

Solving Strong-Substitutes Product-Mix Auctions

Elizabeth Baldwin* Paul W. Goldberg† Paul Klemperer‡ Edwin Lock§

5th July 2023

Abstract

This paper develops algorithms to solve strong-substitutes product-mix auctions: it finds competitive equilibrium prices and quantities for agents who use this auction’s bidding language to truthfully express their strong-substitutes preferences over an arbitrary number of goods, each of which is available in multiple discrete units. Our use of the bidding language, and the information it provides, contrasts with existing algorithms that rely on access to a valuation or demand oracle.

We compute market-clearing prices using algorithms that apply existing submodular minimisation methods. Allocating the supply among the bidders at these prices then requires solving a novel constrained matching problem. Our algorithm iteratively simplifies the allocation problem, perturbing bids and prices in a way that resolves tie-breaking choices created by bids that can be accepted on more than one good. We provide practical running time bounds on both price-finding and allocation, and illustrate experimentally that our allocation mechanism is practical.

Keywords: bidding language, product-mix auction, competitive equilibrium, Walrasian equilibrium, convex optimisation, strong substitutes, submodular minimisation

1 Introduction

This paper develops algorithms that solve product-mix auctions in which participants can make bids that represent any strong-substitutes preferences for an arbitrary number of distinct goods. (These preferences are also known, in other literatures, as M^{\natural} -concave, as matroidal, and as well-layered maps). It thus allows bidders to express more general preferences than could previously be permitted in these auctions, and finds competitive equilibrium prices and quantities consistent with these, and the auctioneer’s preferences.

Importantly, our algorithms for finding equilibrium differ from existing ones in that they directly use the information that the product-mix auction ‘language’ provides. This information is in a very different form from the information provided by a valuation or demand oracle. This creates additional complexities, as well as simplifications which we can exploit. However, the language is conceptually simple, and easy for bidders to use in a (product-mix) auction.

The *product-mix auction* was developed in 2007-8 for the Bank of England to provide liquidity to financial institutions by auctioning loans to them [Klemperer, 2008]. It is now used at least monthly by the Bank, and more often when institutions are more likely to be under stress. (After the 2016 vote for ‘Brexit’, and starting again in March 2019, for instance, the auction was run weekly.)

*Dept. of Economics, Oxford University, UK, elizabeth.baldwin@economics.ox.ac.uk.

†Dept. of Computer Science, Oxford University, UK, paul.goldberg@cs.ox.ac.uk.

‡Dept. of Economics, Oxford University, UK, paul.klemperer@nuffield.ox.ac.uk.

§Dept. of Computer Science, Oxford University, UK, edwin.lock@cs.ox.ac.uk.




Figure 1: Price vectors divide \mathbb{R}^n into regions that specify which good is allocated to bids lying in these region. This example works in the setting with two goods. The price vector $\mathbf{p} = (4, 4)$ (marked by a cross) divides \mathbb{R}^2 into three regions, each of which labelled according to the good that bids in this region are allocated.

The original implementation of the auction allowed bidders to each submit a list of bids, where each bid has $n + 1$ elements: a price for each of the n goods available, and a total quantity of goods sought by that bid.¹ The auction sets a uniform price for each good (i.e. every recipient of any particular good pays the same price per unit for that good), and gives each bid an allocation that maximises that bid’s ‘utility’, assuming preferences with quasilinear utilities. That is, each bid is allocated its desired quantity of the good on which its price strictly exceeds the auction’s price by most, if there is a unique such good, and is allocated nothing if all its prices are strictly below the auction’s corresponding prices. (So if a bid only wants one particular good, it simply sets prices of zero for all the other ($n - 1$) goods.) A bid that creates a tie (i.e. two or more of its prices exceed the auction’s corresponding prices by most, or none of its prices exceed, but at least one equals, the auction’s corresponding prices) can be allocated in any way consistent with equilibrium (see below). Figure 1 illustrates how the good that a bid is allocated depends on the bid’s relative position to the price vector. In particular, we note that individual bids $(b_1, \dots, b_n; +1)$ that seek exactly one item of a good can be interpreted as a unit-demand bidder with valuation vector $\mathbf{b} = (b_1, \dots, b_n)$ and quasilinear utilities. We motivate the bidding language in the following example.

Example 1 (Figure 2a). *Alice participates in an auction held with two goods, say apples (good 1) and bananas (good 2). For her breakfast, she requires a single item of fruit and is willing to pay up to £6 for either an apple or a banana. This is expressed by the bid $\mathbf{b} := (6, 6; +1)$; the last element in the bid denotes the quantity of goods sought. Suppose the auction’s prices are set at $\mathbf{p} = (p_1, p_2)$. If at least one good is priced lower than the respective bid entry, she will demand the good $i \in \{1, 2\}$ that maximises utility $b_i - p_i$. Otherwise, if both apples and bananas are priced higher than her valuation of £6, she demands neither. In addition, as she is fond of bananas, Alice wishes to pick up a banana as a second fruit item, but only if the price of a banana does not exceed £4. This is expressed by the bid $\mathbf{b}^\flat := (0, 4; +1)$.*

To illustrate Alice’s demand, consider the following two possible auction prices. If the auction prices are set at $\mathbf{p} = (1, 3)$, then Alice demands one apple and one banana: \mathbf{b} demands an apple, as its utility for an apple or a banana is 5 and 3, respectively, while \mathbf{b}^\flat demands a banana, as its price is less than £4. On other hand, if auction prices are $\mathbf{p} = (6, 5)$, Alice demands only one banana: \mathbf{b} demands a banana and \mathbf{b}^\flat demands nothing, as the price of a banana is too high. Alice’s two bids \mathbf{b} and \mathbf{b}^\flat , together with the demand correspondence they induce, are shown in Figure 2a.

¹In the Bank of England’s auction, the bidders are commercial banks, etc., each good is a loan secured against one of n different specified qualities of collateral (so the prices are interest rates), and the quantity is the amount of the loan (in £).




Figure 2: Examples of strong-substitutes demand correspondences on two goods $\{1,2\}$ belonging to two bidders, Alice and Bob. Alice’s demand is specified by bid list $B = f(6, 6; +1), (0, 4; +1)g$ and Bob’s bid list is $B^0 = f(2, 4; +1), (4, 2; +1), (4, 4; -1), (6, 6; +1)g$. Positive and negative bids are depicted as solid and hollow circles, respectively. Price space is divided into regions corresponding to demanded bundles (x_1, x_2) , with x_i denoting the number of items of good i . At $\mathbf{p} = (4, 4)$, Alice demands bundles $f(1, 0), (0, 1), (1, 1), (0, 2)g$ and Bob demands bundles $f(1, 0), (0, 1), (1, 1)g$, the discrete convex hulls of bundles demanded in the regions surrounding \mathbf{p} .

The auction’s prices are chosen to maximise the sum of the bids’ and auctioneer’s welfare – that is, it finds competitive equilibrium prices and quantities.² Moreover, the auctioneer expresses her preferences about which goods to allocate in the form of supply functions which can themselves be equivalently represented as (and thus converted into) a list of bids of the kind described above.³

Positive and negative bids. In versions of the product-mix auction thus far implemented, all bids are for positive quantities of goods. In the case with transferable utilities, the market-clearing prices can then be found by solving straightforward linear programs, and finding an allocation of the auctioneer’s supply to the individual bidders is similarly straightforward. However, some strong-substitutes⁴ preferences (over bundles of non-negative quantities of goods) can

²Specifically, the auction finds the lowest such price vector. Since the bids automatically express “strong substitutes” preferences for all bidders (see [Baldwin and Klemperer, 2021]), there exist equilibrium price vectors, and also a unique one among them at which every good’s price is lowest.

³See Appendix 1E of [Klemperer, 2018] for how to convert supply functions into bid lists. (The bids that the auctioneer’s supply functions would be converted into would ensure that it *sells more* units on a good when prices are high, analogous to a buyer *buying fewer* units in this case.) Although the auctioneer’s preferences *could* equivalently have been represented as a list of bids of the kind made by the bidders, describing them simply as two-dimensional graphs of ‘supply schedules’ was an important feature of the auction design: participants’ ability to express their preferences in the ways that are most natural for them is crucial to getting an auction accepted for practical use, to getting bidders to participate, and to the auction working efficiently.

The Bank of England’s original program restricted to $n = 2$. Since 2014 its program permits much larger n (it is currently being run with $n = 3$), and it also allows richer forms of preferences to be expressed by the auctioneer (but not by the bidders, whose preference expression has, by contrast, been restricted in recent auctions).

A variant that allowed for bidders’ budget constraints (hence non-quasilinear preferences) was programmed for the Government of Iceland in 2015-16. For additional information about the strong-substitutes product-mix auction and other variants of the product-mix auction that allow for divisible goods and budget constraints, etc., we refer to <http://pma.nuff.ox.ac.uk> and [Baldwin and Klemperer, 2021, Klemperer, 2008, 2010, 2018].

⁴ Strong-substitutes preferences are those that would be still be gross substitutes preferences if we treated every unit of every good as a separate good. Such preferences have many attractive properties; they mean, for example, that if the price of any one good increases, and the demand for it decreases, then the demand for all other goods can increase by at most the amount of that decrease. Strong substitutability is equivalent to M^1 -concavity [Shioura and Tamura, 2015] in our setting with valuations defined on non-negative and bounded

only be expressed by using lists of bids for both positive *and negative* quantities,⁵ and it has been shown that *any* strong-substitutes preferences can be represented using lists of positive and negative bids.⁶ The rules for accepting, or not accepting, negative bids are identical to those for positive bids: if a negative bid is accepted at some prices, it is allocated a negative quantity of the good it demands, (partially) cancelling the allocated quantities of positive bids that are accepted at the same prices.

To illustrate why negative bids are useful, we consider the setting with two distinct goods 1 and 2 in which an agent has strong-substitutes valuation u and is interested in at most 2 units. In other words, the agent only has a positive value for bundles $(0, 1)$, $(1, 0)$, $(1, 1)$, $(0, 2)$ and $(2, 0)$, where the 2-dimensional vectors (x_1, x_2) denote bundles containing x_1 units of good 1 and x_2 units of good 2. It is known [Shioura, 2017] that the strong-substitutes condition is equivalent to M^\sharp -concavity which, in our setting, is equivalent to discrete concavity together with the following three additional conditions:

- i) $u((1, 1)) < u((0, 1)) + u((1, 0))$,
- ii) $u((2, 0)) < u((1, 0)) + u((1, 1)) + u((0, 1))$.
- iii) $u((0, 2)) < u((0, 1)) + u((1, 1)) + u((1, 0))$.

It is easy to see (by infinitesimally and independently perturbing bundle values) that the generic case has strict inequalities in all three conditions. Intuitively, a strict inequality in the first condition implies that a single unit of good 1 is valued strictly less if the agent also has a unit of 2 than if he does not, reflecting strict substitutability as distinct from complementarity or independence between the two goods. A strict inequality in the second or third condition implies that the marginal utility of a good is reduced more by receiving an additional unit of the same good than by receiving a unit of a different good; ‘more variety’ is valuable. In contrast, it is easy to check that any demand correspondence generated only by positive bids is non-generic, in that at least one of the three conditions must be an equality. Thus a demand correspondence such as the one illustrated in Example 2 and Figure 2b, in which all three conditions are strict inequalities, cannot be represented by positive bids only. We refer to [Klemperer, 2008, p. 15] and [Klemperer, 2018, Appendix IC] for further details and examples of the usefulness of negative bids.

Example 2 (Figure 2b). *A combination of positive and negative bids is required to express Bob’s demand correspondence depicted in Figure 2b, which divides price space into six demand regions for bundles with at most two units. The valuation function u is given by $u((0, 1)) = u((1, 0)) = 6$, $u((0, 2)) = u((2, 0)) = 8$ and $u((1, 1)) = 10$. The corresponding list of bids is given by $B^\theta = f(2, 4; +1), (4, 2; +1), (4, 4; -1), (6, 6; +1)g$.*

Our contributions. This paper addresses the computational challenges of determining uniform component-wise minimal market-clearing prices (at which total market demands equal the quantities of each good that are available) and allocating a fixed collection of the goods at these prices, to bidders whose demands are defined by lists of positive and negative bids that express strong-substitute preferences. Our algorithms exploit the fact that the use of the product-mix bidding language allows for the efficient computation of a demanded bundle as well as of the indirect utility derived at any given prices.

domains.

⁵The ability to express such preferences was not thought necessary in the Bank of England’s application, but might be useful in closely related environments (see [Klemperer, 2018]).

⁶Klemperer [2010] stated this result for the case of multiple units of each of two goods. See [Baldwin and Klemperer, 2021] and [Lin and Tran, 2017] for working papers proving (different) generalisations of this result. (As cited by Lin and Tran [2017], a prior proof of the result itself was presented at a workshop; see [Baldwin et al., 2016].)

Section 2 introduces the product-mix auction’s strong-substitutes bidding language in more detail, and develops some of its properties. A first contribution of our paper is to show that it is coNP-complete to determine whether a given list of positive and negative bids constitute a valid demand correspondence. However, when the number of goods, or the number of negative bids, is bounded by a constant, we present a polynomial-time algorithm for checking validity.

In Section 3, we consider algorithms for finding component-wise minimal market-clearing prices that have practical running time bounds. We adopt a discrete steepest descent method from the discrete optimisation literature [Chapter 10 in Murota, 2003b, Shioura, 2017] that employs submodular minimisation to find the steepest descent directions and discuss two techniques (one novel, and one suggested by Shioura [2017]) that reduce the number of iterations required by taking long steps in the steepest descent direction. Full details are provided in Appendix C. We show that the resulting algorithms run in time polynomial in the input size of the bid lists and, since submodular minimisation is rapid in practice [Chakrabarty et al., 2017], we expect our approach to have a fast running time. Indeed, preliminary experiments (given in the online companion to this paper) suggest our method performs well in practice.

Our main contribution, given in Section 4, is an efficient polynomial-time algorithm that allocates the auctioneer’s chosen supply among the bidders at given market-clearing prices. The difficulty in developing this lies in handling bids whose utility is maximised on more than one alternative good (or whose utility from its most-preferred good is exactly zero), as tie-breaking choices interact with each other. This gives rise to a novel matching problem. Our algorithm proceeds by iteratively simplifying the allocation problem at hand, allocating unambiguous bids and perturbing bids and prices in a way that resolves a subset of the tie-breaks and yields a simplified allocation problem. Progress is measured in terms of reductions in the number of edges of a multigraph associated with the allocation problems.

Software implementations. Our price-finding and allocation algorithms are conceptually simple. They are also straightforward to implement, which has allowed us to develop two practical implementations in Haskell and Python; these can be found at <http://pma.nuff.ox.ac.uk> and <https://github.com/edwinkl/rock/product-mix>, respectively. Some effort was made to optimise for speed in the Python implementation by exploiting fast matrix operations provided by the NumPy package [van der Walt et al., 2011]. Furthermore, in an effort to implement submodular minimisation efficiently, the Fujishige-Wolfe algorithm [Chakrabarty et al., 2014] was implemented in combination with a memoization technique to reduce the number of submodular function queries.

In order to evaluate the practical running time of our allocation algorithm, we used our Python implementation to run experiments on auctions with various numbers of goods, bidders and bids. The results of these experiments, given in the online companion to this paper, demonstrate that our allocation algorithm is efficient in practice for realistic auction sizes.

Related work. Our work continues a long literature, spanning economics and discrete convex analysis, on ‘gross’ and ‘strong substitutes’ [Kelso and Crawford, 1982, Milgrom and Strulovici, 2009] and ‘ M^{\natural} -concavity’ [Murota and Shioura, 1999]. Whereas the discrete convex analysis literature generally allows for multiple units of each good, much focus in economics has been on the case in which there is only one unit of each good; Milgrom and Strulovici [2009] showed that ‘strong substitutes’ provide the suitable generalisation of gross substitutes to the multi-unit case, retaining existence of equilibrium while insisting that any two units of the same good should have the same price.

There is now a substantial body of related work on the two algorithmic problems we address in this paper: finding market-clearing prices and computing an equilibrium allocation of supply to the buyers.

Market-clearing prices are commonly found either by performing an iterative discrete steepest-descent search, or by solving a convex optimisation problem with a cutting plane method.

The method of steepest descent goes back to Ausubel [2006], Kelso and Crawford [1982], Murota et al. [2013] and Murota et al. [2016]; see [Murota, 2003b, Chapter 11], [Murota, 2016] and [Paes Leme, 2017] for recent surveys. However, while the steepest-descent methods in the above references run in pseudo-polynomial time (in the valuation and demand oracle settings), we describe two steepest descent algorithms based on [Shioura, 2017] that take longer steps by exploiting the bid representation of bidder demand, and show that they have polynomial running time in the size of the bid lists. (An algorithm is *polynomial* if its running time is polynomial in the input size, and *pseudo-polynomial* if its running time is required to be polynomial only in numeric values within the input. These differ when numerical values are encoded in binary.)

