On the Power and Limits of Dynamic Pricing in Combinatorial Markets

Ben Berger[1], Alon Eden[2(✉)], and Michal Feldman[1]

[1] Tel-Aviv University, Tel Aviv, Israel
benberger1@tauex.tau.ac.il, michal.feldman@cs.tau.ac.il
[2] Harvard University, Cambridge, USA
aloneden@seas.harvard.edu

Abstract. We study the power and limits of *optimal dynamic pricing* in combinatorial markets; i.e., dynamic pricing that leads to optimal social welfare. Previous work by Cohen-Addad *et al.* [EC'16] demonstrated the existence of optimal dynamic prices for unit-demand buyers, and showed a market with coverage valuations that admits no such prices. However, finding the most general class of markets (i.e., valuation functions) that admit optimal dynamic prices remains an open problem. In this work we establish positive and negative results that narrow the existing gap.

On the positive side, we provide tools for handling markets beyond unit-demand valuations. In particular, we characterize all optimal allocations in multi-demand markets. This characterization allows us to partition the items into equivalence classes according to the role they play in achieving optimality. Using these tools, we provide a poly-time optimal dynamic pricing algorithm for up to 3 multi-demand buyers.

On the negative side, we establish a maximal domain theorem, showing that for every non-gross substitutes valuation, there exist unit-demand valuations such that adding them yields a market that does not admit an optimal dynamic pricing. This result is the dynamic pricing equivalent of the seminal maximal domain theorem by Gul and Stacchetti [JET'99] for Walrasian equilibrium. Yang [JET'17] discovered an error in their original proof, and established a different, incomparable version of their maximal domain theorem. En route to our maximal domain theorem for optimal dynamic pricing, we provide the first complete proof of the original theorem by Gul and Stacchetti.

1 Introduction

We study the power and limitations of pricing schemes for social welfare optimization in combinatorial markets. We consider combinatorial markets with m

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), and by the Israel Science Foundation (grant number 317/17). The work of A. Eden was supported by the National Science Foundation under Grant No. CCF-1718549.

X. Chen et al. (Eds.): WINE 2020, LNCS 12495, pp. 206–219, 2020.
https://doi.org/10.1007/978-3-030-64946-3_15

heterogeneous, indivisible goods, and n buyers with publicly known valuation function $v_i : 2^{[m]} \to \mathbb{R}_{\geq 0}$ over bundles of items. The goal is to allocate items to buyers in a way that maximizes the social welfare.

Apart from being simple, pricing schemes are attractive since they do not require an all powerful central authority. Once the prices are set, the buyers arrive and simply choose a desired set of items from the available inventory. This is the mechanism we see everywhere, from supermarkets to online stores. Formally, the seller sets items prices $\mathbf{p} = (p_1, \ldots, p_m) \in \mathbb{R}^m_{\geq 0}$, buyers arrive sequentially in an arbitrary order, and every buyer chooses a bundle T from the remaining items that maximizes the utility: $u_i(T, \mathbf{p}) = v_i(T) - \sum_{j \in T} p_j$, breaking ties arbitrarily.

A reader familiar with the fundamental notion of Walrasian equilibrium, may conclude that the problem is solved for any market that admits a Walrasian equilibrium. A Walrasian equilibrium is a pair of an allocation $\mathbf{S} = (S_1, \ldots, S_n)$ and prices $\mathbf{p}$, such that for every buyer i, S_i maximizes i's utility given $\mathbf{p}$. By the first welfare theorem, every Walrasian equilibrium maximizes social welfare.

Are Walrasian prices a solution to our problem? The answer is no [1,3]. Walrasian prices cannot resolve a market without coordinating the tie breaking. If a buyer is faced with multiple utility-maximizing bundles, it is crucial that a central authority coordinates the tie breaking in accordance with the corresponding optimal allocation. In real-world markets, however, buyers are only faced with prices and choose a desired bundle by themselves without caring about global efficiency. [1] demonstrated that lacking a tie-breaking coordinator, Walrasian pricing can lead to arbitrarily bad welfare. Moreover, they showed that no fixed prices whatsoever can guarantee more than 2/3 of the optimal social welfare, even when restricted to unit-demand buyers.[1]

In order to circumvent this state of affairs, [1] proposed a more powerful pricing scheme, *dynamic pricing*, in which the seller updates prices in between buyer arrivals. The updated prices are set based on the remaining buyers and the current inventory. The main result of [1] is that every unit-demand market admits an optimal dynamic pricing. They also showed an example of a market with coverage valuations (a strict sub-class of submodular valuations) in which dynamic prices cannot guarantee optimal welfare. A natural question arises:

> *What markets (i.e., what valuation classes) can be resolved optimally using dynamic pricing?*

A similar question was considered for Walrasian equilibrium, where it was shown that every market with gross-substitutes buyers admits a Walrasian equilibrium [4]. Moreover, [2] show that gross-substitutes valuations are also maximal with respect to guaranteed existence of a Walrasian equilibrium:

Theorem 1 (Maximal Domain Theorem for Walrasian Equilibrium [2]). *Let v_1 be a non gross-substitutes valuation. Then, there exist unit-demand*

[1] A unit demand buyer has a value for every item, and the value for a set is the maximum value of any item in the set.

valuations $v_2, \ldots, v_\ell$ for some ℓ such that the valuation profile $(v_1, v_2, \ldots, v_\ell)$ does not admit a Walrasian equilibrium.

