



Faster Dynamic Auctions via Polymatroid Sum

KATHARINA EICKHOFF, School of Business and Economics, RWTH Aachen University, Aachen, Germany

MEIKE NEUWOHNER, Department of Mathematics, London School of Economics and Political Science, London, United Kingdom

BRITTA PEIS, School of Business and Economics, RWTH Aachen University, Aachen, Germany

NIKLAS RIEKEN, School of Business and Economics, RWTH Aachen University, Aachen, Germany

LAURA VARGAS KOCH, Research Institute for Discrete Mathematics and Hausdorff Center for Mathematics, University of Bonn, Bonn, Germany and School of Business and Economics, RWTH Aachen University, Aachen, Germany

LÁSZLÓ A. VÉGH, London School of Economics and Political Science, London, United Kingdom, Institute for Advanced Studies, Corvinus University of Budapest, Budapest, Hungary, and Hertz Chair for Algorithms and Optimization, University of Bonn, Bonn, Germany

We consider dynamic auctions for finding Walrasian equilibria in markets with indivisible items and strong gross substitutes valuation functions. Each price adjustment step in these auction algorithms requires finding an inclusion-wise minimal maximally overdemanded set or an inclusion-wise minimal maximally underdemanded set at the current prices. Both can be formulated as a submodular function minimization problem.

We observe that minimizing this submodular function corresponds to a polymatroid sum problem, and using this viewpoint, we give a fast and simple push-relabel algorithm for finding the required sets. This improves on the previously best running time of Murota, Shioura and Yang (ISAAC 2013). Our algorithm is an adaptation of the push-relabel framework by Frank and Miklós (JJIAM 2012) to the particular setting. We obtain a further improvement for the special case of unit-supplies.

We further show the following monotonicity properties of Walrasian prices: both the minimal and maximal Walrasian prices can only increase if supply of goods decreases, or if the demand of buyers increases. This

An abstract of this article appeared in the proceedings of WINE 2023 [27].

Katharina Eickhoff was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 2236/2. Meike Neuwöhner was supported by the Engineering and Physical Sciences Research Council, part of UK Research and Innovation, grant ref. EP/X030989/1.

László A. Végh received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement no. ScaleOpt– 757481).

Authors’ Contact Information: Katharina Eickhoff, School of Business and Economics, RWTH Aachen University, Aachen, Germany; e-mail: eickhoff@oms.rwth-aachen.de; Meike Neuwöhner, Department of Mathematics, London School of Economics and Political Science, London, United Kingdom; e-mail: m.neuwohner@lse.ac.uk; Britta Peis, School of Business and Economics, RWTH Aachen University, Aachen, Germany; e-mail: peis@oms.rwth-aachen.de; Niklas Rieken, School of Business and Economics, RWTH Aachen University, Aachen, Germany; e-mail: rieken@oms.rwth-aachen.de; Laura Vargas Koch, Research Institute for Discrete Mathematics and Hausdorff Center for Mathematics, University of Bonn, Bonn, Germany and School of Business and Economics, RWTH Aachen University, Aachen, Germany; e-mail: vargaskoch@gdm.rwth-aachen.de; László A. Végh, London School of Economics and Political Science, London, United Kingdom, Institute for Advanced Studies, Corvinus University of Budapest, Budapest, Hungary, Hertz Chair for Algorithms and Optimization, University of Bonn, Bonn, Germany; e-mail: lvegh@uni-bonn.de.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2167-8375/2025/06-ART13

<https://doi.org/10.1145/3729429>

is derived from a fine-grained analysis of market prices. We call *packing prices* a price vector such that there is a feasible allocation where each buyer obtains a utility maximizing set. Conversely, by *covering prices* we mean a price vector such that there exists a collection of utility maximizing sets of the buyers that include all available goods. We show that for strong gross substitutes valuations, the component-wise minimal packing prices coincide with the minimal Walrasian prices and the component-wise maximal covering prices coincide with the maximal Walrasian prices. These properties in turn lead to the price monotonicity results.

CCS Concepts: • **Theory of computation** → **Design and analysis of algorithms**; • **Mathematics of computing** → **Matroids and greedoids**;

Additional Key Words and Phrases: Dynamic auctions, walrasian prices, strong gross substitutes, polymatroid sum, push-relabel, monotonicity

ACM Reference Format:

Katharina Eickhoff, Meike Neuwöhner, Britta Peis, Niklas Rieken, Laura Vargas Koch, and László A. Végh. 2025. Faster Dynamic Auctions via Polymatroid Sum. *ACM Trans. Econ. Comput.* 13, 3, Article 13 (June 2025), 47 pages. <https://doi.org/10.1145/3729429>

1 Introduction

We study computational and structural aspects of Walrasian market equilibria. Let us consider a market involving a set of buyers with quasi-linear utilities that are interested in buying a set of goods. The main objective is to find an allocation of the goods to buyers along with prices for the goods which should fulfill the following properties: (i) every buyer gets a preferred bundle, i.e., a set of items maximizing her utility at the given prices, (ii) no good is oversold, (iii) all goods are sold, and (iv) the allocation is Pareto-optimal. Properties (i)–(iii) together are the *Walrasian equilibrium* conditions, which coincidentally imply property (iv); this is often referred to as *First Welfare Theorem*. In this article, we are interested in finding such an equilibrium allocation and prices via a *dynamic auction*. In simple terms, in a dynamic auction, the auctioneer iteratively announces prices for goods and the buyers report their demand correspondence. If there is a feasible way to satisfy all the demands (i.e., properties (i) and (ii) are satisfied) and all goods are distributed to the buyers (i.e., property (iii) is satisfied in addition) then the auction terminates with the announced prices and an allocation that is feasible. Otherwise, prices are adjusted in an intuitive way: goods that are overdemanded increase in price, whereas underdemanded goods decrease in price.

It is easy to see that such a dynamic auction can take longer than a static *sealed-bid auction*, in which the auctioneer receives all the buyers' information at the start instead of incrementally after announcing new prices. There are several reasons why a dynamic auction may be preferable (see [4, 20] for a detailed discussion). First of all, dynamic auctions as described above are more transparent; buyers observe how some goods increase or decrease in price, whereas in a sealed-bid auction, the auctioneer just sets a price whose origin remains opaque to the buyers as the auctioneer is the only agent with full information. Typically, the buyers' valuation functions (and hence, the utilities upon receiving a bundle at a given price) are private information and buyers do not want to disclose this to an auctioneer in full if it is not necessary. Moreover, while one can argue that sealed-bid auctions might run faster given complete information of buyers (see [54]), this is only true under the assumption that there are no communication costs between buyers and auctioneer; after all, a valuation function of a buyer has to specify a value for every subset of goods (i.e., exponentially many in the number of goods) and all these values have to be communicated to the auctioneer. In fact, one might even argue that the exact valuation functions might not be known to the buyers themselves if they are only implicitly given in form of some combinatorial structure, so some of the computational tasks are just outsourced to the buyers.

As mentioned above, a dynamic auction iteratively adjusts prices by raising prices on underdemanded, or overdemanded sets. Informally speaking, overdemanded sets are those sets of items for which the demand exceeds the supply, and maximally overdemanded sets are those overdemanded sets that maximize the difference between demand and supply (analogously, we can define underdemanded sets and maximally underdemanded sets). A precise definition is given in Definition 1.1. We study four different types of dynamic auctions that are listed below (see Section 2.2 for a precise description). They are known to terminate with Walrasian price vectors, assuming gross substitutes valuations. The first [second] type of auction computes the unique component-wise minimal [maximal] Walrasian price vector by raising [decreasing] prices on the inclusion-wise minimal set among all maximally overdemanded [underdemanded] sets in each iteration. We call this set the *minimal maximally overdemanded [underdemanded] set*. In this article, we present fast algorithms for the problem of finding the minimal maximally overdemanded [underdemanded] set.

- (I) An ascending auction that finds minimal Walrasian prices from any lower bound on prices,
- (II) a descending auction that finds maximal Walrasian prices from any upper bound on prices,
- (III) a two-phase auction, that finds some Walrasian prices from any initial price vector, and
- (IV) a greedy auction, that finds some Walrasian prices from any initial price vector.

The auctions (I)–(III) can be viewed as efficient implementations of the corresponding auctions by Ausubel [6], whereas the auction (IV) is a speeded-up version of the auction by Murota et al. [52].

Moreover, we are able to show that the minimal and maximal Walrasian prices have the following properties, where we refer to the component-wise ordering among price vectors.

- (1) Every price vector satisfying (i) and (ii) is at least as large as the minimal Walrasian price vector,
- (2) every price vector satisfying (i) and (iii) is at most as large as the maximal Walrasian price vector, and
- (3) minimal and maximal Walrasian price vectors react naturally to changes in demand and supply, i.e., they cannot decrease if demand increases or supply decreases and they cannot increase if demand decreases or supply increases.

To start with, we first have a closer look into the economic model.

1.1 The Economic Model

In this article, we consider the following market. Let $E = \{1, \dots, m\}$ be a set of m item types (we will just refer to those as *items*), and $N = \{1, \dots, n\}$ be a set of n buyers. Each item $e \in E$ is available in $b(e)$ indivisible units for some $b \in \mathbb{Z}_+^E$, and each buyer $i \in N$ has an individual valuation function $v_i: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}_+$ over the set $[0, b]_{\mathbb{Z}} := \{z \in \mathbb{Z}^m \mid \mathbf{0} \leq z \leq b\}$ of possible *bundles*. For $z \in [0, b]_{\mathbb{Z}}$, the value $z(e)$ represents the number of items of type $e \in E$. The value $v_i(z) \in \mathbb{Z}_+$ can be interpreted as buyer i 's maximum willingness to pay for a bundle $z \in [0, b]_{\mathbb{Z}}$. Throughout the article, we assume that all valuation functions are *non-decreasing*, i.e., for all $z, z' \in [0, b]_{\mathbb{Z}}$ with $z \leq z'$, we have $v_i(z) \leq v_i(z')$. We let $B := \max_{e \in E} b(e)$. We refer to the case where $b(e) = 1$ for all $e \in E$ as the *unit-supply* case, and to the general problem as the *multi-supply* case.

For a given price vector $p \in \mathbb{R}_+^E$, the *utility function* of buyer i is given by $u_i(z) = v_i(z) - \langle p, z \rangle$. Buyer i 's *indirect utility function* is given by

$$V_i(p) := \max \{v_i(z) - \langle p, z \rangle \mid z \in [0, b]_{\mathbb{Z}}\}.$$

Let us denote by $\mathcal{D}_i(p)$ the *set of preferred bundles* (also called *demand set*), which are the maximizers of u_i , that is,

$$\mathcal{D}_i(p) := \arg \max \{v_i(z) - \langle p, z \rangle \mid z \in [0, b]_{\mathbb{Z}}\}.$$

Further let $\check{\mathcal{D}}_i(p)$ denote the *set of minimal preferred bundles* and $\hat{\mathcal{D}}_i(p)$ denote the *set of maximal preferred bundles* of buyer i for a given price vector p , i.e.,

$$\check{\mathcal{D}}_i(p) := \{z \in \mathcal{D}_i(p) \mid \text{there is no } y \in \mathcal{D}_i(p) \setminus \{z\} \text{ with } y \leq z\}$$

and

$$\hat{\mathcal{D}}_i(p) := \{z \in \mathcal{D}_i(p) \mid \text{there is no } y \in \mathcal{D}_i(p) \setminus \{z\} \text{ with } y \geq z\}.$$

Using $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$, we can already quantify whether the supply for a set $S \subseteq E$ is too high or too low with respect to the current prices p . To do so, we define

$$\check{\theta}_i^p(S) := \min\{z(S) \mid z \in \check{\mathcal{D}}_i(p)\} \quad \text{and} \quad \hat{\rho}_i^p(S) := \max\{z(S) \mid z \in \hat{\mathcal{D}}_i(p)\}.$$

Informally speaking, $\check{\theta}_i^p(S)$ denotes the minimum number of items buyer i has to receive from S , whereas $\hat{\rho}_i^p(S)$ denotes the maximum number of items buyer i is willing to take from set S . With this at hand, we can define the *overdemandedness* of a set S to be

$$\text{overd}^p(S) := \sum_{i \in N} \check{\theta}_i^p(S) - b(S).$$

Similarly, we can define the *underdemandedness* of a set S to be

$$\text{underd}^p(S) := \sum_{i \in N} b(S) - \hat{\rho}_i^p(S).$$

Definition 1.1. We call a set $S \subseteq E$ *overdemanded* [*underdemanded*] w.r.t. prices p if $\text{overd}^p(S) > 0$ [$\text{underd}^p(S) > 0$]. Moreover, a set $S \subseteq E$ is called *maximally overdemanded* [*maximally underdemanded*] if it is overdemanded [*underdemanded*] and maximizes the overdemandedness $\text{overd}^p(S)$ [*underdemandedness* $\text{underd}^p(S)$] among all bundles $S \subseteq E$.

We call S *minimal maximally overdemanded* [*underdemanded*] if S is the inclusion-wise minimal set among all maximally overdemanded [*underdemanded*] sets.

The concept of overdemanded sets has already been introduced in [38], where overdemandedness is called *excess demand*.¹ This concept generalizes the notion overdemandedness introduced in [23] for the case of buyers with unit demands. In [61], the notion of underdemandedness has been introduced, again for the setting of buyers with unit demand.

An allocation (z_1, \dots, z_n) is called *packing* with respect to the prices $p \in \mathbb{R}_+^E$ if

$$z_i \in \mathcal{D}_i(p) \text{ for all } i \in N \quad \text{and} \quad \sum_{i \in N} z_i(e) \leq b(e) \text{ for all } e \in E. \quad (\text{packing})$$

Packing refers to the fact that one preferred bundle per buyer can be chosen such that the sets can be packed into the set of items (no item is assigned too often). A price vector p that admits a packing allocation is also called *packing*.² We will see in Section 3 that a price vector p is packing if and only if there is no overdemanded set w.r.t. p . Note that it is easy to get packing prices by just setting them sufficiently high such that no buyer is interested in any item anymore.

An allocation (z_1, \dots, z_n) is called *covering* with respect to the prices $p \in \mathbb{R}_+^E$ if

$$z_i \in \mathcal{D}_i(p) \text{ for all } i \in N \quad \text{and} \quad \sum_{i \in N} z_i(e) \geq b(e) \text{ for all } e \in E. \quad (\text{covering})$$

Here, covering refers to the fact that there is a preferred bundle per buyer such that all items are covered by these sets. We may also call a price vector *covering* if there is a supporting covering

¹Note that Gul and Stacchetti [38] consider the unit-supply setting, meaning that there is one item of each object type ($b(e) = 1$ for all $e \in E$).

²Note that [9] call packing prices *envy-free prices*.

allocation. In Section 3, we will see that a price vector p is covering if and only if there is no underdemanded set.

A price vector $p \in \mathbb{R}_+^E$ is called a *Walrasian price vector* [67] or a *competitive equilibrium price vector* if there is an allocation (z_1, \dots, z_n) with

$$z_i \in \mathcal{D}_i(p) \text{ for all } i \in N \quad \text{and} \quad \sum_{i \in N} z_i(e) = b(e) \text{ for all } e \in E. \quad (\text{Walrasian})$$

So there is one allocation that is both packing and covering. We show in Appendix B.1 that in auctions with strong gross substitutes valuations, a price vector is Walrasian if and only if it is both packing and covering, i.e., that there is an allocation with both properties and not just different allocations, one packing and one covering.

Walrasian prices may not always exist (see Example A.1 in Appendix A.1). However, Walrasian prices are guaranteed to exist if all buyers' valuation functions satisfy the *strong gross substitutes* property [42] (see Definition 2.2). Intuitively, gross substitutes expresses the property that the demand of an item does not decrease if only the prices on other items go up. For strong gross substitutes valuations, the Walrasian price vectors form a lattice w.r.t. the component-wise ordering, see Gul and Stacchetti [37] and Ausubel [6]. Consequently, there exists a unique component-wise minimal and a unique component-wise maximal Walrasian price vector. We also call the component-wise minimal Walrasian prices *buyer-optimal Walrasian prices* and the component-wise maximal Walrasian prices *seller-optimal Walrasian prices*.

In this article, we focus on computational aspects of the ascending auctions described by Gul and Stacchetti [38] (for unit-supply) and by Ausubel [6] (for multi-supply) for computing the minimal [maximal] Walrasian price vector in a market with (strong) gross substitutes valuations. In both cases, the maximally overdemanded [underdemanded] sets are the minimizers of a submodular function. Thus, the minimal maximally overdemanded [underdemanded] set can be found by general submodular function minimization algorithms.

An exciting connection between market equilibrium computation and discrete optimization was established by Fujishige and Yang [30], who pointed out the equivalence of the gross substitutes property and M^\natural -concavity, a key concept in the field of *discrete convex analysis*. This theory combines convex analysis and matroid theory, summarized in the monograph by Murota [50]. Thus, powerful tools from discrete convex analysis became available to obtain more efficient equilibrium computations. In particular, Murota et al. [51] obtained faster algorithms for the minimal maximally overdemanded [underdemanded] set computation by exploiting the special structure of the submodular function and taking advantage of demand oracle queries; this will be discussed in Section 1.3.

1.2 Our Contributions

Our contributions are twofold: we provide faster auction algorithms and we present structural results on price monotonicity.

In the same vein as Murota et al. [51], we exploit the special structure of the *Lyapunov function* to obtain even faster running times. We view the submodular minimization problem from a dual viewpoint as a *matroid union problem* (for the unit-supply case) or a *polymatroid sum problem* (for the multi-supply case), and use a fast (and also simple) push-relabel algorithm for this problem. Let us denote by DO the time needed for each *demand oracle query* and by ExO the time of each *exchange oracle query* (see Section 2.4 for definitions and discussion). Then, our algorithm runs in time $O(n \cdot \text{DO} + nm^3 \cdot \text{ExO})$ in the multi-supply case. In comparison, the algorithm in [51] runs in time $O(n \cdot \text{DO} + nm^4 \log(nmB) \cdot \text{ExO})$. For the unit-supply case, the problem becomes the classical *matroid union problem*, and we obtain an even better running time of $O(n \cdot \text{DO} + (m^3 + nm^2) \cdot \text{ExO})$.

We note that our algorithm is significantly simpler than the one in [51]. We also show that given an optimal solution to the polymatroid sum problem, the minimal maximally overdemanded [underdemanded] set can be found by a simple breadth-first search in an exchange graph.

Our push-relabel algorithm is a special implementation of the more general submodular flow feasibility algorithm by Frank and Miklós [28]. In this context, our contribution is giving an efficient implementation in terms of the number of oracle calls. The description in [28] is generic and gives bounds in terms of *basic operations*. Implementing a single such operation may take $O(nm \cdot \text{ExO})$; however, we show that this can be amortized over a sequence of basic operations. See Section 1.3 on further discussion of push-relabel algorithms.

Our second main contribution is on price monotonicity. We make a clear conceptual distinction between packing, covering and Walrasian prices. For example, it is not clear *a priori* whether there may exist a packing price vector q with $q(e) < p_*(e)$ for some items $e \in E$, where p_* denotes the buyer-optimal Walrasian price vector, or whether there may exist a covering price vector q' with $q'(e) > p^*(e)$ for some items $e \in E$, where p^* is the maximal Walrasian price vector. In Theorem 5.3, we show that in fact $p_* \leq q$ for all packing price vectors q and in Theorem 5.6, we show that $p^* \geq q'$ for all covering price vectors q' .

Building on this result, we can also prove in Section 6 that the minimal [maximal] Walrasian prices react naturally to changes in supply and demand, i.e., if total supply of items decreases or the total demand of buyers increases, the minimal [maximal] Walrasian prices can only increase. Independently, Raach [56] proved the same monotonicity results.

1.3 Related Work

In this section, we give a literature review over the topics related to our economic model and equilibrium concept. For an excellent overview over (strong) gross substitutes, Walrasian equilibria, and ascending auctions we refer to the survey by Paes Leme [53].

Gross substitutes and Walrasian equilibria. Kelso and Crawford [42] showed in their seminal article that a Walrasian equilibrium is guaranteed to exist if all valuation functions satisfy the gross substitutes condition that in layman's terms can be stated as, "*the demand for an item does not decrease, if only the prices of other items are increased*".³ Roughly speaking, the study of economic models from an algorithmic and complexity theoretical point of view really started off in the 1990s when also the foundations of algorithmic game theory were laid out. Gul and Stacchetti [37] showed with their Maximum Domain Theorem that indeed gross substitutes are a maximal class of valuation functions (containing unit-demand valuations) that guarantee the existence of Walrasian equilibria. Additionally, they provided equivalent definitions for the gross substitutes condition and showed that Walrasian prices form a complete lattice, which implies that there exist unique component-wise minimal and maximal Walrasian price vectors. There are additional equivalent characterizations of gross substitutes in [9]. However, the characterization by Fujishige and Yang [30] turns out to be most useful from a mathematical point of view. It states that a valuation function has the gross substitutes property if and only if it is M^{h} -concave, which allows for the usage of powerful tools from discrete convexity. This result also has been extended to strong gross substitutes valuation functions in [51]. The differences between strong gross substitutes and "weak gross substitutes" is discussed in [47].

There are other classes than (strong) gross substitutes for which Walrasian prices are guaranteed to exist [9, 62] but these classes do not contain the natural unit-demand valuations. On the other hand, there are subclasses of gross-substitute valuations such as unit-demand valuations [23, 49]

³A formal definition is provided in Section 2.

and additive valuations with demand cap [26] for which the structure of Walrasian prices has been described more explicitly.

Dynamic Auctions or Walrasian tâtonnement. Léon Walras, the namesake of our equilibrium concept, already proposed how equilibria may be found, namely by a “tâtonnement”⁴ process. This procedure basically describes a dynamic auction: an auctioneer posts prices and unless these prices are at equilibrium, the auctioneer makes an adjustment. The first modern study of such a process was done by Demange et al. [23] for unit-demand valuations. In Gul and Stacchetti’s follow-up work [38], they gave the framework for ascending auctions for gross substitutes valuations; that is, start at all-zero prices and increase prices on an inclusion-wise minimal maximally overdemanded set until there is no overdemanded set anymore. They proved that such an auction always terminates with minimal Walrasian prices. However, while they showed that an overdemanded set of items has to exist whenever the prices are not yet at equilibrium, they left it open how to compute those sets. This gap was initially closed by Ausubel [5, 6] using submodular function minimization. There is a discrete version of the Lyapunov function (a potential function, introduced by Varian [65] for divisible goods) which is minimized at Walrasian prices [6]. If all valuation functions are gross substitutes, then this function is submodular and its minimum (and every steepest descent direction, which corresponds to a maximally overdemanded set) can be found efficiently using submodular function minimization, which can be done efficiently (see [35, 39–41]). Murota et al. [51] showed that the Lyapunov function is not just submodular but L^1 -convex, which allows for even faster methods than plain submodular function minimization. Note that these methods by Ausubel [6] and Murota et al. [51] allow for descending auctions or other kinds of dynamic auctions, where prices maybe adjusted in a non-monotone fashion. Similar auctions and guarantees for termination with minimal/maximal Walrasian prices are discussed in [2] and [9].

We should mention that the literature mentioned above follows the definition of a dynamic auction as proposed by Gul and Stacchetti [38]. However, the literature also explores other auction designs, e.g., in [4, 7, 13, 55] that consider slightly different market settings or different ways for buyers to report their demands to the auctioneer.

