

RESEARCH ARTICLE

WILEY

A flow-based ascending auction to compute buyer-optimal Walrasian prices

Katharina Eickhoff¹ | S. Thomas McCormick² | Britta Peis¹ |

Niklas Rieken¹ | Laura Vargas Koch³

¹School of Business and Economics, RWTH Aachen University, Aachen, Germany

²Sauder School of Business, University of British Columbia, Vancouver, Canada

³Department of Mathematics, ETH Zurich, Zurich, Switzerland

Correspondence

Katharina Eickhoff, School of Business and Economics, RWTH Aachen University, Aachen, Germany.

Email: eickhoff@oms.rwth-aachen.de

Funding information

K. Eickhoff is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 2236/2.

Abstract

We consider a market where a set of objects is sold to a set of buyers, each equipped with a valuation function for the objects. The goal of the auctioneer is to determine reasonable prices together with a stable allocation. One definition of “reasonable” and “stable” is a Walrasian equilibrium, which is a tuple consisting of a price vector together with an allocation satisfying the following desirable properties: (i) the allocation is market-clearing in the sense that as much as possible is sold, and (ii) the allocation is stable in the sense that every buyer ends up with an optimal set with respect to the given prices. Moreover, “buyer-optimal” means that the prices are smallest possible among all Walrasian prices. In this paper, we present a combinatorial network flow algorithm to compute buyer-optimal Walrasian prices in a multi-unit matching market with truncated additive valuation functions. The algorithm can be seen as a generalization of the classical housing market auction and mimics the very natural procedure of an ascending auction. We use our structural insights to prove monotonicity of the buyer-optimal Walrasian prices with respect to changes in supply or demand.

KEYWORDS

ascending auctions, flow-algorithm, matching markets, network-flow model, price monotonicity, Walrasian equilibrium

1 | INTRODUCTION

We consider a market where m different indivisible types of *objects* $\Omega = \{i_1, \dots, i_m\}$ are sold to n *buyers* $N = \{j_1, \dots, j_n\}$. Every object $i \in \Omega$ has a *supply* of $b_i \in \mathbb{Z}_+$ identical copies of that object. All buyers $j \in N$ have a *demand* $d_j \in \mathbb{Z}_+$, which is the maximum number of items they want to purchase. The goal of the auctioneer is to find a per-unit *price* $p(i)$ for each object $i \in \Omega$ together with an *allocation* $\mathbf{x} \in \mathbb{Z}_+^{\Omega \times N}$ of items to buyers such that the prices \mathbf{p} and the allocation \mathbf{x} satisfy certain desirable properties. Certainly, the allocation \mathbf{x} should be *feasible* in the sense that at most b_i units of each object $i \in \Omega$ are sold, and each buyer $j \in N$ gets awarded at most $d_j \in \mathbb{Z}_+$ items. We consider *truncated additive valuations* $v_j: \mathbb{Z}_+^\Omega \rightarrow \mathbb{R}_+$, based on a value $v_{ij} \in \mathbb{Z}_+$ that a buyer $j \in N$ has per unit of object $i \in \Omega$, and demand $d_j \in \mathbb{Z}_+$,

$$v_j(\mathbf{y}) := \max_{\mathbf{y}' \leq \mathbf{y}} \left\{ \sum_{i \in \Omega} v_{ij} y'_i \mid \sum_{i \in \Omega} y'_i \leq d_j \right\} \quad \text{for all } j \in N, \mathbf{y} \in \mathbb{Z}_+^\Omega.$$

The *utility* of one unit of object $i \in \Omega$ to buyer $j \in N$ is then given by $v_{ij} - p(i)$.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2024 The Authors. *Networks* published by Wiley Periodicals LLC.

A standard and natural equilibrium concept for markets are *Walrasian equilibria*, named after Léon Walras [22]. To properly define those equilibria we first need some more notation. We call a feasible allocation *stable*, when each buyer $j \in N$ gets a best possible allocation $\mathbf{x}_{\bullet j}$ given prices \mathbf{p} , in other words, every buyer is assigned to one of her *preferred bundles*, that is, a vector from the set

$$D_j(\mathbf{p}) = \arg \max_{\mathbf{x}_{\bullet j}} \left\{ \sum_{i \in \Omega} (v_{ij} - p(i))x_{ij} \mid \sum_{i \in \Omega} x_{ij} \leq d_j, 0 \leq x_{ij} \leq b_i, i \in \Omega \right\}. \quad (1)$$

If a price vector \mathbf{p} admits a stable allocation, then \mathbf{p} is called *competitive*. Note that competitive prices always exist as $p(i) = 1 + \max_{j \in N} v_{ij}$ for all $i \in \Omega$ is competitive with stable allocation $\mathbf{x} = \mathbf{0}$. However, these prices are not interesting. We want to find prices that are not only competitive but also *market-clearing*, that is, as much as possible should be sold, which is $\min\{\sum b_i, \sum d_j\}$ in our model. Prices that are competitive and market-clearing are called *Walrasian* and a pair (\mathbf{p}, \mathbf{x}) of Walrasian prices and supporting stable allocation is called *Walrasian equilibrium*. Walrasian prices where the prices are as low as possible, that is, where $\sum p(i)$ is minimum among all Walrasian prices, are called *buyer-optimal Walrasian prices* or *minimum Walrasian prices*.

The definitions above assume that the minimum selling price for an item is 0. If there are positive minimum selling prices, the notion of market-clearing changes. An allocation is market-clearing if everything with a price above the minimum selling price is sold. We can easily adapt our results for this more general case, but for simplicity we assume minimum selling prices to be 0 in the following.

Recall that in our model the buyers' valuations $v_j : \mathbb{Z}_+^\Omega \rightarrow \mathbb{R}$ belong to the class of truncated additive valuations. This is an important special case of *gross substitutes* [13], which form the biggest class of valuation functions that are in some sense well-behaved, which will be covered in Section 1.2. A fundamental property is that Walrasian prices are guaranteed to exist if all buyers have gross substitute valuations (see [1] for multi-supply and [13] for unit-supply settings) and that the (unique) buyer-optimal Walrasian price vector can be found using an ascending auction ([1, 11]).

Truncated additive valuations are a natural class of valuation functions. It will turn out that they allow for a simple flow-based algorithm to execute the auction. Moreover, truncated additive valuations generalize the special case of single-unit matching markets.

Single-unit matching markets. The model we consider generalizes the classical matching market model, in the literature also commonly referred to as housing market (see e.g., [6]). In those markets buyers have only unit-demand, that is, $d_j = 1$ for all $j \in N$, or more formally, $v_j(\mathbf{x}_{\bullet j}) = \max\{v_j(\chi_i) \mid i \in \Omega \text{ and } x_{ij} > 0\}$, and also each object is only available in one copy, that is, $b_i = 1$ for all $i \in \Omega$. Here, we use $\chi_X \in \{0, 1\}^\Omega$ to denote the characteristic vector of X , that is, the zero-one vector in $\{0, 1\}^\Omega$ with a 1 precisely in the entries corresponding to objects in X .

For this setting, a seminal paper by Demange et al. [5] describes an ascending auction which starts at the minimum possible selling price, that is, $p(i) = 0$ for all $i \in \Omega$, and then iteratively raises the prices on a set of items that are overdemanded (which is a set of neighbors of a Hall set, that is, a subset of buyers that are jointly interested in a set of items without enough supply for the demand in the subset) until the prices are market-clearing. If one only raises prices on an inclusion-wise minimal overdemanded set, they show that this process yields component-wise minimum (in fact unique) competitive prices, and that these prices are market-clearing, and hence, Walrasian. It is also well known that one can determine a socially optimal allocation together with buyer-optimal market-clearing prices with an adaption of the Hungarian algorithm, a primal-dual algorithm (see [19]). For a good overview on these topics, we refer to the book [6] or the short but self-contained paper by Kern et al. [14].

The duplication method. A naïve approach to reduce our more general multi-unit auction to a single-unit auction is via the following *duplication method*: we replace the b_i copies of object $i \in \Omega$ by b_i individual objects with unit-supply, and we replace a buyer j demanding d_j items by d_j unit-demand buyers, with the same valuations. Certainly, an ascending auction of the single-unit instance will return Walrasian prices, but these prices are in general not buyer-optimal:

Example 1. Consider one buyer $N = \{1\}$ with a demand of $d_1 = 2$, and two different items $\Omega = \{\alpha, \beta\}$ with a supply of $b_\alpha = b_\beta = 1$ which are valued differently by the sole buyer, say $v_{\bullet 1} = (5, 1)$. If we copy the buyer, there is no stable allocation if $p(\alpha) < 4$ as both copies compete for object α . If $p(\alpha) = 4$, both copies of the sole buyer are indifferent between the objects and thus, $\mathbf{p} = (4, 0)$ and $\mathbf{x} = ((1, 0), (0, 1))$ is a Walrasian equilibrium. However, considering the original situation, since the buyer is alone $\mathbf{p} = (0, 0)$ and $\mathbf{x} = (1, 1)$ is the buyer-optimal Walrasian equilibrium. That is, intuitively, the duplication method is oblivious to the fact that the two copies represent the same buyer and hence, the computed Hall sets are wrongly interpreted as competition between two buyers. Thus, the prices computed by the duplication method are not buyer-optimal.

As we just saw in the example above, an algorithm to compute buyer-optimal prices in single-unit matching markets does not trivially give an algorithm for buyer-optimal prices when buyers have truncated additive valuations. The good news is that these valuation functions still allow for a nice flow-based algorithm to compute overdemanded sets. This generalizes the method

to use a Hall set which works in the single-unit case and is more intuitive than the submodular function minimization methods used for strong gross substitute valuations.

1.1 | Our contribution

In this article, we provide a flow-based ascending auction for multi-unit markets where all buyers have truncated additive valuations. This auction enables us to compute buyer-optimal Walrasian prices and a stable allocation. It generalizes the matching-based ascending auction introduced by [5].

Although there already exist ascending auctions that compute the minimum Walrasian prices by for example, Ausubel [1] and Murota et al. [15], these algorithms strongly rely on the theory of submodular function minimization and discrete convexity to compute the set of objects on which prices should be increased in the ascending auction. In contrast, we can compute this set with a simple and fast algorithm based on a single max flow min cut computation. This enables us to prove all our results independent from the literature on the more general case of auctions where the buyers have gross substitute valuation functions. Moreover, it allows us to show that minimum Walrasian prices react to changes in supply and demand in a natural way, that is, they can only increase when supply decreases or demand increases. This is a very natural monotonicity property, but it was not addressed in the literature yet.

More concretely, in Section 2, we present an ascending auction which iteratively raises prices on the objects in the left-most (i.e., inclusion-wise minimal) min s - t cut in an associated auxiliary flow network. In Section 3, we prove that the auction indeed terminates with minimum Walrasian prices and a stable allocation. Section 4 shows how one can obtain price monotonicity results purely from structural insights of the flow-based auction. Section 5 compares our work to previous work.

