



TECHNISCHE UNIVERSITÄT MÜNCHEN
TUM School of Computation, Information and Technology

Algorithms for Computing Equilibria in Auctions

Maximilian Fichtl

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Matthias Nießner

Prüfer*innen der Dissertation: 1. Prof. Dr. Martin Bichler
2. Prof. Dr. Stefan Weltge
3. Prof. Dr. Kemal Göler

Die Dissertation wurde am 27.02.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 19.06.2023 angenommen.

Abstract

Walrasian or competitive equilibria are the preferred solution concept in many markets and auctions, due to their preferable properties, like welfare maximization, envy-freeness, and core stability. The existence of such equilibria is, however, a highly nontrivial problem. Classical results from microeconomic theory give a positive answer to the existence question in convex markets. In such markets, goods are assumed to be perfectly divisible. However, in many high-stakes auctions observed in practice, goods are indivisible, so these classical results cannot be applied. In markets with indivisible items, very strong assumptions on the structure of the bidders' valuation functions have to be made to ensure the existence of Walrasian equilibrium. From a computational perspective, the problem of finding a Walrasian equilibrium in indivisible markets is a nonlinear discrete optimization problem. Such problems are usually considered to be much harder to solve than their continuous counterparts. Another problem, particularly in single-round auctions, is the bidders' communication complexity of reporting their preferences to the auctioneer: the number of all possible bundles of items is exponential in the number of different goods sold. Thus, providing an explicit list of values for all those bundles is only possible in markets where the number of items is very small. This dissertation focuses on the problems of computing Walrasian equilibrium prices and dealing with the communication complexity in markets with indivisibilities.

The first contribution of this dissertation surveys known conditions for the existence of Walrasian equilibria in markets with indivisibilities, focusing on the gross and strong substitutes conditions, which have been established as a quasi-standard assumption in the field. We explain connections to notions from convex analysis. We also provide a new perspective on a well-known iterative clock auction format by giving an interpretation as a primal-dual algorithm.

Several bid languages have been proposed to solve the communication problem in sealed-bid auctions. A bid language specifies a format bidders must use to report their preferences and aims to satisfy two often conflicting goals: on the one hand, it should be possible to express a large class of valuation functions. On the other hand, the communication complexity should be low, and the language should be intuitive to understand. In the second contribution of this dissertation, we study the Assignment Messages bid language, which was specifically designed to express strong substitutes preferences. We prove that the expressiveness of this bid language does not suffice to express all strong substitutes valuations.

In our third contribution, we jointly consider the two problems of equilibrium computation and preference communication. We propose a new algorithm to compute Walrasian equilibrium prices in markets with strong substitutes bidders, where bidders report their preferences in the Strong Substitutes Product-Mix Auction (SSPMA) bid language. This bid language has originally been developed by Paul Klemperer for the Bank of England during the 2007 financial crisis. It is the only known one that can precisely express all strong substitutes preferences. Our algorithm is based on a novel duality result and uses methods from DC (difference of convex functions) programming. We also provide new geometric insights into strong substitutes valuation functions.

While this dissertation focuses on the results achieved in the field of Walrasian equilibrium computation, the author has also worked on the problem of computing Bayes-Nash equilibria in auction games and the equilibrium problem with budget-constrained buyers. We briefly discuss our results in these areas.

Zusammenfassung

Walrasianische Gleichgewichte sind das bevorzugte Lösungskonzept in vielen Märkten und Auktionen aufgrund ihrer vorteilhaften Eigenschaften, wie Wohlfahrtsmaximierung, Envy-Freeness und Core-Stabilität. Die Existenz solcher Gleichgewichte ist jedoch ein höchst nicht-triviales Problem. Klassische Ergebnisse aus der Mikroökonomie geben eine positive Antwort auf die Existenzfrage in konvexen Märkten, wobei angenommen wird, dass die Güter perfekt teilbar sind. In vielen Auktionen, wie sie in der Praxis durchgeführt werden, sind die Güter jedoch unteilbar, so dass diese klassischen Ergebnisse nicht anwendbar sind. Auf solchen Märkten müssen sehr starke Annahmen über die Struktur der Präferenzen der Bieter getroffen werden, um die Existenz eines Walrasianischen Gleichgewichts zu gewährleisten. Aus algorithmischer Sicht ist die Berechnung eines Walrasianischen Gleichgewichts in Märkten mit unteilbaren Gütern ein nichtlineares diskretes Optimierungsproblem. Derartige Probleme gelten als wesentlich schwieriger zu lösen als ihre kontinuierlichen Gegenstücke. Ein weiteres Problem, insbesondere bei einstufigen Auktionen, ist die Kommunikationskomplexität der Bieter beim Übermitteln ihrer Präferenzen an den Auktionator: Die Anzahl aller möglichen Bündel von Gütern ist exponentiell in der Anzahl der verschiedenen verkauften Güter. Eine explizite Liste von Werten für alle diese Bündel anzugeben ist daher nur in Umgebungen möglich, in denen die Anzahl der Güter sehr gering ist. In dieser Dissertation betrachten wir die Probleme der Berechnung von Walrasianischen Gleichgewichtspreisen und der Kommunikationskomplexität in Märkten mit unteilbaren Gütern.

Der erste in dieser Publikation vorgestellte Forschungsbeitrag gibt einen Überblick über bekannte Bedingungen für die Existenz Walrasianischer Gleichgewichte in Märkten mit Unteilbarkeiten und konzentriert sich dabei auf die Gross- und Strong-Substitutes Bedingungen, die sich als Quasi-Standardannahme in diesem Bereich etabliert haben. Wir zeigen Verbindungen zu Begriffen aus der konvexen Analysis auf. Weiterhin liefert die Interpretation als primal-dualer Algorithmus eine neue Perspektive auf ein bekanntes Clock-Auktionsformat.

Mehrere Bietsprachen wurden in der Literatur vorgeschlagen, um das Kommunikationsproblem bei Ein-Runden-Auktionen zu lösen. Eine Bietsprache spezifiziert ein spezielles Kommunikationsformat, das die Bieter verwenden, um dem Auktionator ihre Präferenzen mitzuteilen. Sie zielt darauf ab, zwei oft widersprüchliche Ziele zu erfüllen: Einerseits soll sie ermöglichen, eine möglichst große Klasse verschiedener Präferenzen auszudrücken. Andererseits soll die Kommunikationskomplexität gering und die Sprache intuitiv zu verstehen sein. Im zweiten in dieser Dissertation vorgestellten Artikel untersuchen wir die Bietsprache der “Assignment Messages”, die speziell für die Kommunikation von Strong-Substitutes Präferenzen entwickelt wurde. Wir beweisen, dass die Ausdruckskraft dieser Gebotssprache nicht ausreicht, um alle Strong-Substitutes Präferenzen auszudrücken.

In der dritten Publikation betrachten wir die beiden Probleme der Gleichgewichtsbe-
rechnung und der Präferenzkommunikation gemeinsam. Wir entwickeln einen neuen
Algorithmus zur Berechnung von Walrasianischen Gleichgewichtspreisen auf Märkten
mit Strong-Substitutes Bietern, bei dem die Bieter ihre Präferenzen mithilfe der Strong-
Substitutes Product-Mix Auction (SSPMA) Bietsprache kommunizieren. Diese Bietspra-
che ist die einzige bekannte, die die Eigenschaft hat, alle Strong-Substitutes Präferenzen
ausdrücken zu können. Unser Algorithmus basiert auf einem neuartigen Dualitätsergeb-
nis und verwendet Methoden der DC-Programmierung (Difference of Convex Functions).
Außerdem präsentieren wir neue geometrische Einsichten in die Strong-Substitutes Be-
dingung.

Während sich diese Dissertation auf die Ergebnisse im Bereich der Berechnung von Wal-
rasianischen Gleichgewichten konzentriert, hat sich der Autor dieser Dissertation auch
mit dem Problem der Berechnung von Bayes-Nash-Gleichgewichten in Auktionsspielen,
sowie dem Gleichgewichtsproblem mit budgetbeschränkten Käufern beschäftigt. Wir dis-
kutieren kurz unsere Ergebnisse in diesen Bereichen.

Acknowledgements

First and foremost, I would like to express my gratitude to my supervisor Martin Bichler for his guidance and support throughout my dissertation journey. I am also very thankful for the opportunity to work on a joint project with Paul Klemperer and Elizabeth Baldwin. I learned a lot from this collaboration.

I also want to thank the second and third reviewers of this dissertation, Stefan Weltge and Kemal Göler, and the chairperson of the doctoral committee, Matthias Nießner.

I would like to express my appreciation to Gregor Schwarz, Laura Mathews, Sören Merting, and all my colleagues at the Chair of Decision Sciences & Systems for their collaboration and the enjoyable working experience.

Working on a Ph.D. during a pandemic is not always easy, and I want to thank my friends Jérôme, Johannes, Lorenzo, Manuel, Philipp, and Raphael for their moral support.

My family's support has been invaluable, and I am grateful to have them in my life. Finally, a special thank you goes to my girlfriend, Daniela, for her constant love and for always being there for me.

Contents

1	Introduction	1
2	Theoretical Background	7
2.1	Markets with Indivisibilities	7
2.2	Desiderata of Auctions and Markets	9
2.3	Walrasian Equilibria	12
2.4	Existence of Walrasian Equilibria	14
2.4.1	Gross and Strong Substitutes	15
2.4.2	Further Conditions for Existence	19
2.5	Computation of Equilibria	21
2.6	Bid Languages for Substitutes Valuations	23
3	Publication 1: Walrasian Equilibria from an Optimization Perspective	27
4	Publication 2: On the Expressiveness of Assignment Messages	47
5	Publication 3: Strong Substitutes: Structural Properties, and a New Algorithm for Competitive Equilibrium Prices	53
6	Discussion and Conclusion	87
	References	91

1 Introduction

Walrasian or competitive equilibria, named after the French economist and pioneer of general equilibrium theory Léon Walras (Walras 1874), are particularly desirable outcomes of markets, as they simultaneously capture multiple economic design desiderata. In a Walrasian equilibrium, prices for goods are such that the demand of the buyers participating in the market is equal to the available supply of goods. Notably, prices in a Walrasian equilibrium are linear and anonymous, meaning that the price for a bundle of goods is equal to the sum of prices of the individual goods, and prices are the same for each market participant. Moreover, Walrasian equilibria are strongly connected to efficient outcomes of the market. An outcome is efficient when the goods are distributed in a way that maximizes the social welfare of all market participants. It is well-known that the equilibrium states of a market precisely coincide with its efficient outcomes.

Classical competitive equilibrium theory considers convex markets (Arrow and Debreu 1954; McKenzie 1959), where goods are assumed to be perfectly divisible. In many real-world markets, however, one cannot justify the assumption of perfect divisibility. Important examples include spectrum auctions (Bichler et al. 2014), electricity markets (Liberopoulos and Andrianesis 2016), truckload transportation (Caplice and Sheffi 2003), industrial procurement (Bichler et al. 2005), and markets for bus routes (Cantillon, Pendorfer, et al. 2005). These applications are all instances of combinatorial auctions, where a central decision maker, the auctioneer, collects information about the bidders' preferences and then decides on a suitable assignment of bundles to bidders, together with the prices they have to pay for them. The importance of understanding such auctions has recently been emphasized by the Nobel prize in economics, awarded to Paul Milgrom and Robert B. Wilson in 2020 for “improvements to auction theory and inventions of new auction formats” (Committee for the Prize in Economic Sciences in Memory of Alfred Nobel 2020).

Since preferences in markets with indivisibilities are defined on a subset of the integer lattice instead of a convex set, we cannot apply classical existence results on Walrasian

equilibria. Consequently, the existence of Walrasian equilibria in markets with indivisibilities is highly non-trivial, and one has to make strong assumptions on the buyers' preference structure to ensure it. In the 1980s, the first general existence condition, called the gross substitutes condition, was introduced by Kelso and Crawford (1982). Intuitively, a bidder has gross substitutes preferences if a price increase of some items does not lead to a decrease in demand for other items. In particular, it does not allow for complementarities between goods. Several generalizations of this condition have been proposed since then. For example, the strong substitutes condition (Milgrom and Strulovici 2009) covers multi-unit markets, and the gross substitutes and complements condition (Sun and Yang 2006; Teytelboym 2014) allows for complementarities between goods.

Quite simultaneously, there have been significant advances in the theory of nonlinear discrete optimization, generalizing results from matroid theory to more general classes of functions that one can efficiently optimize over the integer lattice. In particular, this resulted in the definition of two types of discrete convex functions, called M^{\natural} -convex and L^{\natural} -convex, that share well-known properties like duality and conjugacy with "usual" convex functions (Fujishige and Murota 2000; Murota 2003; Murota and Shioura 1999). Surprisingly, the notions of gross substitutes valuations and M^{\natural} -concave functions turn out to be equivalent (Fujishige and Yang 2003), providing a new perspective on the gross substitutes condition. This insight inspired the development of faster algorithms for solving the equilibrium computation problem – state-of-the-art algorithms are based on ideas from classical convex optimization, like subgradient descent or cutting plane methods (Leme and Wong 2020; Shioura 2017).

Apart from the equilibrium existence question, the preference elicitation problem is a crucial bottleneck when deploying combinatorial auctions (Parkes 2005). To determine a reasonable outcome of an auction, the auctioneer must, of course, collect information about the bidders' preferences. There are two fundamentally different methods of how this can be done. On the one hand, single-round auctions can be understood as a simple two-stage process. First, bidders report their full preferences over bundles of items to the auctioneer. Based on these reported preferences, the auctioneer then determines an outcome. On the other hand, iterative auctions consist of multiple repeated rounds, where in each round partial information about the preferences is requested, for example, the bidders' most preferred bundles at given prices (Segal 2005). The communication complexity of such an iterative auction is relatively low compared to sealed-bid auctions.

Still, there are also downsides: iterative auctions may require many rounds and can thus take weeks or months to terminate. As a recent example, the 3.45 GHz spectrum auction conducted by the Federal Communications Commission of the United States of America took three months for 151 auction rounds (Federal Communications Commission 2022). This also gives bidders more time for potential collusion (Klemperer 2008; Klemperer 2010). A main drawback of single-round auctions is, on the other hand, that it is impractical or even impossible for the buyers to communicate their preferences by explicitly reporting the monetary values they have for every possible bundle of items, as this list grows exponentially in the number of goods sold. A bid language aims to make this communication more efficient, allowing bidders to report primitive bids that express their valuations over bundles more compactly. There are general-purpose bid languages that usually aim at letting bidders express all possible preferences while making the representation of “common” preferences compact (Goetzendorff et al. 2015; Nisan 2005). In contrast, other bid languages try to express a specific class of valuation functions as compactly and intuitively as possible. For the gross substitutes condition and its variants, several bid languages have been proposed (Baldwin and Klemperer 2019; Hatfield and Milgrom 2005; Klemperer 2010; Milgrom 2009). While they all have the property that each valuation expressed in these bid languages satisfies the gross or strong substitutes condition, there is only one known bid language, called the Strong Substitutes Product-Mix Auction (SSPMA) bid language, allowing bidders to express *all* strong substitutes preferences (Baldwin and Klemperer 2021). The SSPMA is a generalization of an auction invented by Paul Klemperer for the Bank of England to quickly sell loans to private Banks during the 2007 financial crisis and is still in use (Klemperer 2008; Klemperer 2010).

Contribution

In this dissertation, we mainly consider the question of computing Walrasian equilibria in markets with indivisibilities. As mentioned, the fact that goods are indivisible makes the existence question much harder to answer. In Publication 1 included in this dissertation, we survey different conditions for existence from the literature. We focus on markets where the bidders’ valuations fulfill the strong substitutes preferences. A natural iterative auction mechanism for determining a Walrasian equilibrium, introduced by Ausubel (2006), exists in this setting. We present a novel interpretation of this iter-

ative auction as a primal-dual algorithm for solving the winner determination problem, making it more accessible to experts in linear programming.

In Publications 2 and 3, we consider markets with strong substitutes bidders, where the preferences are expressed via a specific bid language.

In Publication 2, we analyze the expressive power of Assignment Messages, a bid language introduced by Milgrom (2009), where bidders report preferences in terms of a linear program with structured side conditions. We prove that strong substitutes valuations exist that are not expressible via Assignment Messages. To do so, we provide a novel interpretation of Assignment Messages as min-cost flow problems.

In Publication 3, we study the SSPMA bid language, the only known bid language that can express precisely all strong substitutes preferences. We provide a new algorithm for computing Walrasian equilibria when bidders report their preferences in this bid language. The algorithm uses methods from DC (difference of convex functions) programming. It is based on a novel duality result, expressing the welfare in terms of SSPMA bids in a natural way. We also derive new geometric insights into strong substitutes valuation functions.

Beyond Walrasian Equilibria

The publications included in this thesis focus on questions regarding Walrasian equilibria in indivisible markets. Within the scope of the author’s Ph.D. project, research on different but related topics was also conducted.

Most existence results on Walrasian equilibria with indivisible goods not only require a particular structure of the bidders’ preferences. They also implicitly assume that bidders can pay an arbitrarily large amount for their preferred bundle of goods. In many real-world scenarios, however, bidders have a limited budget. In our paper “Core-Stability in Assignment Markets with Financially Constrained Buyers” (Batziou et al. 2022), we consider the case where bidders have hard budget constraints. We study assignment markets, where each bidder is interested in winning at most one good, and a different seller sells every good. Even in this simple market setting, Walrasian equilibria do generally not exist – the best one can hope for is to find a market outcome satisfying some of the desirable properties of Walrasian equilibria. We propose a new auction format that always terminates in a core-stable outcome. Under additional assumptions, this outcome is also welfare-maximizing among all core outcomes. Moreover, we prove

that without further assumptions, the problem of maximizing welfare among all core-stable outcomes is NP-complete.

While in larger markets, the assumption that bidders report their preferences truthfully is often justified, strategic considerations of individual bidders significantly impact the outcome of auctions with only a few participants. In these settings, a Bayes-Nash equilibrium analysis is more appropriate to predict the outcome of an auction. The famous result by Daskalakis et al. (2009) states that even in standard two-player matrix games, the computation of a Nash equilibrium is PPAD-hard. Exact Nash equilibrium solvers are thus not applicable to auction games due to their high dimensionality. We must rely on more heuristic methods that do not necessarily converge in all games. We use techniques from online learning to approximate Bayes-Nash equilibria in auction games. In “Learning Equilibria in Symmetric Auction Games Using Artificial Neural Networks” (Bichler et al. 2021a), we propose the NPGA (Neural Pseudo-Gradient Ascent) algorithm, which uses neural networks and a pseudo-gradient method to model the bidders’ strategy spaces. In “Computing Bayes Nash Equilibrium Strategies in Auction Games via Simultaneous Online Dual Averaging” (Fichtl et al. 2022), we present the SODA (Simultaneous Online Dual Averaging) algorithm. In contrast to other existing approaches, this algorithm uses the space of distributional strategies (Milgrom and Weber 1985), which allows bidders to randomize over the bid they submit. Experimental results of both approaches are encouraging: in a broad class of different auction games, convergence to a Bayes-Nash equilibrium can be observed. Our experimental results are in sharp contrast to the negative results on the convergence of learning dynamics in general games (see, e.g., Daskalakis et al. (2010) and Vlatakis-Gkaragkounis et al. (2020)).

Outline

This dissertation is structured as follows. Chapter 2 introduces the necessary notation and definitions of auctions and markets. Next, we discuss Walrasian equilibria in markets with indivisibilities and their equivalence to efficient market outcomes. Afterward, we turn to the question of the existence of Walrasian equilibria and present the gross and strong substitutes conditions, which are very commonly used in the literature on indivisible markets. We further give a brief overview of algorithmic methods to compute Walrasian equilibria and discuss common bid languages for gross and strong substitutes valuation functions that are used to reduce the bidders’ communication complexity. In

Chapter 3 of the dissertation, we present our project on the existence and computation of Walrasian equilibria from a linear programming perspective. The next project, which studies the expressive properties of a particular bid language, called Assignment Messages, is included in Chapter 4. Chapter 5 contains our work on computing Walrasian equilibria in markets where bidders use the SSPMA bid language. We conclude this dissertation by discussing our results and possible future research directions in Chapter 6.

2 Theoretical Background

2.1 Markets with Indivisibilities

Auctions, as they are studied in this dissertation, are mechanisms where a single seller (or auctioneer) sells a given bundle of items (or goods) to interested bidders (or buyers). We assume that the number n of bidders participating in the auction is fixed and write $\mathcal{I} = \{1, \dots, n\}$ for the set of bidders. Items can be categorized into m different types, where items of the same type are considered to have equal value for all buyers. We denote the set of different types of goods by $\mathcal{K} = \{1, \dots, m\}$. A bundle of items containing $s(k)$ units of item k will be represented as a vector $\mathbf{s} = (s(1), \dots, s(m)) \in \mathbb{Z}_{\geq 0}^m$. We denote the bundle of items sold in the auction by $\mathbf{t} \in \mathbb{Z}_{\geq 0}^m$, which we call the target supply.

The goal of an auction is to determine an assignment $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_n)$ of bundles $\mathbf{s}_i \in \mathbb{Z}_{\geq 0}^m$ to bidders i , together with prices the bidders have to pay for their bundles. Of course, the total number of goods assigned to bidders cannot be larger than the supply of goods, and we call an assignment \mathbf{S} with $\sum_{i \in \mathcal{I}} \mathbf{s}_i \leq \mathbf{t}$ a feasible assignment. In this dissertation, the relations \leq and \geq for vectors are always meant to hold coefficient-wise. A price vector $\mathbf{p} = (p(1), \dots, p(m)) \in \mathbb{R}_{\geq 0}^m$ comprises the per-unit prices $p(k)$ for all available items $k \in \mathcal{K}$. We will always assume that prices are linear, meaning that bundle $\mathbf{s} \in \mathbb{Z}_{\geq 0}^m$ costs $\sum_{k \in \mathcal{K}} p(k)s(k) = \langle \mathbf{p}, \mathbf{s} \rangle$. Prices will also be anonymous, i.e., the same price vector \mathbf{p} is observed by all agents. We call a tuple (\mathbf{S}, \mathbf{p}) consisting of an assignment \mathbf{S} and a price vector \mathbf{p} an outcome of the auction.

To determine a “good” outcome, the bidders must report information about their preferences to the auctioneer. A bidder’s preference is modeled by a valuation function $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R} \cup \{-\infty\}$, sometimes only called valuation, mapping bundles \mathbf{s} to values $v(\mathbf{s})$. This value can be interpreted as the maximum amount of money the bidder is willing to pay for receiving \mathbf{s} . Throughout this dissertation, we will make the following assumptions on valuation functions v .

1. The value for the empty bundle is 0, i.e., $v(\mathbf{0}) = 0$.
2. Monotonicity: if $\mathbf{s} \in \text{dom } v_i$ and $\mathbf{r} \leq \mathbf{s}$, then $v(\mathbf{r}) \leq v(\mathbf{s})$.
3. The effective domain $\text{dom } v = \{\mathbf{s} \in \mathbb{Z}_{\geq 0}^m : v(\mathbf{s}) > -\infty\}$ is bounded or, equivalently, finite: $|\text{dom } v| < \infty$.

Note that by Property 1, the effective domain of a valuation function is never empty. Property 3 simplifies our exposition in some places. For our purposes, it is without loss of generality: since the target supply \mathbf{t} consists of a finite number of items, it makes no sense for the bidders to report values for bundles containing more items than \mathbf{t} . Our definition of valuation functions implicitly assumes that a bidder's value only depends on the bundle \mathbf{s} they receive, and neither on the values other bidders have for this bundle nor on the bundles other bidders receive. Valuations with this property are referred to as independent private valuations in the literature.

If a bidder receives bundle \mathbf{s} and prices are set to \mathbf{p} , their utility is equal to their value for bundle \mathbf{s} minus the price they have to pay for it:

$$\pi(\mathbf{s}, \mathbf{p}) = v(\mathbf{s}) - \langle \mathbf{p}, \mathbf{s} \rangle.$$

This utility model is called the quasi-linear model in the literature. Given any price vector, one or more bundles maximize the buyer's utility. The demand set consists of exactly those utility-maximizing bundles:

$$D(\mathbf{p}) = \operatorname{argmax}_{\mathbf{s} \in \mathbb{Z}_{\geq 0}^m} \pi(\mathbf{s}, \mathbf{p}).$$

We say that \mathbf{s} is demanded at prices \mathbf{p} if $\mathbf{s} \in D(\mathbf{p})$. The *indirect utility* denotes the bidder's utility for receiving one of the bundles in their demand set:

$$u(\mathbf{p}) = \max_{\mathbf{s} \in \mathbb{Z}_{\geq 0}^m} v_i(\mathbf{s}) - \langle \mathbf{p}, \mathbf{s} \rangle.$$

Whenever we consider multiple bidders, we will write v_i for bidder i 's valuation function and similarly π_i , u_i , and D_i for their (indirect) utilities and demand sets.

Example 1. A well-studied type of valuation functions are unit-demand valuations. A bidder with a unit-demand valuation is interested in winning at most one good: if they

get assigned bundle \mathbf{s} , their value for \mathbf{s} is equal to the maximum value of one of the items contained in \mathbf{s} . Formally, there exists a vector $\mathbf{w} \in \mathbb{R}_{\geq 0}^m$, such that

$$v(\mathbf{s}) = \max_{k: s(k)=1} w(k)$$

if $\mathbf{s} \neq \mathbf{0}$, and $v(\mathbf{0}) = 0$, where $\text{dom } v = \{0, 1\}^m$. One can check that

$$u(\mathbf{p}) = \max \left\{ 0, \max_{k \in \mathcal{K}} w(k) - p(k) \right\}.$$

A market with only unit-demand bidders can be visualized as a weighted bipartite graph between bidders and items, where the edge between bidder i and item k has weight $v_i(\mathbf{e}_k)$, \mathbf{e}_k denoting the k -th standard unit vector. Since assigning more than one item to each bidder does not make sense, outcomes of such markets can be interpreted as matchings between bidders and items. Markets, where all bidders have a unit-demand valuation are called assignment markets.

2.2 Desiderata of Auctions and Markets

This section briefly discusses some typical design desiderata for auctions and markets. The notation mainly follows Bichler (2017) and Blumrosen and Nisan (2007).

By running an auction, the auctioneer aims to determine an outcome that fulfills some specific objective as well as possible. In many applications, the goal is to maximize social welfare, which is defined as the sum of all bidders' valuations: the assignment $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_n)$ generates social welfare $\sum_{i \in \mathcal{I}} v_i(\mathbf{s}_i)$. Social welfare can also be expressed in terms of the utilities of all market participants, including the seller. Assuming that the seller's utility is the total amount of money received from the buyers, the sum of utilities of all agents is equal to

$$\left\langle \mathbf{p}, \sum_{i \in \mathcal{I}} \mathbf{s}_i \right\rangle + \sum_{i \in \mathcal{I}} \pi_i(\mathbf{s}_i, \mathbf{p}) = \left\langle \mathbf{p}, \sum_{i \in \mathcal{I}} \mathbf{s}_i \right\rangle + \sum_{i \in \mathcal{I}} v_i(\mathbf{s}_i) - \langle \mathbf{p}, \mathbf{s}_i \rangle = \sum_{i \in \mathcal{I}} v_i(\mathbf{s}_i).$$

The optimization problem of maximizing social welfare is often called the Winner Determination Problem (WDP):

$$\begin{aligned}
 \max \quad & \sum_{i \in \mathcal{I}} v_i(\mathbf{s}_i) \\
 \text{s.t.} \quad & \sum_{i \in \mathcal{I}} s_i(k) \leq t(k) \quad \forall k \in \mathcal{K} \\
 & \mathbf{s}_i \in \mathbb{Z}_{\geq 0}^m \quad \forall i \in \mathcal{I}.
 \end{aligned} \tag{WDP}$$

Definition 2.2.1 (Efficiency). An allocation \mathbf{S} is called efficient if it maximizes social welfare, i.e., it is an optimal solution to the (WDP).

Since the bidders' valuation functions are their private information and are not known by the auctioneer in advance but are reported during the auction process, the auctioneer has to rely on the reported information to be truthful. To incentivize the truthful behavior of the agents, a typical design desideratum of an auction mechanism is strategyproofness, meaning that bidders can never increase their utility by reporting wrong information about their preferences. We must use a slightly more general notation than in Section 2.1 for a formal definition. A mechanism \mathcal{M} is described by a function f , taking as input the bidders' reported valuations (v_1, \dots, v_n) , and returning an assignment \mathbf{S} of bundles to bidders, together with a payment $\mathbf{q} \in \mathbb{R}_{\geq 0}^n$, q_i denoting the amount of money bidder i has to pay.¹ By abuse of notation, denote by $\pi_i(v_i; f(v'_1, \dots, v'_n)) = v_i(\mathbf{s}_i) - q_i$ the utility obtained by bidder i , if i 's true valuation is v_i and bidders report valuations (v'_1, \dots, v'_n) .

Definition 2.2.2 (Strategyproofness). A mechanism \mathcal{M} is strategyproof, if for all bidders i , all valuation profiles (v_1, \dots, v_n) and all valuation functions v'_i it holds that

$$\pi_i(v_i; f(v_1, \dots, v_n)) \geq \pi_i(v_i; f(v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_n)).$$

Strategyproof mechanisms are also equivalently called dominant-strategy incentive-compatible.

Bidders are assumed to be interested in maximizing their utility. Thus, an outcome (\mathbf{S}, \mathbf{p}) of an auction might be considered unfair if bidders receive bundles \mathbf{s}_i that are

¹Such a mechanism is usually called a direct revelation mechanism in the literature, as bidders directly report their preferences. We restrict our attention to such mechanisms here for simplicity. Generally, a mechanism can be defined as functions taking some general actions of the bidders as input.

suboptimal at given prices, meaning that $\mathbf{s}_i \notin D_i(\mathbf{p})$. In particular, they might envy other bidders for the bundles they receive.

Definition 2.2.3 (Envy-freeness). An outcome (\mathbf{S}, \mathbf{p}) is called envy-free if all bidders i receive a bundle in their demand sets, i.e., $\mathbf{s}_i \in D_i(\mathbf{p})$ for all $i \in \mathcal{I}$.

The fourth design desideratum of auction mechanisms we want to introduce here is core-stability. Let $\mathcal{I}_0 = \mathcal{I} \cup \{0\}$ denote the set consisting of all buyers and the 0-th agent, representing the seller. For any subset $C \subseteq \mathcal{I}_0$, called a coalition, define the coalitional value $w(C)$ to be the maximal social welfare that the agents in C can generate on their own: if $0 \notin C$, there are no goods to be distributed among the buyers, so $w(C) = 0$. Otherwise, if the auctioneer is part of the coalition C , $w(C)$ is just the objective value of the (WDP) where the set of all bidders \mathcal{I} is restricted to the bidders in $C \setminus \{0\}$. We say that a payoff vector $\pi = (\pi_0, \dots, \pi_n) \in \mathbb{R}_{\geq 0}^{n+1}$, assigning a payoff to every agent $i \in \mathcal{I}_0$, lies in the core, or is core-stable, if the following two conditions hold. First, the payoff $\sum_{i \in C} \pi_i$ of every coalition $C \subseteq \mathcal{I}_0$ must be at least as high as the maximal social welfare $w(C)$ the coalition can generate on its own. Second, the sum of all payoffs $\sum_{i \in \mathcal{I}_0} \pi_i$ cannot be larger than the social welfare $w(\mathcal{I}_0)$ all agents generate together. The core thus comprises all possible ways of distributing the jointly generated welfare among all agents, such that no subcoalition C could profit from “ignoring” the agents that are not in C .

Definition 2.2.4 (Core). The core is the following set of payoff vectors.

$$\text{Core}(\mathcal{I}_0, w) = \left\{ \pi \in \mathbb{R}_{\geq 0}^{n+1} : \sum_{i \in C} \pi_i \geq w(C) \forall C \subseteq \mathcal{I}_0 \text{ and } \sum_{i \in \mathcal{I}_0} \pi_i = w(\mathcal{I}_0) \right\}.$$

According to this definition, we call an outcome (\mathbf{S}, \mathbf{p}) , consisting of an assignment \mathbf{S} and linear prices \mathbf{p} , a core outcome, if

$$\left(\left\langle \mathbf{p}, \sum_{i \in \mathcal{I}} \mathbf{s}_i \right\rangle, \pi_1(\mathbf{s}, \mathbf{p}), \dots, \pi_n(\mathbf{s}, \mathbf{p}) \right) \in \text{Core}(\mathcal{I}_0, w).$$

2.3 Walrasian Equilibria

A Walrasian equilibrium is a market state where the total demand for items equals the target supply \mathbf{t} of available items. After giving a formal definition, we will present the essential properties of Walrasian equilibria, explaining why such market states are particularly desirable.

Definition 2.3.1. A Walrasian equilibrium is an outcome (\mathbf{S}, \mathbf{p}) , where \mathbf{S} is a feasible assignment and \mathbf{p} is a price vector, such that $\sum_{i \in \mathcal{I}} \mathbf{s}_i = \mathbf{t}$ and $\mathbf{s}_i \in D_i(\mathbf{p})$ for all $i \in \mathcal{I}$.

By definition, Walrasian equilibria are envy-free. In the following, we will explain the relation between Walrasian equilibria and the other desiderata introduced in Section 2.2.

It turns out that Walrasian equilibria are inherently connected to the market's efficient outcomes. Indeed, the famous welfare theorems state that the set of Walrasian equilibria is equal to the set of solutions of the (WDP) whenever the set of Walrasian equilibria is not empty.

To formally state the welfare theorems, we need to introduce a well-known linear programming relaxation of the (WDP): recall that we assume the sets $\text{dom } v_i$ to be finite. For each $\mathbf{s} \in \text{dom } v_i$, introduce the variable $x_i(\mathbf{s}) \in \{0, 1\}$, indicating whether bidder i wins bundle \mathbf{s} . Relaxing the constraints to $x_i(\mathbf{s}) \in [0, 1]$, we get the following linear programming relaxation of the winner determination problem (RWDP):

$$\begin{aligned}
 \max \quad & \sum_{i \in \mathcal{I}} \sum_{\mathbf{s} \in B_i} v_i(\mathbf{s}) x_i(\mathbf{s}) \\
 \text{s.t.} \quad & \sum_{i \in \mathcal{I}} \sum_{\mathbf{s} \in \text{dom } v_i} s(k) x_i(\mathbf{s}) \leq t(k) \quad \forall k \in \mathcal{K} \\
 & \sum_{\mathbf{s} \in \text{dom } v_i} x_i(\mathbf{s}) \leq 1 \quad \forall i \in \mathcal{I} \\
 & x_i(\mathbf{s}) \geq 0 \quad \forall i \in \mathcal{I} \quad \forall \mathbf{s} \in \text{dom } v_i.
 \end{aligned} \tag{RWDP}$$

The First Welfare Theorem states that every Walrasian equilibrium maximizes the (RWDP). Since assignments \mathbf{S} are always integral, this implies that there is an optimal integral solution to the (RWDP). In particular, if a Walrasian equilibrium exists,

the objective values of the (WDP) and (RWDP) are equal. Blumrosen and Nisan (2007) state the welfare theorems for single-unit markets with indivisible goods, which can directly be extended to the multi-unit setting.

Theorem 2.3.2 (First Welfare Theorem). *Let (\mathbf{S}, \mathbf{p}) constitute a Walrasian equilibrium. Then $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_n)$ is an optimal solution to the winner determination problem. Moreover, by setting $x_i(\mathbf{s}_i) = 1$, we obtain an optimal integral solution to the (RWDP).*

The Second Welfare Theorem states that every integral optimal solution of the (RWDP) is a Walrasian equilibrium.

Theorem 2.3.3 (Second Welfare Theorem). *Suppose there exists an integral optimal solution of the relaxed winner determination problem (RWDP) and let \mathbf{S} be the induced assignment, where every buyer receives the unique bundle \mathbf{s}_i with $x_i(\mathbf{s}_i) = 1$. Then there exists a price vector \mathbf{p} , such that (\mathbf{S}, \mathbf{p}) is a Walrasian equilibrium.*

The welfare theorems imply that a Walrasian equilibrium exists if and only if the integrality gap of the (WDP) is zero.

Having established the efficiency properties of Walrasian equilibria, it is also a folklore result that they lie in the core.

Theorem 2.3.4. *Every Walrasian equilibrium is core-stable.*

The reverse direction is not true: generally, not every payoff vector that lies in the core can be implemented by a Walrasian equilibrium – for example, there may exist core outcomes where bidders pay different prices for identical items (Kaneko 1976).

While Walrasian equilibria are at the same time efficient, envy-free, and core-stable, mechanisms implementing a Walrasian equilibrium are, in general, not strategyproof, as can be seen as follows.

Example 2. Consider a market with two identical goods to be sold and two buyers. Buyer 1's valuation function is defined by $v_1(0) = 0$ and $v_1(1) = 3$, while buyer 2's valuation function is given by $v_2(0) = 0$, $v_2(1) = 2$ and $v_2(2) = 4$. It is easy to see that the unique equilibrium price is $\mathbf{p} = 2$ – at a lower price, the bidders demand 3 items in total, while for $\mathbf{p} > 2$, bidder 2 does not demand any item, and bidder 1 demands at most 1 item.

At $\mathbf{p} = 2$, both bidders receive one item, bidder 1 obtaining a utility of $v_1(1) - 2 = 1$, and bidder 2 obtaining a utility of 0. Now suppose that bidder 2 reported the valuation $v'_2(0) = 0$, $v'_2(1) = 1$ and $v'_2(2) = 2$, instead. Then the unique equilibrium price for v_1 and v'_2 is $\mathbf{p} = 1$, where again both bidders receive one item. Note that bidder 2's utility is now $v_2(1) - 1 = 1$. Thus, by misreporting, bidder 2 can increase their utility.

Nevertheless, it can be shown that with increasing market size, i.e., a larger number of bidders, the potential of a single bidder to manipulate prices becomes arbitrarily small, so that Walrasian equilibria are “almost” strategyproof in large markets. For precise statements on the approximate strategyproofness of Walrasian equilibria in large markets, we refer to Azevedo and Budish (2018), Jackson and Manelli (1997), and Roberts and Postlewaite (1976).

2.4 Existence of Walrasian Equilibria

As demonstrated above, Walrasian equilibria have many advantageous properties. The question of the existence of a Walrasian equilibrium, however, is highly non-trivial when goods are indivisible. It is easy to construct examples where no Walrasian equilibrium exists.

Example 3. Consider a market with two bidders, two types of items, and target supply $\mathbf{t} = (1, 1)$. Bidder 1 is a unit-demand bidder with $v_1(1, 0) = 1$, $v_1(0, 1) = 2$ and $v_1(1, 1) = 2$. Bidder 2 is only interested in winning both items: $v_2(1, 1) = 2$ and $v_2(\mathbf{s}) = 0$ for $\mathbf{s} \neq (1, 1)$. It is easy to see that an optimal solution to the (WDP) is given by $\mathbf{s}_1 = (0, 0)$ and $\mathbf{s}_2 = (1, 1)$ with objective value 2. However, the solution $x_1(1, 0) = x_1(0, 1) = x_2(1, 1) = 0.5$ and all other variables $x_i(\mathbf{s}) = 0$ is feasible for the (RWDP) and has objective value $0.5(1 + 2 + 2) = 2.5 > 2$. Thus, the (WDP) has a non-zero integrality gap, so a Walrasian equilibrium does not exist by the welfare theorems.

To ensure the existence of Walrasian equilibria, additional structural assumptions on the bidders' valuation functions have been extensively studied in the literature. The so-called gross substitutes condition, introduced by Kelso and Crawford (1982), and its multi-unit extension called strong substitutes condition (Milgrom and Strulovici 2009), have been established as a quasi-standard in the literature. Since it is also the underlying assumption of the work presented in this dissertation, we will give a more detailed

overview of these conditions in Section 2.4.1 before briefly describing further existence conditions in Section 2.4.2.

2.4.1 Gross and Strong Substitutes

The gross substitutes condition has been developed for single-unit markets, where exactly one unit of each item is available, i.e., the target supply equals $\mathbf{t} = (1, \dots, 1)$. Consequently, we consider valuation functions v with effective domain $\text{dom } v \subseteq \{0, 1\}^m$. It is often more natural to interpret bundles $\mathbf{s} \in \{0, 1\}^m$ as subsets: we identify $A \subseteq \mathcal{K}$ with the vector $\mathbf{e}_A \in \{0, 1\}^m$ whose k -th entry is 1, if and only if $k \in A$. For singleton sets $\{k\}$, we write \mathbf{e}_k instead of $\mathbf{e}_{\{k\}}$. Intuitively, the gross substitutes condition states that whenever the price for one or more items is raised, the demand for all other items does not decrease.

Definition 2.4.1 (Gross substitutes). A valuation function $v : \{0, 1\}^m \rightarrow \mathbb{R}_{\geq 0} \cup \{-\infty\}$ satisfies the gross substitutes condition, if for every pair of price vectors $\mathbf{p} \leq \mathbf{q}$, and for all $A \in D(\mathbf{p})$, there exists $B \in D(\mathbf{q})$ with $\{k \in A : p(k) = q(k)\} \subseteq B$.