A weak point of dynamic auctions would be their running time if we assume full information for the auctioneer; e.g., in an ascending auction $\|p_*\|_\infty$ price increase steps are needed, where p_* is the (minimal) Walrasian price vector computed by the auction [51]. This process can be speeded up by increasing prices on an overdemanded set not only by one but by the maximal possible amount before the steepest descent direction of the Lyapunov function changes. This results in at most nmB price adjustment steps (c.f. [59, Proposition 4.17]). Further speed ups of the ascending auction as a rounding scheme are discussed in [53, Section 10.1.]. There are other algorithms which can compute Walrasian prices efficiently [54]. However, in contrast to dynamic auctions, these algorithms need direct information of the valuation functions, i.e., by a value oracle.

Oracles and communication complexity. To compute Walrasian prices, it is crucial to get some information on buyers’ valuation functions. Usually, the necessary communication between auctioneer and buyers is modeled via oracle calls. Depending on which oracle type is used, one can get different information. A comparison of the computational power of some oracle types can be found, e.g., in [15]. In auction literature, typically, value oracles and demand oracles are used. For our dynamic auctions, we require a demand oracle and an exchange oracle as those work naturally in our proposed algorithm to find the required overdemanded or underdemanded sets. The exchange oracle is related to the dynamic rank oracle [14], as both provide information on sets

⁴French: trial-and-error.

close to already requested ones. A detailed overview which oracles we use and how they can be compared is given in Section 2.4.

VCG prices. As a corner stone of auction theory and mechanism design, we have the VCG Theorem [18, 36, 66] that states that there is a sealed-bid mechanism that has truthful bidding as a dominant strategy. However, as briefly discussed in the introduction, ascending auctions are often preferred as they yield a more natural and transparent price-finding process. As a consequence, mathematical economists went to great lengths to implement the VCG mechanism in an ascending auction [4, 6, 11, 12, 16, 22, 48, 63]. Kern et al. [43] showed that for unit-demand valuations, the auction by Demange et al. [23] mentioned above indeed returns VCG prices. If we go beyond unit-demand valuations, VCG prices cannot be achieved using a single price trajectory in a dynamic auction, as shown by Gul and Stacchetti [38] since the auction does not gain enough information during the iterations to determine VCG prices. However, there exist variants of dynamic auctions that circumvent this and can obtain VCG prices [6, 16, 63], e.g., by the use of proxy auctions, which compute multiple price trajectories.

In a multi-supply setting, [26] show that VCG prices that are Walrasian and in which every copy of an item has the same cost, do not necessarily exist.

Discrete convexity: matroid union and polymatroid sum. We refer the interested reader to the monographs by Murota [50] and Fujishige [29]. A survey on the connections between discrete convexity and auction theory is [60].

As pointed out above, all buyers having (strong) gross substitutes valuation functions means that we can use methods from discrete convex optimization to solve the computational problem that arises during a dynamic auction. In particular, the whole auction can be viewed as an iterative process to minimize a submodular (and even L^h -convex) function [6, 51, 52, 59]. As we will see in the sequel, these methods can be quite complicated. For the unit-supply case and gross substitutes valuations, the problem of finding a maximally overdemanded set becomes equivalent to the matroid union problem on n matroids, which is efficiently solvable as shown in [25]. Over time, more efficient algorithms were developed [17, 21, 34, 44] and most recently by Terao [64] and Blikstad et al. [14]. Note that some of these methods are hard to compare due to the use of different oracles.

In the more general setting with multi-supply and strong gross substitutes valuations, we need to consider the more general polymatroid sum problem.

Push-relabel algorithms. The algorithms mentioned above to solve the matroid union and the polymatroid sum problem are so called *augmenting path algorithms*. They start by taking a trivially feasible solution (e.g., the empty set or the all-zero vector, respectively) and augment it by finding additional elements that can be added to the current solution or $(k + 1)$ -to- k -swaps to increase the size of the solution until it is optimal. However, more recently, so called *push-relabel algorithms* have been studied for various combinatorial optimization problems and they can outperform augmenting path algorithms both in theory as well as in practice. The push-relabel paradigm was introduced for the maximum flow problem by Goldberg and Tarjan [33]. The idea is in some sense dual to augmenting path algorithms: One starts with a possibly infeasible solution which otherwise satisfies some optimality condition. The push and relabel operations then maintain this optimality condition while making the solution “less infeasible” over time. We use a push-relabel algorithm by Frank and Miklós [28] for matroid union. Our algorithm in Section 4 gives an efficient implementation in the auction setting and also generalizes to polymatroid sum and M -convex sets in general. It is interesting to note that the level functions and local exchanges of push-relabel algorithms have some resemblance to auction algorithms, as already pointed out by Bertsekas [10].

In our push-relabel algorithm in Section 4, one could interpret the level $\Theta(e)$ of an item as a small marginal price discount used for tie-breaking.

1.4 Overview

In Section 2, we start by giving some basic definitions (Section 2.1) and the well-known dynamic auctions that iteratively increases [decrease] prices on items that are overdemanded [underdemanded] (Section 2.2). This is followed by a summary of known definitions and facts from discrete convex analysis and (poly-)matroids in Section 2.3. Finally, we discuss oracle models in Section 2.4.

In Section 3, we describe how to find minimal maximally overdemanded/underdemanded sets, given an optimal solution to the corresponding polymatroid sum problem. Afterward, in Section 4, we present the push-relabel framework and show how to implement it in an efficient way for both the matroid union and the polymatroid sum problem. Combining these results enables us to compute the minimal maximally overdemanded [underdemanded] sets fast, while using only oracle calls that are natural in the auction setting. In Section 5, we show that there exists a component-wise minimal packing price vector and that it coincides with the minimal Walrasian price vector. Further, we show that there also exists a component-wise maximal covering price vector, which is equal to the maximal Walrasian price vector. In Section 6, we use these facts to show that Walrasian prices fulfill natural monotonicity properties, i.e., when the supply decreases or the demand increases, Walrasian prices can only increase.

2 Preliminaries

In this section, we give some basic definitions and known facts on the connection between strong gross substitutes valuations, polymatroids and discrete convex optimization. We introduce the ascending auction presented by Gul and Stacchetti [38] for the unit-supply setting, which was extended by Ausubel [6] for the multi-supply setting.

Notation. For a vector $x \in \mathbb{R}^E$, we let $\text{supp}^+(x) := \{e \in E \mid x(e) > 0\}$ and $\text{supp}^-(z) := \{e \in E \mid x(e) < 0\}$. For $S \subseteq E$, we let $x(S) := \sum_{e \in S} x(e)$. For $x, y \in \mathbb{R}^E$, we define the meet $x \wedge y \in \mathbb{R}^E$ as the component-wise minimum, and the join $x \vee y \in \mathbb{R}^E$ as the component-wise maximum of x and y . Furthermore, for $e \in E$, we write $\chi_e \in \mathbb{R}^E$ to denote the incidence vector of e , i.e., the entry at position e is 1 and all other entries are 0. Similarly, we define $\chi_S := \sum_{e \in S} \chi_e$ to be the incidence vector of a set $S \subseteq E$.

2.1 Strong Gross Substitutes Valuations

Recall from the introduction that $\mathcal{D}_i(p)$ is the set of preferred bundles of buyer $i \in N$ with respect to prices p . Kelso and Crawford [42] introduced the concept of gross substitutes valuations which guarantee the existence of Walrasian prices. Intuitively, gross substitutes expresses the property that the demand for an item will not decrease if only prices of other items are raised.

Definition 2.1. A valuation function $v_i: 2^E \rightarrow \mathbb{Z}$ is *gross substitutes* if for all price vectors $p, q \in \mathbb{R}^E$ with $p \leq q$, it holds that for all $x_i \in \mathcal{D}_i(p)$, there exists $y_i \in \mathcal{D}_i(q)$ such that $x_i(e) \leq y_i(e)$ for all $e \in E$ with $p(e) = q(e)$.

Note that this definition of gross substitutes (see [42]) only considers the unit-supply case (i.e., where $b(e) = 1$ for all $e \in E$). It can be generalized to the multi-supply setting by treating each copy of an item type as an individual item [6]. The following direct definition (see [51]) is a bit more involved. In particular, it includes the additional requirement that higher prices cannot lead to larger preferred bundles.

Definition 2.2. A valuation function $v_i: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ is *strong gross substitutes* if it is concave-extendible⁵ and for all price vectors $p, q \in \mathbb{R}^E$ with $p \leq q$, it holds that for all $x_i \in \mathcal{D}_i(p)$, there exists $y_i \in \mathcal{D}_i(q)$ such that $x_i(e) \leq y_i(e)$ for all $e \in E$ with $p(e) = q(e)$ and such that $x_i(E) \geq y_i(E)$.

Strong gross substitutes valuations form a natural and well-studied class of valuation functions. Such functions are always submodular, i.e., have the diminishing marginal returns property (see Definition 2.4 below). Simple examples of strong gross substitutes valuations include unit-demand valuations (i.e., $v(z) = \max_{e \in \text{supp}^+(z)} w(e)$ for $w \in \mathbb{R}_+^m$), weighted matroid and polymatroid rank functions (i.e., $v(z) = \max\{\langle w, y \rangle : y \leq z, y(S) \leq \rho(S), S \subseteq E\}$ for a submodular set function ρ and weights $w \in \mathbb{R}_+^m$), and OXS valuations (i.e., the maximum weight of a matching in a bipartite graph, restricted to a subset of the items).

It is well-known that strong gross substitutes are a maximal class of functions (containing the unit-demand valuations) that guarantee existence of Walrasian equilibria [37, Theorem 2].

In the sequel, we assume that all valuation functions are strong gross substitutes and non-decreasing. In this case, it is known that the collection of minimal or maximal preferred bundles forms an M-convex set (see Section 2.3). This allows us to use tools from discrete optimization and integral polymatroids.

2.2 Dynamic Auctions

In this section, we present multiple dynamic auctions that can be used to compute Walrasian prices, if they exist. All of them crucially rely on the observation that a price vector p is Walrasian (if such a price vector exists) if and only if it constitutes a minimizer of the Lyapunov function [6]

$$L(p) = \sum_{i \in N} V_i(p) + \langle p, b \rangle.$$

Kelso and Crawford [42] show that Walrasian prices exist if the valuation functions of the buyers are gross substitutes. Ausubel [6] observed in addition that if Walrasian prices exist, they are minimizers of the Lyapunov function. These observations of Kelso and Crawford [42] and Ausubel [6] are summarized in the following lemma:

LEMMA 2.3 ([6]). *If all buyers have strong gross substitutes valuation functions, then prices p are Walrasian if and only if they minimize the Lyapunov function.*

In light of Lemma 2.3, a natural approach to compute a Walrasian price vector is to start with an arbitrary price vector p and to iteratively replace p by $p \pm \chi_S$ for some $S \subseteq E$ if this decreases the value of the Lyapunov function.

It turns out that the sets $S \subseteq E$ minimizing $L(p + \chi_S)$ [$L(p - \chi_S)$] are precisely the maximally overdemanded [underdemanded] sets. For the unit-supply case, this is shown in [8]. With a standard copy trick, we can obtain the same result for the multi-supply case, see Appendix B.4.

Ausubel [6] shows that an auction, starting with an arbitrary price vector p and then alternating increasing and decreasing price adjustments, will return a Walrasian price vector. Further, Sun and Yang [62] show that by first iteratively increasing prices on a minimal maximally overdemanded

⁵A function $v_i: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ is *concave-extendible* if there exists a concave function $\tilde{v}_i: \{x \in \mathbb{R}^n \mid 0 \leq x \leq b\} \rightarrow \mathbb{R}$ such that $v_i(x) = \tilde{v}_i(x)$ for every $x \in [0, b]_{\mathbb{Z}}$.

set, and then iteratively decreasing prices on a minimal maximally underdemanded set, the auction terminates with a Walrasian price vector as well. This is what we call a [Two-Phase Auction](#).

Two-Phase Auction

Start with an arbitrary price vector p

while *an overdemanded set exists* **do**

┌ choose minimal maximally overdemanded set $S \subseteq E$
└ $p(e) := p(e) + 1$ for all $e \in S$

while *an underdemanded set exists* **do**

┌ choose minimal maximally underdemanded set $S \subseteq E$
└ $p(e) := p(e) - 1$ for all $e \in S$

return p

In the unit-supply setting, the [Two-Phase Auction](#) is also known as the Vickrey-English-Dutch auction [3]. However, note that in our general multi-supply setting, no such auction can be guaranteed to return Vickrey prices: In [26], it is shown that Vickrey prices might not exist if just one price per object type is allowed.

Murota et al. [51] shows that one can reduce the number of iterations by interleaving both phases and always performing an increasing or decreasing step depending on which of the two provides the greater reduction of the Lyapunov function. This results in the following [Greedy Auction](#).

Greedy Auction

Start with an arbitrary price vector p

while *an overdemanded or underdemanded set exists* **do**

┌ choose minimal maximally overdemanded set $S_o \subseteq E$
┌ choose minimal maximally underdemanded set $S_u \subseteq E$
└ **if** $\text{overd}(S_o) \geq \text{underd}(S_u)$ **then**
└ | $p(e) := p(e) + 1$ for all $e \in S_o$
└ **else**
└ | $p(e) := p(e) - 1$ for all $e \in S_u$

return p

As the Lyapunov function is submodular and the minimizers of any submodular function form a lattice, by Lemma 2.3, this is also true for the Walrasian price vectors. In particular, there exists a *unique* buyer-optimal, i.e., minimal [seller-optimal, i.e., maximal] Walrasian price vector p_* [p^*]. Ausubel [6] has shown that p_* [p^*] can be obtained via the following [Ascending Auction](#) [[Descending Auction](#)].

Ascending Auction

$p(e) := 0$ for all $e \in E$

while *an overdemanded set exists* **do**

┌ choose minimal maximally overdemanded set $S \subseteq E$
└ $p(e) := p(e) + 1$ for all $e \in S$

return $p_* := p$

Descending Auction

$p(e) := \max_{i \in N} v_i(E) + 1$ for all $e \in E$

while *an underdemanded set exists* **do**

┌ choose minimal maximally underdemanded set $S \subseteq E$
└ $p(e) := p(e) - 1$ for all $e \in S$

return $p^* := p$

Note that both the [Ascending Auction](#) and the [Descending Auction](#) constitute special cases of both the [Two-Phase Auction](#) and the [Greedy Auction](#): When starting with $p(e) := 0$ for all

$e \in E$ [$p(e) := \max_{i \in N} v_i(E) + 1$ for all $e \in E$], throughout the algorithm, there will never be an underdemanded [overdemanded] set (see, e.g., [8]).

We further point out that for the [Ascending Auction](#) [[Descending Auction](#)] to return the buyer-optimal [seller-optimal] Walrasian price vector, it suffices to start with a vector $\leq p_*$ or $\geq p^*$, respectively (see [4]). The running time of the above auctions depends crucially on the time required to find a minimal maximally over- or underdemanded set in each iteration. In this article, we show how to obtain faster algorithms to compute minimal maximally over- and underdemanded sets by leveraging an algorithm for the polymatroid sum problem.

2.3 Discrete Convex Optimization and Polymatroids

In this section, we summarize some basics from discrete convex optimization and their connection to polymatroids. In the following sections, we show that in markets with strong gross substitutes valuations, the set of minimal preferred bundles $\check{\mathcal{D}}_i(p)$, and also the set of maximal preferred bundles $\hat{\mathcal{D}}_i(p)$, form the sets of bases of integral polymatroids. As we will see in Section 3.1, we can use the properties stated here to compute the sets on which we increase or decrease the prices in the auction using a polymatroid sum algorithm.

Definition 2.4. A set function $f: 2^E \rightarrow \mathbb{Z}$ is *submodular* if

$$f(S) + f(T) \geq f(S \cup T) + f(S \cap T) \quad \text{for all } S, T \subseteq E.$$

A function $v: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ is (*lattice*) *submodular* if

$$v(x) + v(y) \geq v(x \vee y) + v(x \wedge y) \quad \text{for all } x, y \in [0, b]_{\mathbb{Z}}.$$

A special class of submodular functions are M^{h} -concave functions, see Murota [50].

Definition 2.5. A valuation function $v: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}$ is M^{h} -concave, if for any $x, y \in [0, b]_{\mathbb{Z}}$, and for any $e \in \text{supp}^+(x - y)$, one of the following holds:

(M1) $v(x) + v(y) \leq v(x - \chi_e) + v(y + \chi_e)$, or

(M2) there exists $f \in \text{supp}^-(x - y)$ such that $v(x) + v(y) \leq v(x - \chi_e + \chi_f) + v(y + \chi_e - \chi_f)$.

LEMMA 2.6 (MUROTA ET AL. [51]). *A valuation function is strong gross substitutes if and only if it is M^{h} -concave.*

The utility function u_i of a buyer i is defined as the valuation minus the price, i.e., $u_i(z) := v_i(z) - \langle p, z \rangle$ for some $z \in [0, b]_{\mathbb{Z}}$ and some given price vector p . Note that the above lemma also implies that a buyer's utility function u_i is M^{h} -concave if her valuation function v_i is strong gross substitutes. This gives us some nice structure of the minimal and maximal preferred bundles.

Definition 2.7. A set of bundles $\mathcal{B} \subseteq [0, b]_{\mathbb{Z}}$ is M -convex if for any $x, y \in \mathcal{B}$, and for any $e \in \text{supp}^+(x - y)$, there exists $f \in \text{supp}^-(x - y)$ such that $x - \chi_e + \chi_f \in \mathcal{B}$ and $y + \chi_e - \chi_f \in \mathcal{B}$.

LEMMA 2.8. *If v_i is strong gross substitutes, then $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$ are M -convex sets for any price vector $p \in \mathbb{R}_+^E$.*

The simple proof of Lemma 2.8 is given in Appendix B.2.

As $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$ are M -convex, it follows (see Appendix B.2) that

$$\check{\mathcal{D}}_i(p) = \arg \min\{\|z\|_1 \mid z \in \mathcal{D}_i(p)\} \quad \text{and} \quad \hat{\mathcal{D}}_i(p) = \arg \max\{\|z\|_1 \mid z \in \mathcal{D}_i(p)\}.$$

An M -convex set forms the set of integer points in a corresponding polymatroid base polytope as the next lemma shows. This leads to a new interpretation of $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$.

LEMMA 2.9 ([50, THEOREM 4.15]). For an M-convex set $\mathcal{B} \subseteq [0, b]_{\mathbb{Z}}$, we define

$$\rho(S) := \max\{z(S) \mid z \in \mathcal{B}\} \quad \text{for } S \subseteq E. \quad (1)$$

Then, $\rho(S)$ is an integer valued submodular set function, and \mathcal{B} is the set of integer points in the corresponding base polytope, that is

$$\mathcal{B} = \{z \in \mathbb{Z}_+^m \mid z(S) \leq \rho(S) \text{ for all } S \subseteq E, z(E) = \rho(E)\}. \quad (2)$$

Summarizing, we observed that if v_i is strong gross substitutes, then for every price vector p , buyer i 's set of minimal preferred bundles $\check{\mathcal{D}}_i(p)$ as well as the set of maximal preferred bundles $\hat{\mathcal{D}}_i(p)$ is an M-convex set. Moreover, $\check{\mathcal{D}}_i(p)$ is the set of integer points in a polymatroid base polytope and the same holds for $\hat{\mathcal{D}}_i(p)$.

Definition 2.10. For an M-convex set $\mathcal{B} \subseteq [0, b]_{\mathbb{Z}}$ with associated rank function ρ defined as in (1), and a vector $z \in \mathcal{B}$, we say that set $S \subseteq E$ is *tight* if $z(S) = \rho(S)$.

A set of item types S is tight with respect to a minimal [maximal] preferred bundle z of (individual) items if among all minimal [maximal] preferred bundles, the bundle z contains as many items as possible from S .

Note that the collection of tight sets $\mathcal{T}_{\mathcal{B}}(z) := \{S \subseteq E \mid z(S) = \rho(S)\}$ with respect to $z \in \mathcal{B}$ is closed with respect to intersection and union. To see this, observe that for any two tight sets $S, T \in \mathcal{T}_{\mathcal{B}}(z)$, we have $z(S) + z(T) = \rho(S) + \rho(T) \geq \rho(S \cup T) + \rho(S \cap T) \geq z(S \cup T) + z(S \cap T) = z(S) + z(T)$ by submodularity of ρ and since $z \in \mathcal{B}$. Hence, we get $\rho(S \cup T) = z(S \cup T)$ and $\rho(S \cap T) = z(S \cap T)$. As a consequence, there exists a *unique* minimal tight set among those tight sets containing e , which we call $\mathcal{T}_{\mathcal{B}}(e, z)$.

Definition 2.11. Let \mathcal{B} be an M-convex set and $z \in \mathcal{B}$. We denote the unique minimal tight set containing e by $\mathcal{T}_{\mathcal{B}}(e, z)$, i.e.,

$$\mathcal{T}_{\mathcal{B}}(e, z) := \bigcap \{S \in \mathcal{T}_{\mathcal{B}}(z) \mid e \in S\}.$$

We will show later in Lemma 2.14 that the set $\mathcal{T}_{\check{\mathcal{D}}_i}(e, z_i)$ contains exactly those items $f \in E$ which can be exchanged against e , i.e., those f for which $z_i + \chi_e - \chi_f$ remains a minimal preferred bundle for buyer i . Analogously, $\mathcal{T}_{\hat{\mathcal{D}}_i}(e, z_i)$ contains those items $f \in E$ which can be exchanged against e such that $z_i + \chi_e - \chi_f$ remains a maximal preferred bundle.

Let $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n \subseteq [0, b]_{\mathbb{Z}}$ be M-convex sets, and let $\mathcal{B}^{\times} := \times_{i \in N} \mathcal{B}_i$ be the collection of all n -tuples $\mathbf{z} = (z_1, \dots, z_n)$ with $z_i \in \mathcal{B}_i$. The following Min-Max Theorem holds for general polymatroids and implies that prices are packing if and only if there is no overdemanded set (see Lemma 3.1) and that they are covering if and only if there is no underdemanded set (Lemma 3.5). The Min-Max Theorem follows, e.g., from [58, (44.8), (44.9), Theorem 44.6]; the sum of polymatroids was first studied by McDiarmid [46].

THEOREM 2.12 (MIN-MAX THEOREM FOR POLYMATROID SUM). Let $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n \subseteq [0, b]_{\mathbb{Z}}$ be M-convex sets with associated rank functions ρ_i , $i \in N$ and let $\mathcal{B}^{\times} := \times_{i \in N} \mathcal{B}_i$. Then

$$\max_{\mathbf{z} \in \mathcal{B}^{\times}} \left\{ \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} \right\} = \min_{S \subseteq E} \left\{ \sum_{i \in N} \rho_i(E \setminus S) + b(S) \right\}. \quad (3)$$

The problem of finding an optimal solution $\mathbf{z} \in \mathcal{B}^{\times}$ to the left-hand side

$$\max_{\mathbf{z} \in \mathcal{B}^{\times}} \left\{ \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} \right\}, \quad (4)$$

is also known as the polymatroid sum problem. The push-relabel algorithm to solve this problem is given in Section 4 and provides a direct proof of the Min-Max Theorem for polymatroid sum.

The easy direction $\max \leq \min$ is shown in Lemma 2.13 below. Moreover, the lemma also formulates the complementary slackness conditions that can be used to certify optimality of a pair of solutions $\mathbf{z} \in \mathcal{B}^\times$ and $S \subseteq E$.