1.2 | Known results and related work

In this subsection, we give a short summary on what is known about ascending auctions (single- and multi-unit), in particular with valuation functions that are less restrictive. We also refer the reader to the excellent survey paper by Paes Leme [17]. The basic setup remains the same as above, that is, an auctioneer wants to determine prices $\mathbf{p} \in \mathbb{Z}_+^\Omega$ and an allocation $\mathbf{x} \in \mathbb{Z}_+^{\Omega \times N}$ satisfying $\sum_{j \in N} x_{ij} \leq b_i$ for all $i \in \Omega$. However, each buyer $j \in N$ might not have a constant v_{ij} for each object $i \in \Omega$ but instead a valuation function $v_j(\mathbf{x}_{\bullet j})$. The net utility buyer j gets given a price vector \mathbf{p} and allocation \mathbf{x} is then $v_j(\mathbf{x}_{\bullet j}) - \mathbf{p}^T \mathbf{x}_{\bullet j}$ and hence, a buyer's preferred bundle under prices \mathbf{p} is given by

$$D_j(\mathbf{p}) := \arg \max \{v_j(\mathbf{x}_{\bullet j}) - \mathbf{p}^T \mathbf{x}_{\bullet j} \mid \mathbf{0} \leq \mathbf{x}_{\bullet j} \leq \mathbf{b}, \mathbf{x}_{\bullet j} \in \mathbb{Z}^\Omega\}.$$

In the sequel, it is also convenient to define a buyer's *indirect utility function*, the maximum utility a buyer can get under a given price vector:

$$V_j(\mathbf{p}) := \max \{v_j(\mathbf{x}_{\bullet j}) - \mathbf{p}^T \mathbf{x}_{\bullet j} \mid \mathbf{0} \leq \mathbf{x}_{\bullet j} \leq \mathbf{b}, \mathbf{x}_{\bullet j} \in \mathbb{Z}^\Omega\}.$$

A valuation function $v_j : \mathbb{Z}_+^\Omega \rightarrow \mathbb{Z}$ is *gross substitute* if for all price vectors $\mathbf{p}, \mathbf{q} \in \mathbb{R}^\Omega$ with $\mathbf{p} \leq \mathbf{q}$ it holds that for all $\mathbf{x}_{\bullet j} \in D_j(\mathbf{p})$ there exists $\mathbf{y}_{\bullet j} \in D_j(\mathbf{q})$ such that $x_{ij} \leq y_{ij}$ for all $i \in \Omega$ with $p(i) = q(i)$. If additionally $\sum_{i \in \Omega} x_{ij} \geq \sum_{i \in \Omega} y_{ij}$ holds for these allocations \mathbf{x}, \mathbf{y} , then v_j is *strong gross substitute*. Intuitively, this condition expresses that if one increases the price on one object, the demand on other objects (whose prices did not increase) does not diminish. What makes this class of functions (introduced by Kelso and Crawford [13]) particularly interesting is that Walrasian prices are guaranteed to exist if all buyers have gross substitute valuation functions (see also [1]). This is not true for more general valuation functions, which makes gross substitutes essentially the widest class of valuation functions that can be handled.¹ Gul and Stacchetti [10] proved some equivalent characterizations of gross substitutes (such as the *single improvement* or *no complementarities* condition) which are insightful on their own, but which also make gross substitute more convenient to consider from an algorithmic point of view. Fujishige and Yang [8] also showed that there is a very fundamental connection of gross substitutes to discrete convex analysis, namely, a valuation function is gross substitute if and only if it is \mathbf{M}^\natural -concave. A few more characterizations of gross substitutes can be found in the more recent work of Ben-Zwi [2], which also provides more fundamental insights to ascending auctions and *overdemanded sets*. However, the aforementioned publications only handle the single-unit case, that is, each object is only available in one copy. Murota et al. [15] showed that these results all transfer when going to the multi-unit setting if the valuation functions are *strong* gross substitutes (going to the multi-unit setting in mathematical terms is to go from the Boolean lattice to the integer lattice as the domain of the valuation functions).

¹However, we should point out that there are other classes that are in some sense orthogonal to gross substitutes that have different interesting applications. In particular, *complementarities* form one such interesting class.

The connection to discrete convex analysis is helpful to define an ascending auction that can find Walrasian prices. For the single-unit case, Gul and Stacchetti [11] laid out the framework for an ascending auction which naturally increases prices on subsets of objects that are overdemanded. However, they only showed that overdemanded sets must exist if the current price vector is not Walrasian using matroid theory (and giving a Hall-type condition), they did not show how to efficiently find those sets. Ausubel [1] and Murota et al. [15, 16] also considered auctions for the multi-unit setting. Their auctions follow essentially the same idea, that is, increasing prices on overdemanded sets but also allow different start prices and price reduction steps (on underdemanded sets, which do not occur if the starting prices are low enough and the price increment step is implemented correctly). The key contribution of their work is the algorithm to find overdemanded sets efficiently. All of those auctions rely on a potential function (the *Lyapunov* function)

$$L(\mathbf{p}) = \sum_{j \in N} V_j(\mathbf{p}) + \mathbf{p}^T \mathbf{b}.$$

The main features of the Lyapunov function are that it is minimized exactly at Walrasian prices, and that it is L^{\natural} -convex (and in particular, submodular) if all valuation functions are M^{\natural} -concave (or equivalently, strong gross substitute). A buyer-optimal Walrasian price vector can be found if one selects an (inclusion-wise) minimal minimizer of $X \mapsto L(\mathbf{p} + \chi_X)$. Hence, a Walrasian price vector can be found efficiently by using submodular function minimization but also a steepest descent direction of L (which corresponds to a maximum² overdemanded set) can be computed in strongly polynomial time [15].

In [7], a follow-up paper of the paper at hand, it is described how to find the overdemanded sets for general strong gross substitute valuation functions using polymatroid sum. This leads to the so far best known running time $\mathcal{O}(|N| \cdot \text{DO} + |N| |\Omega|^3 \cdot \text{ExO})$ for one iteration where buyers have gross substitute valuations, and where DO is the running time of a demand oracle and ExO is the running time of an exchange oracle. For a detailed description of the oracle models see Section 2.1.

Murota et al. [15] also showed that the ascending auction needs $\|\mathbf{p} - \mathbf{p}_0\|$ iterations to terminate, where \mathbf{p}_0 is the initial price vector and $\mathbf{p} \geq \mathbf{p}_0$ a Walrasian price vector.

2 | A FLOW-BASED ASCENDING AUCTION

We first sketch our flow-based ascending auction, called Algorithm 2. The auction starts with the all-zero vector $\mathbf{p}_0(i) = 0$ for all $i \in \Omega$ (or with any initial price vector \mathbf{p}_0 known to be a lower bound on the minimum Walrasian price vector \mathbf{p}^*). In each iteration, given the current price vector $\mathbf{p} \in \mathbb{R}_+^{\Omega}$, the algorithm computes an integral s - t -flow f of maximum value in an auxiliary flow network $G(\mathbf{p})$ (described below). If the value $\text{val}(f)$ of flow f equals the sum of capacities on the s -leaving arcs in $G(\mathbf{p})$, denoted as $D_{\mathbf{p}}$, the algorithm stops. Otherwise ($\text{val}(f) < D_{\mathbf{p}}$), the prices on all objects in the left-most min cut are raised by one unit, and the algorithm iterates with the updated price vector. Note that the notation *left-most* is motivated from the natural way of drawing a layered s - t -network with the source on the left, the sink on the right and all remaining vertices layer-by-layer in between. However, we can also define the left-most min cut without using any kind of topology on the vertices by defining it as the (unique) *inclusion-wise minimal* set S , with $s \in S$ that minimizes the sum of capacities on the outgoing arcs, $\text{cap}(S)$. We can compute this set via a breadth-first-search (BFS) from s in the residual network which corresponds to a max flow. Then Theorem 2 shows that the final price vector \mathbf{p}^* returned by the algorithm is the minimum (and thus buyer-optimal) competitive price vector.

For computing a corresponding stable and market-clearing allocation \mathbf{x}^* such that $(\mathbf{p}^*, \mathbf{x}^*)$ is a buyer-optimal Walrasian equilibrium, we modify $G(\mathbf{p}^*)$ slightly to network $H(\mathbf{p}^*)$, which allows us to find a stable allocation where $\min\{\sum_{i \in \Omega} b_i, \sum_{j \in N} d_j\}$ items are sold. Furthermore, we show that each object $i \in \Omega$ with $\mathbf{p}^*(i) > 0$ is completely sold.

2.1 | Structure of preferred bundles and oracle models

For a fixed price vector \mathbf{p} and a fixed buyer j , a minimal preferred bundle of buyer j can be computed with the following greedy approach. Let \prec_j be a total ordering of the items by nondecreasing payoffs, that is, satisfying

$$v_{ij} - p(i) \geq v_{kj} - p(k) \quad \text{whenever} \quad i \prec_j k. \quad (2)$$

For ease of notation, assume that $\prec_j = (1, \dots, m)$. Note that the ordering \prec_j might not be unique due to ties in the payoffs. However, each ordering \prec_j satisfying (2) uniquely defines a minimum preferred bundle of buyer j which can be constructed as follows:

²Here *maximum* does not refer to the size of the set but to the *overdemandedness*, that is, the difference between demand and supply of that set.

Algorithm 1. Algorithm to construct a minimum preferred bundle for buyer j **Input:** Supplies b_i , demand d_j , valuations v_{ij} for all $i \in \Omega$ and an ordering $\prec_j = (1, \dots, m)$ satisfying (2)**Output:** A minimal preferred bundle for buyer j and the last item k_j selected for this bundleInitially, no items are selected, the residual demand d_j^{res} is equal to d_j , and $i := 1$ **while** $d_j^{res} > 0$, $v_{ij} > p(i)$ and $i \leq m$ **do** Set $x_{ij} := \min\{b_i, d_j^{res}\}$ Iterate with residual demand $d_j^{res} := d_j^{res} - x_{ij}$, and $i := i + 1$ **return** x_j and $k_j := i - 1$

We observe that the greedy algorithm selects exactly b_i copies of all items in $\{1, \dots, k_j - 1\}$ and $d_j - \sum_{\ell=1}^{k_j-1} b_\ell$ copies from item k_j , and that the set of *all* minimal preferred bundles obeys the following structure. Consider the two item sets Ω'_j and Ω''_j consisting of all items of larger or equal payoff than the payoff of item k_j , which was selected last by the greedy algorithm, that is,

$$\Omega'_j(\mathbf{p}) := \{i \in \Omega \mid v_{ij} - p(i) > v_{kj} - p(k_j)\}, \text{ and}$$

$$\Omega''_j(\mathbf{p}) := \{i \in \Omega \mid v_{ij} - p(i) = v_{kj} - p(k_j)\}.$$

Note that $\Omega'_j(\mathbf{p})$ can be empty. This happens when all elements in the preferred bundle have the same payoff.³

If $\sum_{i \in \Omega'_j(\mathbf{p}) \cup \Omega''_j(\mathbf{p})} b_i < d_j$, a minimal preferred bundle contains less than d_j items. In such a case, and if there are objects $i \in \Omega$ of payoff $v_{ij} - p(i) = 0$, we assume what is called “free disposal” in the literature. Namely, that buyer j is indifferent between choosing a preferred bundle as described, or filling the demand up with items in

$$\Omega'''_j(\mathbf{p}) := \{i \in \Omega \mid v_{ij} - p(i) = 0\}.$$

Thus, in every preferred bundle, buyer j buys exactly $d_j'(\mathbf{p}) := \sum_{i \in \Omega'_j(\mathbf{p})} b_i$ items from objects in $\Omega'_j(\mathbf{p})$ and $d_j''(\mathbf{p}) := \min \left\{ \sum_{i \in \Omega''_j(\mathbf{p})} b_i, d_j - d_j'(\mathbf{p}) \right\}$ items from objects in $\Omega''_j(\mathbf{p})$. In addition, there might be up to $d_j'''(\mathbf{p}) := \min \left\{ \sum_{i \in \Omega'''_j(\mathbf{p})} b_i, d_j - d_j'(\mathbf{p}) - d_j''(\mathbf{p}) \right\}$ items of objects with zero payoff in a preferred bundle.

To shorten notation, we omit the price when \mathbf{p} is clear from the context, for example, we write Ω'_j instead of $\Omega'_j(\mathbf{p})$, or d_j' instead of $d_j'(\mathbf{p})$.

Oracles. Since we consider this market in an auction setting, it is natural (or even necessary) to assume that the auctioneer has some sort of communication protocol with the buyers. Ascending auctions, such as those by [1, 11], require the buyers to reveal their demand correspondence $D_j(\mathbf{p})$ in every iteration for some price vector \mathbf{p} (which can have size exponential in $|\Omega|$). In the model with truncated additive valuations, we can ask a buyer to provide her demand d_j and her demand sets $\Omega'_j(\mathbf{p})$, $\Omega''_j(\mathbf{p})$, and $\Omega'''_j(\mathbf{p})$ (the latter one is not required for computing overdemanded sets, however). Asking the buyers for those sets is quite intuitive as Ω' denotes the objects which she wants strongly, Ω'' denotes the objects of which she wants to get as many as she can to fill up the demand, and Ω''' denotes the objects she is indifferent to buy (as buying them results in payoff zero). Hence, it is natural in our setting that the auctioneer only asks the buyers for those three sets and her total demand to get a much more compact representation of the demand sets. We denote such an oracle call, in other words a question from the auctioneer to one buyer for Ω' , Ω'' , Ω''' , and d , as *tier oracle* and use TO for short. In sequel, in our runtime analysis we include TO for these oracle calls. Note that a buyer, who knows her valuations can answer the oracle call in $O(|\Omega| \log(|\Omega|))$ (see Algorithm 1).