Simple examples of gross substitutes valuation functions are unit-demand valuations that we introduced in Section 2.1, and additive valuation functions, which are of the form $v(\mathbf{s}) = \langle \mathbf{w}, \mathbf{s} \rangle$ for some fixed $\mathbf{w} \in \mathbb{R}_{\geq 0}^m$. Weighted matroid rank functions give a more sophisticated example, generalizing the aforementioned types of valuations. Let I be the family of independent sets of a matroid on \mathcal{K} . Then

$$v(\mathbf{s}) = \max \{ \langle \mathbf{w}, \mathbf{e}_A \rangle : A \in I, \mathbf{e}_A \leq \mathbf{s} \}$$

is a gross substitutes valuation function for any fixed $\mathbf{w} \in \mathbb{R}_{\geq 0}^m$ (Shioura and Tamura 2015).

Gul and Stacchetti (1999) give equivalent characterizations of the gross substitutes condition, providing additional intuition.

Theorem 2.4.2. *Let v be a single-unit valuation function. Then the following statements are equivalent:*

- v satisfies the gross substitutes condition.

- *v* has the single improvement property: For every price \mathbf{p} and every bundle $A \subseteq \mathcal{K}$ with $A \notin D(\mathbf{p})$, there exists a bundle B with $|A \setminus B| \leq 1$ and $|B \setminus A| \leq 1$, such that $\pi(B, \mathbf{p}) > \pi(A, \mathbf{p})$.
- *v* satisfies the no complementarities property: For every price \mathbf{p} , all bundles $A, B \in D(\mathbf{p})$ and every bundle $X \subseteq A$, there exists a bundle $Y \subseteq B$ such that $(A \setminus X) \cup Y \in D(\mathbf{p})$.

Kelso and Crawford (1982) prove that if all buyers have gross substitutes valuation functions, a Walrasian equilibrium exists.

Theorem 2.4.3. *Suppose that the single-unit valuation functions v_i satisfy the gross substitutes condition for all buyers $i \in \mathcal{I}$. Given that for each item $k \in \mathcal{K}$ there is some bidder i with $v_i(\mathbf{e}_k) > -\infty$, there exists a Walrasian equilibrium for the target supply $\mathbf{t} = (1, \dots, 1)$.²*

Importantly, as was shown by Gul and Stacchetti (1999), the class of gross substitutes valuations is maximal in this regard. If we add any additional single-unit valuation function to the set of gross substitutes valuations, the equilibrium existence property from Theorem 2.4.3 is violated.

Theorem 2.4.4. *Let $v_1 : \{0, 1\}^m \rightarrow \mathbb{R}_{\geq 0}$ be a single-unit valuation function violating the gross substitutes condition. Then there exist $\ell - 1$ buyers³ with unit-demand (and thus gross substitutes) valuation functions $v_2, \dots, v_\ell : \{0, 1\}^m \rightarrow \mathbb{R}_{\geq 0}$ such that a Walrasian equilibrium does not exist for the target supply $\mathbf{t} = (1, \dots, 1)$.*

The gross substitutes condition has a straightforward generalization to multi-unit markets introduced by Milgrom and Strulovici (2009) (see also Murota (2016) for the definition we give here). Let $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}_{\geq 0} \cup \{-\infty\}$ be an arbitrary valuation function and $\mu \in \mathbb{Z}_{\geq 0}$ be such that $\text{dom } v \subseteq [0, \mu]^m$. We can transform v into a single-

²In their original paper, Kelso and Crawford (1982) only consider the case where $\text{dom } v = \{0, 1\}^m$. We also allow $\text{dom } v \subsetneq \{0, 1\}^m$ here for technical reasons. This slightly more general statement follows, for example, from Murota (2003, Theorem 11.13).

³The proof by Gul and Stacchetti (1999) implies that $\ell \in \mathcal{O}(m)$.

unit valuation, where each unit of each good is considered as a separate good: define $\tilde{v} : \{0, 1\}^{\mu m} \rightarrow \mathbb{R}_{\geq 0} \cup \{-\infty\}$ by

$$\tilde{v}(s(1, 1), \dots, s(1, \mu), \dots, s(m, 1), \dots, s(m, \mu)) = v \left(\sum_{j=1}^{\mu} s(1, j), \dots, \sum_{j=1}^{\mu} s(m, j) \right).$$

Definition 2.4.5. A valuation function $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R} \cup \{-\infty\}$ satisfies the strong substitutes condition if its single-unit representation \tilde{v} satisfies the gross substitutes condition.

Milgrom and Strulovici (2009) show that the equilibrium existence result directly extends to strong substitutes valuations. When all buyers have strong substitutes valuations, a Walrasian equilibrium exists for every target supply $\mathbf{t} \in \sum_{i \in \mathcal{I}} \text{dom } v_i$.

The definitions of gross and strong substitutes provided above are in rather “economic” terms, requiring that the bidders’ demands behave in a certain way. As the welfare theorems from Section 2.3 suggest, however, there is an alternative approach that is rooted in discrete optimization. The theory of discrete convex analysis transfers notions and methods from convex analysis, like duality and optimization algorithms, to functions defined on the integer lattice \mathbb{Z}^m . For an extensive presentation of discrete convex analysis, we refer to Murota (2003). For a survey focusing on economical applications, see also Murota (2016). Central to discrete convex analysis are different notions of discrete convexity, called M^{\natural} -convexity and L^{\natural} -convexity, which are dual to each other. They were originally introduced in Fujishige and Murota (2000) and Murota and Shioura (1999). For a vector $\mathbf{x} \in \mathbb{Z}^m$, define $\text{supp}^+(\mathbf{x}) = \{i : x(i) > 0\}$ and $\text{supp}^-(\mathbf{x}) = \{i : x(i) < 0\}$.

Definition 2.4.6 (M^{\natural} -convexity). A function $f : \mathbb{Z}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ is called M^{\natural} -convex, if for all $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^m$ and for all $i \in \text{supp}^+(\mathbf{x} - \mathbf{y})$ there exists a $j \in \text{supp}^-(\mathbf{x} - \mathbf{y}) \cup \{0\}$ such that

$$f(\mathbf{x}) + f(\mathbf{y}) \geq f(\mathbf{x} - \mathbf{e}_i + \mathbf{e}_j) + f(\mathbf{y} + \mathbf{e}_i - \mathbf{e}_j)$$

where $\mathbf{e}_0 = (0, \dots, 0) \in \mathbb{Z}^m$.

Definition 2.4.7 (L^{\natural} -convexity). A function $g : \mathbb{Z}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ is called L^{\natural} -convex, or discrete midpoint convex, if for all $\mathbf{p}, \mathbf{q} \in \mathbb{Z}^m$ we have

$$g(\mathbf{p}) + g(\mathbf{q}) \geq g\left(\left\lfloor \frac{\mathbf{p} + \mathbf{q}}{2} \right\rfloor\right) + g\left(\left\lceil \frac{\mathbf{p} + \mathbf{q}}{2} \right\rceil\right).$$

The operators $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ are understood coefficient-wise.

Similar to “usual” convex analysis, we can define the discrete convex conjugate of f by $f^*(\mathbf{p}) = \sup_{\mathbf{x} \in \mathbb{Z}^m} \langle \mathbf{p}, \mathbf{x} \rangle - f(\mathbf{x})$. Central to the theory of discrete convex analysis is the discrete duality theorem: if $f : \mathbb{Z}^m \rightarrow \mathbb{Z} \cup \{+\infty\}$ is proper, integer-valued and $M^{\natural}/L^{\natural}$ -convex, its discrete conjugate f^* is $L^{\natural}/M^{\natural}$ -convex, and $f^{**} = f$. We remark that these results can be extended, such that the assumption of integer-valued functions is not required (Murota 2003). Discrete convex functions also share other characteristic properties with continuous convex functions: M^{\natural} -convex functions are closed under infimal convolution, while L^{\natural} -convex functions are closed under addition. For computational purposes, it is important to note that both types of discrete convex functions can be minimized efficiently via steepest descent algorithms (Shioura 2017).

As usual, we say that a function $f : \mathbb{Z}^m \rightarrow \mathbb{R} \cup \{-\infty\}$ is M^{\natural} - or L^{\natural} -concave, if $-f$ is M^{\natural} - or L^{\natural} -convex. The following theorem, which was originally proven by Fujishige and Yang (2003) for gross substitutes valuations, is the central connection between discrete convex analysis and strong substitutes valuations.

Theorem 2.4.8. *A valuation function $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R} \cup \{-\infty\}$ is strong substitutes if and only if it is M^{\natural} -concave.*

Using discrete duality, this also implies that if a valuation function $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{Z}_{\geq 0} \cup \{-\infty\}$ is strong substitutes, its indirect utility function $u(\mathbf{p})$ is L^{\natural} -convex (see Murota (2003, Chapter 11) and Publication 1 for details). This allows us to use minimization algorithms for L^{\natural} -convex functions to solve the equilibrium computation problem. We provide more details in Section 2.5.

The strong substitutes condition also has an intuitive geometric interpretation, as presented by Baldwin and Klemperer (2019), that we briefly want to sketch here. Given a valuation function v , define the Locus of Indifference Prices (LIP) by $\mathcal{L} = \{\mathbf{p} \in \mathbb{R}_{\geq 0}^m : |D(\mathbf{p})| > 1\}$, consisting of the set of all prices where more than one bundle of items is demanded. Since every price $\mathbf{p} \in \mathcal{L}$ satisfies $v(\mathbf{s}) - \langle \mathbf{p}, \mathbf{s} \rangle = v(\mathbf{r}) - \langle \mathbf{p}, \mathbf{r} \rangle$ for at least two




Figure 2.1: The LIP of a strong substitutes valuation function divides the price spaces into connected regions where a unique bundle is demanded. The normal vectors (gray) are all multiples of $(1, 0)$, $(0, 1)$ or $(1, -1)$.

different bundles $\mathbf{s} \neq \mathbf{r}$, one can show that \mathcal{L} is a union of parts of $(m - 1)$ -dimensional hyperplanes, splitting the price space $\mathbb{R}_{\geq 0}^m$ into disjoint regions (see Figure 2.1). It is possible to characterize strong substitutes valuations by the set of normals of \mathcal{L} . We say that $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}_{\geq 0} \cup \{-\infty\}$ is concave-extensible if there exists a concave function on $\mathbb{R}_{\geq 0}^m$ coinciding with v on every integral point $\mathbf{s} \in \mathbb{Z}_{\geq 0}^m$.

Theorem 2.4.9 (Baldwin and Klemperer (2019), Proposition 3.10). *A valuation $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R}_{\geq 0} \cup \{-\infty\}$ is strong substitutes, if and only if v is concave-extensible and all normal directions of \mathcal{L} can be scaled to lie in the set*

$$\{\pm \mathbf{e}_k : k \in \mathcal{K}\} \cup \{\mathbf{e}_k - \mathbf{e}_l : k \neq l\}.$$

2.4.2 Further Conditions for Existence

A generalization of the gross substitutes condition, which is called the (generalized) gross substitutes and complements (GSC) condition and allows for certain types of complementarities between goods has been proposed by Sun and Yang (2006) and was generalized by Teytelboym (2014). Intuitively, a valuation function satisfies the GSC condition if

one can split the items into two groups, where items of the same group are substitutes and items of different groups are complements. One can show that if all bidders have valuations of this type, a Walrasian equilibrium exists. As for strong substitutes preferences, there exists an alternative characterization of such valuation functions in terms of twisted M^{\natural} -concavity (Ikebe and Tamura 2015).

The geometric interpretation of strong substitutes preferences from above can be generalized, providing a complete characterization of valuation functions admitting a Walrasian equilibrium (Baldwin and Klemperer 2019). Let $\mathcal{D} \subseteq \mathbb{Z}_{\geq 0}^m$ be a set of non-zero primitive vectors with $-\mathcal{D} = \mathcal{D}$. A valuation v is of demand type \mathcal{D} if every facet of its LIP \mathcal{L} has a normal vector in \mathcal{D} . A demand type is called unimodular if every linear independent subset of vectors in \mathcal{D} can be extended to a basis of integer vectors of \mathbb{R}^m with determinant ± 1 . The Unimodularity Theorem by Baldwin and Klemperer (2019) states that a Walrasian equilibrium exists for any target supply $\mathbf{t} \in \sum_{i=1}^n \text{dom } v_i$ if the bidders' valuations v_1, \dots, v_n are concave-extensible, and they all have the same unimodular demand type \mathcal{D} . On the other hand, if \mathcal{D} is not unimodular, there exists a finite set of agents with concave-extensible valuation functions v_i , such that no equilibrium exists for some $\mathbf{t} \in \sum_{i \in \mathcal{I}} \text{dom } v_i$.

So far, all presented results on equilibrium existence define a class of valuation functions and show that Walrasian equilibrium exists whenever all valuations are contained in the respective class. Another stream of literature studies existence conditions given that the exact form of the bidders' valuation functions is known.

Of course, one way to evaluate whether a Walrasian equilibrium exists when the valuation functions are known is to check if the (RWDP) has an integral optimal solution. This approach is originally due to Bikhchandani and Mamer (1997). Ma (1998) provides another existence condition based on the core: they suitably transform the market and prove that a Walrasian equilibrium exists in the original market if and only if the core of the transformed market is nonempty. Since this condition takes the form of a set of linear inequalities, it provides another algorithmic existence test. Finally, Baldwin and Klemperer (2019) use methods from tropical geometry to prove their Intersection Count Theorem. They show that the existence of an equilibrium can be checked by only counting specific intersection points of the bidders' LIPs. Since the precise statement is very technical, we omit further details here.

2.5 Computation of Equilibria

A straightforward way to compute a Walrasian equilibrium, implied by the welfare theorems, is to solve the (RWDP) directly. However, due to a large number of variables, this is only feasible for relatively small instances. We will briefly overview more efficient methods and focus on algorithms for gross and strong substitutes preferences.

The original proof of equilibrium existence for gross substitutes bidders by Kelso and Crawford (1982) is constructive and thus provides an algorithm to compute a Walrasian equilibrium. This algorithm can be interpreted as an iterative auction. We present a simplified version of this auction, which has been introduced by Leme (2017). The auction starts with a price vector $\mathbf{p} = (0, \dots, 0)$ and the assignment \mathbf{S} assigning all items to bidder 1. Subsequently, as long as some buyer i receives a bundle $\mathbf{s}_i \notin D_i(\mathbf{p})$, a bundle $\mathbf{r}_i \in D_i(\mathbf{p})$ with $\mathbf{r}_i \geq \mathbf{s}_i$ is determined. By removing the respective items from the other bidders' assignments, bidder i gets assigned bundle \mathbf{r}_i , and the prices for all items in $\mathbf{r}_i - \mathbf{s}_i$ are raised by a small real number δ . Due to the properties of gross substitutes valuations, this process is well-defined and can be shown to terminate in a Walrasian equilibrium.

While this algorithm is quite intuitive, it is not computationally efficient. Leme (2017) provides the upper bound of nM/δ price increases in the algorithm, where $M = \max_{i \in \mathcal{I}} \max_{\mathbf{s} \in \{0,1\}^m} v_i(\mathbf{s})$. Thus, this algorithm is, in general, not polynomial.

More efficient algorithms are based on solving the (RWDP) and deploy methods from convex analysis to make the computation more efficient. By inspecting the dual of the (RWDP), one can show that the set of solutions of the convex minimization problem

$$\min_{\mathbf{p} \in \mathbb{R}_{\geq 0}^m} L(\mathbf{p}) = \sum_{i \in \mathcal{I}} u_i(\mathbf{p}) + \langle \mathbf{p}, \mathbf{t} \rangle$$

coincides with the set of Walrasian equilibrium prices, given that an equilibrium exists. For a formal derivation of this equivalence, we refer to Publication 1. The convex function $L(\mathbf{p})$ is often called the Lyapunov function in the literature (Ausubel 2006). Efficient algorithms for Walrasian equilibrium computation often minimize $L(\mathbf{p})$ instead of solving the primal (RWDP).

Equilibrium computation algorithms can be classified by their access to bidders' preferences. A standard assumption in the literature is that the algorithm has access to a demand oracle, meaning that the algorithm can request a bundle in $D_i(\mathbf{p})$ for any price vector $\mathbf{p} \in \mathbb{R}_{\geq 0}^m$. Another common assumption is that the algorithm has access to a value oracle, which returns $v_i(\mathbf{s})$ for any bidder i and bundle \mathbf{s} . Importantly, bundles that bidders demand at a certain price can be interpreted as subgradients of the indirect utility function: it holds that $\text{conv } D_i(\mathbf{p}) = -\partial u_i(\mathbf{p})$ (see Publication 1 for details). Thus, given bundles $\mathbf{s}_i \in D_i(\mathbf{p})$ for every bidder i , we have that $\mathbf{t} - \sum_{i \in \mathcal{I}} \mathbf{s}_i \in \partial L(\mathbf{p})$.

Based on this observation, Leme and Wong (2020) present the first polynomial-time algorithm for computing a Walrasian equilibrium that only uses an aggregate demand oracle, that is, an oracle returning an arbitrary element $\mathbf{s} \in \sum_{i \in \mathcal{I}} D_i(\mathbf{p})$. Their algorithm combines a cutting plane method (Lee et al. 2015), that approximately minimizes arbitrary convex functions in polynomial time, with structural results on the set of Walrasian equilibria to solve the (RWDP) exactly. Note that their algorithm works for general markets with indivisibilities and does not make further assumptions on the structure of valuations functions. In addition, the authors provide optimized variants of the algorithm for gross substitutes preferences.

For markets with strong substitutes bidders, specialized steepest descent algorithms exist. To simplify the exposition, we assume that the bidders' valuations are integer-valued, i.e., are of the form $v : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{Z}_{\geq 0} \cup \{-\infty\}$.

If all agents have integer-valued strong substitutes valuations, the function L is L^\natural -convex. This implies that the functions $\mathbf{q} \mapsto L(\mathbf{p} + \mathbf{q}) - L(\mathbf{p})$ and $\mathbf{q} \mapsto L(\mathbf{p} - \mathbf{q}) - L(\mathbf{p})$ are submodular for every fixed price vector $\mathbf{p} \in \mathbb{Z}_{\geq 0}^m$ (Murota 2003). Hence, a direction of steepest descent of L at \mathbf{p} can be found in polynomial time by a submodular minimization algorithm (see, for example, Chakrabarty et al. (2017), Iwata et al. (2001), and Schrijver (2000)). This leads to the following algorithm. First, initialize prices by $\mathbf{p} = (0, \dots, 0)$. In each iteration, find a minimizer $\mathbf{q}^* \in \{0, 1\}^m$ of $\mathbf{q} \mapsto L(\mathbf{p} + \mathbf{q}) - L(\mathbf{p})$, and update $\mathbf{p} \leftarrow \mathbf{p} + \mathbf{q}^*$. While each iteration of the algorithm can be executed in polynomial time, the algorithm requires $\mathcal{O}(\min \{\|\mathbf{p}^*\|_\infty : \mathbf{p}^* \in \text{argmin } L\})$ iterations in the worst case (Shioura 2017). This makes the runtime only pseudo-polynomial. Note that values of L are not directly accessible via demand or value oracles. However, as Shioura (2017) points out, it is still possible to minimize $L(\mathbf{p} + \mathbf{q}) - L(\mathbf{p})$ with the help of a demand oracle and a membership oracle. Given a bundle \mathbf{s} and a price \mathbf{p} , a membership oracle answers whether $\mathbf{s} \in D_i(\mathbf{p})$ in constant time. Under different

assumptions on the available oracles, it is possible to make the algorithm more efficient by using larger step sizes: Baldwin et al. (2020) provide a polynomial-time version of the steepest descent algorithm, given that bidders express their valuations via the SSPMA bid language (see Section 2.6 below). An advantage of the steepest descent algorithm is its natural interpretation as an ascending auction. This auction has been introduced by Ausubel (2006), who also provides a strategyproof version of it.

2.6 Bid Languages for Substitutes Valuations

In single-round auctions where many different items are sold simultaneously, efficient communication of the bidders' preferences becomes a significant challenge. If there are m types of items to be sold, the number of all bundles of items is 2^m even in the single-unit case. Thus, communicating preferences by reporting $v_i(\mathbf{s})$ for every bundle $\mathbf{s} \in \text{dom } v_i$ becomes practically impossible. A bid language is a formalized way for bidders to report their valuation functions to the auctioneer. In this formalized communication process, bidders submit one or multiple bids to the auctioneer, each containing a logical unit of information about the preferences. There is an apparent trade-off between the simplicity of a bid language and the complexity of the preferences it can express.

Example 4. Recall that unit-demand preferences are defined as $v(\mathbf{s}) = \max_{k: s(k) \geq 1} w(k)$ for all $\mathbf{s} \neq \mathbf{0}$. Thus, a simple bid language for unit-demand preferences would be to report the m numbers $w(1), \dots, w(m)$. The communication complexity is much lower than explicitly reporting the values $v(\mathbf{s})$ for each bundle $\mathbf{s} \in \{0, 1\}^m$. However, it enables buyers to express only a small set of valuation functions.

Due to their practical and theoretical importance in indivisible markets, various suggestions for bid languages that express gross or strong substitutes preferences have been made. In the following, we present Endowed Assignment Valuations for gross substitutes, as well as Assignment Messages and the Strong Substitutes Product-Mix Auction (SSPMA) bid language for strong substitutes valuations. While every preference expressible via those bid languages is always a gross or strong substitutes valuation, only the SSPMA can express *all* strong substitutes preferences.

Endowed Assignment Valuations, introduced by Hatfield and Milgrom (2005), define a class of single-unit valuation functions (i.e., with domain $\text{dom } v = \{0, 1\}^m$). They were initially motivated by a hospital having a set of jobs \mathcal{J} to fill and an initial endowment

of doctors \mathcal{T} . There is an additional set \mathcal{K} of doctors that the hospital can hire. Each doctor-job pair (d, j) has a weight $w(d, j)$, describing how well doctor d performs in job j . The value for a subset of doctors $S \subseteq \mathcal{K}$ is the maximum weight of a matching from the doctors in $\mathcal{T} \cup S$ to the jobs \mathcal{J} , minus the maximum weight of a matching from \mathcal{T} to \mathcal{J} (so that the value of the empty set $S = \emptyset$ is 0). Consequently, an Endowed Assignment valuation can be reported by providing the $(|\mathcal{T}| + |\mathcal{K}|) \times |\mathcal{J}|$ -matrix of weights $w(d, j)$. While the set of Endowed Assignment valuations has a very intuitive meaning and it is quite efficient to report to the auctioneer, Ostrovsky and Leme (2015) prove that it captures only a strict subset of gross substitutes preferences, i.e., there exist gross substitutes preferences that one cannot express in this way.

Assignment Messages, introduced by Milgrom (2009), form a class of strong substitutes valuations for multi-unit markets. A formal definition of Assignment Messages would go beyond the scope of the discussion here. On a high level, Assignment Messages define valuation functions by linear programs where the constraints have a tree-like structure. As we explain in Publication 2, an Assignment Message valuation can alternatively be interpreted as a min-cost flow problem. Like Endowed Assignment valuations, they are quite efficient to report to the auctioneer. Milgrom (2009) shows that all valuations expressible via Assignment Messages are strong substitutes valuations. However, as we prove in Publication 2, there are strong substitutes valuations not expressible this way.

The SSPMA bid language was originally introduced by Klemperer (2008) during the 2007 financial crisis for the bank of England for auctioning loans to private banks. In the original version of the Product-Mix Auction, a bid is a tuple $(v(1), \dots, v(m); q)$, where $v(1), \dots, v(m) \in \mathbb{R}_{\geq 0}$ and $q \in \mathbb{Z}_{\geq 0}$. Such a bid describes the willingness to buy at most q units of goods, where one unit of good k has the value $v(k)$. Each bidder can submit an arbitrary number of bids. To evaluate a bidder's valuation $v(\mathbf{s})$ for bundle \mathbf{s} , distribute the items in \mathbf{s} among the bids, such that the total value is maximized subject to the quantity constraints. More precisely, suppose a bidder submits the bids $(v^b(1), \dots, v^b(m); q^b)$ for $b \in \mathcal{B}$. Then their value for bundle \mathbf{s} is

$$\begin{aligned}
 v(\mathbf{s}) = & \max \sum_{b \in \mathcal{B}} \sum_{k \in \mathcal{K}} v^b(k) x^b(k) \\
 \text{s.t.} & \sum_{k \in \mathcal{K}} x^b(k) \leq q^b \quad \forall b \in \mathcal{B} \\
 & \sum_{b \in \mathcal{B}} x^b(k) \leq s(k) \quad \forall k \in \mathcal{K} \\
 & x^b(k) \geq 0 \quad \forall b \in \mathcal{B} \quad \forall k \in \mathcal{K}.
 \end{aligned}$$




Figure 2.2: Left: LIP of a positive bid at $(3, 2)$. Center: Positive bids at $(2, 0)$, $(0, 2)$ and $(4, 4)$. Right: Adding a negative bid at $(2, 2)$ removes the respective parts of the LIP generated by the positive bids. Positive bids are indicated with black dots, and the negative bid with a gray dot.

Unfortunately, expressing all strong substitutes preferences is impossible with these bids alone. For this reason, additional negative bids are introduced. These are bids $(v(1), \dots, v(m); q)$ where the quantity $q \in \mathbb{Z}_{<0}$ is negative. Bidders can use negative bids to cancel some part of the demand that positive bids have created. For more intuition and a formal treatment of negative bids, we refer to Publication 3. The distinguishing property of the SSPMA bid language is that the set of valuation functions expressible via SSPMA bids is exactly equal to the set of all strong substitutes valuation functions (Baldwin and Klemperer 2021). In the worst case, exponentially many bids are required to express a given strong substitutes valuation function. One can, however, argue that “typical” valuations only require a small number of bids (see Publication 3).

SSPMA bids have an intuitive geometric interpretation in terms of the Locus of Indifference Hyperplanes introduced in Section 2.4.1 (see also Baldwin and Klemperer (2019) and Publication 3). For simplicity, consider a market with 2 different types of goods, where each bid is of the form $(v(1), v(2); q)$. Let us consider the LIP generated by such bids (Figure 2.2a). Each bid, positive or negative, can be interpreted as a star-like shape in price space, with three rays emanating from the point $(v(1), v(2))$. Given a collection of positive bids, the LIP generated by those bids is just the union of these shapes (Figure 2.2b). A negative bid, on the other hand, removes those parts of the LIP that coincide with the shape of the bid (Figure 2.2c).

3 Publication 1: Walrasian Equilibria from an Optimization Perspective

Peer-Reviewed Journal Paper

Title: Walrasian equilibria from an optimization perspective: A guide to the literature

Authors: Martin Bichler, Maximilian Fichtl, Gregor Schwarz

In: Naval Research Logistics

Abstract: An ideal market mechanism allocates resources efficiently such that welfare is maximized and sets prices in a way so that the outcome is in a competitive equilibrium and no participant wants to deviate. An important part of the literature discusses Walrasian equilibria and conditions for their existence. We use duality theory to investigate existence of Walrasian equilibria and optimization algorithms to describe auction designs for different market environments in a consistent mathematical framework that allows us to classify the key contributions in the literature and open problems. We focus on auctions with indivisible goods and prove that the relaxed dual winner determination problem is equivalent to the minimization of the Lyapunov function. This allows us to describe central auction designs from the literature in the framework of primal-dual algorithms. We cover important properties for existence of Walrasian equilibria derived from discrete convex analysis, and provide open research questions.

Contribution of dissertation author: Methodology, formal analysis, visualization, joint paper management

Copyright Notice: © 2020 Wiley Periodicals LLC.

Reference: Bichler et al. (2021b)

Walrasian equilibria from an optimization perspective: A guide to the literature

Martin Bichler¹ | Maximilian Fichtl | Gregor Schwarz

Department of Informatics, Technical University of Munich, Munich, Germany

Correspondence

Martin Bichler, Department of Informatics, Technical University of Munich, Munich, Germany.
Email: bichler@in.tum.de

Funding information

Deutsche Forschungsgemeinschaft (DFG), Grant/Award Number: BI 1057/1-8.

History

Accepted by Saša Pekeč, algorithms, computation, and economics.

Abstract

An ideal market mechanism allocates resources efficiently such that welfare is maximized and sets prices in a way so that the outcome is in a competitive equilibrium and no participant wants to deviate. An important part of the literature discusses Walrasian equilibria and conditions for their existence. We use duality theory to investigate existence of Walrasian equilibria and optimization algorithms to describe auction designs for different market environments in a consistent mathematical framework that allows us to classify the key contributions in the literature and open problems. We focus on auctions with indivisible goods and prove that the relaxed dual winner determination problem is equivalent to the minimization of the Lyapunov function. This allows us to describe central auction designs from the literature in the framework of primal-dual algorithms. We cover important properties for existence of Walrasian equilibria derived from discrete convex analysis, and provide open research questions.

KEYWORDS

duality, primal–dual algorithms, Walrasian equilibrium

1 | INTRODUCTION

Many markets match supply and demand for multiple goods or services (which we also refer to as items) via optimization. Typically, the auctioneer computes an allocation and linear (i.e., item-level), anonymous prices. Linear and anonymous competitive equilibrium prices are often referred to as Walrasian prices in honor of Léon Walras, a French mathematical economist, who pioneered the development of general equilibrium theory. Prominent examples include financial markets (Klemperer, 2010), day-ahead electricity markets (Meeus et al., 2009; Triki et al., 2005), environmental markets (Bichler et al., 2019), logistics (Caplice & Sheffi, 2003; Bichler et al., 2006; Ağralı et al., 2008) or spectrum auctions (Bichler & Goeree, 2017). In some of these markets the auctioneer computes prices that are in a competitive equilibrium with

linear and anonymous prices (aka. a Walrasian equilibrium),¹ in others Walrasian prices even lead to efficiency losses (Özer & Özturan, 2009; Lessan & Karabatı, 2018; Bichler et al., 2018; Meeus et al., 2009; Madani & Van Vyve, 2015). This raises the question, which market characteristics admit Walrasian equilibria.

While this is an established and central question in the economic sciences, there have been a number of significant contributions in computer science, economics, and operations research in recent years. The literature on auction algorithms initiated by Bertsekas (1988) is one of the early examples of the fruitful interplay between optimization and equilibrium

¹There are also competitive equilibria with nonlinear prices (Bikhchandani & Ostroy, 2002). However, some authors only use competitive equilibrium to refer to one with linear and anonymous prices.

theory. In this paper, we survey the literature and describe established and more recent results. We primarily draw on convex analysis and linear programming duality, and provide a consistent mathematical optimization framework to position and explain the key results of this broad literature.

1.1 | Competitive equilibrium

Early in the study of markets, general equilibrium theory was used to understand how markets could be explained through the demand, supply, and prices of multiple commodities or objects. The Arrow–Debreu model shows that under convex preferences, perfect competition, and demand independence there must be a set of competitive equilibrium prices (Arrow & Debreu, 1954; McKenzie, 1959; Gale, 1963; Kaneko, 1976). Market participants are price-takers, and they sell or buy goods in order to maximize their value subject to their budget or initial wealth in this model. The results derived from the Arrow–Debreu model led to the well-known welfare theorems, important arguments for markets as efficient or welfare-maximizing ways to allocate resources. Stability in the form of competitive equilibria where each participant maximizes his utility at the prices is central to this theory. More specifically, the theory focuses on Walrasian equilibria where there is one equilibrium price per good (aka. linear prices) and the price is the same for all bidders (aka. anonymous prices). The first theorem states that any Walrasian equilibrium leads to a Pareto efficient allocation of resources. The second theorem states that any efficient allocation can be attained by a Walrasian equilibrium under the Arrow–Debreu model assumptions.

However, general equilibrium theory assumes divisible goods and convex preferences, and the results do not carry over to markets with indivisible goods and complex (nonconvex) preferences and constraints. Also, in general equilibrium models money does not have outside value and bidders maximize value subject to a budget constraint (Cole et al., 2016). More importantly, bidders are assumed to be nonstrategic price-takers. Based on the work by Vickrey (1961), attention in economics shifted to auction theory, which focuses on small and imperfectly competitive markets, where strategic players can influence prices. These bidders have a quasilinear utility function, that is, they aim to maximize payoff (i.e., value minus price) (Krishna, 2009). Bayesian Nash equilibria (rather than competitive equilibria) are the central equilibrium solution concept in the auction literature, a branch of noncooperative and incomplete information game theory which led to remarkable results. Most importantly, the Vickrey–Clarke–Groves (VCG) mechanism was shown to be incentive-compatible, and truthful bidding to be a dominant strategy for bidders (Vickrey, 1961).

Many markets that have been implemented for trading financial products, electricity, or environmental access rights as discussed earlier are large markets involving many items and many market participants. Participants want to maximize

payoff, but they might not be able to influence prices on such markets. As a consequence, much of the literature is based on a complete-information game-theoretical analysis where bidders are price-takers rather than an incomplete-information game (Baldwin & Klemperer, 2019). Competitive equilibria are the main design desideratum. Unfortunately, it is well known that in many of these markets linear (i.e., block-level) prices might not allow for a welfare-maximizing trade and that there might not be competitive equilibria (Meeus et al., 2009; Madani & Van Vyve, 2015b).

Such new markets have led to a renewed interest in the question of existence and computation of competitive equilibria (Kim, 1986; Bikhchandani & Mamer, 1997; Bikhchandani & Ostroy, 2002; Baldwin & Klemperer, 2019; Leme, 2017). The problem is fundamentally rooted in mathematical optimization, as we will show. In this survey, we will focus on central and recent results in competitive equilibrium theory and multiobject auction design and reformulate them in the language of optimization, specifically duality theory and primal-dual algorithms.

1.2 | Outline

There are various ways how surveys are written. Some articles collect and categorize a larger number of papers in a new and emerging field (Herroelen & Leus, 2005; Galindo & Batta, 2013; Olafsson et al., 2008), others provide a guide to a larger literature and introduce important concepts in a unified framework. Examples include a survey on bilevel programming by Colson et al. (2005) or a survey on the gross substitutes condition in economics by Leme (2017). We follow the latter path and discuss competitive equilibrium theory using duality theory and linear programming as a framework. While most of the literature on this subject is published in economics journals, key insights of this literature can be introduced conveniently using the mathematical framework of optimization. Fundamentally, auctions are algorithms for optimal resource allocation and there are plenty of questions where the OR community can contribute as we discuss in the last section.

The survey starts with markets for divisible goods and shows that the concave conjugate to the aggregate value function of all bidders yields prices, and that the minimizer of the Lyapunov function results in Walrasian prices if the aggregate value function is concave. A condition for concavity of the aggregate value function is concavity of the individual value functions, which is equivalent to diminishing marginal returns. The Lyapunov function is convex so that a simple subgradient algorithm finds the minimum efficiently. This algorithm has an interpretation as an auction.

We will next show that the same principles from duality theory carry over to markets with indivisible objects. For this, we describe the allocation problem as a binary program.

Whenever the linear programming relaxation of this binary program has integer solutions, then the dual variables of the capacity constraints have an interpretation as Walrasian prices for the respective resources. We prove that the dual of the linear programming relaxation of this binary program is equivalent to the Lyapunov function. Economic literature discusses conditions on individual value functions that allow for Walrasian equilibria. This is the case if the convolution of these individual functions results in a discrete concave aggregate value function.

As in the continuous case with divisible goods, we can use a steepest descent algorithm to find the minimizer of the Lyapunov function, which is equivalent to determining Walrasian prices for the market. This is exactly what the auction mechanism by Ausubel (2005) does, a central contribution to auction design. Primal-dual algorithms are well-known algorithms to solve linear programs, and they have a nice interpretation as a market with an auctioneer and the bidders optimizing alternatively. The steepest descent algorithm that minimizes the Lyapunov function is equivalent and we show the connections.

We contribute the equivalence of the Lyapunov function and the dual linear programming relaxation of the allocation problem in markets with indivisible goods, as well as the equivalence of primal-dual algorithms with central auction designs for selling multiple indivisible goods. These two results allow us to organize the material and use duality theory to discuss the literature on existence of Walrasian equilibria, and linear programming algorithms to discuss auction designs leading to Walrasian equilibria if it exists. The survey helps scholars with a background in mathematical optimization to understand central results in competitive equilibrium theory and draws important connections between competitive equilibrium theory, mathematical optimization, and discrete convexity.

In Section 2 we introduce the notation and standard assumptions in the economic literature for readers from operations research. Then we introduce important concepts for the understanding of Walrasian equilibria such as the Lyapunov function for markets with divisible goods in Section 3. The same concepts play a role for markets with indivisible goods and discrete value functions in Section 4. In Section 5 we use primal-dual algorithms and show that these are equivalent to important auction designs discussed in economics. Finally, we provide a research agenda and discuss open research problems for the operations research community.

2 | NOTATION AND ECONOMIC ENVIRONMENT

In the auction market, there are m types of items or goods, denoted by $k \in \mathcal{K} = \{1, \dots, m\}$, and n bidders $i \in \mathcal{I} = \{1, \dots, n\}$. In the multi-unit case, we have $s \in \mathbb{Z}_{\geq 0}^m$ units available, that is, $s(k)$ homogeneous units for each of

the heterogeneous m items $k \in \mathcal{K}$. A bundle for bidder i is described by a vector $x_i \in \mathbb{Z}_{\geq 0}^m$. In case of single-unit supply the vector is binary, that is, $x_i \in \{0, 1\}^m$. We will sometimes omit the subscript i for convenience. Each bidder i has a value function $v_i : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{Z}_{\geq 0}$ over bundles of items or objects x_i . We assume integer-valued functions v_i as it will be more convenient to analyze the optimality of auction algorithms. Moreover, integer-valued functions v_i allow to use integral prices in ascending auctions without losing efficiency.

Unless stated otherwise this paper we assume that bidders have preferences described via a valuation function with the following properties:

- Pure private values: Bidder i 's value $v_i(x_i)$ does not change when she learns other bidder's information.
- Quasilinearity: Bidder i 's (direct) utility from bundle x_i is given by $\pi_i(x_i, p) = v_i(x_i) - \langle p, x_i \rangle$, where $\langle \cdot, \cdot \rangle$ is the dot product.
- Monotonicity: The function $v_i : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{Z}_{\geq 0}$ is weakly increasing with $v_i(0) = 0$ and, if $x_i \geq x_i'$, then $v_i(x_i) \geq v_i(x_i')$.

An auctioneer wants to find an allocation of items to bidders. Such an allocation is *feasible* when the supply suffices to serve the aggregate demand of the bidders. Furthermore, the auctioneer aims for *allocative efficiency*. This means the auctioneer wants to maximize *social welfare* which is the sum of the utilities of all participants (the bidders and the auctioneer). Maximization of welfare is also referred to as a *utilitarian* welfare function. In case of quasilinear utility functions, prices cancel and the social welfare is defined as $\sum_{i \in \mathcal{I}} v_i(x_i)$.

For the remainder of this survey we assume that the auctioneer's valuation for all items is zero. As a consequence, the auctioneer would sell items to bidders for a price of zero. In some auction scenarios, however, the auctioneer may want to set *reserve prices* which are the minimum prices at which the auctioneer would be willing to sell the goods. Often these reserve prices can be implemented by introducing a dummy bidder who simply bids the reserve prices on behalf of the auctioneer in the auction. In case the dummy bidder wins any items in the auction, these items remain unsold.

The goal of the auctioneer is to find an efficient allocation that yields linear (i.e., item-level) and anonymous market clearing prices $p = \{p(k)\}_{k \in \mathcal{K}} \in \mathbb{R}^m$. The *linearity of prices* refers to the property that individual prices are set for each item $k \in \mathcal{K}$; the price for a bundle x is then simply the sum of the prices of its components, that is, it is given by the dot product $\langle p, x \rangle$. *Anonymity* means that the resulting prices p are the same for all bidders and there is no price differentiation. Furthermore, prices p are *market clearing* when the aggregate demand of all bidders at the given prices p meets the supply s .

With linear and anonymous prices $p = (p(1), \dots, p(k), \dots, p(m))$, the bidder's *indirect utility function* is defined as

$$u_i(p) = \max_{x \in \mathbb{Z}_{\geq 0}^m} \{v_i(x) - \langle p, x \rangle\}.$$

The indirect utility function is widely used in economics and returns the maximal utility that bidder i can obtain for any bundle at prices p . The *demand correspondence* $D_i(p)$ is the set of bundles that maximize the indirect utility function at prices p , that is,

$$D_i(p) = \arg \max_{x \in \mathbb{Z}_{\geq 0}^m} \{v_i(x) - \langle p, x \rangle\}.$$

If in an outcome (consisting of an allocation and prices) all bidders are allocated a bundle from their demand set, then the outcome is *envy-free*. No bidder would want to get another bundle, as a bidder cannot increase her utility at these prices. Envy-free prices always exist. For example, if prices were higher than the valuations, then every bidder would only want the empty set. If in addition to envy-freeness all items are allocated, $\sum_{i \in \mathcal{I}} x_i = s$, then the outcome is a competitive equilibrium.

Definition 1 (Competitive equilibrium, CE). A price vector p^* and a feasible allocation (x_1, \dots, x_n) form a *competitive equilibrium* if $\sum_{i \in \mathcal{I}} x_i = s$ and $x_i \in D_i(p^*)$ for every bidder $i \in \mathcal{I}$.

If there were unsold items, an auctioneer could always add unsold units to the allocation of a bidder without decreasing welfare as bidders are assumed to have monotone value functions v_i .