LEMMA 2.13. *Let $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n \subseteq [0, b]_{\mathbb{Z}}$ be M-convex sets with associated rank functions ρ_i , $i \in N$. Then, for any $\mathbf{z} = (z_1, \dots, z_n) \in \mathcal{B}^\times := \times_{i \in N} \mathcal{B}_i$, and any set $S \subseteq E$, we have*

$$\sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} \leq \sum_{i \in N} \rho_i(E \setminus S) + b(S). \quad (5)$$

Moreover, equality holds if and only if

$$\sum_{i \in N} z_i(e) \geq b(e) \quad \text{for all } e \in S, \quad (6)$$

$$\sum_{i \in N} z_i(e) \leq b(e) \quad \text{for all } e \in E \setminus S, \quad (7)$$

$$z_i(E \setminus S) = \rho_i(E \setminus S) \quad \text{for all } i \in N. \quad (8)$$

PROOF. Let $S \subseteq E$ be an arbitrary set. We have

$$\sum_{i \in N} \rho_i(E \setminus S) + b(S) \geq \sum_{i \in N} z_i(E \setminus S) + b(S) = \sum_{e \in E \setminus S} \sum_{i \in N} z_i(e) + \sum_{e \in S} b(e) \geq \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\}.$$

The first inequality holds with equality if and only if (8) holds. The second inequality holds with equality if and only if both (6) and (7) hold. \square

2.4 Oracle Models

Auction style algorithms (in contrast to direct methods as described in [54]) have to be evaluated differently when we discuss computational efficiency. The overall running time is heavily dependent on the valuation functions and hence, only pseudo-polynomial in the number of items and buyers in the worst case. However, determining an overdemanded set to increase prices on (or more generally, any single price-adjustment step) is a computational problem which we should analyze under two different aspects: (1) how much information does the auctioneer require from the buyers to perform a price update, and (2) given this information, how fast can this update be computed? In summary, we should analyze the time to perform just a single step of the auction and the communication cost incurred in such a step instead of the total running time of the auction. This seems reasonable as the total running time heavily depends on the number of steps, i.e., on something the auctioneer cannot influence with the limited information that is available in a dynamic auction.

Most types of dynamic auctions involve an auctioneer who communicates item prices to the buyers, who are then asked for a bundle z of items, which maximizes their utilities $v_i(z) - \langle p, z \rangle$. In [15], these so called *demand oracles* are compared against other natural types of oracles that partially reveal the private information of the buyers. It should be mentioned that even for (strong) gross substitutes valuations, the size of the set of preferred bundles can be exponential in $|E|$ and there is no more compact way to encode these sets as shown by Knuth [45].

In this section, we describe the oracles we use in our algorithm and we compare them to the oracles used in the literature. As our oracle models only rely on the polymatroid properties of the minimal, respectively maximal, preferred bundles at given prices p , we will use \mathcal{B}_i as a placeholder for $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$.

We first define some weights, that play a key role in our algorithm. Given $\mathbf{z} = (z_1, \dots, z_n)$ where $z_i \in \mathcal{B}_i$ for $i \in N$, let us define the following weights for each pair $\{e, f\} \subseteq E$:

$$w_i(e, f) := \max\{\alpha \in \mathbb{Z} \mid z_i - \alpha\chi_f + \alpha\chi_e \in \mathcal{B}_i\}.$$

Their importance is due to the following connection to i -tight sets:

LEMMA 2.14. *For an M-convex set \mathcal{B}_i , $i \in N$ with associated rank function ρ_i , $z_i \in \mathcal{B}_i$, and $e, f \in E$, we have*

$$w_i(e, f) = \min\{\rho_i(S) - z_i(S) \mid e \in S, f \notin S\}. \quad (9)$$

Thus, $w_i(e, f) > 0$ if and only if $f \in \mathcal{T}_{\mathcal{B}_i}(e, z_i)$. Further, a set S is i -tight if and only if $\mathcal{T}_{\mathcal{B}_i}(e, z_i) \subseteq S$ for all $e \in S$.

PROOF. Let $\alpha := w_i(e, f)$ and let β denote the minimum value on the right-hand side of (9), and S a minimizer. By definition, $z'_i = z_i - \alpha\chi_f + \alpha\chi_e \in \mathcal{B}_i$. We first show $\alpha \leq \beta$. This follows since $z'_i(S) \leq \rho_i(S)$, and $z'_i(S) = z_i(S) + \alpha$, and therefore $\beta = \rho_i(S) - z_i(S) \geq z'_i(S) - z_i(S) = z'_i(S) - (z'_i(S) - \alpha) = \alpha$.

Let us now turn to $\alpha \geq \beta$. For a contradiction, assume that $z''_i = z_i - \beta\chi_f + \beta\chi_e \notin \mathcal{B}_i$, that is, $z''_i(T) > \rho_i(T)$ for some $T \subseteq E$. Thus, $z''_i(T) > z_i(T)$; by definition, this means that $e \in T$, $f \notin T$. Hence, T is in the scope of the minimization problem in (9), giving $\rho_i(T) - z_i(T) \geq \beta$. But this means that $z''_i(T) = z_i(T) + \beta \leq \rho_i(T)$, leading to a contradiction. This completes the proof of (9).

From here, the second statement is immediate. For the final statement, since $\mathcal{T}_{\mathcal{B}_i}(e, z_i)$ is the inclusion-wise minimal tight set containing e , it is clearly contained in every i -tight set $S \subseteq E$ with $e \in S$. Conversely, if $\mathcal{T}_{\mathcal{B}_i}(e, z_i) \subseteq S$ is true for every $e \in S$, then $\bigcup_{e \in S} \mathcal{T}_{\mathcal{B}_i}(e, z_i) = S$, which is i -tight as the union of i -tight set is i -tight by submodularity of the polymatroid rank function ρ_i . \square

In the market context, we can interpret the weights $w_i(e, f)$ as part of the information required by the auctioneer from buyer i . Given the current price vector p , each buyer i is asked to report one minimal/maximal preferred bundle $z_i \in \mathcal{B}_i(p)$. We denote the time of such a ‘demand oracle query’ as DO. Further, we may query values $w_i(e, f)$, which is the maximum number $\alpha \in \mathbb{Z}$ such that buyer i is willing to exchange α units of item f against α units of item e .

Our algorithm needs access to these values $w_i(e, f)$; we denote the time required for computing one such value as ExO. Their computational complexity for the buyers depends on the particular representation of the \mathcal{B}_i sets.

A common model to represent the valuation functions v_i is by a value oracle. Given a value oracle, a minimal/maximal preferred bundle $z_i \in \mathcal{B}_i(p)$ can be computed using the greedy algorithm by m^2 value oracle calls. Given $z_i \in \mathcal{B}_i$ and $\alpha \in \mathbb{Z}$, checking if $z_i - \alpha\chi_f + \alpha\chi_e \in \mathcal{B}_i(p)$ takes a single value oracle call. Thus, one can compute $w_i(e, f)$ in $O(\log(B))$ calls to the value oracle. We can also formulate the computation of $w_i(e, f)$ as the submodular function minimization problem in (9). Hence, one can use a strongly polynomial submodular function minimization algorithm. Nevertheless, this requires access to $\rho_i(S)$, which may lead to running a greedy algorithm for each call. In Appendix A.5, we provide an example that shows that the particular running time of DO and ExO heavily depend on the representation of the valuation function, using the example of OXS functions.

For the unit-supply setting, the sets $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$ are matroid base sets and our task is to solve the matroid union problem. For matroid optimization, the usual oracle settings are via independence or rank oracle queries. Such oracles may be expensive if our primary oracle is for the valuation function v_i : each independence or rank oracle query may require running a greedy algorithm.

To take the actual computation time to answer the oracle in some underlying structure into account, Blikstad et al. [14] propose to use the *dynamic rank oracle*. The idea is that given an independent set S , it is easy to find an independent set close to S , but much harder to find an independent set which is farther away (with respect to the symmetric difference of the sets). To be precise, a dynamic rank oracle starts with the empty set S , then three different operations are denoted as oracle calls (1) adding an element to S , (2) deleting an element from S , and (3) obtaining the rank of S (see also [14, Definition 1.2]). We denote the time of such a ‘dynamic rank oracle query’ by DRO.

Our exchange oracle is efficient from this point of view, since given an independent set, we always query sets with symmetric difference two. We can perform an exchange oracle query by using three dynamic rank oracle calls.⁶ Note that we assume that we do not have to build the first independent set z_i with the dynamic rank oracle but that we can perform the queries starting from the set that we obtained from the demand oracle. A demand oracle query can be answered by $\mathcal{O}(m)$ dynamic oracle queries. The other way around, the dynamic rank oracle is stronger than the exchange oracle, in the sense that the dynamic rank oracle can compute the rank of a given set S in $\mathcal{O}(|S|)$ computations, while the exchange oracle cannot answer this question. Note that in the algorithm for the unit-supply setting, we will use exchange operations that do not change the rank of the sets, so in fact we know the ranks as well. It can therefore be said that the two models behave very similarly in the matroid case and our exchange oracle model can be generalized to polymatroids.

If we analyze our matroid union algorithm in the unit-supply setting with the dynamic rank oracle, we get a running time of $\mathcal{O}((m^3 + nm^2) \cdot \text{DRO})$. Hence, the running time of our matroid union algorithm in the dynamic rank oracle model is not comparable to the running time $\mathcal{O}(\text{poly}(n \log(m))m^{3/2} \cdot \text{DRO})$ by Blikstad et al. [14] as we are linear in n , while Blikstad et al. [14] achieve a better dependency on m .

3 Finding the Overdemanded and Underdemanded Sets

As described earlier, minimal maximally overdemanded and minimal maximally underdemanded sets play an important role in the context of dynamic auctions. In this section, we discuss their connection to the corresponding polymatroids and their rank functions. Remember that $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$ are the base sets of integral polymatroids with associated rank functions $\check{\rho}$ and $\hat{\rho}$.

Moreover, we will see in Sections 3.2 and 3.4 how the minimal maximally over-/underdemanded sets can be computed using an optimal solution to the corresponding polymatroid sum problem.

3.1 Overdemanded Sets

Recall that we defined $\check{\theta}_i^p(S) := \min\{z(S) \mid z \in \check{\mathcal{D}}_i(p)\}$ as the minimum number of items buyer i has to receive from S (in any preferred bundle). By Lemma 2.9, the set $\check{\mathcal{D}}_i(p)$ is implicitly given by $\check{\rho}_i^p$, and also $\check{\theta}_i^p$ can be expressed in terms of $\check{\rho}_i^p$ via $\check{\theta}_i^p(S) = \check{\rho}_i^p(E) - \check{\rho}_i^p(E \setminus S)$. Hence, we can rewrite the overdemandness of a set S w.r.t. prices p as

$$\text{overd}^p(S) := \sum_{i \in N} \check{\theta}_i^p(S) - b(S) = \sum_{i \in N} (\check{\rho}_i^p(E) - \check{\rho}_i^p(E \setminus S)) - b(S).$$

We may omit the superscript p in $\text{overd}^p(S)$ when it is clear from the context.

Note that in the presence of an overdemanded set, the prices cannot be packing. In Lemma 3.1, we show that the converse is also true, namely, a price vector is packing if and only if no overdemanded set exists.

⁶If we cannot perform the queried exchange, one can also argue that we need two dynamic oracle calls more to return to the original set, i.e., five dynamic oracle calls for one exchange oracle query.

LEMMA 3.1. *Given a market instance with strong gross substitutes valuations, prices p are packing if and only if there is no overdemanded set.*

In the unit-supply setting, this lemma was already shown by Gul and Stacchetti [38].

PROOF. Let p be a packing price vector and $\mathbf{z} = (z_1, \dots, z_n)$ the corresponding packing allocation. Hence $z_i \in \check{\mathcal{D}}_i(p)$, i.e., $z_i(E) = \check{\rho}_i(E)$ holds. This yields

$$\sum_{i \in N} \check{\rho}_i(E) = \sum_{i \in N} z_i(E) = \sum_{e \in E} \sum_{i \in N} z_i(e) = \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\}$$

where the last equality holds as \mathbf{z} is packing ($\sum_{i \in N} z_i(e) \leq b(e)$ for all $e \in E$). Hence, Theorem 2.12 allows us to conclude that

$$\begin{aligned} \sum_{i \in N} \check{\rho}_i(E) &= \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} \leq \max_{\mathbf{z}' \in \check{\mathcal{D}}^\times} \sum_{e \in E} \min \left\{ \sum_{i \in N} z'_i(e), b(e) \right\} \\ &= \min_{S \subseteq E} \left\{ \sum_{i \in N} \check{\rho}_i(E \setminus S) + b(S) \right\}. \end{aligned} \quad (10)$$

The outer inequality can be transformed into the following formula, which can be seen as a *generalized Hall formula*:

$$\sum_{i \in N} (\check{\rho}_i(E) - \check{\rho}_i(E \setminus S)) \leq b(S) \quad \text{for all } S \subseteq E. \quad (11)$$

This is again equivalent to the fact that there is no overdemanded set S . To show the reverse direction, assume that there is no overdemanded set at prices p . Thus, we apply Equation (11), which yields

$$\begin{aligned} \sum_{i \in N} \check{\rho}_i(E) &\leq \min_{S \subseteq E} \left\{ \sum_{i \in N} \check{\rho}_i(E \setminus S) + b(S) \right\} = \max_{\mathbf{z} \in \check{\mathcal{D}}^\times} \left\{ \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} \right\} \\ &\leq \max_{\mathbf{z} \in \check{\mathcal{D}}^\times} \left\{ \sum_{e \in E} \sum_{i \in N} z_i(e) \right\} = \max_{\mathbf{z} \in \check{\mathcal{D}}^\times} \left\{ \sum_{i \in N} z_i(E) \right\} = \sum_{i \in N} \check{\rho}_i(E). \end{aligned}$$

Hence, equality must hold throughout the equation. In particular, this means that there exists an allocation $\mathbf{z} \in \check{\mathcal{D}}^\times$ with $\sum_{i \in N} z_i(e) \leq b(e)$ for all $e \in E$, i.e., the allocation \mathbf{z} is packing and the price vector is packing as well. \square

Next, we will consider the connection of maximally overdemanded sets to an optimal solution \mathbf{z} of the polymatroid sum problem

$$\max_{\mathbf{z} \in \check{\mathcal{D}}^\times} \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\}. \quad (12)$$

This is done in several steps. First, we characterize which properties a minimal maximally overdemanded set needs to satisfy and introduce a technical lemma. Then, we combine these two statements to describe a nice way to compute an actual minimal maximally overdemanded set by constructing an auxiliary digraph and choosing all items from which we can reach an oversold item. We say that an item $e \in E$ is *undersold* if $\sum_{i \in N} z_i(e) < b(e)$, and *oversold* if $\sum_{i \in N} z_i(e) > b(e)$. For $i \in N$ and $z_i \in \check{\mathcal{D}}_i$, let us call a set $S \subseteq E$ *i -tight* if $z_i(S) = \check{\rho}_i(S)$. To shorten notation slightly, we write, for $e \in E$,

$$\mathcal{T}_i(e) := \mathcal{T}_{\check{\mathcal{D}}_i}(e, z_i)$$

for the minimal i -tight set containing e when z_i and $\check{\mathcal{D}}_i$ are clear from the context.

LEMMA 3.2. Assume $\mathbf{z} = (z_1, \dots, z_n)$ is an optimal solution to (12). Then, a set S is maximally overdemanded if and only if it fulfills the following properties:

- $E \setminus S$ is i -tight for every $i \in N$;
- S does not include any undersold items;
- S includes all oversold items.

Moreover, there is a unique inclusion-wise minimal set fulfilling these properties.

PROOF. Let $S \subseteq E$. We can write

$$\text{overd}(S) = \sum_{i \in N} \check{\rho}_i(E) - \left(\sum_{i \in N} \check{\rho}_i(E \setminus S) + b(S) \right).$$

Hence, S is a maximally overdemanded set if and only if it is a minimizer of the right-hand side of (3). Since \mathbf{z} is optimal, we get, by Theorem 2.12,

$$\sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} = \sum_{i \in N} \check{\rho}_i(E \setminus S) + b(S).$$

By Lemma 2.13, this is the case if and only if S satisfies all three properties stated in the lemma. Moreover, if S and S' satisfy these three properties, then $S \cap S'$ and $S \cup S'$ also satisfy them. Hence, there exists a unique inclusion-wise minimal set satisfying these three properties. \square

In Section 3.2, we will see that we can compute a minimal maximally overdemanded set by performing a breadth-first search in the exchange graph corresponding to an optimal solution \mathbf{z} of the polymatroid sum problem (12).

3.2 Finding Minimal Maximally Overdemanded Sets via the Exchange Graph

We define the exchange graph $\check{G}_i(z_i) = (E, \check{A}_i)$ for $i \in N$ by setting

$$\check{A}_i := \{(e, f) \in E \times E \mid \check{w}_i(e, f) > 0\},$$

where $\check{w}_i(e, f) := \max\{\alpha \in \mathbb{Z} \mid z_i - \alpha \chi_f + \alpha \chi_e \in \check{\mathcal{D}}_i\}$ is the weight which can be computed by an exchange oracle (ExO) call as described in Section 2.4. We let $\check{G}(\mathbf{z}) = (E, \check{A})$, where $\check{A} := \bigcup_{i \in N} \check{A}_i$.

THEOREM 3.3. Assume \mathbf{z} is an optimal solution to (12). Then the minimal maximally overdemanded set is the set R of items from which we can reach an oversold item in $\check{G}(\mathbf{z})$.

PROOF. Let \mathbf{z} be an optimal solution to (12) and R the set of items from which an oversold item can be reached in $\check{G}(\mathbf{z})$. We will show that R is a maximally overdemanded set and, moreover, included in any other maximally overdemanded set.

To this end, let S be a maximally overdemanded set. Then by Lemma 3.2, S contains all oversold items and no undersold item. Moreover, $E \setminus S$ is tight and thus, by Lemma 2.14, there is no arc leaving $E \setminus S$, i.e., $\delta^-(S) = \emptyset$. Now, consider the set R . A sketch of S and R can be found in Figure 1. By definition, R contains every oversold item. Moreover, as all undersold items are in $E \setminus S$, all oversold items are in S , and $\delta^-(S) = \emptyset$, there is no path from an undersold item to an oversold item. Hence, the undersold items are in $E \setminus R$. Furthermore, by definition of R , there is no arc entering R , i.e., the set $E \setminus R$ is tight. Hence R is maximally overdemanded.

Finally, observe that $R \subseteq S$, as otherwise, there would be an arc entering S which yields a contradiction. This means that the set R is included in every maximally overdemanded set and is thus the inclusion-wise minimal maximally overdemanded set. \square

LEMMA 3.4. Overdemanded sets can be computed in time $O(nm^2 \cdot \text{ExO})$, given an optimal solution \mathbf{z} to the polymatroid sum problem (12).

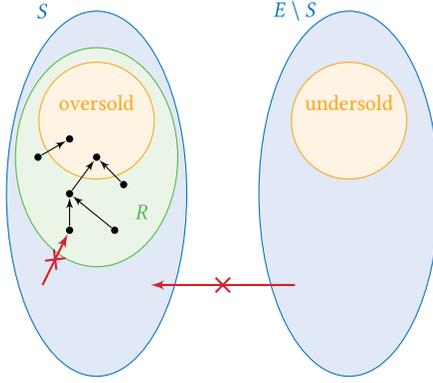


Fig. 1. A maximally overdanded set S and the set of items R from which an oversold item can be reached.

PROOF. First, we construct the exchange graph. We consider each of the m items $e \in E$ once and do an edge-weight query for each of the n buyers $i \in N$ and her currently assigned items. Thus, we have a running time of $O(nm^2 \cdot \text{ExO})$ for this step. Then we do a breadth-first search in the graph, starting from all oversold items in reverse direction of the arcs. This takes time $O(m + m^2)$. Hence, we need $O(nm^2 \text{ExO} + m^2)$ time to compute the set of items from which we can reach any oversold item, which is, by Theorem 3.3, equal to the minimal maximally overdanded set. \square

In Section 4, we will present an algorithm to compute an optimal solution z of the polymatroid sum problem (4). This will lead to a total running time of $O(n \cdot \text{DO} + nm^3 \cdot \text{ExO})$ to compute an overdanded set in the multi-supply setting and to an improved running time of $O(n \cdot \text{DO} + (m^3 + nm^2) \cdot \text{ExO})$ in the unit-supply setting.

3.3 Underdanded Sets

Recall that we formally defined the underdandedness of a set S by

$$\text{underd}^p(S) := b(S) - \sum_{i \in N} \hat{\rho}_i^p(S),$$

where $\hat{\rho}_i^p(S) := \max\{z(S) \mid z \in \hat{\mathcal{D}}_i(p)\}$ is the rank function of the corresponding polymatroid base polytope given by the maximal demand sets $\hat{\mathcal{D}}_i(p)$.

In the following, we will omit the superscript p to improve readability.

LEMMA 3.5. *Given a market instance with strong gross substitutes valuations, prices p are covering if and only if there is no underdanded set.*

The lemma follows, similarly to Lemma 3.1, from the Min-Max Theorem for polymatroid sum (Theorem 2.12). The proof can be found in Appendix B.3. There, we rephrase

$$\text{underd}(S) = b(S) - \sum_{i \in N} \hat{\rho}_i(S) = b(E) - \left(b(E \setminus S) + \sum_{i \in N} \hat{\rho}_i(E \setminus (E \setminus S)) \right).$$

By this, we obtain the following corollary.

COROLLARY 3.6. *$S \subseteq E$ is a maximally underdanded set if and only if*

$$E \setminus S \in \arg \min \left\{ b(T) + \sum_{i \in N} \hat{\rho}_i(E \setminus T) \mid T \subseteq E \right\}.$$

Moreover, S is a minimal maximally underdemanded set if and only if $E \setminus S$ is a maximal minimizer of $T \mapsto b(T) + \sum_{i \in N} \hat{\rho}_i(E \setminus T)$.

In order to determine a minimal maximally underdemanded set S , we follow the same approach as in Section 3.1 using the Min-Max Theorem for polymatroid sum (Theorem 2.12). Given an optimum solution $\mathbf{z} = (z_1, \dots, z_n)$ to the polymatroid sum problem

$$\max_{\mathbf{z} \in \widehat{\mathcal{D}}^\times} \left\{ \sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} \right\}, \quad (13)$$

we explain how to find a minimal maximally underdemanded set S . Recall that for $S \subseteq E$ and $i \in N$, we call S i -tight if $z_i(S) = \hat{\rho}_i^p(S)$. We obtain the following analogue of Lemma 3.2:

LEMMA 3.7. *Assume $\mathbf{z} = (z_1, \dots, z_n)$ is an optimal solution to (13). Then, a set S is maximally underdemanded if and only if it fulfills the following properties:*

- S is i -tight for every $i \in N$;
- S includes all undersold items;
- S does not include any oversold items.

Moreover, there is a unique inclusion-wise minimal set fulfilling this properties.

It can be derived from Lemma 2.13, similar to the proof of Lemma 3.2 (see Appendix B.3).

3.4 Finding Minimal Maximally Underdemanded Sets

Similar to Section 3.2, we can compute the minimal maximally underdemanded set S via reachability in an exchange graph. Given prices $p, z_i \in \widehat{\mathcal{D}}_i(p)$ and $e, f \in E$, we define weights, which can be computed by an exchange oracle as described in Section 2.4, by

$$\widehat{w}_i^p(e, f) := \max\{\alpha \in \mathbb{Z} \mid z_i - \alpha \chi_f + \alpha \chi_e \in \widehat{\mathcal{D}}_i(p)\}.$$

Let further

$$\widehat{A}_i := \{(e, f) \in E \times E \mid \widehat{w}_i^p(e, f) > 0\} \quad \text{for } i \in N,$$

and define $\widehat{G}(\mathbf{z}) := (E, \widehat{A})$ with $\widehat{A} := \bigcup_{i \in N} \widehat{A}_i$.

THEOREM 3.8. *The minimal maximally underdemanded set S is the set R of items which are reachable from an undersold item in $\widehat{G}(\mathbf{z})$.*

This proof follows from Lemmas 2.14 and 3.7. It is similar to the one of Theorem 3.3 and given in Appendix B.3.