Baldwin et al. [2022] develop a conceptually similar descending-price algorithm using DC (difference of convex functions) programming.⁷ Extensive computational experiments in [Baldwin et al., 2022] suggest that this algorithm is competitive with the steepest descent algorithm we describe, and neither algorithm consistently outperforms the other. However, no polynomial running time bound is known for the algorithm of [Baldwin et al., 2022], whereas we provide a polynomial bound (in the bid list representation size) for our method.

Paes Leme and Wong [2017] take the alternative approach of finding market-clearing prices by solving a convex optimisation problem. Their solution builds on an improved cutting plane method of Lee et al. [2015] to achieve a polynomial-time algorithm. However, to the best of our knowledge, the cutting plane method has not yet been implemented and is computationally expensive in practice; moreover, it is not guaranteed to find component-wise minimal prices.

The harder algorithmic problem is finding a valid allocation of supply to buyers. Most previous studies have failed to address this problem in the case when multiple units of each good are available. The two existing algorithms for finding multi-unit allocations of which we are aware, in [Murota and Tamura, 2001] and [Paes Leme and Wong, 2017], both work in the valuation oracle model. That is, their algorithms rely on oracle access to the valuation function of each bidder. In order to apply these algorithms in our bidding language setting, we would require the computational overhead of simulating a valuation oracle. Section 2.3 describes a method for this simulation which takes time $O(n^2jBj^3 + njBjT(n))$ per valuation query (or $O(n^2jBj^2 \log M + njBjT(n))$ for an alternative variant), in which n is the number of goods, jBj is the number of bids, M is an upper bound on the bid entries, and $T(n)$ is the time it takes to minimise an n -dimensional submodular set function.

The allocation algorithm of Murota and Tamura [2001], also described in [Murota, 2003b, Chapter 11], works by reducing the allocation problem to an M -convex network flow problem using access to bidders' valuation functions. Following [Murota, 2003b, Chapters 10 and 11] and [Shioura, 1998] (for a polynomial-time algorithm for minimising M -convex functions), this approach requires $O((nm)^6)$ queries to the valuation function, where m is the number of bidders. With the additional overhead from simulating the valuation oracle this leads to a running time of $O(m^6n^8jBj^3 + m^6n^7jBjT(n))$.

Paes Leme and Wong [2017] provide a different algorithm, also in the valuation oracle setting, that solves the allocation problem only if the target bundle is *uniquely* demanded at some prices. Such prices are guaranteed in the single-unit case (i.e. when there is only one unit of every good) but need not exist in the multi-unit case. (Consider, for instance, the aggregate demand of Alice and Bob in Figure 2, also shown in Figure 3a of Section 3. The bundle $(1, 1)$ is not uniquely demanded at any prices.) However, if the target bundle is uniquely demanded at some prices, [Paes Leme and Wong, 2017]'s algorithm makes $O(mn + n^3)$ queries to the valuation functions to allocate the target bundle. With the simulation of valuation queries, this leads to running-time complexity $O((mn^2 + n^4)(njBj^3 + jBjT(n)))$.

⁷This exploits a result that market-clearing prices minimise the difference between two linear programs (one for the positive bids and one for the negative bids).

In contrast to the algorithms of [Murota and Tamura, 2001] and [Paes Leme and Wong, 2017], our allocation algorithm directly exploits the specific representation provided by the product-mix auction’s bidding language. Theorem 19 establishes that our algorithm takes time $O(mn^3\alpha(n)jBj + mn^2T(n))$. As the near-constant Ackermann function α satisfies $\alpha(n) \leq 4$ for any realistic number of goods n ,⁸ we see that our algorithm outperforms those of [Murota and Tamura, 2001] and [Paes Leme and Wong, 2017] (when applicable) in the product-mix auction setting. Moreover, our algorithm has been implemented and is efficient in practice.

We note that the valuation oracle setting of [Murota et al., 2013, 2016] and others provides no straightforward way to compute a demanded bundle or the indirect utility function at given prices. In contrast, our bidding-language setting is somewhat analogous to, but more informative than, the demand oracle setting used in [Ausubel, 2006], for example, which presupposes access to an oracle returning a demanded bundle at given prices.

This paper assumes that bidders are able to construct and submit bid lists that represent their strong-substitutes preferences. When the number of goods is large or the preferences are otherwise complex, this may present difficulties for bidders. Goldberg et al. [2022] address this issue by presenting algorithms that generate bid lists on behalf of bidders using only access to a demand or valuation oracle to elicit bidders’ preferences.

2 Preliminaries

We denote $[n] := \{1, \dots, n\}$ and $[n]_0 := \{0, \dots, n\}$. In our auction model, there are n distinct goods $[n]$; a single copy of a good is an *item*. A *bundle* of goods, typically denoted by \mathbf{x}, \mathbf{y} or \mathbf{z} in this paper, is a vector in \mathbb{Z}^n whose i -th entry denotes the number of items of good i . The *target bundle* \mathbf{t} is a bundle the auctioneer wants to allocate amongst the bidders. Vectors $\mathbf{p}, \mathbf{q} \in \mathbb{R}^n$ typically denote vectors of prices, with a price entry for each of the n goods. We write $\mathbf{p} \leq \mathbf{q}$ when the inequality holds component-wise. It is often convenient to regard a rejected bid as being accepted on a notional *reject good* for which bids and prices are always zero. Letting the reject good be 0, the set of goods is then $[n]_0$. In this setting, we identify bundles and prices with the $n + 1$ -dimensional vectors obtained by adding a 0-th entry of value 0.

A *valuation* u is a function that maps (non-negative) bundles to non-negative real numbers. The valuations we consider are concave-extensible and have bounded effective domain.⁹ Bidders have *quasilinear utilities*, i.e. the utility derived from bundle \mathbf{x} at price \mathbf{p} by a bidder with valuation u is given by

$$u(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x}. \quad (1)$$

Any valuation u is associated with a demand correspondence D_u that maps prices \mathbf{p} to the set of bundles that maximise (1). We omit the subscript u when it is clear from context.

For any subset $X \subseteq [n]$, \mathbf{e}^X denotes the characteristic vector of X , i.e. an n -dimensional vector whose i -th entry is 1 if $i \in X$ and 0 otherwise. Furthermore, \mathbf{e}^i denotes the vector whose i -th entry is 1 and other entries are 0.

A set function $f : 2^{[n]} \rightarrow \mathbb{R}$ is *submodular* if it satisfies $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$ for all $S, T \subseteq [n]$. Submodular function minimisation (SFM) is the task of finding a minimiser of a submodular function. It is well-known that the minimisers of submodular functions form a lattice; that is, if S and T are minimisers of f , then so are $S \cup T$ and $S \cap T$. SFM can be solved in polynomial time [Iwata, 2008, Jiang, 2021, McCormick, 2005]. For SFM that is efficient in practice, we refer to two SFM algorithms from the literature. The subgradient descent approach by Chakrabarty et al. [2017] determines a minimiser in time $O(nF^3\gamma \log n)$, while the Fujishige-Wolfe algorithm described by Chakrabarty et al. [2014] takes time $O(F^2(n^2\gamma + n^3))$, where F

⁸We have $\alpha(n) \leq 4$ for any $n \leq 2^{2^{2^{2^{16}}}}$.

⁹A function is concave-extensible if its concave closure agrees with the function at every integral point. See [Shioura and Tamura, 2015] for details, and how this assumption is natural for valuation functions.

is an upper bound on the absolute value of f and γ denotes the time it takes to query f . Experimental results [Chakrabarty et al., 2014] indicate that the running time of the Fujishige-Wolfe algorithm depends less on F than suggested by the above bound. For the price-finding algorithms described in Section C, we require a subroutine that finds the *inclusion-wise minimal minimiser*. Note that such a subroutine can be obtained by calling any SFM algorithm $n + 1$ times. Indeed, let S be a minimiser of f and, for every $v \succeq S$, let S_v be a submodular minimiser of the function f restricted to $[n] \cap \bar{v}g$. Then if S_0 denotes the minimal minimiser of f , we have $v \succeq S_0$ if and only if $f(S_v) > f(S)$, as the minimisers of a submodular function form a lattice. Hence, we obtain $S_0 := \bar{v} \succeq S \mid f(S_v) > f(S)g$.

2.1 Strong-substitutes valuation functions

In this paper, bidders have *strong-substitutes* valuations. We review some basic properties of strong-substitutes (SS) valuations. A SS valuation u divides price space into regions corresponding to bundles: any bundle \mathbf{x} has a region of prices in which \mathbf{x} is demanded, possibly along with other bundles (see Figure 2). It is known [Murota, 2003b, Theorem 11.16] that each such region is a convex lattice. When a region for some bundle \mathbf{x} has full dimensionality, then \mathbf{x} is uniquely demanded at any prices in the region’s interior; we call this interior a *unique demand region* (UDR). Moreover, the *Locus of Indifference Prices (LIP)* L_u comprising the set of prices at which demand is not unique consists of the union of $(n - 1)$ -dimensional linear pieces, or *facets*, that separate the UDRs. In Figure 2, the LIPs associated with the demand correspondences of Alice and Bob in Examples 1 and 2 are illustrated with dashed lines. We refer to Appendix B for further details on the LIP, including how every facet of the LIP can be endowed with a weight to capture the change in demand between any two price vectors.

Definition 1. A valuation u is *gross substitutes* if, for any prices $\mathbf{p}^0 \preceq \mathbf{p}$ and any $\mathbf{x} \succeq D_u(\mathbf{p})$, there exists $\mathbf{x}^0 \succeq D_u(\mathbf{p}^0)$ such that $x_k^0 = x_k$ for all k such that $p_k = p_k^0$. A valuation u is *strong substitutes (SS)* if, when we consider every unit of every good to be a separate good, u is gross substitutes.¹⁰

Importantly, Shioura and Tamura [2015, Theorem 4.1] establish that strong-substitutes is also equivalent to M^\natural -concavity [Murota and Shioura, 1999] in our setting of valuations with bounded non-negative domains. (Both strong substitutes and M^\natural -concavity imply concave-extensibility, see [Milgrom and Strulovici, 2009, Theorem 9] and [Murota, 2003b, Theorem 6.42] respectively.)

While GS guarantees that a competitive equilibrium exists in single-unit auction markets, the condition is not sufficient for the existence of such an equilibrium in the multi-unit case (Shioura and Tamura [2015] give an example). SS represents a natural generalisation of single-unit GS that provides a general sufficient condition for the existence of an equilibrium. (See [Baldwin and Klemperer, 2019] and [Shioura and Tamura, 2015] for discussion of the distinction between GS and SS.)

Definition 2. The *indirect utility function* f_u of valuation u maps a price vector \mathbf{p} to the utility that a bidder with demand D_u has for receiving her preferred bundle at a given price vector \mathbf{p} in the following way.

$$f_u(\mathbf{p}) := \max_{\mathbf{x}} (u(\mathbf{x}) \mid \mathbf{p} \succeq \mathbf{x}) = u(\mathbf{x}^0) \mid \mathbf{p} \succeq \mathbf{x}^0 \text{ with } \mathbf{x}^0 \succeq D_u(\mathbf{p}). \quad (2)$$

We note that $f_u(\mathbf{p})$ is convex, piecewise-linear and continuous for SS valuations u .

¹⁰Formally, we can associate every multi-unit valuation function u on a bounded effective domain with a single-unit valuation function $\hat{u}: \mathbb{R}^N \rightarrow \mathbb{R}$ as follows. Let U be an upper bound on the entries of the bundles in the domain of u . We define the domain \hat{u} as $N = \{j \mid j \in [n], j \in [U]g\}$, and let $\hat{u}(\mathbf{x}) = u(\mathbf{x})$, where \mathbf{x} is given by $x_i = \sum_{j=1}^U x_{ij}$. Valuation u is strong substitutes if \hat{u} is gross substitutes. See [Shioura and Tamura, 2015] for details.

2.2 Representing strong-substitutes valuations with weighted bids

We describe how every strong-substitutes valuation function can be represented by a finite list B of positive and negative bids. A *bid* \mathbf{b} consists of an $(n + 1)$ -dimensional integral vector $(b_1, \dots, b_n; b_{n+1})$, where the first n components b_1, \dots, b_n of the vector denote the bid's ‘valuation’ for the goods $[n]$ and the $(n + 1)$ -th component b_{n+1} is interpreted as the bid's *weight* (the quantity of goods sought by the bid, which may be positive or negative). For the reader's convenience, we also define the alternative notation $w(\mathbf{b}) := b_{n+1}$ to denote \mathbf{b} 's weight. Moreover, we restrict ourselves to positive and negative unit weights $\in \pm 1$. This is without loss of generality, as any bid with a weight of $w(\mathbf{b}) \geq 1$ can be represented by $jw(\mathbf{b})/j$ unit bids with the same vector and of the same sign. Finally, when working with the notional reject good 0 as described in Section 2, we identify the bid \mathbf{b} with the $(n + 2)$ -dimensional vector $(0, b_1, \dots, b_n; b_{n+1})$ obtained by prepending a 0-th entry of value 0 (and thus indexing from 0 instead of 1). This allows us to define a demand correspondence D_B that maps each price vector \mathbf{p} to a set of bundles demanded at \mathbf{p} , and an indirect utility function f_B associated with B .

A bid \mathbf{b} *demands good* $i \in [n]_0$ at \mathbf{p} if the surplus $b_i - p_i$ at price \mathbf{p} is maximal, that is if we have $i \in \arg \max_{i \in [n]_0} (b_i - p_i)$; recall that good 0 is the notional ‘reject’ good and $b_0 = p_0 = 0$ by definition. We say that $\arg \max_{i \in [n]_0} (b_i - p_i)$ is the set of *demanded goods of \mathbf{b} at \mathbf{p}* . A bid \mathbf{b} is *marginal* (on the set of its demanded goods) at \mathbf{p} if \mathbf{b} demands more than one good at \mathbf{p} . It is *marginal on good i* at \mathbf{p} if it is marginal at \mathbf{p} and i is in its set of demanded goods. Moreover, we say that price \mathbf{p} is *marginal* with respect to a given bid list B if there are bids in B that are marginal at \mathbf{p} , and non-marginal otherwise. We illustrate these definitions using Example 1 (Figure 2a). Recall that Alice's bid list B consists of two positively-weighted bids $\mathbf{b} := (6, 6; 1)$ and $\mathbf{b}^\ell := (0, 4; 1)$. At prices $\mathbf{p} = (2, 4)$, the bid \mathbf{b} is not marginal as it demands only good 1 while \mathbf{b}^ℓ is marginal between the reject good and good 2. At prices $\mathbf{p} = (6, 6)$, the bid \mathbf{b} is marginal between the reject good, good 1, and 2 while \mathbf{b}^ℓ is rejected (it only demands the reject good). Finally, at prices $\mathbf{p} = (6, 2)$, both bids are non-marginal, as they each demand only good 2. More generally, in Figure 2a, all prices along the dashed lines are marginal, and all other prices are non-marginal. Note that these marginal prices coincide with the LIP L_u , i.e. the set of prices at which demand is not unique, of Alice's underlying valuation u .

For any bid list B , we define the *demand correspondence* $D_B(\mathbf{p})$ at prices \mathbf{p} as follows, distinguishing between the cases that \mathbf{p} is marginal and non-marginal. If \mathbf{p} is non-marginal, every bid $\mathbf{b} \in B$ uniquely demands some good $i(\mathbf{b})$. In this case, $D_B(\mathbf{p})$ is a singleton set consisting of the bundle \mathbf{x} that is obtained by adding up an amount $w(\mathbf{b})$ of good $i(\mathbf{b})$ for each $\mathbf{b} \in B$, i.e. $D_B(\mathbf{p}) = f\mathbf{x}g$ with $\mathbf{x} := \sum_{\mathbf{b} \in B} w(\mathbf{b})\mathbf{e}^{i(\mathbf{b})}$. In Example 1, the demand correspondence at non-marginal prices $\mathbf{p} = (6, 2)$ is $D_B(6, 2) = f(0, 2)g$, as both bids demand good 2. At $\mathbf{p} = (1, 2)$, the demand correspondence is $D_B(1, 2) = f(1, 1)g$, as \mathbf{b} uniquely demands good 2 and \mathbf{b}^ℓ uniquely demands good 1. Figure 2a shows the demand in the different polyhedron interiors corresponding to non-marginal prices.

If \mathbf{p} is marginal, $D_B(\mathbf{p})$ consists of the discrete convex hull of the bundles demanded at non-marginal prices arbitrarily close to \mathbf{p} , where the discrete convex hull of a set of bundles X is defined as $\text{conv}(X) \setminus \mathbb{Z}^n$. In Figure 2a, we see that bundles demanded in the local neighbourhood of $\mathbf{p} = (0, 4)$ are $(1, 0)$ and $(1, 1)$, which implies $D_B(\mathbf{p}) = f(1, 0), (1, 1)g$. At $\mathbf{p} = (4, 4)$, we have $D_B(\mathbf{p}) = f(1, 0), (0, 1), (1, 1), (0, 2)g$.