In the literature often a demand oracle DO is used (i.e., asking for an minimal preferred bundle) or an exchange oracle ExO (i.e., asking how many items of one object the buyer is willing to exchange against the same amount of items of a different object). Each of these oracles can be used in turn, and it depends on the objective which one is most efficient to use. In our case the tier oracle is the most suitable.

2.2 | Construction of auxiliary flow networks

In each iteration of our flow-based ascending auction with current price \mathbf{p} , buyer j reveals to the auctioneer the following information: the two sets Ω'_j and Ω''_j , together with the amounts d_j' and d_j'' they want to buy from these sets, respectively, and the set $\Omega'''_j = \Omega'''_j(\mathbf{p})$ of items of payoff 0 with the amount d_j''' .

³Note that this also holds even if there is no item that is dropped (i.e., the supply for these objects does not exceed the demand). This way, the formal definition is straightforward and we do not have to deal with case distinctions in the analysis.




FIGURE 1 The network $G(\mathbf{p})$ with two buyers j_1, j_2 and three objects α, β, γ . We have valuations $v_{j_1} = (3, 2, 1)$ and $v_{j_2} = (0, 2, 0)$, demands $d_{j_1} = 4, d_{j_2} = 2$, supply $b_\alpha = b_\beta = 1, b_\gamma = 4$, and current prices $p(\alpha) = p(\beta) = p(\gamma) = 0$. In this example, the item sets at prices $\mathbf{p} = (0, 0, 0)$ are given by $\Omega'_1 = \{\alpha, \beta\}$, $\Omega''_1 = \{\gamma\}$, $\Omega'_2 = \emptyset$ and $\Omega''_2 = \{\beta\}$, $\Omega'_2 = \{\beta\}$, $\Omega''_2 = \{\alpha, \gamma\}$.

Given this information, our algorithm constructs a network $G(\mathbf{p})$ and uses a max-flow computation to compute the set of objects on which the price has to be increased. In the following we describe the construction of this network. An example with two buyers and three objects is given in Figure 1.

The vertex set of $G(\mathbf{p})$ consists of a source s , a sink t , one vertex for each object $i \in \Omega$, and two vertices j' and j'' for each buyer $j \in N$. We denote the collection of the j' vertices as N' and the one of j'' as N'' . Vertices j' and j'' correspond to the sets $\Omega'_j(\mathbf{p})$ and $\Omega''_j(\mathbf{p})$, respectively.

The arcs (with positive capacity) are defined as follows:

- (s, j') with capacity $d_{j'}$ for all $j \in N$,
- (s, j'') with capacity $d_{j''}$ for each $j \in N$,
- (j', i) with capacity $c_{j'i} := b_i$ for all $j \in N, i \in \Omega'_j$
- (j'', i) with capacity $c_{j''i} := \min\{b_i, d_{j''}\}$ for all $j \in N, i \in \Omega''_j$, and
- (i, t) with capacity b_i for all $i \in \Omega$.

We denote the total capacity on the s -leaving arcs by $\text{cap}(s)$, and observe that $\text{cap}(s) = \sum_{j \in N} (d_{j'} + d_{j''})$. That is, $\text{cap}(s)$ is equal to the sum of sizes of the buyers' minimal preferred bundles. The following lemma states that prices \mathbf{p} are competitive if and only if the value of a max flow in $G(\mathbf{p})$ is equal to $\text{cap}(s)$.

Lemma 1. *The prices \mathbf{p} are competitive if and only if there is a flow f in $G(\mathbf{p})$ of value $\text{val}(f) = \text{cap}(s)$. Moreover, given a competitive price vector \mathbf{p} , an associated stable allocation \mathbf{x} can be computed via a single max flow computation in $G(\mathbf{p})$.*

Proof. There is a clear one-to-one correspondence between an integral s - t -flow and a feasible assignment of objects. Since all capacities in $G(\mathbf{p})$ are integral, there exists a max flow with integral values. By assigning $f_{j'i} + f_{j''i}$ items of object i to buyer j we obtain a feasible assignment. Note that at least one of the summands is zero since at most one of the arcs has positive capacity ($\Omega'_j \cap \Omega''_j = \emptyset$). This assignment is competitive if and only if the demand of every buyer at prices \mathbf{p} is satisfied. This is equivalent to the requirement that the flow satisfies all s -leaving arcs, that is, $\text{val}(f) = \sum_{j \in N} (d_{j'} + d_{j''}) = \text{cap}(s)$. ■

However, in this allocation \mathbf{x} , no item with payoff zero is included. Since we aim for an allocation where as much as possible is sold, we have to allocate the other items as well. This is easy for those items which have price zero, since then every buyer with unsatisfied demand is willing to buy them. We show in Theorem 3 that there are enough buyers who are willing to buy the items with a positive price as well.

We extend $G(\mathbf{p}^*)$ to a flow network $H(\mathbf{p}^*)$ such that the assignment of buyers to objects of payoff zero is possible. To do so, we first balance the supply and the demand. If $\sum_{i \in \Omega} b_i < \sum_{j \in N} d_j$, we add a dummy object i_0 with supply $b_{i_0} = \sum_{j \in N} d_j - \sum_{i \in \Omega} b_i$ and valuations $v_{ij_0} = 0$ for all $j \in N$. If $\sum_{i \in \Omega} b_i > \sum_{j \in N} d_j$, we add a dummy buyer j_0 with demand $d_{j_0} = \sum_{i \in \Omega} b_i - \sum_{j \in N} d_j$ and valuations $v_{ij_0} = 0$ for all $i \in \Omega$. Now we can assume that $\sum_{i \in \Omega} b_i = \sum_{j \in N} d_j$. Note that Algorithm 2 computes the same prices with the dummy object or buyer as without. To construct $H(\mathbf{p}^*)$ from $G(\mathbf{p}^*)$, additionally we add a vertex j''' for each buyer $j \in N$ and an arc (s, j''') with capacity $d_{j''}$. Furthermore, we add for each $j \in N$ the arc (j''', i) with capacity b_i for $i \in \Omega''_j$.

Algorithm 2. Price-raising algorithm

Input: Supplies b_i , demands d_j , valuations v_{ij} for all $i \in \Omega$ and all $j \in N$, initial prices $\mathbf{p}_0 := \mathbf{0}$

Output: Buyer-optimal Walrasian prices \mathbf{p}^*

Initialize: $\ell := 0$

Construct the auxiliary network $G_0 := G(\mathbf{p}_0)$ and find an integral max flow f^0 in it

while $\text{val}(f^\ell) < \text{cap}_{G_\ell}(s)$ **do**

 Determine the left-most min cut C_ℓ

 Increase prices on $I_\ell := C_\ell \cap \Omega$, that is, $\mathbf{p}_{\ell+1} := \mathbf{p}_\ell + \chi_{I_\ell}$

 Construct the network $G_{\ell+1} := G(\mathbf{p}_{\ell+1})$

 Find an integral max flow in $G_{\ell+1}$

$\ell := \ell + 1$

return $\mathbf{p}^* := \mathbf{p}_\ell$

Algorithm 3. Allocation algorithm

Input: Supplies b_i , demands d_j , $\Omega'_j(\mathbf{p}^*)$, $\Omega''_j(\mathbf{p}^*)$, $\Omega'''_j(\mathbf{p}^*)$ for all buyers $j \in B$, prices \mathbf{p}^*

Output: Stable allocation \mathbf{x}^* where as much as possible is sold

Construct the auxiliary network $H(\mathbf{p}^*)$

Find an integral max flow f in $H(\mathbf{p}^*)$

Determine the allocation \mathbf{x}^* given by the flow on the arcs (j, i)

return \mathbf{x}^*

Proposition 1. A max flow in network $H(\mathbf{p})$ and its corresponding allocation satisfy:

1. A feasible flow in $G(\mathbf{p})$ is a feasible flow in $H(\mathbf{p})$.
2. If for buyer $j \in N$ the flow on the arcs (s, j') and (s, j'') is saturated, j is assigned to one of her preferred bundles at prices \mathbf{p} .
3. If \mathbf{p} is a Walrasian price vector, the allocation induced by a max flow is stable, that is, every buyer obtains a preferred bundle.

Note that part 3 of the proposition is due to the fact that there exists an allocation where everything is sold (including dummy items), so the max flow saturates all s -leaving arcs, since demand and supply are balanced. Hence in the corresponding allocation every buyer gets a preferred bundle.

2.3 | Computation of the buyer-optimal Walrasian equilibrium

Here we formally describe Algorithm 2. Each of its iterations can be done in polynomial time since the network can be constructed in polynomial time and only one max flow computation is needed.

The intuition of Algorithm 2 is to increase the price of a set of objects whenever the demand on this set exceeds the supply. It is natural to increase the prices of the objects of an overdemanded subset until the buyers that were interested in these objects get interested in other objects as well. This is exactly what happens in the following algorithm.

In Section 3.1 we show that the prices computed by Algorithm 2 are the component-wise minimum competitive prices.

Given component-wise minimum competitive prices \mathbf{p}^* , Algorithm 3 constructs the auxiliary network $H(\mathbf{p}^*)$ and its max flow, leading to allocation \mathbf{x}^* . Hence \mathbf{x}^* can be found in polynomial time. Section 3.2 shows that the value of the maximum flow is $\max \left\{ \sum_{i \in \Omega} b_i, \sum_{j \in N} d_j \right\}$, since we include a dummy buyer resp. a dummy object. With Proposition 1 this implies that each buyer is assigned to one of her preferred bundles.

If the supply does not exceed the demand, everything is sold. Otherwise, only the items which are allocated to the dummy buyer j_0 are not sold. Since $v_{ij_0} = 0$ for all $i \in \Omega$, the price of an object allocated to j_0 has to be zero. Thus, all objects with positive price are completely sold.

Theorem 1. The prices \mathbf{p}^* computed by Algorithm 2 are the component-wise minimum competitive prices. Moreover, under the stable allocation \mathbf{x}^* returned by Algorithm 3 as much as possible is sold, and every item with positive price is sold. Thus, the prices \mathbf{p}^* coincide with the buyer-optimal Walrasian price vector.

We prove Theorem 1 in Section 3. Note that the fact that the flow-based ascending auction returns the buyer-optimal Walrasian prices can also be shown directly by combining results of the literature on strong substitute valuation functions [1] with some observation (see Section 5). Ausubel [1] showed that the buyer-optimal prices are computed by an ascending auction, where prices are always raised on the inclusion-wise minimal set corresponding to the steepest descent direction of the Lyapunov function. We observe that these sets correspond to minimal maximum overdemanded sets which are exactly the sets we compute with a left-most min cut computation. However, the proof given in the subsequent sections has two advantages. First, we show that the minimum Walrasian prices \mathbf{p}^* coincide with the minimum competitive prices, which will turn out to be crucial when proving our monotonicity results in Section 4. Second, our proof is independent from the literature on strong substitute and L^1 -convex functions and uses only network flow arguments.

Regarding the running time, our considerations above yield the following lemma.

Lemma 2. *One iteration of Algorithm 2 requires at most $|N|$ tier oracle calls, $\mathcal{O}(|N||\Omega|)$ time to construct the network and one max flow min cut computation.*

For the max flow min cut computation, one can use, for example, the algorithm of Chen et al. [3] to find a max flow in $\mathcal{O}((|N||\Omega|)^{1+o(1)})$, and to find the left most min cut, one can use a breadth first search (BFS) from s in the residual network in $\mathcal{O}(|N||\Omega|)$ time. This results in total running time of $\mathcal{O}(|N|TO + (|N||\Omega|)^{1+o(1)})$.