In particular, we obtain the following analogue of Lemma 3.4:

LEMMA 3.9. *The minimal maximally underdemanded set can be computed in time $O(nm^2 \cdot \text{ExO})$, given an optimal solution \mathbf{z} to (13).*

In the following section, we present an algorithm to compute an optimal solution \mathbf{z} to the general polymatroid sum problem (4). In particular, this algorithm allows us to compute optimal solutions to (12) or to (13), respectively.

4 Push-Relabel Algorithm to Solve the Polymatroid Sum Problem

Push-relabel algorithms are a well-studied algorithmic paradigm to solve submodular function minimization problems or subclasses of those. We start by presenting a general push-relabel framework for polymatroids and prove its correctness. Afterward, we show how to implement this framework first for matroids and then for polymatroids and analyze the running time. Note that, in

principle, any algorithm could be used to solve the matroid union (polymatroid sum) problem in order to compute the minimal maximally over-/underdemanded set with the help of the exchange graph as described in Section 3.2 above. In this section, we present an efficient implementation of a push-relabel algorithm for matroid union and polymatroid sum with respect to the number of ExO oracle queries. Our push-relabel algorithm is based on Frank and Miklós [28] who present a simpler algorithm for matroid union and a more general one for submodular flow feasibility; our algorithm generalizes the first and is a special case of the second. In contrast to our analysis, Frank and Miklós [28] only give a complexity bound in terms of “basic operations”, i.e., push-relabel steps. We give a self-contained presentation of the algorithm and the analysis.

4.1 The Push-Relabel Framework

Let $\mathcal{B}_1, \dots, \mathcal{B}_n \in [0, b]_{\mathbb{Z}}$ be M-convex sets. Furthermore, we consider the weights $w_i(e, f)$ with respect to an $i \in N$ and a fixed $z_i \in \mathcal{B}_i$ as described in Section 2.4. We define the level function $\Theta: E \rightarrow \{0, 1, \dots, m\}$, and denote $\Theta_{\min}(S) := \min\{\Theta(e) : e \in S\}$ for $S \subseteq E$. Recall that $E = \{1, 2, \dots, m\}$, i.e., items are labeled by integers and the same holds for buyers. Thus, we may use the usual order \leq on each of them. Further recall that we use the shorthand notation $\mathcal{T}_i(e)$ instead of $\mathcal{T}_{\mathcal{B}_i}(e, z_i)$ for the unique minimal i -tight set containing e .

Definition 4.1. Given $z_i \in \mathcal{B}_i$ for all $i \in N$, a mapping $\Theta: E \rightarrow \{0, 1, \dots, m\}$ is a *valid level function* if it satisfies the following properties:

- (L1) $\Theta(e) = 0$ whenever $e \in E$ is oversold,
- (L2) $\Theta_{\min}(\mathcal{T}_i(e)) \geq \Theta(e) - 1$ for all $i \in N$ and $e \in E$.

The general idea of the push-relabel algorithm is as follows. Initially, all $e \in E$ have $\Theta(e) = 0$. As long as there still is an undersold item with label below m , we pick one such item e . Now, we aim at pushing at this item if possible or relabel otherwise, where the push-relabel operations are described as follows:

Push: If we find an item f with $f \in \mathcal{T}_i(e)$ and $\Theta(f) = \Theta(e) - 1$, then we perform a *push* at e with respect to i and f . In such a push, we replace z_i by $z'_i := z_i - \alpha \chi_f + \alpha \chi_e$ for $\alpha := \min\{b(e) - \sum_{\ell \in N} z_\ell(e), w_i(e, f)\}$. That is, we select the largest value of α such that e does not get oversold, and such that $z'_i \in \mathcal{B}_i$.

Relabel: If no such item exists, we *relabel* by setting $\Theta(e) := \Theta(e) + 1$.

If in the push operation, the minimum in $\alpha := \min\{b(e) - \sum_{\ell \in N} z_\ell(e), w_i(e, f)\}$ is attained in the second entry, this is called a *saturating push at e w.r.t. f and i* ; otherwise, if the minimum is attained only in the first entry, it is called a *non-saturating push at e w.r.t. f and i* .

Push-Relabel Algorithm

$z_i :=$ any vector from \mathcal{B}_i for all $i \in N$

$\Theta(e) := 0$ for all $e \in E$

while there is an undersold item e with $\Theta(e) < m$ **do**

choose an undersold item e with $\Theta(e) < m$

Push if there is some f with $f \in \mathcal{T}_i(e)$ and $\Theta(f) = \Theta(e) - 1$

Relabel otherwise

return z

The following two lemmata state that the level invariants are maintained, and that the push-relabel framework indeed computes an optimal solution.

LEMMA 4.2. *The level invariants (L1) and (L2) are maintained throughout the Push-Relabel Algorithm.*

PROOF. The invariants (L1) and (L2) clearly hold at initialization. To see that (L1) is maintained, note that labels are only modified for undersold items, and no undersold item gets oversold during the execution of the algorithm (by the choice of α in every push operation). Note that it may happen that an initially oversold item e with $\Theta(e) = 0$ gets undersold, but then it never gets oversold again.

Let us now turn to (L2). First, consider a relabel operation of some item e . Since no push was performed, it follows that for all $f \in \mathcal{T}_i(e)$, it holds that $\Theta(f) > \Theta(e) - 1$ (smaller is not possible, since (L2) is fulfilled before the relabeling operation). Thus, $\Theta_{\min}(\mathcal{T}_i(e)) \geq \Theta(e)$, and hence, (L2) is still fulfilled after the level of e is increased by one.

Consider now a push operation that changes z_i to $z'_i := z_i - \alpha\chi_f + \alpha\chi_e$ for $\alpha > 0$. Let $\mathcal{T}'_j(g)$ denote the j -tight sets for each $g \in E$ after the push operation. We need to show

$$\Theta_{\min}(\mathcal{T}'_j(g)) \geq \Theta(g) - 1 \quad \text{for all } j \in N, g \in E. \quad (14)$$

This inequality is immediate for all $j \neq i$ and $g \in E$ since the bundle z_j and the levels do not change, therefore $\mathcal{T}'_j(g) = \mathcal{T}_j(g)$. It also remains true for $j = i$, and for all $g \in E$ such that $f \notin \mathcal{T}_i(g)$, since in this case $\mathcal{T}_i(g)$ is also i -tight for z'_i , and thus $\mathcal{T}'_i(g) \subseteq \mathcal{T}_i(g)$.

Hence, it is left to show (14) for the case $j = i$, and $f \in \mathcal{T}_i(g)$. We have $\mathcal{T}'_i(g) \subseteq \mathcal{T}_i(g) \cup \mathcal{T}_i(e)$, since $\mathcal{T}_i(g) \cup \mathcal{T}_i(e)$ is tight before the push operation as the union of two tight sets, and remains tight. Thus,

$$\Theta_{\min}(\mathcal{T}'_i(g)) \geq \min \{ \Theta_{\min}(\mathcal{T}_i(g)), \Theta_{\min}(\mathcal{T}_i(e)) \} = \min \{ \Theta_{\min}(\mathcal{T}_i(g)), \Theta(f) \} = \Theta_{\min}(\mathcal{T}_i(g)),$$

where the penultimate equality follows since $\Theta(f) = \Theta_{\min}(\mathcal{T}_i(e))$ if f is selected for a push (by the selection property and since (L2) is fulfilled before the push). The last equality follows by $f \in \mathcal{T}_i(g)$. This shows that (14) remains true after the push operation. \square

LEMMA 4.3. *The Push-Relabel Algorithm returns an optimal solution to the polymatroid sum problem (4).*

PROOF. First, we observe that when the algorithm stops, there is a value $\ell \in \{0, 1, \dots, m\}$ such that there is no item e with $\Theta(e) = \ell$. To see this, note that we have only m items but $m + 1$ levels, so by pigeonhole principle, there exists an empty level $\ell \leq m$.

By Lemma 2.13, Equation (5) always holds. Using the equality criteria of Lemma 2.13, we show that the equation is fulfilled with equality for our computed bundle z and $S := \{e \in E \mid \Theta(e) < \ell\}$, implying that z is an optimal solution to the polymatroid sum problem (4). We note that $S = \emptyset$ and $S = E$ are both possible.

Since there is no undersold item on any level below m by the stopping criterion, there is also no undersold item on any level below ℓ and the equality criterion (6) holds by definition of S . Optimality criterion (7) is satisfied because all items in $E \setminus S$ have a label greater than 0 while all oversold items have label 0 by (L1). To show (8), note that $\mathcal{T}_i(e) \subseteq E \setminus S$ for all $e \in E \setminus S$ by (L2) and the choice of ℓ . Hence, $E \setminus S = \bigcup_{e \in E \setminus S} \mathcal{T}_i(e)$ is i -tight for each $i \in N$, as it is the union of i -tight sets. \square

We have seen that the Push-Relabel Algorithm returns an optimal solution z to the polymatroid sum problem (4). In the next two sections, we explain how to find the items to push or relabel in a fast way. Therefore, we distinguish the unit-supply setting, i.e., where $b(e) = 1$ for each $e \in E$ and the general multi-supply setting.

4.2 Running Time of the Unit-supply Case

We now consider the unit-supply case, i.e., $b(e) = 1$ for all $e \in E$. In this case, each \mathcal{B}_i is an M -convex subset of $[0, 1]_{\mathbb{Z}}$, which precisely corresponds to the set of bases of a matroid. Hence, $z_i \in \mathcal{B}_i$ is the indicator vector of a base; we will also use the notation $B_i = \text{supp}(z_i)$. In this case,

it is easy to see that $\mathcal{T}_i(e) = \{e\}$ if $e \in B_i$, and $\mathcal{T}_i(e)$ is the fundamental circuit $C(B_i, e)$ in the i -th matroid if $e \in E \setminus B_i$.

For the special case of unit-supply, the [Push-Relabel Algorithm](#) can be described as follows: In each iteration with current bases B_i for $i \in N$, and levels $\Theta: E \rightarrow \{0, \dots, m\}$, the algorithm checks whether there exists an undersold, thus unsold, item $e \in E$ of level $\Theta(e) < m$, and if so, chooses one such item.

In the body of the **while** loop, given an unsold item e , the algorithm checks if it can perform a push operation. For this, the algorithm might need to go through all buyers $i \in N$, and check whether there exists an item f in the fundamental circuit $C(B_i, e)$ of level $\Theta(f) = \Theta(e) - 1$. This may cost nm queries, but only if e has level $\Theta(e) = 1$. If e has level $\Theta(e) = 0$, we can immediately relabel. Whenever e has level $\Theta(e) \geq 2$, then it suffices to go through each item f at most once, since each item f of level $\Theta(f) = \Theta(e) - 1$ that is sold is owned by exactly one buyer (by [L1](#)) and since there is only one unit per item). The proof of the following theorem makes use of the described procedure to bound the running time.

THEOREM 4.4. *The [Push-Relabel Algorithm](#) can be implemented in running time $O(n \cdot \text{DO} + (m^3 + nm^2) \cdot \text{ExO})$ for the unit-supply case.*

PROOF. The algorithm starts with finding an initial allocation z in $n \cdot \text{DO}$ time. Given an allocation z , we can compute how many buyers own an item e in $O(nm)$. These values can be updated in time $O(1)$ at every push operation. Moreover, these values allow us to maintain a list of all unsold items, which needs $O(m)$ time for initialization and can be maintained in $O(1)$ at each push. Note that at a relabel operation, nothing changes for these values and the list.

We find an unsold item e and check its level $\Theta(e)$ in $O(1)$. If $\Theta(e) = 0$, we can relabel immediately. If $\Theta(e) = 1$, we will go through all buyer-item combinations of items in level 0 to determine a push (if possible, else we relabel). This takes $O(mn)$ exchange oracle calls. For every other level $\ell \geq 2$, any item on level $\ell - 1$ belongs to at most one buyer (by [L1](#)), and since there is only one unit per item). Thus, we just need to go through the items once to determine a push or a relabel operation. This can be done in $O(m)$ exchange oracle calls.

The number of relabel operations per level $\ell \geq 1$ is bounded by $O(m)$. The number of pushes per level is also bounded by $O(m)$. To see this, we show that for each level ℓ the value $\Phi(\ell) = \sum_{i \in N} |B_{i, \geq \ell}|$ is non-decreasing in every relabel or push operation, where $|B_{i, \geq \ell}|$ is the number of items in the base of buyer i that were added to the base when their level was at least ℓ . Relabel operations are only used for unsold items, and hence they do not affect the $B_{i, \geq \ell}$ sets; thus, $\Phi(\ell)$ does not change. At any push at e with respect to f and i , $|B_{j, \geq \ell}|$ does not change for $j \neq i$. For buyer i , the item f which leaves the base B_i is exactly one level lower than the level of the item e which enters the base. Hence, $|B_{i, \geq \ell}|$ is not decreasing. When pushing at e with respect to i and f with $\Theta(e) = \ell$, we strictly increase $|B_{i, \geq \ell}|$ and thus, $\Phi(\ell)$.

Since $\Phi(\ell) \leq m$ for each $\ell \geq 1$ (as each item with level $\ell \geq 1$ appears in just one base), we get a bound of m on the number of pushes per level.

This gives a total running time of $O((m^2n + m^3) \cdot \text{ExO})$ after finding the initial allocation, where the first term is generated by the pushes and relabel operations on items at level 1 and the second term captures all remaining pushes. \square

COROLLARY 4.5. *Given prices p , we can compute the minimal maximally overdemanding [underdemanded] set in time $O(n \cdot \text{DO} + (m^3 + nm^2) \cdot \text{ExO})$ in the unit-supply setting.*

This follows by Lemma [3.4](#) [Lemma [3.9](#), respectively] and Theorem [4.4](#). Note that the demand and exchange oracle calls are based on the matroids corresponding to the preferred bundles $\check{\mathcal{D}}_i(p)$ or $\hat{\mathcal{D}}_i(p)$, respectively.

4.3 A Polymatroid Sum Algorithm

In this section, we provide a fast implementation of the push-relabel framework for the multi-supply setting. We will need a different approach here than in the unit-supply setting. There, the key property of the approach is that an item which is not oversold is owned by at most one buyer. This is not true any longer.

Recall that the generic push-relabel algorithm selects an undersold item e and then searches for a buyer-item pair (i, f) satisfying $w_i(e, f) > 0$ and $\Theta(f) = \Theta(e) - 1$. If such a pair exists, a push is performed. If not, item e gets relabeled. In this section, we show that it is possible to go through the buyer-item pairs in a structured way. Namely, for a fixed item e we consider the buyer-item pairs in lexicographically increasing order. This (total) lexicographical ordering is based on the ordering of the buyers $N = \{1, \dots, n\}$ and the items $E = \{1, \dots, m\}$ by natural numbers, i.e., (i, f) is lexicographically smaller than (i', f') if $i < i'$, or if $i = i'$ and $f < f'$.

In Lemma 4.7, we show a crucial monotonicity property: For a fixed item e , the lexicographically minimal buyer item pair (i, f) which fulfills all necessary conditions for a push will lexicographically only increase until e is relabeled. This implies that for a fixed item e , we can find all possible pushes until e is relabeled by going in lexicographically increasing order through all item-buyer combinations.

We note that lexicographic selection rules are commonly used in submodular optimization. This idea traces back to the first polynomial-time algorithm for polymatroid intersection by Schönsleben [57], and was used in the push-relabel algorithms by Fujishige and Zhang [31, 32]. Notably, the algorithm by Frank and Miklós [28] does not rely on this rule. It is beneficial in our implementation as it enables to decrease the number of exchange oracle queries, leading to an implementation of the push-relabel framework in $\mathcal{O}(m^3 n \cdot \text{ExO})$.

4.3.1 Algorithm. We implement the general push-relabel algorithm with two subtleties. First of all, we always select an undersold item e of highest level among those of level strictly smaller than m . Second, for some given item e , when selecting the buyer-item pair (i, f) for the next push, we always select the pair which is lexicographically smallest among all buyer-item pairs that fulfill the necessary requirements for a push at item e .

To implement this rule efficiently, we exploit the monotonicity property stated in Lemma 4.7. This means that for every item e , we keep track of the next buyer-item pair which we have to consider as a candidate for a push in case e gets selected. Starting from this pair, we search for a feasible candidate in lexicographically increasing order. When a candidate is found, a push is performed and the pointer to the next buyer-item pair to consider is updated. For a non-saturating push, this is the buyer-item pair for which the push was performed, for a saturating one it is the lexicographically next buyer-item pair. If no feasible pair for a push can be found, this means that e gets relabeled and the pointer to the next buyer-item pair to consider is reset to the lexicographically minimal one. We give an implementation of the described procedure in pseudo code.

THEOREM 4.6. *The Polymatroid Sum Algorithm via Push-Relabel returns a solution to (4). The algorithm runs in time $\mathcal{O}(n \cdot \text{DO} + nm^3 \cdot \text{ExO})$.*

To prove this theorem, we first show correctness of the algorithm. More precisely, we have to show the correctness of the procedure to select a buyer-item pair for a push. In particular, we need to prove that with the described procedure, we do not relabel although there is still a feasible buyer-item pair for a push when ignoring buyer-item pairs which are lexicographically below the pointer.

4.3.2 Correctness. We start by proving the above mentioned monotonicity lemma.

Polymatroid Sum Algorithm via Push-Relabel

z_i := any vector from \mathcal{B}_i for all $i \in N$

$\Theta(e) := 0$ for all $e \in E$

initialize pointer to the buyer-item pair for the exchange at item e as $(j_e, g_e) := (1, 1)$ for all $e \in E$

while there is an undersold item e with $\Theta(e) < m$ **do**

$push := false$

 Choose $e \in \arg \max_{e \in E} \{\Theta(e) \mid e \text{ is undersold and } \Theta(e) < m\}$

for buyer-item pairs (i, f) in lexicographically increasing order starting with the pair (j_e, g_e) **do**

if $w_i(e, f) > 0$ **and** $\Theta(e) = \Theta(f) + 1$ **then** // buyer-item pair to perform a push

$push := true$

$\alpha := \min\{b(e) - \sum_{\ell \in N} z_\ell(e), w_i(e, f)\}$

$z_i := z_i - \alpha \chi_f + \alpha \chi_e$

if $b(e) - \sum_{\ell \in N} z_\ell(e) < w_i(e, f)$ **then** // non-saturating push

$(j_e, g_e) := (i, f)$ // set pointer to current pair

 continue with next undersold item e (by aborting the for-loop) // e is not undersold anymore

else // saturating push

$(j_e, g_e) := \begin{cases} (i, f + 1) & \text{if } f + 1 \text{ exists,} \\ (i + 1, 1) & \text{otherwise} \end{cases}$ // set pointer to next pair in lex. order

if not $push$ **then** // relabel if no push was found for e

$\Theta(e) := \Theta(e) + 1$

$(j_e, g_e) := (1, 1)$ // reset pointer to first pair in lex. order

return z and $\Theta(e)$ for all $e \in E$

LEMMA 4.7. Let $e \in E$ with $\Theta(e) = \ell$ and let $i \in N$. Under the assumption that every push in *Push-Relabel Algorithm* is executed with respect to a minimal suitable item,

$$\min\{f \mid f \in \mathcal{T}_i(e), \Theta(f) = \ell - 1\}, \quad (15)$$

is monotonically increasing and strictly increases at every saturating push at e with respect to f and i as long as we do not relabel e , i.e., as long as $\Theta(e)$ remains ℓ .

PROOF. We prove the lemma by considering the different steps a push and relabel algorithm can execute, namely pushing and relabeling and show that after both these operations, the statement of the lemma is fulfilled. A relabel operation does not affect $\mathcal{T}_i(e)$. Moreover, by (L2), there is no item $f \in \mathcal{T}_i(e)$ with $\ell - 2 = \Theta(f) < \Theta(e) - 1 = \ell - 1$. Thus, no additional candidates can appear in $\{f \mid f \in \mathcal{T}_i(e), \Theta(f) = \ell - 1\}$. The label of the currently minimal item could increase from $\ell - 1$ to ℓ though, in which case (15) increases. In total, we obtain the desired monotonicity for a relabel operation.

For push operations with respect to $i' \neq i$, neither $\mathcal{T}_i(e)$ nor the labels change and thus, again, the statement holds since the minimal item remains the same.

Consider a push operation at k with respect to i and g . Let $z'_i = z_i - \alpha \chi_g + \alpha \chi_k$ be the set assigned to i after the push and let $\mathcal{T}'_i(e)$ denote the minimal i -tight set containing e after the push.

First assume that $k = e$. In this case, g is the minimal item in $\mathcal{T}_i(e)$ with $\Theta(g) = \Theta(e) - 1$ by the assumption of the lemma that every push is executed with respect to a minimal suitable item.

- For a saturating push, we have $\alpha = w_i(e, g)$. This implies $z'_i - \chi_g + \chi_e \notin \mathcal{B}_i$ and thus, $g \notin \mathcal{T}'_i(e)$ by Lemma 2.14. Moreover, $\rho_i(\mathcal{T}_i(e)) = z_i(\mathcal{T}_i(e)) = z'_i(\mathcal{T}_i(e))$ since $e, g \in \mathcal{T}_i(e)$. Hence, $\mathcal{T}_i(e)$ is a tight set with respect to z'_i and, thus, $\mathcal{T}'_i(e) \subseteq \mathcal{T}_i(e)$. We get $\mathcal{T}'_i(e) \subseteq \mathcal{T}_i(e) \setminus \{g\}$, which implies that the item fulfilling all requirements is strictly increasing.

- For a non-saturating push, we get $\alpha < w_i(e, g)$. Thus, $z'_i - (w_i(e, g) - \alpha)\chi_g + (w_i(e, g) - \alpha)\chi_e \in \mathcal{B}_i$, which implies $g \in \mathcal{T}'_i(e)$ by Lemma 2.14. Moreover, $\mathcal{T}_i(e)$ is tight with respect to z'_i , hence, $\mathcal{T}'_i(e) \subseteq \mathcal{T}_i(e)$. This implies that the minimal item fulfilling all requirements remains the same.

Next, we show monotonicity for $k \neq e$. To this end, we distinguish three cases:

- If $k, g \notin \mathcal{T}_i(e)$, then $\mathcal{T}'_i(e) = \mathcal{T}_i(e)$ and thus, (15) remains the same.
- If $k \in \mathcal{T}_i(e)$, it holds that $\mathcal{T}_i(k) \subseteq \mathcal{T}_i(e)$ by Lemma 2.14. Thus, we get that also g is in $\mathcal{T}_i(e)$, as $g \in \mathcal{T}_i(k) \subseteq \mathcal{T}_i(e)$. This implies that $z'_i(\mathcal{T}_i(e)) = z_i(\mathcal{T}_i(e))$. Since $\mathcal{T}_i(e)$ is a tight set, we have $z_i(\mathcal{T}_i(e)) = \rho(\mathcal{T}_i(e))$. In other words, $\mathcal{T}_i(e)$ remains tight after the push. This yields $\mathcal{T}'_i(e) \subseteq \mathcal{T}_i(e)$, as $\mathcal{T}'_i(e)$ is defined as the minimal tight set containing e and $\mathcal{T}_i(e)$ is already a valid candidate. Monotonicity of (15) follows immediately.
- If $k \notin \mathcal{T}_i(e)$ and $g \in \mathcal{T}_i(e)$, it holds that $\mathcal{T}'_i(e) \subseteq \mathcal{T}_i(e) \cup \mathcal{T}_i(k)$. This follows since the set $\mathcal{T}_i(e) \cup \mathcal{T}_i(k)$ is tight and remains tight after the push since it contains k and g . Now, we distinguish three cases, based on the level of k . First assume $\Theta(k) > \ell$. Then $\Theta_{\min}(\mathcal{T}_i(k)) \geq \ell$ by (L2) and therefore $\{f \in E \mid \Theta(f) = \ell - 1\} \cap \mathcal{T}'_i(e) \subseteq \mathcal{T}_i(e)$. In other words, every item that is a candidate in the minimum after the push was also in $\mathcal{T}_i(e)$ and thus a candidate before the push since all items in $\mathcal{T}_i(k)$ have a label which is too high. Next, consider the case $\Theta(k) = \ell$. Then we get

$$\min\{f' \mid f' \in \mathcal{T}_i(k), \Theta(f') = \ell - 1\} = g \geq \min\{f' \mid f' \in \mathcal{T}_i(e), \Theta(f') = \ell - 1\},$$

since we are now assuming $g \in \mathcal{T}_i(e)$ and $\Theta(g) = \ell - 1$ and thus, g is a valid candidate for the minimum on the right side. This shows the desired monotonicity of (15) since we showed that $\mathcal{T}'_i(e) \subseteq \mathcal{T}_i(e) \cup \mathcal{T}_i(k)$. Finally, $\Theta(k) < \ell$ is not possible, as on one hand we have $\Theta(g) = \Theta(k) - 1 < \ell - 1$ since we push on k with respect to g and on the other hand we have $g \in \mathcal{T}_i(e)$ by case distinction and thus, $\Theta(g) \geq \Theta(e) - 1 = \ell - 1$. This is a contradiction.