For any bid list B , we can define the *indirect utility function*

$$f_B(\mathbf{p}) = \sum_{\mathbf{b} \in B} w(\mathbf{b}) \max_{i \in [n]_0} (b_i - p_i). \quad (3)$$

From (3) it is clear that we can compute $f_B(\mathbf{p})$ in linear time in the number of goods and the size of the bid list. In our setting, we can also efficiently compute a bundle demanded at a given price \mathbf{p} . This *demand oracle problem* is noted in [Paes Leme, 2017] as an algorithmic primitive needed to implement the Walrasian tâtonnement procedure. If \mathbf{p} is non-marginal, simply allocate

each bid \mathbf{b} a positive or negative item of the good $i(\mathbf{b})$ it uniquely demands and add up these items to obtain the demanded bundle $\mathbf{x} = \sum_{\mathbf{b} \in B} w(\mathbf{b}) \mathbf{e}^{i(\mathbf{b})}$. However, if \mathbf{p} is marginal and the bid list contains both positive and negative bids, we emphasise that independently allocating to each bid one of the goods it demands may not result in a bundle that lies in $D_B(\mathbf{p})$. In Example 2 (depicted in Figure 2b), for instance, we see that at prices $\mathbf{p} = (4, 4)$, allocating a negative item of good 1 to bid $(4, 4; -1)$, a positive item of good 1 to bid $(4, 2; 1)$ and a positive item of good 2 each to bids $(2, 4; 1)$ and $(6, 6; 1)$ leads to a total bundle of $(0, 2) \notin D_B(\mathbf{p})$. Instead, care must be taken to accept bids in a consistent way; one way to do this is to perturb entries of the price vector slightly so as to break ties, then note that the resulting bundle is demanded at the unperturbed prices.

Concurrent work by Baldwin and Klemperer [2021] (see also Footnote 6) shows that any SS demand correspondence can be represented as a finite list of positive and negative bids, and this representation is essentially unique (up to redundancies). Conversely, however, not all lists of positive and negative bids induce a valuation function. We call a bid list *valid* if the indirect utility function f_B defined in (3) is convex; Theorem 1 gives two further equivalent characterisations of validity.¹¹ In particular, it establishes that the indirect utility function and demand correspondence of a valid bid list as defined in Section 2.2 coincide with the indirect utility function and demand correspondence for some strong-substitutes valuation as defined in Sections 2 and 2.1. In Section 4, we will also introduce a weaker local notion of validity that holds in an open, convex neighbourhood of a price \mathbf{p} .

Theorem 1. *Let B be a list of positive and negative bids and f_B be its associated indirect utility function as defined in (3). The following conditions are equivalent.*

1. f_B is convex.
2. For every price vector \mathbf{p} and pair of distinct goods $i, i^0 \in [n]_0$ it holds that $\sum_{\mathbf{b} \in B_{ii^0}} w(\mathbf{b}) = 0$, where $B_{ii^0} \subseteq B$ is the subset of bids marginal on both i and i^0 at \mathbf{p} .
3. There exists a strong-substitutes valuation u satisfying $f_B = f_u$ and $D_B = D_u$.

The equivalence of (1) and (2) is a special case of Proposition 6 in Section 4, which is proved in Appendix B.2. It is well-known that (3) implies (1). To show that (2) implies (3), we use results of Mikhalkin [2004] to verify the existence of a suitable valuation u , and check that u indeed has the properties required. We present the full proof of Theorem 1 in Appendix B.1.

2.3 Simulating a valuation oracle

A bid list B does not directly provide the value $u(\mathbf{x})$ that the bidder has for a given bundle \mathbf{x} . However, the ability to evaluate $u(\mathbf{x})$ for arbitrary bundles is a prerequisite for the allocation algorithms in [Murota, 2003a] and [Paes Leme and Wong, 2017], as mentioned in the related work section. An efficient way to determine $u(\mathbf{x})$ is to compute prices \mathbf{p} at which the bidder demands \mathbf{x} using the price-finding algorithm from Section 3; then computing the bidder’s indirect utility $f_u(\mathbf{p}) = f_B(\mathbf{p})$ (cf. Definitions 2 and 3) at these prices; and finally recovering the bundle value

¹¹Equivalent conditions for validity are provided in concurrent work [Baldwin and Klemperer, 2021, Proposition 4.14], in a more general setting. There, also, it is shown that condition (2) is equivalent to the existence of a strong-substitutes valuation u satisfying $D_B = D_u$ (that is, the second part of condition (3)). [Baldwin and Klemperer, 2021] does not work with the indirect utility function, but does additionally show that validity is equivalent to D_B satisfying the “law of demand”: an increase in the price of only one good, between prices at which demand is unique, leads to a weak decrease in demand for the good whose price has changed, with equality in this demand if and only if the full demand vector is unchanged. Baldwin et al. [2022] additionally show that, when B is valid, the demand D_B can be given as the Minkowski difference between the demand associated with the set of positive bids in B and the demand associated with the set of negative bids in B . (The work of Lin and Tran [2017], mentioned in Footnote 6, does not consider characterisations of validity as, in their framework, invalid combinations do not arise.)

$u(\mathbf{x}) = f_B(\mathbf{p}) + \mathbf{p} \cdot \mathbf{x}$. The first step, finding the prices \mathbf{p} , takes time $O(n^2 j B^3 + n j B T(n))$ with the demand change method (and $O(n^2 j B^2 \log M + n j B T(n))$ with the binary search method) for determining the step lengths of the steepest descent method described in Section 3 and Appendix C. Here n is the number of goods, M is an upper bound on the bid entries and $T(n)$ is the time it takes to minimise an n -dimensional submodular set function. The second and third steps, computing $f_B(\mathbf{p})$ and solving for $u(\mathbf{x})$, take time $O(n j B)$, so the first step dominates the computation time.

2.4 Deciding validity of bid lists

In our auction, bidders submit their lists of positive and negative bids to the auctioneer prior to the auction. We show in Theorem 2 that checking the validity of a given list of bids is computationally hard. The proof, given in Appendix D, exploits definition (2) of validity in Theorem 1. However, in Appendix D we provide a straightforward algorithm to verify the validity of a given bid list in polynomial time if the number of goods, or the number of negative bids, is bounded by a constant.

Theorem 2. *Deciding the validity of a given list of positive and negative bids is coNP-complete.*

Theorem 3. *Let B be a list of positive and negative bids. If the number of goods or the number of negative bids is bounded by a constant, there exists a polynomial-time algorithm for checking the validity of B .*

Furthermore, many useful subclasses of bid lists (of arbitrary size) are easily checkable for validity in practice. Hence, while the coNP-completeness result is disappointing, it does not seem to be an important limitation on the auction, as sensible restrictions on the permitted bids can allow us to check validity efficiently in practice.

2.5 The computational challenges

In this subsection we state the computational problems to be solved. For any bidder $j \in J$, where J is the set of bidders, B^j denotes the bids of bidder j . We assume each B^j is *valid*, as defined in Section 2.2, and provided as a list of vectors encoded in binary. Let B be the aggregate list of bids obtained by concatenating the lists of all bidders. Theorem 1 part (2) implies that B is valid. Figure 3a depicts the aggregate bid list of Alice and Bob from Figure 2. The running times of our algorithms are given in terms of n , $|J|$, $|B|$ and $M := \max_{\mathbf{b} \in B} \|\mathbf{b}\|_1$, the maximum bid vector entry.

Suppose the auctioneer intends to sell target bundle \mathbf{t} . Our aim is to compute a competitive equilibrium: a market-clearing price \mathbf{p} at which \mathbf{t} is demanded and an allocation of \mathbf{t} to the various bidders so that every bidder receives a bundle they demand at \mathbf{p} . In the event that not enough bids are made for the target bundle, some items can go unsold, which is equivalent to the auctioneer buying them back from the market. To reflect this, the auctioneer places a total of $\|\mathbf{t}\|_1 + 1$ positive bids at $\mathbf{0}$.¹² The computation of a competitive equilibrium separates into two problems.

The price-finding problem. Given the aggregate bid list B and target bundle \mathbf{t} , find the coordinate-wise minimal *market-clearing* price \mathbf{p} at which \mathbf{t} is demanded, that is $\mathbf{t} \in D_B(\mathbf{p})$.

The allocation problem. Given a valid bid list B^j for each bidder $j \in J$, a target bundle \mathbf{t} and market-clearing price \mathbf{p} , allocate \mathbf{t} to the bidders so that each bidder receives a bundle they demand at \mathbf{p} , that is, a partition $(\mathbf{t}^j)_{j \in J}$ of \mathbf{t} with $\mathbf{t}^j \in D_{B^j}(\mathbf{p})$ for all $j \in J$.

¹²Of course, any positive reserve price can be reflected by appropriately locating the auctioneer's bids.

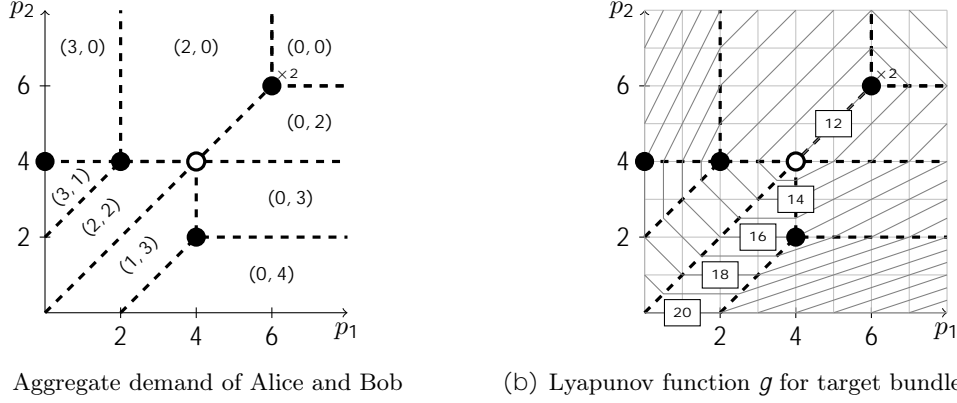


Figure 3: Figure (a) shows the demand of the aggregate bid list obtained by aggregating Alice’s and Bob’s bids from Figure 2, while (b) depicts a contour plot of the Lyapunov function g for target bundle $\mathbf{t} = (1, 1)$. Note that $\mathbf{p} = (4, 4)$ is the minimal market-clearing price for \mathbf{t} . Starting at $\mathbf{0}$, MinUp repeatedly moves in direction $d = (1, 1)$ until it reaches \mathbf{p} , whereas LongStepMinUp makes a single long step from $\mathbf{0}$ to \mathbf{p} .

Note that we do not ask for a breakdown of which of j ’s bids are accepted on which goods. Indeed, we show that a solution to the allocation problem can be obtained without this knowledge.

Theorem 4. *Suppose that each bidder $j \in J$ has SS demand correspondence D^j . Given a target allocation \mathbf{t} and a market-clearing price \mathbf{p} for \mathbf{t} , there exists a partition $\mathbf{t} = \sum_{j \in J} \mathbf{t}^j$ such that $\mathbf{t}^j \in D^j(\mathbf{p})$.*

Theorem 4 (cf. [Danilov et al., 2001, Murota, 2003b, Milgrom and Strulovici, 2009]) ensures that if the target bundle \mathbf{t} is aggregately demanded at market-clearing price \mathbf{p} , then there exists an allocation of \mathbf{t} among the bidders so that every bidder receives a bundle she demands. However, care must be taken in finding such an allocation, as bidders cannot simply be allocated an arbitrary bundle they demand at the market-clearing price. (Indeed, in Figure 2 the lowest prices at which Alice and Bob aggregately demand $(1, 1)$ is given by $\mathbf{p} = (4, 4)$. If we allocate Alice the whole bundle $(1, 1)$, Bob does not demand the empty remainder $(0, 0)$ at \mathbf{p} .)

If we had (oracle) access to the valuations of the bidders, we could perturb the valuations such that the target bundle is uniquely demanded at some price by subtracting carefully constructed functions that are discrete-convex in one variable.¹³ However, it is not clear how to perturb the bids individually to achieve a suitable ‘indirect’ perturbation of the valuations. Instead, our allocation algorithm takes the approach of perturbing bids bidder by bidder. Some intuition can be gained by considering the addition of a small ‘random’ perturbation vector \mathbf{v}^j to all the bids in bid set B^j (a different \mathbf{v}^j for each $j \in J$) and recomputing the new market-clearing prices. This has the effect of breaking ties between bids in two different bid lists. (This also uses the plausible fact, established in Proposition 14, that small perturbations make only small changes to the market clearing price(s).) However, the perturbations need to be exponentially small in the number of goods in order to ensure that all possible ties are broken. (One can check, noting the pigeonhole principle, that if perturbations are multiples of some inverse-polynomial, then they will have distinct subsets having the same sum, losing the guarantee that all possible ties are broken.) Our algorithm of Section 4 avoids this by systematically perturbing and un-perturbing bid sets, making decisions on marginal bids as it proceeds.

3 Finding market-clearing prices

In order to determine component-wise minimal prices at which the market is cleared, we apply an iterative steepest-descent approach from the discrete convex optimisation literature [Shioura, 2017]. This approach generalises an algorithm used by Ausubel’s ascending auction design [Ausubel, 2006] and Gul and Stacchetti [2000] for the task of finding equilibrium prices in single-unit markets. Ausubel [2006] and others define the *Lyapunov function* with regard to indirect utility function f_u for strong-substitutes valuation u and target bundle \mathbf{t} as $g_{\mathbf{t}}(\mathbf{p}) := f_u(\mathbf{p}) + \mathbf{t} \cdot \mathbf{p}$.

A function $f : Z^n \rightarrow \mathbb{R}$ is L^{\natural} -convex if it satisfies the *translation submodularity* property,

$$f(\mathbf{p}) + f(\mathbf{q}) \leq f((\mathbf{p} \vee \alpha \mathbf{1}) \wedge \mathbf{q}) + f(\mathbf{p} \wedge (\mathbf{q} + \alpha \mathbf{1})) \quad (\forall \mathbf{p}, \mathbf{q} \in Z_+^n, \forall \alpha \in Z_+). \quad (4)$$

Here, \vee and \wedge denote the component-wise maximum and minimum, respectively. It is well-known that the Lyapunov function $g_{\mathbf{t}}$ is L^{\natural} -convex if the valuation is strong substitutes (see [Murota et al., 2013, 2016], or Proposition 29 in Appendix C) and thus has a unique component-wise minimal minimiser [Shioura, 2017]. Moreover, this minimiser is integral and corresponds to the component-wise minimal market clearing prices (cf. Lemma 30 in Appendix C), and can be found using the following discrete steepest descent procedure. Given a starting point $\mathbf{p} := \mathbf{p}^0$ that is dominated by the market-clearing prices \mathbf{p}^* , repeatedly find the component-wise minimal subset $S \subseteq [n]$ minimising $g_{\mathbf{t}}(\mathbf{p} + \mathbf{e}^S) - g_{\mathbf{t}}(\mathbf{p})$ and move a unit step in this direction by updating $\mathbf{p} := \mathbf{p} + \mathbf{e}^S$. The method terminates once a local minimum is reached. As a local minimum of an L^{\natural} -convex function is a global minimum (see, e.g., [Shioura, 2017]), we have found \mathbf{p}^* . It is straightforward that the L^{\natural} -convexity of the Lyapunov function as defined in (4) guarantees the submodularity of function $g^{\theta}(\mathbf{p}; S) := g_{\mathbf{t}}(\mathbf{p} + \mathbf{e}^S) - g_{\mathbf{t}}(\mathbf{p})$ with respect to argument S (for any fixed \mathbf{p}). Hence we can apply standard submodular minimisation algorithms (cf. Section 2) to find the component-wise minimal steepest descent direction at any prices \mathbf{p} .

Adapting this method to our bidding-language setting, we can use our direct knowledge of the bids and (3) to express $g_{\mathbf{t}}$ for any valid bid list B as

$$g_{\mathbf{t}}(\mathbf{p}) := f_B(\mathbf{p}) + \mathbf{t} \cdot \mathbf{p} = \sum_{\mathbf{b} \in B} w(\mathbf{b}) \max_{i \in [n]_0} (b_i - p_i) + \mathbf{t} \cdot \mathbf{p}. \quad (5)$$

Theorem 1 guarantees that the function f_B defined in (3) is the indirect utility function f_u of some strong-substitutes valuation u , so (5) is a well-defined Lyapunov function. (Figure 3b depicts a contour plot of the Lyapunov function defined for the aggregate demand of Alice and Bob as given in Examples 1 and 2 and shown in Figure 3a.) Moreover, (5) implies that $g_{\mathbf{t}}$ can be evaluated at any \mathbf{p} in time $O(n|B|)$.