Although the notions of dynamic pricing and Walrasian are incomparable, Cohen-Addad *et al.* [1] conjectured that GS valuations are also maximal and sufficient with respect to the existence of dynamic prices. For the special case of markets with a unique optimal allocation, they showed that every GS market admits static prices guaranteeing optimal welfare, and there exists a market with non-GS (though submodular) valuations such that no pricing, even dynamic, guarantee optimal welfare.

1.1 Our Results and Techniques

In this work we shrink the known gap between markets that can and cannot be resolved optimally via dynamic pricing, from both ends.

A natural extension of unit-demand valuations is *multi-demand* valuations, where every buyer i has a public *cap* $k_i \in \mathbb{N}$ on the number of desired items, and the value for a set is the sum of the values for the k_i most valued items in the set. The case of $k_i = 1$ is simply unit-demand. Every multi-demand valuation is gross-substitutes. Our main positive result is the following:

Theorem 2. *Every market with up to* 3 *buyers, each with a multi-demand valuation function, admits an optimal dynamic pricing. Moreover, the prices can be computed in polynomial (in the number of items m) time, using value queries*[2]*.*

On the negative side, we show the first general negative result for dynamic prices, which takes the form of a maximal domain result in the spirit of [2]:

Theorem 3. *Let v_1 be a non gross-substitutes valuation. Then, there are unit-demand valuations $v_2, \ldots, v_\ell$ such that the valuation profile $(v_1, v_2, \ldots, v_\ell)$ does not admit an optimal dynamic pricing.*

En route, we provide *the first complete proof of the maximal domain theorem by Gul-Stacchetti* (Theorem 1 above), whose original proof was imprecise.

Techniques: Positive Results. We first review the solution of Cohen-Addad *et al.* [1] for unit-demand valuations, and show why we need a more fundamental technique in order to get past unit-demand bidders. Their scheme computes an optimal allocation $\mathbf{X} = (x_1, \ldots, x_n)$, where item x_i is allocated to buyer i, and then constructs a *complete*, weighted directed graph in which the vertices are the items. An edge $x_i \to x_j$ in this graph represents a *preference constraint*, requiring that buyer i *strongly* prefers item x_i over x_j, relative to the output prices. Hereafter, we term this graph the *preference graph*.

If there exist prices $\mathbf{p}$ that satisfy all edge constraints, then all buyers strongly prefer their items over the rest, and the allocation obtained after the last buyer

[2] A value query for a valuation v receives a set S as input, and returns $v(S)$.

leaves the market is precisely $\mathbf{X}$, which is optimal. Unfortunately, in some cases such prices do not exist. In order to circumvent this problem, [1] proves the following two claims:

- An edge $x_i \rightarrow x_j$ participates in a 0-weight cycle iff there is an alternative optimal allocation in which x_j is allocated to buyer i.
- If 0-weight cycles are removed from the graph, then one can compute prices that satisfy the remaining edge constraints.

Their pricing scheme removes every edge that participates in a 0-weight cycle, and then computes the prices as per the second bullet above. Relative to these prices, every buyer strongly prefers her allocated item to every other item, except perhaps for the set of items that are allocated to her in some alternative optimal allocation. Since every buyer takes at most one favorite item, as the buyers are unit-demand, this property guarantees that allocating this item to the buyer is consistent with an optimal allocation (not necessarily $\mathbf{X}$), as desired. When agents are multi-demand, they might take multiple items, and this breaks the solution by [1].

To illustrate this, consider the example given in Fig. 1, which serves as a running example throughout the paper. Removing the given 0-weight cycles could result in buyer 1 taking c and d instead of a and b, and the only remaining item that gives buyer 2 any positive value is e. This decreases the maximum attainable welfare from 5 to 4. The reason for this is that the two cycles intersect, and item e acts as a bottleneck for the two cycles. The machinery developed in [1] cannot identify the special role of item e, which is crucial for resolving this instance.

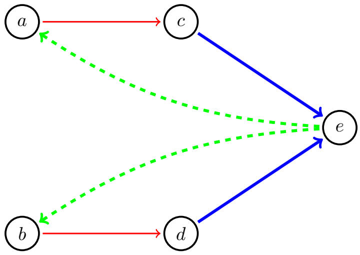

Fig. 1. Consider a market with 5 items a, b, c, d, e and 3 buyers, 1, 2, 3. Buyers 1 and 2 are both 2-demand, and buyer 3 is unit-demand. Buyer 1 values a, b, c, d at 1 and e at 0, Buyer 2 values c, d, e at 1 and a, b at 0, and buyer 3 values a, b, e at 1 and c, d at 0. One can verify that allocating a, b to 1, c, d to 2 and e to 3 maximizes social welfare. The figure depicts two 0-weight cycles in the preference graph constructed in the running example (edge weights are omitted). The thin red, thick blue and dashed green arrows correspond to the constraints of buyers 1, 2, 3 respectively.

Our first step is to gain a better structural understanding of optimal allocations in multi-demand markets. This is cast in the following theorem that characterizes the set of optimal allocations in multi-demand markets with any number of buyers. For the sake of simplicity, we present the theorem for markets in which all m items are allocated in every optimal allocation, and in which the total demand of the players equals supply, i.e. $m = \sum_{i=1}^{n} k_i$, where k_i is the cap of buyer i. In the full version, we show that an analogous result holds in the general case.