Hence monotonicity of (15) holds in every case. \square

Note that when pushes are only executed with respect to lexicographically minimal item-buyer pairs, the condition of Lemma 4.7 is fulfilled. Thus, we can apply the lemma to obtain monotonicity of (15) for every buyer i . This implies that, when fixing i , the potential items for a push at e are monotonically increasing, i.e., not below the items considered so far. Note that this implies that whenever there was no push for a buyer i at item e , there will not be any push for this buyer until e is relabeled. Moreover, at a saturating push (15), increases strictly and thus also the buyer-item pair. Hence, the procedure in [Polymatroid Sum Algorithm via Push-Relabel](#) is indeed a valid implementation of the general push-relabel framework.

COROLLARY 4.8. *The [Polymatroid Sum Algorithm via Push-Relabel](#) returns an optimal solution to (4).*

Next, we aim at bounding the running time.

4.3.3 Running Time Analysis. As a first step to show Theorem 4.6, we bound the number of non-saturating pushes. This is crucial since for these pushes, we cannot proceed with the next lexicographically higher item-buyer pair, but instead have to reconsider a pair that we looked at already. The key ingredient of the proof is that we always choose an undersold item of maximal level below m for a push and the fact that an item is not undersold anymore after a non-saturating push and can only get undersold again by a push on an item from a higher level.

LEMMA 4.9. *The number of non-saturating pushes in the [Polymatroid Sum Algorithm via Push-Relabel](#) is at most m^3 .*

PROOF. A *phase* of the algorithm is the set of iterations between two relabel operations. Once every item $e \in E$ has reached $\Theta(e) = m$, the algorithm stops since then there is no item left we can choose for a push. Since an item with level m is never selected, there are at most m^2 phases (m items, that are relabeled at most m times).

We will show that for each $e \in E$, there can be at most one non-saturating push in each phase. Here, we exploit that the algorithm picks an undersold item e with $\Theta(e) = \ell$ maximal. After a non-saturating push, item e is not undersold anymore. Thus, it cannot be picked again before it becomes undersold again. Push operations of an item at level ℓ' may create new undersold items but only on level $\ell' - 1$. Therefore, item e could become undersold again due to a push at an undersold item f with $\Theta(f) = \ell + 1$. However, no undersold item at level $\ell + 1$ may appear without a relabeling after e was picked, i.e., within the same phase.

Hence, there can be at most one non-saturating push per item, i.e., at most m non-saturating pushes in a phase. As there are at most m^2 phases, there are at most m^3 non-saturating pushes during the algorithm. \square

LEMMA 4.10. *The Polymatroid Sum Algorithm via Push-Relabel runs in $O(n \cdot \text{DO} + nm^3 \cdot \text{ExO})$ time.*

PROOF. We can find an initial allocation in $n \cdot \text{DO}$ time by asking every buyer i for a vector $z_i \in \mathcal{B}_i$.

According to the previous lemma, the total number of non-saturating pushes is $O(m^3)$. We estimate the running time of all relabel and saturating push operations induced by a fixed item e while it is on level $\Theta(e) = \ell$. Given e , we go through all buyer-item combinations in a structured way such that we consider every buyer-item combination at most once unless there is a non-saturating push. So, without the non-saturating pushes an item e on a fixed level ℓ induces a running time of mn arithmetic operations and edge-weight queries, in particular one ExO query for each item-buyer pair. Since there are m items, that can be on at most m different levels with a total addition of m^3 overall, the running time is $O(nm^3 + m^3) = O(nm^3)$ edge-weight queries and arithmetic operations.

In the analysis so far, we ignored the time needed to pick an undersold item e on the maximal possible level below m . We show now that this is not a bottleneck operation. We do so by using a standard network flow push-relabel construction (see e.g., [1, Section 7.8]). In the beginning, we initialize a list per level giving all undersold items that are on this level. In the beginning, all lists are empty except for the level 0 list, which contains all undersold items. This construction takes $O(m)$ but is only needed once. Moreover, we keep track of the maximal level that has a non-empty list and is below m . The list is easily maintainable, since only one item changes the level and the property undersold per iteration. The total increase of the maximal relevant level is bounded by $O(m^2)$ (m items on up to m levels). Thus, the total decrease is at most $O(m^2)$. Hence, scanning the lists to find the first non-empty list is not a bottleneck operation. Having such a data structure, we can access any item from the list on the level of the pointer in $O(1)$. \square

Now, we are ready to prove our main Theorem 4.6 as we considered correctness and running time of the algorithm.

PROOF OF THEOREM 4.6. The statement follows directly by Corollary 4.8 and Lemma 4.10. \square

Using the Polymatroid Sum Algorithm via Push-Relabel and a breadth-first search in the exchange graph, we can compute the minimal maximally over-/underdemanded sets efficiently.

COROLLARY 4.11. *Given prices p , we can compute the minimal maximally overdemanding [underdemanded] set in time $O(n \cdot \text{DO} + nm^3 \cdot \text{ExO})$ in the multi-supply setting.*

This follows by Lemma 3.4 [Lemma 3.9, respectively] and Theorem 4.6. Recall that the demand and exchange oracle calls are based on the matroids corresponding to the preferred bundles $\check{\mathcal{D}}_i(p)$ or $\hat{\mathcal{D}}_i(p)$, respectively.

5 Minimal Packing Prices and Maximal Covering Prices are Walrasian

It is well-known that Walrasian prices form a complete lattice using the component-wise minimum and component-wise maximum as meet and join operations, respectively. Hence, there also exists a component-wise minimal and maximal Walrasian price vector p_* and p^* , which we already discussed in the previous sections.

Packing and covering prices are already interesting on their own as they both constitute weaker notions of equilibria (packing: every buyer gets a preferred bundle, covering: all items are sold). Moreover, packing and covering prices are guaranteed to exist, also for non strong gross substitutes valuation functions, e.g., by setting the prices very high (such that no buyer is interested in any good) or to zero, respectively. It is a natural question whether these relaxations inherit those lattice properties and whether minimal packing prices or maximal covering prices (if they exist) are guaranteed to be Walrasian. Particularly, the question whether minimal packing and maximal covering prices exist is interesting.

For non strong gross substitutes valuations, we show in Appendix A.2 and A.3 that packing and covering prices do not form a lattice. Moreover, the example in Appendix A.2 even shows that there is no unique minimal packing price vector.

However, for strong gross substitutes valuations, the answer is positive: in the following, we show that there is a unique minimal packing price vector and a unique maximal covering price vector. Moreover, we show that these vectors are equal to the minimal, respectively maximal, Walrasian price vector.

5.1 Reduction to Unit-supply Case

Recall that we are given some market consisting of items (types) $e \in E$, where $b(e)$ is the number of units which are available of item e and where each buyer $i \in N$ has a strong gross substitutes valuation function $v_i: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}_+$. Moreover, we set prices $p(e)$ for each unit of item e .

Now, we will transfer this setting to a market with unit-supply, i.e., where only one unit of each item is available. This unit-supply setting is used in the following proofs since with an individual price for every copy, it is easier to show the statements and we can directly use well-known results for the unit-supply case. This simple trick of copying the items is also used for example in [51, Appendix A.2].

Given an instance in the multi-supply setting, we construct an instance in the unit-supply setting by defining $b(e)$ copies $\tilde{e}_1, \dots, \tilde{e}_{b(e)}$ of item $e \in E$. We call the set of all copies of items \tilde{E} , i.e., $\tilde{E} := \{\tilde{e}_1, \dots, \tilde{e}_{b(e)} : e \in E\}$. This means that \tilde{e}_k is a copy of item e . For simplicity, we may omit the index if we pick an item of \tilde{E} . To convert the bundles z to the unit-supply setting, we go through the buyers one by one and allocate any $z_i(e)$ (preferably unsold) different copies of e to buyer i . Note that if z is packing there are enough unsold copies available, otherwise, we have to allocate some item copies to multiple buyers. We convert a bundle $S \subseteq \tilde{E}$ to a bundle $z \in [0, b]_{\mathbb{Z}}$ by setting $z(e) = |\{\tilde{e} \in S \mid \tilde{e} \text{ is a copy of } e\}|$. Thus, the valuation of S is given by $v(S) := v(z)$.

Note that in the unit-supply setting, it is allowed that two different copies \tilde{e}_j and \tilde{e}_k of item e have different prices. However, in the following lemma, we show that in the minimal and maximal Walrasian price vector no two copies of an item have a different price. Given this, converting a price vector from the unit-supply setting to the multi-supply setting becomes straight-forward. For the reverse direction, given a multi-supply price vector p , we can construct a unit-supply price

vector by just setting the price of each copy of e to $p(e)$. We stick to calling these prices p since it will be clear from the context whether we are talking about the multi-supply or the unit-supply setting.

LEMMA 5.1. *A vector p is the minimal [maximal] Walrasian price vector in the multi-supply setting if and only if p is also the minimal [maximal] Walrasian price vector in the corresponding unit-supply setting, i.e., the prices of each copy of an item e are given by $p(e)$.*

PROOF. It suffices to show that the minimal [maximal] Walrasian price vector \tilde{p} in the unit-supply setting has the same prices for each copy of an item. If two item copies \tilde{e}_j and \tilde{e}_k of item type e do not have the same price, we will show that this price vector is not the minimal, respectively maximal, Walrasian price vector.

If $\tilde{p}(\tilde{e}_j) \neq \tilde{p}(\tilde{e}_k)$, we have that \tilde{p}' with $\tilde{p}'(\tilde{f}) = \tilde{p}(\tilde{f})$ for $\tilde{f} \in \tilde{E} \setminus \{\tilde{e}_j, \tilde{e}_k\}$, $\tilde{p}'(\tilde{e}_j) = \tilde{p}(\tilde{e}_k)$, and $\tilde{p}'(\tilde{e}_k) = \tilde{p}(\tilde{e}_j)$ is Walrasian by symmetry. Hence, as Walrasian prices form a lattice [37], $\tilde{p} \wedge \tilde{p}'$ and $\tilde{p} \vee \tilde{p}'$ are Walrasian as well. This is a contradiction to \tilde{p} being the minimal, respectively maximal, Walrasian price vector and thus, all item copies have the same price given the minimal [respectively maximal] Walrasian price vector \tilde{p} . \square

The following lemma states that when using the described reduction to the unit supply setting (see Section 5.1), the valuation functions of the buyers remain gross substitutes.

LEMMA 5.2 ([51, PROPOSITION A.1]). *The valuation function $v: 2^{\tilde{E}} \rightarrow \mathbb{Z}_+$ is gross substitutes.*

5.2 Minimal Packing Prices are Walrasian

It is already known that in the unit-supply setting, packing prices are also Walrasian if they lie below a Walrasian price vector [8]. However, this is not sufficient to show that there exists a unique minimal packing price vector (which is Walrasian), as it is not *a priori* clear whether there can be a packing price vector that is in some component smaller (but in another larger) than the minimal Walrasian price vector. In this section, we show that this is never the case as the minimal Walrasian price vector is a componentwise lower bound on every packing price vector.

THEOREM 5.3. *Given an instance where all valuation functions are strong gross substitutes, let p_* be the minimal Walrasian price vector. Then $p_* \leq q$ for any packing price vector q .*

Remember that every Walrasian price vector is also packing, but not vice versa. Thus, we can reformulate the main theorem of this section as follows.

COROLLARY 5.4. *There exists a component-wise minimal packing price vector and it is equal to the minimal Walrasian price vector.*

To show Theorem 5.3, we will use the following lemma, that compares $p_*(e)$ and $q(e)$ for the items e which are at least partially sold in a packing allocation w.r.t. prices q .

LEMMA 5.5. *Let p_* be the minimal Walrasian price vector and let q be any packing price vector. Moreover, let (z_1, \dots, z_n) be a packing allocation w.r.t. prices q . Then $p_*(e) \leq q(e)$ for all $e \in E$ with $\sum_{i \in N} z_i(e) > 0$.*

PROOF. To prove this lemma, we will use the reduction to the unit-supply case. Therefore, let \tilde{E} denote the set containing all item copies and let (S_1, \dots, S_n) be the sets of bundles in the unit-supply setting corresponding to (z_1, \dots, z_n) .

By Lemma 5.1 and since p_* is a minimal Walrasian price vector in the multi-supply setting, it is also a minimal Walrasian price vector in the unit-supply setting. Hence, in the unit-supply setting,

it minimizes the submodular Lyapunov function $L(p) := \sum_{i \in N} V_i(p) + \sum_{\tilde{e} \in \tilde{E}} p(\tilde{e})$ (see Lemma 2.3), i.e., p_* is an optimal solution to

$$\alpha := \min \left\{ \sum_{i \in N} V_i(p) + \sum_{\tilde{e} \in \tilde{E}} p(\tilde{e}) \mid p(\tilde{e}) \geq 0 \text{ for all } \tilde{e} \in \tilde{E} \right\}. \quad (16)$$

Let $S = \bigcup_{i \in N} S_i$ be the set of item copies chosen in the preferred bundles. Now, we define a new price vector q' by

$$q'(\tilde{e}) := \begin{cases} q(\tilde{e}) & \text{if } \tilde{e} \in S, \\ \infty & \text{otherwise.} \end{cases}$$

Note that $V_i(q) = V_i(q')$ holds by definition. The price vector q restricted to S is Walrasian. The same holds for q' , thus, it minimizes the Lyapunov function restricted to S . In other words, q' is an optimal solution to

$$\beta := \min \left\{ \sum_{i \in N} V_i(p) + p(S) \mid p(\tilde{e}) \geq 0 \text{ for all } \tilde{e} \in S \text{ and } p(\tilde{e}) = \infty \text{ for all } \tilde{e} \in \tilde{E} \setminus S \right\}. \quad (17)$$

Next, we define the join $\check{p} := p_* \vee q'$ and the meet $\hat{p} := p_* \wedge q'$ of the two price vectors p_* and q' (i.e., $\check{p}(\tilde{e}) := \max\{p_*(\tilde{e}), q'(\tilde{e})\}$ and $\hat{p}(\tilde{e}) := \min\{p_*(\tilde{e}), q'(\tilde{e})\}$). Then, the following holds:

$$\begin{aligned} \alpha + \beta &= \sum_{j \in N} V_j(p_*) + p_*(\tilde{E}) + \sum_{j \in N} V_j(q') + q'(S) = \sum_{j \in N} (V_j(p_*) + V_j(q')) + (p_*(\tilde{E}) + q'(S)) \\ &\geq \sum_{j \in N} (V_j(\hat{p}) + V_j(\check{p})) + (\hat{p}(\tilde{E}) + \check{p}(S)) = \sum_{j \in N} V_j(\hat{p}) + \hat{p}(\tilde{E}) + \sum_{j \in N} V_j(\check{p}) + \check{p}(S) \\ &\geq \alpha + \beta, \end{aligned}$$

where the first inequality follows from submodularity of $p \mapsto V_j(p)$ (cf. [7, Theorem 10]) and since $p_*(\tilde{e}) + q'(\tilde{e}) = \check{p}(\tilde{e}) + \hat{p}(\tilde{e})$ for all $\tilde{e} \in S$ and $\hat{p}(\tilde{e}) = p_*(\tilde{e})$ for all $\tilde{e} \in \tilde{E} \setminus S$. The second inequality follows since \hat{p} is a feasible price vector for (16) and since \check{p} is a feasible price vector for (17). Therefore, equality must hold in the equation system. This implies that \hat{p} is an optimal solution for (16) and thus a Walrasian price vector. Since p_* is the minimal Walrasian price vector, it holds that $p_*(\tilde{e}) \leq \hat{p}(\tilde{e}) = \min\{p_*(\tilde{e}), q'(\tilde{e})\} \leq q'(\tilde{e}) = q(\tilde{e})$ for $\tilde{e} \in S$. By definition of S , it follows that $p_*(e) \leq q(e)$ holds for all $e \in E$ with $\sum_{i \in N} z_i(e) > 0$. \square

PROOF OF THEOREM 5.3. Let q be a minimal packing price vector and consider a packing allocation (z_1, \dots, z_n) w.r.t. prices q . By Lemma 5.5, we already know that $p_*(e) \leq q(e)$ for all e with $\sum_{i \in N} z_i(e) > 0$. For e with $z_i(e) = 0$ for all $i \in N$, we consider two cases.

If $z'_i(e) > 0$ for some $z'_i \in \mathcal{D}_i(q)$ for some buyer $i \in N$, we show that there is $z''_i \in \mathcal{D}_i(q)$ with $z''_i(e) > 0$ and $z''_i \leq z_i + \chi_e$: Since both z_i and z'_i are preferred bundles of buyer i , they maximize the utility function u_i . Since the utility function u_i is M^{\natural} -concave, for $e \in \text{supp}^+(z'_i - z_i)$ one of the following cases holds (see Definition 2.5):

- It holds that $u_i(z'_i) + u_i(z_i) \leq u_i(z'_i - \chi_e) + u_i(z_i + \chi_e)$. Hence, equality holds since z'_i and z_i are maximizers of the utility and thus, $z''_i = z_i + \chi_e$ is also a preferred bundle.
- There exists an $f \in \text{supp}^-(z'_i - z_i)$ with $u_i(z'_i) + u_i(z_i) \leq u_i(z'_i - \chi_e + \chi_f) + u_i(z_i + \chi_e - \chi_f)$. Again, equality must hold and thus $z''_i = z_i + \chi_e - \chi_f$ is a preferred bundle.

In both cases, there exists a packing allocation $\bar{z} = (z_1, \dots, z_{i-1}, z''_i, \dots, z_n)$ with respect to prices q with $\sum_{i \in N} \bar{z}_i(e) > 0$. Thus, it follows by Lemma 5.5 that $p_*(e) \leq q(e)$.

If $e \notin \text{supp}^+(z'_i)$ for any $z'_i \in \mathcal{D}_i(q)$ and any buyer $i \in N$, we can reduce the price of e slightly,⁷ which results in a smaller packing price vector and thus in a contradiction to the choice of q . \square

5.3 Maximal Covering Prices are Walrasian

In this section, we establish that maximal covering prices are Walrasian. More precisely, we show that first of all, there exists a unique component-wise maximal covering price vector, and second that this price vector coincides with the maximal Walrasian price vector.

The main result of this section can be formulated as follows:

THEOREM 5.6. *Given an instance where all valuation functions are strong gross substitutes, let p^* be the maximal Walrasian price vector. Then $p^* \geq q$ for any covering price vector q .*

COROLLARY 5.7. *There exists a component-wise maximal covering price vector and it is equal to the maximal Walrasian price vector.*

The remainder of this section is dedicated to the proof of Theorem 5.6. By Lemma 5.1, it suffices to establish Theorem 5.6 in the unit-supply setting because for any covering vector q in the multi-supply setting, q will also be covering in the corresponding unit-supply instance.

Fix a unit-supply instance with item set E and buyers $i \in N$ with gross substitutes valuation functions $v_i: 2^E \rightarrow \mathbb{Z}_+$. Denote the maximal Walrasian price vector by p^* and let q be covering. Fix preferred bundles $(S_i)_{i \in N}$ for prices q with $\bigcup_{i \in N} S_i = E$. For $e \in E$, let $k_e := |\{i \in N \mid e \in S_i\}|$ denote the multiplicity of item e among the bundles $(S_i)_{i \in N}$.

Our strategy to prove Theorem 5.6 is to construct an instance in which we have k_e copies of each item, and argue that q will be Walrasian prices for that instance (assigning to each copy of e a price of $q(e)$).

To this end, let $E' := \{e'_1, \dots, e'_{k_e} \mid e \in E\}$ and let $\pi: E' \rightarrow E, e'_i \mapsto e$ be the projection of E' onto E . We define new valuation functions $v'_i: E' \rightarrow \mathbb{Z}_+$ for $i \in N$ by setting

$$v'_i(S') := v_i(\pi(S')).$$

For a price vector $p \in \mathbb{R}_+^E$ [$p' \in \mathbb{R}_+^{E'}$], we denote the set of preferred bundles of buyer i w.r.t. prices p [prices p'] by $\mathcal{D}_i(p)$ [$\mathcal{D}'_i(p')$].

For a price vector $p' \in \mathbb{R}_+^{E'}$, we denote by $p^\downarrow \in \mathbb{R}_+^E$ the price vector with

$$p^\downarrow(e) := \min\{p'(e'_i) \mid i \in \{1, \dots, k_e\}\}.$$

LEMMA 5.8. *Let $p' \in \mathbb{R}_+^{E'}$ and $S' \subseteq E'$. Then $S' \in \mathcal{D}'_i(p')$ for an $i \in N$ if and only if*

- $\pi(S') \in \mathcal{D}_i(p^\downarrow)$,
- S' only contains cheapest copies of each item, and
- whenever S' contains multiple copies of an item, all of them have price 0.

PROOF. Let $p' \in \mathbb{R}_+^{E'}$ and let $S' \subseteq E'$ be an arbitrary set and $S := \pi(S')$ its projection. Further, let $T \in \mathcal{D}_i(p^\downarrow)$ and let T' arise from T by picking a cheapest copy of each item in T . Then

$$v'_i(S') - p'(S') = v_i(S) - p'(S') \stackrel{(*)}{\leq} v_i(S) - p^\downarrow(S) \stackrel{(**)}{\leq} v_i(T) - p^\downarrow(T) = v'_i(T') - p'(T').$$

The first inequality holds by definition of p^\downarrow and the second inequality holds by definition of T . This chain of inequalities shows that $T' \in \mathcal{D}'_i(p')$ because $S' \subseteq E'$ was chosen arbitrarily. Knowing

⁷Note, that we are talking about preferred bundles which do not need to be minimal. Hence, if $p(e) = 0$, it can be added to every preferred bundle of every buyer.

this, we have $S' \in \mathcal{D}'_i(p')$ if and only if the outer inequality is tight, which is the case if and only if (*) and (**) are tight.

(*) is tight if and only if $p'(S') = p^\perp(S)$, i.e., S' only contains cheapest copies of any item and whenever S contains multiple copies of an item, all of them have price 0. (**) is tight if and only if $S \in \mathcal{D}_i(p^\perp)$. \square

LEMMA 5.9. *The valuation function v'_i is gross substitutes for each $i \in N$.*

PROOF. By Lemma 2.8, it suffices to establish that v'_i is M^h -concave for each $i \in N$.

Let $X', Y' \subseteq E'$ and let $X := \pi(X')$ and $Y := \pi(Y')$. Let further $e' \in X' \setminus Y'$ and define $e := \pi(e')$. Then $e \in X$.