Thus, our Algorithm 2 is faster than the running time of the algorithm for the general strong gross substitute case by Murota et al. [15], which is $\mathcal{O}(|N|DO + |N||\Omega|^4 \log(|N||\Omega|B)ExO)$, where B is the maximum number of copies of an object. Note that different oracle models are used in the two algorithms, but since the running time of the tier oracle can be bounded by $\mathcal{O}(|\Omega| \log(|\Omega|))$ for a value oracle see Algorithm 1, we consider this a fair statement. In a follow up paper of this work [7] we improved the running time for the general case to $\mathcal{O}(|N|DO + |N||\Omega|^3 ExO)$. Again different oracle models are used and thus the running times are not directly comparable. As before, we can afford a running time of $\mathcal{O}(|\Omega| \log(|\Omega|))$ for the tier oracle without exceeding the running time of the general mechanism. If we assume in favor of the general mechanism that the exchange and demand oracle queries takes $\mathcal{O}(1)$ time, our algorithm for the special case is still faster in all instances where $|N|^{o(1)} < |\Omega|^{2-o(1)}$, that is, in all instances where the number of buyers is not drastically larger than the number of items.

As shown in [15], the number of iterations of Algorithm 2 is at most $\|\mathbf{p}^* - \mathbf{p}_0\|_\infty$, where \mathbf{p}^* is the minimum Walrasian price vector. Since $\mathbf{p}^*(i) \leq \max_{j \in \Omega} v_{ij}$ this is pseudo-polynomial. Note, however, that we may as well start with any alternative start vector \mathbf{p}_0 , as long as \mathbf{p}_0 is known to be component-wise smaller or equal to the unique minimum competitive price vector.

In Section 5.2, we present a variation of the auction where the price is raised as far as possible for the same inclusion-wise minimal maximum overdemanded set.

2.4 | Auctioneer with memory: Warm start with flow updates

A very natural idea is to not compute the flow completely from scratch in every iteration but instead to update the flow, when updating the price vector. We will describe this procedure and analyze the resulting running time. Interestingly, in one iteration this may not speed up the algorithm, but when combining this with the adapted step length algorithm described in Section 5.2, then the running time of the complete auction improves.

To update the flow from one iteration to the next, assume we are given a maximum flow f^ℓ in the network induced by price vector \mathbf{p}_ℓ . Moreover, let C be the left most min cut in the network $G(\mathbf{p}_\ell)$ and assume $C \neq \{s\}$ (because in this case the computed prices are already competitive and the auction terminates). Then Algorithm 4 shows how to update the flow to obtain a new feasible flow in $G(\mathbf{p}_{\ell+1})$. The idea is to keep as much flow as possible from the previous assignment. We use that items might change their set, that is, go from $\Omega'_j(\mathbf{p}_\ell)$ to $\Omega''_j(\mathbf{p}_{\ell+1})$ or vice versa, that is, from $\Omega''_j(\mathbf{p}_\ell)$ to $\Omega'_j(\mathbf{p}_{\ell+1})$ but that for a fixed buyer this is possible in at most one direction. A third possibility is that a buyer might lose interest in an item, but in this case this item is not in $\Omega'_j(\mathbf{p}_{\ell+1}) \cup \Omega''_j(\mathbf{p}_{\ell+1})$ and thus no flow is assigned.

This flow update is well defined, since feasibility is shown in the following lemma.

Lemma 3. *The flow f computed in Algorithm 4 is a feasible flow in $G(\mathbf{p}_{\ell+1})$.*

Proof. Note that flow conservation is fulfilled in every node by construction. It remains to show that the capacities in $G(\mathbf{p}_{\ell+1})$ are obeyed. Let C be the left-most min cut in iteration ℓ and $I = C \cap \Omega$ the set of items on which the prices are increased.

- For (s, j') this is given: we assign $\sum_{i \in \Omega'_j(\mathbf{p}_{\ell+1})} (f_{ji}^\ell + f_{ji}^{\ell'})$ units of flow and the capacity is given by $\sum_{i \in \Omega'_j(\mathbf{p}_{\ell+1})} b_i$. Note that $(f_{ji}^\ell + f_{ji}^{\ell'}) \leq b_i$ since the only i leaving arc has capacity b_i .
- For (s, j'') we consider how $d_{j''}(\mathbf{p}_\ell)$ changes if we increase the prices on objects in I .

Algorithm 4. Flow update algorithm

Input: Network $G(\mathbf{p}_\ell)$ with max flow f^ℓ , $\mathbf{p}_{\ell+1} := \mathbf{p}_\ell + \chi_I$, new network $G(\mathbf{p}_{\ell+1})$

Output: A feasible flow f in $G(\mathbf{p}_{\ell+1})$ with $D_{\mathbf{p}_\ell} - \text{val}(f^\ell) \geq D_{\mathbf{p}_{\ell+1}} - \text{val}(f)$

$f := 0$ for all arcs in $G(\mathbf{p}_{\ell+1})$

for $(j, i) \in N \times \Omega$ with $f_{ji}^\ell > 0$ or $f_{ji}^\ell > 0$ **do**

if (j', i) is an arc in $G(\mathbf{p}_{\ell+1})$, that is, $i \in \Omega_j'(\mathbf{p}_{\ell+1})$ **then**

 Add $(f_{ji}^\ell + f_{ji}^\ell)$ flow units to f on the path $s-j'-i-t$

else if (j'', i) is an arc in $G(\mathbf{p}_{\ell+1})$, that is, $i \in \Omega_j''(\mathbf{p}_{\ell+1})$ **then**

 Add $(f_{ji}^\ell + f_{ji}^\ell)$ flow units to f on the path $s-j''-i-t$

return f

- If there are objects in $\Omega_j''(\mathbf{p}_\ell) \cap \Omega_j'(\mathbf{p}_{\ell+1})$, let $M := \Omega_j''(\mathbf{p}_\ell) \cap \Omega_j'(\mathbf{p}_{\ell+1})$ be the objects which move from Ω_j'' to Ω_j' and $S := \Omega_j''(\mathbf{p}_\ell) \cap \Omega_j''(\mathbf{p}_{\ell+1})$ be the items that stay in Ω_j'' .

We know that the objects in $\Omega_j''(\mathbf{p}_\ell) \cap I \subseteq \Omega_j''(\mathbf{p}_{\ell+1})$. Moreover, the objects in $\Omega_j'(\mathbf{p}_\ell)$ stay in $\Omega_j'(\mathbf{p}_{\ell+1})$, that is, $\Omega_j'(\mathbf{p}_\ell) \subseteq \Omega_j'(\mathbf{p}_{\ell+1})$. Hence, the demand d_j increases and the demand $d_{j''}$ decreases by the number of items in M , that is, $\sum_{i \in M} b_i$.

Recall that left-most min cut C is defined by the vertices reachable from s in the residual network corresponding to f^ℓ . By definition $M \cap I = \emptyset$ and $S \subseteq I$.

If j'' is in the left-most min cut, the arcs in the cut which are leaving this node are fully used in any maximum flow. Thus, $f_{ji}^\ell = c_{ji} = b_i$ for $i \in M$. That $c_{ji} = \min\{b_i, d_{j''}\} = b_i$ follows since otherwise with this item the complete demand could be satisfied. But this is a contradiction since buyer j wants to buy items with less payoff at prices $\mathbf{p}_{\ell+1}$, in other words $i \in \Omega_j'(\mathbf{p}_{\ell+1})$ and not $i \in \Omega_j''(\mathbf{p}_{\ell+1})$.

Since this flow is shifted to j' after the price update, the flow is reduced by the number of items in M as well, thus the capacities are still obeyed.

If j'' is not in the left-most min cut, we cannot reach j'' from items in $S \subseteq I$. Hence, the flow on the edges $(j''i)$ with $i \in S$ is zero (otherwise the backwards arc exists in the residual network and j'' is reachable). Hence, in Algorithm 4 we do not assign any flow to an edge through the vertex j'' . Thus, the capacity of (s, j'') is still not exceeded.

- Consider the case where $\Omega_j''(\mathbf{p}_\ell) \cap \Omega_j'(\mathbf{p}_{\ell+1}) = \emptyset$.

If Ω_j'' does not lose any objects by the price update, we know that

$$d_j''(\mathbf{p}_{\ell+1}) = d_j''(\mathbf{p}_\ell) + \sum_{i \in \Omega_j''(\mathbf{p}_\ell) \cap \Omega_j''(\mathbf{p}_{\ell+1})} b_i.$$

Hence, the capacity is not exceeded on (s, j'') .

It remains to show that the capacity is not exceeded if Ω_j'' loses some objects (it might get new ones from Ω_j' as well). The only situation which can cause problems is if the demand of objects in Ω_j'' decreases. The demand d_j'' only decreases if there are objects moving from Ω_j'' to Ω_j' (which does not happen by assumption) or if there are not enough objects with a positive payoff. The latter case implies that all items available in $\Omega_j''(\mathbf{p}_{\ell+1})$ are demanded and thus, the capacity constraint on (s, j'') is fulfilled.

- For (i, t) the capacity does not change, so the capacity constraints are fulfilled.
- For (j', i) and (j'', i) it follows directly by the definition of the capacity and since the capacity on the s -leaving and t -entering arcs is not exceeded. ■

Lemma 4. The adaption of the network and the check whether the set of objects I in the left-most min cut changes runs in time $\mathcal{O}(|N|TO + |N||\Omega|)$.

Proof. After asking every buyer to report their preferences for the new prices, every arc is checked and then in constant time the flow is adapted (see Algorithm 4). Afterwards, the left-most min cut can be computed by a breadth first search (BFS) in the residual network. As there are $\mathcal{O}(|N||\Omega|)$ many arcs, the statement follows. ■

We can use the structure of the left-most min cut to show that if the flow decreases in an update step, then the demand decreases by at least the same amount. This will help us to show that the prices returned by the auction are market-clearing (see Section 3.2). Moreover, this enables us to give an upper bound on the overall running time of the ascending auction with adapted step length (see Section 5.2).

Lemma 5. *Given a max flow in $G(\mathbf{p}_\ell)$ with corresponding left-most min cut C . Let $I = C \cap \Omega$ be the overdemanded set and $\mathbf{p}_{\ell+1}$ be the price vector after the price update. Then, the demand $d_j + d_{j'}$ of buyer j will decrease at least by*

$$\sum_{\substack{i \in \Omega_j(\mathbf{p}_\ell): \\ i \notin \Omega_j'(\mathbf{p}_{\ell+1}) \cup \Omega_j''(\mathbf{p}_{\ell+1})}} f_{j''i}^\ell. \quad (3)$$

Proof. If (3) is zero, the statement is directly fulfilled. Thus, from now on, we consider the cases when flow is removed. This can only occur if some objects get a payoff of zero after the price update, that is, if

$$S := \Omega_j''(\mathbf{p}_\ell) \setminus (\Omega_j'(\mathbf{p}_{\ell+1}) \cup \Omega_j''(\mathbf{p}_{\ell+1})) \neq \emptyset.$$

Note that S describes the set of objects which move from Ω_j'' to Ω_j''' . Since Ω_j''' just contains objects with utility zero and Ω_j'' only those with a positive payoff, all objects in S are contained in I .

By definition, the demand $d_j + d_{j'}$ will decrease by

$$\max \left\{ 0, d_{j'}(\mathbf{p}_\ell) - \sum_{i \in \Omega_j''(\mathbf{p}_\ell) \setminus I} b_i \right\}.$$

If j'' is contained in the left-most min cut, all arcs from j'' to $\Omega_j''(\mathbf{p}_\ell) \setminus I$ are fully satisfied (since they are not reachable from j''). Hence, there can be at most $d_{j'}(\mathbf{p}_\ell) - \sum_{i \in \Omega_j''(\mathbf{p}_\ell) \setminus I} b_i \geq 0$ units of flow going through j'' to vertices in S . Thus, for buyer j we reduce the demand at least by the removed flow units traveling through a vertex of j .