In order to improve the running time of the steepest-descent method, Shioura [2017] proposes aggregating several steps in the same direction into a single ‘long’ step. After finding the component-wise minimal direction \mathbf{e}^S , we move $\lambda(\mathbf{p}, S)$ unit steps in this direction at once by updating $\mathbf{p} := \mathbf{p} + \lambda(\mathbf{p}, S)\mathbf{e}^S$, where $\lambda(\mathbf{p}, S)$ is defined as

$$\lambda(\mathbf{p}, S) := \max\{\lambda \in Z_+ \mid g^{\theta}(\mathbf{p}; S) = g^{\theta}(\mathbf{p} + (\lambda - 1)\mathbf{e}^S; S)\} \quad (6)$$

and $g^{\theta}(\mathbf{p}; S) := g_{\mathbf{t}}(\mathbf{p} + \mathbf{e}^S) - g_{\mathbf{t}}(\mathbf{p})$. Shioura [2017] shows that the number of long steps before reaching \mathbf{p}^* is at most $n \max_{S \subseteq [n]} \{g^{\theta}(\mathbf{0}; S) \mid S \subseteq [n]\} / j$. In our setting, we see that $g^{\theta}(\mathbf{0}; S) \leq |B|j$, and obtain an upper bound of $n|B|j$ on the number of long steps that our steepest-descent method takes (see Appendix C).

In Appendix C, we describe two methods for computing the step length (6) at every iteration of the steepest-descent procedure, and analyse their complexity in our bidding-language setting. The first method uses binary search, as suggested by Shioura [2017], to find each $\lambda(\mathbf{p}, S)$; we show that determining each $\lambda(\mathbf{p}, S)$ takes time $O(n|B| \log M)$. (Recall that $M = \max_{\mathbf{b} \in B} \|\mathbf{b}\|_1$.)

¹³See [Murota, 2003b, Theorem 6.13 (4)].

Our second method for determining $\lambda(\mathbf{p}, S)$ (the ‘demand change’ method) directly exploits our knowledge of the individual bids to determine the step length in time $O(njBj^2)$. We see that the best method in practice depends on the magnitudes of $\log M$ and jBj . Overall, both methods lead to price-finding algorithms that are polynomial in the size of the bid lists, as we state in Theorem 5.

Theorem 5. *The steepest descent approach takes time $O(n^2jBj^2 \log M + njBjT(n))$ and $O(n^2jBj^3 + njBjT(n))$ to find the component-wise smallest market-clearing prices using the binary search and demand change long step methods, respectively.*

The computation time of the steepest descent method, even with long steps, is dominated by the task of finding a component-wise minimal submodular minimiser. In practice, we can exploit our knowledge of the bids to decrease the dimensionality of the submodular function to be minimised, which leads to performance improvements. For full details on these practical improvements, and on the ‘unit-step’ and ‘long-step’ steepest-descent methods for computing market-clearing prices in the bidding-language setting, we refer to Appendix C.

4 Allocations to the separate bidders

Suppose we are given a market-clearing price \mathbf{p} for target bundle \mathbf{t} . We now present an algorithm that solves the allocation problem, i.e. finds a partition $(\mathbf{t}^j)_{j \in J}$ of the target bundle \mathbf{t} such that \mathbf{t}^j is demanded by bidder j at price \mathbf{p} . We note that while the market-clearing price \mathbf{p} returned by our steepest descent approach in Section 3 is component-wise minimal, our allocation algorithm works for any integral market-clearing price.

Algorithm 1 Allocate

- 1: **Input:** Initial allocation problem A .
 - 2: **Output:** Target bundle allocation.
 - 3: Allocate all non-marginal bids by applying NonMarginals (Algorithm 3).
 - 4: **while** A has a non-empty bid list **do**
 - 5: Run the FindParams subroutine (Algorithm 2).
 - 6: **if** FindParams returns a demand cluster (I, j) with at most one link good i **then**
 - 7: Process unambiguous marginal bids with UnambiguousMarginals (Algorithm 4).
 - 8: **else**
 - 9: Simplify the allocation problem with ShiftProjectUnshift (Algorithm 5).
 - 10: Allocate all (newly) non-marginal bids with NonMarginals (Algorithm 3).
 - 11: **return** the partial allocation bundle \mathbf{m}^j for each bidder $j \in J$.
-

Our algorithm Allocate, stated in Algorithm 1, repeatedly simplifies the problem until it becomes vacuous. It generates a sequence of *allocation problems* by iteratively allocating parts of the target bundle to bidders and removing satisfied bids from the bid lists until the residual supply bundle (initially the target bundle) is empty. In order to capture this formally, we define a generalised allocation problem in Section 4.1; the definition introduces a weaker notion of local validity for bid lists and also features a partial allocation bundle for each bidder as well as a residual target bundle. With every allocation problem we associate a corresponding *marginal bids multigraph* and *derived graph*. These graphs, and key terminology such as demand clusters and link goods, are defined in Section 4.2. The marginal bids multigraph is used to quantify progress and establish the running time of the algorithm, while the derived graph is required for Allocate to decide whether to apply the UnambiguousMarginals or ShiftProjectUnshift procedure.

The algorithm starts by removing all non-marginal bids and transferring their demanded items from the residual supply to the partial allocation bundle of the bid’s owner. This is

performed by the procedure `NonMarginals`, which is defined in Section 4.3. Next, the algorithm calls the `FindParams` subroutine (defined in Section 4.2), which constructs and works on the derived graph. If `FindParams` identifies a subset of marginal bids that can be allocated unambiguously, `Allocate` invokes `UnambiguousMarginals` (defined in Section 4.3) to process these bids. Otherwise `FindParams` identifies parameters that `Allocate` uses to call the `ShiftProjectUnshift` procedure. `ShiftProjectUnshift` is defined in Section 4.4; it does not allocate items and delete bids but instead simplifies the allocation problem by shifting and projecting the bids in order to strictly reduce the demand ties between bids. Finally, `NonMarginals` is invoked again to process any bids that may have become non-marginal. The algorithm repeats this process until the allocation problem become vacuous in the sense that all bid lists are empty, at which point the partial allocation bundles constitute a valid allocation of the target bundle \mathbf{t} at price vector \mathbf{p} . Section 4.5 gives the running time and a proof of correctness for `Allocate`.

4.1 The generalised allocation problem

We now introduce allocation problems formally and define the terminology and technical tools required to reason about how `Allocate` successively reduces the initial allocation problem using the three procedures `NonMarginals`, `UnambiguousMarginals` and `ShiftProjectUnshift`. In general, we note that applying these procedures may result in bid lists that are no longer valid in the sense of Theorem 1. (Figure 5 in Section 4.3 gives an example of how applying `NonMarginals` can result in an invalid bid list.) Instead, we introduce the notion of local validity that requires bid lists to be valid only in the neighbourhood of the market-clearing price \mathbf{p} . For prices \mathbf{p} in this neighbourhood, we can define the demand correspondence $D_B(\mathbf{p})$ in the same way as for globally valid bid lists. Let $B(\mathbf{p}, \varepsilon) := \{\mathbf{y} \in \mathbb{R}^n \mid \|\mathbf{y} - \mathbf{p}\|_1 < \varepsilon\}$ denote the open ε -ball with regard to the L_1 -norm.

Definition 3. A list of bids B is ε -valid at \mathbf{p} if the indirect utility function f_B restricted to $B(\mathbf{p}; \varepsilon)$ is convex. The bid list is *locally valid* at \mathbf{p} if there exists some $\varepsilon > 0$ for which it is ε -valid.

This generalises the notion of valid bid lists introduced in Section 2.2, and we correspondingly generalise Theorem 1 with the following proposition (proved in Appendix B.2).

Proposition 6. Let B be an ε -valid bid list at \mathbf{p} , and let $P = B(\mathbf{p}; \varepsilon)$. The following are equivalent.

1. The indirect utility function f_B restricted to P is convex.
2. There is no price vector $\mathbf{p} \in P$ and pair of distinct goods $i, i^0 \in [n]_0$ at which the weights $w(\mathbf{b})$ of the bids $\mathbf{b} \in B$ marginal on i and i^0 sum to a negative number.

Definition 4. An allocation problem is a 4-tuple $A = [\mathbf{p}, (B^j)_{j \in J}, (\mathbf{m}^j)_{j \in J}, \mathbf{r}]$, where

1. $\mathbf{p} \in \mathbb{R}^{n+1}$ is a price vector with $p_0 = 0$,
2. for each bidder $j \in J$, B^j is a list of bids in $f_0 g \quad (Z + f_0, \frac{1}{10}g)^n \quad Z$ that is locally valid at \mathbf{p} ,
3. $\mathbf{m}^j \in \mathbb{Z}_+^{n+1}$ for each $j \in J$ is the partial allocation,
4. $\mathbf{r} \in \mathbb{Z}_+^{n+1}$ is the residual supply,

and there exists a valid allocation $(\mathbf{r}^j)_{j \in J}$ of the residual supply \mathbf{r} to bidders, i.e. $\mathbf{r}^j \in D_{B^j}(\mathbf{p})$ and $\sum_{j \in J} \mathbf{r}^j = \mathbf{r}$. For any such valid allocation, we say that $\mathbf{t}^j := \mathbf{r}^j + \mathbf{m}^j$ is a *solution* to A .

Note that the bids need not be integral and the bid lists for the bidders are only required to be locally valid at \mathbf{p} . The intuition behind this definition is to capture the action of successively allocating items of the target bundle \mathbf{t} to bidders, which may break global validity but preserves local validity at \mathbf{p} . For each bidder $j \geq J$, the partial allocation \mathbf{m}^j is a bundle denoting the subset of \mathbf{t} already allocated to j , while \mathbf{r} denotes the remaining part of \mathbf{t} not yet allocated to any bidder. The 0-th coordinates of \mathbf{m}^j and \mathbf{r} denote the number of reject goods (not yet) allocated.

Definition 5. We say that A^θ is a *valid reduction* from A if A^θ is an allocation problem and all solutions to A^θ are also solutions to A .

In the *initial problem*, the residual supply is given by the target bundle and the partial allocation vectors are $\mathbf{0}$. Note that t_0 denotes the total number of rejected bids and can be computed as the total weight of bids minus the total number of items in \mathbf{t} . In the *vacuous problem*, the bid lists are empty and the residual supply is $\mathbf{0}$.

Technical tools. The following observations and lemma will be used in Section 4.4 to prove that applying the procedure `ShiftProjectUnshift` to an allocation problem A produces an allocation problem that is a valid reduction from A . We define the *surplus gap* of a bid \mathbf{b} at \mathbf{p} as the difference between the utility derived from a demanded good and the maximum utility derived from a non-demanded good. This allows us to describe the goods that \mathbf{b} demands when we perturb the price or the bid by a small amount. Moreover, we see in Lemma 9 that the size of the neighbourhood around \mathbf{p} in which a bid list is valid can be lower bounded by the surplus gaps of the bids in B . Formally, let $I = \arg \max_{i \in [n]_0} b_i - p_i$ denote the goods demanded by \mathbf{b} at \mathbf{p} . Then if $I \not\subseteq [n]_0$, we define the *surplus gap* of \mathbf{b} at \mathbf{p} as

$$\max_{i \in [n]_0} (b_i - p_i) - \max_{i \in [n]_{0 \setminus I}} (b_i - p_i).$$

Observation 7. If bid \mathbf{b} demands goods I at \mathbf{p} with a surplus gap of at least ε , then for any price $\mathbf{p}^\theta \geq B(\mathbf{p}, \varepsilon/2)$, the goods I^θ demanded by \mathbf{b} at \mathbf{p}^θ form a subset of I .

Observation 8. If bid \mathbf{b} demands goods I at \mathbf{p} with a surplus gap of at least ε , then for any bid $\mathbf{b}^\theta \geq B(\mathbf{b}, \varepsilon/4)$ and price $\mathbf{p}^\theta \geq B(\mathbf{p}, \varepsilon/4)$, the goods I^θ demanded by \mathbf{b}^θ at \mathbf{p}^θ form a subset of I .

Lemma 9. If B is δ -valid at \mathbf{p} for some $\delta > 0$ and all bids in B have a surplus gap of at least ε at \mathbf{p} , then B is $\varepsilon/2$ -valid at \mathbf{p} .

We prove this lemma in Appendix B.3.

4.2 The marginal bids graph and the derived graph

With every allocation problem A we associate a marginal bids multigraph and a derived graph. As mentioned above, the marginal bids multigraph allows us to provide upper bounds on the running time of our allocation algorithm. Moreover, in order to decide whether to apply `UnambiguousMarginals` or `ShiftProjectUnshift` and to determine the input parameters for these procedures, our `Allocate` algorithm calls the subroutine `FindParams`, defined in this section, which first constructs the derived graph and then attempts to find a maximal open or closed path in this graph.

Definition 6. The *marginal bids graph* G_A associated with allocation problem A is an undirected edge-labelled multigraph whose vertices are the goods $[n]_0$ (including the ‘reject’ good). G_A has an edge (i, i^θ) labelled with j if bidder j has a bid that is marginal at \mathbf{p} between i and i^θ .




Figure 4: Example of a marginal bids multigraph (a) and derived graph (b) with three bidders A, B, C and six goods $1, \dots, 6$ together with reject good 0. Goods $1, 2, 5, 6$ are link goods. Goods $2, 3, 4$ form a demand cluster $(\{2, 3, 4\}, A)$ with one link good, represented by a leaf demand cluster in the derived graph (bolded in both graphs).

For any bidder j , let $G_A[j]$ denote the subgraph of G_A induced by all j -labelled edges; note that this graph is simple, as there is at most one edge labelled j connecting two goods i and i^θ in G_A . A pair (I, j) with $I \subseteq [n]_0$ and $j \in J$ is a *demand cluster* if I is the set of vertices of some connected component of $G_A[j]$. Intuitively, this captures the dependencies between the demands of (and thus the possible allocations to) bidder j 's bids. A vertex is a *link good* if its incident edges are labelled with more than one bidder. (Effectively, a link good expresses a dependency between demand clusters of different bidders.) We call a cycle in G_A a *multi-bidder cycle* if it contains edges labelled by different bidders. A vertex in such a cycle is called a *cycle-link good* if its two edges in the cycle are labelled differently. Note that any multi-bidder cycle has at least two cycle-link goods. Like link goods, multi-bidder cycles and their cycle-link goods capture the dependencies between demands of (and thus possible allocations to) different bidders.

Definition 7. The *derived graph* D_A associated with allocation problem A is a simple bipartite graph whose two disjoint and independent vertex sets are the set of link goods and the set of demand clusters of G_A . There is an edge between link good i and demand cluster (I, j) if $i \in I$.

Note that by the definition of link goods, every link good is adjacent to at least two demand clusters in the derived graph. This is consistent with our observation that a link good can be interpreted as a link between connected components of $G_A[j]$ for different bidders j . Figure 4 gives an example of a marginal bids multigraph and its corresponding derived graph.

We describe a procedure to construct the derived graph in near-linear time $O(\alpha(n)n/B)$ using a disjoint-union data structure. Here $\alpha(\cdot)$ is the inverse Ackermann function, which grows extremely slowly and is near-constant in our context, as $\alpha(n) \leq 4$ for any $n \leq 2^{2^{2^{16}}}$. The disjoint-union data structure (cf. [Tarjan and van Leeuwen, 1984]) maintains a representation of a set partition and admits an $\alpha(n)$ -time operation to merge two subsets of the partition. Internally, it assigns a distinct label to each subset of the partition and provides $\alpha(n)$ -time access to the label of the subset in which a given element lies.

First we show how to compute the demand clusters for each bidder. Fix a bidder j and initialise the disjoint-union data structure. For each bid $\mathbf{b} \in B^j$, compute the marginal goods S at prices \mathbf{p} and, fixing any $i \in S$, merge the sets containing i and i^θ for all $i^\theta \in S \cap \hat{f}ig$. Now the data structure has learnt the vertex partition induced by the connected components of G_A in time $O(\alpha(n)n/B^j)$. In order to recover this partition and express it as a family of sets, we initialise an empty family \mathcal{K} of sets, where each set in \mathcal{K} will have an associated label. For each good $i \in [n]$, determine the label l of i 's subset in the data structure. If there already is a set in \mathcal{K} with label l , add i to this set. Otherwise, add the new singleton set $\hat{f}ig$ with label l to \mathcal{K} . Finally, iterate through \mathcal{K} and delete all singleton sets. This takes time $O(\alpha(n)n)$. Now \mathcal{K} represents the family of demand clusters of bidder j . Hence in total it takes time $O(\alpha(n)n/B)$ to compute the demand clusters for all bidders.

In order to compute the link goods, iterate through the demand clusters of all bidders and count the number of times each good appears. If a good appears at least twice, it is a link

good. Once we know which good is a link good, we can compute the edges of the derived graph: for each demand cluster (I, j) , add an edge between (I, j) and each link good in I . Iterating through the demand clusters of all bidders takes $O(nm) = O(njBj)$ time.

The FINDPARAMS subroutine. In every iteration, the Allocate algorithm employs a subroutine to decide, using the derived graph, whether to invoke UnambiguousMarginals or ShiftProjectUnshift, and to compute the input parameters for the respective procedure. This subroutine FindParams, given in Algorithm 2, takes as input an allocation problem and returns one of three possible outputs: a demand cluster (I, j) with no link goods, a demand cluster (I, j) with one link good denoted i , or a cycle-link good i and the label j of one of its incident edges in the multi-bidder cycle. The Allocate algorithm invokes UnambiguousMarginals if FindParams returns a demand cluster and ShiftProjectUnshift if it returns a cycle-link good and edge label.

