeness in multi-unit markets. *Artif. Intell.*, *269*, 1–26.

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

Groves, T. (1973). Incentives in teams. *Econometrica*, *41*, 617–631.

Gul, F., & Stacchetti, E. (1999). Walrasian Equilibrium with gross substitutes. *Journal of Economic Theory*, *87*, 95–124.

Guruswami, V., Hartline, J. D., Karlin, A. R., Kempe, D., Kenyon, C., & McSherry, F. (2005). On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1164–1173. ACM Press.

Hartline, J., & Yan, Q. (2011). Envy, truth, and profit. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pp. 243–252. ACM Press.

Karp, J., Kazachkov, A. M., & Procaccia, A. D. (2014). Envy-free division of sellable goods. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 728–734. AAAI Press.

Kelso, A. S., & Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes. *Econometrica*, *50*, 1483–1504.

Kim, J. K., Kim, H. K., Oh, H. Y., & Ryu, Y. U. (2010). A group recommendation system for online communities. *International Journal of Information Management*, *30*(3), 212–219.

Lawler, E. L. (1979). Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, *44*(4), 339–356.

Leme, R. P., & Wong, S. C. (2017). Computing Walrasian Equilibria: Fast algorithms and structural properties. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 632–651. ACM Press.

Martello, S., & Toth, P. (1990). *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, Inc., USA.

Monaco, G., Sankowski, P., & Zhang, Q. (2015). Revenue maximization envy-free pricing for homogeneous resources. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 90–96. ijcai.org.

Nagle, T. T., & Holden, R. K. (2002). *The strategy and tactics of pricing : A guide to profitable decision making*. Prentice-Hall, USA.

Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press.

Nyman, K., Su, F. E., & Zerbib, S. (2020). Fair division with multiple pieces. *Discrete Applied Mathematics*, *283*, 115–122.

Ortega, J. (2018). Social integration in two-sided matching markets. *Journal of Mathematical Economics*, *78*, 119–126.

Searns, A. (2020). Rethinking resource allocation: Fairness and computability. Master's thesis, Rochester Institute of Technology. Accessed from https://scholarworks.rit.edu/theses/10651.

Segal-Halevi, E. (2018). Redividing the cake. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, p. 498–504. ijcai.org.

Segal-Halevi, E., & Nitzan, S. (2019). Fair cake-cutting among families. *Social Choice and Welfare*, *53*, 709–740.

Segal-Halevi, E., & Suksompong, W. (2019). Democratic fair allocation of indivisible goods. *Artificial Intelligence*, *277*, 103167.

Varian, H. R. (1974). Equity, envy, and efficiency. *Journal of Economic Theory*, *9*, 63–91.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*, 8–37.

Walras, L. (1954). *Elements of Pure Economics*. Allen and Unwin.

Weng, S. S., & Liu, M. J. (2004). Feature-based recommendations for one-to-one marketing. *Expert Systems with Applications*, *26*(4), 493–508.

Zuckerman, D. (2007). Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, *3*(1), 103–128.