If j' is not contained in the left-most min cut, there is no flow on the arcs from j'' to I . Thus, no flow through buyer j is removed and we are done in this case as well. ■

Corollary 1. *The flow f computed in Algorithm 4 satisfies*

$$D_{\mathbf{p}_\ell} - \text{val}(f^\ell) \geq D_{\mathbf{p}_{\ell+1}} - \text{val}(f).$$

3 | THE ASCENDING AUCTION RETURNS THE BUYER-OPTIMAL WALRASIAN PRICES

We will show that Algorithm 2 returns buyer-optimal Walrasian prices by showing separately that the prices are the component-wise minimum competitive ones and that the prices are market clearing.

3.1 | The ascending auction returns component-wise minimum competitive prices

We now analyze Algorithm 2 by considering the behavior of the buyers at given prices \mathbf{p} using the structure of $G(\mathbf{p})$. As usual, for a given digraph (V, A) , we use $\Gamma^-(v)$ to refer to all nodes that are the starting node of an incoming arc into v , that is, $\Gamma^-(v) := \{u \in V \mid (u, v) \in A\}$, analogously for $\Gamma^+(v)$. We extend this definition also to sets, so that $\Gamma^-(I) := \{u \in V \mid \text{there exists } v \in I \text{ with } (u, v) \in A\}$.

For a given set of items $I \subseteq \Omega$, we denote the sum of demands of all buyers which cannot be fulfilled by items which are not in I by

$$d_p(I) := \sum_{j \in \Gamma^-(I)} \max \left\{ 0, d_{j \cdot} - \sum_{i \in \Gamma^+(j \cdot) \setminus I} c_{ij} \right\}.$$

Note that $d_p(I)$ is a natural lower bound on the number of copies from I that are needed in every stable allocation. We call an item set I *overdemanded* if $\sum_{i \in I} b_i < d_p(I)$. The following condition (4) is a consequence of Hall's Theorem for b -matching (cf. [12]).

Lemma 6. *Prices \mathbf{p} are competitive if and only if*

$$\sum_{i \in I} b_i \geq d_p(I) \quad \text{for all } I \subseteq \Omega. \quad (4)$$

Moreover, if prices \mathbf{p} are not competitive, and C is the left-most min cut in $G(\mathbf{p})$, then $C \cap \Omega$ is an overdemanded set.

Proof. Recall from Lemma 1 that prices \mathbf{p} are competitive if and only if there exists a flow f in $G(\mathbf{p})$ of value $\text{val}(f) = \text{cap}(\{s\}) = \sum_{j \in N} d_{j'} + d_{j''}$. Thus, by the Max Flow Min Cut Theorem, prices \mathbf{p} are competitive if and only if $\{s\}$ is a min cut in $G(\mathbf{p})$.

First, we show that (4) is necessary, that is, that it holds if \mathbf{p} is competitive. Consider a flow f in $G(\mathbf{p})$ of value $\text{val}(f) = \sum_{j \in N} d_{j'} + d_{j''}$ and some arbitrary set $I \subseteq \Omega$. According to the flow conservation, for each vertex j^\bullet corresponding to a buyer $j \in N$ (i.e., either j' or j'') we have

$$\sum_{i \in \Gamma^+(j^\bullet) \cap I} f_{j^\bullet i} = \overbrace{f_{sj^\bullet}}^{=d_{j'} \text{ resp. } d_{j''}} - \sum_{i \in \Gamma^+(j^\bullet) \setminus I} \overbrace{f_{j^\bullet i}}^{\leq c_{j^\bullet i}} \geq \max \left\{ 0, d_{j^\bullet} - \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i} \right\}.$$

If we sum over all $j^\bullet \in \Gamma^-(I)$ we get

$$d_p(I) = \sum_{j^\bullet \in \Gamma^-(I)} \max \left\{ 0, d_{j^\bullet} - \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i} \right\} \leq \sum_{j^\bullet \in \Gamma^-(I)} \sum_{i \in \Gamma^+(j^\bullet) \cap I} f_{j^\bullet i} = \sum_{i \in I} \sum_{j^\bullet \in \Gamma^-(I)} f_{j^\bullet i} = \sum_{i \in I} f_{it} \leq \sum_{i \in I} b_i.$$

Thus, I is not an overdemanded set. Since we chose I arbitrary, we have that (4) is fulfilled.

Now, we show that condition (4) implies that prices \mathbf{p} are competitive. To see this, suppose that \mathbf{p} is not competitive, and let C be the left-most min cut of $G(\mathbf{p})$. Since \mathbf{p} is not competitive, we know that

$$\text{cap}(C) < \text{cap}(\{s\}) = \sum_{j^\bullet \in N' \cup N''} d_{j^\bullet}. \quad (5)$$

Define $I := C \cap \Omega$ and $T := C \cap (N' \cup N'')$. The capacity of C is given by

$$\text{cap}(C) = \sum_{i \in I} b_i + \sum_{j^\bullet \in T} \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i} + \sum_{j^\bullet \in (N' \cup N'') \setminus T} d_{j^\bullet}. \quad (6)$$

By combining and rearranging (6) and (5) we get the following chain of inequalities:

$$\sum_{i \in I} b_i < \sum_{j^\bullet \in T} d_{j^\bullet} - \sum_{j^\bullet \in T} \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i} = \sum_{j^\bullet \in T} \left(d_{j^\bullet} - \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i} \right) \leq \sum_{j^\bullet \in T} \max \left\{ 0, d_{j^\bullet} - \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i} \right\} = d_p(I).$$

Hence, $I = C \cap \Omega$ is an overdemanded set and (4) does not hold in this case. ■

Theorem 2. The prices \mathbf{p}^* returned by our Algorithm 2 are the (unique) component-wise minimum competitive prices, that is, if \mathbf{q} is a competitive price vector then $p^*(i) \leq q(i)$ for all $i \in \Omega$.

Proof. Assume toward a contradiction that for a competitive price vector \mathbf{q} , there is an object $i \in \Omega$ such that $p^*(i) > q(i)$. Let \mathbf{p}_τ denote the price vector in iteration τ . Then, for the start vector \mathbf{p}_0 we have $\mathbf{p}_0 \leq \mathbf{q}$. Let t be the last iteration where $\mathbf{p}_t \leq \mathbf{q}$, that is, there is an object $i \in \Omega$ such that $p_{t+1}(i) > q(i)$. Let I be the overdemanded set chosen by the algorithm in iteration t and split it into two parts $I^= \cup I^<$ as follows:

$$I = \{i \in \Omega \mid p_t(i) < p_{t+1}(i)\}, \quad I^= = \{i \in I \mid p_t(i) = q(i)\}, \quad I^< = \{i \in I \mid p_t(i) < q(i)\}.$$

We will derive a contradiction by showing that $I^<$ induces a min cut \tilde{C} which is a strict subset of C , since $I^=$ is non-empty by choice of t . To do so, we start with analyzing the behavior of buyer $j \in N$ at prices \mathbf{q} by comparing it with the behavior at prices \mathbf{p}_t . For this purpose, we fix the network properties, that is, we talk about everything w.r.t prices \mathbf{p}_t if not stated otherwise. In the following lemma, we will show that

$$\tilde{C} = \{s\} \cup \left\{ j^\bullet \in N' \cup N'' \mid d_{j^\bullet} > \sum_{i \in \Gamma^+(j^\bullet) \setminus I^<} c_{j^\bullet i} \right\} \cup I^<$$

is also a min cut with $\tilde{C} \subsetneq C$ (see Lemma 7 below). This, however, is a contradiction to C being a left-most min cut, implying that the assumption that there is an object i with price $p^*(i) > q(i)$ cannot be true. Therefore, Algorithm 2 finds the component-wise minimum competitive price vector. ■

The following lemma is needed in the proof of Theorem 2 and uses the notation which is described there.

Lemma 7. The cut $\tilde{C} = \{s\} \cup \left\{ j^\bullet \in N' \cup N'' \mid d_{j^\bullet} > \sum_{i \in \Gamma^+(j^\bullet) \setminus I^<} c_{j^\bullet i} \right\} \cup I^<$ is a min cut and $\tilde{C} \subsetneq C$.




FIGURE 2 Proof sketch. Illustration of a minimum s - t -cut C and the induced sets I (light blue) and T (light green). The set $I^=$ is red, $I^<$ is orange, T_1 is yellow and T_2 is petrol.

Proof. A sketch of the left-most min cut is given in Figure 2. Consider the buyers who need to buy items of I such that I becomes overdemand. Let $T = C \cap (N' \cup N'')$, T_1 be the subset of buyers $(N' \cup N'')$ demanding some objects in $I^=$, and T_2 are those who do not, that is,

$$T = \left\{ j^\bullet \in N' \cup N'' \mid d_{j^\bullet} > \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i} \right\}, \quad T_1 = \{ j^\bullet \in T \mid c_{j^\bullet i} > 0 \text{ for an } i \in I^= \}, \quad T_2 = T \setminus T_1.$$

To see that this fits to the definition of T , we only need to check which vertices in $N' \cup N''$ are contained in the cut C . Recall that the capacity $\text{cap}(C)$ of cut C is defined as the sum of capacities on all outgoing arcs of C . If a vertex j^\bullet is in the cut, $\text{cap}(C)$ includes the sum over all capacities of arcs (j^\bullet, i) where $i \notin I$. If vertex j^\bullet is not in the cut, $\text{cap}(C)$ includes the capacity of the arc (s, j^\bullet) which is d_{j^\bullet} (Figure 2). Because C is a left-most min cut, C contains j^\bullet if and only if

$$d_{j^\bullet} > \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i}.$$

Claim. It holds that

$$\sum_{i \in I^=} b_i \geq \sum_{j^\bullet \in T_1} \min \left\{ d_{j^\bullet} - \sum_{i \in \Gamma^+(j^\bullet) \setminus I} c_{j^\bullet i}, \sum_{i \in \Gamma^+(j^\bullet) \cap I^=} c_{j^\bullet i} \right\}. \quad (7)$$

Proof. First, consider the demand that j' has on objects in $I^=$ at price p_t , that is,

$$\sum_{i \in \Gamma^+(j') \cap I^=} c_{j' i}.$$

Comparing q and p_t , remember that the price of any item only increases, that is, $p_t(i) \leq q(i)$, while the prices in $I^=$ remain the same. That is why the buyer j likes to buy at prices q at least the same amount of items from objects in $I^=$ as at prices p_t .