Theorem ***(Informal. See Theorem 4).*** *In a market with multi-demand buyers, an allocation is optimal if and only if the following hold:*

- *Every buyer i receives k_i items.*
- *If item x is allocated to buyer i, then there exists an optimal allocation where x is allocated to i.*

Put informally, the above states that one can mix-and-match items given to a buyer in *different* optimal allocations, and as long as each buyer i receives *exactly* k_i items, the resulting allocation is also optimal. While the only if direction is straightforward, it is not a-priori clear that the other direction holds as well. We prove this direction by reducing the problem to unit-demand valuations and proving for this case.

This characterization significantly simplifies the problem. It allows us to ignore the concrete values, and consider for each item only the set of buyers that receives it in some optimal allocation. Two items are essentially equivalent if their corresponding sets of buyers coincide. Thus, we group items into *equivalence classes*, providing a compact view of the market. For example, in markets with up to 3 multi-demand buyers, there are at most 8 (non-empty) equivalence classes corresponding to the possible subsets of players, while the total number of items can be arbitrarily large. We construct a new directed graph, termed the *item-equivalence graph*, where the vertices are these equivalence classes (refined after intersecting them with the bundles from the initial optimal allocation $\mathbf{X}$), and there is an edge $C \rightarrow D$ whenever the buyer that receives the items in C in the allocation $\mathbf{X}$ also receives every item in D in some optimal allocation. Figure 2 depicts the item-equivalence graph for the running example.

We show that there is a correspondence between cycles in the item-equivalence graph and 0-weight cycles in the preference graph. Thus our challenge is reduced to removing enough edges from the first (and translating these removals back to the second), in a way that eliminates all cycles, but also guarantees the following: every deviation by any buyer from her prescribed bundle, implied by the edge removals, allows the other buyers to simultaneously compensate for their "stolen" items by replacing them with items from other relevant equivalence classes. The optimal allocation characterization theorem then guarantees that the obtained allocation is indeed optimal. We devise an edge-removal method satisfying these requirements whenever the number of buyers is at most 3.

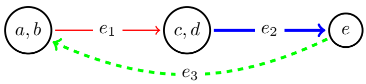

Fig. 2. The item-equivalence graph for the running example. E.g., the items a, b are equivalent since the set of buyers that receive any of them in some optimal allocation is the same ($\{1, 3\}$).

We believe this characterization theorem and the item equivalence graph may prove useful in other problems related to multi-demand markets.

Techniques: Negative Results. The original proof of Theorem 1 by Gul and Stacchetti considers two cases, and for each case, they construct a different market that does not admit a Walrasian equilibrium. Yang [7] showed one of the constructions does not work by finding an instance such that the constructed market does admit a Walrasian equilibrium. The error could not be easily fixed, and Yang proceeded by establishing an alternative, incomparable theorem; namely, that for every non gross-substitutes valuation there is a (single) gross-substitutes valuation for which the obtained market has no Walrasian equilibrium. While Yang's version of the assertion requires only a single valuation, this valuation has a complex structure compared with the simple unit-demand valuations in the original version. In the full version, we prove the maximal domain theorem *as originally stated*. The proof relies on a theorem which allows us to consider only the case with the correct construction in the original proof.

Our proof of Theorem 3, which is deferred to the full version, is driven by the following lemma from Cohen-Addad *et al.* [1]—in the case of a unique optimal allocation, the existence of optimal dynamic prices implies the existence of Walrasian prices. We modify the construction of Gul and Stacchetti to a market with an optimal allocation that is "almost" unique. This market still does not admit a Walrasian equilbrium. We then adapt the lemma in [1] to show that the existence of optimal dynamic prices in this market also implies the existence of Walrasian prices. The non-existence of Walrasian prices now implies non-existence of dynamic prices.

2 Preliminaries

We consider a setting with a finite set of indivisible items M (with $m := |M|$) and a set of n buyers (or players). Every buyer has a valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$. As standard, we assume monotonicity and normalization of all valuations, i.e. $v(S) \leq v(T)$ whenever $S \subseteq T$, and $v(\emptyset) = 0$. A valuation profile of n buyers is denoted $\mathbf{v} = (v_1, \ldots, v_n)$ and we assume that it is known by all. An *allocation* is a vector $\mathbf{A} = (A_1, \ldots, A_n)$ of disjoint subsets of M, indicating the bundles of items given to each player (not all items have to be allocated). The *social welfare* of an allocation $\mathbf{A}$ is given by $\mathsf{SW}(\mathbf{A}) = \sum_{i=1}^{n} v_i(A_i)$. An *optimal allocation* is an allocation that achieves the maximum social welfare among all allocations.

A *pricing* or a *price vector* is a vector $\mathbf{p} \in \mathbb{R}^n_{\geq 0}$ indicating the price of each item. We assume a quasi-linear utility, i.e. the *utility* of a buyer i from a bundle $S \subseteq M$ given prices $\mathbf{p}$ is $u_i(S, \mathbf{p}) = v_i(S) - \sum_{x \in S} p_x$. The *demand correspondence* of buyer i given $\mathbf{p}$ is the collection of utility maximizing bundles $D_{\mathbf{p}}(v) := \arg\max_{S \subseteq M}\{u(S, \mathbf{p})\}$.

Dynamic Pricing. In the dynamic pricing problem buyers arrive to the market in an arbitrary and unknown order. Before every buyer arrival new prices are set to the items that are still available, and these prices are based only on the set of buyers that have not yet arrived (their arrival order remains unknown). The arriving buyer then chooses an arbitrary utility-maximizing bundle based on the current prices and available items. The goal is to set the prices so that for any arrival order and any tie breaking choices by the buyers, the obtained social welfare is optimal.