FindParams works by walking through the derived graph D_A to find a maximal path. A *path* in a graph is a sequence of distinct vertices v_1, \dots, v_k such that v_i and v_{i+1} are connected by an edge for all $i \in [k-1]$; it is *maximal* in a given graph if we cannot extend the path with any vertex v_{k+1} of the graph. A path is *closed* if there is an edge between v_1 and v_k , and *open* otherwise. (Hence a closed path is a cycle.) Note that a path in D_A alternates between link good and demand cluster vertices.

Algorithm 2 FindParams


- 1: Compute the derived graph D_A .
 - 2: **if** D_A contains an isolated demand cluster **then**
 - 3: **return** this demand cluster (without a link good).
 - 4: Starting from any link good, take a walk in D_A to find a maximal path.
 - 5: Let i and (I, j) be the last link good and demand cluster visited, respectively.
 - 6: **if** the path is open **then**
 - 7: **return** (I, j) and i .
 - 8: **else**
 - 9: **return** i and j .
-

Lemma 10. *In time $O(\alpha(n)njBj)$, FindParams returns a demand cluster (I, j) with no link goods, a demand cluster (I, j) with one link good i , or a cycle-link good i and the label j of one of its incident edges in the multi-bidder cycle.*

Proof. Note that any path in the derived graph D_A has length at most $2n$. Hence constructing the derived graph, which takes time $O(\alpha(n)njBj)$, dominates the running time. By construction of the derived graph, an isolated demand cluster in D_A contains no link goods, while a leaf demand cluster contains exactly one link good. Hence if there is an isolated demand cluster, the subroutine returns a demand cluster without a link good. If the path found is open, the last vertex must be a leaf demand cluster and the subroutine returns a demand cluster with single link good. Now suppose the path is closed and consider only the cycle formed by alternating link good and demand cluster vertices. Firstly, note that there is a path between consecutive link goods in the marginal bids graph G_A using only j -labelled edges. Secondly, consecutive demand clusters have different bidder labels. These two observations imply that G_A contains a simple multi-bidder cycle with cycle-link good i and an incident edge labelled with j . \square

4.3 Allocating unambiguous bids

Non-marginal bids. Suppose bidder j has a non-marginal bid \mathbf{b} at market-clearing prices \mathbf{p} that demands good $i \in [n]_0$ in the allocation problem A . Then this bid contributes exactly



(a) Alice's bid list after applying NONMARGINALS. (b) Bob's list after applying NONMARGINALS.

Figure 5: The resulting bid lists after applying NonMarginals (Algorithm 3) to Alice's and Bob's bid lists from Figure 2 at prices $\mathbf{p} = (6, 4)$ (marked by a cross). Note that Alice's bid list remains globally valid, whereas Bob's bid list is only locally 2-valid at \mathbf{p} (indicated by the hatched square).

$w(\mathbf{b}) \geq f_{i,1}g$ items of good i to any solution of A . Hence we can unambiguously allocate these items to \mathbf{m}^j and remove them from the residual supply \mathbf{r} , thus accepting the non-marginal bid on the appropriate good. NonMarginals, given in Algorithm 3, processes all non-marginal bids in this way. Note that while this operation may not preserve global validity of bid lists, the resulting lists remain locally valid at \mathbf{p} , so that the result is a valid allocation problem. Figure 5 gives an example. Allocate calls NonMarginals in every iteration in order to process non-marginal bids that can be allocated unambiguously.

Algorithm 3 NonMarginals (accept non-marginal bids)

- 1: **Input:** Allocation problem $A = [\mathbf{p}, (B^j)_{j \in J}, (\mathbf{m}^j)_{j \in J}, \mathbf{r}]$.
 - 2: **Output:** Reduced allocation problem A^θ without non-marginal bids.
 - 3: **for all** bidders $j \in J$ **do**
 - 4: **for all** non-marginal bids $\mathbf{b} \in B^j$ **do**
 - 5: Determine unique good i demanded by \mathbf{b} and remove \mathbf{b} from B^j .
 - 6: Increment m_i^j by $w(\mathbf{b})$.
 - 7: Decrement r_i by $w(\mathbf{b})$.
-

Lemma 11. *Given a valid allocation problem A , NonMarginals (Algorithm 3) outputs a reduction A^θ of A in linear time. Moreover, we have $G_A = G_{A^\theta}$.*

Proof. Let $A^\theta = [\mathbf{p}, (B^j)_{j \in J}^\theta, (\mathbf{m}^j)_{j \in J}^\theta, \mathbf{r}^\theta]$ be the output of NonMarginals. First we show that A^θ is a valid allocation problem. Fix some $j \in J$. Since B^j is locally valid by assumption, the indirect utility function f_{B^j} is convex in some small neighbourhood of \mathbf{p} . Allocating a non-marginal bid to bidder j corresponds to subtracting from the utility function $f_{B^j}(q)$ the term $\max_{i \in [n]_0} (b_i - q_i)$. In a sufficiently small neighbourhood of \mathbf{p} , this term is linear, so the resulting utility function is also convex in some open neighbourhood of \mathbf{p} . It is straightforward to see that $(\mathbf{t}^j)_{j \in J}$ is a solution to A^θ if and only if it is a solution to A , so A^θ is a reduction of A . To see that $G_A = G_{A^\theta}$, note that the marginal bids are unchanged. \square

Unambiguous marginal bids. Let A be an allocation problem without non-marginal bids and suppose FindParams returns a demand cluster (I, j) with at most one link good. Unambiguous-Marginals is invoked by Allocate to unambiguously allocate the bids specified by this demand cluster. By the definition of a demand cluster, none of bidder j 's bids are marginal between items in both I and $[n]_0 \setminus I$. Let B_I^j denote the bids marginal on goods in I . All bids

$\mathbf{b} \in B_I^j$ contribute an item of a good in I to bidder j 's final allocation bundle, while all other bids contribute an item of a good not in I .

If the demand cluster (I, j) has no link goods, none of the bids by bidders other than j are marginal on I , so all items of goods $i \in I$ in the residual supply must be allocated to j . Hence we reduce our allocation problem by transferring r_i items from the residual supply \mathbf{r} to bidder j 's partial allocation bundle \mathbf{m}^j for each $i \in I$ and removing all bids marginal on items in I from B^j .

If the demand cluster (I, j) has a single link good i , all items of goods in $I \cap \bar{f}_i$ must be allocated to j by the same argument as above. Secondly, the bids in B_I^j must be allocated a total of $\sum_{\mathbf{b} \in B_I^j} w(\mathbf{b})$ units of goods from I . Hence the difference $\sum_{\mathbf{b} \in B_I^j} w(\mathbf{b}) - \sum_{i \in I \cap \bar{f}_i} r_i$ gives us the number of items of i that must be allocated to bidder j . This yields Algorithm 4.

Algorithm 4 UnambiguousMarginals (process all unambiguous marginal bids)

- 1: **Input:** Allocation problem $A = [\mathbf{p}, (B^j)_{j \in J}, (\mathbf{m}^j)_{j \in J}, \mathbf{r}]$ and demand cluster (I, j) with at most one link good i .
 - 2: **Output:** Valid reduction A^θ of A with no bids marginal on goods in I .
 - 3: **if** (I, j) has no link goods **then**
 - 4: **for all** $i \in I$ **do**
 - 5: Increment m_i^j by r_i and set r_i to 0.
 - 6: **else**
 - 7: Compute $d = \sum_{\mathbf{b} \in B_I^j} w(\mathbf{b}) - \sum_{i \in I \cap \bar{f}_i} r_i$.
 - 8: Increment m_i^j by d and decrement r_i by d .
 - 9: **for all** $i \in I \cap \bar{f}_i$ **do**
 - 10: increment m_i^j by r_i and set r_i to 0.
 - 11: Remove all bids marginal on I from B^j .
-

Lemma 12. UnambiguousMarginals, given in Algorithm 4, returns a valid reduction of the input allocation problem in time $O(njB^j)$. Moreover, the marginal bids graph G_{A^θ} of A^θ has strictly fewer edges than G_A .

Proof. In order to see that A^θ is a valid allocation problem, we verify that the new bid list $(B^j)^\theta$ of bidder j after applying UnambiguousMarginals is locally valid by checking criterion (2) of Definition 3. Fix a price \mathbf{p} and goods i, i^θ . Recall that bids in B^j cannot be marginal on goods in I and $[n]_0 \cap I$, by the definition of demand clusters. If $i \in I$ and $i^\theta \in [n]_0 \cap I$, then B^j has no bids marginal on i and i^θ . If both i, i^θ are goods in I , all bids marginal on i and i^θ are removed by UnambiguousMarginals. If neither i nor i^θ is a good in I , then none of the bids marginal on i and i^θ are removed. As B^j is locally valid by assumption, this implies that $(B^j)^\theta$ is also locally valid. It is straightforward to see that a solution to A^θ is a solution to A , as the partial allocation performed by UnambiguousMarginals is unambiguous. Finally, let $i, i^\theta \in I$ and note that the marginal bids graph G_A contains an edge between i and i^θ that is not present in G_{A^θ} . As removing bids does not introduce new edges to the graph, G_{A^θ} has strictly fewer edges. \square

4.4 The shift-project-unshift reduction

Suppose that FindParams returns a cycle-link good i and the label j of one of its adjacent edges in the cycle. In this case Allocate invokes the ShiftProjectUnshift procedure to obtain a reduction. Unlike NonMarginals and UnambiguousMarginals, this procedure does not make progress by allocating items and deleting the satisfied bids. Instead, it temporarily shifts the bids of bidder j and perturbs the market-clearing price \mathbf{p} using SFM. We show that

the marginal bids graph for this new allocation problem has strictly fewer edges, reducing the overall dependencies between the demands of bids. Finally, the procedure projects the bids of bidder j in such a way that it can unperturb the market-clearing price while retaining the dependency structure of the new marginal bids graph. We now introduce the shift and project operations.

Shifting bids. Suppose all bids are integral and we shift some bidder's bids by a small quantity $\varepsilon < 1/4$ (we use $\varepsilon = 1/10$ for concreteness) in the direction of i by adding $\varepsilon \mathbf{e}^i$ to each bid vector. Then Lemma 13 and Proposition 14 (proved in Appendix B.4) together show that we can use SFM to find a price $\mathbf{p}^\varepsilon \geq \{\mathbf{p} + \varepsilon \mathbf{e}^S, S = [n]\}$, at which the new allocation problem A^0 with shifted bids and price vector \mathbf{p}^ε is a valid reduction.

Lemma 13. *Fix ε with $j\varepsilon < 1/4$, as well as $i \in [n]$ and $j \in J$. Let A be an allocation problem with integral bids and prices, and let A^0 be obtained by replacing B^j with $(B^j)^0 := \{\mathbf{b} + \varepsilon \mathbf{e}^j \mid \mathbf{b} \in B^j\}$. Then the bid lists of all bidders in A^0 are locally valid at any price $\mathbf{p}^\varepsilon \in B(\mathbf{p}, \varepsilon)$. Moreover, the bundles demanded by any bidder at \mathbf{p}^ε in A^0 form a subset of the bundles they demand at \mathbf{p} in A .*

Proof. As the bids and prices in A are integral, each bid either demands all goods $[n]_0$ or has a surplus gap of at least 1 at \mathbf{p} . By Lemma 9, all bid lists are $1/2$ -valid at \mathbf{p} . This implies that every unshifted bid list is $(1/2 - \varepsilon)$ -valid at \mathbf{p}^ε and the shifted bid list is $(1/2 - 2\varepsilon)$ -valid at \mathbf{p}^ε . By Observations 7 and 8, a non-marginal bid demands the same good for all $\mathbf{q} \in B(\mathbf{p}^\varepsilon, \delta)$ and sufficiently small $\delta > 0$. Hence a bidder will demand the same set of bundles at all non-marginal prices in $B(\mathbf{p}^\varepsilon, \delta)$. As the bundles a bidder demands at \mathbf{p}^ε are by definition the discrete convex hull of bundles they demand at non-marginal prices infinitesimally close to \mathbf{p}^ε , we are done. \square

Proposition 14. *Let A be an allocation problem with price vector \mathbf{p} and fix ε with $j\varepsilon < 1/4$. Suppose we shift all bids of bidder j 's in B by ε to obtain (aggregate) bid list B^0 . Then there exists a price $\mathbf{p}^\varepsilon \geq \{\mathbf{p} - \varepsilon \mathbf{e}^S, S = [n]\}$ at which the residual supply \mathbf{r} is demanded by B^0 . We can determine \mathbf{p}^ε using submodular minimisation on the Lyapunov function g with regard to list B^0 and the residual supply \mathbf{r} .*

Projecting bids. Next we define a projection operation on bids. The idea of this operation is to modify a bid \mathbf{b} so that its surplus gap at \mathbf{p} is increased. As a consequence, if that bid or the price is perturbed slightly, it 'remembers' its preferred goods. Let \mathbf{b} be a bid and $I = \arg \max_{i \in [n]_0} (b_i - p_i)$ be the set of demanded goods at price \mathbf{p} . The projection \mathbf{b}^0 of \mathbf{b} w.r.t. \mathbf{p} is defined as follows.


$$\mathbf{b}^0 := \begin{cases} \mathbf{b} - \mathbf{e}^{[n]_0 \setminus I} & \text{if } 0 \in I, \\ \mathbf{b} + \mathbf{e}^I & \text{otherwise.} \end{cases}$$

Note that we allow for bid vector entries to be negative. Figure 6 illustrates the projection operation.

Observation 15. *If bid \mathbf{b} demands all goods $[n]_0$, we have $\mathbf{b}^0 = \mathbf{b}$. Otherwise, the projection operation on \mathbf{b} w.r.t. \mathbf{p} increases the bid's surplus gap at \mathbf{p} by 1.*

Lemma 16. *Projecting all bids in a bidder's bid list B w.r.t. \mathbf{p} does not change the set of bundles demanded at \mathbf{p} . Moreover, if B is locally valid at \mathbf{p} , the projected bid list is locally $1/2$ -valid.*

Proof. Suppose all bids in B have a surplus gap of at least ε at \mathbf{p} and B is locally δ -valid at \mathbf{p} . Note that at any price \mathbf{q} in the open ball $B(\mathbf{p}, \varepsilon)$ centred at \mathbf{p} , every bid demands the same set of goods before and after projecting. Hence the projected bid list is locally $\min\{\delta, \varepsilon\}$ -valid at \mathbf{p} , and the set of bundles demanded at \mathbf{p} , consisting of the discrete convex hull of bundles demanded in UDRs bordering \mathbf{p} remains the same. The second statement follows from Observations 9 and 15. \square



(a) Five numbered bids before projecting.

(b) The result of projecting w.r.t. prices $\mathbf{p} = (4, 4)$.

Figure 6: Example of the project operation w.r.t. prices $\mathbf{p} = (4, 4)$ on five bids numbered from 1 to 5 shown in (a). The projected bids are shown in (b). Note that only the negative bid is unchanged, as it demands all goods, including the reject good.

The procedure. The procedure `ShiftProjectUnshift` is stated in Algorithm 5. It shifts and projects to reduce the number of edges in marginal goods graph. This effectively reduces the number of demand ties among the bids in our allocation problem. Lemmas 17 (proved in Appendix B.5) and 18 establish that the reduction obtained by `ShiftProjectUnshift` is valid and makes progress.

Algorithm 5 `ShiftProjectUnshift`

- 1: **Input:** Allocation problem $A = [\mathbf{p}, (B^j)_{j \in J}, (\mathbf{m}^j)_{j \in J}, \mathbf{r}]$, cycle-link good i and edge label j .
 - 2: **Output:** Allocation problem A^θ .
 - 3: Add $\frac{1}{10} \mathbf{e}^i$ to each of j 's bids and compute a new market-clearing price $\mathbf{p}^\varepsilon = \mathbf{p} + \frac{1}{10} \mathbf{e}^i$ by solving $S = \arg \min_{S \subseteq [n]} g_{\mathbf{r}}(\mathbf{p} + \frac{1}{10} \mathbf{e}^i)$ using SFM.
 - 4: Project every bid list w.r.t. \mathbf{p}^ε .
 - 5: Subtract $\frac{1}{10} \mathbf{e}^i$ from each of j 's bids and reset price to \mathbf{p} .
-

Lemma 17. *The marginal bids graph G_{A^θ} of A^θ has strictly fewer edges than G_A .*

Lemma 18. *`ShiftProjectUnshift`, given in Algorithm 5, returns a reduction A^θ of A in time $O(T(n) + n|B|f)$, where $T(n)$ is the time it takes to minimise an n -dimensional submodular function.*

Proof. Let A^θ and A^θ denote the allocation problems after executing lines 3 and 4, respectively. First we show that A^θ is a reduction of A . Indeed, by Proposition 14 and Step 1, the residual bundle \mathbf{r} is aggregately demanded at \mathbf{p}^ε in A^θ . Furthermore, for any allocation $(\mathbf{r}^j)_{j \in J} = (\mathbf{t}^j - \mathbf{m}^j)_{j \in J}$ of the residual bundle \mathbf{r} at price \mathbf{p}^ε in A^θ , Lemma 13 implies that $(\mathbf{m}^j + \mathbf{r}^j)_{j \in J}$ is a solution of A . Secondly, Lemma 16 implies that A^θ is a reduction of A^θ . Finally, after line 5, the bid lists are locally valid at \mathbf{p} by Lemma 13. Proposition 14 implies that the residual bundle \mathbf{r} is demanded at some price $\mathbf{p}^\theta \geq \{\mathbf{p}^\varepsilon - \varepsilon \mathbf{e}^i \mid S \subseteq [n]\}$. Note that if a bundle is demanded at \mathbf{p}^θ , it is also demanded at $\lfloor \mathbf{p}^\theta \rfloor$ due to the structure of our integral bid vectors, where $\lfloor \cdot \rfloor$ denotes the operation of rounding each component to the nearest integer. As $\lfloor \mathbf{p}^\theta \rfloor = \mathbf{p}$, the result follows. \square

4.5 The main algorithm

The algorithm `Allocate`, stated in Algorithm 1 above, combines the procedures `NonMarginals`, `UnambiguousMarginals` and `ShiftProjectUnshift` in order to solve the allocation problem described in Section 2.5. Recall that it uses `FindParams` as a subroutine to decide whether

to call `UnambiguousMarginals` or `ShiftProjectUnshift` in each iteration of the loop. Theorem 19 proves correctness and gives a running time bound for `Allocate`. We note that this bound is likely to be pessimistic.