In our setting with linear and anonymous prices, a competitive equilibrium is also called a *Walrasian equilibrium*. If there exists a Walrasian price vector p^* such that $p^* \leq p'$ for any other Walrasian price vector p' , then p^* is called the *bidder-optimal* Walrasian price vector. For Walrasian equilibria the well-known welfare theorems hold:

Theorem 1 *First and second welfare theorem (following Blumrosen and Nisan (2007)) Let $x = (x_1, \dots, x_n)$ be an equilibrium allocation induced by a Walrasian equilibrium price vector p , then x yields the optimal social welfare. Conversely, if x is a Pareto efficient allocation, then it can be supported by a Walrasian price vector p so that the pair (p, x) forms a Walrasian equilibrium.*

3 | WALRASIAN EQUILIBRIA WITH DIVISIBLE GOODS AND CONJUGACY

In this article, we focus on markets with indivisible goods. However, for instructive purposes, we briefly consider the case of divisible goods to introduce relevant concepts. These can then be transferred to the indivisible case. Our aim is to give an intuitive graphical and analytical interpretation of how the aggregate valuation function is connected to the indirect utility function, the Lyapunov function and the market prices.

We consider a market with multiple bidders $i \in \mathcal{I}$ and multiple divisible goods $k \in \mathcal{K}$ with $|\mathcal{I}| = n$ and $|\mathcal{K}| = m$. The aggregate value function $v_{\mathcal{I}}$ is defined as the supremum convolution of concave functions $v_i : \mathbb{R}_{\geq 0}^m \rightarrow \mathbb{R}$ where v_i is the value function of the i th bidder.

$$v_{\mathcal{I}}(s) = \max_{\{x_i\}_{i \in \mathcal{I}}} \left\{ \sum_{i \in \mathcal{I}} v_i(x_i) \mid x_i \in \mathbb{R}_{\geq 0}^m \text{ and } \sum_{i \in \mathcal{I}} x_i = s \right\}.$$

By compactness and continuity, the maximum exists. Concavity implies that $v_i((1 - \alpha)x + \alpha y) \geq (1 - \alpha) v_i(x) + \alpha v_i(y)$ with $x, y \in \mathbb{R}_{\geq 0}^m$ and $\alpha \in (0, 1)$. The economic interpretation of a concave valuation function is that it exhibits decreasing marginal valuations. Since every function v_i is concave, also their convolution $v_{\mathcal{I}}$ is concave.

The aggregate indirect utility is defined as $u_{\mathcal{I}}(p) = \sum_i u_i(p)$ and the aggregate demand set is given by the Minkowski sum $D_{\mathcal{I}}(p) = \sum_i D_i(p)$.

For the sake of simplicity of the following graphical interpretation of indirect utility and the concept of conjugacy, we consider a market with multiple bidders but only a single divisible good $x \in \mathbb{R}_{\geq 0}$. However, our explanations carry over directly to markets with multiple goods. It is also worth mentioning that in the presence of only a single bidder i the aggregate valuation function $v_{\mathcal{I}}$ becomes the individual valuation function v_i of the single bidder. Thus, even though the following example illustrates the aggregate valuation and indirect utility function of multiple bidders, it similarly applies to the valuation and indirect utility function of an individual bidder.

In our example, we assume $v_{\mathcal{I}}(x) = \ln(x + 1)$. It is well known that for concave functions $v_{\mathcal{I}}$ local optimality implies global optimality and this yields efficient optimization algorithms.

At a given price, every rational bidder $i \in \mathcal{I}$ only demands a quantity of good x which maximizes her utility at this price. The utility of such a quantity is described by the indirect utility function $u_i(p) = \max_x \{v(x) - \langle p, x \rangle\}$, which is convex as it is the maximum of affine linear functions. As the aggregate indirect utility function $u_{\mathcal{I}}(p)$ is a sum of convex functions, it must also be convex.

A quantity x^* is demanded at prices p if and only if $v_{\mathcal{I}}(x^*) - \langle p, x^* \rangle \geq v_{\mathcal{I}}(x) - \langle p, x \rangle$ for all $x \in \mathbb{R}$. When rearranging terms to $v_{\mathcal{I}}(x^*) + \langle p, x - x^* \rangle \geq v_{\mathcal{I}}(x)$, it becomes clear that the left-hand side of the inequality describes the tangent at $v_{\mathcal{I}}(x^*)$ (see Figure 1). In other words, a quantity x^* is demanded at prices p whenever the slope of the tangent at $v_{\mathcal{I}}(x^*)$ equals the price p . The aggregate utility of quantity x^* is given by $\pi_{\mathcal{I}}(x^*, p) = v_{\mathcal{I}}(x^*) - \langle p, x^* \rangle$. As $x^* \in D_{\mathcal{I}}(p)$, the aggregate utility $\pi_{\mathcal{I}}(x^*, p)$ equals the aggregate indirect utility $u_{\mathcal{I}}(p)$. The graphical interpretation of the aggregate indirect utility function $u_{\mathcal{I}}(p)$ is the intercept of the tangent at $v_{\mathcal{I}}(x^*)$ with the ordinate.

We can now compute the quantity of good x that generates maximum utility at prices p . In our illustrative example with $v_{\mathcal{I}}(x) = \ln(x + 1)$, the aggregate utility $\pi_{\mathcal{I}}(x, p) = \ln(x + 1) - \langle p, x \rangle$ at given prices p is maximized when




FIGURE 1 Graphical representation of $v_I(x) = \ln(x+1)$ with tangent at $v_I(x^*)$

$\partial\pi_I/\partial x = 1/(x+1) - p = 0$. This means, at a price of $p = 1/3$ for example, the total utility π_I is maximized for a demand of $x^* = 2$. Thus, the aggregate indirect utility function at prices $p = 1/3$ equals $u_I(1/3) = \pi_I(2, 1/3) = \ln(3) - 2/3$. The concave conjugate (or Legendre transformation) of v_I is defined as $v_I^*(p) = \min_x \{ \langle p, x \rangle - v_I(x) \}$, which is the aggregate indirect utility function multiplied by -1 . We also note that convex and concave conjugates are connected via $v_I^*(p) = -(-v_I)^*(-p)$, so $u_I(p) = (-v_I)^*(-p)$. From these results, we can make the following connection: In order to construct the concave conjugate $v_I^*(p)$ of $v_I(x) = \ln(x+1)$ for a fixed p , we must calculate the minimum of $\langle p, x \rangle - \ln(x+1)$. Taking the derivative, we see that a minimizing x must solve $x = 1/p - 1$, so we get $v_I^*(p) = 1 - p + \ln(p)$ and consequently $u_I(p) = -v_I^*(p) = p - \ln(p) - 1$. For a given price of $p = 1/3$ the reader may verify that the bidders' aggregate indirect utility equals $u_I(1/3) = 1/3 - \ln(1/3) - 1 = \ln(3) - 2/3$, which is in line with our calculations above.

Unlike in this single-item example, the price p is not known in an auction setting. Instead, the auctioneer tries to find a price vector p^* for which the supply s is a maximizer of the aggregate utility function $\pi_I(x, p^*)$. Note that such a p^* is a Walrasian equilibrium price vector, because s maximizes $\pi_I(s, p^*) = v_I(s) - \langle p^*, s \rangle$ and the aggregate demand of the bidders equals the supply s .

We will now return to a market with multiple divisible goods $k \in \mathcal{K}$. First, we introduce important notions from convex analysis.

Definition 2 Let $f: \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ be a convex function. The *subdifferential* of f at x is the set of all tangents of f at x :

$$\partial f(x) = \{y \in \mathbb{R}^d \mid f(x') \geq f(x) + \langle y, x' - x \rangle \forall x' \in \mathbb{R}^d\}.$$

Any element of $\partial f(x)$ is called a *subgradient*. The *convex conjugate* or *Legendre transform* of f is the convex function

$$f^*(y) = \sup_{x \in \mathbb{R}^d} \langle y, x \rangle - f(x).$$

Under additional mild assumptions on the convex function f , the conjugate of the conjugate is again f , $f^{**} = f$, and subdifferentials of f and f^* are connected in the following way: $y \in \partial f(x) \Leftrightarrow x \in \partial f^*(y)$. For more details, we refer to Rockafellar (2015). The concave conjugate defined above and the convex conjugate are related as follows: If g is concave, then $g^*(y) = -(-g)^*(-y)$. In particular, we have for the indirect utility function $u_I(p) = (-v_I)^*(-p)$. We make the following important observation: The bundle x is in the demand set $D_I(p)$, if and only if $v_I(x) - \langle p, x \rangle \geq v_I(x') - \langle p, x' \rangle$ for all $x' \in \mathbb{R}^{|\mathcal{K}|}$. By rearranging terms we see that this is equivalent to $-v(x') \geq -v(x) + \langle -p, x' - x \rangle$ and thus to $-p \in \partial(-v_I)(x)$. Convex analysis tells us that this is equivalent to $x \in \partial(-v_I)^*(-p) = -\partial u_I(p)$. Consequently, demand sets are equal to subdifferentials of the indirect utility function—a fact that allows us to interpret auctions as descent algorithms.

The Lyapunov function was a central concept already in the early literature on general equilibrium theory (Arrow & Hahn, 1971). The same function plays a central role in more recent auction designs for markets with indivisible goods (Ausubel, 2006). Since this function plays such a central role, we introduce it in detail for the continuous case.

Definition 3 (Lyapunov function). The *Lyapunov function* is defined as $L(p) = \sum_{i \in \mathcal{I}} u_i(p) + \langle p, s \rangle$, where s is the supply and $u_i(p)$ is the indirect utility function of bidder $i \in \mathcal{I}$ at prices p .

The Lyapunov function has its roots in the dynamical systems literature (La Salle & Lefschetz, 2012). Since the indirect utility $u_i(p)$ is convex in p , also the Lyapunov function is convex, because it is the sum of convex functions. For convex functions such as $L(p)$ the vector p^* minimizes L iff 0 is a subgradient at p^* . The first-order condition for $L(p)$ yields $-\sum_{i \in \mathcal{I}} x_i + s = 0$, where $x_i \in D_i(p)$.

$\forall i \in \mathcal{I}$. In words, the prices are minimized when supply equals demand:

Proposition 1 A vector $p^* \in \mathbb{R}^m$ is a Walrasian equilibrium price vector for supply s if and only if it is a minimizer of the Lyapunov function $L(p) = u_I(p) + \langle p, s \rangle$.

Proof If there is a Walrasian equilibrium, then $\sum_{i \in \mathcal{I}} x_i = s$ and $x_i \in D_i(p^*)$ need to hold. The minimizer p^* of $L(p)$ requires that $\partial L(p) = s - \sum_{i \in \mathcal{I}} x_i = 0$, which is equivalent to the first condition of a Walrasian equilibrium. Also, when $L(p) = \sum_{i \in \mathcal{I}} \max_{x_i} \{v_i(x_i) - \langle p, x_i \rangle\} + \langle p, s \rangle$ attains the minimum, then each bidder is assigned a bundle x_i that maximizes her utility $v_i(x_i) - \langle p, x_i \rangle$. This implies $x_i \in D_i(p^*)$ for all i , so that the second condition of a Walrasian equilibrium is fulfilled. Thus, if $L(p)$ is minimized then both conditions of a Walrasian

equilibrium are satisfied. By reversing the argument it becomes evident that any price vector p^* supporting s in a Walrasian equilibrium is also a minimizer for $L(p)$. ■

Similar results can be found in Ausubel and Milgrom (2006) or later in Murota (2016). One way to find Walrasian equilibria is now to minimize the Lyapunov function. Since we can interpret the subdifferential of u_i at price p as the demand set at this price—for an auction setting it is natural to utilize standard subgradient methods for (approximately) minimizing $L(p)$ —computing subgradients is then equivalent to asking bidders for their demand sets at a given price. Note that it is in general not possible to compute *exact* minimizers to general convex functions—algorithms for minimizing a convex function f can in general only provide complexity bounds for finding an ε -approximate solution x' , in the sense that

$$f(x') \leq \varepsilon + \min_x f(x).$$

Note that in general x' does not even have to be close to the true minimizer x without additional assumption on f . Since the aim of our treatment of divisible economies is mainly to motivate the ideas in the indivisible case, we will not go into more detail here. If no additional regularity assumptions on L are imposed, it can be shown that finding ε -approximate solutions has a worst-case running time of $\Theta(1/\varepsilon^2)$ (Nesterov, 2018). Interestingly, for markets with indivisible goods where Walrasian equilibria exist, we will show that the Lyapunov function equals the dual of the allocation problem.

Central results of convex economic theory with divisible goods are reasonable approximations to large economies where nonconvexities vanish in the aggregate (Starr, 1969). However, most markets are such that indivisibilities and nonconvexities matter. As one would assume, the analysis of markets with indivisible items has proven much harder.

4 | EXISTENCE OF WALRASIAN EQUILIBRIA WITH INDIVISIBLE GOODS

In this section, we discuss sufficient and necessary conditions for the individual value functions of bidders such that Walrasian equilibria exist in markets with indivisible goods.

4.1 | Conditions on aggregate value functions

A simple multi-item market with remarkable properties is the assignment market by Shapley and Shubik (1971). In assignment markets each bidder can bid on multiple items but wants to win at most one (aka. *unit-demand*). As a consequence, the allocation problem reduces to an assignment problem, that is, the problem of finding a maximum weight matching in a weighted bipartite graph. On an aggregate level, the LP relaxation of the assignment problem is always integral. This is a consequence of the unit demand on an individual

level and the resulting total unimodularity of the constraint matrix, and this is a sufficient condition for the existence of Walrasian prices. The environment of assignment markets allows for incentive-compatible auctions. Besides, simple ascending clock auctions yield bidder-optimal Walrasian prices (Demange et al., 1986).

4.1.1 | The allocation problem

Let us first extend the assignment market to a more general multi-item, multi-unit market which allows for package bids. Let $\mathcal{X}_i \subseteq \mathbb{Z}_{\geq 0}^m$ denote all bundles for which bidder i submitted a bid. For simplicity, we make the natural assumption that every bidder submits a bid with value 0 for the empty bundle. Let $z_i(x) \in \{0, 1\}$ be a binary decision variable denoting whether bidder i wins bundle $x \in \mathcal{X}_i$. The allocation or winner determination problem WDP can then be written as an integer program as follows:

$$\begin{aligned} \max \quad & \sum_{i \in \mathcal{I}} \sum_{x \in \mathcal{X}_i} v_i(x) z_i(x) && \text{(WDP)} \\ \text{s.t.} \quad & \sum_{x \in \mathcal{X}_i} z_i(x) \leq 1 && \forall i \in \mathcal{I} \quad (\pi_i) \\ & \sum_{i \in \mathcal{I}} \sum_{x \in \mathcal{X}_i} x(k) z_i(x) \leq s(k) && \forall k \in \mathcal{K} \quad (p(k)) \\ & z_i(x) \in \{0, 1\} && \forall i \in \mathcal{I}, \forall x \in \mathcal{X}_i \end{aligned}$$

For a given supply s the WDP determines an allocation of bundles to bidders maximizing social welfare. The LP relaxation RWDP in standard form replaces $z_i(x) \in \{0, 1\}$ by $z_i(x) \geq 0$ and introduces additional slack variables. We use the standard form with slack variables (a_i, b_k) because it will be helpful in our algorithmic treatment of the subject.

$$\begin{aligned} \max \quad & \sum_{i \in \mathcal{I}} \sum_{x \in \mathcal{X}_i} v_i(x) z_i(x) && \text{(RWDP)} \\ \text{s.t.} \quad & \sum_{x \in \mathcal{X}_i} z_i(x) + a_i = 1 && \forall i \in \mathcal{I} \quad (\pi_i) \\ & \sum_{i \in \mathcal{I}} \sum_{x \in \mathcal{X}_i} x(k) z_i(x) + b_k = s(k) && \forall k \in \mathcal{K} \quad (p(k)) \\ & z_i(x), a_i, b_k \geq 0 && \forall i \in \mathcal{I}, \forall x \in \mathcal{X}_i, \forall k \in \mathcal{K} \end{aligned}$$

In contrast to the assignment problem where bidders have unit demand, the RWDP does not yield integer solutions in general.

Example 1 Consider a market with three items $\mathcal{K} = \{A, B, C\}$ and two bidders with valuations v_1 and v_2

	x_\emptyset	x_A	x_B	x_C	x_{AB}	x_{AC}	x_{BC}	x_{ABC}
x	(0, 0, 0)	(1, 0, 0)	(0, 1, 0)	(0, 0, 1)	(1, 1, 0)	(1, 0, 1)	(0, 1, 1)	(1, 1, 1)
$v_1(x)$	0	1	2	1	2	2	2	2
$v_2(x)$	0	1	2	2	3	2	3	3

The optimal solution of the RWDP given these valuations is fractional:

$z_1(x_B) = z_1(x_{AC}) = z_2(x_C) = z_2(x_{AB}) = 0.5$ with all other decision variables set to 0. The optimal value of the RWDP with respect to this fractional solution is 4.5. An optimal integral solution (e.g., assigning bundle x_{AC} to the first and x_B to the second bidder) only leads to a social welfare of 4.

Let us also introduce the dual DRWDP of the RWDP.

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{I}} \pi_i + \sum_{k \in \mathcal{K}} s(k)p(k) && \text{(DRWDP)} \\ \text{s.t.} \quad & \pi_i + \sum_{k \in \mathcal{K}} x(k)p(k) \geq v_i(x) && \forall i \in \mathcal{I}, \forall x \in \mathcal{X}_i \quad (z_i(x)) \\ & \pi_i \geq 0 && \forall i \in \mathcal{I} \quad (a_i) \\ & p(k) \geq 0 && \forall k \in \mathcal{K} \quad (b_k) \end{aligned}$$

We will draw on these models in the subsequent sections.

4.1.2 | Integrality of the linear program

Bikhchandani and Mamer (1997) describe a multi-item, single-unit market. Their central theorem shows that there exist clearing prices for the indivisible single-unit problem if and only if the RWDP has an integer solution. In this case, the set of equilibrium prices is the set of solutions to the dual LP projected to the price coordinates. The result can be proven via the strong duality theorem in linear programming (Blumrosen & Nisan, 2007). As was already noted by Bikhchandani and Mamer (1997), the result for multi-item, *multi*-unit markets also directly follows from their result, by considering each of the multiple units as separate items. As the proof is a particularly nice application of duality theory, we provide a direct proof in the Appendix. Note that this theorem proves the welfare theorems from general equilibrium theory (see Theorem 1).

Theorem 2 *Walrasian prices exist for the supply s if and only if the RWDP has an optimal integral solution.*

The proof can be found in Appendix A.

As indicated, the RWDP typically does not yield an integral solution, and there can be a significant integrality gap between the objective function value of the RWDP and that of the optimal integer program WDP. In the next sections, we will discuss conditions on the individual value functions, which yield integral solutions of the RWDP and Walrasian prices.

Before we do this, let us return to the Lyapunov function that has proven so helpful in our analysis of markets with divisible goods. A minimizer to this function yielded the Walrasian prices in Section 3, where we analyzed markets with divisible goods. It turns out that the Lyapunov function is actually equivalent to the DRWDP, as we show in the following proposition.

Proposition 2 *A vector $p^* \in \mathbb{R}^m$ minimizes the DRWDP if and only if it is a minimizer of the Lyapunov function $L(p) = u_I(p) + \langle p, s \rangle$.*

Proof We can substitute the utilities π_i in the dual objective function $\min \sum_{i \in \mathcal{I}} \pi_i + \sum_{k \in \mathcal{K}} s(k)p(k)$ by the tight dual constraints $\pi_i = v_i(x) - \sum_{k \in \mathcal{K}} x(k)p(k)$ of the optimal DRWDP and get the following convex function:

$$\min_p \sum_{i \in \mathcal{I}} \max_{x \in \mathbb{Z}_{\geq 0}^m} \left[v_i(x) - \sum_{k \in \mathcal{K}} x(k)p(k) \right] + \sum_{k \in \mathcal{K}} s(k)p(k). \quad (4.1)$$

Note that this is equivalent to minimizing the Lyapunov function $L(p) = \sum_{i \in \mathcal{I}} u_i(p) + \langle p, s \rangle$. Obviously, $\langle p, s \rangle$ in $L(p)$ is equal to $\sum_{k \in \mathcal{K}} s(k)p(k)$, and $u_i(p)$ equals $\max_{x \in \mathbb{Z}_{\geq 0}^m} [v_i(x) - \sum_{k \in \mathcal{K}} x(k)p(k)]$ for every bidder i . Since the equivalence of the Lyapunov function and the DRWDP holds for any price vector p , minimizing prices of the Lyapunov function also constitutes a minimal solution to the DRWDP and vice versa. ■

In summary, both the Lyapunov function and the LP approach yield equilibrium prices, and such prices are minimizers of both problems. We will leverage this insight, when we analyze auction algorithms to solve the RWDP in Section 5.

4.2 | Conditions for individual value functions

In practical applications a market designer often wants to understand which assumptions on the individual value functions v_i allow for integer solutions of the LP relaxation and Walrasian prices. Discrete convex analysis identifies classes of convex functions defined on a subset of the discrete lattice \mathbb{Z}^m , which allow for integrality and efficient optimization algorithms.

First, we discuss single-unit, multi-item auctions. There are several classes of integrally convex functions such as separable-convex functions on \mathbb{Z}^m or gross substitutes set functions on $\{0, 1\}^m$, which yield a discrete concave aggregate value function v_I and integral solutions of the RWDP, such that Walrasian equilibria exist.

4.2.1 | Single-unit multi-item auctions

Let us first define monotonicity and submodularity, two well-known properties of set functions that allow for efficient function minimization.

Definition 4 For a finite set \mathcal{K} of items, the set function $v : 2^{\mathcal{K}} \rightarrow \mathbb{R}$ is

- *monotone* if $v(S) \leq v(T)$ for all $S, T \subseteq \mathcal{K}$ with $S \subseteq T$,

- *submodular* if $v(S \cup \{k\}) - v(S) \geq v(T \cup \{k\}) - v(T)$ for all $S, T \subseteq \mathcal{K}$ with $S \subseteq T$ and for all $k \notin T$.

In the above definition, submodularity can be understood as diminishing marginal values. Alternatively, submodularity can be defined as $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$ for all S, T . The vector notation $v: \{0, 1\}^m \rightarrow \mathbb{R}$ in the single-unit case maps a set S to a vector $x \in \{0, 1\}^m$ by setting $x(k) = 1$ whenever $k \in S$ and $x(k) = 0$ otherwise.

It is well-known that the minimization of unconstrained submodular functions can be done in polynomial time, for example via the ellipsoid method (Grötschel et al., 1981). The ellipsoid method is notoriously slow in practice. However, there are also more effective algorithms such as the Fujishige-Wolfe algorithm (Chakrabarty et al., 2014) and specialized subgradient methods (Chakrabarty et al., 2017). Unfortunately, even when submodularity and monotonicity are satisfied, this does not guarantee the integrality of a welfare maximization problem such as the RWDP.

Example 2 The reader may verify that the valuation functions of both bidders in example 1 satisfy monotonicity and submodularity. However, the optimal solution of the RWDP is not integral.

The subset of submodular valuations called gross substitutes valuations, however, has this desirable property. Gross substitutes roughly means that a bidder regards the items as substitute goods or independent goods but not complementary goods.

Definition 5 (Gross substitutes, GS). Let p denote the prices on all items, with item k demanded by bidder i if there is some bundle S , with $k \in S$, for which S maximizes the utility $v_i(S') - \sum_{j \in S'} p(j)$ across all bundles $S' \subseteq \mathcal{K}$. The gross substitutes condition requires that, for any prices $p' \geq p$ with $p'(k) = p(k)$, if item $k \in \mathcal{K}$ is demanded at the prices p then it is still demanded at p' .

The definition includes both substitute goods and independent goods, but rules out complementary goods.²

Example 3 Consider a market with three items $\mathcal{K} = \{A, B, C\}$ and a single bidder with a valuation function v fulfilling the gross substitutes condition

x_\emptyset	x_A	x_B	x_C	x_{AB}	x_{AC}	x_{BC}	x_{ABC}	
x	(0, 0, 0)	(1, 0, 0)	(0, 1, 0)	(0, 0, 1)	(1, 1, 0)	(1, 0, 1)	(0, 1, 1)	(1, 1, 1)
$v(x)$	0	1	2	3	3	3	5	5

²Sometimes the word “gross” used by Kelso and Crawford (1982) is omitted, but it is useful to distinguish the single-unit case from substitutes valuations in other environments, such as the strong substitutes definition that we will introduce later.

At prices $p = (0, 1, 2)$ the bidder’s indirect utility is $u(p) = 2$ and the bidder’s demand set is given by $D(p) = \{x_{AB}, x_{BC}, x_{ABC}\}$, that is, items A, B , and C are demanded as for each item there exists at least one bundle in the demand set containing the item. If the price for item A is raised to 1 but stays constant for items B and C , then the gross substitutes condition implies that items B and C must still be demanded at the new prices $p' = (1, 1, 2)$. This is obviously true as the demand set at the new prices p' is given by $D(p') = \{x_{BC}\}$. Note that price vectors p and p' were only chosen for illustrative purposes. In fact, valuation function v satisfies the gross substitutes condition for any price vectors $p, p' \in \mathbb{R}_{\geq 0}^3$ with $p' \geq p$.

Kelso and Crawford (1982) show that if all agents have GS valuations, then a Walrasian equilibrium always exists, which implies that the RWDP has an optimal integral solution. Ausubel and Milgrom (2002) prove that a bidder has GS valuations if and only if the indirect utility function u is submodular. Gross substitutes appear to be a rather restricted type of valuations, but it contains important subclasses such as unit-demand valuations (Shapley & Shubik, 1971) and additive valuations. Gul and Stacchetti (1999) show that GS excludes complementarity between goods and show equivalence with the so called single improvement property. The latter property states that whenever a bundle is not optimal at the given prices, then a better bundle can be found which is derived from the original one by performing any of the following operations: removing an item, adding an item, or doing both. Leme (2017) provides a survey of the extensive literature on the gross substitutes condition and its alternative definitions for multi-item, single-unit markets, and show that additive valuations \subset GS \subset submodular valuations \subset subadditive valuations. We also refer to Shioura and Tamura (2015) for an extensive survey of GS.

Sun and Yang (2006) identify the gross substitutes and complements (GSC) condition, which also guarantees for Walrasian equilibria in single-unit, multi-item markets. It describes an exchange economy with two classes of goods, where each class only contains substitutes, but there are complements across these classes of goods. Teytelboym (2014) generalizes the GSC condition in the sense that goods are partitioned into more than two classes. His generalized version of the GSC condition is satisfied if it is possible to partition goods into several classes so that whenever considering the bidders’ valuations for items contained in only two of these classes in isolation, there exist some bidders for which these valuations satisfy the GSC condition.

4.2.2 | Multi-unit multi-item auctions

Let us now concentrate on more general conditions for $x \in \mathbb{Z}_{\geq 0}^m$ rather than $x \in \{0, 1\}^m$. $A \subset \mathbb{Z}^m$ is *integrally convex* if

$A = (\text{conv } A) \cap \mathbb{Z}^m$. First, we define the *convex closure* \bar{f} of f as

$$\bar{f}(x) = \sup_{p \in \mathbb{R}^m, \alpha \in \mathbb{R}} \{ \langle p, x \rangle + \alpha \mid \langle p, y \rangle + \alpha \leq f(y) \ \forall y \in \mathbb{Z}^m \}.$$

Geometrically, the epigraph of \bar{f} is the convex hull of the epigraph of f . If the convex closure coincides with f on the set of integer vectors, that is, if $f(x) = \bar{f}(x)$ for all $x \in \mathbb{Z}^m$, f is called *convex-extensible*. In the same way, we can define the concave closure of f by $\overline{-f}$. The definition can be restricted to the integral neighborhood of a bundle $x \in \mathbb{R}^m$ and is then referred to as a *local convex extension* \tilde{f} of f (Murota, 2003, Chap. 3). Formally, set $N(x) = \{y \in \mathbb{Z}^m \mid \lfloor x(k) \rfloor \leq y(k) \leq \lceil x(k) \rceil \ \forall k = 1, \dots, m\}$. Then the local convex extension is given by

$$\tilde{f}(x) = \sup_{p \in \mathbb{R}^m, \alpha \in \mathbb{R}} \{ \langle p, x \rangle + \alpha \mid \langle p, y \rangle + \alpha \leq f(y) \ \forall y \in N(x) \}.$$

Definition 6 A function $f: \mathbb{Z}^m \rightarrow \mathbb{R}$ is called *integrally convex* if the local convex extension of f is convex, or *integrally concave* if the function $-f$ is integrally convex.

Integrally convex functions share with convex functions the property that local minima are also global minima (Murota, 2016). We have already seen in the divisible case that concavity of the valuation functions is necessary for equilibrium prices to exist. We also want to make this connection here in the indivisible case, by explaining how convexity is related to integrality of the WDP—which is necessary and sufficient for the existence of equilibrium prices. To start with, consider the aggregate valuation function $v_I(s)$, given by the value of the WDP for the supply s , and the “relaxed” aggregate valuation function $\tilde{v}_I(s)$, given by the value of the RWDP at s . Note \tilde{v}_I is well-defined for all *real* supply vectors $s \geq 0$ and attains finite values at each such s . A central observation is the following: \tilde{v}_I is the concave extension of v_I . This shows that v_I is concave-extensible, and thus $v_I = \tilde{v}_I$ if and only if for every integral supply vector s , the RWDP has an integral solution, which—as we have seen—is equivalent to the existence of equilibrium prices. While the stronger assumption of integral concavity is not necessary for the existence of equilibrium prices, it is not hard to imagine, that this property is of importance for the algorithmic problem of computing equilibrium prices. Loosely speaking, since the value of the concave extension can then be evaluated at any point s by considering an easy to characterize neighborhood of s , the computation of subgradients of v_I gets much simpler. Unfortunately, concave extensibility, and even integral concavity of the individual valuation functions does not suffice to guarantee concave extensibility of the aggregate valuation function, or equivalently, existence of equilibrium prices. It is thus of central importance to identify conditions on the individual valuations that imply concave extensibility of the aggregate valuation, or equivalently integrality of the RWDP.

Definition 7 A function $f: \mathbb{Z}^m \rightarrow \mathbb{R} \cup \{\infty\}$ is said to be M^{\natural} -convex if for $x, y \in \text{dom} f$ and $j \in \text{supp}^+(x - y)$

- (i) $f(x) + f(y) \geq f(x - \mathbb{1}_j) + f(y + \mathbb{1}_j)$ or
- (ii) $f(x) + f(y) \geq f(x - \mathbb{1}_j + \mathbb{1}_k) + f(y + \mathbb{1}_j - \mathbb{1}_k)$ for some $k \in \text{supp}^-(x - y)$.

A function f is M^{\natural} -concave if the function $-f$ is M^{\natural} -convex. A set $X \subseteq \mathbb{Z}^m$ is an M^{\natural} -convex set if its indicator function δ_X is M^{\natural} -convex.

Here $\mathbb{1}_j$ denotes the j th unit vector, whereas the positive and negative support are defined as $\text{supp}^+(x) = \{k \in \mathcal{K} \mid x(k) > 0\}$ and $\text{supp}^-(x) = \{k \in \mathcal{K} \mid x(k) < 0\}$, respectively. The effective domain is $\text{dom} f = \{z \in \mathbb{Z}^m \mid f(z) \neq \infty\}$. An M^{\natural} -convex function is integrally convex, and thus convex-extensible (Murota, 2003, Theorem 6.42). Since the exchange property (ii) is closely related to the exchange axiom of a matroid, the M stands for “matroid”. It means that if we add the j th unit-vector to one point x and exchange it with the i th unit vector of another point y , then the function value decreases or stays the same. Fujishige & Yang, 2003 showed that for the single-unit case the GS condition is equivalent to M^{\natural} -concavity.

Theorem 3 (Fujishige and Yang (2003)). A value function $v: \{0, 1\}^m \rightarrow \mathbb{R}$ satisfies the GS condition if and only if it is an M^{\natural} -concave function.

This equivalence extends to multi-unit extensions of the gross substitutes property. Milgrom and Strulovici (2009) distinguish between *weak* and *strong* substitutes. The weak substitutes condition can be seen as the natural extension of the original gross substitutes property to the multi-unit setting by simply quantifying the demand for every item. Note however, that weak substitutes do not correspond to M^{\natural} functions anymore (Shioura & Tamura, 2015). The strong substitutes condition, on the other hand, transforms a multi-unit to a single-unit valuation function by treating each copy of a good as an individual item. Whenever the corresponding single-unit valuation function fulfills the original gross substitutes property (as defined by Kelso and Crawford (1982)), the multi-unit valuation function satisfies the strong substitutes condition.

Definition 8 (Strong substitutes, SS). Let $\mathcal{K} = \{k_1, k_2, \dots, k_m\}$ be the set of items with $d_i \in \mathbb{N}$ denoting the number of units available of item k_i . Treating each copy of a good as an individual item leads to the definition of a set $\mathcal{K}_s = \{(k_i, z) \mid k_i \in \mathcal{K}, 1 \leq z \leq d_i\}$. A multi-unit valuation function $v: \mathbb{N}_0^m \rightarrow \mathbb{R}$ can then be transformed to a single-unit valuation function $v_s: \{0, 1\}^{\mathcal{K}_s} \rightarrow \mathbb{R}$ by setting $v_s(x_s) = v(x)$

for $x_s \in \{0, 1\}^{\mathcal{K}_s}$ where $x(i) = \sum_{z=1}^{d_i} x_s(k_i, z)$. The valuation v fulfills the strong substitutes condition if v_s is a gross substitutes valuation function.

There exist many equivalent definitions of the strong substitutes condition, among them the binary single-improvement property as shown by Milgrom and Strulovici (2009).

Danilov et al. (2001) and Milgrom and Strulovici (2009) show that a Walrasian equilibrium exists for every finite set of strong substitutes valuations. Ausubel (2006) shows that in case of strong substitutes valuations the Lyapunov function is submodular which ensures the existence of a bidder-optimal Walrasian price vector. While the strong substitutes property is a sufficient condition for the existence of Walrasian equilibria, it is not a necessary one and alternatives exist.

Shioura and Yang (2015) extend the gross substitutes and complements (GSC) condition to a multi-unit and multi-item economy with two classes of items, where units of the same type are substitutable, whereas goods across two classes are complementary. When there is only one class of indivisible goods, their generalized gross substitutes and complements (GGSC) condition becomes identical to the strong-substitute valuation of Milgrom and Strulovici (2009). Further, if each type of good has only one unit, it becomes the gross substitute condition of Kelso and Crawford (1982).

Baldwin and Klemperer (2019) provide an innovative approach characterizing preferences where Walrasian equilibria exist. Instead of working with the value functions, their framework is based on properties of the geometric structure of the regions in the price space where a bidder demands different bundles. A demand type is defined by a list of vectors that give the possible ways in which the individual or aggregate demand can change in response to a small price change. Intuitively, given a valuation v_i , consider the set $\mathcal{L}_i = \{p | D_i(p) > 1\}$ of all prices at which more than one bundle is in the bidder's demand set. \mathcal{L}_i can be shown to form a so-called *polyhedral complex*, and in particular is a union of hyperplanes, which splits price space into multiple full-dimensional regions where a unique bundle is demanded, which are called *unique demand regions (UDRs)*. Now given a set \mathcal{D} of integral vectors, v_i is of the demand type defined by \mathcal{D} if all normals of all hyperplanes in \mathcal{L}_i are integral multiples of vectors in \mathcal{D} .³ We say that the demand type defined by \mathcal{D} is *unimodular* if any linear independent subset of vectors in \mathcal{D} can be extended with integral vectors to a basis with determinant in $\{-1, 1\}$. It can be shown, that if participants' valuations




FIGURE 2 Illustration of \mathcal{L}_i (gray). For each indifference hyperplane, we indicate one of the two normal vectors associated with this hyperplane. We can directly see that these normals all lie in \mathcal{D} as defined in Example 4. The tuples (x_1, x_2) indicate the bundles that are demanded in the respective UDRs

are concave and all have the same unimodular demand type \mathcal{D} , then a Walrasian equilibrium exists. There are several proofs for the unimodularity theorem, see Baldwin and Klemperer (2019); Danilov et al. (2001); Tran and Yu (2015). The authors further show that an equilibrium is guaranteed for more classes of pure complements than of pure substitutes preferences. Note that while all agents being drawn from an equal certain valuation type (SS, GGSC, pure complements) allows for Walrasian equilibria, agent valuations drawn from a mixture of these types in general do not allow for one. Unimodularity of the demand types is a sufficient condition for the existence of Walrasian equilibria. Remarkably, it is also *necessary*: Given valuations of the agents, there exist equilibrium prices for *every* given supply if and only if the agents' demand types are unimodular. Again, whenever the unimodularity condition holds, the optimal solution to the RWDP is integral.

Example 4 Consider a market with two items $\mathcal{K} = \{A, B\}$ and a single bidder with a valuation function v , given by the following table

	x_\emptyset	x_A	x_B	x_{AB}
x	(0, 0)	(1, 0)	(0, 1)	(1, 1)
$v(x)$	0	2	3	4

The set \mathcal{L} is shown in Figure 2. We can see that v is of the demand type given by $\mathcal{D} = \{\pm(1, 0), \pm(0, 1), \pm(1, -1)\}$. It can be checked that \mathcal{D} is actually unimodular.

4.2.3 | From individual to aggregate value functions

We now want to understand when we can expect individual value functions v_i to yield aggregate value functions v_I that are integrally concave. The aggregation of value functions is referred to as *convolution* (see Section 3).

³The normals of these hyperplanes have the following economic meaning: Consider a path in price space starting in some UDR. Each time the path crosses an indifference hyperplane, and thus entering another UDR, the demanded bundle changes by the normal vector of the crossed hyperplane, which points into the opposite direction of the price path. In Figure 2 for example, if the price path goes from the UDR (0, 1) to the UDR (1, 0) in a straight line, we cross the hyperplane with normal (1, -1), and of course (1, 0) = (0, 1) + (1, -1).




FIGURE 3 For a market with two units of a single indivisible item x , the figure shows the aggregate valuation function $v_I(x)$, the aggregate utility function u_I , and the Lyapunov function $L_I(p)$. The Lyapunov function is minimized at $p = 2$, denoting the Walrasian equilibrium prices. Note that $p = 2$ is also the supergradient of $v_I(x)$ at $x = 2$

Murota (2016)[p. 196] shows that if the individual value functions v_i of all bidders $i \in I$ are M^{\natural} -concave, also their convolution is M^{\natural} -concave. Similarly, one can define the aggregate demand correspondence $D_I(p)$, which is equal to the Minkowski sum $\sum_{i \in I} D_i(p)$.

For M^{\natural} -concave functions there is a supergradient at any point that determines a Walrasian price p . To show this, let us consider an arbitrary bounded, integrally convex set $A \subset \mathbb{Z}_{\geq 0}^m$. Let $v_I : A \rightarrow \mathbb{Z}$ be an M^{\natural} -concave valuation on this set. A bundle $x \in A$ is demanded at price $p \in \mathbb{R}^m$ iff $v_I(x) - \langle p, x \rangle \geq v_I(x') - \langle p, x' \rangle \quad \forall x' \in A$, which is equivalent to $v_I(x) + \langle p, x' - x \rangle \geq v_I(x') \quad \forall x' \in A$ (as for divisible goods in Section 3). Figure 3 now illustrates an integrally concave value function on the left and the resulting indirect utility function $u_I(p)$ as well as the Lyapunov function $L_I(p)$ for a single item on the right.

With indivisible items and an integrally concave aggregate value function v_I , bundle x is demanded at p if and only if p is a supergradient of v_I at x . The superdifferential $\partial v_I(x)$ of an integrally concave function $v_I : \mathbb{Z}_{\geq 0}^m \rightarrow \mathbb{R} \cup \{-\infty\}$ at $x \in \text{dom } v_I$ is defined as

$$\partial v_I(x) = \{p \in \mathbb{R}_{\geq 0}^m \mid v_I(y) - v_I(x) \leq \langle p, y - x \rangle \quad \forall y \in \mathbb{Z}_{\geq 0}^m\}.$$

The individual and aggregate value functions are nondecreasing such that the gradient p^* of the superdifferential is $p^* \geq 0$. With an integrally concave value function v_I there exists an integral equilibrium price vector p^* (Murota et al., 2016). The integrality of the prices follows from the fact that an integer-valued M^{\natural} -concave function v_I on $\mathbb{Z}_{\geq 0}^m$ has an integral subgradient at every point x in $\text{dom } v_I$. As both $v_I(x)$ and the subgradient at x are integral, the tangent at $v_I(x)$ has an integral slope p , which can be verified in Figure 3.

An underlying assumption in the study of competitive equilibria is that agents are price-takers, that is, agents

honestly report their true demand in response to prices in each round of an auction. Mechanism design, a line of research initiated by Hurwicz (1972), wants to understand how such markets perform when agents are strategic about their demands. Unfortunately, Gul and Stacchetti (1999) showed that even if goods are substitutes, Walrasian markets are not incentive-compatible. The assignment market, where bidders have unit-demand is an exception where straightforward bidding is actually an ex post equilibrium (Shapley & Shubik, 1971; Demange et al., 1986).

5 | ALGORITHMIC AUCTION MODELS

Auctions can be understood as algorithms to solve a welfare maximization problem. Some algorithms provide models that allow us to understand when an auction can be expected to be efficient and when it yields a Walrasian equilibrium.

The auction proposed by Ausubel (2005) for strong substitutes valuations follows a greedy steepest descent algorithm to minimize the (integrally convex) Lyapunov function (Murota & Tamura, 2003). This algorithm has an intuitive interpretation as an ascending auction: subgradients of the Lyapunov function at p are oversupplies at this price: $\partial L(p) = s - D_I(p)$.⁴ Knowing that the Lyapunov function is equivalent to the DRWDP (see Proposition 1), the overall auction can now be described as a primal-dual algorithm to solve the RWDP. For the price minimization, both algorithms require all subgradients at each point, that is, the entire demand set needs to be revealed. A specific version of a primal-dual algorithm yields the same steps.