Next, we consider the demand j'' has on objects in I at prices p_t . Recall that j'' gets the same utility from every item in Ω_j'' . Comparing p and q we have that the prices on $I^=$ remain constant while the prices on $I^<$ strictly increase. Thus, j likes to fill up the preferred bundle with items in $I^=$ and maybe with items outside of I before buying the items in $I^<$. Note however, that for prices q it is not clear to which copy of buyer j this demand is assigned. Since furthermore, we have that buyer j cannot buy more objects than available in $I^=$ we get the following lower bound on the demand reassigned from a buyer j'' for prices p_t to buyer j' or j'' for prices q :

$$\min \left\{ d_{j''} - \sum_{i \in \Gamma^+(j'') \setminus I} c_{j'' i}, \sum_{i \in \Gamma^+(j'') \cap I^=} c_{j'' i} \right\}.$$

Finally, if we sum up the demands of all $j' \in N'$ and $j'' \in N''$ at prices q we obtain the following lower bound of the total demand of all buyers $j \in T_1$

$$\begin{aligned} & \sum_{j' \in T_1 \cap N'} \sum_{\substack{i \in I^+(j') \\ i \in I^\neq}} c_{j'i} + \sum_{j'' \in T_1 \cap N''} \min \left\{ d_{j''} - \sum_{\substack{i \in I^+(j'') \\ i \notin I}} c_{j''i}, \sum_{\substack{i \in I^+(j'') \\ i \in I^\neq}} c_{j''i} \right\} \\ & \geq \sum_{j^* \in T_1} \min \left\{ d_{j^*} - \sum_{\substack{i \in I^+(j^*) \\ i \notin I}} c_{j^*i}, \sum_{\substack{i \in I^+(j^*) \\ i \in I^\neq}} c_{j^*i} \right\}. \end{aligned}$$

The set I^\neq is not overdemanded at price q . Thus, the demand of items in I^\neq at price q is smaller or equal than the total supply of all objects in I^\neq . Using the bound of the total demand of all buyers in T_1 , we obtain the inequality of the claim. ■

Next, we show that \tilde{C} is a min cut by comparing the capacities of \tilde{C} and of the left-most min cut C .

$$\begin{aligned} & \text{cap}(C) - \text{cap}(\tilde{C}) \\ &= \left(\sum_{i \in I} b_i + \sum_{j^* \in N' \cup N''} \min \left\{ d_{j^*}, \sum_{i \in I^+(j^*) \setminus I} c_{j^*i} \right\} \right) - \left(\sum_{i \in I^<} b_i + \sum_{j^* \in N' \cup N''} \min \left\{ d_{j^*}, \sum_{i \in I^+(j^*) \setminus I^<} c_{j^*i} \right\} \right) \\ &= \sum_{i \in I^\neq} b_i + \sum_{j^* \in N' \cup N''} \underbrace{\left(\min \left\{ d_{j^*}, \sum_{i \in I^+(j^*) \setminus I} c_{j^*i} \right\} - \min \left\{ d_{j^*}, \sum_{i \in I^+(j^*) \setminus I_2} c_{j^*i} \right\} \right)}_{\substack{=: \alpha \\ =: \beta}} \\ & \quad \underbrace{= 0 \text{ for } j^* \notin T, \text{ since } d_{j^*} \leq \alpha \text{ (by definition) and } \alpha \leq \beta}_{=0} \\ &= \sum_{i \in I^\neq} b_i + \sum_{j^* \in T} \left(\sum_{i \in I^+(j^*) \setminus I} c_{j^*i} - \min \left\{ d_{j^*}, \sum_{i \in I^+(j^*) \setminus I^<} c_{j^*i} \right\} \right) \\ &= \sum_{i \in I^\neq} b_i - \sum_{j^* \in T} \min \left\{ d_{j^*} - \sum_{i \in I^+(j^*) \setminus I} c_{j^*i}, \sum_{i \in I^+(j^*) \cap I^\neq} c_{j^*i} \right\} \stackrel{(7)}{\geq} 0. \end{aligned}$$

Therefore \tilde{C} is a min cut too.

It remains to show that $\tilde{C} \subsetneq C$. This follows directly by the definitions: $\tilde{C} \cap (N' \cup N'') \subseteq T = C \cap (N' \cup N'')$ and $\tilde{C} \cap \Omega = I^< \subsetneq I = C \cap \Omega$, since I^\neq is not empty. ■

Observation 1. The proof of Theorem 2 does not use that Algorithm 2 starts at prices zero. As long as $p_0(i) \leq p^*(i)$ for all $i \in \Omega$, the proof works. Thus, the algorithm finds the minimum competitive price vector p^* if it starts at prices $p_0 \leq p^*$.

3.2 | The auction returns market-clearing prices

To complete the proof of Theorem 1 it remains to show that the prices computed by Algorithm 2 are market-clearing, that is, that an allocation exists where $D = \min \left\{ \sum_{i \in \Omega} b_i, \sum_{j \in N} d_j \right\}$ is sold.

Theorem 3. Given prices p^* computed by Algorithm 2, there exists an allocation $x^* \in \mathbb{Z}_+^{\Omega \times N}$ such that:

1. Any buyer $j \in N$ gets a preferred bundle, that is, $x_{\bullet j}^* \in D_j(p^*)$.
2. If there is an item which is not sold, it has price zero, that is, $\sum_{j \in N} x_{ij}^* = b_i$ for all $i \in \Omega$ with $p^*(i) > 0$.
3. As much as possible is sold, that is, $\sum_{i \in \Omega} \sum_{j \in N} x_{ij}^* = D$.

To show this theorem, we will use the following notation. An allocation \bar{x} is an *extension* of x if $\bar{x}_{ij} \geq x_{ij}$ for all $j \in N$ and $i \in \Omega$. A flow f in $G(p)$ induces an allocation x^f by defining $x_{ij} = f_{ij'} + f_{ij''}$. Note that an allocation induced by a flow in $G(p)$ always distributes subsets of a preferred bundle w.r.t. prices p .

Proof. We construct an allocation with the desired properties by using the flow obtained during Algorithm 2 which uses Algorithm 4 as described in Section 2.4. The constructed allocation x^* will be induced by the max flow in $G(p^*)$, thus property 1 is fulfilled by definition.

We extend x^* to an allocation where every item with a positive price is sold. For $i \in \Omega$ with $\sum_{j \in N} x_{ij}^* < b_i$ we assign the remaining items to buyers j who demanded items of this objects in the last iteration where the price on

i was increased, that is, in the last iteration with $i \in I$. More precisely, we assign it to those buyers for which the flow was removed. Since the price is not increased in later iterations, the object is still in $\Omega_j'''(\mathbf{p}^*)$. By Lemma 5 we know that the removed flow did not exceed $d_j - d_j(\mathbf{p}^*) - d_{j'}(\mathbf{p}^*)$, since $d_j + d_{j'}$ just decreases during the algorithm. Hence, the demand is not exceeded if we assign items according to the removed flow.

This means all items with a positive price are assigned. This is due to the fact that objects in I are fully assigned, since I is given by a left-most min cut. The flow on (b_i, t) remains if we do not remove flow through i which can only happen if $i \in I$. Thus, all items i with a positive price are assigned, either induced by the max flow or by the removed-flow in the last iteration with $i \in I$.

It remains to extend this assignment to an allocation where as much as possible is sold. All unsold items have price 0 and all players with left over demand have payoff 0 on all these items. This means we get a complete bipartite graph in which every maximum flow assigns either all remaining objects or satisfies all demands. Since all prices and payoffs are 0, any such induced allocation can just be added to the computed assignment. ■

Note that the allocation can also be computed by Algorithm 3. We mainly use the described procedure to prove that there is an allocation with the desired properties. This answers the question why the algorithm returns market-clearing prices.

Proof of Theorem 1. The prices are the minimum competitive prices by Theorem 2. By Theorem 3 there is an associated stable allocation where D is sold and where each item with positive price is sold. This existence together with the construction of the network $H(\mathbf{p}^*)$ gives us that Algorithm 3 computes such an allocation. The existence of an allocation where as much as possible is sold also implies that the minimum competitive prices computed by Algorithm 2 are buyer-optimal Walrasian prices. ■

4 | MONOTONICITY

Algorithm 2 determines for each instance the unique component-wise minimum Walrasian price vector \mathbf{p}^* . In this section, we analyze the monotonicity of the auction with respect to changes of the demand and supply. This research question belongs to the field of comparative statics, which is commonly used in economics to analyze markets. In particular, we show that, as intuitively expected, the auction is monotone in the sense that the returned prices can only increase if the demand increases or the supply decreases.

Theorem 4. *Given an instance with valuations v , demands \mathbf{d} , supplies \mathbf{b} and the corresponding buyer-optimal Walrasian prices \mathbf{p} , consider a second instance with the same valuation functions but increased demands and decreased supplies, that is, demands \mathbf{d}^{new} with $0 \leq d_j \leq d_j^{\text{new}}$ for all $j \in N$ and supplies \mathbf{b}^{new} with $0 \leq b_i^{\text{new}} \leq b_i$ for all $i \in \Omega$ with buyer-optimal Walrasian prices \mathbf{p}^{new} . Then we have $p(i) \leq p^{\text{new}}(i)$ for all $i \in \Omega$ with $b_i^{\text{new}} > 0$.*

The proof idea is that prices remain competitive when the supply increases or the demand decreases. Since for truncated additive valuation functions minimum competitive prices are unique and market-clearing, the buyer-optimal Walrasian prices are smaller or equal to any competitive prices.

To prove the theorem we show two lemmas, analyzing the change of the demand and the supply separately.

Lemma 8. *Let \mathbf{d} and \mathbf{d}^{new} be two demand vectors with $d_j \leq d_j^{\text{new}}$ for all $j \in N$ (here it is possible that $d_j = 0$ for some $j \in N$). Then the buyer-optimal Walrasian prices \mathbf{p} at demand \mathbf{d} are not greater than the buyer-optimal Walrasian prices \mathbf{p}^{new} at demand \mathbf{d}^{new} , that is, $p(i) \leq p^{\text{new}}(i)$ for all $i \in \Omega$.*

Proof. Consider an integral max flow f' at prices \mathbf{p}^{new} in the auxiliary flow network $G(\mathbf{p}^{\text{new}})$. Recall that flow f' corresponds to a feasible allocation from buyers to their preferred bundles. Now adapt this flow as follows. For each buyer j with $d_j < d_j^{\text{new}}$, among the paths going through j' or j'' , select one with a currently lowest payoff and reduce flow on that path, until flow through j' and j'' gets reduced to d_j . This procedure terminates with a flow meeting demands d_j . Furthermore, since we still allocate the items with the highest payoff to a buyer, each buyer is allocated to a preferred bundle at prices \mathbf{p}^{new} at demand \mathbf{d} . Thus, \mathbf{p}^{new} is a competitive price vector for demand \mathbf{d} . By Theorem 2 it follows that \mathbf{p} is not only the minimum Walrasian price vector but also the component-wise minimum competitive one. Thus, we get $p(i) \leq p^{\text{new}}(i)$ for all $i \in \Omega$. ■

Next we show that the minimum competitive prices are bigger if the supply is smaller.

Lemma 9. Let \mathbf{b} and \mathbf{b}^{new} be two supply vectors with $b_i^{\text{new}} \leq b_i$ for all $i \in \Omega$ (here it is possible that $b_i^{\text{new}} = 0$ for some $i \in \Omega$). Then for the corresponding buyer-optimal Walrasian prices \mathbf{p} and \mathbf{p}^{new} , it holds that $p(i) \leq p^{\text{new}}(i)$ for all i with $b_i^{\text{new}} > 0$.

Proof. Assume without loss of generality that \mathbf{b} and \mathbf{b}^{new} only differ in the supply of object ℓ by one item. We fix the allocation computed by the minimum competitive prices at supply \mathbf{b}^{new} and we consider two cases.

First, we consider the case $b_\ell^{\text{new}} > 0$. Given the assigned bundles, we analyze the behavior of the buyers when the additional item of object ℓ arrives. If there is a buyer j who is not assigned to one of her preferred bundles at supply \mathbf{b} , we knew that she is the only one who is assigned to items of object ℓ since otherwise the prices \mathbf{p}^{new} are not competitive. Thus, all other buyers are assigned to one of their preferred bundles at supply \mathbf{b} . Hence we can change the preferred bundle of buyer j by assigning one more item of ℓ to her, if necessary by omitting the least profitable item. This change does not harm the other buyers, thus in the new allocation everyone is assigned to a preferred bundle at prices \mathbf{p}^{new} . Thus, prices \mathbf{p}^{new} are competitive for the instance with supply \mathbf{b} . The prices \mathbf{p} are the minimum competitive prices by Theorem 2, which yields $p(i) \leq p^{\text{new}}(i)$.

Next, we consider the case $b_\ell^{\text{new}} = 0$. We adapt the prices \mathbf{p}^{new} to prices $\bar{\mathbf{p}}$ by setting $\bar{p}(i) = p^{\text{new}}(i)$ for $i \in \Omega \setminus \{\ell\}$ and $\bar{p}(\ell) = \max_{j \in B} v_{\ell j} + 1$. Thus, no buyer wants to buy an item of object ℓ . Therefore, the given allocation is an assignment of buyers to one of their preferred bundles at prices $\bar{\mathbf{p}}$ for supply \mathbf{b} . Using again that \mathbf{p} is the minimum competitive price vector at supply \mathbf{b} , we get $p(i) \leq \bar{p}(i) = p^{\text{new}}(i)$ for all $i \in \Omega \setminus \{\ell\}$. ■

Proof of Theorem 4. For a given modified instance we can construct an intermediate instance where only the demand is changed and the supply remains as in the original instance. With Lemma 8 this implies that prices only increase compared to the original instance. Now applying Lemma 9 to the intermediate instance gives the statement of the theorem. ■

The monotonicity allows for faster reoptimization by starting with the old Walrasian price vector.