Theorem 19. *Allocate solves the allocation problem in time $O(n^2 j J j(\alpha(n) n j B j + T(n)))$, where $T(n)$ is the time required to minimise an n -dimensional submodular set function.*

Proof. By construction, the marginal bids graph of the initial allocation problem has at most $j J j \binom{n+1}{2}$ edges. Every call to `UnambiguousMarginals` or `ShiftProjectUnshift` strictly reduces the number of edges (by Lemmas 12 and 17). Hence after at most $j J j \binom{n+1}{2}$ iterations of Step 2, the marginal bids graph of the current allocation problem is an empty graph, implying that there are no more marginal bids. In particular, at this point a single call to `NonMarginals` allocates all remaining non-marginal bids and returns a vacuous allocation problem. As all three procedures `NonMarginals`, `UnambiguousMarginals` and `ShiftProjectUnshift` return a reduction in the sense of Definition 5, the solution to the final vacuous allocation problem is also a solution to the original allocation problem. To see the running time guarantee, note that `FindParams` and the procedures `NonMarginals`, `UnambiguousMarginals` and `ShiftProjectUnshift` are each called at most $j J j \binom{n+1}{2} = O(n^2 j J j)$ times, and `FindParams` dominates `NonMarginals` and `UnambiguousMarginals`. \square

Incorporating priorities. The `FindParams` subroutine as stated in Section 4.2 is not fully specified, and different implementations may lead to different inputs for `UnambiguousMarginals` or `ShiftProjectUnshift` when given the same allocation problem. If the input for `ShiftProjectUnshift` depends on the implementation used, ties may be broken differently and thus the target bundle is allocated differently among the bidders.

In order to control which input for `ShiftProjectUnshift` is returned, we propose the use of a priority list consisting of a permutation of all good-bidder pairs $(i, j) \in [n]_0 \times J$. The priority list can be chosen to favour certain types of bidder; for example, if it is considered desirable to bring ‘small’ bidders (having low demand) into the market, we can prioritise their bids, making it slightly more likely that they will be allocated. Moreover, the list may be given as an additional input parameter at the start of `Allocate` or be generated and updated dynamically as `Allocate` runs. The subroutine `FindParams` is replaced by `PriorityParams`, which returns the highest pair (i, j) in the priority list that constitutes a valid input to `ShiftProjectUnshift`, or an input to `UnambiguousMarginals` if no such pair exists. Recall that a pair (i, j) is a valid input for `ShiftProjectUnshift` if i is a cycle-link good in some multi-bidder cycle and j is the label of one of its edges in this cycle. This is the case if the derived graph D_A has a cycle containing the edge from i to the demand cluster (I, j) satisfying $i \succeq I$. For any (undirected) graph G and edge vw , one can check whether G has a cycle containing vw by determining whether its endpoints v and w are connected in the graph $G - vw$ obtained by deleting vw , using breadth first search (BFS). This implies the following subroutine.

Algorithm 6 `PriorityParams`

- 1: Compute the derived graph D_A .
 - 2: **for all** pairs (i, j) in the priority list **do**
 - 3: **if** i is a link good and $i \succeq I$ for some demand cluster (I, j) **then**
 - 4: Check whether the edge from i to (I, j) lies in a cycle by temporarily removing the edge and verifying (using BFS) whether the two endpoints are still connected in the graph.
 - 5: **return** (i, j) if the edge lies in a cycle.
 - 6: **return** some leaf demand cluster (I, j) and the adjacent link good i .
-

To see that `PriorityParams` is well-defined, note that the priority list is a permutation of all possible pairs (i, j) . Hence if the derived graph contains a cycle, the subroutine will return

it, otherwise there exist at least two leaves for Step 2 to choose from. In practice, we can prune the priority list dynamically by removing pairs (i, j) once i is no longer a link good. This is possible because `UnambiguousMarginals` and `ShiftProjectUnshift` do not add edges to the derived graph, and so a good cannot become a link good again at some later point.

5 Conclusions and further work

This paper provides a practical process for running auctions in which bidders can express strong-substitutes preferences using positive and negative bids. Recent working papers [Baldwin and Klemperer, 2021, Lin and Tran, 2017] show that *all* strong-substitutes preferences can be represented using appropriate combinations of these bids. While we show that verifying the validity of bid lists is coNP-complete, useful subclasses of collections of bids are efficiently checkable for validity in practice, and can be checked offline, bidder by bidder, prior to the auction.

Our approach to finding market-clearing prices builds on the long-step steepest-descent procedure proposed by Shioura [2017]. We both analyse his binary search method for determining long steps in our product-mix auction bidding language context, and also propose a novel second method for determining step lengths.

The structure of the product-mix bidding language also facilitates the design of an efficient algorithm for computing an equilibrium allocation of goods to bidders. Our approach contrasts with the previous literature on the problem of allocating supply to bidders at given prices ([Murota, 2003b, Murota and Tamura, 2001] and [Paes Leme and Wong, 2017]), which instead assumes that bidders’ preferences are accessed through an abstract valuation oracle.

All our algorithms approaches are polynomial-time in the size of the bid list inputs and, as we show in the online companion, efficient in practice.

One question for future research is whether we can also exploit the information provided by the bidding language to improve the computational efficiency of the pre-existing submodular function minimisation subroutine that both our price-finding and allocation algorithms use. Another obvious question is the extent to which our methods can be extended to broader classes of valuations.

Acknowledgements

Baldwin and Klemperer were supported by ESRC Grant ES/L003058/1. We thank the reviewers for valuable comments and suggestions, which have been helpful in improving this paper. We are particularly grateful to a referee who both identified a problem in an earlier draft of our paper, and took the trouble to develop an elegant solution to it.

We also thank our colleagues, including Meg Meyer and Duncan Coutts, as well as Tobias Dammers and Andres Löh, who undertook the Haskell implementation and provided much helpful feedback as the algorithms were developed.

[Further acknowledgements to be completed later.]

References

- L. M. Ausubel. An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, 96(3):602–629, June 2006.
- E. Baldwin and P. Klemperer. Understanding preferences: “demand types”, and the existence of equilibrium with indivisibilities. *Econometrica*, 87(3):867–932, May 2019.

- E. Baldwin and P. Klemperer. Proof that the product-mix auction bidding language can represent any substitutes preferences. working paper, 2021. URL https://www.nuffield.ox.ac.uk/economics/Papers/2021/2021-W05_subproof.pdf.
- E. Baldwin, P. W. Goldberg, and P. Klemperer. Tropical intersections and equilibrium (day 2 slides). Hausdorff School on Tropical Geometry and Economics, 2016. URL <http://people.math.gatech.edu/~jyu67/HCM/Baldwin2.pdf>.
- E. Baldwin, M. Bichler, M. Fichtl, and P. Klemperer. Strong substitutes: structural properties, and a new algorithm for competitive equilibrium prices. *Mathematical Programming*, pages 1–33, 2022.
- D. Chakrabarty, P. Jain, and P. Kothari. Provable submodular minimization using Wolfe’s algorithm. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 802–809. Curran Associates, Inc., Cambridge, MA, USA, 2014.
- D. Chakrabarty, Y. T. Lee, A. Sidford, and S. C.-W. Wong. Subquadratic submodular function minimization. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 1220–1231, New York, NY, USA, 2017. ACM.
- V. Danilov, G. Koshevoy, and K. Murota. Discrete convexity and equilibria in economies with indivisible goods and money. *Mathematical Social Sciences*, 41(3):251 – 273, 2001.
- P. W. Goldberg, E. Lock, and F. Marmolejo-Cossío. Learning strong substitutes demand via queries. *ACM Trans. Econ. Comput.*, 10(2), oct 2022. ISSN 2167-8375. doi: 10.1145/3546604. URL <https://doi.org/10.1145/3546604>.
- F. Gul and E. Stacchetti. The English auction with differentiated commodities. *Journal of Economic theory*, 92(1):66–95, 2000.
- S. Iwata. Submodular function minimization. *Mathematical Programming*, 112(1):45–64, 2008.
- H. Jiang. Minimizing convex functions with integral minimizers. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 976–985. SIAM, 2021.
- M. Joswig. *Essentials of tropical combinatorics*, volume 219 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2021.
- A. S. Kelso and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50:1483–1504, 1982.
- P. Klemperer. A new auction for substitutes: Central bank liquidity auctions, the U.S. TARP, and variable product-mix auctions. working paper, Nuffield College, 2008. URL <https://www.nuffield.ox.ac.uk/economics/Papers/2008/substsauc.pdf>.
- P. Klemperer. The product-mix auction: A new auction design for differentiated goods. *Journal of the European Economic Association*, 8(2–3):526–536, 2010.
- P. Klemperer. Product-mix auctions. working paper 2018-W07, Nuffield College, 2018. URL <https://www.nuffield.ox.ac.uk/economics/Papers/2018/2018W07productmix.pdf>.
- Y. T. Lee, A. Sidford, and S. C.-W. Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, FOCS ’15, pages 1049–1065, Washington, DC, USA, 2015. IEEE Computer Society. doi: 10.1109/FOCS.2015.68. URL <http://dx.doi.org/10.1109/FOCS.2015.68>.

- B. Lin and N. M. Tran. Linear and rational factorization of tropical polynomials. *ArXiv preprint*, 2017. URL <https://arxiv.org/abs/1707.03332>.
- D. Maclagan and B. Sturmfels. *Introduction to Tropical Geometry*, volume 161 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2015.
- A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*, volume 1. Oxford University Press, New York, 1995.
- S. T. McCormick. Submodular function minimization. In K. Aardal, G. Nemhauser, and R. Weismantel, editors, *Discrete Optimization*, volume 12 of *Handbooks in Operations Research and Management Science*, pages 321–391. Elsevier, 2005. doi: [https://doi.org/10.1016/S0927-0507\(05\)12007-6](https://doi.org/10.1016/S0927-0507(05)12007-6). URL <https://www.sciencedirect.com/science/article/pii/S0927050705120076>.
- G. Mikhalkin. Decomposition into pairs-of-pants for complex algebraic hypersurfaces. *Topology*, 43(5):1035–1065, 2004. ISSN 0040-9383. doi: <https://doi.org/10.1016/j.top.2003.11.006>.
- P. Milgrom and B. Strulovici. Substitute goods, auctions, and equilibrium. *Journal of Economic Theory*, 144(1):212–247, 2009.
- K. Murota. On steepest descent algorithms for discrete convex functions. *SIAM J. on Optimization*, 14(3):699–707, March 2003a. ISSN 1052-6234. doi: 10.1137/S1052623402419005.
- K. Murota. *Discrete Convex Analysis*. SIAM Monographs on Discrete Mathematics and Applications, 2003b.
- K. Murota. Discrete convex analysis: A tool for economics and game theory. *Journal of Mechanism and Institution Design*, 1(1):151–273, 2016.
- K. Murota and A. Shioura. M-convex function on generalized polymatroid. *Mathematics of Operations Research*, 24(1):95–105, 1999.
- K. Murota and A. Shioura. Exact bounds for steepest descent algorithms of L-convex function minimization. *Operations Research Letters*, 42(5):361–366, 2014.
- K. Murota and A. Tamura. Application of M-convex submodular flow problem to mathematical economics. In *Proceedings of the 12th International Symposium on Algorithms and Computation*, ISAAC '01, page 14–25, Berlin, Heidelberg, 2001. Springer-Verlag. ISBN 3540429859.
- K. Murota, A. Shioura, and Z. Yang. Computing a Walrasian equilibrium in iterative auctions with multiple differentiated items. In L. Cai, S.-W. Cheng, and T.-W. Lam, editors, *Algorithms and Computation*, pages 468–478, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. ISBN 978-3-642-45030-3.
- K. Murota, A. Shioura, and Z. Yang. Time bounds for iterative auctions: A unified approach by discrete convex analysis. *Discrete Optimization*, 19:36–62, 2016. doi: 10.1016/j.disopt.2016.01.001.
- R. Paes Leme. Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, 106:294–316, 2017.
- R. Paes Leme and S. C.-W. Wong. Computing Walrasian equilibria: Fast algorithms and structural properties. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 632–651, Philadelphia, PA, USA, 2017. Society for Industrial and Applied Mathematics.

- A. Schrijver. *Theory of Linear and Integer Programming*. Interscience series in discrete mathematics and optimisation. John Wiley & Sons, 1998.
- A. Shioura. Minimization of an M-convex function. *Discrete Applied Mathematics*, 84(1):215–220, 1998. ISSN 0166-218X. doi: [https://doi.org/10.1016/S0166-218X\(97\)00140-6](https://doi.org/10.1016/S0166-218X(97)00140-6). URL <https://www.sciencedirect.com/science/article/pii/S0166218X97001406>.
- A. Shioura. Algorithms for L-convex function minimization: Connection between discrete convex analysis and other research fields. *Journal of the Operations Research Society of Japan*, 60(3): 216–243, 2017.
- A. Shioura and A. Tamura. Gross substitutes condition and discrete concavity for multi-unit valuations: a survey. *Journal of the Operations Research Society of Japan*, 58(1):61–103, 2015.
- R. E. Tarjan and J. van Leeuwen. Worst-case analysis of set union algorithms. *Journal of the ACM*, 31(2):245–281, March 1984. doi: 10.1145/62.2160.
- S. van der Walt, S. C. Colbert, and G. Varoquaux. The NumPy array: A structure for efficient numerical computation. *Computing in Science Engineering*, 13(2):22–30, March 2011.

A Experiments

In order to evaluate the practical running time of our allocation algorithm, we run experiments on various numbers of goods, bidders and bids. We use our own Python implementation of the product-mix auction, available at <https://github.com/edwijnlock/product-mix>. Some effort was made to optimise for speed by exploiting fast matrix operations provided by the NumPy package [van der Walt et al., 2011]. Furthermore, in an effort to implement submodular minimisation efficiently, the Fujishige-Wolfe algorithm was implemented in combination with a memoization technique to reduce the number of submodular function queries.

A.1 Generating test data

We describe a procedure to generate a valid list of positive and negative bids at points within the lattice $[M]^n$ and a bundle \mathbf{x} that is demanded in aggregate by these bids at prices $\mathbf{p} = \frac{1}{2}M\mathbf{e}^{[n]}$ in the centre of the lattice. In our experiments, we fix $M = 100$. The bids are generated in such a way that for parameter q at least q bids are marginal between two or more goods at \mathbf{p} and we pick \mathbf{x} such that \mathbf{p} is the component-wise minimal market-clearing price vector. For any permutation π of $[n]$, let $\mathbf{p}_\pi = \mathbf{p} + \sum_{i \in [n]} \frac{\varepsilon}{2i} \mathbf{e}^{\pi(i)}$ for some $\varepsilon < 0.1$. Note that a unique bundle is demanded at \mathbf{p}_π for any permutation π .

Generated in this way, every bid list has $2.5q$ bids in expectation. The procedure `GenerateList` is repeated for each bidder, so the total number of bids generated is $B = 2.5qm$.

Lemma 20. *The bid list generated by `GenerateList` is valid.*

Proof. Note that the union of finitely many valid bid lists is again a valid bid list. Hence it suffices to show that the list consisting of one negative bid \mathbf{b} and three positive bids $\mathbf{b} - \lambda_i \mathbf{e}^i$, $\mathbf{b} - \lambda_j \mathbf{e}^j$ and $\mathbf{b} + \lambda \mathbf{1}$ is valid. To see this, apply Theorem 19 from the main paper. \square

A.2 Testing the algorithm

We run `Allocate` on allocation problems with bid lists generated by `GenerateList` for different numbers of goods n , bidders m and bids B . It should be noted that, as expected, the value of M has no discernable impact on the running time of the allocation algorithm and thus can be fixed to $M = 100$. We perform three pairs of experiments in which we vary one of the parameters n, m, q and fix the other two to realistic values.