We focus on primal-dual algorithms as a consistent algorithmic framework to model Walrasian auction mechanisms. Let us first describe the auction by Ausubel (2005) as a steepest descent algorithm before we introduce the overall primal-dual auction framework.

5.1 | The auction by Ausubel (2005)

The auction algorithm starts with an arbitrary price vector p below the bidder-optimal Walrasian prices, possibly $p(k) = 0$ for all $k \in \mathcal{K}$. The algorithm then searches iteratively in each round $t \in T$ for a subset of goods $S \subseteq \mathcal{K}$ such that $L(p^t) - L(p^t + \mathbb{1}_S)$ is maximized. Here, p^t denotes the prices in round t . This is equivalent to determining the direction of steepest descent to find the global minimum of this function:

⁴Note that subgradient and steepest descent algorithms for convex minimization are equivalent for differential functions but not for the minimization of discrete functions as in the case of markets with indivisible goods. The difference between the two algorithms is that the steepest descent algorithm evaluates all subgradients at a point, while subgradient algorithms use only a single subgradient. This is equivalent to eliciting the entire demand correspondence or only a single bundle from the demand correspondence. As a result, the primal-dual algorithm needs fewer iterations to converge to the exact solution (de Vries et al., 2007).

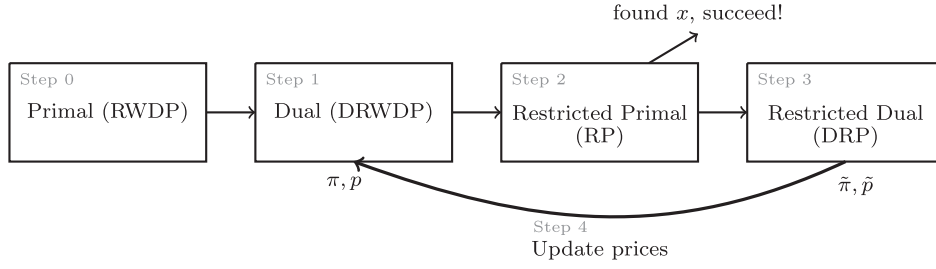


FIGURE 4 A primal-dual algorithm following Papadimitriou and Steiglitz (1998)

- (i) At p^t the auctioneer asks each bidder $i \in \mathcal{I}$ for her entire demand set $D_i(p^t)$.
- (ii) For all potential price update vectors $\tilde{p} \in \{\mathbb{1}_S : S \subseteq \mathcal{K}\}$ the auctioneer determines each bidder's decrease of the indirect utility. The auctioneer chooses the price update $\tilde{p} \in \{\mathbb{1}_S : S \subseteq \mathcal{K}\}$ such that the Lyapunov function is decreased the most, that is, $L(p^t) - L(p^t + \tilde{p})$ is maximized. In case there are multiple such minimizers, the \tilde{p} with the smallest number of positive entries is selected. This price vector is referred to as the minimal minimizer and is guaranteed to be unique.
- (iii) If no nonempty subset S can be found satisfying $L(p^t) - L(p^t + \mathbb{1}_S) > 0$, then the submodularity of the Lyapunov function guarantees that p^t is the bidder-optimal Walrasian price vector and the algorithm terminates. Otherwise the price p^{t+1} is set to $p^t + \tilde{p}$ and the algorithm continues.

With integer valuations, $L(p)$ decreases by at least 1 in each iteration and therefore converges after finitely many steps. Murota et al. (2016) analyze the convergence and number of iterations of this steepest descent algorithm. In particular, if the auction algorithm is initialized with $p(k) = 0$ for all $k \in \mathcal{K}$ and p^* is the minimal equilibrium price, the algorithm terminates in exactly $\|p^*\|_\infty = \max_{k \in \mathcal{K}} |p^*(k)|$ iterations. The price update step described in this subsection can now be interpreted as an operation in a primal-dual algorithm to solve the WDP, as we show next.

5.2 | The primal-dual auction framework

Let us now describe the auction by Ausubel (2005) in the context of the more general primal-dual framework. Primal-dual algorithms (Papadimitriou & Steiglitz, 1998) can be used to compute solutions of the RWDP and DRWDP (see Section 4.1.1). Based on a feasible solution of the DRWDP, one derives a restricted primal RP that determines whether supply equals demand at these prices or not. If this is not the case, the dual restricted primal DRP determines the price increment, which is then added to the current price vector of the dual DRWDP, before a new restricted primal is

computed. The overall process is illustrated in Figure 4. There is some flexibility in choosing each iteration's direction of price adjustment. In this primal-dual auction framework, we compute the price update that yields the steepest descent of the DRWDP.

Instead of solving the RWDP and the DRWDP directly, the primal-dual algorithm replaces these linear programs by a series of other linear programs known as the restricted primal RP and the dual of the restricted primal DRP. As the primal dual algorithm follows the same price trajectory as Ausubel's auction as we will show below, exactly $\|p^*\|_\infty$ iterations must be executed where p^* is the minimal equilibrium price vector (Murota et al., 2016). In each iteration two linear programs (the RP and DRP) must be solved which both are of exponential size in the number of goods. Clearly, the primal dual algorithm does not give any runtime benefits over solving the RWDP and DRWDP directly. However, executing the primal-dual algorithm instead of solving the RWDP and DRWDP directly allows to interpret the auction by Ausubel (2005) in terms of a primal-dual framework. Moreover, unlike the solution obtained by solving the RWDP and DRWDP directly, the allocation and prices computed by the primal-dual algorithm are guaranteed to constitute the Walrasian equilibrium with bidder-optimal prices.

Let us discuss the algorithm in more detail. In an ascending auction the components of the initial price vector are set to $p(k) = 0$ for all $k \in \mathcal{K}$. To obtain an initial feasible dual solution, the dual is solved with these prices to find initial values for the indirect utility π_i of every bidder i .

With a feasible dual solution, one can exploit the complementary slackness conditions to derive an optimal primal solution which defines a welfare-maximizing allocation of bundles to bidders. Naturally, not every feasible dual solution allows for an optimal primal solution. To check this, one solves an optimization problem known as the restricted primal RP problem.

$$\begin{aligned}
 & \max - \sum_{i \in \mathcal{I}} \lambda_i c_i - \sum_{k \in \mathcal{K}} \mu_k d_k & \text{(RP)} \\
 & \text{s.t. } \sum_{x \in \mathcal{X}_i} z_i(x) + a_i + c_i = 1 \quad \forall i \in \mathcal{I}(\tilde{\pi}_i) \\
 & \sum_{i \in \mathcal{I}} \sum_{x \in \mathcal{X}_i} x(k) z_i(x) + b_k + d_k = s(k) \quad \forall k \in \mathcal{K}(\tilde{p}(k)) \\
 & z_i(x), a_i, b_k \geq 0 \quad \forall z_i(x) \in \mathcal{J}_z, \forall a_i \in \mathcal{J}_a, \forall b_k \in \mathcal{J}_b \\
 & z_i(x) = 0, a_i = 0, b_k = 0 \quad \forall z_i(x) \notin \mathcal{J}_z, \forall a_i \notin \mathcal{J}_a,
 \end{aligned}$$

$$\begin{aligned} \forall b_k \notin \mathcal{J}_b \\ c_i, d_k \geq 0 \quad \forall i \in \mathcal{I}, \forall k \in \mathcal{K} \end{aligned}$$

Given a feasible dual solution for the DRWDP, any tight dual constraint $\pi_i \geq v_i(x) - \sum_{k \in \mathcal{K}} x(k)p(k)$ corresponds to a bundle x that maximizes the utility of bidder i at prices p . Thus, the set of tight dual constraints \mathcal{J}_z corresponds to the bidders' demand sets. In case the given dual solution is optimal, the complementary slackness conditions mandate that whenever the dual constraint has slack, that is, $\pi_i > v_i(x) - \sum_{k \in \mathcal{K}} x(k)p(k)$, the corresponding primal variable $z_i(x)$ defining whether bidder i is allocated bundle x equals zero. The interpretation of this is that a bidder is never allocated a bundle not being part of her demand set. Of course, if the given dual solution is not optimal, there might not exist an allocation such that each bidder receives a bundle from her demand set. Therefore, additional slack variables c_i and d_k are introduced to the RP that measure by how much the complementary slackness conditions are violated. A violation may either occur due to bidder i not being allocated a bundle from her demand set ($c_i > 0$) or an item k remaining (partially) unsold ($d_k > 0$). The restricted primal problem tries to find an allocation in which these violations are minimized. In fact, when the optimal solution of the RP equals 0, the complementary slackness conditions are fulfilled so that the current primal and dual solution constitute a Walrasian equilibrium. Otherwise, the price of some items needs to be raised.

Complementary slackness conditions must also hold for the dual constraints $\pi_i \geq 0$ and $p(k) \geq 0$. We denote the set of tight dual constraints by \mathcal{J}_a and \mathcal{J}_b respectively. Due to complementary slackness, the primal variable a_i must equal zero whenever the corresponding dual constraint $\pi_i \geq 0$ has slack. In other words this means that whenever a bidder's indirect utility is positive, she must be allocated a nonempty bundle from her demand set. Similarly, complementary slackness implies that when a price of an item $k \in \mathcal{K}$ is greater than zero, then slack variable b_k must equal zero, which guarantees that all units of item k are allocated in an optimal solution.

In the primal-dual framework of Papadimitriou and Steiglitz (1998) all coefficients λ_i and μ_k in the objective function of the restricted primal RP equal 1. Note that as long as λ_i and μ_k are chosen to be strictly positive, their specific values do not influence the termination criterion of the primal-dual algorithm as one only checks whether the objective equals zero. However, the particular choice of λ_i and μ_k affects the constraints in the dual of the restricted primal DRP, and we will take advantage of this to find a particular price update vector when solving the DRP.

In case the RP objective does not equal zero, the current dual solution of the DRWDP is updated using the solution to the dual of the restricted primal DRP. Solving the DRP essentially means computing a direction $\tilde{\pi}, \tilde{p}$ in which the dual objective function can be improved the most. We set $\tilde{\pi}, \tilde{p}$ such that it minimizes the function $\sum_{i \in \mathcal{I}} (\pi_i + \tilde{\pi}_i) + \sum_{k \in \mathcal{K}} s(k)p(k) +$

$\tilde{p}(k)$). This is equivalent to finding a subgradient to the Lyapunov function as we will show below.

As there may exist multiple potential directions $(\tilde{\pi}, \tilde{p})$ that minimize the Lyapunov function, we need to make small adaptations to the DRP such that the gradient found by the DRP is equivalent to the minimal minimizer in Ausubel's auction. For this purpose we introduce additional constraints $0 \leq \tilde{p}(k) \leq 1$ for all $k \in \mathcal{K}$. As proven in Ausubel (2005), the Lyapunov function restricted to the unit $|\mathcal{K}|$ -dimensional cube $\{p + \tilde{p} : 0 \leq \tilde{p}(k) \leq 1 \forall k \in \mathcal{K}\}$ is minimized on the vertices of this cube. Thus, limiting price updates $\tilde{p}(k)$ to the interval $[0, 1]$ for all $k \in \mathcal{K}$ ensures that the same potential price updates as in Ausubel's auction (i.e., $\{\mathbb{1}_S : S \subseteq \mathcal{K}\}$) are considered. Note that this also implies that in each iteration of our primal-dual auction framework the respective prices and price updates are integer valued.

Another adaption to be made is to choose λ_i suitably large for all $i \in \mathcal{I}$ so that the decrease of utility for each bidder i is unrestricted when raising prices. To guarantee that the gradient found by the DRP is not only a minimizer of the Lyapunov function but a minimal minimizer, price penalties $\tau_k > 0$ are added to the objective function that are small enough so that their impact on the objective value is negligible.

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{I}} \tilde{\pi}_i + \sum_{k \in \mathcal{K}} (s(k) + \tau_k) \tilde{p}(k) & \text{(DRP)} \\ \text{s.t.} \quad & \tilde{\pi}_i + \sum_{k \in \mathcal{K}} x(k) \tilde{p}(k) \geq 0 \quad \forall i, x : z_i(x) \in \mathcal{J}_z \quad (z_i(x)) \\ & \tilde{\pi}_i \geq 0 \quad \forall i : a_i \in \mathcal{J}_a & (a_i) \\ & \tilde{\pi}_i \geq -\lambda_i \quad \forall i : a_i \notin \mathcal{J}_a & (c_i) \\ & \tilde{p}(k) \geq 0 \quad \forall k : b_k \in \mathcal{J}_b & (b_k) \\ & \tilde{p}(k) \geq -\mu_k \quad \forall k : b_k \notin \mathcal{J}_b & (d_k) \\ & 0 \leq \tilde{p}(k) \leq 1 \quad \forall k \in \mathcal{K} \end{aligned}$$

In the following we make the connection between the DRP and the price update step of Ausubel's ascending auction explicit by demonstrating how to transform one approach into the other. Recall that in Ausubel (2006) the goal is to find a $\tilde{p} \in \{\mathbb{1}_S : S \subseteq \mathcal{K}\}$ leaving all entries of $p + \tilde{p}$ nonnegative and minimizing

$$L(p + \tilde{p}) - L(p).$$

Ausubel (2006) shows that for a fixed \tilde{p} it holds that

$$\begin{aligned} L(p + \tilde{p}) - L(p) = \sum_{i \in \mathcal{I}} \max_{x \in D_i(p)} \left\{ - \sum_{k \in \mathcal{K}} x(k) \tilde{p}(k) \right\} \\ + \sum_{k \in \mathcal{K}} s(k) \tilde{p}(k). \end{aligned}$$

The term $\max_{x \in D_i(p)} \{- \sum_{k \in \mathcal{K}} x(k) \tilde{p}(k)\}$ is clearly equal to

$$\begin{aligned} \min \quad & \tilde{\pi}_i \\ \text{s.t.} \quad & \tilde{\pi}_i \geq - \sum_{k \in \mathcal{K}} x(k) \tilde{p}(k) \quad \forall x \in D_i(p) \end{aligned}$$

Consequently, by adjusting notation and noting that \mathcal{J}_z represents the demand set $D_i(p)$, we can rewrite the problem of

minimizing $L(p + \tilde{p}) - L(p)$:

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{I}} \tilde{\pi}_i + \sum_{k \in \mathcal{K}} s(k) \tilde{p}(k) \\ \text{s.t.} \quad & \tilde{\pi}_i + \sum_{k \in \mathcal{K}} x(k) \tilde{p}(k) \geq 0 \quad \forall i, x : z_i(x) \in \mathcal{J}_z \\ & p(k) + \tilde{p}(k) \geq 0 \quad \forall k \in \mathcal{K} \\ & 0 \leq \tilde{p}(k) \leq 1 \quad \forall k \in \mathcal{K} \end{aligned}$$

As argued above, all price updates and consequently also the prices are integral in each step of our primal-dual auction framework. Hence, the second last set of inequalities can be replaced by

$$\tilde{p}(k) \geq 0 \quad \forall k : b_k \in \mathcal{J}_b$$

since \mathcal{J}_b represents all indices where $p(k)$ equals 0.

The only remaining difference to the DRP is that we are missing the inequalities $\tilde{\pi}_i \geq 0$ for $a_i \in \mathcal{J}_a$. From the definition we see, however, that $a_i \in \mathcal{J}_a$ if and only if the utility of bidder i at price p is 0. But this means that the empty bundle is in her demand set. Hence, $\tilde{\pi}_i \geq 0$ is actually one of the constraints $\tilde{\pi}_i + \sum_{k \in \mathcal{K}} x(k) \tilde{p}(k) \geq 0$. As a result we get that one step of the Lyapunov minimization approach is exactly the same as one step of the primal-dual algorithm.

We restricted our attention so far on explaining the relationship between the primal-dual algorithm and the ascending version of the tâtonnement process described by Ausubel (2005). However, similar observations can also be made for the descending version. The only adaptations to be made in our argument concern the formulation of the DRP. Instead of applying positive price penalties τ_k in the objective function, negative ones have to be used to ensure that a maximal minimizer is found in each iteration. Furthermore, the price updates $\tilde{p}(k)$ need to be bounded to the interval $[-1, 0]$ instead of $[0, 1]$. Of course, this also implies that μ_k must be chosen suitably large, that is, $\mu_k \geq 1$, in order to allow for price updates of -1 .

While the auction described by Ausubel (2005) requires the bidders' valuations to satisfy the strong substitutes condition, the primal-dual algorithm also works for other environments, in particular for economies where the preferences of the bidders fulfill the more general GGSC condition. Sun and Yang (2006) propose the *dynamic double-track auction* (DDT) that terminates in a Walrasian equilibrium if bidders bid straightforwardly and have GSC valuations. Given two sets S_1 and S_2 describing two classes of goods, the auctioneer announces start prices of zero for items in S_1 and suitable high start prices in S_2 such that items in S_1 are overdemanded while items in S_2 are underdemanded. In the course of the auction the auctioneer simultaneously adjusts prices of items S_1 upwards but those of items in S_2 downwards.

Shioura and Yang (2015) introduce the *generalized double-track auction* which is an extension of the DDT to multi-item multi-unit economies where bidders' valuations satisfy the GGSC condition. Their auction starts with an arbitrary integral price vector and then proceeds in two phases.

While in the first phase the auctioneer adjusts prices of items in S_1 upwards and prices in S_2 downwards, the price update directions are reversed in the second phase.

Similar to the auction proposed by Ausubel (2005), the price updates in the generalized double-track auction correspond to the steepest descent direction of the Lyapunov function, which can be embedded into a primal-dual algorithm. Essentially, the primal-dual algorithm for the generalized double-track auction combines the DRP adaptations for the ascending and descending version of the auction by Ausubel (2005) as described above. Let the set S_1 and S_2 denote the set of items with an upward and downward moving price trajectory, respectively. While price updates for items in S_1 are bounded to the interval $[0, 1]$, they are restricted to interval $[-1, 0]$ for items in S_2 . Similarly, the price penalties in the objective of the DRP are positive for items in S_1 and negative for items in S_2 . Once the generalized double-track auction moves from the first to the second phase, the price trajectories of items in S_1 and S_2 are inverted so that the adaptations made to the DRP for items in S_1 now apply for items in S_2 and vice versa.

5.3 | Allocation of items

While our paper focuses on the process of determining equilibrium prices, of course, the auctioneer must determine an equilibrium allocation as well. That is, given a target supply s and an equilibrium price vector p^* , we must find allocations $x_i \in D_i(p^*)$ for every bidder, such that $\sum_{i \in \mathcal{I}} x_i = s$. Since we assume access to demand oracles, that is, each bidder i reports her whole demand set $D_i(p^*)$ in each iteration of the auction, and as demand sets only contain integer points, we could just try every of the finitely many combinations of allocations $x_i \in D_i(p^*)$ in order to match the target supply. This approach is however not very efficient: the number of combinations we possibly have to check is $\prod_{i \in \mathcal{I}} |D_i(p^*)|$, which can clearly be exponential.

The allocation problem can also be interpreted as a flow problem: Consider the directed graph $G = (V, A)$ consisting of $|\mathcal{I}| \cdot |\mathcal{K}|$ vertices $b_i(k)$, describing bidder i 's demand of good k , and $|\mathcal{K}|$ vertices $t(k)$, describing the total supply of good k . For each $i \in \mathcal{I}$ and $k \in \mathcal{K}$, there is an arc pointing from $t(k)$ to $b_i(k)$. Now consider a flow x on this graph, where $x_i(k)$ denotes the amount of flow from vertex $b_i(k)$ to vertex $t(k)$. We interpret $x_i(k)$ as the number of units of good k bidder i receives. As usually, given a flow x , and a node v in the graph, the *excess* at node v is the difference of the flow entering the node and the flow leaving the node:

$$\partial x(v) = \sum_{(w,v) \in A} x(w,v) - \sum_{(v,w) \in A} x(v,w).$$

We call the vector ∂x the *boundary* of x . In our above defined graph, we have $\partial x(b_i(k)) = x_i(k)$ and $\partial x(t(k)) = -\sum_{i \in \mathcal{I}} x_i(k)$. The total number of goods of type k should be equal to the supply of good k . Hence, we have the constraint $\partial x(t(k)) = -s(k)$.

Also, each bidder should receive an allocation in her demand set $D_i(p^*)$, so $(\partial x(b_i(1)), \dots, \partial x(b_i(|\mathcal{K}|))) \in D_i(p^*)$ should hold. Thus, the allocation problem can be interpreted as finding a feasible flow with respect to these constraints on the boundary. In the case of strong-substitutes valuations, the demand sets $D_i(p^*)$ are all M^{\downarrow} -convex, so this is an instance of the *M-convex submodular flow problem*. Polynomial-time algorithms have been developed for this problem, many of them are based on well-known algorithms for min-cost flows. For an overview, see for example (Murota, 2003, Ch. 10).

6 | SUMMARY AND RESEARCH AGENDA

A number of assumptions are crucial for the existence of Walrasian equilibria. Apart from (a) *integral concavity of the aggregate value function*, (b) *the bidders' valuations need to be independent* of each other, and all bidders need to be pure payoff maximizers, that is, have a (c) *quasilinear utility function*. Also, we assume that (d) *the bidders are price-takers and truthfully reveal their demand correspondence* in each round. With these assumptions we can guarantee Walrasian equilibria. However, these are strong assumptions, which might not hold in the field.

- (i) Bidder valuations in real-world auctions include complements and substitutes such that Walrasian equilibria might not even exist. Competitive equilibria with nonlinear and personalized prices always exist in ascending auctions under the assumptions above.⁵
- (ii) Quasilinearity is not always given as there might exist budget constraints, spitefulness, or market-power effects. For example, if bidders have financial constraints, quasilinearity is violated, and ascending auctions with budget constrained bidders have only been analyzed recently (Gerard van der Laan, 2016; Yang et al., 2018). Even if one tries to set budget constraints endogenously for bidders, it might not always be possible to implement an efficient outcome via an auction (Bichler & Paulsen, 2018).
- (iii) Finally, bidders might not bid straightforward in a simple clock auction and behave strategically. A number of papers discusses

variations or extensions of simple clock auctions, which yield incentive compatibility (Ausubel, 2006). These are, however, quite different from the simple clock auctions we see in the field.

The assumptions (i)–(iii) above also lead to corresponding research challenges for the operations research community.

1. Most resource allocation problems analyzed in operations research (e.g., scheduling or packing problems) do not satisfy the assumptions that allow for Walrasian equilibria. Duality breaks for nonconvex integer programming problems and new concepts for competitive equilibrium prices need to be derived. The literature on integer programming duality can provide useful insights and guidance how to derive equilibrium prices for such nonconvex allocation problems (Wolsey, 1981).
2. Budget constraints play a major role in many markets. We need to understand equilibria in markets where bidders maximize payoff, but are financially restricted. Very recent results suggest that budget constraints have a substantial impact on the computational complexity of the allocation and pricing problem and require bilevel integer programs which are known to be Σ_2^P -hard (Bichler & Waldherr, 2019). Overall, it will be useful to analyze utility models different from the standard quasi-linear utility function as they have been observed in advertising and other domains where bidders might not maximize payoff but their net present value or return on investment (Fadaei & Bichler, 2017; Baisa, 2017; Baldwin et al., 2020). Effective ways to compute market equilibria in such an environment still need to be developed.
3. Finally, incentive-compatibility plays an important role in small markets where participants can influence the price. Recent research tries to design simple ascending auction and pricing rules that are incentive-compatible (Baranov, 2018). Incentive-compatibility is very restrictive in most environments. For example, in markets with purely quasilinear utilities, the Vickrey–Clarke–Goves mechanism is unique (Green & Laffont, 1979). For larger markets it can also be useful to understand weaker notions of robustness against strategic manipulation (Azevedo & Budish, 2018).

⁵For example, Sun and Yang (2014) introduces an ascending and incentive-compatible auction in markets with only complements using non-linear and anonymous prices. Ausubel and Milgrom (2002), Parkes and Ungar (2000) and de Vries et al. (2007) discuss ascending auctions for markets where bidders have substitutes and complements and allow for discriminatory and non-linear prices. These auctions are incentive-compatible if the bidders' valuations were gross substitutes.

Overall, competitive equilibrium theory is closely related to mathematical optimization and it provides a rich field for operations research to contribute.

ACKNOWLEDGMENT

The financial support from the Deutsche Forschungsgemeinschaft (DFG) (BI 1057/1-8) is gratefully acknowledged. Open access funding enabled and organized by Projekt DEAL.

ORCID

Martin Bichler  <https://orcid.org/0000-0001-5491-2935>

REFERENCES

- Ağralı, S., Tan, B., & Karaesmen, F. (2008). Modeling and analysis of an auction-based logistics market. *European Journal of Operational Research*, 191(1), 272–294.
- Arrow, K. J., & Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, 22, 265–290.
- Arrow, K. J., & Hahn, F. H. (1971). General competitive equilibrium. *Advanced Textbooks in Economics*, 12.
- Ausubel, L. (2005). *Walrasian tatonnement for discrete goods* (Unpublished paper).
- Ausubel, L., & Milgrom, P. (2006). *The lovely but lonely vickrey auction*. In P. Cramton, Y. Shoham, & R. Steinberg (Eds.), *Combinatorial auctions* (pp. 17–40). Cambridge, MA, Ch. 1: MIT Press.
- Ausubel, L. M. (2006). An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, 96(3), 602–629.
- Ausubel, L. M., & Milgrom, P. R. (2002). Ascending auctions with package bidding. *Advances in Theoretical Economics*, 1(1).
- Azevedo, E. M., & Budish, E. (2018). Strategy-proofness in the large. *The Review of Economic Studies*, 86(1), 81–116.
- Baisa, B. (2017). Auction design without quasilinear preferences. *Theoretical Economics*, 12(1), 53–78.
- Baldwin, E., Edhan, O., Jagadeesan, R., Klemperer, P., & Teytelboym, A. (2020). *The equilibrium existence duality: Equilibrium with indivisibilities & income effects*. arXiv preprint arXiv:2006.16939.
- Baldwin, E., & Klemperer, P. (2019). Understanding preferences: “Demand types,” and the existence of equilibrium with indivisibilities. *Econometrica*, 87(3), 867–932.
- Baranov, O. (2018). An efficient ascending auction for private valuations. *Journal of Economic Theory*, 177, 495–517.
- Bertsekas, D. P. (1988). The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research*, 14(1), 105–123.
- Bichler, M., Davenport, A., Hohner, G., & Kalagnanam, J. (2006). Industrial procurement auctions. *Combinatorial Auctions*, 1, 593–612.
- Bichler, M., Fux, V., & Goeree, J. (2018). A matter of equality: Linear pricing in combinatorial exchanges. *Information Systems Research*, 29(4), 1024–1043.
- Bichler, M., Fux, V., & Goeree, J. K. (2019). Designing combinatorial exchanges for the reallocation of resource rights. *Proceedings of the National Academy of Sciences of the United States of America*, 116(3), 786–791.
- Bichler, M., & Goeree, J. K. (2017). *Handbook of spectrum auction design*. Cambridge, MA: Cambridge University Press.
- Bichler, M., & Paulsen, P. (2018). A principal-agent model of bidding firms in multi-unit auctions. *Games and Economic Behavior*, 111, 20–40.
- Bichler, M., & Waldherr, S. (2019). *Computing core-stable outcomes in combinatorial exchanges with financially constrained bidders*. In Proceedings of the 2019 ACM Conference on Economics and Computation (pp. 747–747). New York, NY: ACM.
- Bikhchandani, S., & Mamer, J. W. (1997). Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, 74(2), 385–413.
- Bikhchandani, S., & Ostroy, J. M. (2002). The package assignment model. *Journal of Economic Theory*, 107(2), 377–406.
- Blumrosen, L., & Nisan, N. (2007). Combinatorial auctions. *Algorithmic Game Theory*, 267, 300.
- Caplice, C., & Sheffi, Y. (2003). Optimization-based procurement for transportation services. *Journal of Business Logistics*, 24(2), 109–128.
- Chakrabarty, D., Jain, P., & Kothari, P. (2014). *Provable submodular minimization using Wolfe’s algorithm*. In Proceedings of the 27th International Conference on Neural Information Processing Systems—Volume 1. NIPS’14 (pp. 802–809). Cambridge, MA: MIT Press.
- Chakrabarty, D., Lee, Y. T., Sidford, A., & Wong, S. C.-w. (2017). *Subquadratic submodular function minimization*. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing. STOC 2017 (pp. 1220–1231). New York, NY: Association for Computing Machinery. <https://doi.org/10.1145/3055399.3055419>.
- Cole, R., Devanur, N. R., Gkatzelis, V., Jain, K., Mai, T., Vazirani, V. V., & Yazdanbod, S. (2016). *Convex program duality, fisher markets, and nash social welfare*. arXiv preprint arXiv:1609.06654.
- Colson, B., Marcotte, P., & Savard, G. (2005). Bilevel programming: A survey. *4OR*, 3(2), 87–107.
- Danilov, V., Koshevoy, G., & Murota, K. (2001). Discrete convexity and equilibria in economies with indivisible goods and money. *Mathematical Social Sciences*, 41(3), 251–273.
- de Vries, S., Schummer, J., & Vohra, R. V. (2007). On ascending vickrey auctions for heterogeneous objects. *Journal of Economic Theory*, 132(1), 95–118.
- Demange, G., Gale, D., & Sotomayor, M. (1986). Multi-item auctions. *Journal of Political Economy*, 94(4), 863–872.
- Fadaei, S., & Bichler, M. (2017). Truthfulness with value-maximizing bidders: On the limits of approximation in combinatorial markets. *European Journal of Operational Research*, 260(2), 767–777.
- Fujishige, S., & Yang, Z. (2003). A note on Kelso and Crawford’s gross substitutes condition. *Mathematics of Operations Research*, 28(3), 463–469.
- Gale, D. (1963). A note on global instability of competitive equilibrium. *Naval Research Logistics Quarterly*, 10(1), 81–87.
- Galindo, G., & Batta, R. (2013). Review of recent developments in or/ms research in disaster operations management. *European Journal of Operational Research*, 230(2), 201–211.
- Gerard van der Laan, Z. Y. (2016). An ascending multi-item auction with financially constrained bidders. *Journal of Mechanism and Institution Design*, 1, 109–149.
- Green, J., & Laffont, J. (1979). *Incentives in public decision making*. Amsterdam, Netherlands: North Holland.
- Grötschel, M., Lovász, L., & Schrijver, A. (1981). The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2), 169–197.
- Gul, F., & Stacchetti, E. (1999). Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87, 95–124.

- Herroelen, W., & Leus, R. (2005). Project scheduling under uncertainty: Survey and research potentials. *European Journal of Operational Research*, 165(2), 289–306.
- Hurwicz, L. (1972). On informationally decentralized systems. *Decision and Organization*, 12.
- Kaneko, M. (1976). On the core and competitive equilibria of a market with indivisible goods. *Naval Research Logistics Quarterly*, 23(2), 321–337.
- Kelso, A. S., & Crawford, V. P. (1982). Job matching, coalition formation, and gross substitute. *Econometrica*, 50, 1483–1504.
- Kim, S. (1986). Computation of a large-scale competitive equilibrium through optimization. *Computers and Operations Research*, 13(4), 507–515.
- Klemperer, P. (2010). The product-mix auction: A new auction design for differentiated goods. *Journal of the European Economic Association*, 8(2–3), 526–536.
- Krishna, V. (2009). Auction theory. Boston, MA: Academic Press.
- La Salle, J., & Lefschetz, S. (2012). Stability by Liapunov's direct method with applications by Joseph L Salle and Solomon Lefschetz (Vol. 4). Amsterdam, Netherlands: Elsevier.
- Leme, R. P. (2017). Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, 106, 294–316.
- Lessan, J., & Karabati, S. (2018). A preference-based, multi-unit auction for pricing and capacity allocation. *Computers and Operations Research*, 91, 237–246.
- Madani, M., & Van Vyve, M. (2015). Computationally efficient mip formulation and algorithms for european day-ahead electricity market auctions. *European Journal of Operational Research*, 242(2), 580–593.
- McKenzie, L. W. (1959). On the existence of general equilibrium for a competitive market. *Econometrica: Journal of the Econometric Society*, 27, 54–71.
- Meeus, L., Verhaegen, K., & Belmans, R. (2009). Block order restrictions in combinatorial electric energy auctions. *European Journal of Operational Research*, 196, 1202–1206.
- Milgrom, P., & Strulovici, B. (2009). Substitute goods, auctions, and equilibrium. *Journal of Economic Theory*, 144(1), 212–247.
- Murota, K. (2003). *Discrete convex analysis*. Society for Industrial and Applied Mathematics. Retrieved from <https://epubs.siam.org/doi/abs/10.1137/1.9780898718508>.
- Murota, K. (2016). Discrete convex analysis: A tool for economics and game theory. *Journal of Mechanism and Institution Design*, 1(1), 151–273.
- Murota, K., Shioura, A., & Yang, Z. (2016). Time bounds for iterative auctions: A unified approach by discrete convex analysis. *Discrete Optimization*, 19, 36–62.
- Murota, K., & Tamura, A. (2003). New characterizations of m-convex functions and their applications to economic equilibrium models with indivisibilities. *Discrete Applied Mathematics*, 131(2), 495–512.
- Nesterov, Y. (2018). Lectures on convex optimization. Berlin, Heidelberg: Springer.
- Olafsson, S., Li, X., & Wu, S. (2008). Operations research and data mining. *European Journal of Operational Research*, 187(3), 1429–1448.
- Özer, A. H., & Özturan, C. (2009). A model and heuristic algorithms for multi-unit nondiscriminatory combinatorial auction. *Computers and Operations Research*, 36(1), 196–208 Part Special Issue: Operations Research Approaches for Disaster Recovery Planning.
- Papadimitriou, C. H., & Steiglitz, K. (1998). Combinatorial optimization: Algorithms and complexity. Chelmsford, MA: Courier Corporation.
- Parkes, D. C., & Ungar, L. H. (2000). Iterative combinatorial auctions: Theory and practice (p. 7481). AAAI/IAAI.
- Rockafellar, R. T. (2015). Convex analysis. Princeton, NJ: Princeton University Press.
- Shapley, L. S., & Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, 1(1), 111–130.
- Shioura, A., & Tamura, A. (2015). Gross substitutes condition and discrete concavity for multi-unit valuations: A survey. *Journal of the Operations Research Society of Japan*, 58(1), 61–103.
- Shioura, A., & Yang, Z. (2015). Equilibrium, auction, and generalized gross substitutes and complements. *Journal of the Operations Research Society of Japan*, 58(4), 410–435.
- Starr, R. M. (1969). Quasi-equilibria in markets with non-convex preferences. *Econometrica: Journal of the Econometric Society*, 37, 25–38.
- Sun, N., & Yang, Z. (2006). Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica*, 74(5), 1385–1402.
- Sun, N., & Yang, Z. (2014). An efficient and incentive compatible dynamic auction for multiple complements. *Journal of Political Economy*, 122(2), 422–466.
- Teytelboym, A. (2014). Gross substitutes and complements: A simple generalization. *Economics Letters*, 123(2), 135–138.
- Tran, N. M., & Yu, J. (2015). *Product-mix auctions and tropical geometry*.
- Triki, C., Beraldi, P., & Gross, G. (2005). Optimal capacity allocation in multi-auction electricity markets under uncertainty. *Computers and Operations Research*, 32(2), 201–217.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1), 8–37.
- Wolsey, L. A. (1981). Integer programming duality: Price functions and sensitivity analysis. *Mathematical Programming*, 20(1), 173–195.
- Yang, Z., Yu, J., et al. (2018). *Efficient ascending menu auctions with budget constrained bidders* (Tech. rep).

How to cite this article: Bichler M, Fichtl M, Schwarz G. Walrasian equilibria from an optimization perspective: A guide to the literature. *Naval Research Logistics* 2021;68:496–513. <https://doi.org/10.1002/nav.21963>

APPENDIX A.

Proof of Theorem 2. First, let $\{z_i^*(x)\}_{i \in I, x \in \mathcal{X}_i}$ be an optimal solution to the RWDP and $(\{\pi_i^*\}_{i \in I}, \{p^*(k)\}_{k \in \mathcal{K}})$ be an optimal solution to the DRWDP. By assumption, the optimal value of the WDP is equal to the one of the RWDP, so we may assume that all $z_i^*(x)$ are in $\{0, 1\}$. We may further assume without loss of generality that for each bidder i , there exists exactly one x with $z_i^*(x) = 1$: If $z_i^*(x) = 0$ for all $x \in \mathcal{X}_i$, we can just set $z_i^*(\mathbf{0}) = 1$, where $\mathbf{0}$ is the empty bundle, without altering the value of the WDP, since $v_i(\mathbf{0}) = 0$. Similarly, if for some $k \in \mathcal{K}$, $\sum_{i \in I} \sum_{x \in \mathcal{X}_i} x(k) z_i^*(x) < s(k)$, we may distribute the remaining items of type k arbitrarily among the agents. This does not decrease the value of the WDP because of monotonicity of the agents' valuations. The (possibly altered) variables $z_i^*(x)$ thus constitute an allocation where the whole

supply is distributed among the agents—so the first criterion of a Walrasian equilibrium is satisfied. Let us now check that every bidder receives a bundle in her demand set: If $z_i^*(\bar{x}) = 1$, that is, bidder i receives bundle \bar{x} , we have by complementary slackness $\pi_i = v_i(\bar{x}) - \sum_{k \in \mathcal{K}} \bar{x}(k)p^*(k)$. Since π_i^* is part of an optimal solution,

$$\pi_i^* = \max_{x \in \mathcal{X}_i} v_i(x) - \sum_{k \in \mathcal{K}} x(k)p^*(k).$$

Otherwise, we could decrease π_i^* , making the value of the DRWDP smaller. Consequently, $v_i(\bar{x}) - \sum_{k \in \mathcal{K}} \bar{x}(k)p^*(k) = \max_{x \in \mathcal{X}_i} v_i(x) - \sum_{k \in \mathcal{K}} x(k)p^*(k)$, so \bar{x} is in the demand set of bidder i at prices $\{p^*(k)\}_{k \in \mathcal{K}}$. The second condition of a Walrasian equilibrium is thus satisfied, and $\{p^*(k)\}_{k \in \mathcal{K}}$ are equilibrium prices.

For the other direction, let $\{p^*(k)\}_{k \in \mathcal{K}}$ be equilibrium prices together with an allocation, described by binary variables $\{z_i^*(x)\}_{i \in I, x \in \mathcal{X}_i}$. Let \bar{x} be the bundle with $z_i^*(\bar{x}) = 1$. Set $\pi_i^* = v_i(\bar{x}) - \sum_{k \in \mathcal{K}} \bar{x}(k)p^*(k)$. Since \bar{x} is in the demand set of bidder i , $\pi_i^* \geq v_i(x) - \sum_{k \in \mathcal{K}} x(k)p^*(k)$ for all bundles x , so $(\{p^*(k)\}, \{\pi_i^*\})$ is feasible for the DRWDP ($\pi_i^* \geq 0$ follows from choosing $x = \mathbf{0}$ in the above inequality). By definition of the Walrasian equilibrium, $\{z_i^*(x)\}$ is also feasible for the (R)WDP. All inequalities in the WDP actually hold with equality—so complementary slackness of the primal problem is trivially fulfilled. From the choice of π_i^* we also directly see, that complementary slackness is satisfied for the dual problem. It follows that the optimal value of the WDP equals the optimal value of the DRWDP.

4 Publication 2: On the Expressiveness of Assignment Messages

Peer-Reviewed Journal Paper

Title: On the expressiveness of assignment messages

Authors: Maximilian Fichtl

In: Economics Letters

Abstract: In this note we prove that the class of valuation functions representable via integer assignment messages is a proper subset of strong substitutes valuations. Thus, there are strong substitutes valuations not expressible via assignment messages.

Contribution of dissertation author: Methodology, formal analysis, visualization, paper management

Copyright Notice: © 2021 Elsevier B.V.

Reference: Fichtl (2021)



On the expressiveness of assignment messages

Maximilian Fichtl

Decision Sciences & Systems (DSS), Department of Informatics (I18), Technische Universität München, Boltzmannstr. 3, 85748 Garching, Germany



ARTICLE INFO

Article history:

Received 22 July 2021

Accepted 16 August 2021

Available online 27 August 2021

Keywords:

Strong substitutes

Assignment messages

Auctions

Min-cost flows

ABSTRACT

In this note we prove that the class of valuation functions representable via integer assignment messages is a proper subset of strong substitutes valuations. Thus, there are strong substitutes valuations not expressible via assignment messages.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

Strong substitutes valuations are an important class of valuation functions for indivisible markets, guaranteeing existence of a Walrasian equilibrium. They were introduced by Milgrom and Strulovici (2009) as a multi-unit generalization of the gross substitutes condition for single-unit markets (Kelso and Crawford, 1982).

In sealed-bid auctions where bidders may be assumed to have strong substitutes preferences, it is of major practical importance that bidders can efficiently report their preferences to the auctioneer. Thus, there is need for a bidding language, allowing bidders to express their preferences in a compact and intuitive way, while not further restricting the class of expressible valuations. Milgrom (2009) introduces *integer assignment messages*, in the following only called *assignment messages*, and proves that every valuation function expressible via assignment messages fulfills the strong substitutes condition. However, the question if bidders can express *arbitrary* strong substitutes valuations with assignment messages remained open. In this note we give a negative answer by proving that there are strong substitutes valuations not expressible via assignment messages. Our proof follows the lines of Ostrovsky and Paes Leme (2015), who showed that a related bidding language, called *endowed assignments*, for single unit markets cannot express arbitrary gross substitutes valuations.