We are interested in proving the guaranteed existence of an optimal dynamic pricing for any market composed entirely of buyers from a given valuation class C. It can be easily shown by induction that the problem is reduced to proving the guaranteed existence of item prices $\mathbf{p}$ such that any utility-maximizing bundle of any buyer can be completed to an optimal allocation. In other words, we can rephrase dynamic pricing as follows:

Definition 1. *An optimal dynamic pricing (hereafter, dynamic pricing) for the buyer profile* $\mathbf{v} = (v_1, \ldots, v_n)$ *is a price vector* $\mathbf{p} \in \mathbb{R}^m_{\geq 0}$ *such that for any buyer* i *and any* $S \in D_{\mathbf{p}}(v_i)$ *there is an optimal allocation in which player* i *receives* S.

3 Dynamic Pricing for Multi-demand Buyers

In this section we prove Theorem 2, namely we establish a dynamic pricing scheme for up to $n = 3$ multi-demand buyers that runs in $poly(m)$ time. As we shall see most of the tools we use hold for any number of buyers n. We fix a multi-demand buyer profile $\mathbf{v} = (v_1, \ldots, v_n)$ over the item set M, where each v_i is k_i-demand. We assume w.l.o.g. that all items are essential for optimality (i.e. all items are allocated in every optimal allocation) since otherwise we can price all unnecessary items at ∞ in every round to ensure that no player takes any of them (and price the rest of the items as if the unnecessary items do not exist). Note that under this assumption, each optimal allocation gives buyer i at most k_i items, for every i. In particular we have $m \leq \sum_{i=1}^n k_i$. For the sake of simplicity we further assume for the rest of this section that every optimal allocation gives each buyer i exactly k_i items, and thus $m = \sum_{i=1}^n k_i$. The case $m < \sum k_i$ introduces substantial technical difficulty. We show the solution for the general case in the full version. We first go over the tools used in our dynamic pricing scheme. We then present the dynamic pricing scheme for $n = 3$ buyers.

3.1 Tools and Previous Solutions

We start by presenting the main combinatorial construct of our solution, namely the preference-graph, which generalizes the construct given by [1] in their solution for unit-demand buyers. Then we explain the obstacles for generalizing the approach of [1] to the multi-demand setting. Finally, we develop the necessary machinery needed to overcome these obstacles. All the tools we develop and their properties hold for any number of buyers n.

The Preference Graph and an Initial Pricing Attempt. Let $\mathbf{O}$ be an arbitrary optimal allocation. The preference graph based on $\mathbf{O}$ is the directed graph H whose vertices are the items in M. Furthermore there is a special 'source' vertex denoted s. For any two different players i, j and items $x \in O_i, y \in O_j$ we have a directed edge $e = x \rightarrow y$ with weight $w(e) = v_i(x) - v_i(y)$. We also have a 0-weight edge $s \to x$ for every item $x \in M$. Since an optimal allocation can be computed in $poly(n, m)$ time with value queries (since the valuations are gross substitutes, see [5]), it follows that the preference graph can also be computed in $poly(n, m)$ time with value queries. When $|O_i| = 1$ for every i, the graph is exactly the one introduced by [1] in their unit-demand solution[3]. The proofs of the following two lemmas and corollary are deferred to the full version.

Lemma 1. *Let $C := x_1 \to x_2 \to \cdots x_k \to x_1$ be a cycle in H, where x_i is allocated to player i in $\mathbf{O}$ and $x_i \neq x_j$ for every $i \neq j$. Then the weight of the cycle is $w(C) = \mathsf{SW}(\mathbf{O}) - \mathsf{SW}(\mathbf{A})$ where $\mathbf{A}$ is the allocation obtained from $\mathbf{O}$ by transferring x_{i+1} to player i for every i (we identify player $k+1$ with player 1).*

Corollary 1. *Every cycle in H has non-negative weight.*

Corollary 1 implies that the weight of the min-weight path from s to x, denoted $\delta(s, x)$, is well-defined for any item x.

Lemma 2. *Let $p_x := -\delta(s, x)$ for every item x. Let i be some player, and let x, y be items such that $x \in O_i, y \notin O_i$. Then:* (1) $p_x \geq 0$. (2) $v_i(x) - p_x \geq v_i(y) - p_y$ (3) $v_i(x) - p_x \geq 0$.

Note that the utility player i obtains from any bundle of size at most k_i is the sum of the individual utilities obtained by the individual items. Thus, Lemma 2 shows that setting the prices $p_x = -\delta(s, x)$ almost achieves the requirements of dynamic pricing. However, since the inequalities in Lemma 2 are not strict, the incoming player might deviate from the designated bundle.

Solution for Unit-Demand Valuations and its Failure to Generalize. The inequalities of Lemmas 2 can be made strict by decreasing the weight of all edges by an appropriately selected $\varepsilon > 0$, but in the case H has zero-weight cycles, this can introduce negative cycles to H, in which case $\delta(s, x)$ is not defined for any x in such cycle. To circumvent this issue, [1] remove every edge that participates in a 0-weight cycle in H. Therefore, by choosing a small enough ϵ to decrease from the remaining edges, the remaining cycles are guaranteed to be strictly positive. Removing an edge $x \to y$ for $x \in O_i$, $y \notin O_i$ cancels the preference guarantee of Lemma 2 (part 2), leading to a possible deviation by buyer i from taking x to taking y. However, since 0-weight cycles correspond to alternative optimal allocations (see Lemma 1 with $w(C) = 0$), then this is not a problem: if the edge $x \to y$ was removed, then there is an optimal allocation in which player i receives y instead of x. As for the edges $x \to y$ that were not removed, the ϵ decrement causes i to strongly prefer x over y. The other inequalities of Lemma 2 would also be strict, and we are thus guaranteed that the incoming player indeed takes a one-item bundle that is part of some optimal allocation, as desired.