Algorithm 7 GenerateList(n, M, q)

- 1: **Initialise:** Empty bid list \mathbf{b} and bundle $\mathbf{x} = \mathbf{0}$.
 - 2: **Repeat** the following q times:
 - 3: Pick a subset S of goods with $|S| = 2$ from $[n]_0$ and flip a fair coin.
 - 4: **if** the coin lands on heads **then**
 - 5: Pick a positive bid that is marginal on goods S at \mathbf{p} and add it to B .
 - 6: Pick any good $i \in S$ uniformly at random and increment x_i by one.
 - 7: **else**
 - 8: Generate a negative bid \mathbf{b} that is marginal on goods S at \mathbf{p} , as well as the following three positive bids.
 - 9: Pick two goods $i, j \in [n]$ and add a bid at points $\mathbf{b} + \lambda_i \mathbf{e}^i$ and $\mathbf{b} + \lambda_j \mathbf{e}^j$ for some $1 \leq \lambda_i \leq b_i - 1$ and $1 \leq \lambda_j \leq b_j - 1$.
 - 10: Add a bid at $\mathbf{b} + \lambda \mathbf{1}$ for some $1 \leq \lambda \leq \min_{i \in [n]} M - b_i$.
 - 11: Pick a permutation π of $[n]$ and increment \mathbf{x} by the bundle aggregately demanded at \mathbf{p}_π by the four bids just generated.
 - 12: **if** \mathbf{p} is the component-wise minimal market-clearing price vector of \mathbf{x} **then**
 - 13: **return** B .
 - 14: **else**
 - 15: Repeat the algorithm.
-

1. For the first pair of experiments, we vary the average number of bids by running GenerateList with values $q = 20, 40, 60, \dots, 500$. We fix the number of bidders to $m = 5$ and the number of goods to $n = 2$ and $n = 10$, respectively.
2. For the second pair of experiments, we vary the number of goods from $n = 10$ to 50 in steps of 5 and fix q to 50 and 100, respectively.
3. Finally, the last pair of experiments varies the number of bidders from $m = 2$ to $m = 20$ in steps of 1 with two goods $n = 2$, while the bid numbers are fixed by setting q to 50 and 100, respectively.


For each data point (n, m, q) , 50 allocation problems with n goods, m bidders and $2.5q$ bids per bidder (in expectation) are generated. The Allocate algorithm is then timed on each allocation problem and the average over all 50 times is recorded.

A.3 Results


The outcomes of the three pairs of experiments are shown in Figures 7 to 9. The experimental data corroborates the running time bounds for Allocate given in Theorem 9: the algorithms is linear in the number of bids and quadratic in the number of goods. Figure 9 suggests that our algorithm runs in quadratic time on our generated allocation problems, which is in line with our theoretical bound, as the total number of bids $B = 2.5qm$ is linear in m . Overall, we see that our allocation algorithm runs quickly even when presented with a large number of bids.

B Additional proofs

We give some definitions that are used in the proofs below. Here and in the following, we use terminology and material developed in [Baldwin and Klemperer, 2019], drawing on the literature in tropical geometry, especially [Mikhalkin, 2004]; see also the textbooks by Maclagan and Sturmfels [2015] and Joswig [2021].




(a) Parameters: $n = 2; m = 5; M = 100$.




(b) Parameters: $n = 10; m = 5; M = 100$.

Figure 7: Testing `Allocate` by varying the number B of bids for each bidder while keeping all other parameters fixed. We increase q from 20 to 500 in steps of 20, fix the number of bidders at $m = 5$ and set the number of goods to $n = 2$ (a) and $n = 10$ (b), respectively.




(a) Parameters: $m = 5; M = 100; q = 50$.




(b) Parameters: $m = 5; M = 100; q = 100$.

Figure 8: Testing `Allocate` by varying the number n of goods. We increase n from 5 to 50 in steps of 5, fix the number of bidders at $m = 5$ and $M = 100$ and set the number of bids to $q = 50$ (a) and $q = 100$ (b), respectively.



(a) Parameters: $n = 2; M = 100; q = 50$.



(b) Parameters: $n = 2; M = 100; q = 100$.

Figure 9: Testing `Allocate` by varying the number m of bidders. We increase m from 2 to 20 in steps of 1, fix the number of goods at $n = 2$ and set q to 50 (a) and 100 (b), respectively.

Geometric objects for valuations. The Locus of Indifference Prices (LIP) L_u for valuation u with demand correspondence D_u is defined as

$$L_u = \{p \in \mathbb{R}^n \mid \exists j \in D_u(p) \text{ s.t. } p_j > 1\}.$$

This set was identified as an object of mathematical interest by Mikhalkin [2004] and is also known as a “tropical hypersurface” [Maclagan and Sturmfels, 2015, Joswig, 2021].

Recall that a *rational polyhedral complex* is a finite collection of polyhedrons, closed under

taking faces; non-empty intersections of these polyhedrons must be faces of them, and these faces must have rational slope [Maclagan and Sturmfels, 2015, Joswig, 2021]. L_u can be decomposed into an $(n - 1)$ -dimensional *rational polyhedral complex* \mathcal{C}_u [Mikhalkin, 2004, Proposition 2.1]. We call the $(n - 1)$ -dimensional polyhedra in this complex its *facets*. Let w_u be the weight function obtained as follows: if \mathbf{x} and \mathbf{y} denote the two bundles demanded on either side of a facet F , then the weight $w_u(F)$ of F is the greatest common divisor of the entries of $\mathbf{x} - \mathbf{y}$. We recall the following proposition from [Baldwin and Klemperer, 2019].

Proposition 21 ([Baldwin and Klemperer, 2019, Proposition 2.4]). *The change in demand, as prices change between the UDRs on either side of facet F , is $w_u(F)$ times the primitive integer vector that is normal to F and points in the opposite direction to the change in price.*

In particular we identify the *strong-substitutes vectors*: those non-zero vectors in \mathbb{Z}^n which have at most one $+1$ entry, at most one -1 entry, and no other non-zero entries. (These vectors are also distinguished in the literature on discrete convex analysis, being the edges of discrete generalised polymatroids; see e.g. [Shioura and Tamura, 2015, Theorem 3.9]). Recall that we assume throughout this paper that valuations are concave-extensible. This implies the following theorem (see also [Baldwin and Klemperer, 2019, Proposition 3.8]).

Theorem 22 ([Shioura and Tamura, 2015, Theorems 4.1 and 4.10]). *A valuation u is strong-substitutes if and only if the facets of L_u are normal to strong-substitutes vectors.*

Following [Mikhalkin, 2004, Definition 3], we say that a weighted polyhedral complex (\mathcal{C}, w) is *balanced* if, for every $(n - 2)$ -cell G of \mathcal{C} , the weights $w(F_j)$ on the facets F_1, \dots, F_l that contain G , and primitive integer normal vectors \mathbf{v}_{F_j} for these facets that are defined by a fixed rotational direction about G , satisfy $\sum_{j=1}^l w(F_j)\mathbf{v}_{F_j} = 0$. Balancing is essential to what [Baldwin and Klemperer, 2019, Theorem 2.14] call the ‘‘Valuation-Complex Equivalence Theorem’’:

Theorem 23 ([Mikhalkin, 2004, Remark 2.3 and Proposition 2.4]). *Suppose that (\mathcal{C}, w) is an $(n - 1)$ -dimensional weighted rational polyhedral complex in \mathbb{R}^n with positive weights only, and that L is the union of the cells in \mathcal{C} . There exists a finite set $A \subset \mathbb{Z}^n$ and a concave-extensible valuation $u : A \rightarrow \mathbb{R}$ such that $L_u = L$ and $w_u = w$, if and only if (\mathcal{C}, w) is balanced. Moreover in this case, then for any price vector \mathbf{p} not contained in L , there exists a unique concave-extensible u satisfying also $u(\mathbf{0}) = 0$ and $D_u(\mathbf{p}) = f\mathbf{0}g$.*

Geometric objects for bids. In analogy to the definitions above, we now define geometric objects for product-mix bids and bid lists. The Locus of Indifference Prices (LIP) for a bid \mathbf{b} is

$$L_{\mathbf{b}} = \{ \mathbf{p} \in \mathbb{R}^n \mid \arg \max_{i \in [n]_0} (b_i - p_i) > 1 \}.$$

Similarly, the LIP of a list of bids B is given by

$$L_B = \{ \mathbf{p} \in \mathbb{R}^n \mid \arg \max_{j \in B} D_B(\mathbf{p})_j > 1 \}.$$

Any L_B can be decomposed into an $(n - 1)$ -dimensional rational polyhedral complex \mathcal{C}_B . We call the $(n - 1)$ -dimensional polyhedra in this complex its *facets*. Endow every facet F of \mathcal{C}_B with a weight, as follows:

$$w_B(F) := \sum_{\mathbf{b} \in B \text{ and } L_{\mathbf{b}} \subset F} f w(\mathbf{b}) g.$$

B.1 Proof of Theorem 1

Proof of Theorem 1. The equivalence of (1) and (2) follows from Proposition 6 if we set $P = \mathbb{R}^n$. The implication (3) \Rightarrow (1) is well-known; Shioura and Tamura [2015, Theorem 7.1] and Murota et al. [2016, Theorem 4.3], for instance, show that f_u is a polyhedral L^\flat -convex function and these functions are defined as a subclass of convex functions (cf. [Murota et al., 2016]). We now show (2) \Rightarrow (3) using Theorem 23. Let (\mathcal{C}_B, w_B) be the weighted polyhedral complex associated with B . Note that as well as showing that balancing holds, we must also demonstrate that all weights of the polyhedral complex are positive, which is the setting for Theorem 23.

Observe first that every facet of \mathcal{C}_B has a strong-substitutes vector as its normal vector, as this property holds for any individual $L_{\mathbf{b}}$. Moreover, the weighted polyhedral complex $(\mathcal{C}_{\mathbf{b}}, w_{\mathbf{b}})$ associated with a single bid is balanced. If $w(\mathbf{b}) = 1$, then this follows because $(\mathcal{C}_{\mathbf{b}}, w_{\mathbf{b}})$ is the weighted polyhedral complex associated with a simple valuation for at most one unit of any good, and apply Theorem 23. This then extends to the case $w(\mathbf{b}) = -1$ because changing the sign of the weights of all facets does not affect balancing. Then (\mathcal{C}_B, w_B) is also balanced, as it is the complex corresponding to the union of balanced LIPs.

Next we see that all weights of (\mathcal{C}_B, w_B) are non-negative. Let F be a facet of \mathcal{C}_B and fix a point $\boldsymbol{\rho}$ in the relative interior of F . Then, since the normal vector to F is a strong-substitutes vector and by Proposition 21, every bid is either non-marginal at $\boldsymbol{\rho}$ or marginal between the same two distinct goods, $i, i^\flat \in [n]_0$. Let B_{ii^\flat} denote the bids that are marginal between i, i^\flat at $\boldsymbol{\rho}$. Note that we have $F \in L_{\mathbf{b}}$ if and only if \mathbf{b} is marginal on i, i^\flat , so $w_B(F) = \sum_{\mathbf{b} \in B_{ii^\flat}} w(\mathbf{b}) \geq 0$ (by assumption (2)).

(\mathcal{C}_B, w_B) may have zero-weighted facets: let the set of these be \mathcal{C}_B^0 and write $\mathcal{C}_B^+ := \mathcal{C}_B \setminus \mathcal{C}_B^0$. Then \mathcal{C}_B^+ inherits from \mathcal{C}_B the structure of a polyhedral complex. Moreover, if we write w_B^+ for the restriction of w_B to the facets of \mathcal{C}_B^+ , then (\mathcal{C}_B^+, w_B^+) is balanced: removing 0-weighted facets does not affect this criterion.

So, we write L_B^+ for the union of the cells in \mathcal{C}_B^+ . Let $\bar{\boldsymbol{\rho}}$ be a price vector whose components are strictly larger than the components of any bid, so that $D_B(\bar{\boldsymbol{\rho}}) = \mathbf{0}$. By Theorem 23, there exists a unique concave-extensible valuation u such that $L_u = L_B^+$, $w_u = w_B^+$, $u(\mathbf{0}) = 0$ and $D_u(\bar{\boldsymbol{\rho}}) = \mathbf{0}$. We observed above that all facets of \mathcal{C}_B , and hence of L_u , are normal to a strong-substitutes vector, so u is a strong-substitutes valuation by Theorem 22.

By Proposition 21, the facets and weights of L_u define all changes in demand between UDRs associated with valuation u . If $w(\mathbf{b})$ is positive then $L_{\mathbf{b}}$ has the property that facets and weights define all changes in demand between UDRs (see also [Baldwin and Klemperer, 2019, Section 2.2]), and it is easy to see that the same follows for negative-weighted bids, and thus extends by aggregation to our positive-weighted polyhedral complex (\mathcal{C}_B^+, w_B^+) . So $D_u(\bar{\boldsymbol{\rho}}) = D_B(\bar{\boldsymbol{\rho}})$ implies $D_u(\boldsymbol{\rho}) = D_B(\boldsymbol{\rho})$ for all UDR prices $\boldsymbol{\rho} \in \mathbb{R}^n$, and hence for all prices, since u is concave-extensible and $D_B(\boldsymbol{\rho})$ is defined to be discrete convex.

We now show that $f_B = f_u$. For any price $\boldsymbol{\rho}$ at which demand is unique for every bid $\mathbf{b} \in B$, let \mathbf{x} be the bundle demanded and let $(i(\mathbf{b}))_{\mathbf{b} \in B}$ be the list of goods allocated to the bids. That is, each bid $\mathbf{b} \in B$ uniquely demands $i(\mathbf{b})$ and $\mathbf{x} = \sum_{\mathbf{b} \in B} w(\mathbf{b}) \mathbf{e}^{i(\mathbf{b})}$. Then we claim that

$$u(\mathbf{x}) = \sum_{\mathbf{b} \in B} w(\mathbf{b}) b_{i(\mathbf{b})}. \quad (7)$$

This then immediately implies $f_B(\boldsymbol{\rho}) = f_u(\boldsymbol{\rho})$ for any price $\boldsymbol{\rho} \in \mathbb{R}^n$ at which demand is unique for every bid, since

$$f_B(\boldsymbol{\rho}) = \sum_{\mathbf{b} \in B} w(\mathbf{b}) \max_{i \in [n]_0} (b_i - p_i) = \sum_{\mathbf{b} \in B} w(\mathbf{b}) (b_{i(\mathbf{b})} - p_{i(\mathbf{b})}) = u(\mathbf{x}) \quad \mathbf{x} = \boldsymbol{\rho} = f_u(\boldsymbol{\rho}).$$

But these prices are dense in \mathbb{R}^n so it follows by continuity of these piecewise-linear functions that $f_B(\boldsymbol{\rho}) = f_u(\boldsymbol{\rho})$ for all $\boldsymbol{\rho} \in \mathbb{R}^n$.

To see that (7) holds, first observe that it holds for the demand and allocation at price $\bar{\mathbf{p}}$: here, by definition $i(\mathbf{b}) = 0$ for all $\mathbf{b} \succeq B$, so $\mathbf{x} = \mathbf{0}$; and $u(\mathbf{0}) = 0$.

Next, consider price vectors $\mathbf{p}^\theta, \mathbf{p}^{\theta\theta}$ at which bundles $\mathbf{x}^\theta, \mathbf{x}^{\theta\theta}$ are uniquely demanded, via unique allocations to bids $(i^\theta(\mathbf{b}))_{\mathbf{b} \succeq B}$ and $(i^{\theta\theta}(\mathbf{b}))_{\mathbf{b} \succeq B}$ respectively. Suppose moreover that there exists a price \mathbf{p} satisfying $\mathbf{x}^\theta, \mathbf{x}^{\theta\theta} \succeq D_B(\mathbf{p}) = D_u(\mathbf{p})$ with $i^\theta(\mathbf{b}), i^{\theta\theta}(\mathbf{b}) \succeq \arg \max_{i \in [n]_0} (b_i - p_i)$ for all $\mathbf{b} \succeq B$: this holds if \mathbf{p}^θ and $\mathbf{p}^{\theta\theta}$ are in the same or neighbouring connected components of the complement of ∂_B in \mathbb{R}^n . Then we have:

$$u(\mathbf{x}^\theta) - \mathbf{p} \cdot \mathbf{x}^\theta = u(\mathbf{x}^{\theta\theta}) - \mathbf{p} \cdot \mathbf{x}^{\theta\theta} \quad (8)$$

and

$$f_B(\mathbf{p}) = \sum_{\mathbf{b} \succeq B} w(\mathbf{b})(b_{i^\theta(\mathbf{b})} - p_{i^\theta(\mathbf{b})}) = \sum_{\mathbf{b} \succeq B} w(\mathbf{b})(b_{i^{\theta\theta}(\mathbf{b})} - p_{i^{\theta\theta}(\mathbf{b})}). \quad (9)$$

Suppose (7) holds for bundle \mathbf{x}^θ and $(i^\theta(\mathbf{b}))_{\mathbf{b} \succeq B}$. Then (8) and (9) imply that (7) holds for $\mathbf{x}^{\theta\theta}$ and $(i^{\theta\theta}(\mathbf{b}))_{\mathbf{b} \succeq B}$. But now we can use the fact that (7) holds for the demand and allocation at price $\bar{\mathbf{p}}$ to infer that (7) for the demand and allocation at any price at which demand is unique, by induction. \square

B.2 Proof of Proposition 6

We make use of the following technical observation to prove Proposition 6.