As has recently been shown by Baldwin and Klemperer (2021) the Strong Substitutes Product-Mix Auction (SSPMA) (Klemperer, 2008, 2010; Baldwin and Klemperer, 2019) is capable of expressing arbitrary strong substitutes preferences. Thus, our result implies that the SSPMA remains the only known bidding language that allows bidders to express all such (and only such) preferences.

E-mail address: max.fichtl@tum.de.

<https://doi.org/10.1016/j.econlet.2021.110051>

0165-1765/© 2021 Elsevier B.V. All rights reserved.

2. Economic setting

We consider a market with $n \geq 2$ types goods $i \in \{1, \dots, n\}$. A bundle $\mathbf{q} \in \mathbb{Z}^n$ contains q_i items of good i , where a negative value of q_i expresses a willingness to sell. Bidders' preferences are given by *valuation functions* $v : \mathcal{Q} \rightarrow \mathbb{R}$, where $\mathcal{Q} \subset \mathbb{Z}^n$ is a finite set of feasible bundles with $0 \in \mathcal{Q}$. For a price vector $\mathbf{p} \in \mathbb{R}^n$, p_i denotes the cost per unit of good i . Given \mathbf{p} , bidders seek to maximize their quasi-linear utility by choosing a bundle from their *demand set*

$$D(\mathbf{p}) = \arg \max_{\mathbf{q} \in \mathcal{Q}} v(\mathbf{q}) - \langle \mathbf{p}, \mathbf{q} \rangle.$$

The utility of receiving such a bundle is called the *indirect utility* and is denoted by

$$u(\mathbf{p}) = \max_{\mathbf{q} \in \mathcal{Q}} v(\mathbf{q}) - \langle \mathbf{p}, \mathbf{q} \rangle.$$

3. Assignment messages

An integer assignment message (Milgrom, 2009) expresses a bidder's valuation via a linear program. It is determined by a set of $m \in \mathbb{N}$ variables x_j for $j \in J = \{1, \dots, m\}$, where each variable is associated with one of the $n \geq 2$ types of goods $k_j \in \{1, \dots, n\}$, and with a value $v_j \in \mathbb{R}$. We assume that for each good i there is at least one variable associated with it – if not, introduce dummy variables with a value of 0. We define $R_i = \{j \in J : k_j = i\}$ as the set of all variables associated with good i . Additionally, the bidder provides a set $\mathcal{I} \subset \mathcal{P}(J)$ of inequalities. Each inequality $I \in \mathcal{I}$ is a subset of the variables J and is associated with integral upper and lower bounds $u(I) \geq 0$, $\ell(I) \leq 0$, describing the linear constraint $\ell(I) \leq x(I) \leq u(I)$, where $x(I) = \sum_{j \in I} x_j$. The value $v(\mathbf{q})$ for a bundle $\mathbf{q} \in \mathcal{Q}$ is given by

$$v(\mathbf{q}) = \max \sum_{j=1}^m v_j x_j \quad (\text{VAL})$$

$$\begin{aligned} \text{s.t. } \ell(I) \leq x(I) \leq u(I) \forall I \in \mathcal{I} \\ x(R_i) = q_i \forall i = 1, \dots, n. \end{aligned}$$

Here, $\mathcal{Q} \subset \mathbb{Z}^n$ is the set of all \mathbf{q} for which (VAL) has a feasible solution. The indirect utility $u(\mathbf{p}) = \max_{\mathbf{q} \in \mathcal{Q}} v(\mathbf{q}) - \langle \mathbf{p}, \mathbf{q} \rangle$ can be expressed via

$$\begin{aligned} u(\mathbf{p}) = \max \sum_{i=1}^n \sum_{j \in R_i} (v_j - p_i) x_j \quad (\text{IU}) \\ \text{s.t. } \ell(I) \leq x(I) \leq u(I) \forall I \in \mathcal{I}. \end{aligned}$$

The demand set $D(\mathbf{p})$ of maximizers of $v(\mathbf{q}) - \langle \mathbf{p}, \mathbf{q} \rangle$ is the set of all $\mathbf{q} \in \mathcal{Q}$ that can be written as $q_i = x(R_i)$ where \mathbf{x} is an integral solution to (IU). The set of inequalities \mathcal{I} may not be chosen arbitrarily, but must possess a certain tree structure. The following two definitions are taken from Milgrom (2009).

Definition 1. A nonempty subset $\mathcal{T} \subseteq \mathcal{P}(J)$ is called a *tree*, if for any $K, L \subseteq \mathcal{T}$ with $K \cap L \neq \emptyset$ there holds $K \subseteq L$ or $L \subseteq K$. For $K \in \mathcal{T}$, we call the inclusion-minimal set $L \in \mathcal{T}$ with $L \supseteq K$ the *predecessor* of K , if such L exists. Conversely, we call each K , such that L is the predecessor of K , a *successor* of L . We write $s_{\mathcal{T}}(L) = \{K : L \text{ predecessor of } K \text{ in } \mathcal{T}\}$ for the set of successors of L in \mathcal{T} .

Definition 2. The variables J and inequalities \mathcal{I} define an *assignment message*, if $\mathcal{I} = \mathcal{T}_0 \cup \dots \cup \mathcal{T}_n$ is the union of $n + 1$ trees, such that

- for $i = 1, \dots, n$, \mathcal{T}_i only contains inequalities in variables associated with good i : $\mathcal{T}_i \subseteq \mathcal{P}(R_i)$. Furthermore, $R_i \in \mathcal{T}_i$ and $\{j\} \in \mathcal{T}_i$ for all $j \in R_i$.
- $J \in \mathcal{T}_0$ and $\{j\} \in \mathcal{T}_0$ for all $j \in J$. We also write $R_0 = J$.

For each tree \mathcal{T}_i there is a unique element $R_i \in \mathcal{T}_i$ without predecessor, called the *root* of the tree. The only elements in \mathcal{T}_i that are no predecessors of any other element are the singletons $\{j\}$, which we also call the *terminal nodes*. In the following, we write $s_i(L) := s_{\mathcal{T}_i}(L)$ for the set of successors of L .

Since $n \geq 2$, the trees can always be chosen such that they intersect only in terminal nodes.

4. Strong exchangeability

Ostrovsky and Paes Leme (2015) show that there are gross substitutes valuations that are not expressible via *endowed assignments* (Hatfield and Milgrom, 2005). They observe that all endowed assignment valuations satisfy a certain property, called *strong exchangeability*, and provide a gross substitutes valuation that is not strongly exchangeable, which is then consequently not expressible via endowed assignments. For two vectors \mathbf{q} and \mathbf{r} , denote by $\text{supp}_+ \mathbf{q} - \mathbf{r}$ the set of indices i with $q_i - r_i > 0$.

Definition 3 (*Single-Unit Strong Exchangeability* (Ostrovsky and Paes Leme, 2015)). A valuation $v : \{0, 1\}^n \rightarrow \mathbb{R}$ satisfies *strong exchangeability*, if for every price vector \mathbf{p} and all bundles $\mathbf{q}, \mathbf{r} \in D(\mathbf{p})$ with a minimal number of items, i.e., $\sum_i q_i = \sum_i r_i = \min_{\mathbf{q}' \in D(\mathbf{p})} \sum_i q'_i$, there is a bijection $\sigma : \text{supp}_+ \mathbf{q} - \mathbf{r} \rightarrow \text{supp}_+ \mathbf{r} - \mathbf{q}$, such that $\mathbf{q} - \mathbf{e}_i + \mathbf{e}_{\sigma(i)} \in D(\mathbf{p})$ and $\mathbf{r} - \mathbf{e}_{\sigma(i)} + \mathbf{e}_i \in D(\mathbf{p})$ for all $i \in \text{supp}_+ \mathbf{q} - \mathbf{r}$.

Theorem 4.1 (Ostrovsky and Paes Leme, 2015). *There are gross substitutes valuations not satisfying strong exchangeability.*

Our proof follows the same lines: first, we provide a multi-unit extension of strong exchangeability, and then we show that all valuations induced by assignment messages satisfy this property.

Definition 4 (*Multi-Unit Strong Exchangeability*). A valuation $v : \mathcal{Q} \rightarrow \mathbb{R}$ satisfies *strong exchangeability*, if for every price vector \mathbf{p} and all bundles $\mathbf{q}, \mathbf{r} \in D(\mathbf{p})$ with a minimal number of items, there is a correspondence $\sigma \in \text{supp}_+ \mathbf{q} - \mathbf{r} \times \text{supp}_+ \mathbf{r} - \mathbf{q}$, such that

1. For each $(i, j) \in \sigma$, $\mathbf{q} - \mathbf{e}_i + \mathbf{e}_j \in D(\mathbf{p})$ and $\mathbf{r} + \mathbf{e}_i - \mathbf{e}_j \in D(\mathbf{p})$
2. For each $i \in \text{supp}_+ \mathbf{q} - \mathbf{r}$ and $j \in \text{supp}_+ \mathbf{r} - \mathbf{q}$, we have $1 \leq |\{j' : (i, j') \in \sigma\}| \leq q_i - r_i$ and $1 \leq |\{i' : (i', j) \in \sigma\}| \leq r_j - q_j$.

Remark. In single-unit markets Property 2 says that for every $i \in \text{supp}_+ \mathbf{q} - \mathbf{r}$, there is exactly one $j \in \text{supp}_+ \mathbf{r} - \mathbf{q}$ such that $(i, j) \in \sigma$ and vice-versa. Thus, σ represents a bijection $\sigma : \text{supp}_+ \mathbf{q} - \mathbf{r} \rightarrow \text{supp}_+ \mathbf{r} - \mathbf{q}$, so for single-unit markets Definitions 3 and 4 are equivalent.

In order to prove that every assignment message satisfies strong exchangeability, we show that computing the indirect utility of an assignment message valuation can be interpreted as a min-cost flow problem. Given the tree structure of assignment messages from Definitions 1 and 2, we can transform the indirect utility problem (IU) by variable substitution as follows: for each $I \in \mathcal{I}$ introduce a variable y_I representing $y_I = x(I)$. Note that since $\{j\} \in \mathcal{I}$ for all $j \in J$, there are variables $y_{\{j\}}$ corresponding to the variables x_j . If $I \in \mathcal{T}_i$ is not a singleton, I is the disjoint union of all its successors $K \in \mathcal{T}_i$, so

$$y_I = x(I) = \sum_{K \in s_i(I)} x(K) = \sum_{K \in s_i(I)} y_K.$$

Similarly, we have

$$y_{R_0} = x(R_0) = \sum_{i=1}^n x(R_i) = \sum_{i=1}^n y_{R_i}.$$

The constraints $\ell(I) \leq x(I) \leq u(I)$ translate to $\ell(I) \leq y_I \leq u(I)$. Using these observations, one can see that Problem (IU) can equivalently be formulated as

$$\min \sum_{i=1}^n \sum_{j \in R_i} (p_i - v_j) y_{\{j\}} \quad (\text{MCF})$$

$$\text{s.t. } y_I - \sum_{K \in s_0(I)} y_K = 0 \forall I \in \mathcal{T}_0 \setminus \{\{j\} : j \in R_0\} \quad (1)$$

$$\sum_{K \in s_i(I)} y_K - y_I = 0 \forall I \in \mathcal{T}_i \setminus \{\{j\} : j \in R_i\} \forall i = 1, \dots, n \quad (2)$$

$$\sum_{i=1}^n y_{R_i} - y_{R_0} = 0 \quad (3)$$

$$\ell(I) \leq y_I \leq u(I) \forall I \in \mathcal{I} \quad (4)$$

where instead of maximizing the objective function of (IU), we minimize the negative objective function to be consistent with literature on min-cost flows. The following lemma is a simple consequence of our variable substitution, so the proof is omitted.

Lemma 1. *Let $v : \mathcal{Q} \rightarrow \mathbb{R}$ be an assignment message. Then $\mathbf{q} \in D(\mathbf{p})$ if and only if there is an integral solution to (MCF) with $q_i = y_{R_i}$ for all $i \geq 1$.*

One can check that each variable y_I for $I \in \mathcal{I}$ appears exactly twice in the set of equality constraints of (MCF), once with coefficient 1, and once with coefficient -1 . For example, consider $I \in \mathcal{T}_i$ with $i \geq 1$ and $I \neq R_i$ not a singleton. Since I is no singleton, y_I appears with negative sign in Eq. (2). On the other hand, I is the successor of exactly one element, so y_I also appears with positive sign exactly once in (2). All other $I \in \mathcal{I}$ can be checked similarly.




Fig. 1. Directed graph from Example 4.1. The labels on the vertices indicate the equality constraint in (MCF) they correspond to.

Hence, if we collect the variables y_l in the vector $\mathbf{y} = (y_l)_{l \in \mathcal{I}}$ and write constraints (1)–(3) in matrix form as $A\mathbf{y} = \mathbf{0}$, A is the incidence matrix of a directed graph, where \mathcal{I} is the set of arcs, and each of the constraints from (1)–(3) corresponds to a vertex in the graph. For any such vertex, l is an ingoing arc if y_l appears with coefficient 1, and an outgoing arc if it appears with coefficient -1 . Consequently, Problem (MCF) is a min-cost flow problem where y_l denotes the flow along arc l .

Example 4.1. Suppose a bidder submits an assignment message in four variables $J = R_0 = \{1, 2, 3, 4\}$, where $R_1 = \{1, 2, 3\}$ and $R_2 = \{4\}$. The submitted inequalities induce the trees $\mathcal{T}_0 = \{R_0, \{2, 3, 4\}, \{1\}, \dots, \{4\}\}$, $\mathcal{T}_1 = \{R_1, \{1, 2\}, \{1\}, \dots, \{3\}\}$ and $\mathcal{T}_2 = \{R_2\}$. The directed graph corresponding to the incidence matrix A is shown in Fig. 1.

We recall some properties of min-cost flows.

Lemma 2 (Properties of Flows (Ahuja et al., 1993)). Let $G = (V, A)$ be a directed graph with vertex set V and arc set A . Let $\mathbf{f} : A \rightarrow \mathbb{Z}$ be a flow on G .

1. If \mathbf{f} is balanced at every vertex, i.e.,

$$\sum_{a=(w,v) \in A} f_a - \sum_{a=(v,w) \in A} f_a = 0 \forall v \in V,$$

then \mathbf{f} can be decomposed into finitely many cycles: there are subsets $C^1, \dots, C^m \subseteq A$ of arcs, such that each C^k is an undirected cycle in G , and balanced flows $\mathbf{c}^k : C^k \rightarrow \{-1, 1\}$, such that $\mathbf{f} = \sum_{k=1}^m \mathbf{c}^k$. Moreover, we have $c_a^k > 0 \Rightarrow f_a > 0$ and $c_a^k < 0 \Rightarrow f_a < 0$ for all $a \in A$ and all $k = 1, \dots, m$.

2. Suppose \mathbf{f} is an optimal solution to the min-cost flow problem

$$\begin{aligned} & \min \sum_{a \in A} w_a f_a \\ & \text{s.t. } \sum_{a=(w,v) \in A} f_a - \sum_{a=(v,w) \in A} f_a = s(v) \forall v \in V \\ & \ell(a) \leq f_a \leq u(a) \forall a \in A \end{aligned}$$

for some given weights w_a , supplies $s(v)$ and bounds $\ell(a), u(a)$. Then \mathbf{f} does not contain any negative cycle: for $C \subseteq A$ an undirected cycle and a balanced flow $\mathbf{c} : C \rightarrow \mathbb{Z}$ such that $\mathbf{f} + \mathbf{c}$ is feasible, we have $\sum_{a \in A} w_a c_a \geq 0$.

Remark. In their book, Ahuja et al. (1993) consider only non-negative flows, as arbitrary flow problems can be easily transformed into non-negative ones. For the sake of brevity, we allow negative flows here. Note that the proofs given in their book for the mentioned flow properties do actually not require non-negativity. Property 1 follows from the construction in the proof

of Theorem 3.5 in their book, while Property 2 follows from Theorem 3.8.

We now prove our main result.

Theorem 4.2. Let v be a valuation induced by an assignment message. Then v satisfies the strong exchangeability property.

Proof. Let $\mathbf{q}, \mathbf{r} \in D(\mathbf{p})$ be bundles containing a minimal number of goods and $\mathbf{y}^q, \mathbf{y}^r$ corresponding integral solutions to (MCF) with $q_i = y_{R_i}^q$ and $r_i = y_{R_i}^r$ for all i . We are going to construct a correspondence σ satisfying the properties from Definition 4. Since \mathbf{y}^q and \mathbf{y}^r are balanced, i.e., $A\mathbf{y}^q = A\mathbf{y}^r = \mathbf{0}$, so is $\mathbf{y}^q - \mathbf{y}^r$, and we can write $\mathbf{y}^q - \mathbf{y}^r = \mathbf{c}^1 + \dots + \mathbf{c}^m$ where each \mathbf{c}^k is supported on a cycle C^k by Lemma 2. The flows \mathbf{c}^k have the following properties:

- (i) $\mathbf{y}^r + \mathbf{c}^k$ and $\mathbf{y}^q - \mathbf{c}^k$ are optimal solutions to (MCF) for every k .
- (ii) $|\{i \geq 1 : c_{R_i}^k = 1\}| = |\{i \geq 1 : c_{R_i}^k = -1\}| \in \{0, 1\}$ for every k .

To see (i), we first note that $\mathbf{y}^r + \mathbf{c}^k$ is feasible for Problem (MCF): as \mathbf{y}^r and \mathbf{c}^k are balanced, so is $\mathbf{y}^r + \mathbf{c}^k$. Concerning the inequality constraints (4), if $c_j^k > 0$, then it follows from Property 1 in Lemma 2 that $u(l) \geq y_l^q \geq y_l^r + c_l^k > \ell(l)$. With a similar argument we can treat the case $c_j^k < 0$. Consequently, by Property 2, we have

$$\sum_{i=1}^n \sum_{j \in R_i} (p_i - v_j) c_{ij}^k \geq 0.$$

With the same argument applied to $\mathbf{y}^q - \mathbf{c}^k$, we get

$$\sum_{i=1}^n \sum_{j \in R_i} (p_i - v_j) c_{ij}^k \leq 0,$$

so $\sum_{i=1}^n \sum_{j \in R_i} (p_i - v_j) c_{ij}^k = 0$. Hence, the objective values in (MCF) of the flows $\mathbf{y}^q, \mathbf{y}^r, \mathbf{y}^r + \mathbf{c}^k$ and $\mathbf{y}^q - \mathbf{c}^k$ are all equal and thus optimal.

Let us now prove (ii). To that goal, note that, since \mathbf{q} and \mathbf{r} are bundles with a minimum number of elements, we have $\mathbf{y}_{R_0}^q = \mathbf{y}_{R_0}^r$, so by Property 1 of Lemma 2 we have $c_{R_0}^k = 0$ for all k . Consider the flow of \mathbf{c}^k through the vertex corresponding to constraint (3), i.e., representing the equality

$$\sum_{i=1}^n c_{R_i}^k - c_{R_0}^k = 0.$$

As \mathbf{c}^k is supported on a cycle, at most two of the appearing variables $c_{R_i}^k$ can be nonzero. Thus, since $c_{R_0}^k = 0$, either no or exactly two of the c_{R_i} are nonzero, and since their sum equals 0, one must be 1, and the other must be -1 .

Now define $\sigma \in \text{supp}_+ \mathbf{q} - \mathbf{r} \times \text{supp}_+ \mathbf{r} - \mathbf{q}$ by

$$\sigma = \left\{ (i, j) : \exists k : c_{R_i}^k = 1 \wedge c_{R_j}^k = -1 \right\}.$$

σ has the required properties from Definition 4: let $(i, j) \in \sigma$. Then there is some \mathbf{c}^k , such that $c_{R_i}^k = 1, c_{R_j}^k = -1$ and $c_{R_l}^k = 0$ for $l \notin \{i, j\}$. From observation (i) above we have that $\mathbf{y}^r + \mathbf{c}^k$ is an optimal solution to problem (MCF), and the demanded bundle corresponding to that solution is $\mathbf{r} + \mathbf{e}_i - \mathbf{e}_j$. Similarly, the requested bundle corresponding to $\mathbf{y}^q - \mathbf{c}^k$ is $\mathbf{y}^q - \mathbf{e}_i + \mathbf{e}_j$, so Property 1 from Definition 4 is satisfied.

For Property 2 from Definition 4, let $i \in \text{supp}_+ \mathbf{q} - \mathbf{r}$. We need to show that $1 \leq |\{j' : (i, j') \in \sigma\}| \leq q_i - r_i$. Since $q_i > r_i$ and $\mathbf{y}^q - \mathbf{y}^r = \mathbf{c}^1 + \dots + \mathbf{c}^m$, there must be some k with $c_{R_i}^k = 1$.

Consequently, there is some j with $c_{R_j}^k = -1$, which proves the lower bound. Moreover, by Property 1 of Lemma 2, there is no flow \mathbf{c}^k with $c_{R_i}^k = -1$. Thus, there are at most $q_i - r_i$ flows with $c_{R_i}^k = 1$, proving the upper bound. \square

Theorem 4.2 together with Theorem 4.1 directly imply that assignment messages do not cover all strong substitutes valuations.

Corollary 1. *There are strong substitutes valuations that are not representable via assignment messages.*

Proof. Each assignment message satisfies the strong exchangeability property from Definition 4. However, by Theorem 4.1 by Ostrovsky and Paes Leme (2015), there exist gross substitutes valuations that are not strongly exchangeable. Since gross substitutes valuations are a subset of strong substitutes valuations, and Definitions 3 and 4 are equivalent for single-unit markets, the result follows. \square

Acknowledgments

I would like to thank Edwin Lock from the University of Oxford for his very valuable comments and suggestions.

References

- Ahuja, R.K., Magnanti, T.L., Orlin, J.B., 1993. *Network Flows - Theory, Algorithms, and Applications*. Prentice Hall, London.
- Baldwin, E., Klemperer, P., 2019. Understanding preferences: demand types, and the existence of equilibrium with indivisibilities. *Econometrica* 87 (3), 867–932.
- Baldwin, E., Klemperer, P., 2021. Proof that the strong substitutes product-mix auction bidding language can represent any strong substitutes preferences. <https://elizabeth-baldwin.me.uk/papers/strongsubsproof.pdf>.
- Hatfield, J.W., Milgrom, P.R., 2005. Matching with contracts. *Amer. Econ. Rev.* 95 (4), 913–935.
- Kelso, A.S., Crawford, V.P., 1982. Job matching, coalition formation, and gross substitutes. *Econometrica* 50 (6), 1483–1504.
- Klemperer, P., 2008. A new auction for substitutes: Central bank liquidity auctions, the u.s. tarp, and variable product-mix auctions. <https://www.nuffield.ox.ac.uk/economics/Papers/2008/substsauc.pdf>.
- Klemperer, P., 2010. The product-mix auction: a new auction design for differentiated goods. *J. Eur. Econom. Assoc.* 8, 526–536.
- Milgrom, P., 2009. Assignment messages and exchanges. *Am. Econ. J. Microecon.* 1 (2), 95–113.
- Milgrom, P., Strulovici, B., 2009. Substitute goods, auctions, and equilibrium. *J. Econom. Theory* 144 (1), 212–247.
- Ostrovsky, M., Paes Leme, R., 2015. Gross substitutes and endowed assignment valuations. *Theor. Econ.* 10 (3), 853–865.

5 Publication 3: Strong Substitutes: Structural Properties, and a New Algorithm for Competitive Equilibrium Prices

Peer-Reviewed Journal Paper

Title: Strong substitutes: structural properties, and a new algorithm for competitive equilibrium prices

Authors: Elizabeth Baldwin, Martin Bichler, Maximilian Fichtl, Paul Klemperer

In: Mathematical Programming

Abstract: We show the Strong Substitutes Product-Mix Auction bidding language provides an intuitive and geometric interpretation of strong substitutes as Minkowski differences between sets that are easy to identify. We prove that competitive equilibrium prices for agents with strong substitutes preferences can be computed by minimizing the difference between two linear programs for the positive and the negative bids with suitably relaxed resource constraints. This also leads to a new algorithm for computing competitive equilibrium prices which is competitive with standard steepest descent algorithms in extensive experiments.

Contribution of dissertation author: Methodology, formal analysis, visualization, programming, joint paper management

Copyright Notice: Published under a Creative Commons license:
<https://creativecommons.org/licenses/by/4.0/>

Reference: Baldwin et al. (2022)



Strong substitutes: structural properties, and a new algorithm for competitive equilibrium prices

Elizabeth Baldwin¹ · Martin Bichler² · Maximilian Fichtl² · Paul Klemperer³

Received: 27 January 2021 / Accepted: 17 February 2022

© The Author(s) 2022

Abstract

We show the strong substitutes product-mix auction bidding language provides an intuitive and geometric interpretation of strong substitutes as Minkowski differences between sets that are easy to identify. We prove that competitive equilibrium prices for agents with strong substitutes preferences can be computed by minimizing the difference between two linear programs for the positive and the negative bids with suitably relaxed resource constraints. This also leads to a new algorithm for computing competitive equilibrium prices which is competitive with standard steepest descent algorithms in extensive experiments.

Keywords Competitive equilibrium · Walrasian equilibrium · Strong substitutes · Product-Mix auction · Envy-free prices · Indivisible goods · Equilibrium computation · DC programming · Auction theory · Algorithms · product mix auction

Mathematics Subject Classification 91-08

Maximilian Fichtl
max.fichtl@tum.de

Elizabeth Baldwin
elizabeth.baldwin@economics.ox.ac.uk

Martin Bichler
bichler@in.tum.de

Paul Klemperer
paul.klemperer@nuffield.ox.ac.uk

¹ Department of Economics and Hertford college, Oxford University, Oxford, UK

² Department of Computer Science, Technical University of Munich, Munich, Germany

³ Department of Economics and Nuffield College, Oxford University, Oxford, UK

1 Introduction

This paper shows that for an important and widely-studied class of problems—those for which agents have strong substitutes valuations over multiple units of multiple differentiated goods—competitive-equilibrium prices can be found by considering two linear programs. Specifically, we relax resource constraints on both programs in the same way, and find the relaxation that minimizes the difference between the objectives of the two programs; the dual prices of one of these relaxed programs are competitive equilibrium prices. We derive this result by using the geometric representation of preferences provided by the Strong Substitutes Product-Mix Auction (SSPMA) bidding language. This then allows us to develop an efficient algorithm to find the competitive equilibrium prices when preferences are represented this way. Since, as we detail below, the SSPMA language is a natural way for agents to express their preferences, our algorithm is a practical way to find competitive equilibrium prices for strong substitutes.

Our paper also provides a novel algorithm to find the prices in an SSPMA, since these are prices that would be competitive equilibrium prices for the given aggregate supply if bidders had bid their actual values.¹ Participants in SSPMAs make bids that express “strong-substitutes” preferences for multiple units of multiple, differentiated, indivisible goods. Strong substitutes preferences are those that would be ordinary substitutes preferences if every unit of every good were treated as a separate good [35]. They are an extension of gross substitutes preferences [25] from single-unit to multi-unit, multi-item markets, and are equivalent to M^{\sharp} -concavity [17, 37, 45]. They have many attractive properties. In particular, all agents having strong substitutes preferences is a sufficient condition for the existence of competitive equilibrium prices in markets with indivisible goods.

Furthermore, even though strong substitutes are a small subset of the set of all possible valuation functions of a bidder, they are practically relevant for various applications such as auctions used by the Bank of England [26, 28]. So a significant amount of theoretical literature has been devoted to markets where participants have these valuations [4, 8, 38, 41].

Importantly, valid bids in the SSPMA bidding language permits the specification of *precisely* the set of preferences that are strong substitutes, and indeed is the *only* language that is known to do this.² As we will see, it is also parsimonious, or “compact”, in that many valuations can be expressed using only a small number of simple bids.³ Finally, it expresses valuations in a natural way, which can be understood and analyzed

¹ Product-Mix Auctions give envy-free allocations to bidders who express their valuations truthfully. The auctioneer can express its own preferences, and if all the bidders and the auctioneer express their true valuations (the Bank of England does in its role as a product-mix auctioneer, and bidders approximate this if no one bidder is too large) then the auction yields a competitive equilibrium.

² See [7, 9]. By contrast, [40] show [22]’s endowed assignment messages cannot express all strong substitute valuations, [18] likewise shows [34]’s (integer) assignment messages cannot express all strong substitute valuations, and [48] shows that it is not possible to express all strong substitute valuations as combinations of weighted ranks of matroids on a ground set bounded by the number of goods.

³ See [19] for a discussion of compactness.

geometrically; we show aggregate demand is the Minkowski difference between two easily identified demand sets.

1.1 The strong substitutes product-mix auction (SSPMA)

There is significant literature on computing competitive equilibria with strong substitutes valuations. See, for example, [4, 11, 13, 21, 25, 35, 38, 39, 41, 42].⁴ The interest in strong substitutes is due to the fact that it captures practically relevant valuations for indivisible goods, but the allocation problem can be solved in polynomial time and Walrasian competitive equilibrium prices always exist, which is not the case for general valuations [14].

Prior literature requires either value oracles for exponentially many bundles, or demand oracles. Demand oracles can be understood as indirect or iterative mechanisms, where bidders reveal their demand correspondence for a set of prices specified by the auctioneer. So in a large market with many goods that is organized as a sealed-bid auction, the auctioneer needs to perform an exponential number of value queries for each bidder before the allocation algorithm can be run. Such enumerative (XOR) bid languages are used in spectrum auctions, but can lead to “missing bids” problems, which can significantly affect prices, and also create efficiency losses [12].⁵

The SSPMA was developed by [26] for the Bank of England to provide liquidity to financial institutions by auctioning loans to them. The SSPMA is neither based on a value nor a demand oracle.⁶ A collection of bids specifies a large number of package values, which mitigates the missing bids problem. This type of preference elicitation permits efficient ways to compute Walrasian prices, and allows us to uncover new properties of strong substitutes valuations.

Each bidder makes a set of bids, each of which is a vector \mathbf{b} , incorporating an integer weight $w(\mathbf{b})$. Each bidder’s set of bids is interpreted as specifying a quasi-linear utility function over multiple units of each of n goods plus money. A bid in which $w(\mathbf{b}) > 0$ (a “positive” bid) is simply interpreted as a bid for up to, but not more than, $w(\mathbf{b})$ units, in total, of the goods $i = 1, \dots, n$, in which the expressed value of receiving x_i units of good i is $x_i \cdot b_i$.

Example 1 A bidder might be interested in 2 units, and be willing to pay up to price 2 for each unit of good 1, but only up to price 1 for each unit of good 2. These preferences can be implemented by a single bid $\mathbf{b} = (2, 1)$ with $w(\mathbf{b}) = 2$. Figure 1 shows how the bid \mathbf{b} divides price space into three regions: for example, if the price vector $\mathbf{p} = (p_1, p_2)$ lies in the region labeled as “(2, 0) demanded” then, at this price vector, receiving the bundle (2, 0) maximizes the bidder’s utility among all feasible

⁴ [42] provides the fastest algorithm for value oracles and a new algorithm for aggregate demand queries. However, the latter is different in spirit to our paper which addresses a market design for applications such as the Bank of England’s.

⁵ Bidders who do not submit the very large number of bids required to fully specify their valuations are treated as if they place no value on the packages they fail to bid for.

⁶ If a demand oracle is what is available, a conversion to SSPMA is available via [20]’s algorithm which computes the (unique) list of bids corresponding to a bidder’s demand preferences, given access to either a demand or a valuation oracle.

Fig. 1 Example of using a single bid to represent preferences in a Product-Mix auction. The single bid with weight 2 implements the preferences of Example 1. The total demand generated by the bid is indicated in each region of price space


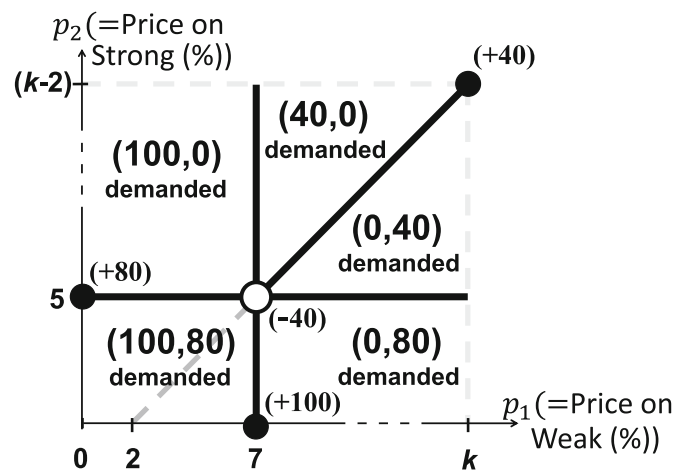


Fig. 2 Example of using positive and negative bids to represent preferences in a Product-Mix auction. The set of bids implements the preferences of Example 2. The sizes of the bids (\$millions) are shown next to the black and white circles that represent the positive and negative bids, respectively. The total demand generated by the complete set of bids, (\$millions of weak, \$millions of strong), is indicated in each region of price space



bundles. The black lines mark the borders at which the demanded bundle changes. If prices lie exactly on the boundary of two or more regions, then the set of demanded bundles is given by the discrete convex hull of the bundles demanded in the adjacent regions. For example, if $\mathbf{p} = (2, 4)$, the bidder demands bundles $(0, 0)$, $(1, 0)$ or $(2, 0)$.

Bids in which $w(\mathbf{b}) < 0$ (“negative” bids) are interpreted as “cancellation” bids that cancel part of the demand created by positive bids. But this means that all bids can be treated by the auctioneer in exactly the same way: a bid is accepted whenever at least one of its prices exceeds the corresponding auction price and, if it is accepted, then it is allocated the good on which its price exceeds the corresponding auction price by most.⁷ The following example from [26, 27] demonstrates the potential usefulness of negative bids in the context of the Bank of England’s auctions, in which the different goods were “weak collateral” and “strong collateral”, and the prices were the interest rates that the winning bidders paid:⁸

⁷ Note that negative dot bids cannot be understood as offers to sell—an offer to sell would be accepted whenever its price is sufficiently low, whilst a negative bid cancels a purchase whenever one of its prices is sufficiently high.

⁸ Although negative bids were offered as an option to the Bank of England in [26], its Product-Mix auctions have not used them. Prior to 2014, bidders could make any set of positive bids, and the auctioneer (the Bank of England) expressed its own preferences using a supply function that was equivalent to using any set of positive bids (see Appendix E1 of [28]). Since 2014, the auctions run by the Bank have allowed the

Example 2 A bidder might be interested in \$100 million of weak collateral (good 1) at up to a 7% interest rate, and \$80 million of strong collateral (good 2) at up to a 5% interest rate. However, even if prices are high, the bidder wants an absolute minimum of \$40 million (see Fig. 2). These preferences can be implemented by making all of the following four bids:

- I \$100 million of weak at 7%.
- II \$80 million of strong at 5%.
- III \$40 million of { weak at maximum permitted bid OR strong at maximum permitted bid less 2% }.
- IV *minus* \$40 million of { weak at 7% OR strong at 5% }.

Note that the bids lead to an arrangement of hyperplanes, at each of which the agent is indifferent among more than one bundle. Bids (I) and (II) together generate the demand shown in the three quadrants to the left of $(7, 0)$ and/or below $(0, 5)$, but would on their own imply zero demand in the top right quadrant. Adding the high positive bid, (III), at $(k, k - 2)$, in which k is the maximum permitted bid on either good (we assume k is large), would add demand of \$40 million of weak everywhere above the 45 deg diagonal line through $(2, 0)$, and \$40 million of strong everywhere below this line; the negative bid, (IV), at $(7, 5)$ then cancels this bid everywhere to the left of, and below, $(7, 5)$.

Preferences of the kind illustrated in Example 2 are very natural for a liquidity-constrained bidder, but cannot be accurately represented without the use of a negative bid.⁹ However, with positive and negative bids, valid bids in the bidding language can precisely represent any “strong substitutes” preferences.¹⁰ Moreover, the way in which positive and negative bids define demand sets has a nice geometric interpretation as Minkowski differences, as we will show. And, importantly, as we discuss below, in practical settings expressing valuations with SSPMA bids is likely to be *much* more compact than listing valuations explicitly as assumed in [13] or subsequent literature. For all these reasons, the SSPMA is a natural choice for applications.

To make practical use of SSPMAs, however, requires that we can find competitive equilibrium prices among participants using the bid language.¹¹ That is, given the

auctioneer to use richer preferences than this, but have restricted to bidders to sets of bids “on the axes” (that is, to sets of bids each of which has $b_i > 0$ for only one i).

⁹ One way to understand a negative bid for a unit is that it is the highest price at which you would cancel a bid for one unit. Reducing your purchases only at *low* prices makes no sense on its own. However, in two dimensions, for example, it *does* make sense in conjunction with a positive bid north-east of the negative bid which gives higher prices at which you would buy (and that the negative bid therefore cancels when prices are low) and *also* other bids to the west and south of it, at least one of which is accepted when the cancellation operates (and without which there would be no reason for the cancellation).

¹⁰ [27] stated this result for the case of multiple units of each of two goods. [7] and [9] show the general result, and also show that any preferences represented by this language that are valid (i.e., the demand for a good cannot decrease if its price falls while no other price changes—see discussion below Definition 2) must be strong substitutes.

¹¹ Bidders in a Product-Mix auction simultaneously make sets of bids that express their preferences. The auctioneer then chooses the aggregate supply and allocates each participant its competitive-equilibrium allocation at competitive-equilibrium prices, assuming that all the expressed preferences are accurate. Ties between bids can be broken arbitrarily, since participants who express their preferences accurately are

collection of the sets of bids made by all the participants, we need to be able to find a price vector at which any given quantity vector of goods would be exactly demanded if all the bids expressed participants' actual preferences.

If all the bids are positive, the competitive equilibrium price vectors are just the shadow price vectors in the solution to a simple linear program, more specifically a network flow problem, in which the number of variables is linear in the number of bids and distinct goods. The reason is that competitive equilibrium maximizes social surplus in our setting, so the relevant linear program allocates the bids among participants to maximize the sum of their surpluses, subject to allocating exactly the available supply.¹² With negative bids, however, the allocation problem cannot be modeled with only a single linear program, and the computation of prices is then more challenging.

1.2 Our contribution

We study characteristics of strong substitutes by using the SSPMA language. First, we show that the positive and negative bids in the SSPMA allow us to interpret strong substitutes as Minkowski differences between sets that are easy to identify. This gives new insight into the geometric structure of strong substitutes, a valuation class that is difficult to characterize. We then illustrate the SSPMA language's expressiveness using [40]'s notorious example of strong substitutes that other languages cannot represent. We also explain that the language is compact for realistic settings, since the bidder need not explicitly give a value for every bundle which it might be allocated.

Our main contribution is an equivalence result for different mathematical formulations of the pricing problem. We show that minimizing the difference between the maximum social surpluses attained by solving certain pairs of allocation problems—each of which is a simple problem—provides the information we need to compute the equilibrium prices. Specifically, the correct quantity of negative bids, s , accepted by the auctioneer minimizes the difference between the objective function of the linear program that would be solved to allocate the available supply increased by s if only the positive bids were available (we call this the “positive program”), and the objective of a corresponding linear program that would be solved to allocate a quantity of s using only the negative bids (the “negative program”). Moreover, the competitive equilibrium price vectors are the shadow price vectors for the positive program for this value of s .¹³ We prove these results by showing that minimizing the difference between the positive and negative programs is dual to minimizing a Lyapunov function $L(\mathbf{p})$. More precisely, we show that the Toland-Singer dual [33, 47] of $L(\mathbf{p})$ is the minimum difference between the positive and negative linear programs.

indifferent. If there are multiple competitive equilibria, the Bank of England's Product-Mix auctions choose the best one for bidders (this is uniquely defined—see discussion below Definition 2). See [26, 27] for more details.

¹² This is the solution method currently used by the Bank of England's Product-Mix program, which does not allow bidders to use negative bids.

¹³ These shadow price vectors are a subset (often a proper subset) of the shadow price vectors for the negative program for this s .

[6] have recently shown that a standard steepest-descent algorithm based on the Lyapunov function (following [38]) can solve the SSPMA pricing problem, but their method takes only limited advantage of the special features of the geometric representation.¹⁴ By taking fuller advantage of the structure of strong substitutes analyzed in this paper, we find an alternative to steepest descent on the Lyapunov function. Our algorithm draws on DC (difference of convex functions) programming.¹⁵ Steepest descent algorithms on the Lyapunov function are known to be very efficient. But we find that the DC algorithm is at least similarly fast in all our experiments. Neither algorithm is consistently faster, and in environments with only a low number of negative bids (which we conjecture are the most likely ones in practice—see Section 2.3), the DC algorithm is the faster one. So, while both algorithms terminate in a few seconds even for large problem instances, the DC algorithm provides a valuable new alternative by taking advantage of the structural properties of strong substitutes.

1.3 Outline

We proceed as follows. Section 2 introduces the SSPMA bidding language. We illustrate its expressiveness, and explain that it is a compact language that expresses all strong substitutes valuations (and no others) as the Minkowski difference of positive and negative bids. Section 3 proves that the pricing problem can be solved by minimizing the difference between the objectives of the two linear programs, by showing that this is dual to minimizing the Lyapunov function. Section 4 takes advantage of this result to use “DC programming” (difference of convex functions programming) to specify an algorithm to solve the problem, and uses numerical experiments to compare our algorithm to a steepest-descent algorithm based on the Lyapunov function. Section 5 concludes. All proofs are in the Appendix.