Case 1: It holds that $e \notin Y$.

Then, as v_i is gross substitutes, it is M^h -concave (see Lemma 2.8), so at least one of the Equations ((M1)) or ((M2)) holds:

Case 1.1: Equation ((M1)) holds, i.e., $v_i(X) + v_i(Y) \leq v_i(X \setminus \{e\}) + v_i(Y \cup \{e\})$.

As $X \setminus \{e\} \subseteq \pi(X' \setminus \{e'\})$ and $\pi(Y' \cup \{e'\}) = Y \cup \{e\}$, it holds by monotonicity of v_i that $v'_i(X') + v'_i(Y') = v_i(X) + v_i(Y) \leq v_i(X \setminus \{e\}) + v_i(Y \cup \{e\}) \leq v'_i(X' \setminus \{e'\}) + v'_i(Y' \cup \{e'\})$.

Case 1.2: Equation ((M2)) holds, i.e., $v_i(X) + v_i(Y) \leq v_i(X \setminus \{e\} \cup \{f\}) + v_i(Y \setminus \{f\} \cup \{e\})$ for an $f \in Y \setminus X$. Let $f' \in Y'$ with $\pi(f') = f$. Then $f' \in Y' \setminus X'$. Moreover, $X \setminus \{e\} \cup \{f\} \subseteq \pi(X' \setminus \{e'\} \cup \{f'\})$ and $Y \setminus \{f\} \cup \{e\} \subseteq \pi(Y' \setminus \{f'\} \cup \{e'\})$. Hence, by monotonicity of v_i , we get

$$\begin{aligned} v'_i(X') + v'_i(Y') &= v_i(X) + v_i(Y) \leq v_i(X \setminus \{e\} \cup \{f\}) + v_i(Y \setminus \{f\} \cup \{e\}) \\ &\leq v'_i(X' \setminus \{e'\} \cup \{f'\}) + v'_i(Y' \setminus \{f'\} \cup \{e'\}). \end{aligned}$$

Hence, either (M1) or (M2) holds for v'_i .

Case 2: It holds that $e \in Y$.

Case 2.1: There exists an $e'' \in X' \setminus \{e'\}$ with $\pi(e'') = e$.

Then we have $\pi(X' \setminus \{e'\}) = \pi(X') = X$ and $\pi(Y' \cup \{e'\}) = \pi(Y') = Y$. This yields

$$v'_i(X') + v'_i(Y') = v_i(X) + v_i(Y) = v'_i(X' \setminus \{e'\}) + v'_i(Y' \cup \{e'\}).$$

Case 2.2: There is no $e'' \in X' \setminus \{e'\}$ with $\pi(e'') = e$.

As $e \in Y$, there exists an $f' \in Y'$ with $\pi(f') = e$. It holds that $e' \notin Y'$, so $f' \neq e'$. In particular, it follows that $f' \notin X'$. Hence, $f' \in Y' \setminus X'$. We further have $\pi(X' \setminus \{e'\} \cup \{f'\}) = X$ and $\pi(Y' \setminus \{f'\} \cup \{e'\}) = Y$. This results in

$$v'_i(X') + v'_i(Y') = v_i(X) + v_i(Y) = v'_i(X' \setminus \{e'\} \cup \{f'\}) + v'_i(Y' \setminus \{f'\} \cup \{e'\}).$$

Hence, also either (M1) or (M2) holds for v'_i .

Thus, v'_i is M^h -concave and thus gross substitutes by Lemma 2.8. \square

LEMMA 5.10. *Let $q \in \mathbb{R}_+^E$ be a covering price vector. Then $q' \in \mathbb{R}_+^E$ with $q'(e') := q(\pi(e'))$ for $e' \in E'$ is a Walrasian price vector with respect to supply E' and valuation functions v'_i for $i \in N$.*

PROOF. By our definition of the numbers $(k_e)_{e \in E}$, we can partition E' into sets $(S'_i)_{i \in N}$ with $\pi(S'_i) = S_i$ and $|S'_i| = |S_i|$ for $i \in N$. By Lemma 5.8, we have $S'_i \in \mathcal{D}'_i(q')$ for all $i \in N$. \square

Now, we are ready to prove the main theorem of this section.

PROOF OF THEOREM 5.6. As mentioned before, it suffices to show this statement in the unit-supply setting, i.e., we assume $b(e) = 1$ for all $e \in E$. Let p^* be the maximal Walrasian price vector and let q be any covering price vector.

In the following, we will use that Walrasian prices are minimizers of the Lyapunov function (see Lemma 2.3). We distinguish the settings by using L for the Lyapunov function with respect to E and v_i for all $i \in N$, and L' for the Lyapunov function with respect to E' and v'_i . Analogously, let $V_i(p) = \max_{S \subseteq E} v_i(S) - p(S)$ and $V'_i(p') = \max_{S' \subseteq E'} v'_i(S') - p'(S')$.

We know that p^* is a minimizer of the Lyapunov function in the original setting, i.e., it minimizes

$$L(p) = \sum_{i \in N} V_i(p) + p(E).$$

Moreover, q' with $q'(e') := q(\pi(e'))$ for $e' \in E'$ is a Walrasian price vector for E' and $(v'_i)_{i \in N}$ (see Lemma 5.10), i.e., it minimizes

$$L'(p') = \sum_{i \in N} V'_i(p') + p'(E').$$

Note that $q^\downarrow = q$ by definition of q' . By Lemma 5.8, we have

$$V_i(q') = \max_{S' \subseteq E'} v'_i(S') - q'(S') = \max_{S \subseteq E} v_i(S) - q^\downarrow(S) = V_i(q^\downarrow) = V_i(q),$$

for all $i \in N$. Thus,

$$L'(q') = \sum_{i \in N} V_i(q) + q'(E').$$

Define $\check{p} \in \mathbb{R}_+^E$ and $\hat{q}' \in \mathbb{R}_+^{E'}$ via

$$\check{p}(e) := \max\{p^*(e), q(e)\} \quad \text{and} \quad \hat{q}'(e'_\ell) := \begin{cases} \min\{p^*(e), q(e)\} & \text{for } \ell = 1, \\ q(e) & \text{otherwise.} \end{cases}$$

This yields

$$\begin{aligned} L(p^*) + L'(q') &= \sum_{i \in N} V_i(p^*) + p^*(E) + \sum_{i \in N} V_i(q) + q'(E') \\ &= \sum_{i \in N} (V_i(p^*) + V_i(q)) + \sum_{e \in E} (p^*(e) + k_e \cdot q(e)) \\ &\geq \sum_{i \in N} (V_i(\max\{p^*, q\}) + V_i(\min\{p^*, q\})) \\ &\quad + \sum_{e \in E} (\max\{p^*(e), q(e)\} + \min\{p^*(e), q(e)\} + (k_e - 1) \cdot q(e)) \\ &= \sum_{i \in N} V_i(\check{p}) + V_i(\hat{q}'^\downarrow) + \check{p}(E) + \hat{q}'(E') \\ &= L(\check{p}) + L'(\hat{q}'), \end{aligned}$$

where in the first inequality, we used that V_i is lattice submodular (cf. [7, Theorem 10]). Hence, \check{p} and \hat{q}' are minimizers of L and L' , respectively, as well. Thus, \check{p} is Walrasian. As p^* is the maximal Walrasian price vector, $q(e) \leq \max\{p^*(e), q(e)\} = \check{p}(e) \leq p^*(e)$ for all $e \in E$. This yields $q \leq p^*$. \square

6 Monotonicity Analysis

In this section, we show monotonicity results in supply and demand for the (unique) buyer-optimal, as well as for the (again unique) seller-optimal Walrasian prices, provided that all buyers' valuations are strong gross substitutes. Independently, Raach [56] proved the same monotonicity results. We also provide an example showing that monotonicity cannot be guaranteed if the valuation functions are changed. We model the decrease in supply by restricting the market to less copies of items $b' \leq b$ with $b'(e) < b(e)$ for at least one $e \in E$ and by restricting all valuations to $v'_i = v_i|_{[0, b']_{\mathbb{Z}}}$.

The decrease in demand is modeled as item-truncation, i.e., we assume that a buyer i is only interested in d_i items in total. Since it is known that strong gross substitutes functions are closed under item-truncation (cf. [19]⁸), all valuation functions remain (strong) gross substitutes after this operation. Formally, this leads to the following definition.

Definition 6.1. Let $v_i: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}_+$ be a valuation function and $d \in \mathbb{Z}_+^N$. The *truncation* v_i^d of buyer i to demand d is given by $v_i^d(z) := \max\{v_i(y) \mid y \leq z, \sum_{e \in E} y(e) \leq d_i\}$.

We write $\mathcal{D}_i^d(p)$ [respectively $\check{\mathcal{D}}_i^d(p)$ and $\hat{\mathcal{D}}_i^d(p)$] to denote the set of [minimal/maximal] preferred bundles of buyer i at prices p with valuation function v_i^d . Note that if we choose d_i large enough, it holds that $v_i(z) = v_i^d(z)$ for all $z \in [0, b]_{\mathbb{Z}}$. Hence, without loss of generality, we assume that each buyer indeed has a demand and denote the demand vector by $d = (d_1, \dots, d_n)$.

The monotonicity properties we want to show are the following:

- *in supply:* Let p be the buyer-optimal [seller-optimal] Walrasian price vector. If the supply of some items decreases, i.e., if there are only $b' \leq b$ items to be sold, then the corresponding buyer-optimal [seller-optimal] Walrasian prices p' for items $e \in E' = \{f \in E \mid b'(f) > 0\}$ will only increase, i.e., $p(e) \leq p'(e)$ for all $e \in E'$.
- *in demand:* Consider two demand vectors d' and d in \mathbb{Z}_+^N with $d' \leq d$, and let p and p' be the buyer-optimal [seller-optimal] Walrasian price vectors if the valuation functions $v_i: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}_+$ of the buyers i are truncated at demands d and d' , respectively. Then the buyer-optimal [seller-optimal] Walrasian prices p' will only decrease if the demand decreases, i.e., $p' \leq p$.

THEOREM 6.2. *Monotonicity in supply and demand with respect to the minimal and the maximal Walrasian price vector can be guaranteed if all buyers' valuations are strong gross substitutes.*

The main idea of the proof is to use the equivalence between minimal packing [maximal covering] and minimal [maximal] Walrasian prices. We first show monotonicity in supply.

LEMMA 6.3. *Monotonicity in supply with respect to the minimal Walrasian price vector can be guaranteed if all buyers' valuations are strong gross substitutes.*

PROOF. Let b and b' be two supply vectors with $b' \leq b$. Assume without loss of generality that b and b' only differ in one item f and only in one unit, i.e., $b(f) - b'(f) = 1$ and $b(e) = b'(e)$ for all $e \in E \setminus \{f\}$. We fix a Walrasian allocation (z_1, \dots, z_n) given the minimal Walrasian prices p' at supply E' . Now, we use the reduction to the unit-supply setting. Therefore, let \tilde{E} denote the set containing all item copies with respect to b . Further let (S_1, \dots, S_n) be the packing allocation corresponding to (z_1, \dots, z_n) .

We fix an arbitrary copy $\tilde{f} \in \tilde{E}$ of the unique item $f \in \text{supp}^+(b - b')$. Next, we adapt the prices p' to prices \tilde{p} by setting $\tilde{p}(\tilde{e}) = p'(\tilde{e})$ for $\tilde{e} \in \tilde{E} \setminus \{\tilde{f}\}$ and $\tilde{p}(\tilde{f}) = \max_{i \in N} v_i(E) + 1$. Thus, no buyer wants to buy item \tilde{f} , so any packing allocation with respect to supply $\tilde{E} \setminus \{\tilde{f}\}$ and prices p' is also packing with respect to supply \tilde{E} and prices \tilde{p} . As there is a packing allocation in the first setting, we know that the same allocation is packing for prices \tilde{p} at supply \tilde{E} . Using that p is the component-wise minimal packing price vector at supply \tilde{E} , we get by Theorem 5.3 that $p(\tilde{e}) \leq \tilde{p}(\tilde{e}) = p'(\tilde{e})$ for all $\tilde{e} \in \tilde{E} \setminus \{\tilde{f}\}$.

Thus, $p(e) \leq p'(e)$ holds for every item $e \in E$ with a copy in $\tilde{E} \setminus \{\tilde{f}\}$, i.e., for the items $e \in E$ with $b'(e) > 0$. \square

⁸Collina and Weinberg [19] showed it in the unit-supply setting. The statement for the multi-supply setting follows by copying the items as presented in Section 5.1.

LEMMA 6.4. *Monotonicity in supply with respect to the maximal Walrasian price vector can be guaranteed if all buyers' valuations are strong gross substitutes.*

PROOF. As before, let b and b' be two supply vectors with $b' \leq b$ and assume without loss of generality that there is $f \in E$ such that $b(f) = b'(f) + 1$ and $b(e) = b'(e)$ for all $e \in E \setminus \{f\}$. Let (z_1, \dots, z_n) be a Walrasian allocation for the maximal Walrasian prices p at supply b . Again, we reduce to the unit-supply setting. Let \tilde{E} denote the set containing all item copies with respect to b and let $\tilde{p} \in \mathbb{R}_+^{\tilde{E}}$ be the price vector with $\tilde{p}(\tilde{e}) = p(e)$ for every copy \tilde{e} of $e \in E$. Let further (S_1, \dots, S_n) be the partition of \tilde{E} corresponding to (z_1, \dots, z_n) .

We fix an arbitrary copy \tilde{f} of f and define a price vector $\bar{p} \in \mathbb{R}_+^{\tilde{E}}$ by setting $\bar{p}(\tilde{e}) = \tilde{p}(\tilde{e})$ for all $\tilde{e} \in \tilde{E} \setminus \{\tilde{f}\}$ and $\bar{p}(\tilde{f}) = \max_{i \in N} v_i(E) + 1$. As the valuation functions $(\tilde{v}_i)_{i \in N}$ are gross substitutes, we know that there exist preferred bundles $(\tilde{S}_i)_{i \in N}$ with respect to \bar{p} such that $S_i \setminus \{\tilde{f}\} \subseteq \tilde{S}_i$. Moreover, by our choice of $\bar{p}(\tilde{f})$, we can infer that $\tilde{f} \notin \tilde{S}_i$ for $i \in N$. In particular, the bundles $(\tilde{S}_i)_{i \in N}$ also constitute preferred bundles for supply $\tilde{E} \setminus \{\tilde{f}\}$ and the restriction of \tilde{p} to $\tilde{E} \setminus \{\tilde{f}\}$. Furthermore, $\bigcup_{i \in N} S_i = \tilde{E}$ implies $\bigcup_{i \in N} \tilde{S}_i = \tilde{E} \setminus \{\tilde{f}\}$.

Going back to the multi-supply setting, this tells us that the prices p are covering for supply b' . By Theorem 5.6, we may conclude $p \leq p'$ where p' is the minimal Walrasian price vector for supply b' . \square

Now, we examine the effect of changes in the buyer's demand on the minimal [maximal] Walrasian prices.

LEMMA 6.5. *Monotonicity in demand with respect to the minimal Walrasian price vector can be guaranteed if all buyers' valuations are strong gross substitutes.*

LEMMA 6.6. *Monotonicity in demand with respect to the maximal Walrasian price vector can be guaranteed if all buyers' valuations are strong gross substitutes.*

Clearly, Theorem 6.2 follows directly from Lemmas 6.3–6.6. To show the last two lemmas, we will use that in the unit-supply setting, gross substitutes valuations are well-layered (cf. [24, 53]). Intuitively, well-layered means that the greedy algorithm (always selecting an item with the highest marginal value) computes in each iteration a bundle with the highest utility among all bundles of that size.

Definition 6.7 ([24]). In the unit-supply setting, let $v: 2^E \rightarrow \mathbb{Z}_+$ and $p \in \mathbb{Z}^E$. A *greedy sequence* for v and p is an ordering $e_1, \dots, e_{|E|}$ of the elements of E such that for every $j \in \{1, \dots, |E|\}$, we have

$$e_j \in \arg \max_{e \notin \{e_1, \dots, e_{j-1}\}} \{v(\{e_1, \dots, e_{j-1}, e\}) - v(\{e_1, \dots, e_{j-1}\}) - p(e)\}.$$

Based on our reduction to the unit-supply case, we obtain the following analogous notion of a greedy sequence for the multi-supply setting:

Definition 6.8. In the multi-supply setting, let $v: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{Z}_+$ and $p \in \mathbb{Z}^E$. A *greedy sequence* for v and p is a sequence $e_1, \dots, e_{b(E)}$ of elements of E with the following properties: For $j \in \{0, \dots, b(E)\}$, let $z^j \in [0, b]_{\mathbb{Z}}$ such that $z^j(e)$ equals the number of occurrences of e among e_1, \dots, e_j . Then $z^{b(E)} = b$ and for each $j \in \{1, \dots, b(E)\}$, we have $e_j \in \arg \max_{e \in \text{supp}^+(b - z^{(j-1)})} \{v(z^{(j-1)} + \chi_e) - v(z^{(j-1)}) - p(e)\}$.

Definition 6.9 ([24]). A valuation function $v: 2^E \rightarrow \mathbb{Z}_+$ is called *well-layered* if for any price vector $p \in \mathbb{Z}^E$ and any greedy sequence $e_1, \dots, e_{|E|}$ for v and p , we have

$$\{e_1, \dots, e_j\} \in \arg \max \{v(S) - p(S) : S \subseteq E, |S| = j\} \text{ for all } j \in \{1, \dots, |E|\}.$$

LEMMA 6.10 ([24]). *In the unit-supply setting, a non-decreasing gross substitutes function $v: 2^E \rightarrow \mathbb{Z}_+$ is well-layered.*

For non-decreasing gross substitutes valuations, it is further known that every bundle that is optimum for a given size can be obtained via a greedy sequence.

LEMMA 6.11 ([53, REMARK 6.5]). *In the unit-supply setting, let $v: 2^E \rightarrow \mathbb{Z}_+$ be a non-decreasing gross substitute valuation function and let $p \in \mathbb{Z}^E$. Let further $j \in \{1, \dots, |E|\}$ and $S^* \in \arg \max\{v(S) - p(S) : S \subseteq E, |S| = j\}$. Then there exists a greedy sequence $e_1, \dots, e_{|E|}$ for v and p such that $S^* = \{e_1, \dots, e_j\}$.*

In the multi-supply setting, we can obtain a similar result by using the reduction to the unit supply case as described in Section 5.1, in particular, that the corresponding valuation functions are gross substitutes and thus well-layered.

COROLLARY 6.12. *In the multi-supply setting, for a strong gross substitutes valuation function v , demand d and prices $p \in \mathbb{Z}_+^E$, and a preferred bundle $z \in \mathcal{D}_i^d(p)$, there exist a greedy sequence $e_1, \dots, e_{b(E)}$ for v^d and p , as well as $j^* \in \{0, \dots, b(E)\}$, such that*

- $z^{j^*} = z$,
- $\max_{e \in \text{supp}^+(b - z^{(j-1)})} \{v(z^{(j-1)} + \chi_e) - v(z^{(j-1)}) - p(e)\} \geq 0$ for all $j \in \{1, \dots, j^*\}$ and
- $\max_{e \in \text{supp}^+(b - z^{(j-1)})} \{v(z^{(j-1)} + \chi_e) - v(z^{(j-1)}) - p(e)\} \leq 0$ for all $j \in \{j^* + 1, \dots, b(E)\}$.

With this knowledge, we can prove Lemma 6.5.

PROOF OF LEMMA 6.5. Let d and d' be two demand vectors with $d' \leq d$. Let p_* and p_*' be the minimal Walrasian prices for demand d and d' , respectively, and let (z_1, \dots, z_n) be a Walrasian allocation for prices p_* and demand d .

We will show that for each buyer i , there is a vector $y_i \leq z_i$ with $y_i \in \mathcal{D}_i^{d'}(p_*)$. If $z_i(E) \leq d'_i$ we are done (we set $y_i = z_i$). Thus, we assume $z_i(E) > d'_i$ in the sequel. Since v_i^d is strong gross substitutes and $z_i \in \mathcal{D}_i^d(p_*)$, we can obtain z_i using the greedy algorithm (Corollary 6.12). Let y_i be the bundle of the d'_i items that were selected first. Hence, $y_i \leq z_i$. Moreover, all marginal returns in the greedy sequence leading to y_i are non-negative and by definition of $v_i^{d'}$, every further item has a non-positive marginal utility w.r.t. y_i . Hence, $y_i \in \mathcal{D}_i^{d'}(p_*)$.

Thus, for demand d' , we know that (y_1, \dots, y_n) is a packing allocation at prices p_* . Hence, p_* is a packing price vector. Since p_*' is the buyer-optimal Walrasian price vector at demand d' , it holds, by Theorem 5.3, that $p_*' \leq p_*$. \square

PROOF OF LEMMA 6.6. Let d and d' be two demand vectors with $d' \leq d$. Let p' be Walrasian prices for demand d' and let (z'_1, \dots, z'_n) be a Walrasian allocation for prices p' and demand d' . By Corollary 6.12, we can extend each bundle z'_i to a preferred bundle $z_i \in \mathcal{D}_i^d(p')$ with $z_i \geq z'_i$ (component-wise) by continuing the greedy allocation process until the marginal returns become negative. In particular, (z_1, \dots, z_n) will be a covering allocation for prices p' and demand d . By Theorem 5.6, we can conclude that for the maximal Walrasian prices p for demand d , we have $p' \leq p$. \square

In contrast to the monotonicity in supply and demand, there is no monotonicity if we change the valuation functions for some buyers and some items. We provide an example showing this for minimal Walrasian prices with purely additive valuation functions (which are strong gross substitutes):

Example 6.13. Let there be three items $E = \{e_1, e_2, e_3\}$ and three buyers with unit-demand valuations $v_1 = (2, 3, 0)$, $v_2 = (0, 1, 1)$ and $v_3 = (0, 1, 1)$, i.e., $v_i(S) = \max_{e \in S} v_{ie}$ for all $S \subseteq E$. The buyer-optimal Walrasian prices are given by $p_* = (0, 1, 1)$. Now, assume that the valuation of the first buyer for the second item e_2 is decreased by one, i.e., $v'_1 = (2, 2, 0)$. In this setting, the prices $p_*' = (0, 0, 0)$ is the buyer-optimal Walrasian price vector. Hence, the price of the items whose valuation was not changed are the same or reduced by one unit.

Consider the same setting again and assume that the first buyer decreases her valuation for item e_1 as well, i.e., $v''_1 = (1, 2, 0)$. The minimal Walrasian price vector is $p_*'' = (0, 1, 1)$. Thus, the price of the items whose valuation was not changed are increased by one unit.

The example shows that the minimal Walrasian prices do not change in a monotone way when the valuations change (an example for maximal Walrasian prices can be found in Appendix A.4). This seems to be intuitive since if the valuation of one item is decreased, it could increase the demand for other items which then increases the price. On the other hand, the item itself is less attractive, such that the demand can reduce (and thus also potential conflicts).

Note also that the question regarding price monotonicity w.r.t. changes in valuations is a somewhat very restricted problem, as the strong gross substitutes property might get lost if one is not careful with the changes in the valuation function.

7 Conclusion

We provide a simple, combinatorial algorithm to compute minimal maximally overdemanded sets and minimal maximally underdemanded sets for markets where items are available in multiple copies and all buyers have strong gross substitutes valuation functions. The algorithm is essentially an implementation of a polymatroid sum algorithm using an appropriate oracle model. It turns out to be quite useful as it allows for a fast execution of the dynamic auction step.