Corollary 2. Given prices \mathbf{p} we can compute \mathbf{p}^{new} by applying at most $\|\mathbf{p} - \mathbf{p}^{\text{new}}\|_\infty$ iterations of Algorithm 2 with start prices $p(i)$ for all $i \in \Omega$ with $b_i^{\text{new}} > 0$.

Theorem 1 allows starting with any initial price vector which is in every component at most as large as the minimum competitive price vector. Thus, Theorem 4 allows us to start Algorithm 2 at the price $p(i)$ for all $i \in \Omega$ with $b_i^{\text{new}} > 0$. Murota et al. [15] show that the number of iterations is then bounded by $\max\{p(i) - p^{\text{new}}(i) \mid i \in \Omega, b_i^{\text{new}} > 0\}$. However, the following example shows that we cannot bound $\|\mathbf{p} - \mathbf{p}^{\text{new}}\|_\infty$ even if the demand or supply is only slightly changed:

Example 2. Consider an instance with two buyers and two objects. The valuation of both buyers are (M, M) , the demand for both is two and the supply of both objects is two. In this instance $\mathbf{p} = (0, 0)$ are buyer-optimal competitive prices. Now assume the demand of one buyer is increased by one. In this case, the unique buyer-optimal prices are $\mathbf{p}^{\text{new}} = (M, M)$ and thus $\|\mathbf{p} - \mathbf{p}^{\text{new}}\|_\infty = M$. The same happens if the supply of one item is decreased by one.

It is worth pointing out that monotonicity is not guaranteed if the valuation functions are not gross substitute.

Example 3. Consider an instance with two buyers $N = \{1, 2\}$ and two objects $\Omega = \{\alpha, \beta\}$. The supply of each object is 1, the demand of both buyers is unbounded. The valuations are not given element-wise but as a set function by

$$v_i(\emptyset) = v_i(\{\alpha\}) = v_i(\{\beta\}) = 0, \quad v_i(\Omega) = 1$$

for $i \in N$. The valuations are *complementarities*, that is, here buyers want to buy items only in a bundle with the other object otherwise they have no value to the buyers (think of left-hand gloves and right-hand gloves). Gross substitutes always have the no-complementarities condition (see [10]). Any price vector that yields a total price of 1, together with an allocation that gives both objects to the same buyer is at equilibrium. However, if we reduce the supply of either object to 0, the unique Walrasian price vector is $(0, 0)$ as each object on its own is worthless to both buyers.

5 | EMBEDDING INTO EXISTING LITERATURE

The goal of this section is to show the connection between our algorithm and the existing work. In the first part we consider the connection to ascending auctions for strong gross substitute valuations. It turns out that computing a left-most min cut in our network directly corresponds to finding the inclusion-wise minimal set defining a steepest descent direction of the Lyapunov

function in case of truncated additive valuation functions. In the second part we address the question whether the computed prices can be considered as VCG prices (see below for definition). While this does not hold for gross substitute valuations it is shown for single-unit demands. Unfortunately, for truncated additive valuations no mechanism based on prices per item can determine VCG prices as we show later on.

5.1 | Comparison to ascending auction

Note that in our model the Lyapunov function can be rewritten to

$$L(\mathbf{p}) = \max \sum_{j \in N} \sum_{i \in \Omega} (v_{ij} - p(i))x_{ij} + \sum_{i \in \Omega} b_i \cdot p(i)$$

subject to $\sum_{i \in \Omega} x_{ij} \leq d_j$ and $x_{ij} \in [0, b_i]$ for all $j \in N$ and all $i \in \Omega$.

Proposition 2. *In our model with truncated additive valuation functions the difference of the Lyapunov function in an augmentation step equals the difference between the capacity of the s -leaving arcs and the min cut value. More formally*

$$L(\mathbf{p}) - L(\mathbf{p} + \chi_X) = \sum_{j^* \in \Gamma^-(X)} \max \left\{ 0, d_{j^*} - \sum_{i \in \Gamma^+(j^*) \setminus X} c_{j^*i} \right\} - \sum_{i \in X} b_i.$$

The first sum is the sum over the differences $V_j(\mathbf{p}) - V_j(\mathbf{p} + \chi_X)$ for each buyer. The difference corresponds to the total amount of items that j wants to buy from X under prices \mathbf{p} without any alternative outside of X . This yields the equation by definition of the Lyapunov function.

Lemma 10. *Given prices \mathbf{p} , the overdemanded set I determined by the left-most min cut in $G(\mathbf{p})$ minimizes $L(\mathbf{p} + \chi_X)$ among all $X \subseteq \Omega$.*

Proof. A set $X \subseteq \Omega$ minimizes $L(\mathbf{p} + \chi_X)$ if and only if it maximizes

$$L(\mathbf{p}) - L(\mathbf{p} + \chi_X) = \sum_{j \in N} (V_j(\mathbf{p}) - V_j(\mathbf{p} + \chi_X)) - \sum_{i \in X} b_i.$$

Now we consider again the constructed auxiliary networks $G(\mathbf{p})$ and $G(\mathbf{p} + \chi_X)$ and the induced changes in capacities. With Proposition 2, we obtain that $X \subseteq \Omega$ maximizes

$$\sum_{j^* \in \Gamma^-(X)} \max \{0, d_{j^*} - \sum_{i \in \Gamma^+(j^*) \setminus X} c_{j^*i}\} - \sum_{i \in X} b_i \quad (8)$$

Given an X that minimizes $L(\mathbf{p} + \chi_X)$, we construct the cut

$$C_X = \{s\} \cup \{j^* \in N' \cup N'' \mid d_{j^*} > \sum_{i \in \Gamma^+(j^*) \setminus X} c_{j^*i}\} \cup X.$$

The structure of the cut, determined by our algorithm from set I , is given by

$$C = \{s\} \cup \{j^* \in N' \cup N'' \mid d_{j^*} > \sum_{i \in \Gamma^+(j^*) \setminus I} c_{j^*i}\} \cup I.$$

We can show that C_X is also a min cut:

$$\begin{aligned} 0 &\geq \text{cap}(C) - \text{cap}(C_X) \\ &= \sum_{i \in I} b_i + \sum_{j^* \in N' \cup N''} \min \left\{ d_{j^*}, \sum_{i \in \Gamma^+(j^*) \setminus I} c_{j^*i} \right\} - \sum_{i \in X} b_i - \sum_{j^* \in N' \cup N''} \min \left\{ d_{j^*}, \sum_{i \in \Gamma^+(j^*) \setminus X} c_{j^*i} \right\} \\ &= \sum_{i \in I} b_i + \sum_{j^* \in N' \cup N''} \left(d_{j^*} - \max \left\{ 0, d_{j^*} - \sum_{i \in \Gamma^+(j^*) \setminus I} c_{j^*i} \right\} \right) \\ &\quad - \sum_{i \in X} b_i - \sum_{j^* \in N' \cup N''} \left(d_{j^*} - \max \left\{ 0, d_{j^*} - \sum_{i \in \Gamma^+(j^*) \setminus X} c_{j^*i} \right\} \right) \\ &= \left(\sum_{j^* \in N' \cup N''} \max \left\{ 0, d_{j^*} - \sum_{i \in \Gamma^+(j^*) \setminus X} c_{j^*i} \right\} - \sum_{i \in X} b_i \right) - \left(\sum_{j^* \in N' \cup N''} \max \left\{ 0, d_{j^*} - \sum_{i \in \Gamma^+(j^*) \setminus I} c_{j^*i} \right\} - \sum_{i \in I} b_i \right) \\ &\geq 0, \end{aligned}$$

Algorithm 5. Adapted step length algorithm**Input:** Buyer with valuations and demand, objects with supply**Output:** Buyer optimal Walrasian prices $p := 0$ f max flow in network $G(p)$ C left most min cut in $G(p)$ **while** $val(f) < D_p$ **do** $\alpha_{up} := v_{\max}$ $\alpha_{lo} := 1$ $\alpha := \lfloor \frac{\alpha_{up} + \alpha_{lo}}{2} \rfloor$ **while** $\alpha_{up} - \alpha_{lo} > 0$ **do** Adapt network and flow to price $p + \alpha \cdot \chi_{C \cap \Omega}$

Use breadth first search (BFS) in residual network

if min cut is equal to C **then** $\alpha_{lo} := \alpha$ $\alpha_{up} := \alpha$ Set $p := p + \alpha \cdot \chi_{C \cap \Omega}$ Adapt network and flow to p Augment flow f as long as possible in $G(p)$

(for each augmentation step we use BFS in residual network and augment along the path found)

 C new left most min cut (this is all nodes reachable by BFS in residual network)**return** p

where the first inequality follows from the fact that C is a min cut and the last inequality follows from (8), that is, that X maximizes the term. Hence, C_X is a min cut. Moreover the choice $X = C \cap \Omega$ minimizes $L(p + \chi_X)$, since equality holds in the chain of inequalities. ■

5.2 | Adapted step length

One natural approach to speed up the computation of buyer-optimal prices is to increase the length of the augmentation steps. In Algorithm 2 this means given a left-most min cut, increase the prices of the corresponding objects until the min cut changes. In the more general case of strong substitute valuation functions, this means a steepest descent direction of the Lyapunov function is used as long as it remains the steepest descent direction (see [20] section 3, Theorem 4.17). The fact that increasing the prices on the objects in a left-most min cut as long as possible is indeed a special case of the algorithm of Shioura follows by Lemma 10.

How to find this step length is not immediate. It is not enough to ask all buyers for the information when the preferred sets Ω' , Ω'' and Ω''' change. It is well possible that these sets change but the items on which the prices should be increased may not. One possibility to find the correct step length is using a binary search. This comes with the drawback that it destroys the natural process of increasing prices but with the advantage that it allows us to use the same oracle model as before.

That a binary search can be used for determining the step length is possible by the following proposition of Shioura [20].

Proposition 3 ([20, Proposition 4.16]). *In an ascending auction, whenever a steepest descent direction of the Lyapunov function becomes infeasible one of the following two things happens: either the support of the direction increases, or the slope with which the Lyapunov function changes decreases.*

For us, this means whenever the set determined by a cut is not the inclusion-wise minimal maximum overdemanded set any longer, then either the number of objects in the left-most min cut increases or the overdemandedness decreases. Due to the resulting monotonicity we can apply binary search to determine the step length.

This observation allows to bound the number of price raising steps as done by [20] in a similar way for the general setting.

Lemma 11. *The running time of an ascending flow auction with adapted step length as described in Algorithm 5 is given by $\mathcal{O}(\log(v_{\max})|\Omega|D_0|N|(|\Omega| + TO))$, where $v_{\max} = \max_{j \in N, i \in \Omega} v_{ij}$ and $D_0 = \sum_{j \in N} d_j$.*

Proof. To prove the statement we will show two claims from which it follows directly. We start by bounding the time of an update step. As an update step, we summarize a flow augmentation (by a value of one), a min cut computation, and an update of the network and the flow from one iteration to the next. In the next step we will bound the total

number of update steps needed. Note that the number of update steps does not correspond to the iterations of the while-loop in Algorithm 5.

Claim. Every update step takes time at most $\mathcal{O}(|N|TO + |N||\Omega|)$.

Proof. To increase the flow by one unit (or more), we use a breadth first search (BFS) in the residual network. The same is true for the computation of a min cut. The BFS runs in time $\mathcal{O}(|N||\Omega|)$.

By Lemma 4 the adaption of the network runs in time $\mathcal{O}(|N|TO + |N||\Omega|)$. ■

Now, we bound the number of update steps.

Claim. The number of update steps is bounded by $\mathcal{O}(\log(v_{\max})|\Omega|D_0)$.

Proof. In total the auction finds some prices p^* and some flow f^* such that $\text{val}(f^*) = D_{p^*}$ or in other words $\text{val}(f^*) - D_{p^*} = 0$. It starts with the all 0-flow and D_0 . During the algorithm $\text{val}(f) - D_p$ is nonincreasing (see Corollary 1). If it decreases in an update step it decreases by at least by one.