2 The SSPMA bid language

2.1 Formal description of the SSPMA language

In the SSPMA, each of m bidders $j \in \{1, \dots, m\}$ submit an arbitrary number of bids for distinct goods $i \in \{1, \dots, n\}$. A bid is a vector $\mathbf{b} = (b_1, \dots, b_n; b_{n+1}) \in \mathbb{Z}_{\geq 0}^n \times (\mathbb{Z} \setminus \{0\})$. Here, for $i = 1, \dots, n$, coordinate b_i gives the value for good i . The final coordinate $b_{n+1} \in \mathbb{Z} \setminus \{0\}$ is the *weight* of the bid; we write $w(\mathbf{b})$ for the projection to this final coordinate. We refer to *positive* and *negative* bids according to the sign of $w(\mathbf{b})$. Prices $\mathbf{p} = (p_1, \dots, p_n) \in \mathbb{R}^n$ are linear. Our bundles of indivisible goods will be written $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_{\geq 0}^n$. We write \mathbf{e}^i for the coordinate vectors $i = 1, \dots, n$ in \mathbb{Z}^n .

¹⁴ Unlike [6] we focus on the structural properties of strong substitutes that arise from the SSPMA bid language as well as the economic interpretation of the Toland-Singer dual of the widely used Lyapunov function.

¹⁵ Minimizing the difference between two M^{\natural} -convex functions is in general *NP*-hard [29, 32]: the difference between the positive and negative programs is neither convex nor concave. However, this specific problem can be solved in polynomial time, as is clear from the relationship to the Lyapunov function.

A positive bid \mathbf{b} expresses the willingness of the bidder to pay at most b_i for units of good $i = 1, \dots, n$, and for up to $w(\mathbf{b})$ units in total. It defines a *valuation* $v_{\mathbf{b}}$ on the domain $\Delta_{w(\mathbf{b})}$ of bundles of at most $w(\mathbf{b})$ units, that is, $\Delta_{w(\mathbf{b})} = \{\mathbf{y} \in \mathbb{Z}_{\geq 0}^n : \sum_{i=1}^n y_i \leq w(\mathbf{b})\}$, with $v_{\mathbf{b}}(\mathbf{y}) = \sum_{i=1}^n b_i y_i$. The utility associated with this bid is quasi-linear, $v_{\mathbf{b}}(\mathbf{y}) - \langle \mathbf{p}, \mathbf{y} \rangle$, so the *indirect utility* associated with such a bid is just

$$u_{\mathbf{b}}(\mathbf{p}) = w(\mathbf{b}) \max_{i \in \{1, \dots, n\}} (b_i - p_i, 0), \quad (1)$$

where we include 0 because the bid may instead be rejected. Any combination of $w(\mathbf{b})$ units of utility-maximizing goods is acceptable, as are fewer units when $u_{\mathbf{b}}(\mathbf{p}) = 0$, so the *demand set* is

$$D_{\mathbf{b}}(\mathbf{p}) := \left\{ \mathbf{y} \in \Delta_{w(\mathbf{b})} : \sum_{i=1}^n y_i (b_i - p_i) = u_{\mathbf{b}}(\mathbf{p}) \right\}. \quad (2)$$

This set comprises all integer bundles in the convex polytope in which the bundles $w(\mathbf{b})\mathbf{e}^i$, where i maximizes $b_i - p_i \geq 0$, are vertices, and $\mathbf{0}$ is also a vertex if $\max_{i \in \{1, \dots, n\}} (b_i - p_i, 0) = 0$. If $D_{\mathbf{b}}(\mathbf{p})$ contains more than one bundle, we say all goods $i = 1, \dots, n$ maximizing $b_i - p_i$ are *marginal* for bid \mathbf{b} at \mathbf{p} . If $\{\mathbf{0}\} \subsetneq D_{\mathbf{b}}(\mathbf{p})$ then we say the bid is *marginal to be accepted*. If $D_{\mathbf{b}}(\mathbf{p}) = \{\mathbf{0}\}$ we say the bid is *rejected*.

Now consider a multiset \mathcal{B} of positive bids, which could be all those placed by a single bidder, or could, for example, be all bids from *all* bidders. The aggregate indirect utility $u_{\mathcal{B}}(\mathbf{p})$ is just the sum of indirect utilities: $u_{\mathcal{B}}(\mathbf{p}) = \sum_{\mathbf{b} \in \mathcal{B}} u_{\mathbf{b}}(\mathbf{p})$, and the aggregate demand set $D_{\mathcal{B}}(\mathbf{p})$ is the Minkowski sum of demand sets $D_{\mathcal{B}}(\mathbf{p}) = \sum_{\mathbf{b} \in \mathcal{B}} D_{\mathbf{b}}(\mathbf{p})$.

However, we also allow negative bids: those for which $w(\mathbf{b}) < 0$. These do not represent a meaningful economic valuation on their own, but do so in “valid” combinations with positive bids. Given a collection \mathcal{B} of bids, write respectively \mathcal{B}_+ and \mathcal{B}_- for the positive and negative bids in \mathcal{B} . Write $|\mathbf{b}|$ for the bid $(b_1, \dots, b_n; |w(\mathbf{b})|)$, and write $|\mathcal{B}_-|$ for the set of bids $|\mathbf{b}|$ where $\mathbf{b} \in \mathcal{B}_-$. Now the *aggregate indirect utility* is an appropriately signed sum of indirect utilities:

$$u_{\mathcal{B}}(\mathbf{p}) := \sum_{\mathbf{b} \in \mathcal{B}_+} u_{\mathbf{b}}(\mathbf{p}) - \sum_{\mathbf{b} \in |\mathcal{B}_-|} u_{\mathbf{b}}(\mathbf{p}). \quad (3)$$

We say that the set \mathcal{B} is *valid* when the indirect utility $u_{\mathcal{B}}$ is concave. (See Theorem 1 of [6]; further discussion of this notion is given below after Definition 2.)

To define the *aggregate demand set* with positive and negative bids, first define the demand $D_{\mathbf{b}}(\mathbf{p})$ associated with an individual negative bid \mathbf{b} as $D_{\mathbf{b}}(\mathbf{p}) = -D_{|\mathbf{b}|}(\mathbf{p}) = \{-\mathbf{x} \mid \mathbf{x} \in D_{|\mathbf{b}|}(\mathbf{p})\}$. Let \mathcal{Q} comprise all price vectors \mathbf{q} in a small neighborhood of \mathbf{p} , and such that $D_{\mathbf{b}}(\mathbf{q}) = \{\mathbf{x}_{\mathbf{b}}(\mathbf{q})\}$ are singletons for all $\mathbf{b} \in \mathcal{B}$. Then the aggregate

demand set is equal to the discrete convex hull

$$D_{\mathcal{B}}(\mathbf{p}) = \text{conv} \left\{ \sum_{\mathbf{b} \in \mathcal{B}} D_{\mathbf{b}}(\mathbf{q}) : \mathbf{q} \in \mathcal{Q} \right\} \cap \mathbb{Z}^n.$$

In particular, if $D_{\mathbf{b}}(\mathbf{p})$ is a singleton for all $\mathbf{b} \in \mathcal{B}$, then $D_{\mathcal{B}}(\mathbf{p})$ is just $\sum_{\mathbf{b} \in \mathcal{B}} D_{\mathbf{b}}(\mathbf{p}) = \sum_{\mathbf{b} \in \mathcal{B}_+} D_{\mathbf{b}}(\mathbf{p}) - \sum_{\mathbf{b} \in |\mathcal{B}_-|} D_{\mathbf{b}}(\mathbf{p})$: negative bids are used to “cancel” part of the demand arising from positive bids. We cannot extend this rule to prices at which the demand set is non-unique simply by taking the Minkowski sum of demand sets associated with all bids; negative bids which are marginal between goods must be treated consistently with positive bids marginal on those same goods.¹⁶ However, if the bids \mathcal{B}^j of each bidder $j = 1, \dots, m$ are valid, then the full aggregate demand set $D_{\mathcal{B}}(\mathbf{p})$ defined by $\mathcal{B} = \bigcup_{j=1}^m \mathcal{B}^j$ is indeed the Minkowski sum: $D_{\mathcal{B}}(\mathbf{p}) = \sum_{j=1}^m D_{\mathcal{B}^j}(\mathbf{p})$.

When \mathcal{B} contains only positive bids, we can aggregate the simple valuations implied by individual bids, to obtain the *aggregate valuation* $v_{\mathcal{B}} : \Delta_W \rightarrow \mathbb{Z}$, where $W = \sum_{\mathbf{b} \in \mathcal{B}} w(\mathbf{b})$:

$$v_{\mathcal{B}}(\mathbf{y}) = \max \left\{ \sum_{\mathbf{b} \in \mathcal{B}} \sum_{i=1}^n x_{i\mathbf{b}} b_i : \sum_{\mathbf{b} \in \mathcal{B}} x_{i\mathbf{b}} \leq y_i \forall i \text{ and } \sum_{i=1}^n x_{i\mathbf{b}} \leq w(\mathbf{b}) \forall \mathbf{b} \in \mathcal{B} \right\}. \quad (4)$$

As usual, the relations $u_{\mathcal{B}}(\mathbf{p}) = \max_{\mathbf{x} \in \mathbb{Z}_{\geq 0}^n} v_{\mathcal{B}}(\mathbf{x}) - \langle \mathbf{p}, \mathbf{x} \rangle$ and $v_{\mathcal{B}}(\mathbf{x}) = \min_{\mathbf{p} \in \mathbb{R}^n} u_{\mathcal{B}}(\mathbf{p}) + \langle \mathbf{p}, \mathbf{x} \rangle$ hold. The latter equation also gives us an indirect way to identify the aggregate valuation if \mathcal{B} is a valid set of positive and negative bids. However, one of our main results, which is also the starting point to our equilibrium pricing algorithm, is a purely primal expression for the aggregate valuation in the presence of negative bids (Theorem 1).

The valuation implied by such bids is for strong substitutes:

Definition 1 (Ordinary and strong substitutes, [8] and [35]) A valuation v is *ordinary substitutes*, if for any price vectors $\mathbf{p}' \geq \mathbf{p}$ with singleton demand sets $D_v(\mathbf{p}') = \{\mathbf{x}'\}$ and $D_v(\mathbf{p}) = \{\mathbf{x}\}$, we have $\mathbf{x}'_k \geq \mathbf{x}_k$ for all k with $p'_k = p_k$. A valuation v is *strong substitutes*, if, when we consider every unit of every good to be a separate good, v is ordinary substitutes.

The SSPMA only expresses preferences of this kind, and can express *any* strong substitutes valuation [7, 9]¹⁷. It is, to our knowledge, the only bidding language that provably has this feature.

In an SSPMA such as the Bank of England’s, total supply is not pre-determined; the auctioneer represents its preferences as supply schedules, and the auction finds competitive equilibrium given the auctioneer’s and bidders’ expressed preferences.

¹⁶ For example, (140, 40) is *not* in the demand set at $\mathbf{p} = (3, 1)$ in the right-hand side of Figure 1; the bids for -40 and 40 units must be treated consistently.

¹⁷ [27] stated this result for the case of multiple units of two goods. [31] describes how any valuation can be analyzed tropical-geometrically and can be decomposed into a combination of simpler pieces, but if the valuation is not strong substitutes, these simpler pieces do not correspond to positive and negative bids.

However, the auctioneer can equivalently auction the maximum quantity of each good that it would ever sell at any price vector, and place appropriate bids to buy back quantities at lower prices.¹⁸ So we index the auctioneer as agent 0 and include its bids in the set \mathcal{B} of all bids from all bidders. This paper therefore addresses the following problem:

Definition 2 (Equilibrium pricing problem) Given a valid set \mathcal{B} of all bids from all bidders (including the auctioneer) and a *target supply* \mathbf{t} , find a price vector $\mathbf{p} \in \mathbb{R}^n$, such that \mathbf{t} is demanded at \mathbf{p} , that is, $\mathbf{t} \in D_{\mathcal{B}}(\mathbf{p})$. Such a price vector is called an equilibrium price.¹⁹

It is well-known that a competitive equilibrium does indeed exist, given our assumptions of strong substitutes and a seller who will retain units of any underdemanded good at a price of zero [17, 35]. Indeed, this also implies that equilibrium price in $\mathbb{R}_{\geq 0}^n$ exists. The set of equilibrium prices forms a lattice with respect to the Euclidean ordering [21, 36], i.e., for any valuations v_1, \dots, v_m , if \mathbf{p} and \mathbf{p}' are equilibrium prices for such valuations, then $\mathbf{p} \wedge \mathbf{p}'$ and $\mathbf{p} \vee \mathbf{p}'$ are also equilibrium prices. This implies that there exists a unique minimal equilibrium price vector. It is possible to modify the algorithm we will develop to find the minimal equilibrium price vector rather than an arbitrary price vector.

To understand the “validity” of bids in the SSPMA, we briefly outline some geometric ideas from [8]. First, the collection \mathcal{B} of bids induces a set of prices at which the aggregate demand is not unique: the “locus of indifference prices” (LIP), notated $\mathcal{L}_{\mathcal{B}} := \{\mathbf{p} : |D_{\mathcal{B}}(\mathbf{p})| > 1\}$. For a price \mathbf{p} to be in the LIP, at least one bid must be marginal, so some equality of the form $b_i = p_i$ or $b_i - p_i = b_j - p_j$ must hold, where $i, j \in \{1, \dots, n\}$ and $j \neq i$. Therefore, $\mathcal{L}_{\mathcal{B}}$ consists of a union of pieces of hyperplanes with normals in $\{\mathbf{e}^i, \mathbf{e}^i - \mathbf{e}^j : 1 \leq i < j \leq n\}$. These pieces of hyperplanes are known as *facets*. To each facet F , we assign a *weight* $w(F)$, given by the sum of the weights of bids that are marginal at a price in the relative interior of F . Facets always have nonzero weight; if the sum of weights of marginal bids is zero then one may see that demand is in fact unique.

The LIP $\mathcal{L}_{\mathcal{B}}$ splits price space into multiple *unique demand regions* (UDRs) at which a unique bundle is demanded. Let \mathbf{p} be a price vector in an UDR for which the demand is known (for example, for \mathbf{p} large, the demand is 0). If the price \mathbf{p} changes along a curve, and crosses a facet F of $\mathcal{L}_{\mathcal{B}}$, then the demand changes by $w(F)\mathbf{n}$, where \mathbf{n} is the normal of F pointing into the opposite direction of the path. For an illustration, see Fig. 3. Thus, the LIP fully encodes the aggregate demand at every UDR-price, and so – by taking convex hulls – at every price.

Now, a negative-weighted facet cannot arise from a quasi-linear preference relation: when the price of one good decreases, the demand for that good must not also decrease.

¹⁸ [28] (Appendix E1) illustrates how the auctioneer can do this for general supply schedules; for the special case in which it just wishes to sell a bundle $\mathbf{t} = (t_1, \dots, t_n)$ at any non-negative prices, it will simply auction supply \mathbf{t} and enter a bid $(0, \dots, 0; \sum_i t_i)$ into the auction.

¹⁹ This paper takes competitive behavior as given; we do not address the extent to which bidders may distort their preferences. In an SSPMA it is rational for bidders whose demand is not too large relative to aggregate demand to make bids that approximately reflect their true preferences. Throughout this article, we assume that bids reflect true preferences.




Fig. 3 Finding demand in each UDR of a LIP. The black circles represent positive bids with weight 1, namely $(2, 2; 1)$, $(1, 0; 1)$ and $(0, 1; 1)$; the white circle represents a negative bid, $(1, 1; -1)$. Note that all facets emanating from this negative bid are canceled by parts of facets arising from positive bids. A curve which determines demand in every UDR is shown as a dashed line. The curve starts at a high price, where the demand is $(0, 0)$. The vectors where the path intersects the LIP indicate the correctly oriented normals of the facets with respect to the path. For example, inspecting the crossings of facets reveals that the demand at $(0.5, 0.5)$ is $1 \cdot (1, 0) + 1 \cdot (-1, 1) + 1 \cdot (1, 0) = (1, 1)$

So negative bids must be placed in such a way that, in the resulting collection of facets, no facet has a negative weight. This condition is equivalent to concavity of the indirect utility function.²⁰ From now on we assume that *our bid collections are always valid*. Note that if each individual bidder's bid set is valid, then so is the set of all bidders.

The following two subsections discuss geometrical interpretations and properties of this bid language. While they do not contribute directly to our main Theorem 1 and the algorithm, they provide useful background on the role of negative bids in the SSPMA and intuition for the overall approach.

2.2 Interpretation via Minkowski differences

There appears to be a contrast between the intuitive definition of the aggregate demand set when all bids are positive (so the aggregate demand set is just the Minkowski sum of the individual demand sets) and the more involved definition when negative bids are present. Recall from Sect. 2.1 that in this case we defined the aggregate demand set to be the discrete convex hull of bundles which are demanded uniquely when we slightly change the price vector \mathbf{p} . We cannot simply take Minkowski sums because we must ensure that negative bids are treated in a valid way with their associated positive bids (see the discussion after Definition 2). However, if \mathcal{B} is valid, then we can provide a more parsimonious novel definition by using the Minkowski difference operation. We recall that $A - B$ consists of all points $\mathbf{x} \in \mathbb{R}^n$, such that $\mathbf{x} + B \subseteq A$.

²⁰ See [6] Theorem 1, which shows that, at any price, the sum of the weights of bids marginal between any pair of goods, or between any good and being rejected, must be non-negative. The failure of this condition is equivalent to the existence of a negative-weighted facet.




Fig. 4 The Minkowski sum $A + (-B)$ and Minkowski difference $A - B$ of rectangle A and a triangle B . Left: the rectangle A (in gray); four instances of $a + (-B)$, in which $a \in A$ and $-B$ is the triangle (dashed line and its interior); and $A + (-B)$ (black line, including its interior). Right: the same rectangle A (in gray); six instances of $a + B$. For five of these, such as $a = x$, we have $a + B \subseteq A$ and so $a \in A - B$, but $y + B \not\subseteq A$ and so $y \notin A - B$. The full set $A - B$ is given by the black line and its interior

The geometric effect of this operation is illustrated in Fig. 4. Note in particular that in general $A + (-B) \neq A - B$.

Proposition 1 *Let \mathcal{B} be a valid collection of bids in an SSPMA. Then for every price vector \mathbf{p} the demand set $D_{\mathcal{B}}(\mathbf{p})$ is equal to $D_{\mathcal{B}_+}(\mathbf{p}) - D_{|\mathcal{B}_-|}(\mathbf{p})$.*

We prove this result in Appendix A.1.

2.3 Expressiveness and compactness of the SSPMA

To illustrate the expressive power of negative bids, we consider [40]’s notorious example of a valuation, v_r , that shows that prior bid languages such as the endowed assignment valuations by [22] are strictly less expressive than the set of gross substitutes (and so also strong substitutes) valuations.²¹ (We discuss the construction of v_r in Appendix A.2.) However,

Proposition 2 *The valuation v_r in [40] can be represented by 8 positive and 6 negative SSPMA bids.*

This proposition illustrates [9]’s more general result that *all* strong substitutes valuations can be depicted in the SSPMA. We prove Proposition 2 without relying on this general result by explicitly providing the list of SSPMA bids.

Moreover, an important feature of the SSPMA language is that it is parsimonious: the valuations that are most used in practice can be expressed very simply, using far fewer bids than the number of different bundles valued. Let $W := \sum_{\mathbf{b} \in \mathcal{B}} w(\mathbf{b})$. Note that W equals the maximum number of units that a bidder who makes bids $\mathbf{b} \in \mathcal{B}$ is interested in. Then SSPMA bids can assign a “non-trivial” value to $\Omega(W^n)$ bundles:

Proposition 3 *Consider an SSPMA with n goods, and suppose a bidder makes bids \mathcal{B} . Let $D := \bigcup \{D_{\mathcal{B}}(\mathbf{p}) : \mathbf{p} \in \mathbb{R}^n\}$ and let $W := \sum_{\mathbf{b} \in \mathcal{B}} w(\mathbf{b})$. Then $D = \Delta_W$ and so $|D| = \binom{n+W}{n} \geq (1 + W/n)^n$.*

²¹ [18] uses the same example to show [34]’s (integer) assignment messages cannot express all strong substitute valuations.

Moreover, bidders in practical applications are likely to need to make far fewer bids than Proposition 3 suggests: Expressing a demand function for each good independently is trivial—it just requires providing a separate list of bids for each i with, for each i , $b_j = 0$ for all $j \neq i$. In many settings these bids will express much of the information about bidders' valuations.

At a second, higher, level of complexity, any bid which selects the “best value” among any number of goods can be expressed using only positive bids. Observe that W is the maximum number of bids that a bidder who is interested in winning at most W units, and who uses only positive bids, needs to make—and if any of her bids have greater weight than 1, she will need fewer bids. So such a bidder can express her valuations of all possible bundles with only a few bids.

More complex features of preferences require negative bids to express, but these features seem less likely to arise frequently. Example 2 is one example, and there are others,²² but we expect most bidders would be unlikely to have to handle more than a very small number of these special issues. In fact, in the Bank of England's auctions, bidders showed relatively little interest even in bids of the “second level” of complexity, and they used such bids only rarely—perhaps because they are only very important to banks in times of real crisis.²³ So bidders are unlikely to need many negative bids in most practical auction settings.²⁴ The number of bids needed by a bidder who is interested in winning at most W units, and who needs only a small number of negative bids, cannot much exceed W . Moreover, such a bidder will need to use many fewer than W bids unless most or all of her bids are of weight only 1. So these bidders, too, are likely to be able to express their full valuations with only a few bids.

The number of bundles valued by a set of bids of mixed sign can be much smaller than the lower bound on the number of bundles valued by the same number of bids that are all positive.²⁵ But we expect the SSPMA bidding language will be much more “compact”²⁶ in most practical cases than—say—listing valuations for all bundles explicitly.

²² For example, choosing the best N_1 out of N_2 options requires the use of negative bids if $N_2 > N_1 > 1$.

²³ The need for “second-level” bids is likely to grow, however, as technology develops—they are most useful for banks who can coordinate different parts of their operations in a sophisticated way, and “big investment programmes are already underway in many [banks], to ensure that [they] have real-time information on the collateral they have available globally across all their business lines, that the collateral they deliver is cost effective, and that the cost of delivering (and financing) that collateral is factored into their risk and business decisions. These programmes involve sometimes relatively advanced technology; indeed, as some of our contacts remark, somewhat alarmed, ‘for the first time in living memory, pointy heads are sitting on the repo desk’.” (Andrew Hauser [Executive Director of the Bank of England], 2013) [23].

²⁴ Moreover, [28] shows how to enhance the SSPMA with additional “words”, each of which refers to a particular configuration of positive and negative bids. This can greatly reduce the number of bids required to express special situations. For example, our Example 2 could be expressed by a single “word” from a parameterised class of words. Preferences of the kind described in note 22 could also be expressed as “words”.

²⁵ Of course, no language can express every possible valuation using fewer pieces of information than the number of bundles that can be independently valued. However, in extreme cases the number of bids required to express a full valuation for up to W units in the SSPMA can exceed the number of different possible bundles of up to W units.

²⁶ See [19] for a discussion of compactness.

So in the settings that we believe are most likely to arise, bidders can probably build up their valuations bit by bit through adding individual bids. We expect most of these bids will be positive; moreover, software can make entering negative bids easier, by checking feasibility and by allowing the bidding of “words,”²⁷ and by checking feasibility.²⁸

At least with human participants, it seems less likely that a valuation function with all exponentially many package values would be available, but in this case the prices can be computed directly, using steepest descent or linear programming algorithms. Alternatively, bidders could use [20]’s algorithm to generate bids from an arbitrary value function.²⁹

3 The SSPMA pricing problem

With only positive bids, our equilibrium pricing problem (Definition 2) can be solved via a simple linear program that maximizes the total welfare given the target bundle \mathbf{t} . We know that $\mathbf{t} \in \Delta_W$, where W is the total weight of bids placed, by our assumption about the bids of the auctioneer. But recall from Equation (4) that, given the collection \mathcal{B} of *all* bids of *all* bidders, the aggregate valuation of any bundle $\mathbf{t} \in \Delta_W$ is given - in LP notation - by

$$\begin{aligned}
 v_{\mathcal{B}}(\mathbf{t}) = \max & \sum_{\mathbf{b} \in \mathcal{B}} \sum_{i \in [n]} b_i x_{\mathbf{b}i} && \text{(LP)} \\
 \text{s.t.} & \sum_{i \in [n]} x_{\mathbf{b}i} \leq w(\mathbf{b}) && \forall \mathbf{b} \in \mathcal{B} && (\pi_{\mathbf{b}}) \\
 & \sum_{\mathbf{b} \in \mathcal{B}} x_{\mathbf{b}i} = t_i && \forall i \in [n] && (p_i) \\
 & x_{\mathbf{b}i} \geq 0 && \forall \mathbf{b} \in \mathcal{B}, i \in [n].
 \end{aligned}$$

Here $\pi_{\mathbf{b}}$ and p_i denote the respective dual variables. This program always has an integral optimal solution, as may be seen either by properties of strong substitutes valuations, or by recognizing that it is an instance of the min-cost flow problem. The number of constraints and variables is polynomial in the number of bids and goods, in contrast to the formulation of [13]. The set of equilibrium prices can be computed directly:

²⁷ See note 24.

²⁸ In more complex cases, we can use the [20] algorithm to generate bids from an arbitrary value function.

²⁹ [20]’s algorithm has linear query complexity for preferences that require positive bids only. (An asymptotic lower bound of $\Omega(B \log M)$ queries are required to learn a list of B positive bids, where M is the magnitude of the bid vectors w.r.t. the L_{∞} -norm.) It has exponential query complexity in the worst case when negative bids are required. (The query complexity of learning bid lists corresponding to strong substitutes demand has a rate of growth of $\Theta(B \log M + B^n)$.) However, if the number of goods is not too large, the algorithm still performs well, even though [31] observe that breaking a general valuation up into constituent simpler parts can be NP-hard.

Proposition 4 *In an SSPMA with only positive bids, the equilibrium prices for the target supply \mathbf{t} are the optimal dual variables $\mathbf{p} = (p_1, \dots, p_n)$ of the network linear program (LP) which can be solved in polynomial time in the number of goods and bids.*

Proposition 4 simply follows from writing down the complementary slackness conditions of (LP), so we do not provide an explicit proof. If \mathcal{B} also contains negative bids, the problem of efficiently computing equilibrium prices is less obvious. One route, taken by [6], is to minimize the *Lyapunov function* $L : \mathbb{R}^n \rightarrow \mathbb{R}$ [4], defined for target \mathbf{t} as

$$L(\mathbf{p}) = u_{\mathcal{B}}(\mathbf{p}) + \langle \mathbf{p}, \mathbf{t} \rangle$$

where aggregate indirect utility $u_{\mathcal{B}}(\mathbf{p})$ is as defined in Eq. (3). The set of minimizers of L coincides with the set of equilibrium prices, and structural properties of L allow for polynomial-time steepest descent algorithms to find these minima [6, 36, 42]. However, this approach works by invoking a rather generic submodular function minimization algorithm, under the assumption that a demand oracle is available.

By contrast, with only positive bids we can build upon much more specialized algorithms to solve network linear programs. And, as we now show, taking advantage of the economic structure of the problem allows us to incorporate negative bids into this approach:

Recall that the total allocation in an SSPMA is equal to that assigned to positive bids *minus* that assigned to negative bids. So, to assign \mathbf{t} units in total, we must assign $\mathbf{t} + \mathbf{s}$ units to positive bids and \mathbf{s} to negative bids, for some “supplementary” bundle \mathbf{s} . Recall also that we write \mathcal{B}_+ for the positive bids in \mathcal{B} , and $|\mathcal{B}_-|$ for the negative bids $\mathbf{b} \in \mathcal{B}$ endowed with weights $|w(\mathbf{b})|$. We introduce two additional SSPMAs: that with bids \mathcal{B}_+ and target $\mathbf{t} + \mathbf{s}$, which we call the “positive auction”; and that with (positive) bids $|\mathcal{B}_-|$ and target \mathbf{s} , which we call the “negative auction”. Write W_+ and W_- for the total weights of bids in these respective auctions, so that Δ_{W_+} and Δ_{W_-} are the sets of bundles that may be sold by each of them. Note that for each $\mathbf{t} \in \Delta_{W_+}$ and $\mathbf{s} \in \Delta_{W_-}$, $\mathbf{t} + \mathbf{s}$ lies in Δ_{W_+} (see Appendix Lemma 4).

If we pick \mathbf{s} correctly, then this is equivalent to allocating \mathbf{t} units in the auction with bids \mathcal{B} . Moreover, since both \mathcal{B}_+ and $|\mathcal{B}_-|$ are sets of positive bids, their respective aggregate valuations and equilibrium prices can be evaluated using the linear program above. We now show how to find \mathbf{s} :

Theorem 1 *If \mathcal{B} represents all bids from all bidders, then the aggregate valuation at the target supply $\mathbf{t} \in \Delta_{W_+}$ can be written as*

$$v_{\mathcal{B}}(\mathbf{t}) = \min_{\mathbf{s} \in \Delta_{W_-}} (v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - v_{|\mathcal{B}_-|}(\mathbf{s})).$$

Moreover, given a minimizer $\bar{\mathbf{s}}$, each equilibrium price $\bar{\mathbf{p}}$ of the auction with bids \mathcal{B}_+ and target supply $\mathbf{t} + \bar{\mathbf{s}}$ is an equilibrium price for the auction with bids $|\mathcal{B}_-|$ and target supply $\bar{\mathbf{s}}$, and also for the complete auction with bids \mathcal{B} and target supply \mathbf{t} .

We prove Theorem 1 in Appendix A.4.

To understand the economic intuition underlying Theorem 1 assume that the set of equilibrium prices is n -dimensional and consider a price \mathbf{p} in its interior. (Although the SSPMA would choose the minimum of the equilibrium prices, choosing an interior price simplifies the intuition.) Let $\bar{\mathbf{s}}$ be the vector of negative bids accepted in the equilibrium. Initially set the target \mathbf{s} of the negative auction to be $\bar{\mathbf{s}}$, which means that \mathbf{p} is an equilibrium price for both the positive and negative auctions.

Consider the effect of changing \mathbf{p} on the weighted sum of bids accepted in these two auctions. Recall that the full set \mathcal{B} of positive and negative bids in the original SSPMA is valid. So for any price at which additional negative bids are marginal to be accepted, positive bids with at least as great a weight must also be marginal to be accepted—see the discussion of validity of \mathcal{B} at the end of Sect. 2.1. (The converse does not hold: positive bids can be marginal at prices at which no negative bid is marginal.) So, any change in price from \mathbf{p} would alter the total weight of bids accepted in the positive auction by weakly more than it would alter the total weight of bids accepted in the negative auction.

Now consider an increase in one coordinate of the supplementary bundle, from $\bar{\mathbf{s}}$ to $\mathbf{s} \geq \bar{\mathbf{s}}$, in both the positive and negative auctions. The additional bids that will be accepted in the positive auction with target $\mathbf{t} + \mathbf{s}$ will, because of our observation above, have weakly greater value than the additional bids accepted in the negative auction. That is, $v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - v_{\mathcal{B}_+}(\mathbf{t} + \bar{\mathbf{s}}) \geq v_{|\mathcal{B}_-|}(\mathbf{s}) - v_{|\mathcal{B}_-|}(\bar{\mathbf{s}})$. Similarly, if we decrease one coordinate to $\mathbf{s} \leq \bar{\mathbf{s}}$, then bids which are now rejected from the positive auction will have weakly lower value than the bids rejected from the negative auction. So, again, $v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - v_{\mathcal{B}_+}(\mathbf{t} + \bar{\mathbf{s}}) \geq v_{|\mathcal{B}_-|}(\mathbf{s}) - v_{|\mathcal{B}_-|}(\bar{\mathbf{s}})$. General changes in \mathbf{s} may be understood as a sequence of these two operations.

It follows that $\bar{\mathbf{s}}$ can be identified by minimizing $v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - v_{|\mathcal{B}_-|}(\mathbf{s})$.

The formal proof of Theorem 1 rests on applying a version of Toland-Singer duality [47] to the valuations in the positive and negative auctions, and relating this to the Lyapunov function $L(\mathbf{p})$. [33] provide a theoretical treatment of discrete DC-functions, establishing (their Theorem 4.6) Toland-Singer duality in discrete DC-functions.

First recall that, for a function $f : \text{dom } f \rightarrow \mathbb{R}$, where $\text{dom } f \subseteq \mathbb{R}^n$, the convex conjugate $f^* : \text{dom } f^* \rightarrow \mathbb{R}$ is defined by $f^*(\mathbf{p}) = \sup_{\mathbf{x} \in \text{dom } f} (\langle \mathbf{p}, \mathbf{x} \rangle - f(\mathbf{x}))$, where $\text{dom } f^* \subseteq \mathbb{R}^n$ comprises those \mathbf{p} at which $f^*(\mathbf{p})$ is finite-valued. The *subdifferential* of f is the set-valued function

$$\partial f(\mathbf{x}) = \{\mathbf{p} \in \mathbb{R}^n : \langle \mathbf{p}, \mathbf{y} \rangle - f(\mathbf{y}) \leq \langle \mathbf{p}, \mathbf{x} \rangle - f(\mathbf{x}) \quad \forall \mathbf{y} \in \mathbb{R}^n\}.$$

The *domain* $\text{dom } \partial f$ of the subdifferential consists of all points $\mathbf{x} \in \text{dom } f$ with $\partial f(\mathbf{x}) \neq \emptyset$. It turns out that in our application the convex conjugates and subdifferentials have an intuitive economic meaning.

Lemma 1 *Let \mathcal{B} be a collection of positive bids. Then $-v_{\mathcal{B}}$ can be naturally extended to a convex function $f : \text{dom } f \rightarrow \mathbb{R}$ with the following properties:*

1. $\text{dom } \partial f = \text{dom } f = \text{conv } \Delta_W$ and $\text{dom } \partial f^* = \text{dom } f^* = \mathbb{R}^n$
2. $f^*(\mathbf{q}) = u_{\mathcal{B}}(-\mathbf{q})$ and $\partial f^*(\mathbf{q}) = \text{conv } D_{\mathcal{B}}(-\mathbf{q})$
3. $\partial f(\mathbf{x}) = -\{\mathbf{p} \in \mathbb{R}^n : \mathbf{x} \in \text{conv } D_{\mathcal{B}}(\mathbf{p})\}$.

We will use the following version of Toland-Singer duality, which allows for restricted domains:

Theorem 2 (Toland-Singer duality) *Let $f : \text{dom } f \rightarrow \mathbb{R}$ and $g : \text{dom } g \rightarrow \mathbb{R}$ be proper convex lower semi-continuous functions with closed and convex domains $\text{dom } f \subseteq \text{dom } g \subseteq \mathbb{R}^n$ and such that $\text{dom } g^* \subseteq \text{dom } f^* \subseteq \mathbb{R}^n$. If one of the differences $f(\mathbf{x}) - g(\mathbf{x})$ and $g^*(\mathbf{y}) - f^*(\mathbf{y})$ has a minimum in $\text{dom } f$, respectively $\text{dom } g^*$, the other difference also has one, and*

$$\min_{\mathbf{x} \in \text{dom } f} f(\mathbf{x}) - g(\mathbf{x}) = \min_{\mathbf{y} \in \text{dom } g^*} g^*(\mathbf{y}) - f^*(\mathbf{y}).$$

Moreover, if $\bar{\mathbf{x}}$ minimizes $f(\mathbf{x}) - g(\mathbf{x})$, then any $\bar{\mathbf{y}} \in \partial g(\bar{\mathbf{x}})$ minimizes $g^(\mathbf{y}) - f^*(\mathbf{y})$. Conversely, for any minimizer $\bar{\mathbf{y}}$ of $g^*(\mathbf{y}) - f^*(\mathbf{y})$, any $\bar{\mathbf{x}} \in \partial f^*(\bar{\mathbf{y}})$ minimizes $f(\mathbf{x}) - g(\mathbf{x})$.*

For a proof see [46, Theorem 1]. We will apply Theorem 2 to the convex extensions of $-v_{|\mathcal{B}_-|}$ and $-v_{\mathcal{B}_+}$ (to the convex hulls of their domains).

4 The pricing algorithm

Using Theorem 1, we can approach the pricing problem by minimizing the difference $v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - v_{|\mathcal{B}_-|}(\mathbf{s})$. While the valuations $v_{\mathcal{B}_+}$ and $v_{|\mathcal{B}_-|}$ can be extended to concave functions, and can efficiently be evaluated with linear programs at any given pair of bundles, their difference is in general neither concave nor convex. Moreover, as recently shown by [29], minimizing the difference between two M^{\natural} -convex functions is an NP-hard optimization problem. However, there is a class of algorithms on the difference of convex functions (DC algorithms; see [2, 46]), that find at least local minima of such problems and are often very fast in practice.

4.1 A DC auction algorithm

By Theorem 1, we seek $\bar{\mathbf{s}}$ minimizing $v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - v_{|\mathcal{B}_-|}(\mathbf{s})$. We will approach this by minimizing $f(\mathbf{s}) - g(\mathbf{s})$, where $f(\mathbf{s})$ and $g(\mathbf{s})$ are the convex extensions of $-v_{|\mathcal{B}_-|}(\mathbf{s})$, respectively $-v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s})$, to the convex hulls of their domains. A *necessary* condition for such $\bar{\mathbf{s}}$ is that it gives a stationary point, that is, $\bar{\mathbf{s}} \in \text{dom } \partial f$ with $\partial f(\bar{\mathbf{s}}) \cap \partial g(\bar{\mathbf{s}}) \neq \emptyset$. To interpret this in our context, if $\mathbf{q} \in \partial f(\bar{\mathbf{s}}) \cap \partial g(\bar{\mathbf{s}})$ then $\mathbf{p} = -\mathbf{q}$ is a price at which $\mathbf{t} + \bar{\mathbf{s}}$ is demanded in the positive auction, and $\bar{\mathbf{s}}$ is demanded in the negative auction (see Lemma 1).

The DC Algorithm 1 finds a *stationary point* for two convex functions $f : \text{dom } f \rightarrow \mathbb{R}$ and $g : \text{dom } g \rightarrow \mathbb{R}$ with $\text{dom } \partial f \subseteq \text{dom } \partial g$ and $\text{dom } \partial g^* \subseteq \text{dom } \partial f^*$ [46]. Our functions f and g , defined above, satisfy these conditions: By Appendix Lemma 4, $\text{dom } f = \text{conv } \Delta_{W_-} \subseteq \text{conv}\{\mathbf{s} \in \mathbb{Z}^n : \mathbf{t} + \mathbf{s} \in \Delta_{W_+}\} = \text{dom } g$, $\text{dom } g^* = \mathbb{R}^n = \text{dom } f^*$, and by Lemma 1 the domains of the respective functions coincide with the domains of their subdifferentials.

Algorithm 1: A DC-algorithm

Input: Convex functions $f : \text{dom } f \rightarrow \mathbb{R}$, $g : \text{dom } g \rightarrow \mathbb{R}$ with $\text{dom } \partial f \subseteq \text{dom } \partial g$ and $\text{dom } \partial g^* \subseteq \text{dom } \partial f^*$

Output: Stationary points $\bar{s} \in \mathbb{R}^n$ of $f - g$ and \bar{q} of $g^* - f^*$

- 1: Choose an initial $q^0 \in \mathbb{R}^n$
- 2: **for** $k = 0, 1, \dots$ **do**
- 3: Choose $s^k \in \partial f^*(q^k)$
- 4: Choose $q^{k+1} \in \partial g(s^k)$
- 5: **if** $g^*(q^{k+1}) - f^*(q^{k+1}) = g^*(q^k) - f^*(q^k)$ **then**
- 6: **return** (s^k, q^k)
- 7: **end if**
- 8: **end for**

However, \bar{s} being a stationary point for f and g is not a *sufficient* condition for \bar{s} to globally minimize $f - g$. So we check whether a corresponding p is a local – and hence global – minimizer of the Lyapunov function L . If it is, then it is indeed an equilibrium price. If not, we go one step in the direction of steepest descent of the Lyapunov function and then restart the DC-algorithm. This is Algorithm 2 (where lines 1–8 are exactly Algorithm 1 with expressed in their economic interpretation; see Lemma 1 for more details).

Algorithm 2: DC auction algorithm

Input: Valid set \mathcal{B} of SSPMA bids

Output: Equilibrium price p and supplementary bundle \bar{s}

- 1: Choose an initial price p^0
- 2: **for** $k = 0, 1, \dots$ **do**
- 3: Choose a bundle s^k demanded at price p^k in the negative-bids auction
- 4: Choose an integral price vector p^{k+1} at which $t + s^k$ is demanded in the positive-bids auction
- 5: **if** $L(p^{k+1}) = L(p^k)$ **then**
- 6: **return** (s^k, p^k)
- 7: **end if**
- 8: **end for**
- 9: **if** there exists $e \in \pm\{0, 1\}^n$ with $L(p^k + e) < L(p^k)$ **then**
- 10: Restart the algorithm with $p^0 := p^k + e$
- 11: **end if**

The value of $L(p^k)$ decreases by at least one in every iteration 2–8 of the algorithm until the termination criterion in Step 5 is satisfied (we refer to Appendix A.5 for details). Whenever the algorithm is restarted in Step 10, L also decreases by at least one. Since there exists a minimizer for L , the algorithm terminates:

Theorem 3 *Algorithm 2 always terminates in a Walrasian equilibrium price.*

Algorithm 2 does not specify how to choose bundles s^k and prices p^{k+1} . Determining bundles s^k is particularly simple when valuations are expressed in the SSPMA – we just allocate each bid with utility maximizing goods. For finding prices p^{k+1} , an instance of (LP) must be solved. We use a min-cost flow solver to do so. Appendix A.6 explains our implementation in more detail.