Moreover, we prove that unique minimal packing prices as well as unique maximal covering prices always exist. The unique minimal packing prices coincide with the minimal Walrasian prices, while the unique maximal packing prices coincide with the maximal Walrasian prices. The distinction between packing prices and Walrasian prices had been mostly overlooked in the literature. The clear distinction made here and the aforementioned results allow us to prove monotonicity properties of minimal and maximal Walrasian prices w.r.t. changes in demand and supply.

Acknowledgment

We thank András Frank and Georg Loho for fruitful discussions and comments. We also thank the three anonymous reviewers for helpful comments and suggestions to improve the article.

Data availability statement

No data are associated with this article. Data sharing is not applicable to this article.

References

- [1] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. 1988. *Network Flows*. Cambridge, Mass.: Alfred P. Sloan School of Management, Massachusetts.
- [2] Tommy Andersson, Christer Andersson, and A. J. J. Talman. 2013. Sets in excess demand in simple ascending auctions with unit-demand bidders. *Annals of Operations Research* 211, 1 (2013), 27–36. DOI : <https://doi.org/10.1007/s10479-013-1344-1>
- [3] Tommy Andersson and Albin Erlanson. 2013. Multi-item Vickrey–English–Dutch auctions. *Games and Economic Behavior* 81 (2013), 116–129. DOI : <https://doi.org/10.1016/j.geb.2013.05.001>
- [4] Lawrence M. Ausubel. 2004. An efficient ascending-bid auction for multiple objects. *American Economic Review* 94, 5 (2004), 1452–1475. DOI : <https://doi.org/10.1257/0002828043052330>
- [5] Lawrence M. Ausubel. 2005. Walrasian tatonnement for discrete goods. *Unpublished Paper* (2005).

- [6] Lawrence M. Ausubel. 2006. An efficient dynamic auction for heterogeneous commodities. *American Economic Review* 96, 3 (2006), 602–629. DOI : <https://doi.org/10.1257/aer.96.3.602>
- [7] Lawrence M. Ausubel and Paul R. Milgrom. 2002. Ascending auctions with package bidding. *The BE Journal of Theoretical Economics* 1, 1 (2002), 1–42. DOI : <https://doi.org/10.2202/1534-5963.1019>
- [8] Oren Ben-Zwi. 2017. Walrasian’s characterization and a universal ascending auction. *Games and Economic Behavior* 104 (2017), 456–467. DOI : <https://doi.org/10.1016/j.geb.2017.05.013>
- [9] Oren Ben-Zwi, Ron Lavi, and Ilan Newman. 2013. Ascending auctions and walrasian equilibrium. *arXiv preprint* (2013). DOI : <https://doi.org/10.48550/arXiv.1301.1153> arXiv:1301.1153
- [10] Dimitri P. Bertsekas. 1992. Auction algorithms for network flow problems: A tutorial introduction. *Computational Optimization and Applications* 1, 1 (1992), 7–66. DOI : <https://doi.org/10.1007/BF00247653>
- [11] Sushil Bikhchandani, Sven de Vries, James Schummer, and Rakesh V. Vohra. 2001. Linear programming and Vickrey auctions. *IMA Volumes in Mathematics and its Applications* 127 (2001), 75–116.
- [12] Sushil Bikhchandani, Sven de Vries, James Schummer, and Rakesh V. Vohra. 2011. An ascending vickrey auction for selling bases of a matroid. *Operations Research* 59, 2 (2011), 400–413. DOI : <https://doi.org/10.1287/opre.1100.0888>
- [13] Sushil Bikhchandani and Joseph M. Ostroy. 2002. The package assignment model. *Journal of Economic Theory* 107, 2 (2002), 377–406. DOI : <https://doi.org/10.1006/jeth.2001.2957>
- [14] Joakim Blikstad, Sagnik Mukhopadhyay, Danupon Nanongkai, and Ta-Wei Tu. 2023. Fast algorithms via dynamic-oracle matroids. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 1229–1242. DOI : <https://doi.org/10.1145/3564246.3585219>
- [15] Liad Blumrosen and Noam Nisan. 2010. On the computational power of demand queries. *SIAM J. Comput.* 39, 4 (2010), 1372–1391. DOI : <https://doi.org/10.1137/050641181>
- [16] Ozan Candogan, Asuman Ozdaglar, and Pablo A. Parrilo. 2015. Iterative auction design for tree valuations. *Operations Research* 63, 4 (2015), 751–771. DOI : <https://doi.org/10.1287/opre.2015.1388>
- [17] Deeparnab Chakrabarty, Yin Tat Lee, Aaron Sidford, Sahil Singla, and Sam Chiu-Wai Wong. 2019. Faster matroid intersection. In *Proceedings of the 2019 IEEE 60th Annual Symposium on Foundations of Computer Science*. IEEE, 1146–1168. DOI : <https://doi.org/10.1109/FOCS.2019.00072>
- [18] Edward H. Clarke. 1971. Multipart pricing of public goods. *Public Choice* 11, 1 (1971), 17–33. Retrieved from <https://www.jstor.org/stable/30022651>
- [19] Natalie Collina and S. Matthew Weinberg. 2020. On the (in-) approximability of bayesian revenue maximization for a combinatorial buyer. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 477–497. DOI : <https://doi.org/10.1145/3391403.3399496>
- [20] Peter Cramton. 1998. Ascending auctions. *European Economic Review* 42, 3 (1998), 745–756. DOI : [https://doi.org/10.1016/S0014-2921\(97\)00122-0](https://doi.org/10.1016/S0014-2921(97)00122-0)
- [21] William H. Cunningham. 1986. Improved bounds for matroid partition and intersection algorithms. *SIAM J. Comput.* 15, 4 (1986), 948–957. DOI : <https://doi.org/10.1137/0215066>
- [22] Sven de Vries, James Schummer, and Rakesh V. Vohra. 2007. On ascending Vickrey auctions for heterogeneous objects. *Journal of Economic Theory* 132, 1 (2007), 95–118. DOI : <https://doi.org/10.1016/j.jet.2005.07.010>
- [23] Gabrielle Demange, David Gale, and Marilda Sotomayor. 1986. Multi-item auctions. *Journal of Political Economy* 94, 4 (1986), 863–872. DOI : <https://doi.org/10.1086/261411>
- [24] Andreas W. M. Dress and Werner Terhalle. 1995. Well-layered maps – A class of greedily optimizable set functions. *Applied Mathematics Letters* 8, 5 (1995), 77–80. DOI : [https://doi.org/10.1016/0893-9659\(95\)00070-7](https://doi.org/10.1016/0893-9659(95)00070-7)
- [25] Jack Edmonds. 1968. Matroid partition. *Mathematics of the Decision Sciences: Part 1*, G. B. Dantzig and A. F. Veinott (Eds.). American Mathematical Society, 335–345.
- [26] Katharina Eickhoff, S. Thomas McCormick, Britta Peis, Niklas Rieken, and Laura Vargas Koch. 2024. A flow-based ascending auction to compute buyer-optimal Walrasian prices. *Networks* 84, 2 (2024), 161–180. DOI : <https://doi.org/10.1002/net.22218>
- [27] Katharina Eickhoff, Britta Peis, Niklas Rieken, Laura Vargas Koch, and László A. Végh. 2024. Faster ascending auctions via polymatroid sum. In *Web and Internet Economics, Springer Nature Switzerland, Cham*, 682–683. DOI : <https://doi.org/10.1007/978-3-031-48974-7>
- [28] András Frank and Zoltán Miklós. 2012. Simple push-relabel algorithms for matroids and submodular flows. *Japan Journal of Industrial and Applied Mathematics* 29, 3 (2012), 419–439. DOI : <https://doi.org/10.1007/s13160-012-0076-y>
- [29] Satoru Fujishige. 2005. *Submodular Functions and Optimization*. Vol. 58. Elsevier.
- [30] Satoru Fujishige and Zaifu Yang. 2003. A note on Kelso and Crawford’s gross substitutes condition. *Mathematics of Operations Research* 28, 3 (2003), 463–469. DOI : <https://doi.org/10.1287/moor.28.3.463.16393>
- [31] Satoru Fujishige and Xiaodong Zhang. 1992. New algorithms for the intersection problem of submodular systems. *Japan Journal of Industrial and Applied Mathematics* 9, 3 (1992), 369–382. DOI : <https://doi.org/10.1007/BF03167272>
- [32] Satoru Fujishige and Xiaodong Zhang. 1996. A Push/Relabel framework for submodular flows and its definement for 0-1 submodular flows. *Optimization* 38, 2 (1996), 133–154. DOI : <https://doi.org/10.1080/02331939608844244>

- [33] Andrew V. Goldberg and Robert E. Tarjan. 1988. A new approach to the maximum-flow problem. *Journal of the ACM* 35, 4 (1988), 921–940. DOI: <https://doi.org/10.1145/48014.61051>
- [34] Curtis Greene and Thomas L. Magnanti. 1975. Some abstract pivot algorithms. *SIAM J. Appl. Math.* 29, 3 (1975), 530–539. DOI: <https://doi.org/10.1137/0129045>
- [35] Martin Grötschel, László Lovász, and Alexander Schrijver. 1981. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica* 1, 2 (1981), 169–197. DOI: <https://doi.org/10.1007/BF02579273>
- [36] Theodore Groves. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* 41, 4 (1973), 617–631. DOI: <https://doi.org/10.2307/1914085>
- [37] Faruk Gul and Ennio Stacchetti. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic theory* 87, 1 (1999), 95–124. DOI: <https://doi.org/10.1006/jeth.1999.2531>
- [38] Faruk Gul and Ennio Stacchetti. 2000. The English auction with differentiated commodities. *Journal of Economic Theory* 92, 1 (2000), 66–95. DOI: <https://doi.org/10.1006/jeth.1999.2580>
- [39] Satoru Iwata. 2003. A faster scaling algorithm for minimizing submodular functions. *SIAM J. Comput.* 32, 4 (2003), 833–840. DOI: <https://doi.org/10.1137/S0097539701397813>
- [40] Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. 2001. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *Journal of the ACM* 48, 4 (2001), 761–777. DOI: <https://doi.org/10.1145/502090.502096>
- [41] Satoru Iwata and James B. Orlin. 2009. A simple combinatorial algorithm for submodular function minimization. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1230–1237. DOI: <https://doi.org/10.1137/1.9781611973068.133>
- [42] Alexander S. Jr. Kelso and Vincent P. Crawford. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society* 50, 6 (1982), 1483–1504. DOI: <https://doi.org/10.2307/1913392>
- [43] Walter Kern, Bodo Manthey, and Marc Uetz. 2016. Note on VCG vs. price raising for matching markets. *arXiv preprint* (2016). DOI: <https://doi.org/10.48550/arXiv.1604.04157>
- [44] Donald E. Knuth. 1973. *Matroid Partitioning*. Computer Science Department, Stanford University.
- [45] Donald E. Knuth. 1974. The asymptotic number of geometries. *Journal of Combinatorial Theory, Series A* 16, 3 (1974), 398–400. DOI: [https://doi.org/10.1016/0097-3165\(74\)90063-6](https://doi.org/10.1016/0097-3165(74)90063-6)
- [46] Colin J. H. McDiarmid. 1975. Rado’s theorem for polymatroids. In *Mathematical Proceedings of the Cambridge Philosophical Society*. Cambridge University Press, 263–281. DOI: <https://doi.org/10.1017/S0305004100051677>
- [47] Paul Milgrom and Bruno Strulovici. 2009. Substitute goods, auctions, and equilibrium. *Journal of Economic Theory* 144, 1 (2009), 212–247. DOI: <https://doi.org/10.1016/j.jet.2008.05.002>
- [48] Debasis Mishra and David C. Parkes. 2007. Ascending price Vickrey auctions for general valuations. *Journal of Economic Theory* 132, 1 (2007), 335–366. DOI: <https://doi.org/10.1016/j.jet.2005.09.004>
- [49] Debasis Mishra and Dolf Talman. 2010. Characterization of the Walrasian equilibria of the assignment model. *Journal of Mathematical Economics* 46, 1 (2010), 6–20. DOI: <https://doi.org/10.1016/j.jmateco.2009.05.008>
- [50] Kazuo Murota. 2003. *Discrete Convex Analysis*. Society for Industrial and Applied Mathematics (SIAM). DOI: <https://doi.org/10.1137/1.9780898718508>
- [51] Kazuo Murota, Akiyoshi Shioura, and Zaifu Yang. 2013. Computing a Walrasian equilibrium in iterative auctions with multiple differentiated items. Retrieved April 24, 2025 from <https://web.archive.org/web/20250424105153/https://www.keisu.tu-tokyo.ac.jp/data/2013/METR13-10.pdf>
- [52] Kazuo Murota, Akiyoshi Shioura, and Zaifu Yang. 2016. Time bounds for iterative auctions: A unified approach by discrete convex analysis. *Discrete Optimization* 19 (2016), 36–62. DOI: <https://doi.org/10.1016/j.disopt.2016.01.001>
- [53] Renato Paes Leme. 2017. Gross substitutability: An algorithmic survey. *Games and Economic Behavior* 106 (2017), 294–316. DOI: <https://doi.org/10.1016/j.geb.2017.10.016>
- [54] Renato Paes Leme and Sam Chiu-wai Wong. 2020. Computing Walrasian equilibria: Fast algorithms and structural properties. *Mathematical Programming* 179, 1 (2020), 343–384. DOI: <https://doi.org/10.1007/s10107-018-1334-9>
- [55] David C. Parkes and Lyle H. Ungar. 2000. Iterative combinatorial auctions: Theory and practice. In *Proceedings of the AAAI Conference on Artificial Intelligence, Association for the Advancement of Artificial Intelligence*, Washington, DC 20004, USA, 74–82.
- [56] Stephen Raach. 2024. Monotonicity of extremal walrasian prices in gross substitutes markets. *Operations Research Letters* 56 (2024), 107149. DOI: <https://doi.org/10.1016/j.orl.2024.107149>
- [57] Paul Schönsleben. 1980. *Ganzzahlige polymatroid-intersektions-algorithmen*. Ph.D. Dissertation. ETH Zürich.
- [58] Alexander Schrijver. 2003. *Combinatorial Optimization: Polyhedra and Efficiency*. Vol. 24. Springer.
- [59] Akiyoshi Shioura. 2017. Algorithms for L-convex function minimization: Connection between discrete convex analysis and other research fields. *Journal of the Operations Research Society of Japan* 60, 3 (2017), 216–243. DOI: <https://doi.org/10.15807/jorsj.60.216>
- [60] Akiyoshi Shioura and Akihisa Tamura. 2015. Gross substitutes condition and discrete concavity for multi-unit valuations: a survey. *Journal of the Operations Research Society of Japan* 58, 1 (2015), 61–103. DOI: <https://doi.org/10.15807/jorsj.58.61>

- [61] Marilda Sotomayor. 2002. A simultaneous descending bid auction for multiple items and unitary demand. *Revista Brasileira de Economia* 56, 3 (2002), 497–510. DOI: <https://doi.org/10.1590/S0034-71402002000300006>
- [62] Ning Sun and Zaifu Yang. 2009. A double-track adjustment process for discrete markets with substitutes and complements. *Econometrica* 77, 3 (2009), 933–952. DOI: <https://doi.org/10.3982/ECTA6514>
- [63] Ning Sun and Zaifu Yang. 2014. An efficient and incentive compatible dynamic auction for multiple complements. *Journal of Political Economy* 122, 2 (2014), 422–466. DOI: <https://doi.org/10.1086/674550>
- [64] Tatsuya Terao. 2025. Faster matroid partition algorithms. *ACM Trans. Algorithms* 21, 2 (2025). DOI: <https://doi.org/10.1145/3707208>
- [65] Hal R. Varian. 1981. Dynamical systems with applications to economics. *Handbook of Mathematical Economics* 1 (1981), 93–110. DOI: [https://doi.org/10.1016/S1573-4382\(81\)01007-2](https://doi.org/10.1016/S1573-4382(81)01007-2)
- [66] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance* 16, 1 (1961), 8–37. DOI: <https://doi.org/10.2307/2977633>
- [67] Léon Walras. 1874. *Éléments d'économie Politique Pure ou théorie de la Richesse Sociale*. F. Rouge.

Appendices

A Examples

A.1 Walrasian Prices May Not Exist if Valuations are Not Strong Gross Substitutes

Example A.1. Consider an example with two buyers $N = \{1, 2\}$ and three items $E = \{e_1, e_2, e_3\}$. The multiplicity of every item is one, i.e., $b(e) = 1$ for all $e \in E$. The valuation functions of the buyers are as follows (note that they are not strong gross substitutes):

$$v_1(\chi_S) = \begin{cases} 2 & \text{if } \{e_1, e_2\} \subseteq S, \\ 1 & \text{if } \{e_1, e_2\} \not\subseteq S, \{e_3\} \subseteq S, \\ 0 & \text{otherwise,} \end{cases} \quad v_2(\chi_S) = \begin{cases} 2 & \text{if } \{e_2, e_3\} \subseteq S, \\ 1 & \text{if } \{e_2, e_3\} \not\subseteq S, \{e_1\} \subseteq S, \\ 0 & \text{otherwise.} \end{cases}$$

Here $p = (0, 0, 0)$ is not packing since the unique minimal preferred bundle of buyer 1 is $(1, 1, 0)$, while the unique minimal preferred bundle of buyer 2 is $(0, 1, 1)$. The vector $p = (0, 1, 0)$ is packing but it is not covering, since item e_2 is not sold.

Assume that there is a packing and covering allocation z , i.e., $\sum_{i \in \{1, 2\}} z_i(e) = 1$ for all $e \in E$ and $z_i \in \mathcal{D}_i(p)$ for $i \in \{1, 2\}$.

- In case that $z_1 = (0, 0, 0)$, it holds that $p(e_1) + p(e_2) \geq 2$ and $p(e_3) \geq 1$, since $z_1 \in \mathcal{D}_1(p)$. So, the utility for the second buyer is $v_2(\chi_E) - p(E) \leq 2 - 3 = -1$ and thus, $\chi_E = (1, 1, 1) \notin \mathcal{D}_2(p)$.
- In case that $z_1 = (1, 0, 0) \in \mathcal{D}_1(p)$, the utility of buyer 1 needs to be zero and $p(e_1) = 0$, $p(e_2) \geq 2$ and $p(e_3) \geq 1$. Hence, the second buyer receives items e_2, e_3 with utility $v_2(\chi_{\{e_2, e_3\}}) - p(\{e_2, e_3\}) \leq 2 - 3 = -1$ which is a contradiction to $\chi_{\{e_2, e_3\}} \in \mathcal{D}_2(p)$.
- If $z_1 = (0, 1, 0) \in \mathcal{D}_1(p)$, the utility of buyer 1 is zero again and $p(e_1) \geq 2$, $p(e_2) = 0$ and $p(e_3) \geq 1$. Then, buyer 2 does not receive a preferred bundle since the utility of $\chi_{\{e_1, e_3\}} = (1, 0, 1)$ is given by $v_2(\chi_{\{e_1, e_3\}}) - p(\{e_1, e_3\}) \leq 1 - 2 = -2$.
- In case that $z_1 = (0, 0, 1) \in \mathcal{D}_1(p)$, we know that $p(e_3) \leq 1$ and $p(e_1) + p(e_2) \geq p(e_3) + 1$. Since $\chi_{\{e_1, e_2\}} = z_2 \in \mathcal{D}_2(p)$, it holds that $p(e_2) = 0$ and $v_2(\chi_{\{e_1, e_2\}}) - p(\{e_1, e_2\}) \geq v_2(\chi_{\{e_2, e_3\}}) - p(\{e_2, e_3\})$ which is equivalent to $p(e_1) + 1 \leq p(e_3)$. This is a contradiction to $p(e_1) \geq p(e_3) + 1$.

In case that $\|z_1\|_1 \geq 2$, we obtain the same contradictions by using $\|z_2\|_1 \leq 1$ and the symmetry of the valuations. To sum up, there exists no Walrasian price vector p in this market.

A.2 Packing Vectors do not form a Lattice Given Non Strong Gross Substitutes Valuations

In the following, we show that the component-wise minimal vector of packing prices is not necessarily packing if valuations are not strong gross substitutes. Note that Walrasian prices are not guaranteed to exist if buyer valuations are not all strong gross substitutes [42]. Buyer-optimal packing prices are neither unique nor covering in the following example.

Example A.2. Given a market with four items $E = \{e_1, e_2, e_3, e_4\}$ where $b(e) = 1$ for each $e \in E$ and two buyers. The valuations for the first buyer are given by the item-wise valuations $\tilde{v}_1 = (6, 6, 6, 10)$ and the independent sets $\mathcal{I}_1 = \{I \subseteq F \mid F \in \{\{e_1, e_2, e_3\}, \{e_4\}\}\}$ (i.e., not a matroid). Thus, the valuation of a bundle χ_S is given by $v_1(S) = \max_{I \in \mathcal{I}_1} \sum_{e \in I \cap S} \tilde{v}_1(e)$. The valuation function for the second buyer is defined similarly using item-wise valuations $\tilde{v}_2 = (10, 6, 6, 6)$ and the independent sets $\mathcal{I}_2 = \{I \subseteq F \mid F \in \{\{e_1\}, \{e_2, e_3, e_4\}\}\}$.

First, we show that there is no packing price vector p with $\sum_{e \in E} p(e) < 8$. Consider such a vector p and let E' be all elements in $e \in E$ with $p(e) < 6$. By definition, $|E \setminus E'| \leq 1$. The unique preferred bundle for buyer 1 is $\{e_1, e_2, e_3\} \cap E'$ and for buyer 2, it is $\{e_2, e_3, e_4\} \cap E'$. Since $|\{e_1, e_2, e_3\} \cap \{e_2, e_3, e_4\}| = 2$ the intersection of the preferred bundles is non-empty, so p cannot be a packing price vector.

Moreover, there are packing price vectors with $\sum_{e \in E} p(e) = 8$, namely $p = (0, 2 + \alpha, 6 - \alpha, 0)$ with $\alpha \in [0, 4]$ are packing price vectors, since the first buyer is indifferent between $\{e_1, e_2, e_3\}$ and $\{e_4\}$ and the second one is indifferent between $\{e_1\}$ and $\{e_2, e_3, e_4\}$. Thus, the allocation $\{e_4\}$ to the first buyer and $\{e_1\}$ to the second buyer is packing. This family of packing prices gives a component-wise minimal vector $(0, 2, 2, 0)$ which itself is not packing. However, since $\sum_{e \in E} p(e) = 8$ for all price vectors in the defined family, they are all buyer-optimal packing prices. Hence, clearly, the buyer-optimal packing prices are not unique and moreover these buyer-optimal packing price vectors are all not covering.

Note that the valuations functions in this example are not strong gross substitutes (in contrast to weighted matroid rank valuations the independent sets do not fulfill the exchange property and thus, do not form a matroid). Consider the valuation function of the first buyer. It holds that $\{e_1, e_2, e_3\} \in D_1(0)$. But if we increase the prices of items e_1 and e_2 by 6, i.e., $p = (6, 6, 0, 0)$, there is no demand for e_3 anymore since $D_1(p) = \{\{e_4\}\}$. Similarly, the valuation function of the second buyer is also not strong gross substitutes.

A.3 Covering Vectors do not form a Lattice Given Non Strong Gross Substitutes Valuations

Here we give an example that the join of two packing price vectors is not necessarily packing if the valuation functions are not gross substitutes.