This approach works in the unit-demand setting, but poses problems in the multi-demand setting, as illustrated in the running example (presented in the

[3] A similar graph structure has been used by Murota in order to compute Walrasian equilibria in gross-substitutes markets [6].

introduction, see Fig. 1). Therefore, a more sophisticated method of eliminating 0-weight cycles must be employed instead of simply removing all edges that participate in some 0-weight cycle. Our informal goal is:

Remove a set of edges from the preference graph so that no 0-weight cycles are left, and every possible deviation implied by the removed edges is consistent with some optimal allocation.

Legal Allocations.

Definition 2

- *An item $x \in M$ is legal for player i if there is some optimal allocation $\mathbf{X} = (X_1, \ldots, X_n)$ such that $x \in X_i$.*
- *A bundle $S \subseteq M$ is legal for player i if $|S| = k_i$ and every $x \in S$ is legal for player i.*
- *A legal allocation $\mathbf{A} = (A_1, \ldots, A_n)$ is an allocation in which A_i is legal for player i, for every i.*

In a legal allocation every player i receives exactly k_i items, each of which is allocated to her in some optimal allocation. Note that a legal bundle for buyer i might not form a part of any optimal allocation (e.g., the bundle $\{c, d\}$ for buyer 1 in the running example). The following theorem provides a characterization of the collection of optimal allocations in the given market $\mathbf{v}$. The subsequent Corollary follows directly from the theorem and Definition 1. We next provide a proof sketch, for the full proof, we refer the reader to the full version.

Theorem 4. *An allocation is legal if and only if it is optimal.*

Proof (sketch). We first show that the theorem holds for unit-demand valuations, and then reduce the case of multi-demand valuations to unit-demand valuations. The *if* direction is trivial; we show the *only if* direction in this sketch. Let OPT denote the optimal welfare, and let $M_L = \{(i, \ell_i)\}_{i \in [n]}$ be a legal allocation (which is also a perfect matching since $m = \sum_i k_i = n$). For each edge (i, ℓ_i) in M_L, let M^i represent an optimal allocation, which is also a max-weight matching, in which (i, ℓ_i) participates (there must exist such a matching by the definition of legal allocations). Let $G = \bigcup_i M^i$ be the bipartite multigraph that is the union of all the perfect matchings. This is an n-regular bipartite graph which has M_L as a subgraph. We decompose G as follows—we first remove the matching M_L from G, resulting in an $n-1$ regular graph. We then decompose this graph into $n-1$ perfect matchings $M'_1, \ldots, M'_{n-1}$ (which is possible due to the regularity of the graph). For a matching M, we use $w(M)$ to denote its weight. We notice that $w(M_L) + \sum_{i=1}^{n-1} w(M'_i) = \sum_{i=1}^{n} w(M^i) = n \cdot \mathsf{OPT}$. Since $w(M'_i) \le \mathsf{OPT}$ for every i, it follows that $w(M_L) \ge \mathsf{OPT}$. Therefore, M_L is optimal.

We now describe the reduction: each k_i-demand buyer is decomposed into k_i identical unit-demand buyers, each of whom has the same value as the original buyer for every item j. The corresponding allocations are naturally defined: A single multi-demand buyer receiving $k \le k_i$ items in the original market corresponds to allocating these k items to k copies of this agent in the unit-demand market (one item to each copy). In the other direction, all the items allocated to copies of a particular multi-demand buyer are allocated to that buyer in the original market. It is not hard to see that an allocation is legal (resp. optimal) in the original market if and only if the corresponding allocation is legal (resp. optimal) in the corresponding unit-demand market. □

Corollary 2. *A price vector* $\mathbf{p}$ *is a dynamic pricing if for every player* i *and* $S \in D_{\mathbf{p}}(i)$*,* S *is legal for player* i *and there exists an allocation of the items* $M \setminus S$ *to the other players in which every player receives a legal bundle.*

Going back to our informal goal, Theorem 4 determines the deviations from the bundles O_i which are tolerable. A buyer can only deviate to a bundle which is legal for her, in a way that the leftover items can be partitioned "legally" among the rest of the buyers. The proof of Theorem 4 is deferred to the full version.

The Item-Equivalence Graph. Let $\mathbf{O}$ be some optimal allocation and H the corresponding preference graph. For every player i and set of players $C \subseteq [n]\setminus\{i\}$, we denote by $B_{i,C}$ the set of items allocated to buyer i in $\mathbf{O}$, and whose set of players to which they are legal is exactly $\{i\} \cup C$. For example, $B_{1,\{2,3\}}$ is the set of items $x \in O_1$ such that there are optimal allocations $\mathbf{O}'$,$\tilde{\mathbf{O}}$ in which x is allocated to players 2, 3 (respectively), and for any other player $j \notin \{1, 2, 3\}$, there is no optimal allocation in which x is allocated to j. Formally,

$$B_{i,C} := \left\{ x \in O_i \;\middle|\; \begin{array}{l} \forall j \in \{i\} \cup C \ x \text{ is legal for } j \\ \forall j \notin \{i\} \cup C \ x \text{ is not legal for } j \end{array} \right\}$$

We make a few observations:

- The sets $B_{i,C}$ form a partition of M (some of these sets might be empty sets).
- Let $x \in O_i$ and $y \in O_j$ for $i \neq j$. If $x \to y$ participates in a 0-weight cycle in H and $y \in B_{j,C}$, then $i \in C$.