Obtaining sharp worst-case bounds for Algorithm 2 is challenging due to the very generic nature of the DC-Algorithm 1. Note that the class of functions representable as a difference of convex functions is very large - for example, it contains all functions with continuous second derivative [24]. Also recall that [29] shows that minimizing the difference of two general M^{\natural} -convex functions is NP-hard. Intuitively, we expect Algorithm 2 to perform particularly well when the number of negative bids is small. For example, when there are no negative bids at all, the algorithm boils down to solving the min-cost flow problem (LP). For the general case, we provide the following simple bound for Algorithm 2 by the number of negative bids.

First, observe that we may implement Step 3 to choose a bundle \mathbf{s}^k which is uniquely demanded at some price—and indeed we do so in our practical implementation, because the vertices of demand sets $D_{|\mathcal{B}^-|}(\mathbf{p}^k)$ have this property. We also assume that prices in Step 4 are chosen deterministically – for the same bundle, the algorithm always returns the same price.

Second, observe that if $\mathbf{s}^{k+1} = \mathbf{s}^k$, then the chosen prices in Step 4 are also equal, so the termination criterion 5 is satisfied. After a possible restart, the algorithm also can never reach this bundle again – this would contradict the strict monotonicity properties as we explain in the Appendix (Lemma 6). So in the worst case, after each restart of the algorithm, we directly choose bundles $\mathbf{s}^0 = \mathbf{s}^1$ in the first two iterations which immediately causes another restart. It follows that every possible bundle uniquely demanded in the negative auction is chosen at most twice in Step 3 of the algorithm.³⁰ If there is only one single negative bid, these are exactly $n + 1$ bundles, and so the number of iterations, by which we mean the total number of iterations through the loop from Step 2 to Step 8, of Algorithm 2 is in $\mathcal{O}(n)$. Note that after each restart, we iterate at least once through the for loop, so the number of restarts is also in $\mathcal{O}(n)$. More generally, Proposition 3 shows that $\binom{n+|\mathcal{B}^-|}{n}$ bundles are demanded in total in the negative auction if the weights of all negative bids are equal to one. Since the number of uniquely demanded bundles does not change if we increase weights, $\binom{n+|\mathcal{B}^-|}{n}$ bounds the number of uniquely demanded bundles in general negative auctions. This therefore provides an upper bound on the number of bundles demanded uniquely in this auction, so on the number of iterations of Algorithm 2.

Proposition 5 *Algorithm 2 requires at most $\mathcal{O}\left(\binom{n+|\mathcal{B}^-|}{n}\right)$ iterations for solving the equilibrium pricing problem.*

This analysis gives a rather pessimistic worst-case bound for the algorithm, but it suggests that the algorithm performs particularly well with a low number of negative bids. In fact, in our experimental evaluation, we find that the DC algorithm is even faster than steepest descent in these environments.

When there are only positive bids, Algorithm 2 boils down to solving a single linear program, which can be formulated as a min-cost flow problem on a graph with $|V| = \mathcal{O}(n + |\mathcal{B}|)$ vertices and $|E| = \mathcal{O}(n|\mathcal{B}|)$ edges (see Appendix A.6). The enhanced capacity scaling algorithm [1] finds an optimal integral solution in

³⁰ Moreover, if the same bundle is chosen twice, it is unnecessary to repeat step 4 – the most computationally costly part of the algorithm – so checking for $\mathbf{s}^{k+1} = \mathbf{s}^k$ provides a practical runtime improvement, although it does not alter the complexity class.

$$\mathcal{O}(|E| \log |V| (|E| + |V| \log |V|))$$

iterations, which implies that the pricing problem can be solved in time

$$\mathcal{O}(n^2 |\mathcal{B}|^2 \log^2(n + |\mathcal{B}|)) = \tilde{\mathcal{O}}(n^2 |\mathcal{B}|^2)$$

this way. On the other hand, [6] provide the worst-case bound $\mathcal{O}(n^2 |\mathcal{B}|^2 \log M + n |\mathcal{B}| T(n))$ for the steepest descent algorithm, where $M = \max_{\mathbf{b} \in \mathcal{B}} \|\mathbf{b}\|_\infty$ is the maximal price of a bid vector and $T(n)$ is the complexity of minimizing an n -dimensional submodular function. Note that the total asymptotic runtime of the network flow formulation coincides with the first summand $n^2 |\mathcal{B}|^2 \log M$ of the steepest descent formulation up to a logarithmic factor. However, we may expect the second summand $n |\mathcal{B}| T(n)$ to dominate the runtime. To the best of our knowledge, the best known weakly polynomial worst case bound for minimizing a general integral submodular function is $\mathcal{O}(n^2 \log(nU) \cdot EO + n^3 \log^{\mathcal{O}(1)}(nU))$, where U is an upper bound for the maximal value of the submodular function [16, 30]. Thus, the worst-case bound for $n |\mathcal{B}| T(n)$ cannot be better than $\mathcal{O}(n^4 |\mathcal{B}| \log(nU))$ using known methods. Hence, in particular when the number of goods increases, we may expect the min-cost flow formulation to outperform the steepest descent formulation in the absence of negative bids.³¹

[42] present a polynomial time algorithm for computing competitive equilibrium prices for bidders with general preferences and provide a specialization of their algorithm for gross substitutes valuations, which is the fastest algorithm for this setting currently known. They provide the worst-case bound

$$\begin{aligned} mn \cdot T_V + \mathcal{O}(mn \log m + n^3 \log(mnM) + n^3 \log^{\mathcal{O}(1)}(mnM)) \\ = \tilde{\mathcal{O}}(mn \cdot T_V + n^3) \end{aligned}$$

for the runtime of their algorithm, where n is the number of goods, m is the number of bids, M is the maximum value a bid has for a bundle and T_V is the runtime of the value oracle. So the worst-case runtime of [42]'s algorithm grows linearly in the number of bidders, while the worst-case runtime of our algorithm (when all its bids are positive) grows quadratically.³² However, our algorithm performs much better than [42]'s as we increase the number of units available without changing the number of goods (when all the SSPMA bids are positive), because its worst case is unaffected, whilst [42]'s worst case is cubic in the number of units because they must treat each additional unit as an additional good.³³ Moreover, even when there are only small

³¹ [6]'s algorithm's worst case also depends on M , while our algorithm's does not, so our algorithm is more robust to increases in the precision with which valuations can be expressed (e.g. expressing valuations in cents rather than dollars multiplies M by 100).

³² We assume that the runtime T_V of a value query is constant, and that the total number of SSPMA bids grows linearly in the number of agents.

³³ The worst-case bounds suggest our algorithm also performs better as we increase the number of goods, but this comparison is less clear. The analysis of our algorithm is for bidders with strong substitutes preferences expressed via positive SSPMA bids, and bidders may submit more SSPMA bids in markets with more goods. Although a bid in [42]'s algorithm describes the entire valuation function of a bidder with general

Table 1 Runtimes of the DC and the steepest descent (SD)-algorithm for instances where the number of negative bids is low

#pos. bids	#neg. bids	#goods	Time DC (ms)	Time SD (ms)
1020	20	10	31	394
1020	20	30	105	410
1020	20	50	206	665
3020	20	10	115	1152
3020	20	30	445	1278
3020	20	50	970	1366

numbers of units available of each distinct good, we expect the use cases for the two algorithms to be different: in a setting where fast direct access to bidders' valuations is possible, we expect applying [42]'s algorithm would be preferable to first computing the corresponding SSPMA bids for every bidder and then applying our algorithm. On the other hand, when preferences are provided by bidders in the SSPMA bid language, we expect it is faster to use our algorithm directly, rather than translating the SSPMA bids into value oracles first and then using [42].

4.2 Experimental evaluation

We implemented both the DC auction algorithm and a steepest descent algorithm based on the Lyapunov function. The Lyapunov approach and the restart step in the DC algorithm require the minimization of a submodular function. As in [6], we use the Fujishige-Wolfe algorithm [15], which in practice often outperforms other submodular minimization algorithms.

In our experimental evaluation we solved problems with 10-50 goods, 1020/1200/1500/3020/3500 positive and 20/200/500 negative bids. We drew on a specialization of the algorithm by [5] to randomly generate valid groups of bids, each group consisting of 3 positive and 1 negative bids. Algorithm 3 in Appendix A.7 describes this procedure, and Table 2 in Appendix A.8 gives our results.

The DC algorithm appears to be faster if there are not too many negative bids (less than 200, in our experiments). Table 1 shows a selection of our results for the case of 20 negative bids. So if the number of negative bids is small, which we consider the most likely scenario (see Sect. 2.3), our DC algorithm is a particularly good choice.

However, the main conclusion from Table 2 is that both algorithms are very fast, solving even the largest problems in our experiments in less than 3 seconds. Experiments with up to 50 goods and 10,000 bids can also be solved in a few seconds only.³⁴

Footnote 33 continued

gross substitutes preferences, the runtime T_V of a value query may depend on the number of goods. Ignoring both these effects, our algorithm's worst case depends quadratically on the number of goods, while [42]'s has cubic dependence.

³⁴ Obviously our results are sensitive to the details of the implementations. In particular, in a first, textbook-style implementation, the steepest descent algorithm was much slower beyond 50 goods and 4000 bids.

5 Conclusion

Strong substitutes valuations are of central importance for both theory and practical applications. We have developed a new algorithm for computing competitive equilibrium prices when agents' preferences are expressed using the Strong Substitutes Product-Mix Auction bidding language, a compact language that permits the expression of all strong substitutes valuations (and no other valuations). By contrast with a previous approach of using a standard steepest-descent algorithm that tests candidate solutions in turn, we began from the economics of the problem. We used the fact that the shadow prices of two separate linear programs that maximize value for "positive" and "negative" bids, respectively, must be equal, and proved that our model formulation is dual to the Lyapunov function. We also used the bidding language to provide new insight into the geometric structure of strong substitutes valuations.

Acknowledgements We are grateful for useful discussions with, and helpful advice from Michal Feldman, Paul Goldberg, Edwin Lock, Meg Meyer, Jan Ringling, the referees, and the editor.

Funding Open Access funding enabled and organized by Projekt DEAL. Elizabeth Baldwin and Paul Klempner gratefully acknowledge financial support from the UK Economic and Social Research Council's grant number ES/L003058/1; Martin Bichler gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft (DFG) (BI 1057/9-1).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Appendix: Additional proofs

A.1 Proof of Proposition 1

We need the following simple Lemmas.

Lemma 2 (See, e.g. [44] Lemma 3.1.11) *Let $A, B \subseteq \mathbb{R}^n$ be convex. Then $(A + B) - B = A$.*

Lemma 3 *Suppose $(\text{conv } A) \cap \mathbb{Z}^n = A$ and $(\text{conv } B) \cap \mathbb{Z}^n = B$. Then $(\text{conv } A - \text{conv } B) \cap \mathbb{Z}^n = A - B$.*

Proof If $\mathbf{x} \in (\text{conv } A - \text{conv } B) \cap \mathbb{Z}^n$ then $\mathbf{x} \in \mathbb{Z}^n$ and $\mathbf{x} + \text{conv } B \subseteq \text{conv } A$, so $\mathbf{x} + (\text{conv } B) \cap \mathbb{Z}^n \subseteq (\text{conv } A) \cap \mathbb{Z}^n$, and therefore $\mathbf{x} + B \subseteq A$. So $\mathbf{x} \in A - B$. Conversely,

Footnote 34 continued

However, an additional pre-processing step led to significant improvements in the steepest descent algorithm, and we report the results for this improved steepest descent algorithm. With this improvement in the steepest descent algorithm, the differences between the algorithms seem likely to be small in most applications.




Fig. 5 Graph used to construct the valuation v_r from [40]

if $\mathbf{x} \in A - B$ then $\mathbf{x} \in \mathbb{Z}^n$, and $\mathbf{x} + B \subseteq A$ implies $\text{conv}(\mathbf{x} + B) = \mathbf{x} + \text{conv } B \subseteq \text{conv } A$. \square

Proof of Proposition 1 By the strong substitutes property, the sets $D_{\mathcal{B}_+}(\mathbf{p})$ and $D_{|\mathcal{B}_-|}(\mathbf{p})$ are equal to the set of integer points of their respective convex hull, as by definition is $D_{\mathcal{B}}(\mathbf{p})$. So if we can show that $\text{conv } D_{\mathcal{B}}(\mathbf{p}) + \text{conv } D_{|\mathcal{B}_-|}(\mathbf{p}) = \text{conv } D_{\mathcal{B}_+}(\mathbf{p})$, this implies by Lemma 2 that $\text{conv } D_{\mathcal{B}}(\mathbf{p}) = \text{conv } D_{\mathcal{B}_+}(\mathbf{p}) - \text{conv } D_{|\mathcal{B}_-|}(\mathbf{p})$ and by Lemma 3 consequently that $D_{\mathcal{B}}(\mathbf{p}) = D_{\mathcal{B}_+}(\mathbf{p}) - D_{|\mathcal{B}_-|}(\mathbf{p})$.

But, as \mathcal{B} is a valid set of bids, we know by [6] Theorem 2.3 that $u_{\mathcal{B}}$ is the indirect utility of a strong substitutes valuation v such that $D_v(\mathbf{p}) = D_{\mathcal{B}}(\mathbf{p})$ for all $\mathbf{p} \in \mathbb{R}^n$. It follows that each vertex \mathbf{x} of $D_{\mathcal{B}}(\mathbf{p})$ is the unique element of $D_{\mathcal{B}}(\mathbf{q})$ for a price \mathbf{q} close to \mathbf{p} and such that \mathbf{x} minimizes $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ for $\mathbf{x}' \in D_{\mathcal{B}}(\mathbf{p})$. But similarly the minimizers of $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ for $\mathbf{x}' \in D_{\mathcal{B}_+}(\mathbf{p})$ and $\mathbf{x}' \in D_{|\mathcal{B}_-|}(\mathbf{p})$ are, respectively, the unique elements of $D_{\mathcal{B}_+}(\mathbf{q})$ and $D_{|\mathcal{B}_-|}(\mathbf{q})$. By definition, we have $D_{\mathcal{B}}(\mathbf{q}) = D_{\mathcal{B}_+}(\mathbf{q}) - D_{|\mathcal{B}_-|}(\mathbf{q})$, so $D_{\mathcal{B}}(\mathbf{q}) + D_{|\mathcal{B}_-|}(\mathbf{q}) = D_{\mathcal{B}_+}(\mathbf{q})$. As this holds for all extreme points of $D_{\mathcal{B}}(\mathbf{p})$, it follows that $\text{conv } D_{\mathcal{B}}(\mathbf{p}) + \text{conv } D_{|\mathcal{B}_-|}(\mathbf{p}) = \text{conv } D_{\mathcal{B}_+}(\mathbf{p})$, as required. \square

A.2 The valuation v_r from [40]

We now explain the construction of v_r from [40]. Let $G = (V, E)$ be an undirected graph with 4 vertices and 6 edges $E = \{1, \dots, 6\}$, such that every vertex is connected to every other vertex by an edge (see Figure 5).

A subset H of E is called *independent* if it contains no cycles. For any $H \subseteq E$, the *rank* of H is the maximal cardinality of an independent subset contained in H :

$$\text{rank}(H) = \max \{|H'| : H' \subseteq H \text{ is independent}\}.$$

The rank function induces the valuation $v_r : \{0, 1\}^6 \rightarrow \mathbb{Z}$ given by $v_r(\mathbf{x}) = \text{rank}(\{i : x_i = 1\})$. As [40] show, v_r is strong substitutes. However, it does not satisfy the property of *strong exchangeability* which, as [40] show, is a characteristic of every endowed assignment valuation. Consequently, it is not possible to express v_r by endowed assignment messages. We demonstrate, however, that it can be expressed using the SSPMA. Note that valuations induced by SSPMA bids are always defined on a scaled simplex Δ_W for some total weight $W \in \mathbb{Z}_{\geq 0}$. We thus naturally extend v_r to $\Delta_6 \supseteq \{0, 1\}^6$ by assuming free disposal: $v_r(\mathbf{x}) = \text{rank}(\{i : x_i \geq 1\})$.

Proof of Proposition 2 Given $H \subseteq E$, we write $\mathbf{b}^H := \sum_{i \in H} \mathbf{e}_i$. We make the following bids:

0. Place a bid \mathbf{b}^\emptyset with $w(\mathbf{b}^\emptyset) = 3$.
1. For all $H \subseteq E$ with $|H| = 3$ and H^c is a cycle in G , make a bid \mathbf{b}^H with $w(\mathbf{b}^H) = 1$.
2. For all $H \subseteq E$ constituting a cycle of length 4, make a bid \mathbf{b}^H with $w(\mathbf{b}^H) = 1$.
3. For all $H \subseteq E$ with $|H| = 5$ make a bid \mathbf{b}^H with $w(\mathbf{b}^H) = -1$.
4. Make a bid \mathbf{b}^E with $w(\mathbf{b}^E) = 2$.

Denote by $v_r(\mathbf{x}) = \text{rank}(\{i : x_i \geq 1\})$ for $\mathbf{x} \in \Delta_6$ the valuation induced by the rank function, and by $v_{\mathcal{B}}(\mathbf{x})$ the valuation induced by the above bids. Our goal is to show $v_{\mathcal{B}} = v_r$. Note that bid 0 only ensures that the domains of v_r and $v_{\mathcal{B}}$ are equal, and does not “contribute” to the valuations apart from this. So let us check that indeed $\text{dom } v_{\mathcal{B}} = \Delta_6$. There is 1 bid of type 0, 4 bids of type 1, 3 bids of type 2, and 1 bid of type 4. So summing up the weights of these bids gives $W_+ = 12$. On the other hand, there are 6 bids of type 3, so $W_- = 6$, and consequently $\text{dom } v_{\mathcal{B}} = \Delta_{12-6} = \Delta_6$.

We have $u_r(\mathbf{p}) = \max_{\mathbf{x} \in \Delta_6} v_r(\mathbf{x}) - \langle \mathbf{p}, \mathbf{x} \rangle$ and $u_{\mathcal{B}}$ is defined by Equation (3). Recall from Section 2.1 that, for $i \in \{r, \mathcal{B}\}$, we have $v_i(\mathbf{x}) = \min_{\mathbf{p} \in \mathbb{R}^6} u_i(\mathbf{p}) + \langle \mathbf{p}, \mathbf{x} \rangle$, where one can check that $\mathbf{p} \mapsto u_i(\mathbf{p}) + \langle \mathbf{p}, \mathbf{x} \rangle$ always possesses a non-negative minimizer \mathbf{p} for $\mathbf{x} \in \Delta_6$. So in order to prove Proposition 2, it suffices to show that $u_r(\mathbf{p}) = u_{\mathcal{B}}(\mathbf{p})$ for all $\mathbf{p} \in \mathbb{R}_{\geq 0}^6$. By L^1 -convexity of u_r and $u_{\mathcal{B}}$ [36], both are determined uniquely on $\mathbb{R}_{\geq 0}^6$ by the values $u_r(\mathbf{p})$ and, respectively, $u_{\mathcal{B}}(\mathbf{p})$ for $\mathbf{p} \in \mathbb{Z}_{\geq 0}^6$. Moreover, given $\mathbf{p} \in \mathbb{Z}_{\geq 0}^6$, define $\tilde{\mathbf{p}}$ by $\tilde{p}_i = p_i$ if $p_i \leq 1$ and $\tilde{p}_i = 1$, otherwise. Since the marginal value of any good is at most 1 for v_r , and no bid in \mathcal{B} has any value greater than 1, allocating a good i with $\tilde{p}_i = 1$ can never increase utilities, so we have $u_r(\mathbf{p}) = u_r(\tilde{\mathbf{p}})$ and $u_{\mathcal{B}}(\mathbf{p}) = u_{\mathcal{B}}(\tilde{\mathbf{p}})$. So our problem reduces to showing that $u_r(\mathbf{p}) = u_{\mathcal{B}}(\mathbf{p})$ for all $\mathbf{p} \in \{0, 1\}^6$. For $H \subseteq \{1, \dots, 6\}$, denote by $\mathbf{p}^H \in \{0, 1\}^6$ the price vector with $p_i^H = 1$ if and only if $i \in H$. We will show that $u_r(\mathbf{p}^H) = u_{\mathcal{B}}(\mathbf{p}^H)$ for all $H \subseteq \{1, \dots, 6\}$.

We claim that $u_r(\mathbf{p}^H) = \text{rank}(H^c)$. To see this, let \mathbf{x} be a bundle with $u_r(\mathbf{p}^H) = v_r(\mathbf{x}) - \langle \mathbf{p}^H, \mathbf{x} \rangle = v_r(\mathbf{x}) - \sum_{i \in H} x_i$. Let $P = \{i : x_i \geq 1\}$. Then

$$\begin{aligned} u_r(\mathbf{p}^H) &= \text{rank}(P) - |P \cap H| \leq \text{rank}(P \cap H^c) + \text{rank}(P \cap H) - |P \cap H| \\ &\leq \text{rank}(P \cap H^c) + |P \cap H| - |P \cap H| \\ &\leq \text{rank}(H^c) = v_r \left(\sum_{i \in H^c} \mathbf{e}_i \right) - 0 \leq u_r(\mathbf{p}^H) \end{aligned}$$

by properties of matroid rank functions. Consequently, equality must hold everywhere, so $u_r(\mathbf{p}^H) = \text{rank}(H^c)$.

Regarding the indirect utility of our bids, we observe that at prices \mathbf{p}^H , the bid $\mathbf{b}^{\tilde{H}}$ generates a utility of $w(\mathbf{b}^{\tilde{H}})$ if and only if $\tilde{H} \cap H^c \neq \emptyset$, i.e., if and only if $\mathbf{b}^{\tilde{H}}$ has positive value for at least one good not in H . Otherwise it generates utility 0.

We now consider all subsets $H \subseteq \{1, \dots, 6\}$ and show that in each case, $u_r(\mathbf{p}^H) = u_{\mathcal{B}}(\mathbf{p}^H)$.

First, for price vectors \mathbf{p}^H with $|H| < 3$, all bids are accepted, since every placed bid has positive values for at least 3 goods. There are 4 bids of type 1, 3 bids of type 2, 6 bids of type 3 and 1 bid of type 4. In total, we get

$$u_b(\mathbf{p}^H) = 4 \cdot 1 + 3 \cdot 1 + 6 \cdot (-1) + 1 \cdot 2 = 3.$$

On the other hand, one can see from Fig. 5 that every subset containing at least 4 edges contains a cycle-free subset of cardinality 3, and there is no cycle-free subset with more than 3 elements. Consequently, $u_r(\mathbf{p}^H) = \text{rank}(H^c) = 3$.

Now consider \mathbf{p}^H with $|H| = 3$. Obviously, all bids on more than 3 edges get accepted. A bid $\mathbf{b}^{\tilde{H}}$ with $|\tilde{H}| = 3$ is rejected, if and only if $H = \tilde{H}$. In this case, H^c is a cycle of length 3, so $u_r(\mathbf{p}^H) = \text{rank}(H^c) = 2$. We then also clearly have $u_{\mathcal{B}}(\mathbf{p}^H) = 2$, since exactly one bid is rejected, and all others are accepted.

On the other hand, if $|H| = 3$ and no bid is rejected, H^c is cycle free, so $u_r(\mathbf{p}^H) = \text{rank}(H^c) = 3 = u_{\mathcal{B}}(\mathbf{p}^H)$.

Next, suppose $|H| = 4$, so $u_r(\mathbf{p}^H) = \text{rank}(H^c) = 2$, because 2 edges cannot form a cycle. Regarding the bids, if H is a cycle of length 4, one bid of type 2 is rejected. In this case, H^c consists of two non adjacent edges. Consequently, there is no $i \in H$ such that $\{i\} \cup H^c$ is a cycle. Equivalently, for no $\tilde{H} \subseteq H$ with $|\tilde{H}| = 3$ we have that \tilde{H}^c is a cycle, so no bid of type 1 is rejected, and $u_{\mathcal{B}}(\mathbf{p}^H) = 2$.

If, otherwise, H has no cycle of length 4, H^c consists of two adjacent edges. Thus, there is a unique $e \in H$ with $\{e\} \cup H^c$ being a cycle, so a single bid of type 1 is rejected, which means that again $u_{\mathcal{B}}(\mathbf{p}^H) = 2$.

Concerning $|H| = 5$, since the graph is complete, we can assume by symmetry that $H = \{1, 2, 3, 4, 5\}$. Then the bids $\mathbf{b}^{\tilde{H}}$ with

$$\tilde{H} \in \{\{1, 2, 5\}, \{3, 4, 5\}, \{1, 2, 3, 4\}, \{1, 2, 3, 4, 5\}\}$$

are rejected, and $u_{\mathcal{B}}(\mathbf{p}^H) = 2 \cdot 1 + 2 \cdot 1 + 5 \cdot (-1) + 1 \cdot 2 = 1 = \text{rank}(H^c) = u_r(\mathbf{p}^H)$.

Finally, for $H = E$, all bids are rejected, so $u_{\mathcal{B}}(\mathbf{p}^H) = u_r(\mathbf{p}^H) = 0$.

We have shown that for all $p \in \{0, 1\}^6$, $u_{\mathcal{B}}(\mathbf{p}) = u_r(\mathbf{p})$, which proves our statement. \square

A.3 Proof of Proposition 3

Proof of Proposition 3 We will show that $D = \Delta_W = \mathbb{Z}^n \cap W\Delta$, where Δ is the standard simplex in dimension n , spanned by $\mathbf{0}$ and the standard unit vectors \mathbf{e}_i . Since $W\Delta$ contains exactly $\binom{n+W}{n}$ integer points [10, Theorem 2.2], the remaining results follow.

By the strong substitutes property, $D = (\text{conv } D) \cap \mathbb{Z}^n$, so it suffices to show that $\text{conv } D = W\Delta$. To that goal we note that if we set $p_i = -1$ and p_j very large for $j \neq i$, then $D(\mathbf{p}) = \{W\mathbf{e}_i\}$, since every bid \mathbf{b} is allocated with $w(\mathbf{b})$ items of good i and the total weight of all bids is W . Also, for a very large price (in every coordinate) \mathbf{p} , we have $D_{\mathcal{B}}(\mathbf{p}) = \{\mathbf{0}\}$. Consequently, $\text{conv } D \supseteq W\Delta$. To see the reverse inclusion, note that any demanded bundle cannot contain strictly more than W items, as some bid

would have to be allocated with more than $w(\mathbf{b})$ items otherwise. The lower bounds come from the basic inequality $\binom{m}{k} \geq (m/k)^k$. \square

A.4 Proof of Theorem 1

Proof of Lemma 1 The linear program (LP) of Sect. 3 is clearly defined for any $\mathbf{x} = \mathbf{t} \in \text{conv } \Delta_W$, and we can use this to assign a real value to $\tilde{v}_{\mathcal{B}}(\mathbf{x})$ for $\mathbf{x} \in \text{conv } \Delta_W$ and set $f = -\tilde{v}_{\mathcal{B}}$. Since f is a polyhedral convex function according to [43, p. 172], its subdifferential is nonempty at every point of $\text{dom } f$ [43, Theorem 23.10], so $\text{dom } \partial f = \text{dom } f = \text{conv } \Delta_W$. Let us consider the convex conjugate f^* of $f(\mathbf{x}) = -\tilde{v}_{\mathcal{B}}(\mathbf{x})$. By definition, $f^*(\mathbf{q}) = \max_{\mathbf{x} \in \text{conv } \Delta_W} \langle \mathbf{q}, \mathbf{x} \rangle + \tilde{v}_{\mathcal{B}}(\mathbf{x})$, or in LP-form:

$$\begin{aligned} f^*(\mathbf{q}) = \max \quad & \sum_{\mathbf{b} \in \mathcal{B}} \sum_{i \in [n]} (b_i + q_i) y_{\mathbf{b}i} \\ \text{s.t. } \quad & x_i = \sum_{\mathbf{b} \in \mathcal{B}} y_{\mathbf{b}i} \forall i \in [n] \\ & \sum_{i \in [n]} y_{\mathbf{b}i} \leq w(\mathbf{b}) \forall \mathbf{b} \in \mathcal{B} \\ & y_{\mathbf{b}i} \geq 0 \forall \mathbf{b} \in \mathcal{B}, i \in [n]. \end{aligned}$$

Note that since the set of feasible solutions \mathbf{x} is compact, $f^*(\mathbf{q})$ attains a finite value for all $\mathbf{q} \in \mathbb{R}^n$, so $\text{dom } \partial f^* = \text{dom } f^* = \mathbb{R}^n$, since f^* is also polyhedral convex. Let us now derive the expressions for ∂f and ∂f^* . To that goal, note that \mathbf{x} maximizes the above linear program if and only if $\mathbf{x} \in \partial f^*(\mathbf{q})$, which is in turn equivalent to $\mathbf{q} \in \partial f(\mathbf{x})$ [43, Theorem 23.5]. It is not hard to see from Equations (1) and (2) that the variables $y_{\mathbf{b}i}$ constitute an optimal solution for the above linear program, if and only if for every fixed \mathbf{b} the vector $(y_{\mathbf{b}i})_{i=1}^n$ lies in $\text{conv } D_{\mathbf{b}}(-\mathbf{q})$, which can be seen to be equivalent to $\mathbf{x} \in \text{conv } D_{\mathcal{B}}(-\mathbf{q})$ (recall that in the case of only positive bids, the aggregate demand set is just the Minkowski sum of the individual demand sets). It now directly follows that $\partial f^*(\mathbf{q}) = \text{conv } D_{\mathcal{B}}(-\mathbf{q})$ and $\partial f(\mathbf{x}) = -\{\mathbf{p} : \mathbf{x} \in \text{conv } D_{\mathcal{B}}(\mathbf{p})\}$. \square

Lemma 4 Let \mathcal{B} be a valid collection of bids. Let $\mathbf{t} \in \Delta_W$ and $\mathbf{s} \in \Delta_{W_-}$. Then $\mathbf{t} + \mathbf{s} \in \Delta_{W_+}$. Consequently, for the convex extensions f of $\mathbf{s} \mapsto -v_{|\mathcal{B}_-|}(\mathbf{s})$ and g of $\mathbf{s} \mapsto -v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s})$ we have that $\text{dom } f = \text{conv } \Delta_{W_-} \subseteq \text{conv}\{\mathbf{s} \in \mathbb{Z}^n : \mathbf{t} + \mathbf{s} \in \Delta_{W_+}\} = \text{dom } g$.

Proof As $\mathbf{t} \in \Delta_W$, we have $\sum_{i=1}^n t_i \leq W$. Similarly, $\sum_{i=1}^n s_i \leq W_-$. Since $W = W_+ - W_-$ it follows that $\sum_{i=1}^n (t_i + s_i) \leq W_+$, so $\mathbf{t} + \mathbf{s} \in \Delta_{W_+}$. This directly implies the second part of the Lemma. \square

Proof of Theorem 1 Let f be the convex extension of $\mathbf{s} \mapsto -v_{|\mathcal{B}_-|}(\mathbf{s})$ and g the convex extension of $\mathbf{s} \mapsto -v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s})$. Then $\text{dom } f = \text{conv } \Delta_{W_-} \subseteq \text{dom } g$ and $\text{dom } g^* = \mathbb{R}^n = \text{dom } f^*$ by Lemmas 1 and 4. From Lemma 1 we know that $f^*(\mathbf{q}) = u_{|\mathcal{B}_-|}(-\mathbf{q})$. Similarly, $g^*(\mathbf{q}) = u_{\mathcal{B}_+}(-\mathbf{q}) - \langle \mathbf{q}, \mathbf{t} \rangle$. So we can apply Theorem 2 to $f - g$ and get

$$\min_{\mathbf{s} \in \text{conv } \Delta_{W_-}} \tilde{v}_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - \tilde{v}_{|\mathcal{B}_-|}(\mathbf{s}) = \min_{\mathbf{q} \in \mathbb{R}^n} u_{\mathcal{B}_+}(-\mathbf{q}) - u_{|\mathcal{B}_-|}(-\mathbf{q}) - \langle \mathbf{q}, \mathbf{t} \rangle,$$

if any of the two problems has a solution. By substituting $\mathbf{p} = -\mathbf{q}$, we can rewrite the problem on the right as

$$\min_{\mathbf{p} \in \mathbb{R}^n} u_{\mathcal{B}_+}(\mathbf{p}) - u_{|\mathcal{B}_-|}(\mathbf{p}) + \langle \mathbf{p}, \mathbf{t} \rangle = \min_{\mathbf{p} \in \mathbb{R}^n} u_{\mathcal{B}}(\mathbf{p}) + \langle \mathbf{p}, \mathbf{t} \rangle.$$

The expression $u_{\mathcal{B}}(\mathbf{p}) + \langle \mathbf{p}, \mathbf{t} \rangle$ is exactly the Lyapunov function $L(\mathbf{p})$ introduced in Section 3. For strong substitutes valuations, the Lyapunov function always attains a minimum, and the set of minimizers is equal to the set of equilibrium prices for the target \mathbf{t} [3]. Consequently, the problem $\min_{\mathbf{s} \in \text{conv } \Delta_{W_-}} \tilde{v}_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - \tilde{v}_{|\mathcal{B}_-|}(\mathbf{s})$ also has a solution $\mathbf{s} \in \text{conv } \Delta_{W_-}$, and the values of both minimization problems are equal. There exists at least one integral solution $\bar{\mathbf{s}} \in \Delta_W$ to this problem: Let \mathbf{p} be a minimizer of the Lyapunov function. By Theorem 2, each $\mathbf{s} \in \partial f^*(-\mathbf{p}) = \text{conv } D_{\mathcal{B}}(\mathbf{p})$ minimizes $\tilde{v}_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - \tilde{v}_{|\mathcal{B}_-|}(\mathbf{s})$, so in particular each $\bar{\mathbf{s}} \in D_{\mathcal{B}}(\mathbf{p}) \neq \emptyset$ does so. Since the valuations $v_{\mathcal{B}_+}$ and $v_{|\mathcal{B}_-|}$ coincide on integral bundles with $\tilde{v}_{\mathcal{B}_+}$ and $\tilde{v}_{|\mathcal{B}_-|}$ by construction,

$$\min_{\mathbf{s} \in \Delta_{W_-}} v_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - v_{|\mathcal{B}_-|}(\mathbf{s}) = \min_{\mathbf{s} \in \text{conv } \Delta_{W_-}} \tilde{v}_{\mathcal{B}_+}(\mathbf{t} + \mathbf{s}) - \tilde{v}_{|\mathcal{B}_-|}(\mathbf{s}) = \min_{\mathbf{p} \in \mathbb{R}^n} L(\mathbf{p}).$$

Finally, again by Theorem 2, if $\bar{\mathbf{s}} \in \Delta_{W_-}$ is a minimizer, each \mathbf{p} with $-\bar{\mathbf{p}} \in \partial g(\mathbf{t} + \mathbf{s}) = -\{\mathbf{p} \in \mathbb{R}^n : \mathbf{t} + \bar{\mathbf{s}} \in \text{conv } D_{\mathcal{B}_+}(\mathbf{p})\}$ minimizes L . In other words, each equilibrium price for $\mathbf{t} + \mathbf{s}$ for the positive auction is an equilibrium price for the complete auction as well. \square

A.5 Proof of Theorem 3

We now prove Theorem 3 which states that our DC-algorithm always terminates in a global minimum. First, we collect some properties of the DC-algorithm.

Proposition 6 *Algorithm 1 has the following properties:*

1. *The sequences $f(\mathbf{s}^k) - g(\mathbf{s}^k)$ and $g^*(\mathbf{q}^k) - f^*(\mathbf{q}^k)$ are decreasing. Furthermore, $f(\mathbf{s}^k) - g(\mathbf{s}^k) \leq g^*(\mathbf{q}^k) - f^*(\mathbf{q}^k)$ and $g^*(\mathbf{q}^{k+1}) - f^*(\mathbf{q}^{k+1}) \leq f(\mathbf{s}^k) - g(\mathbf{s}^k)$. The sequence $g^*(\mathbf{q}^k) - f^*(\mathbf{q}^k)$ is strictly decreasing until the termination criterion is met.*
2. *If the algorithm terminates with $(\mathbf{s}^k, \mathbf{q}^k)$, then $\mathbf{s}^k \in \partial f^*(\mathbf{q}^k) \cap \partial g^*(\mathbf{q}^k)$ and $\mathbf{q}^k \in \partial f(\mathbf{s}^k) \cap \partial g(\mathbf{s}^k)$. Furthermore, $f(\mathbf{s}^k) - g(\mathbf{s}^k) = g^*(\mathbf{q}^k) - f^*(\mathbf{q}^k)$.*

Proof A proof can be found in [46, Theorem 3]. The sequence $g^*(\mathbf{q}^k) - f^*(\mathbf{q}^k)$ is strictly decreasing because the algorithm terminates as soon as $g^*(\mathbf{q}^k) - f^*(\mathbf{q}^k) = g^*(\mathbf{q}^{k+1}) - f^*(\mathbf{q}^{k+1})$. \square

Next, we show that we can always restart the DC algorithm from a computed stationary point.

Lemma 5 *Suppose that in Algorithm 2 the termination criterion in line 5 is met with supply \mathbf{s} and price vector \mathbf{p} . If \mathbf{p} is no equilibrium price, then there exists a descent direction $\mathbf{e} \in \pm\{0, 1\}^n$ of the Lyapunov function at \mathbf{p} . If we restart the algorithm with $\tilde{\mathbf{p}}^0 := \mathbf{p} + \mathbf{e}$, we have for all elements $(\tilde{\mathbf{p}}^k, \tilde{\mathbf{s}}^k)$ of the new sequence that $L(\tilde{\mathbf{p}}^k) \leq L(\mathbf{p}) - 1$.*




Fig. 6 A flow graph modeling (LP)

Proof If the returned price \mathbf{p} is no equilibrium price, then it is no minimizer of the Lyapunov function [4]. It follows by L^{\natural} -convexity of L that there exists $\mathbf{e} \in \pm\{0, 1\}^n$ with $L(\mathbf{p} + \mathbf{e}) \leq L(\mathbf{p}) - 1$ [36]. By Property 1 in Proposition 6 we have that $L(\tilde{\mathbf{p}}^k) \leq L(\mathbf{p}) - 1$ for the sequence of prices generated after the restart with initial price $\tilde{\mathbf{p}}^0 = \mathbf{p} + \mathbf{e}$. Since L possesses a minimizer [4] and after each restart the value of the Lyapunov function decreases by at least 1, the algorithm terminates with an equilibrium price. \square

This completes the proof of Theorem 3: In each step of the main loop, the value of $L(\mathbf{p}^k)$ strictly decreases by an integer amount, and if we leave the main loop, we either restart with a price vector of a strictly smaller value, or we terminate, if we are at a global minimum already.

Lemma 6 *Suppose that the prices \mathbf{p}^{k+1} in Step 4 of Algorithm 2 are chosen deterministically. Let $R \in \mathbb{Z}_{\geq 0}$ be the number of restarts of the algorithm and let $S_r = (\mathbf{s}_r^0, \mathbf{s}_r^1, \dots, \mathbf{s}_r^{|S_r|})$ denote the sequence of iterates generated in Step 3 after the r -th restart for $r = 0, \dots, R$ (S_0 is the sequence before the first restart). Then for $r_1 \neq r_2$, S_{r_1} and S_{r_2} do not contain any common bundle. Moreover, for each r the bundles $\mathbf{s}_r^0, \dots, \mathbf{s}_r^{|S_r|-1}$ are pairwise distinct.*

Proof Suppose that $\mathbf{s}_{r_1}^k = \mathbf{s}_{r_2}^l$ for some $r_1 \leq r_2$ and k, l . Then we have for the computed prices in Step 4 that $\mathbf{p}_{r_1}^{k+1} = \mathbf{p}_{r_2}^{l+1}$, so $L(\mathbf{p}_{r_1}^{k+1}) = L(\mathbf{p}_{r_2}^{l+1})$. This can only happen if $r_1 = r_2$, since otherwise $L(\mathbf{p}_{r_2}^{l+1}) \leq L(\mathbf{p}_{r_1}^{k+1}) - 1$ by Lemma 5. Now suppose that $r_1 = r_2 = r$ and $k \leq l$. Then again $L(\mathbf{p}_r^{k+1}) = L(\mathbf{p}_r^{l+1})$. By Property 1 of Proposition 6, it follows that $k = l - 1$ and the termination criterion is satisfied in iteration l , so $l = |S_r|$ and the bundles $\mathbf{s}_r^0, \dots, \mathbf{s}_r^{|S_r|-1}$ are pairwise distinct. \square

A.6 DC algorithm

For the DC algorithm, reformulating the (LP) as a min-cost flow problem comes with a significant computational advantage as compared to solving it with a generic LP-solver. We briefly describe the general min-cost flow problem. For more details, we refer to [1]. Given a directed graph, an arc is a tuple (\mathbf{v}, \mathbf{w}) where \mathbf{v} and \mathbf{w} are nodes

of the graph. We denote by $u(\mathbf{v}, \mathbf{w}) \geq 0$ the maximum capacity of this arc and by $c(\mathbf{v}, \mathbf{w}) \in \mathbb{R}$ the cost per unit flow along (\mathbf{v}, \mathbf{w}) . For a node \mathbf{v} , we denote by $\beta(\mathbf{v}) \in \mathbb{R}$ the supply at node \mathbf{v} . Depending on the sign of $\beta(\mathbf{v})$, a total flow of $|\beta(\mathbf{v})|$ must leave (positive sign) or enter (negative sign) \mathbf{v} . If the supply is 0, the inflow must equal the outflow. A flow f assigns a value $f(\mathbf{v}, \mathbf{w}) \in \mathbb{R}$ to each arc, the amount of flow from \mathbf{v} to \mathbf{w} . It is feasible, if $0 \leq f(\mathbf{v}, \mathbf{w}) \leq u(\mathbf{v}, \mathbf{w})$ for each arc (\mathbf{v}, \mathbf{w}) in the network, and $\sum_{\mathbf{w}} f(\mathbf{v}, \mathbf{w}) - \sum_{\mathbf{w}} f(\mathbf{w}, \mathbf{v}) = \beta(\mathbf{v})$ for all nodes \mathbf{v} , where the sums run over all \mathbf{w} such that (\mathbf{v}, \mathbf{w}) , respectively (\mathbf{w}, \mathbf{v}) is an arc in the network. The cost of the flow is equal to $\sum_{(\mathbf{v}, \mathbf{w})} c(\mathbf{v}, \mathbf{w}) f(\mathbf{v}, \mathbf{w})$. The objective of the min-cost flow problem is to find a feasible flow with minimal cost.