Example A.3. Given a market with three items $E = \{e_1, e_2, e_3\}$ and unit-supply, i.e. $b(e) = 1$ for each $e \in E$, and two buyers. Both buyers have the same valuation function, given by the following table:

S	\emptyset	$\{e_1\}$	$\{e_2\}$	$\{e_3\}$	$\{e_1, e_2\}$	$\{e_1, e_3\}$	$\{e_2, e_3\}$	$\{e_1, e_2, e_3\}$
$v_i(S)$	0	7	7	8	14	13	13	18

Then $p = (6, 7, 7)$ is covering, since $\mathcal{D}_i(p) = \{\{e_1\}, \{e_1, e_2\}, \{e_3\}\}$. Moreover, $p' = (6, 7, 7)$ is then covering with $\mathcal{D}_i(p) = \{\{e_2\}, \{e_1, e_2\}, \{e_3\}\}$. However, $q = p \vee p' = (7, 7, 7)$ is not covering, as $\mathcal{D}_i(q) = \{\{e_3\}\}$.

Note there is a Walrasian price vector above all these price vectors, which is $p^* = (7, 7, 8)$.

A.4 Maximal Walrasian Prices are not Monotone with Respect to Changes in the Valuations

As already seen for minimal Walrasian prices, the maximal Walrasian prices do not change in a monotone way when the valuations are changed:

Example A.4. Let there be four items $E = \{e_1, e_2, e_3\}$ and three buyers with unit-demand valuations i.e., $v_i(S) = \max_{e \in S} v_{ie}$. The valuations per item and the corresponding maximal Walrasian prices can be found in the following table:

valuations buyer 1 ($v_{11}, v_{12}, v_{13}, v_{14}$)	valuations buyer 2 ($v_{21}, v_{22}, v_{23}, v_{24}$)	valuations buyer 3 ($v_{31}, v_{32}, v_{33}, v_{34}$)	max Walrasian price p^*
(0, 9, 1, 1)	(6, 10, 0, 0)	(4, 0, 1, 1)	(4, 8, 0, 0)
(0, 9, 2, 1)	(6, 10, 0, 0)	(4, 0, 1, 1)	(3, 7, 0, 0)
(0, 9, 2, 2)	(6, 10, 0, 0)	(4, 0, 1, 1)	(3, 7, 0, 0)
(0, 10, 2, 2)	(6, 10, 0, 0)	(4, 0, 1, 1)	(4, 8, 0, 0)

Hence, also for the maximal Walrasian prices, the price can either stay the same, go up or go down, i.e., there is no monotone behavior.

A.5 Example of Running Time of Different Oracle Calls

Let us consider an OXS valuation function v_i in the unit-supply case. This is described by a weighted bipartite graph $(E \cup T_i, H_i, w^{(i)})$. For a vector $z_i \in \{0, 1\}^E$, we let $S_i = \text{supp}(z_i)$. Then, $v_i(z_i)$ is the maximum weight of a matching between the sets S_i and T_i .

Thus, the time DO needed to find a set $z_i \in \check{\mathcal{D}}_i(p)$ amounts to the time needed to compute a maximum weight matching. This can be done using the Hungarian algorithm; better algorithms exist if the graph is not dense, or in the weakly polynomial regime. However, given $z_i \in \check{\mathcal{D}}_i(p)$ along with an optimal matching and an optimal dual solution, the time ExO of a $w_i(e, f)$ query is the time of finding a single augmenting path.

A.6 Example of an Exchange Graph

Suppose there are six items $E = \{e_1, e_2, e_3, e_4, e_5, e_6\}$, every $e \in E$ is available in quantity $b(e) = 2$, and there are three buyers $N = \{\text{Blue}, \text{Red}, \text{Green}\}$.

Further suppose at some stage of the auction, the allocation $\mathbf{z} = (z_B, z_R, z_G)$ is as follows:

$$z_B = (2, 2, 2, 1, 0, 0),$$

$$z_R = (2, 1, 0, 0, 1, 1),$$

$$z_G = (0, 2, 0, 1, 1, 0).$$

Thus, the items $\{e_1, e_2\}$ are oversold and $\{e_6\}$ is undersold at current prices. The auctioneer also queries each buyer $i \in N$ for the values $w_i(e, f)$ for each $e, f \in E$ with $z_i(f) > 0$ and $z_i(e) < b(e)$. Suppose the buyers answer with the following non-zeros:

$$w_B(e_4, e_1) = 1, \quad w_R(e_3, e_1) = 2, \quad w_G(e_1, e_5) = 1,$$

$$w_B(e_4, e_2) = 1, \quad w_R(e_3, e_2) = 1, \quad w_G(e_3, e_4) = 1,$$

$$w_B(e_4, e_3) = 1, \quad w_R(e_5, e_6) = 1, \quad w_G(e_4, e_2) = 1,$$

$$w_G(e_4, e_5) = 1.$$

Then the exchange graph looks as depicted as in Figure 2. Note that there is no path from an undersold to an oversold item (otherwise the polymatroid sum problem would not have been solved to optimality). The set of items from which we can reach an oversold item is $\{e_1, e_2, e_3, e_4\}$, which then by Theorem 3.3 is a minimal maximally overdemand set (i.e., not just the oversold items themselves but also suitable substitutes for them).

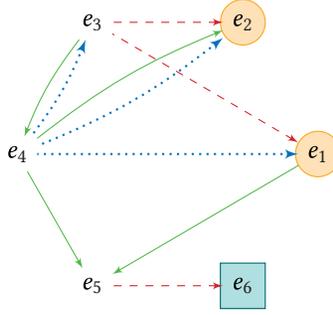


Fig. 2. Exchange graph $G(z) = (E, A)$ where blue (dotted) lines indicate A_B , red (dashed) lines indicate A_R , and green (solid) lines indicate A_G . Oversold items are indicated by orange cycles, undersold items are indicated by turquoise squares.

B Missing Proofs

B.1 Packing, Covering, and Walrasian

A Walrasian price vector is packing and covering. Now, we show the other direction, i.e., that p is Walrasian if it is packing and covering.

LEMMA B.1. *Let p be a price vector in an auction with strong gross substitutes valuations that is both packing and covering. Then p is Walrasian.*

PROOF. Let $\mathbf{z} = (z_1, \dots, z_n)$ be a covering allocation with respect to p . Further, let $\mathbf{y} = (y_1, \dots, y_n)$ be a packing allocation with respect to p that, among all such collections, minimizes $\sum_{i \in N} \sum_{e \in E} |z_i(e) - y_i(e)|$. Note that such allocations exist as p is both packing and covering. We will show that we have $\sum_{i \in N} y_i(e) = b(e)$ for every $e \in E$, i.e., \mathbf{y} is also covering and thus p is Walrasian.

Assume toward a contradiction that this is not the case, i.e., there is an $e \in E$ such that $\sum_{i \in N} y_i(e) \leq b(e) - 1$ (where we used that b and all y_i are integral). As $b(e) \leq \sum_{i \in N} z_i(e)$, we know that there is a $j \in N$ such that $y_j(e) < z_j(e)$. As v_j is gross substitutes, it is M^{\natural} -concave by Lemma 2.6. Hence, u_j is M^{\natural} -concave too, as the difference of an M^{\natural} -concave function and a modular function. As $e \in \text{supp}^+(z_j - y_j)$, we can infer that one of the following holds:

(M1) $u_j(z_j) + u_j(y_j) \leq u_j(z_j - \chi_e) + u_j(y_j + \chi_e)$, or

(M2) there exists $f \in \text{supp}^-(z_j - y_j)$ such that $u_j(y_j) + u_j(z_j) \leq u_j(z_j - \chi_e + \chi_f) + u_j(y_j + \chi_e - \chi_f)$.

First, assume that (M1) holds true. Then $y_j, z_j \in \mathcal{D}_j(p)$ implies that both $z_j - \chi_e$ and $y_j + \chi_e$ are contained in $\mathcal{D}_j(p)$ as well.

Define an allocation $\mathbf{y}' = (y'_1, \dots, y'_n)$ by setting $y'_i = y_i$ for all $i \in N \setminus \{j\}$ and $y'_j = y_j + \chi_e$. Then $y'_i \in \mathcal{D}_i(p)$ holds for all $i \in N$. Furthermore, we have $\sum_{i \in N} y'_i(g) = \sum_{i \in N} y_i(g) \leq b(g)$ for all $g \in E \setminus \{e\}$, and $\sum_{i \in N} y'_i(e) = \sum_{i \in N} y_i(e) + 1 \leq b(e)$. Moreover, we have $|z_i(g) - y'_i(g)| = |z_i(g) - y_i(g)|$ for all $(i, g) \in N \times E \setminus \{(j, e)\}$, and $|z_j(e) - y'_j(e)| = |z_j(e) - y_j(e)| - 1$ because $e \in \text{supp}^+(z_j - y_j)$. But this contradicts our assumption that \mathbf{y} minimizes $\sum_{i \in N} \sum_{g \in E} |z_i(g) - y_i(g)|$.

Next, assume that (M2) holds. Again, $y_j, z_j \in \mathcal{D}_j(p)$ allows us to conclude that both $z_j - \chi_e + \chi_f$ and $y_j + \chi_e - \chi_f$ are contained in $\mathcal{D}_j(p)$ as well. Define a collection of bundles \mathbf{y}' by setting $y'_i = y_i$ for all $i \in N \setminus \{j\}$, and $y'_j = y_j + \chi_e - \chi_f$. Then $y'_i \in \mathcal{D}_i(p)$ for all $i \in N$. Moreover, $\sum_{i \in N} y'_i(g) = \sum_{i \in N} y_i(g) \leq b(g)$ for all $g \in E \setminus \{e, f\}$, $\sum_{i \in N} y'_i(e) = \sum_{i \in N} y_i(e) + 1 \leq b(e)$ and $\sum_{i \in N} y'_i(f) = \sum_{i \in N} y_i(f) - 1 \leq b(f)$. Additionally, for $(i, g) \in N \times E \setminus \{(j, e), (j, f)\}$, we have $|z_i(g) - y'_i(g)| = |z_i(g) - y_i(g)|$. As $e \in \text{supp}^+(z_j - y_j)$ and $f \in \text{supp}^-(z_j - y_j)$, we further know

that $|z_j(e) - y'_j(e)| = |z_j(e) - y_j(e)| - 1$ and $|z_j(f) - y'_j(f)| = |z_j(f) - y_j(f)| - 1$. But again, this contradicts the fact that y minimizes $\sum_{i \in N} \sum_{e \in E} |z_i(e) - y_i(e)|$.

To sum up, we obtained a contradiction in both cases and hence, we have $\sum_{i \in N} y_i(e) = b(e)$ for all $e \in E$. Thus, y is packing and covering. We conclude that p is a Walrasian price vector. \square

B.2 Gross Substitutes Valuations Yield M-convex Demand Sets

LEMMA 2.8. *If v_i is strong gross substitutes, then $\check{\mathcal{D}}_i(p)$ and $\hat{\mathcal{D}}_i(p)$ are M-convex sets for any price vector $p \in \mathbb{R}_+^E$.*

PROOF. By Lemma 2.6, the valuation function v_i is M^h -concave. Thus, the utility function u_i is also M^h -concave as the sum of an M^h -concave and a modular function. Let $x, y \in \check{\mathcal{D}}_i(p)$ and let $e \in \text{supp}^+(x - y)$. Then, since u_i is M^h -concave, either $u_i(x) + u_i(y) \leq u_i(x - \chi_e) + u_i(y + \chi_e)$, or there is $f \in \text{supp}^-(x - y)$ with $u_i(x) + u_i(y) \leq u_i(x - \chi_e + \chi_f) + u_i(y + \chi_e - \chi_f)$. In the first case, $x - \chi_e, y + \chi_e \in \check{\mathcal{D}}_i(p)$, and in the second case, $x - \chi_e + \chi_f, y + \chi_e - \chi_f \in \mathcal{D}_i(p)$. The first case is impossible since $\check{\mathcal{D}}_i(p)$ is the set of inclusion wise minimal vectors in $\mathcal{D}_i(p)$, but $x - \chi_e \leq x$. Hence, the second case holds and it remains to show that there is no preferred bundle strictly included in $x' := x - \chi_e + \chi_f$ or $y' := y + \chi_e - \chi_f$.

W.l.o.g. consider x' and assume that there is a preferred bundle $z \neq x'$ with $z \leq x'$. Hence, there exist an item g with $z(g) < x'(g) \leq x(g)$. As x is a minimal preferred bundle, we must have $z(f) = x(f) + 1$, as otherwise $z \leq x$ as well. Moreover, $z(e) \leq x(e) - 1$ as $z \leq x'$. Hence, $f \in \text{supp}^+(z - x) = \{f\}$ and thus by M^h -concavity of u_i , we have $u_i(z) + u_i(x) \leq u_i(x + \chi_f) + u_i(z - \chi_f)$ or there exists an item $g \in \text{supp}^-(z - x)$ with $u_i(z) + u_i(x) \leq u_i(x + \chi_f - \chi_g) + u_i(z - \chi_f + \chi_g)$. By minimality of x , the first case is not possible since $z - \chi_f \leq x$ with $z - \chi_f \neq x$.

Thus, we can focus on the second case with $z - \chi_f + \chi_g \in \mathcal{D}_i(p)$. By definition of z and g , we know that $z - \chi_f + \chi_g \leq x$. Hence, by minimality of x it follows that $z - \chi_f + \chi_g = x$. As $z(e) < x(e)$, this implies $g = e$. However, this yields $z(e) + 1 \leq x'(e) < x(e)$, a contradiction.

Hence, $x - \chi_e + \chi_f, y + \chi_e - \chi_f \in \check{\mathcal{D}}_i(p)$ holds and we have shown that $\check{\mathcal{D}}_i(p)$ is an M-convex set.

The proof that $\hat{\mathcal{D}}_i(p)$ is an M-convex set follows the same lines. \square

B.3 Underdemanded Sets

LEMMA 3.5. *Given a market instance with strong gross substitutes valuations, prices p are covering if and only if there is no underdemanded set.*

PROOF. Observe that for any price vector p , we have

$$\max_{y \in \mathcal{D}(p)} \left\{ \sum_{e \in E} \min \left\{ \sum_{i \in N} y_i(e), b(e) \right\} \right\} \leq \sum_{e \in E} b(e). \quad (18)$$

Moreover, equality holds in (18) if and only if there is a covering allocation $z \in \hat{\mathcal{D}}(p)$, i.e., p is covering. Now, assume that p is covering. Using the Min-Max Theorem for polymatroid sum (Theorem 2.12), we get

$$\min_{T \subseteq E} \left\{ b(T) + \sum_{i \in N} \hat{p}(E \setminus T) \right\} = \max_{y \in \hat{\mathcal{D}}(p)} \left\{ \sum_{e \in E} \min \left\{ \sum_{i \in N} y_i(e), b(e) \right\} \right\} = \sum_{e \in E} b(e).$$

We obtain

$$\begin{aligned}
& b(T) + \sum_{i \in N} \hat{\rho}(E \setminus T) \geq b(E) && \text{for all } T \subseteq E \\
\Leftrightarrow & b(E \setminus S) + \sum_{i \in N} \hat{\rho}(S) \geq b(E) && \text{for all } S \subseteq E \\
\Leftrightarrow & 0 \geq \text{underd}^p(S) && \text{for all } S \subseteq E
\end{aligned}$$

i.e., there is no underdemanded set S .

Next, we show that there is a covering allocation if there is no underdemanded set. By the same arguments as above, we get

$$b(E) \leq \min_{T \subseteq E} \left\{ b(T) + \sum_{i \in N} \hat{\rho}(E \setminus T) \right\} = \max_{y \in \hat{\mathcal{D}}} \left\{ \sum_{e \in E} \min \left\{ \sum_{i \in N} y_i(e), b(e) \right\} \right\} \leq \sum_{e \in E} b(e), \quad (19)$$

where the first inequality holds as there is no underdemanded set S and the next equality holds by the Min-Max Theorem for polymatroid sum (Theorem 2.12). Thus, (19) holds with equality everywhere, especially for the last equation. But this is only possible if there is a covering allocation $z \in \hat{\mathcal{D}}$. Hence, the prices p are covering if there is no underdemanded set. \square

LEMMA 3.7. *Assume $\mathbf{z} = (z_1, \dots, z_n)$ is an optimal solution to (13). Then, a set S is maximally underdemanded if and only if it fulfills the following properties:*

- S is i -tight for every $i \in N$;
- S includes all undersold items;
- S does not include any oversold items.

Moreover, there is a unique inclusion-wise minimal set fulfilling this properties.

PROOF. Let S be a maximally underdemanded set and let $T := E \setminus S$. By Corollary 3.6, Theorem 2.12 and optimality of \mathbf{z} , we have

$$\sum_{e \in E} \min \left\{ \sum_{i \in N} z_i(e), b(e) \right\} = \sum_{i \in N} \hat{\rho}_i^p(E \setminus T) + b(T).$$

By Lemma 2.13, $S = E \setminus T$ satisfies all three properties stated in this lemma. Moreover, if two sets S and S' satisfy these three properties, then $S \cap S'$ and $S \cup S'$ also satisfy them. Hence, there exists a unique minimal set S satisfying these three properties. By Lemma 2.13, $E \setminus S$ is a minimizer of $T \mapsto \sum_{i \in N} \hat{\rho}_i^p(E \setminus T) + b(T)$, and for every other minimizer T' , we have $S \subseteq E \setminus T'$. Hence, S is the minimal maximally underdemanded set for prices p . \square

THEOREM 3.8. *The minimal maximally underdemanded set S is the set R of items which are reachable from an undersold item in $\hat{G}(\mathbf{z})$.*

PROOF. We first show that $R \subseteq S$. By Lemma 3.7, S includes all undersold items. Moreover, S is i -tight for every $i \in N$, which, by Lemma 2.14, means that whenever $e \in S$ and $(e, f) \in A$, we have $f \in S$ as well. Hence, $R \subseteq S$.

As S does not include any oversold item by Lemma 3.7, R does not include any oversold item, either. By definition, R contains all undersold items and R does not have any leaving arc in $\hat{G}(\mathbf{z})$. By Lemma 2.14, this implies that R is i -tight for every $i \in N$. Hence, R satisfies all three conditions listed in Lemma 3.7. Minimality of S , thus, allows us to conclude that $S = R$. \square

B.4 Maximally Overdemanded [Underdemanded] Sets are Steepest Descent Directions of the Lyapunov Function

Ben-Zwi [8] provided two alternative definitions for gross substitutes, adding to the standard one from Definition 2.2 and the one from discrete convexity (see [30] and [51]) as stated in Lemma 2.6.

Ben-Zwi's definition is in terms of change in a buyer's indirect utility function. We add this lemma here to have a similar definition for the multi-supply setting, which Ben-Zwi did not consider. The lemma holds in the same way as in the unit-supply setting by using Lemma 5.1.

LEMMA B.2. *Given a valuation function $v_i: [0, b]_{\mathbb{Z}} \rightarrow \mathbb{R}_+$, the following are equivalent:*

- (1) For all $p \in \mathbb{R}_+^E$ and $S \subseteq E$, $V_i(p) = V_i(p + \chi_S) + \check{\theta}_i^p(S)$.
- (2) For all $p \in \mathbb{R}_+^E$ and $S \subseteq E$, $V_i(p) = V_i(p - \chi_S) - \hat{\rho}_i^p(S)$.
- (3) The function v_i is non-decreasing and gross substitutes.

With this lemma, we can prove that the steepest descent directions coincide either with the maximally overdemanded sets or with the maximally underdemanded sets.

LEMMA B.3 (C.F. [8]). *Let $p \geq 0$ be a price vector, then*

$$\arg \min\{L(p + \chi_S) \mid S \subseteq E\} = \arg \max\{\text{overd}^p(S) \mid S \subseteq E\}.$$

PROOF OF LEMMA B.3. We show that for all $S, T \subseteq E$, we have

$$L(p + \chi_T) - L(p + \chi_S) = \text{overd}^p(S) - \text{overd}^p(T). \quad (20)$$

This implies the statement of the lemma because for a set $S^* \subseteq E$, we have

$$\begin{aligned} S^* \in \arg \min\{L(p + \chi_S) \mid S \subseteq E\} &\Leftrightarrow \forall T \subseteq E: L(p + \chi_T) - L(p + \chi_{S^*}) \geq 0 \\ &\stackrel{(20)}{\Leftrightarrow} \forall T \subseteq E: \text{overd}^p(S^*) - \text{overd}^p(T) \geq 0 \Leftrightarrow S^* \in \arg \max\{\text{overd}^p(S) \mid S \subseteq E\}. \end{aligned}$$

To verify (20), we calculate

$$\begin{aligned} &L(p + \chi_T) - L(p + \chi_S) \\ &= \sum_{i \in N} V_i(p + \chi_T) + \langle p + \chi_T, b \rangle - \sum_{i \in N} V_i(p + \chi_S) - \langle p + \chi_S, b \rangle \\ &= \sum_{i \in N} (V_i(p + \chi_T) - V_i(p + \chi_S)) + \langle p + \chi_T, b \rangle - \langle p + \chi_S, b \rangle \\ &= \sum_{i \in N} \left((V_i(p) - \check{\theta}_i^p(T)) - (V_i(p) - \check{\theta}_i^p(S)) \right) + \langle \chi_T - \chi_S, b \rangle \\ &= \sum_{i \in N} (\check{\theta}_i^p(S) - \check{\theta}_i^p(T)) + \langle \chi_T - \chi_S, b \rangle \\ &= \left(\sum_{i \in N} \check{\theta}_i^p(S) - \langle \chi_S, b \rangle \right) - \left(\sum_{i \in N} \check{\theta}_i^p(T) - \langle \chi_T, b \rangle \right) \\ &= \text{overd}^p(S) - \text{overd}^p(T). \quad \square \end{aligned}$$

Analogously, we can also show the following result.

LEMMA B.4 (C.F. [8]). *Let $p \geq 0$ be a price vector, then*

$$\arg \min\{L(p - \chi_S) \mid S \subseteq E\} = \arg \max\{\text{underd}^p(S) \mid S \subseteq E\}.$$

PROOF OF LEMMA B.4. We show that for all $S, T \subseteq E$, we have

$$L(p - \chi_T) - L(p - \chi_S) = \text{underd}^p(S) - \text{underd}^p(T). \quad (21)$$

This implies the statement of the lemma because for a set $S^* \subseteq E$, we have

$$S^* \in \arg \min \{L(p - \chi_S) \mid S \subseteq E\} \Leftrightarrow \forall T \subseteq E: L(p - \chi_T) - L(p - \chi_{S^*}) \geq 0$$

$$\stackrel{(21)}{\Leftrightarrow} \forall T \subseteq E: \text{underd}^p(S^*) - \text{underd}^p(T) \geq 0 \Leftrightarrow S^* \in \arg \max \{\text{underd}^p(S) \mid S \subseteq E\}.$$

To verify (21), we calculate

$$\begin{aligned} & L(p - \chi_T) - L(p - \chi_S) \\ &= \sum_{i \in N} V_i(p - \chi_T) + \langle p - \chi_T, b \rangle - \sum_{i \in N} V_i(p - \chi_S) - \langle p - \chi_S, b \rangle \\ &= \sum_{i \in N} (V_i(p - \chi_T) - V_i(p - \chi_S)) + \langle p - \chi_T, b \rangle - \langle p - \chi_S, b \rangle \\ &= \sum_{i \in N} \left((V_i(p) + \hat{\rho}_i^p(T)) - (V_i(p) + \hat{\rho}_i^p(S)) \right) + \langle \chi_S - \chi_T, b \rangle \\ &= \sum_{i \in N} (\hat{\rho}_i^p(T) - \hat{\rho}_i^p(S)) + \langle \chi_S - \chi_T, b \rangle \\ &= \left(\langle \chi_S, b \rangle - \sum_{i \in N} \hat{\rho}_i^p(S) \right) - \left(\langle \chi_T, b \rangle - \sum_{i \in N} \hat{\rho}_i^p(T) \right) \\ &= \text{underd}^p(S) - \text{underd}^p(T). \quad \square \end{aligned}$$

Received 28 November 2023; revised 26 March 2025; accepted 26 March 2025