The second observation holds since if $x \to y$ participates in a 0-weight cycle, then there is an alternative optimal allocation in which y is allocated to player i (see Lemma 1 with $w(C) = 0$).

Definition 3 (Item-Equivalence Graph). *Given an optimal allocation* $\mathbf{O}$*, its associated item-equivalence graph is the directed graph* $B = (T, D)$ *with vertices* $T = \{B_{i,C} \neq \emptyset \mid i \in [n], C \subseteq [n] \setminus \{i\}\}$ *and directed edges* $D = \{B_{i,C_1} \to B_{j,C_2} \mid i \in C_2\}$*.*

For example, $(B_{1,\emptyset} \to B_{2,\{1,4\}})$ and $(B_{2,\{1,5\}} \to B_{6,\{2\}})$ are edges in the item-equivalence graph (assuming that the participating sets are non-empty), whereas, for example, $(B_{1,\emptyset} \to B_{1,\{2\}})$ and $(B_{2,\{1\}} \to B_{3,\{1,4\}})$ are not. Note also that the number of vertices is at most m.

The next Lemma shows that the item-equivalence graph can be computed efficiently. The proof is deferred to the full version.

Lemma 3. *Given an optimal allocation* $\mathbf{O}$*, its associated item-equivalence graph can be computed in poly*(m, n) *time and value queries.*

The following lemma uses Theorem 4 to establish a correspondence between 0-weight cycles in H and cycles in B. Its proof is deferred to the full version.

Lemma 4. *Let* $\mathbf{O}$ *be an optimal allocation and let* H *and* B *be the corresponding preference graph and item-equivalence graph, respectively. Then:*

1. *If* $B_{i_1,C_1} \rightarrow \cdots \rightarrow B_{i_k,C_k} \rightarrow B_{i_1,C_1}$ *is a cycle in* B *then for any items* $x_1 \in B_{i_1,C_1}, \ldots, x_k \in B_{i_k,C_k}$*, the cycle* $C = x_1 \rightarrow x_2 \rightarrow \cdots \rightarrow x_k \rightarrow x_1$ *is a 0-weight cycle in* H*.*
2. *If* $C = x_1 \rightarrow x_2 \rightarrow \cdots \rightarrow x_k \rightarrow x_1$ *is a 0-weight cycle in* H*, and* $x_\ell \in B_{i_\ell,C_\ell}$ *for every* $1 \leq \ell \leq k$ *then* $C' := B_{i_1,C_1} \rightarrow \cdots \rightarrow B_{i_k,C_k} \rightarrow B_{i_1,C_1}$ *is a cycle in* B*.*

As explained before, our main challenge in the dynamic pricing problem is to come up with a method to remove all 0-weight cycles from H in a way that each potential deviation of any player i from the designated bundle O_i, that emanates from the edge removals, is consistent with some optimal allocation. In particular the method must overcome the "bottleneck problem" (as illustrated in Fig. 1). Lemma 4 allows us to shift the focus from removing 0-weight cycles in H to removing cycles in B and translate these removals back to H.

[**Running Example**]**.** Figure 2 shows the item-equivalence graph obtained from the initial optimal allocation. Each of the items a, b is allocated to buyer 3 in some other optimal allocation, and is never allocated to buyer 2. Thus $a, b \in B_{1,\{3\}}$. Similarly we have $c, d \in B_{2,\{1\}}$ and $e \in B_{3,\{2\}}$. Removing any edge of the item-equivalence graph makes it cycle-free. Thus, by Lemma 4, if we choose one of the edges e_1, e_2, e_3 and remove all edges in the preference graph corresponding to the chosen edge, then the preference graph will remain cycle-free. Removing the edges corresponding to e_1 could cause player 1 to take the bundle $\{c, d\}$ instead of the designated bundle $\{a, b\}$, which cannot be completed to an optimal allocation. On the other hand, removing the preference graph edges that correspond to the edges e_2, e_3 is fine. If player 2 arrives first to the market, then the removal of edge e_2 might cause her to take the item e instead of c or d, and both options are consistent with some optimal allocation. Likewise if player 3 arrives first and takes a or b instead of e then this too can be completed to an optimal allocation. The important property here is that $B_{3,\{2\}}$ has minimal size in the cycle, and thus removing its incoming and outgoing edges introduces tolerable potential deviations.

3.2 Solution for Up to 3 Multi-demand Buyers

We are now ready to present the dynamic pricing scheme for up to $n = 3$ multi-demand buyers. The algorithm makes use of the item-equivalence graph. We abuse notation and instead of writing $B_{i,\{j\}}$ ($B_{i,\{j,k\}}$) we write B_{ij} (B_{ijk}). Thus the vertices of the item-equivalence graph for 3 buyers are $B_{1,\emptyset}$, $B_{2,\emptyset}, B_{3,\emptyset}$, B_{12}, B_{21}, B_{31}, B_{123}, B_{213}, B_{312}, B_{13}, B_{23}, B_{32} (only the non-empty sets out of these appear in the graph). For 2 buyers there are at most 4 vertices in the graph: $B_{1,\emptyset}$, $B_{2,\emptyset}$, B_{12}, B_{21}. Step 4 is only relevant for the case of 3 buyers.