Using Proposition 3 we can show that this value decreases all $\mathcal{O}(\log(v_{\max})|\Omega|)$ update steps.

Given some network and some flow, if we can augment the flow, clearly the value $\text{val}(f) - D_p$ decreases. So assume it is not possible to augment the flow. With one update step we can recognize this situation and compute a min cut giving an inclusion-wise minimal maximum overdemanded set. It is possible that we adapt the network in the next update step, but in the resulting graph, the adapted flow is already maximum, that is, the min cut does not change. Using a binary search, we can find in $\mathcal{O}(\log(v_{\max}))$ time a point where the min cut changes.

It is still possible that the left-most min cut changes, but the value of the min cut and thus of the flow remains the same. Using Proposition 3, we know that the left-most min-cut changes monotonically. By the structure of the network, the left-most min cut is determined by the objects in the cut. Thus, a left-most min cut with the same value can be seen at most $|\Omega|$ times. Thus, after $\mathcal{O}(\log(v_{\max})|\Omega|)$ update steps, the flow needs to be augmentable. ■

This finishes the proof since the algorithm does at most $\mathcal{O}(\log(v_{\max})|\Omega|D_0)$ update steps of cost $\mathcal{O}(|N|TO + |N||\Omega|)$ each. ■

5.3 | VCG prices

When assuming that every object is owned by a seller and the revenue of the seller is the sum of prices he collects from buyers, then every Walrasian equilibrium is socially optimal (see e.g., [17]). “Socially optimal” means that the utility of the buyers plus the utility of the seller is maximum among all allocation and price combinations (p, x) . This observation follows since in a Walrasian equilibrium (p^*, x^*) every buyer maximizes her payoff given prices p^* . The sum of the prices cancels out since the sellers get what the buyers pay. Therefore, the Walrasian equilibrium is socially optimal.

A famous mechanism in auction theory is the VCG mechanism (see for an introduction [18] and [4, 9, 21] for the original papers). It takes bids or valuation functions of the buyers and computes an allocation and prices. The VCG mechanism is essential for auction theory, since it has two nice properties. First, the computed allocation, is socially optimal and second, the mechanism is truthful. This means, for a buyer it never pays off to lie, that is, when reporting the true valuations a buyer j is never worse off than with reporting a fake valuation. Moreover, there exist some reported valuations of the other buyers where it is strictly better for buyer j to report her true valuations than to lie.

Single-unit matching markets. As introduced in Section 1, a famous special case of our model is the classical matching market (in the literature also commonly referred to as housing market) model, where buyers have unit demand. Interestingly, in this restricted model, the buyer-optimal Walrasian prices coincide with prices that are computed with the VCG mechanism ([4, 9, 21]) using Clarke’s pivot rule (see [14]). We call these prices *VCG prices*. The VCG mechanism is known for the fact that it is a dominant strategy for all buyers to announce their valuations truthfully. However, note that this is only true for a sealed-bid auction and not necessarily an ascending auction, in which buyers might learn about other buyers’ valuations as the auction progresses and are hence able to react and adjust their bidding strategy. Note that the ascending auction could be seen as a proxy auction to determine VCG prices when buyers submit their valuation for each item to a proxy bidder who answers all queries of the auctioneer truthfully.

The property that minimal Walrasian prices coincide with VCG prices does not hold for markets where buyers have non-unit demand. In Proposition 4, we show that VCG prices in such markets might require that different buyers get charged different prices for the same object. Moreover, it has already been observed in [11] that an auctioneer cannot get enough information in an ascending auction about buyers’ valuation functions to determine VCG prices in those markets.

We show that for truncated additive valuations, no prices that are based on a per-object price can coincide with the VCG prices. But before, we show that the very natural approach to copy buyers and give them unit demand does not lead to VCG prices.

Let us revisit Example 1 from Section 1 and observe that the duplication method is not truthful, that is, it pays off to lie, and thus does not lead to VCG prices. For the buyer it would be strictly better to report the sets according to the valuation $v_{\bullet 1} = (1, 1)$.

Proposition 4. *There is an example with buyers having truncated additive valuations where there exists no mechanism to determine prices per object that coincide with the VCG prices.*

The following example shows that we need prices per bundle to obtain VCG prices, so the proposition holds.

Example 4. Let $N = \{1, 2, 3\}$ and $\Omega = \{\alpha, \beta\}$ with demands $d_1 = d_2 = 2, d_3 = 1$, supplies $b_\alpha = 3, b_\beta = 2$, and valuations $v_1 = (3, 1), v_2 = (2, 0)$, and $v_3 = (0, 1)$. It is easy to see that both feasible allocations (buyer 1 or 2 gets two items of object α , the other one gets one item of each object, buyer 3 gets one item of object β in both allocations) are socially optimal. Say buyer 1 gets two items of α , buyer 2 gets α only once. Then the net effect of buyer 1 on buyer 2 (i.e., the cost that buyer 1 imposes on buyer 2 by her presence) is 2, so the only possible price for object α is 1. Since buyer 2 also has to purchase an item of α once, object β needs to be subsidized (i.e., assigned a price of -1) to give her a total cost of 0. However, this would give buyer 3 total cost of -1 , which is not VCG.

As an immediate observation from the last example, we can also see that it would have been better for buyer 2 to submit her preferred bundles according to a false valuation function $\tilde{v}_2 = (0, 0)$, which would have resulted in prices $p = (0, 0)$ and the same allocation. Hence, she would have gotten a payoff of 2. The prices we would have obtained by the auction with truthful reporting would have been $(2, 0)$ which yields a payoff of 0 for buyer 2. Thus, the ascending auction does not incentivize truthful reporting of demand sets.

6 | CONCLUSION AND OUTLOOK

We present a network interpretation of an ascending auction for a multi-unit market where buyers have truncated additive valuations. We show that by iteratively raising the prices on a left-most min cut we can compute buyer-optimal Walrasian prices via an ascending auction. The new part here is the simple and efficient flow-based algorithm to determine the sets on which prices should be raised in the ascending auction, namely the minimal maximum overdemand sets. For the special case of unit demand a nice matching based algorithm was known to compute these sets. For the general case of strong gross substitute valuations, in prior literature, this question was either not addressed or the computation was done by using tools from convex analysis. We currently work on a more natural and direct approach to compute the minimal maximum overdemand sets also in the general setting.

With our approach we are, moreover, able to reuse computations from previous iterations to speed up the computation. Combining this with the algorithm using adapted step length allows for an improved running time analysis for the complete auction. Still the resulting running time is not polynomial but only pseudo-polynomial. While buyer-optimal Walrasian prices can be computed directly by using an LP, it remains open whether there is an iterative ascending auction that runs in polynomial time. Note that when using an LP approach, the buyers need to reveal their whole valuation function, while they only reveal their preferred bundles in the more natural iterative auction.

Our approach enabled us to show that the minimum Walrasian prices coincide with the minimum competitive prices for truncated additive valuations. This connection seems very natural but was never discussed before. It is the main part of our proof to show monotonicity properties, that is, that buyer-optimal Walrasian prices can only increase when the supply decreases or the demand increases. We are working on achieving a similar result for the general case of strong gross substitute valuations. Also in this case the connection between minimum Walrasian prices and minimum competitive prices seems intuitive and would imply monotonicity results.

Lastly, the number of iterations necessary to reach buyer-optimal Walrasian prices for a slightly perturbed instance is not polynomial bounded for the perturbations we considered. It is an interesting question whether there are perturbations where the algorithm reaches buyer-optimal Walrasian prices again with only constantly many, or only polynomial many update steps.

ACKNOWLEDGMENTS

We would like to thank two anonymous reviewers who raised interesting questions and provided further insights, helping us to improve multiple sections of the article. Open Access funding enabled and organized by Projekt DEAL.

DATA AVAILABILITY STATEMENT

Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

ORCID

Katharina Eickhoff  <https://orcid.org/0000-0002-9809-8821>

S. Thomas McCormick  <https://orcid.org/0000-0002-2828-4100>

Britta Peis  <https://orcid.org/0000-0002-8938-8843>

Niklas Rieken  <https://orcid.org/0000-0001-7018-8692>

Laura Vargas Koch  <https://orcid.org/0000-0002-7499-5958>

REFERENCES

- [1] L. M. Ausubel, *An efficient dynamic auction for heterogeneous commodities*, Am. Econ. Rev. **96** (2006), no. 3, 602–629.
- [2] O. Ben-Zwi, *Walrasian's characterization and a universal ascending auction*, Games Econom. Behavior **104** (2017), 456–467.
- [3] L. Chen, R. Kyng, Y. P. Liu, R. Peng, M. P. Gutenberg, and S. Sachdeva, “Maximum flow and minimum-cost flow in almost-linear time,” 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS), IEEE, Denver, Colorado, 2022, pp. 612–623.
- [4] E. H. Clarke, *Multipart pricing of public goods*, Public Choice **11** (1971), no. 1, 17–33.
- [5] G. Demange, D. Gale, and M. Sotomayor, *Multi-item auctions*, J. Polit. Econ. **94** (1986), no. 4, 863–872.
- [6] D. Easley, J. Kleinberg, et al., *Networks, crowds, and markets*, Vol **8**, Cambridge University press, Cambridge, 2010.
- [7] K. Eickhoff, B. Peis, N. Rieken, L. Vargas Koch, and L. A. Végh. Faster ascending auctions via polymatroid sum. arXiv preprint, arXiv:2310.08454 2023.
- [8] S. Fujishige and Z. Yang, *A note on Kelso and Crawford's gross substitutes condition*, Math. Oper. Res. **28** (2003), no. 3, 463–469.
- [9] T. Groves, *Incentives in teams*, Econometrica **41** (1973), no. 4, 617–631.
- [10] F. Gul and E. Stacchetti, *Walrasian equilibrium with gross substitutes*, J. Econ. Theory **87** (1999), no. 1, 95–124.
- [11] F. Gul and E. Stacchetti, *The English auction with differentiated commodities*, J. Econ. Theory **92** (2000), no. 1, 66–95.
- [12] P. Hall, *On representatives of subsets*, Classic Pap. Combinat. (1987), 58–62.
- [13] A. S. Kelso Jr and V. P. Crawford, *Job matching, coalition formation, and gross substitutes*, Econometrica **50** (1982), no. 6, 1483–1504.
- [14] W. Kern, B. Manthey, and M. Uetz. Note on VCG vs. price raising for matching markets. arXiv preprint, arXiv:1604.04157 2016.
- [15] K. Murota, A. Shioura, and Z. Yang, “Computing a Walrasian equilibrium in iterative auctions with multiple differentiated items,” *International Symposium on Algorithms and Computation*, Springer, Berlin, Heidelberg, 2013, pp. 468–478.
- [16] K. Murota, A. Shioura, and Z. Yang, *Time bounds for iterative auctions: A unified approach by discrete convex analysis*, Discret. Optim. **19** (2016), 36–62.
- [17] R. Paes Leme, *Gross substitutability: An algorithmic survey*, Games Econom. Behavior **106** (2017), 294–316.
- [18] T. Roughgarden. CS269I: Incentives in computer science lecture# 15: The VCG mechanism. 2016.
- [19] L. S. Shapley and M. Shubik, *The assignment game I: The core*, Int. J. Game Theory **1** (1971), no. 1, 111–130.
- [20] A. Shioura, *Algorithms for L-convex function minimization: Connection between discrete convex analysis and other research fields*, J. Oper. Res. Soc. Jpn. **60** (2017), no. 3, 216–243.
- [21] W. Vickrey, *Counterspeculation, auctions, and competitive sealed tenders*, J. Financ. **16** (1961), no. 1, 8–37.
- [22] L. Walras, *Éléments d'économie politique pure ou théorie de la richesse sociale*, F. Rouge, Lausanne, 1874.

How to cite this article: K. Eickhoff, S. T. McCormick, B. Peis, N. Rieken, and L. Vargas Koch, *A flow-based ascending auction to compute buyer-optimal Walrasian prices*, Networks.. (2024), 1–20. <https://doi.org/10.1002/net.22218>