The linear program (LP) is used to solve Step 4 in Algorithm 2 where we need to compute a price vector \mathbf{p}^{k+1} at which the bundle $\mathbf{t} + \mathbf{s}^k$ is demanded. A straightforward network flow model for (LP) is illustrated in Fig. 6. For each good $i \in \{1, \dots, n\}$ there is a node \mathbf{g}^i , and for each of the $m = |\mathcal{B}_+|$ positive bids indexed by $j \in \{1, \dots, m\}$ there is a node \mathbf{b}^j . Finally, there is a destination node \mathbf{d} . In our flow network, there is an arc $(\mathbf{g}^i, \mathbf{b}^j)$ from each good i to each bid j with unlimited capacity $u(\mathbf{g}^i, \mathbf{b}^j) = \infty$ and cost $c(\mathbf{g}^i, \mathbf{b}^j) = -b_i^j$, i.e., the negative value of bid j for good i . The arcs $(\mathbf{b}^j, \mathbf{d})$ from the bids to the destination node have capacity $u(\mathbf{b}^j, \mathbf{d}) = w(\mathbf{b}^j)$ and cost $c(\mathbf{b}^j, \mathbf{d}) = 0$. In Step 4 of Algorithm 2, a supply of $\mathbf{t} + \mathbf{s}^k$ must be distributed among the bids. We set the supply of node \mathbf{g}^i to $\beta(\mathbf{g}^i) = t_i + s_i^k$ and the supply of node \mathbf{d} to $\beta(\mathbf{d}) = -\sum_{i=1}^n t_i + s_i^k$. Finally, the supply of node \mathbf{b}^j is set to $\beta(\mathbf{b}^j) = 0$. Since $\mathbf{t} + \mathbf{s}^k \in \Delta_{W_+}$ (Appendix Lemma 4), $\sum_{i=1}^n t_i + s_i^k \leq \sum_{j=1}^m w(\mathbf{b}^j)$, so a feasible flow f exists. Moreover, since the capacities and supplies are all integral, an integral optimal flow exists. Note that the proposed flow network contains arcs with negative cost. If required by a specific solver, it can however easily be transformed into a network with only non-negative costs [1, p. 40].

We assume that the applied min-cost flow solver provides us with an integral optimal flow f , as well as with an integral optimal dual solution, consisting of node potentials $\pi(\mathbf{v}) \in \mathbb{R}$ for each node \mathbf{v} in the network. These satisfy the following complementary slackness conditions [1, Theorem 9.4].

1. If $c(\mathbf{v}, \mathbf{w}) + \pi(\mathbf{v}) - \pi(\mathbf{w}) > 0$, then $f(\mathbf{v}, \mathbf{w}) = 0$.
2. If $0 < f(\mathbf{v}, \mathbf{w}) < u(\mathbf{v}, \mathbf{w})$, then $c(\mathbf{v}, \mathbf{w}) + \pi(\mathbf{v}) - \pi(\mathbf{w}) = 0$.
3. If $c(\mathbf{v}, \mathbf{w}) + \pi(\mathbf{v}) - \pi(\mathbf{w}) < 0$, then $f(\mathbf{v}, \mathbf{w}) = u(\mathbf{v}, \mathbf{w})$.

From the complementary slackness conditions it is not hard to deduce that \mathbf{p} defined by $p_i = \pi(\mathbf{g}^i) - \pi(\mathbf{d})$ is an equilibrium price vector for the supply $\mathbf{t} + \mathbf{s}^k$, so we can choose $\mathbf{p}^{k+1} = \mathbf{p}$ in Step 4 of the algorithm.

Let us finally consider Step 3 of Algorithm 2, where a bundle \mathbf{s}^k must be chosen that is demanded at price $\mathbf{p}^k \in \mathbb{Z}^n$ in the negative auction. This is particularly easy to do in the Product-Mix Auction (see also [6]): For each bid \mathbf{b} in the negative auction, choose a bundle $\mathbf{s}(\mathbf{b}) \in D_{\mathbf{b}}(\mathbf{p})$. By Eqs. (1) and (2), this can be done in linear time in the number of different goods. Then set $\mathbf{s}^k = \sum_{\mathbf{b} \in |\mathcal{B}_-|} \mathbf{s}(\mathbf{b})$. In our implementation, we choose a bundle \mathbf{s}^k which is a vertex of $D_{|\mathcal{B}_-|}(\mathbf{p})$. This can be achieved by suitably perturbing \mathbf{p} : Let $\mathbf{q} = \mathbf{p} + \Delta$ be a price such that $D_{|\mathcal{B}_-|}(\mathbf{q}) \cap D_{|\mathcal{B}_-|}(\mathbf{p}) \neq \emptyset$ and

$|D_{|\mathcal{B}_-|}(\mathbf{q})| = 1$. For example, $\Delta = (\varepsilon, 2\varepsilon, \dots, n\varepsilon)$ works for $\varepsilon > 0$ small enough. Then simply choose the unique $\mathbf{s}(\mathbf{b}) \in D_{|\mathcal{B}_-|}(\mathbf{q})$ and construct \mathbf{s}^k as above.

A.7 Generating valid bid groups

Algorithm 3: Algorithm for generating valid groups of positive and negative bids, used for experiments.

Input: Random parameters $a_0, \dots, a_n \in \mathbb{Z}_{>0}$, a random permutation $\sigma \in S_n$, weight $w \in \mathbb{Z}_{>0}$, displacement parameter $\mathbf{c} \in \mathbb{Z}_{\geq 0}^n$.

Output: Group of 3 positive and 1 negative bids.

- 1: Generate 2 vectors $\mathbf{v}^1, \mathbf{v}^2 \in \mathbb{Z}_{\geq 0}^n$:
- 2: **for** $i \in \{1, 2\}$ **do**
- 3: Set $v_j^i = 0$ for all j with $\sigma(j) \leq 2$ and $\sigma(j) \neq i$.
- 4: Set $v_j^i = a_j$ if $\sigma(j) = i$.
- 5: Otherwise, choose $v_j^i \in \{0, a_j\}$ uniformly at random.
- 6: **end for**
- 7: Place positive bids at $\mathbf{v}^1, \mathbf{v}^2$.
- 8: Place a negative bid at the coefficient-wise maximum $\mathbf{v}^1 \wedge \mathbf{v}^2$.
- 9: Let $J = \{j : v_j^1 \neq v_j^2\}$. Place a positive bid at $\mathbf{v}^1 \wedge \mathbf{v}^2 + a_0 \mathbf{e}_J$.
- 10: Assign weight w to all these bids and shift them by \mathbf{c} .

A.8 Experimental results

See Table 2.

Table 2 Runtimes of the DC- and the steepest descent (SD)-algorithm. For each experimental setting, we generated 15 sample auctions. The indicated runtimes are the averages over the respective 15 samples

#pos. bids	#neg. bids	#goods	Time DC (ms)	Time SD (ms)
1020	20	10	31	394
1020	20	20	63	352
1020	20	30	105	410
1020	20	40	133	505
1020	20	50	206	665
1200	200	10	60	502
1200	200	20	157	478
1200	200	30	288	522
1200	200	40	453	620
1200	200	50	597	791
1500	500	10	128	649
1500	500	20	313	664

Table 2 continued

#pos. bids	#neg. bids	#goods	Time DC (ms)	Time SD (ms)
1500	500	30	562	682
1500	500	40	916	775
1500	500	50	1163	962
3020	20	10	115	1152
3020	20	20	252	1116
3020	20	30	445	1278
3020	20	40	592	1225
3020	20	50	970	1366
3200	200	10	175	1303
3200	200	20	413	1236
3200	200	30	628	1226
3200	200	40	1082	1421
3200	200	50	1649	1588
3500	500	10	244	1620
3500	500	20	606	1559
3500	500	30	983	1509
3500	500	40	1642	1575
3500	500	50	2803	1919

References

1. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: Network flows: theory, algorithms, and applications. Prentice hall, USA (1993)
2. An, L.T.H., Tao, P.D.: The dc (difference of convex functions) programming and dca revisited with dc models of real world nonconvex optimization problems. *Ann. Oper. Res.* **133**(1), 23–46 (2005). <https://doi.org/10.1007/s10479-004-5022-1>
3. Ausubel, L.M., Milgrom, P.R.: The lovely but lonely Vickrey auction. In: Cramton, P., Shoham, Y., Steinberg, R. (eds.) *Combinatorial Auctions*, chap. 1, pp. 17–40, MIT Press, Cambridge, MA (2006). <https://doi.org/10.7551/mitpress/9780262033428.001.0001>
4. Ausubel, L.M.: An efficient dynamic auction for heterogeneous commodities. *Am. Econ. Rev.* **96**(3), 602–629 (2006). <https://doi.org/10.1257/aer.96.3.602>
5. Baldwin, E., Goldberg, P.W., Klemperer, P.: Valid combinations of bids. Tech. rep. Oxford University (2016)
6. Baldwin, E., Goldberg, P.W., Klemperer, P., Lock, E.: Solving strong-substitutes product-mix auctions. arXiv preprint [arXiv:1909.07313](https://arxiv.org/abs/1909.07313) (2019)
7. Baldwin, E., Klemperer, P.: Proof that any strong substitutes preferences can be represented by the strong substitutes product-mix auction bidding language. Tech. rep., (day 2 slides on Tropical intersections and equilibrium, Baldwin, E., Goldberg, P. and Klemperer, P.) Hausdorff School on Tropical Geometry and Economics (2016), <http://people.math.gatech.edu/~jyu67/HCM/Baldwin2.pdf>
8. Baldwin, E., Klemperer, P.: Understanding preferences: “demand types”, and the existence of equilibrium with indivisibilities. *Econometrica* **87**(3), 867–932 (2019). <https://doi.org/10.3982/ECTA13693>
9. Baldwin, E., Klemperer, P.: Proof that the product-mix auction bidding language can represent any substitutes preferences. Working Paper 2021–W05. Nuffield College (2021)
10. Beck, M., Robins, S.: Computing the continuous discretely—integer-point enumeration in polyhedra. Springer, Berlin Heidelberg (2007). <https://doi.org/10.1007/978-0-387-46112-0>

11. Bichler, M., Fichtl, M., Schwarz, G.: Walrasian equilibria from an optimization perspective: a guide to the literature. *Naval Research Logistics (NRL)*. **68**(4), 496–513 (2021). <https://doi.org/10.1002/nav.21963>
12. Bichler, M., Shabalin, P., Ziegler, G.: Efficiency with linear prices? A game-theoretical and computational analysis of the combinatorial clock auction. *Information Syst. Res.* **24**(2), 394–417 (2013). <https://doi.org/10.1287/isre.1120.0426>
13. Bikhchandani, S., Mamer, J.W.: Competitive equilibrium in an exchange economy with indivisibilities. *J. Econ. Theory* **74**(2), 385–413 (1997). <https://doi.org/10.1006/jeth.1996.2269>
14. Bikhchandani, S., Ostroy, J.M.: The package assignment model. *J. Econ. Theory* **107**(2), 377–406 (2002). <https://doi.org/10.1006/jeth.2001.2957>
15. Chakrabarty, D., Jain, P., Kothari, P.: Provable submodular minimization using Wolfe’s algorithm. In: *Proceedings of the 27th International Conference on Neural Information Processing Systems, NIPS’ 14*, vol. 1, pp. 802–809. MIT Press, Cambridge, MA, USA (2014). <https://doi.org/10.48550/arXiv.1411.0095>
16. Chakrabarty, D., Lee, Y.T., Sidford, A., Wong, S.C.-w.: Subquadratic submodular function minimization. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1220–1231, STOC 2017, Association for Computing Machinery, New York, NY, USA (2017). ISBN 9781450345286. <https://doi.org/10.1145/3055399.3055419>
17. Danilov, V., Koshevoy, G., Murota, K.: Discrete convexity and equilibria in economies with indivisible goods and money. *Math. Soc. Sci.* **41**(3), 251–273 (2001)
18. Fichtl, M.: On the expressiveness of assignment messages. *Economics Letters* **208**, 110051 (2021). <https://doi.org/10.1016/j.econlet.2021.110051>
19. Goetzendorff, A., Bichler, M., Shabalin, P., Day, R.W.: Compact bid languages and core pricing in large multi-item auctions. *Manag. Sci.* **61**(7), 1684–1703 (2015). <https://doi.org/10.1287/mnsc.2014.2076>
20. Goldberg, P.W., Lock, E., Marmolejo-Cossío, F.J.: Learning strong substitutes demand via queries. In: Chen, X., Gravin, N., Hoefer, M., Mehta, R. (eds.) *Web and Internet Economics - 16th International Conference, WINE*, Lecture Notes in Computer Science, vol. 12495, pp. 401–415, Springer (2020). https://doi.org/10.1007/978-3-030-64946-3_28
21. Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. *J. Econ. Theory* **87**, 95–124 (1999). <https://doi.org/10.1006/jeth.1999.2531>
22. Hatfield, J.W., Milgrom, P.R.: Matching with contracts. *Am. Econ. Rev.* **95**(4), 913–935 (2005). <https://doi.org/10.1257/0002828054825466>
23. Hauser, A.: The future of repo: ‘too much’ or ‘too little’? (June 2013), <https://www.bankofengland.co.uk/-/media/boe/files/speech/2013/the-future-of-repo-too-much-or-too-little>, speech at the ICMA Conference on the Future of the Repo Market
24. Horst, R., Thoai, N.V.: Dc programming: overview. *J. Opt. Theory Appl.* **103**(1), 1–43 (1999). <https://doi.org/10.1023/A:1021765131316>
25. Kelso, A.S., Crawford, V.P.: Job matching, coalition formation, and gross substitutes. *Econometrica* **50**, 1483–1504 (1982). <https://doi.org/10.2307/1913392>
26. Klemperer, P.: A new auction for substitutes: central bank liquidity auctions, the U.S. TARP, and variable product-mix auctions. Tech. rep. Oxford University, Oxford (2008)
27. Klemperer, P.: The product-mix auction: a new auction design for differentiated goods. *J. Eur. Econ. Assoc.* **8**(2–3), 526–536 (2010). <https://doi.org/10.1111/j.1542-4774.2010.tb00523.x>
28. Klemperer, P.: Product-mix auctions. Working Paper 2018-W07. Nuffield College (2018)
29. Kobayashi, Y.: The complexity of minimizing the difference of two M-convex set functions. *Oper. Res. Lett.* **43**, 573–574 (2015). <https://doi.org/10.1016/j.orl.2015.08.011>
30. Lee, Y.T., Sidford, A., Wong, S.C.-w.: A faster cutting plane method and its implications for combinatorial and convex optimization. In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pp. 1049–1065 (10 2015), <https://doi.org/10.1109/FOCS.2015.68>
31. Lin, B., Tran, N.M.: Linear and rational factorization of tropical polynomials. arXiv preprint [arXiv:1707.03332](https://arxiv.org/abs/1707.03332) (2017)
32. Maehara, T., Marumo, N., Murota, K.: Continuous relaxation for discrete dc programming. *Math. Program.* **169**(1), 199–219 (2018). <https://doi.org/10.1007/s10107-017-1139-2>
33. Maehara, T., Murota, K.: A framework of discrete dc programming by discrete convex analysis. *Math. Program.* **152**(1), 435–466 (2015). <https://doi.org/10.1007/s10107-014-0792-y>

34. Milgrom, P.R.: Assignment messages and exchanges. *Am. Econ. J. Microecon.* **1**(2), 95–113 (2009). <https://doi.org/10.1257/mic.1.2.95>
35. Milgrom, P.R., Strulovici, B.: Substitute goods, auctions, and equilibrium. *J. Econ. Theory* **144**(1), 212–247 (2009). <https://doi.org/10.1016/j.jet.2008.05.002>
36. Murota, K.: *Discrete convex analysis: Monographs on discrete mathematics and applications*, vol. 10. Society for Industrial and Applied Mathematics, USA (2003). <https://doi.org/10.1137/1.9780898718508>
37. Murota, K.: Discrete convex analysis: a tool for economics and game theory. *J. Mech. Inst. Des.* **1**(1), 151–273 (2016). <https://doi.org/10.22574/jmid.2016.12.005>
38. Murota, K., Tamura, A.: New characterizations of M-convex functions and their applications to economic equilibrium models with indivisibilities. *Discrete Appl. Math.* **131**(2), 495–512 (2003). [https://doi.org/10.1016/S0166-218X\(02\)00469-9](https://doi.org/10.1016/S0166-218X(02)00469-9)
39. Nisan, N., Segal, I.: The communication requirements of efficient allocations and supporting prices. *J. Econ. Theory* **129**(1), 192–224 (2006). <https://doi.org/10.1016/j.jet.2004.10.007>
40. Ostrovsky, M., Paes Leme, R.: Gross substitutes and endowed assignment valuations. *Theor. Econ.* **10**(3), 853–865 (2015). <https://doi.org/10.3982/TE1840>
41. Paes Leme, R.: Gross substitutability: an algorithmic survey. *Games Econ. Behav.* **106**, 294–316 (2017). <https://doi.org/10.1016/j.geb.2017.10.016>
42. Paes Leme, R., Wong, S.C.-w.: Computing Walrasian equilibria: fast algorithms and structural properties. *Math. Program.* **179**(1), 343–384 (2020). <https://doi.org/10.1007/S10107-018-1334-9>
43. Rockafellar, R.T.: *Convex analysis*, princeton mathematical series. Princeton University Press, Princeton, NJ (1970). ISBN 0691015864. <https://doi.org/10.1515/9781400873173>
44. Schneider, R.: *Convex bodies: the Brunn-Minkowski theory*, 2nd edn. Cambridge University Press, Cambridge (2013). <https://doi.org/10.1017/CBO9781139003858>
45. Shioura, A., Tamura, A.: Gross substitutes condition and discrete concavity for multi-unit valuations: a survey. *J. Oper. Res. Soc. Jpn.* **58**(1), 61–103 (2015). <https://doi.org/10.15807/jorsj.58.61>
46. Tao, P.D., An, L.T.H.: Convex analysis approach to dc programming: theory, algorithms and applications. *Acta Mathematica Vietnamica* **22**(1), 289–355 (1997)
47. Toland, J.F.: A duality principle for non-convex optimisation and the calculus of variations. *Arch. Ration. Mech. Anal.* **71**(1), 41–61 (1979). <https://doi.org/10.1007/BF00250669>
48. Tran, N.M.: The finite matroid-based valuation conjecture is false. *SIAM J. Appl. Algebra Geom.* **5**(3), 506–525 (2021). <https://doi.org/10.1137/19M1304295>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

6 Discussion and Conclusion

Due to their preferable properties, Walrasian equilibria are the desired outcomes in many markets and auctions. In real-world applications, it is essential that auctions can be conducted in a time-efficient way, raising the problems of efficiently communicating preferences and computing an outcome of the auction. These two central problems are the main subject of study in this dissertation. We presented scientific contributions on the existence of Walrasian equilibria, the analysis of bid languages, and how we can leverage a specific bid language to design new efficient algorithms.

In the first publication included in this dissertation, we survey existing conditions on the bidders' preferences that guarantee the existence of a Walrasian equilibrium and explain how these conditions are related to notions of discrete convexity. For the case of gross substitutes bidders, we give a new interpretation of the iterative ascending auction of Ausubel (2006) in terms of the primal-dual algorithm. Publication 1 also points to three major open questions in competitive equilibrium theory.

First, usual assumptions on the bidders' valuation functions, like the gross substitutes condition, to guarantee the existence of a Walrasian equilibrium are often not satisfied in real-world markets. It is, therefore, necessary to come up with generalized, weaker notions of equilibria to deal with such markets. Connected to this problem is the design of new iterative auctions for computing these outcomes.

The quasi-linear utility model is the second strong standard assumption of competitive equilibrium theory that is usually not met in practice. Quasi-linear utilities imply, in particular, that bidders do not face any financial constraints: no matter how high the price is, bidders can always afford their preferred bundle. If we deviate from this model and assume that bidders face a hard upper limit on how much money they can spend, it is easy to construct examples where a competitive equilibrium does not exist – even in one of the most basic market environments, namely assignment markets. Recent results suggest that computing outcomes that satisfy a subset of the properties of Walrasian

equilibria – namely core-stability and welfare-maximization – are computationally very hard: in Batziou et al. (2022) we prove that this problem is already NP-complete in assignment markets, while Bichler and Waldherr (2022) show that the very general setting of combinatorial exchanges, where bidders are allowed to have arbitrary valuation functions, is even Σ_2^P -hard.

On the other hand, the question of equilibrium existence has a positive answer in settings with *weak* budget constraints (Baldwin et al. (accepted) 2022), under conditions on the valuation functions that are generalizations to those from the quasi-linear setting. In the model with weak budget constraints, a bidder’s utility for a bundle decreases continuously to $-\infty$ as the price approaches the budget constraint, in contrast to hard budget constraints, where the utility has a jump discontinuity as soon as the budget is exceeded. However, no algorithms have been developed for computing competitive equilibria in this setting, and little is known about the computational complexity of this problem.

Third, incentive-compatibility is a major issue, particularly in smaller markets. As was demonstrated in Section 2.3, Walrasian equilibria can generally not be implemented by a strategyproof mechanism. With quasi-linear utilities, it is well-known that the Vickrey-Clarke-Groves (VCG) mechanism is essentially the unique strategyproof mechanism yielding a welfare-maximizing assignment. In contrast to Walrasian equilibria, the VCG mechanism uses non-linear and non-anonymous prices. Several variants of iterative auctions have been proposed for computing VCG-outcomes, see for example Ausubel (2006), Baranov (2018), and de Vries et al. (2007). VCG-outcomes, however, have their own disadvantages. For example, they are prone to collusion and bid shilling. They also often provide low revenue for the seller – in particular, the revenue is generally non-monotone in the number of participating bidders (Ausubel and Milgrom 2006). As mentioned in Section 2.3, an alternative approach is to consider relaxed notions of strategyproofness for large markets, as done by Azevedo and Budish (2018), Jackson and Manelli (1997), and Roberts and Postlewaite (1976). We think that both questions – the design of incentive-compatible iterative auctions, and developing weaker notions of strategyproofness – are far from being answered satisfactorily. This is particularly the case when considered jointly with other deviations from the standard assumptions, like financially constrained bidders.

In Publication 2, we study the expressiveness of Assignment Messages, a bid language for expressing strong substitutes preferences. We prove that there are strong substitutes

preferences that cannot be expressed in this bid language. This is done by giving a new interpretation of Assignment Messages as min-cost flow problems, leading to the conclusion that all valuations described by Assignment Messages satisfy a strictly stronger property called strong exchangeability.

Our result naturally raises the question of the exact relation between Assignment Messages and strong substitutes valuations. It would be interesting to understand if Assignment Messages form a “natural” subset of strong substitutes. In particular, it is unknown whether Assignment Messages can express all strongly exchangeable valuation functions. Connected to this, we might also ask whether there is a generalization of Assignment Messages capable of expressing a larger set of strong substitutes preferences. A related question concerns the translation of Assignment Messages to SSPMA bids. Note that the reverse direction is impossible for arbitrary SSPMA bids since they can express every strong substitutes valuation function. Whether Assignment Messages or SSPMA bids are more intuitive for a bidder is likely to depend on the particular structure of their valuation function. Thus, a tool for translating one bid language into the other can ease the bidders’ process of reporting their valuations. There exist algorithms that convert arbitrary strong substitutes valuation functions into the corresponding SSPMA bids (Lock et al. 2022), which can also be used to convert Assignment Messages. These algorithms are, however, notoriously slow, and we may hope for significantly faster algorithms when taking into account the structure of Assignment Messages.

In Publication 3, we develop a novel algorithm for computing Walrasian equilibrium prices when bidders express their preferences via SSPMA bids. The algorithm is based on a novel duality result, relating the Lyapunov function to the welfare difference of the positive-bids and negative-bids auction. While the observed number of iterations of the algorithm until convergence is very low, it turns out to be very hard to give a good theoretical worst-case estimate. Note that minimizing the difference of two general strong substitutes valuation functions is NP-hard (Kobayashi 2015), so a good upper bound for our problem would necessarily strongly depend on the additional structure provided by the SSPMA bid language. We believe that research in this direction will deepen the understanding of the strong substitutes condition in general.

As part of this Ph.D. project, the author also conducted research on learning Bayes-Nash equilibria in auction games (Bichler et al. 2021a; Fichtl et al. 2022). The developed algorithms show excellent experimental performance in a large class of practically relevant auction games. However, there are almost no theoretical results regarding the conver-

gence of these algorithms. The only known results consider second-price auctions and first-price auctions with two bidders and particular assumptions on the bidders' prior distributions (Feng et al. 2020). Our experiments suggest that convergence holds for much more general auction settings, and we conjecture that it holds indeed globally in many scenarios. Besides theoretical questions, there is also space for algorithmic improvements of our methods. While the SODA algorithm outperforms NPGA in instances with few (up to 3 – 4) bidders, the execution time becomes infeasible for larger numbers of bidders and interdependent prior distributions. Natural approaches for improving the algorithms running time include adaptive discretization of the bidders' strategy spaces and a sampling-based method for approximating the gradients.

References

- Arrow, K. J. and Debreu, G. (1954). Existence of an Equilibrium for a Competitive Economy. In: *Econometrica* 22.3, p. 265. DOI: 10.2307/1907353.
- Ausubel, L. M. and Milgrom, P. (2006). The Lovely but Lonely Vickrey Auction. In: *Combinatorial Auctions*. Cambridge, MA: MIT Press. Chap. 1, pp. 17–40. DOI: 10.7551/mitpress/9780262033428.003.0002.
- Ausubel, L. M. (2006). An Efficient Dynamic Auction for Heterogeneous Commodities. In: *American Economic Review* 96.3, pp. 602–629. DOI: 10.1257/AER.96.3.602.
- Azevedo, E. M. and Budish, E. (2018). Strategy-Proofness in the Large. In: *The Review of Economic Studies* 86.1, pp. 81–116. DOI: 10.1093/restud/rdy042.
- Baldwin, E., Bichler, M., Fichtl, M., and Klemperer, P. (2022). Strong substitutes: structural properties, and a new algorithm for competitive equilibrium prices. In: *Mathematical Programming*. DOI: 10.1007/s10107-022-01792-w.
- Baldwin, E., Goldberg, P. W., Klemperer, P., and Lock, E. (2020). Solving Strong-Substitutes Product-Mix Auctions. CEPR Discussion Paper No. DP14976.
- Baldwin, E., Jagadeesan, R., Klemperer, P., and Teytelboym, A. ((accepted) 2022). The Equilibrium Existence Duality. In: *Journal of Political Economy*, to appear. DOI: 10.1086/722935.
- Baldwin, E. and Klemperer, P. (2019). Understanding Preferences: “Demand Types”, and the Existence of Equilibrium With Indivisibilities. In: *Econometrica* 87.3, pp. 867–932. DOI: 10.3982/ECTA13693.
- Baldwin, E. and Klemperer, P. (2021). Proof that the Strong Substitutes Product-Mix Auction Bidding Language can represent any Strong Substitutes preferences. URL: <https://elizabeth-baldwin.me.uk/papers/strongsubsproof.pdf>.
- Baranov, O. (2018). An efficient ascending auction for private valuations. In: *Journal of Economic Theory* 177, pp. 495–517. DOI: 10.1016/j.jet.2018.07.005.
- Batziau, E., Bichler, M., and Fichtl, M. (2022). Core-Stability in Assignment Markets with Financially Constrained Buyers. In: *Proceedings of the 23rd ACM Conference on Economics and Computation* 2.22, pp. 473–474. DOI: 10.1145/3490486.3538262.

- Bichler, M. (2017). *Market Design: A Linear Programming Approach to Auctions and Matching*. Cambridge University Press. DOI: 10.1017/9781316779873.
- Bichler, M., Davenport, A., Hohner, G., and Kalagnanam, J. (2005). Industrial Procurement Auctions. In: *Combinatorial auctions*. MIT Press, pp. 593–612. DOI: 10.7551/mitpress/9780262033428.003.0024.
- Bichler, M., Fichtl, M., Heidekrüger, S., Kohring, N., and Sutterer, P. (2021a). Learning Equilibria in Symmetric Auction Games Using Artificial Neural Networks. In: *Nature Machine Intelligence* 3.8, pp. 687–695. DOI: 10.1038/s42256-021-00365-4.
- Bichler, M., Fichtl, M., and Schwarz, G. (2021b). Walrasian equilibria from an optimization perspective: A guide to the literature. In: *Naval Research Logistics (NRL)* 68.4, pp. 496–513. DOI: <https://doi.org/10.1002/nav.21963>.
- Bichler, M., Goeree, J., Mayer, S., and Shabalin, P. (2014). Spectrum auction design: Simple auctions for complex sales. In: *Telecommunications Policy* 38.7, pp. 613–622. DOI: 10.1016/j.telpol.2014.02.004.
- Bichler, M. and Waldherr, S. (2022). Core-Pricing in Combinatorial Exchanges with Financially Constrained Buyers: Computational Hardness and Algorithmic Solutions. In: *Operations Research* 70.1, pp. 241–264. DOI: 10.1287/opre.2021.2132.
- Bikhchandani, S. and Mamer, J. W. (1997). Competitive Equilibrium in an Exchange Economy with Indivisibilities. In: *Journal of Economic Theory* 74.2, pp. 385–413. DOI: 10.1006/JETH.1996.2269.
- Blumrosen, L. and Nisan, N. (2007). Combinatorial Auctions. In: *Algorithmic Game Theory*. Cambridge: Cambridge University Press, pp. 267–300. DOI: 10.1017/CB09780511800481.013.
- Cantillon, E., Pesendorfer, M., et al. (2005). Auctioning bus routes: The London experience. In: *Combinatorial auctions*. MIT Press, pp. 573–592. DOI: 10.7551/mitpress/9780262033428.003.0023.
- Caplice, C. and Sheffi, Y. (2003). Optimization-Based Procurement for Transportation Services. In: *Journal of Business Logistics* 24.2, pp. 109–128. DOI: 10.1002/j.2158-1592.2003.tb00048.x.
- Chakrabarty, D., Lee, Y. T., Sidford, A., and Wong, S. C.-w. (2017). Subquadratic Submodular Function Minimization. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2017. New York, NY, USA: Association for Computing Machinery, pp. 1220–1231. DOI: 10.1145/3055399.3055419.

-
- Committee for the Prize in Economic Sciences in Memory of Alfred Nobel (2020). Improvements to auction theory and inventions of new auction formats. URL: <https://www.nobelprize.org/prizes/economic-sciences/2020/press-release/>.
- Daskalakis, C., Frongillo, R., Papadimitriou, C., Pierrakos, G., and Valiant, G. (2010). On Learning Algorithms for Nash Equilibria. In: pp. 114–125. DOI: 10.1007/978-3-642-16170-4_11.
- Daskalakis, C., Goldberg, P. W., and Papadimitriou, C. H. (2009). The Complexity of Computing a Nash Equilibrium. In: *SIAM Journal on Computing* 39.1, pp. 195–259. DOI: 10.1137/070699652.
- de Vries, S., Schummer, J., and Vohra, R. V. (2007). On ascending Vickrey auctions for heterogeneous objects. In: *Journal of Economic Theory* 132.1, pp. 95–118. DOI: 10.1016/j.jet.2005.07.010.
- Federal Communications Commission (2022). Auction 110: 3.45 GHz Service. Tech. rep. URL: <https://www.fcc.gov/auction/110>.
- Feng, Z., Guruganesh, G., Liaw, C., Mehta, A., and Sethi, A. (2020). Convergence Analysis of No-Regret Bidding Algorithms in Repeated Auctions. In: *AAAI Conference on Artificial Intelligence*. DOI: 10.1609/aaai.v35i6.16680.
- Fichtl, M. (2021). On the expressiveness of assignment messages. In: *Economics Letters* 208, p. 110051. DOI: <https://doi.org/10.1016/j.econlet.2021.110051>.
- Fichtl, M., Oberlechner, M., and Bichler, M. (2022). Computing Bayes Nash Equilibrium Strategies in Auction Games via Simultaneous Online Dual Averaging. *Earlier versions presented at the AAAI-RLG 22 and the NeurIPS21 workshops*. DOI: 10.48550/ARXIV.2208.02036.
- Fujishige, S. and Murota, K. (2000). Notes on L-/M-convex functions and the separation theorems. In: *Mathematical Programming* 88.1, pp. 129–146. DOI: 10.1007/PL00011371.
- Fujishige, S. and Yang, Z. (2003). A Note on Kelso and Crawford’s Gross Substitutes Condition. In: *Mathematics of Operations Research* 28.3, pp. 463–469. DOI: 10.1287/moor.28.3.463.16393.
- Goetzendorff, A., Bichler, M., Shabalin, P., and Day, R. W. (2015). Compact Bid Languages and Core Pricing in Large Multi-item Auctions. In: *Management Science* 61.7, pp. 1684–1703. DOI: 10.1287/mnsc.2014.20760.
- Gul, F. and Stacchetti, E. (1999). Walrasian Equilibrium with Gross Substitutes. In: *Journal of Economic Theory* 87.1, pp. 95–124. DOI: 10.1006/JETH.1999.2531.

- Hatfield, J. W. and Milgrom, P. (2005). Matching with Contracts. In: *American Economic Review* 95.4, pp. 913–935. DOI: 10.1257/0002828054825466.
- Ikebe, Y. T. and Tamura, A. (2015). Stability in Supply Chain Networks: An Approach by Discrete Convex Analysis. In: *Journal of the Operations Research Society of Japan* 58.3, pp. 271–290. DOI: 10.15807/jorsj.58.271.
- Iwata, S., Fleischer, L., and Fujishige, S. (2001). A Combinatorial Strongly Polynomial Algorithm for Minimizing Submodular Functions. In: *J. ACM* 48.4, pp. 761–777. DOI: 10.1145/502090.502096.
- Jackson, M. O. and Manelli, A. M. (1997). Approximately Competitive Equilibria in Large Finite Economies. In: *Journal of Economic Theory* 77.2, pp. 354–376. DOI: 10.1006/jeth.1997.2327.
- Kaneko, M. (1976). On the Core and Competitive Equilibria of a Market with Indivisible Goods. In: *Naval Research Logistics Quarterly* 23.2, pp. 321–337. DOI: 10.1002/NAV.3800230214.
- Kelso, A. S. and Crawford, V. P. (1982). Job Matching, Coalition Formation, and Gross Substitutes. In: *Econometrica* 50.6, p. 1483. DOI: 10.2307/1913392.
- Klemperer, P. (2008). A new auction for substitutes: Central Bank liquidity auctions, the U.S. TARP, and variable product-mix auctions. Tech. rep. 1. Oxford University.
- Klemperer, P. (2010). The Product-Mix Auction: A New Auction Design for Differentiated Goods. In: *Journal of the European Economic Association* 8.2-3, pp. 526–536. DOI: 10.1111/j.1542-4774.2010.tb00523.x.
- Kobayashi, Y. (2015). The complexity of minimizing the difference of two M^h -convex set functions. In: *Operations Research Letters* 43.6, pp. 573–574. DOI: 10.1016/j.orl.2015.08.011.
- Lee, Y. T., Sidford, A., and Wong, S. C.-w. (2015). A Faster Cutting Plane Method and its Implications for Combinatorial and Convex Optimization. In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pp. 1049–1065. DOI: 10.1109/FOCS.2015.68.
- Leme, R. P. (2017). Gross Substitutability: An Algorithmic Survey. In: *Games and Economic Behavior* 106, pp. 294–316. DOI: 10.1016/J.GEB.2017.10.016.
- Leme, R. P. and Wong, S. C.-w. (2020). Computing Walrasian equilibria: fast algorithms and structural properties. In: *Mathematical Programming* 179.1, pp. 343–384. DOI: 10.1007/s10107-018-1334-9.

-
- Liberopoulos, G. and Andrianesis, P. (2016). Critical Review of Pricing Schemes in Markets with Non-Convex Costs. In: *Operations Research* 64.1, pp. 17–31. DOI: 10.1287/opre.2015.1451.
- Lock, E., Goldberg, P. W., and Marmolejo-Cossío, F. (2022). Learning Strong Substitutes Demand via Queries. In: *ACM Trans. Econ. Comput.* 10.2. DOI: 10.1145/3546604.
- Ma, J. (1998). Competitive Equilibrium with Indivisibilities. In: *Journal of Economic Theory* 82.2, pp. 458–468. DOI: 10.1006/jeth.1998.2445.
- McKenzie, L. W. (1959). On the Existence of General Equilibrium for a Competitive Market. In: *Econometrica* 27.1, p. 54. DOI: 10.2307/1907777.
- Milgrom, P. (2009). Assignment Messages and Exchanges. In: *American Economic Journal: Microeconomics* 1.2, pp. 95–113. DOI: 10.1257/MIC.1.2.95.
- Milgrom, P. and Strulovici, B. (2009). Substitute Goods, Auctions, and Equilibrium. In: *Journal of Economic Theory* 144.1, pp. 212–247. DOI: 10.1016/J.JET.2008.05.002.
- Milgrom, P. and Weber, R. J. (1985). Distributional Strategies for Games with Incomplete Information. In: *Mathematics of Operations Research* 10.4, pp. 619–632. DOI: 10.1287/moor.10.4.619.
- Murota, K. (2003). *Discrete Convex Analysis*. Society for Industrial and Applied Mathematics. DOI: 10.1137/1.9780898718508.
- Murota, K. (2016). *Discrete Convex Analysis: A Tool for Economics and Game Theory*. In: *Journal of Mechanism and Institution Design* 1.1, pp. 149–271. DOI: 10.22574/JMID.2016.12.005.
- Murota, K. and Shioura, A. (1999). M-Convex Function on Generalized Polymatroid. In: *Mathematics of Operations Research* 24.1, pp. 95–105. DOI: 10.1287/moor.24.1.95.
- Nisan, N. (2005). *Bidding Languages for Combinatorial Auctions*. In: *Combinatorial Auctions*. The MIT Press. DOI: 10.7551/mitpress/9780262033428.003.0010.
- Ostrovsky, M. and Leme, R. P. (2015). Gross Substitutes and Endowed Assignment Valuations. In: *Theoretical Economics* 10.3, pp. 853–865. DOI: 10.3982/TE1840.
- Parkes, D. C. (2005). *Iterative Combinatorial Auctions*. In: *Combinatorial Auctions*. The MIT Press. DOI: 10.7551/mitpress/9780262033428.003.0003.
- Roberts, D. J. and Postlewaite, A. (1976). The Incentives for Price-Taking Behavior in Large Exchange Economies. In: *Econometrica*, pp. 115–127. DOI: 10.2307/1911385.
- Schrijver, A. (2000). A Combinatorial Algorithm Minimizing Submodular Functions in Strongly Polynomial Time. In: *J. Comb. Theory, Ser. B* 80, pp. 346–355. DOI: 10.1006/jctb.2000.1989.

- Segal, I. (2005). The Communication Requirements of Combinatorial Allocation Problems. In: *Combinatorial Auctions*. The MIT Press. DOI: 10.7551/mitpress/9780262033428.003.0012.
- Shioura, A. and Tamura, A. (2015). Gross Substitutes Condition and Discrete Concavity for Multi-Unit Valuations: A Survey. In: *Journal of the Operations Research Society of Japan* 58.1, pp. 61–103. DOI: 10.15807/jorsj.58.61.
- Shioura, A. (2017). Algorithms for L-convex function minimization: Connection between discrete convex analysis and other research fields. In: *Journal of the Operations Research Society of Japan* 60.3, pp. 216–243. DOI: 10.15807/jorsj.60.216.
- Sun, N. and Yang, Z. (2006). Equilibria and Indivisibilities: Gross Substitutes and Complements. In: *Econometrica* 74.5, pp. 1385–1402. DOI: 10.1111/J.1468-0262.2006.00708.X.
- Teytelboym, A. (2014). Gross Substitutes and Complements: A Simple Generalization. In: *Economics Letters* 123.2, pp. 135–138. DOI: 10.1016/J.ECONLET.2014.01.038.
- Vlatakis-Gkaragkounis, E.-V., Flokas, L., Lianas, T., Mertikopoulos, P., and Piliouras, G. (2020). No-Regret Learning and Mixed Nash Equilibria: They Do Not Mix. In: *Advances in Neural Information Processing Systems*. Vol. 33. Curran Associates, Inc., pp. 1380–1391.
- Walras, L. (1874). Éléments d'économie politique pure. In: *Revue de Théologie et de Philosophie et Compte-rendu des Principales Publications Scientifiques* 7, pp. 628–632.