ALGORITHM 1: Dynamic Pricing Scheme for up to 3 Multi-Demand Buyers.

Input: Multi-demand valuations v_1, v_2, and also v_3 when $n = 3$.
Output: prices $\mathbf{p} = (p_x)_{x \in M}$.

1 Compute some optimal allocation $\mathbf{O}$.
2 Compute the preference graph H and the item-equivalence graph B based on $\mathbf{O}$.
3 Mark all edges that participate in a cycle of size 2 in B.
4 In each of the cycles $B_{13} \to B_{21} \to B_{32} \to B_{13}$ and $B_{12} \to B_{31} \to B_{23} \to B_{12}$ (if these exist) choose a set of minimal size and mark its incoming edge and outgoing edge in the cycle.
5 For every edge $B_{i_1,C_1} \to B_{i_2,C_2}$ in B that was marked, and for every $x \in B_{i_1,C_1}, y \in B_{i_2,C_2}$, remove the edge $x \to y$ from H. Denote the obtained graph by H'.
6 Let $\Delta > 0$ be the difference in social welfare between the optimal and 2nd optimal allocation. Denote $\epsilon := \frac{\Delta}{m+1}$ and for every edge e that was not removed (except for edges starting at the source vertex s) update its weight to $w'(e) = w(e) - \epsilon$.
7 Compute the min-weight paths from s to every x in H', and let $\delta(s,x)$ be its weight. For every item x set the price $p_x = -\delta(s,x) + \epsilon$.
8 **return** $(p_x)_{x \in M}$

When $n = 2$, the only cycle in the item-equivalence graph is $B_{12} \to B_{21} \to B_{12}$ (assuming both of these are non-empty sets), and both of its edges were marked in step 3. Thus, by Lemma 4, all edges that participate in a 0-weight cycle in the preference graph were removed in step 5. Thus for $n = 2$ Algorithm 1 is, effectively, the straightforward generalization of the Cohen-Addad *et al.* [1] unit-demand solution to multi-demand buyers.

As stated before, computing $\mathbf{O}$, H and B can be done in polynomial time. Finding the cycles in B can also be done efficiently (B has a constant number of vertices) as well as computing min-weight paths. Thus the algorithm indeed runs in $poly(m)$ time as desired. The proofs of the following 4 lemmas are deferred to the full version.

Lemma 5. *After step 5 every cycle in H' has strictly positive weight.*

Lemma 6. *For any item x, $p_x > 0$.*

Lemma 7. *For any player i, $x \in O_i$ and $y \notin O_i$, if $e = x \to y \in H'$ then $u_i(x, \mathbf{p}) > u_i(y, \mathbf{p})$.*

Lemma 8. *For any player i and item $x \in O_i$ we have $v_i(x) - p_x > 0$.*

We are now ready to prove that the output of our dynamic pricing scheme meets the requirements of Corollary 2. This is cast in the following lemma:

Lemma 9. *Let $\mathbf{p}$ be the price vector output by Algorithm 1. Then, for every player i and $S \in D_{\mathbf{p}}(i)$, (a) S is legal for player i; and (b) S can be completed to a legal allocation, i.e. there exists an allocation of the items $M \setminus S$ to the other players in which every player receives a bundle that is legal for her.*

Proof. We prove for $i = 1$ (the same proof applies also for $i = 2, 3$). We first prove part (a). We start by showing that every $S \in D_{\mathbf{p}}(1)$ is of size k_1. Since all item prices are positive (Lemma 6) and player 1 is k_1-demand, it cannot be the case that player 1 maximizes utility with a bundle consisting of more than k_1 items. Furthermore, by Lemma 8 there are at least k_1 legal items from which she derives positive utility. Combining, every demanded bundle has exactly k_1 items. Now, for any two items x, y where $x \in O_1$ and y is not legal for player 1, the edge $x \to y$ was not removed in the transition from H to H' (since there is no corresponding edge in the item-equivalence graph that could have been marked). Thus, player 1 strongly prefers x over y (by Lemma 7) and we conclude that every demanded bundle contains only legal items, as desired.

We proceed to prove part (b). Let $S \in D_{\mathbf{p}}(1)$. We refer to the items in $S \setminus O_1$ as the items that player 1 'stole' from players 2 (and 3 if $n = 3$), and to the items in $O_1 \setminus S$ as those player 1 'left behind'. We need to show that players 2 and 3 can compensate for their stolen items in a 'legal manner', that is, by completing their leftover bundles $O_2 \setminus S$ and $O_3 \setminus S$ to k_2 and k_3 legal items, respectively. The first step is to determine where the stolen and left behind items are taken from. Since $B_{1,\emptyset}$ does not participate in any cycle in the item-equivalence graph (as it has no incoming edge), then none of its outgoing edges were marked, implying (by Lemma 7) that player 1 strongly prefers every item of $B_{1,\emptyset}$ over every item $y \notin O_1$. Since buyer 1 derives positive utility from these items (Lemma 8), we conclude that $B_{1,\emptyset}$ is contained in every demanded bundle and in particular in S. In other words, all the items player 1 left-behind are in B_{12} if $n = 2$, or in $B_{12} \cup B_{13} \cup B_{123}$ if $n = 3$. Thus, if $n = 2$ we are done: buyer 2 can compensate for her stolen items by taking the leftover items in B_{12} which are legal for her (the amount of stolen items equals the amount of leftover items since $|O_1| = |S| = k_1$). We assume for the rest of the proof that $n = 3$. Since S is legal for buyer 1, the stolen items $S \setminus O_1$ are contained in $B_{21} \cup B_{213} \cup B_{31} \cup B_{312}$.

Denote $a_2 := |(O_1 \setminus S) \cap B_{12}|, a_3 := |(O_1 \setminus S) \cap B_{13}|, a_{23} := |(O_1 \setminus S) \cap B_{123}|,\ b_2 := |(S \setminus O_1) \cap B_{21}|, b_{23} := |(S \setminus O_1) \cap B_{213}|, b_3 := |(S \setminus O_1) \cap B_{31}|, b_{32} := |(S \setminus O_1) \cap B_{312}|$.

In words, a_2 is the number of items player 1 left behind in B_{12}, b_2 is the number of items she 'stole' from player 2 out of the items in B_{21}, b_{32} is the amount she 'stole' from player 3 out of the items in B_{312}, etc. By the discussion in the previous paragraph, these account for all stolen and leftover items, and we get

$$b_2 + b_{23} + b_3 + b_{32} = |S \setminus O_1| = |O_1 \setminus S| = a_2 + a_{23} + a_3. \quad (1)$$

Consider the bipartite graph G whose left side consists of the items in $S \setminus O_1$ and whose right side consists of the items in $O_1 \setminus S$, with edges (x, y) whenever the stolen item x can be replaced by the leftover item y legally (e.g., if $x \in O_2$, then $y \in B_{12} \cup B_{123}$). Specifically, G is composed of a bi-clique between the stolen items from $B_{21} \cup B_{213}$ (the stolen items of player 2) and the leftover items from $B_{12} \cup B_{123}$ (these are the leftover items that are legal for player 2), and of another bi-clique between the stolen items of $B_{31} \cup B_{312}$ (the stolen items of player 3) and the leftover items of $B_{13} \cup B_{123}$ (the leftover items that are legal for player

3). If there is a perfect matching in G, then every stolen item can be replaced with the item it was matched to in the perfect matching, resulting in a legal allocation, and we are done. Thus we assume that there is no perfect matching in G. Since Hall's condition does not hold, we have that either $b_2+b_{23} > a_2+a_{23}$, or $b_3+b_{32} > a_3+a_{23}$. Assume w.l.o.g. that $b_2+b_{23} > a_2+a_{23}$. Then, by Eq. (1), we have $a_3 > b_3+b_{32} \geq 0$. We claim that this implies $b_{23} = 0$. The reason is that otherwise, player 1 stole some item, denoted y, from B_{213} and left behind some item, denoted x, in B_{13}. But this cannot be the case since this would imply (by Lemma 7) that the edge $x \rightarrow y$ was removed in the transition from H to H', but the edge $B_{13} \rightarrow B_{213}$ was never marked in the pricing scheme. Therefore $b_{23} = 0$ and $b_2 > a_2 + a_{23} \geq 0$. The combination of $b_2 > 0$ and $a_3 > 0$ implies that the edge $B_{13} \rightarrow B_{21}$ was marked in step 4, and so one of B_{13}, B_{21} is of minimal size in the cycle $B_{13} \rightarrow B_{21} \rightarrow B_{32} \rightarrow B_{13}$. In particular,

$$\begin{aligned}|B_{32}| &\geq \min\{|B_{13}|, |B_{21}|\} \geq \min\{a_3, b_2\} \\ &\geq \min\{a_3 - (b_3 + b_{32}), b_2 - (a_2 + a_{23})\} = b_2 - (a_2 + a_{23}),\end{aligned}$$

where the equality holds by Eq. (1). In order to complete S to a legal allocation, player 2 compensates for his stolen b_2 items by taking the $a_2 + a_{23}$ items player 1 left behind in $B_{12} \cup B_{123}$ and by "stealing" $b_2 - (a_2 + a_{23})$ items from B_{32} (indeed there are enough items there for player 2 to steal). Player 3 now has to compensate for the items stolen from her by both players, a total of $(b_{32} + b_3) + (b_2 - (a_2 + a_{23})) = a_3$ items. Since player 1 left precisely this number of items in B_{13}, player 3 can take them. Note that the resulting allocation is indeed legal and thus optimal. □

Acknowledgments. We deeply thank Amos Fiat and Renato Paes-Leme for helpful discussions.

References

1. Cohen-Addad, V., Eden, A., Feldman, M., Fiat, A.: The invisible hand of dynamic market pricing. In: Proceedings of the 2016 ACM Conference on Economics and Computation, EC 2016, Maastricht, NL, 24–28 July 2016, pp. 383–400 (2016)
2. Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. J. Econ. Theory **87**(1), 95–124 (1999)
3. Hsu, J., Morgenstern, J., Rogers, R.M., Roth, A., Vohra, R.: Do prices coordinate markets? In: Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, 18–21 June 2016, pp. 440–453 (2016)
4. Kelso Jr., A.S., Crawford, V.P.: Job matching, coalition formation, and gross substitutes. Econ. J. Econ. Soc., 1483–1504 (1982)
5. Leme, R.P.: Gross substitutability: an algorithmic survey. Games Econ. Behav. **106**, 294–316 (2017)
6. Murota, K.: Valuated matroid intersection I: optimality criteria. SIAM J. Discrete Math. **9**(4), 545–561 (1996)
7. Yang, Y.: On the maximal domain theorem: a corrigendum to "Walrasian equilibrium with gross substitutes". J. Econ. Theory **172**, 505–511 (2017)