Observation 24. *Let $f : [0, 1] \rightarrow \mathbb{R}$ be a continuous, piecewise-linear function. If its slope is non-decreasing as we move from 0 to 1, f is convex. This implies that if f is not convex, there exists a point λ at which two linear segments meet and the slope decreases.*

Proof of Proposition 6. Recall that f_B restricted to P is convex if and only if

$$f_B(\lambda \mathbf{p} + (1 - \lambda) \mathbf{q}) \geq \lambda f_B(\mathbf{p}) + (1 - \lambda) f_B(\mathbf{q}), \quad \forall \mathbf{p}, \mathbf{q} \succeq P \text{ and } \lambda \in [0, 1]. \quad (10)$$

Suppose that \mathbf{p} and \mathbf{q} are two non-marginal prices in neighbouring UDRs such that for some $\lambda \in (0, 1)$, the point $\mathbf{r} = \lambda \mathbf{p} + (1 - \lambda) \mathbf{q}$ lies on the interior of a facet separating the UDRs. This implies two possibilities for each bid $\mathbf{b} \succeq B$. Either \mathbf{b} non-marginally demands the same good at \mathbf{p} and \mathbf{q} (and thus also at \mathbf{r}). Or \mathbf{b} non-marginally demands i at \mathbf{p} and i^θ at \mathbf{q} , and is marginal (only) on i and i^θ at \mathbf{r} . Moreover, all bids of the latter kind are marginal on the same two goods.

Let $B_{i i^\theta}$ denote the bids that demand the set $\{i, i^\theta\}$ at \mathbf{r} . We show that f_B satisfies (10) for \mathbf{p}, \mathbf{q} and any $\lambda \in [0, 1]$ if and only if the weights of the bids marginal on i and i^θ sum to a non-negative number, i.e. if $\sum_{\mathbf{b} \in B_{i i^\theta}} w(\mathbf{b}) \geq 0$. By construction, we have

$$\begin{aligned} f_B(\mathbf{p}) &= \sum_{\mathbf{b} \in B_{i i^\theta}} w(\mathbf{b})(b_i - p_i) + \sum_{\mathbf{b} \notin B_{i i^\theta}} w(\mathbf{b}) \max_{j \in [n]_0} (b_j - p_j), \\ f_B(\mathbf{q}) &= \sum_{\mathbf{b} \in B_{i i^\theta}} w(\mathbf{b})(b_{i^\theta} - q_{i^\theta}) + \sum_{\mathbf{b} \notin B_{i i^\theta}} w(\mathbf{b}) \max_{j \in [n]_0} (b_j - q_j), \\ f_B(\mathbf{r}) &= \sum_{\mathbf{b} \in B_{i i^\theta}} w(\mathbf{b})(b_i - r_i) + \sum_{\mathbf{b} \notin B_{i i^\theta}} w(\mathbf{b}) \max_{j \in [n]_0} (b_j - r_j). \end{aligned}$$

Hence the inequality

$$f_B(\mathbf{r}) \geq \lambda f_B(\mathbf{p}) + (1 - \lambda) f_B(\mathbf{q}) \quad (11)$$

holds if and only if

$$\sum_{\mathbf{b} \in B_{i i^\theta}} w(\mathbf{b}) [(b_i - q_i) - (b_{i^\theta} - q_{i^\theta})] \geq 0. \quad (12)$$

Note that the term $(b_i - q_i) - (b_{i^0} - q_{i^0})$ is strictly negative, as \mathbf{b} demands i^0 and not i at \mathbf{q} . Moreover, this term is the same for all bids $\mathbf{b} \succeq B_{ii^0}$, as i and i^0 are both demanded at \mathbf{r} . This implies by (12) that $\sum_{\mathbf{b} \succeq B_{ij}} w(\mathbf{b}) = 0$ if and only if (11) holds.

Now we prove our main statement. Suppose f_B restricted to P is not convex. Then there exist $\mathbf{p}, \mathbf{q} \succeq P$ and $\lambda \succeq [0, 1]$ that violate (10). Due to continuity of f_B , we can perturb \mathbf{p} and \mathbf{q} slightly so that (10) is still violated, \mathbf{p}, \mathbf{q} are non-marginal and the line segment $[\mathbf{p}, \mathbf{q}]$ crosses only interiors of facets of the LIP L_B . We can assume without loss of generality that \mathbf{p} and \mathbf{q} are in neighbouring UDRs; otherwise, we can apply to find two non-marginal prices on the interior of the line segment for which (10) fails. Thus we can apply the above to see the implication 2) \rightarrow 1).

For the converse (1) \rightarrow 2), suppose there exist \mathbf{r} and $i, i^0 \succeq [n]_0$ such that the weights of bids marginal on i, i^0 at \mathbf{r} sum to a negative number. Without loss of generality, we can assume that all marginal bids at \mathbf{r} are marginal only on goods i, i^0 , by subtracting $\delta \mathbf{e}^{F_i, i^0 \mathbf{g}}$ from \mathbf{r} for some infinitesimal positive value of δ if necessary. Hence \mathbf{r} lies in the interior of a facet F and for small enough $\varepsilon > 0$, the points $\mathbf{p} = \mathbf{r} + \varepsilon \mathbf{e}^i - \varepsilon \mathbf{e}^{i^0}$ and $\mathbf{q} = \mathbf{r} - \varepsilon \mathbf{e}^i + \varepsilon \mathbf{e}^{i^0}$ lie in P and in neighbouring UDRs separated by F . By the above, this implies that \mathbf{p} and \mathbf{q} with $\lambda = 1/2$ violate (10) and we are done. \square

B.3 Proof of Lemma 9

Proof of Lemma 9. We show that f_B is convex on $B(\mathbf{p}, \varepsilon/2)$ by verifying that it satisfies midpoint convexity. Note that f_B is linear on the line segment connecting \mathbf{p} to \mathbf{q} , for any $\mathbf{q} \succeq B(\mathbf{p}, \varepsilon/2)$. Indeed, f_B is the sum of terms $w(\mathbf{b}) \max_{i \succeq [n]_0} (b_i - p_i)$ and it suffices to show that each term is linear. Hence fix \mathbf{b} and define $h(\mathbf{p}) = w(\mathbf{b}) \max_{i \succeq [n]_0} (b_i - p_i)$ as well as $\mathbf{r} = \theta \mathbf{p} + (1 - \theta) \mathbf{q}$. By Observation 7, the goods demanded by \mathbf{b} at \mathbf{q}, \mathbf{r} and \mathbf{p} satisfy $I_{\mathbf{q}} \cap I_{\mathbf{r}} = I_{\mathbf{p}}$. Hence, for any $i \succeq I_{\mathbf{q}}$, we have $h(\mathbf{p}) = w(\mathbf{b})(b_i - p_i) = \theta h(\mathbf{p}) + (1 - \theta)h(\mathbf{q})$. Now fix $\mathbf{q}, \mathbf{q}^0 \succeq B(\mathbf{p}, \varepsilon/2)$ and choose $\theta > 0$ so that $\mathbf{r} = \theta \mathbf{p} + (1 - \theta) \mathbf{q}$, $\mathbf{r}^0 = \theta \mathbf{p} + (1 - \theta) \mathbf{q}^0$ and $(\mathbf{r} + \mathbf{r}^0)/2$ are in $B(\mathbf{p}, \delta)$. As f_B is convex on $B(\mathbf{p}, \delta)$, by assumption, we have $f(\frac{1}{2}(\mathbf{r} + \mathbf{r}^0)) \leq \frac{1}{2}f(\mathbf{r}) + \frac{1}{2}f(\mathbf{r}^0)$. Secondly, we have $f(\mathbf{r}) = \theta f(\mathbf{p}) + (1 - \theta)f(\mathbf{q})$, $f(\mathbf{r}^0) = \theta f(\mathbf{p}) + (1 - \theta)f(\mathbf{q}^0)$ and $f(\frac{1}{2}(\mathbf{r} + \mathbf{r}^0)) = \theta f(\mathbf{p}) + (1 - \theta)f(\frac{1}{2}(\mathbf{q} + \mathbf{q}^0))$ due to the linearity of f_B on the line segments connecting \mathbf{p} to \mathbf{q}, \mathbf{q}^0 and $(\mathbf{q} + \mathbf{q}^0)/2$. This implies midpoint convexity, $f(\frac{1}{2}(\mathbf{q} + \mathbf{q}^0)) \leq \frac{1}{2}f(\mathbf{q}) + \frac{1}{2}f(\mathbf{q}^0)$. \square

B.4 Proof of Proposition 14

In order to prove Proposition 14, we first prove in Lemma 25 that the demand of a bid list that is locally valid around integral prices coincides, in the local neighbourhood of these prices, with the demand of some strong-substitutes valuation (in analogy to the third statement in Theorem 1). Working directly with the valuations that correspond locally to our bid lists, we then establish two technical lemmas (Lemmas 26 and 28) that illustrate the robustness of demand when we perturb valuations by some small value. We recall that a vector is strong-substitutes if it has at most one +1 entry, at most one -1 entry, and no other non-zero entries. In Lemma 28 and the proof of Proposition 14, we will refer to the *aggregate valuation* of two valuations. The aggregation of two bid lists is defined as the concatenation of the two individual lists. For any two strong-substitutes valuations u^1 and u^2 associated with valid bid lists B^1 and B^2 , we define the *aggregate valuation* $u^{1,2}$ as the strong-substitutes valuation associated with the aggregation of the two bid lists B^1 and B^2 ; note that the existence of $u^{1,2}$ is guaranteed by Theorem 1.¹⁴

Lemma 25. *Let B be an integral bid list that is ε -valid at integral prices \mathbf{p} with $\varepsilon < 1/2$. There exists a strong-substitutes valuation u such that $D_B(\mathbf{q}) = D_u(\mathbf{q})$ for all $\mathbf{q} \succeq B(\mathbf{p}; \varepsilon)$.*

¹⁴Note that $u^{1,2}$ can be expressed as $u^{1,2}(\mathbf{x}) = \max_{\mathbf{y} \succeq Z^n} (u^1(\mathbf{y}) + u^2(\mathbf{y} - \mathbf{x}))$ for all $\mathbf{x} \succeq Z^n$. This is stated, for instance, in footnote 20 of [Baldwin and Klempner, 2019], and is related to the *aggregate cost function* in [Murota, 2003b, Chapter 11, page 335]. We do not make use of this way of expressing $u^{1,2}$ in the following proofs.

Proof of Lemma 25. Let H be the collection of hyperplanes obtained by taking the affine hull of all facets in the polyhedral complex \mathcal{C}_B of B that intersect $B(\mathbf{p}; \varepsilon)$. Let $H \supseteq H$ be a hyperplane obtained from facet F of \mathcal{C}_B . As F is normal to a strong-substitutes vector and the bid entries are integral, we can write $H = \{ \mathbf{q} \in \mathbb{R}^n \mid q_i - q_j = c \}$ for some $i, j \in [n]_0$ and some integer c . For any $\mathbf{q} \in B(\mathbf{p}; \varepsilon) \setminus H$ we have $p_i - p_j - 2\varepsilon < q_i - q_j < p_i - p_j + 2\varepsilon$. Due to integrality of c and $\varepsilon < 1/2$, we have $q_i - q_j = p_i - p_j = c$. Hence every $H \in H$ contains \mathbf{p} .

We now construct a second rational polyhedral complex (\mathcal{C}, w) as follows. Take the union of all hyperplanes in H , and decompose this union into a polyhedral complex to obtain \mathcal{C} . The cells of this complex are intersections of hyperplanes in H . Every cell (and so every facet and $(n-2)$ -dimensional cell) of \mathcal{C} contains \mathbf{p} , as \mathbf{p} is contained in every $H \in H$. So, in particular, $F \setminus P \neq \emptyset$ for every facet F of \mathcal{C} . Moreover, for each facet F of \mathcal{C} , either $F \setminus P = F^0 \setminus P$ for some facet F^0 of \mathcal{C}_B , or $F \setminus P$ has at most $(n-2)$ -dimensional intersection with the facets of \mathcal{C}_B ; this holds because the facets of \mathcal{C}_B meet facets with distinct affine spans at every point on their boundaries. This allows us to define the weight function w as follows. For any facet F of \mathcal{C} , we let $w(F) = w_B(F^0)$ if $F \setminus P = F^0 \setminus P$ for some facet F^0 of \mathcal{C}_B , and we let $w(F) = 0$ otherwise. Note that (\mathcal{C}, w) is balanced within P , since (\mathcal{C}_B, w_B) is balanced (see also the proof of Theorem 1): dividing a facet of a balanced polyhedral complex into two equally-weighted facets, and adding zero-weighted facets, does not affect balancing. But every $(n-2)$ -cell of \mathcal{C} , and every facet of \mathcal{C} containing this $(n-2)$ -cell, has non-zero intersection with P . So (\mathcal{C}, w) is balanced.

The facets of \mathcal{C}_B that intersect with P have non-negative weight due to the local validity of B at \mathbf{p} (by the same argument as in the proof of Theorem 1). Hence, the weight function w is non-negative by construction. Now let (\mathcal{C}^+, w^+) be obtained from (\mathcal{C}, w) by removing all zero-weighted facets, and let L^+ be the union of the facets of \mathcal{C}^+ . Note that $L_B \setminus P = L^+ \setminus P$; and if F^0 is a facet of L_B and F^+ is a facet of L^+ such that $\dim(F \setminus F^+ \setminus P) = n-1$, then $w(F^0) = w^+(F^+)$. Fix some prices $\mathbf{q} \in P$, $\mathbf{q} \notin L^+$, and let \mathbf{x} be the bundle uniquely demanded by B at \mathbf{q} . By Theorem 23, there exists a valuation function u such that $D_u(\mathbf{q}) = \mathbf{x}$, $L_u = L^+$ and $w_u = w^+$. Using the same argument as in the proof of Theorem 1, we have that the normal vectors and weights of the facets of L_B fully specify how demand of B changes in P . Hence, by Proposition 21, it follows that $D_B(\mathbf{q}^0) = D_u(\mathbf{q}^0)$ for all $\mathbf{q}^0 \in P$. Moreover, as the facets of L_u are normal to strong-substitutes vectors, the valuation u is strong-substitutes by Theorem 22. \square

Lemma 26. Suppose that $G = \begin{pmatrix} G_1 \\ G_2 \end{pmatrix}$ is an invertible $n \times n$ matrix with strong-substitutes rows.

Moreover, let H_i denote the i -th column of the matrix $H := \begin{pmatrix} 0 \\ G_2 \end{pmatrix}$. Then, for any $i \in [n]$, we have either (i) $G^{-1}H_i = \mathbf{e}^S$ for some $S \subseteq [n]$ with $i \in S$, or (ii) $G^{-1}H_i = -\mathbf{e}^S$ for some $S \subseteq [n]$ with $i \notin S$.

Proof. Suppose G is a lower triangular matrix with diagonal values $g_{ii} = 1$ and $g_{ij} \geq \rho$, $1 \geq \rho$ for all $i > j$. This is without loss of generality, as we can permute columns and multiply by a diagonal matrix U with entries ± 1 on its diagonals otherwise; in this case, replace G by PUG (where P denotes the permutation matrix), below and adjust the proof slightly. Note that the rows of PUG remain strong-substitutes.

In order to prove our lemma, we associate G with a directed graph D on vertices $[n]_0$ as follows. As G is invertible and the rows are strong-substitutes, each row takes the form $\mathbf{e}^j - \mathbf{e}^i$ for some distinct $i, j \in [n]_0$, where we define $\mathbf{e}^0 = \mathbf{0}$. Thus, for each row $\mathbf{e}^j - \mathbf{e}^i$ we add an arc from i to j .

We establish some properties of D . Firstly, D is an arborescence rooted at 0; that is, it is a directed graph in which there is exactly one path from 0 to any other vertex. This can be seen by induction on $k \in [n]$ if we define G_{kk} as the sub-matrix of G consisting of the first k rows and columns and let D_{kk} denote its graph. It is immediate that D_{11} is an arborescence. Secondly,

consider D_{kk} and assume the claim holds for $k - 1$. The last row (and column) of G_{kk} adds a single arc from a good $j < k$ to k , so D_{kk} remains an arborescence.

Secondly, we can express G as $G = I_n - L$, where I_n is the identity matrix and L is a strictly lower triangular matrix. Moreover, L admits the intuitive interpretation of being the transpose of the adjacency matrix of D . (To see this, note that $\mathbf{e}^j - \mathbf{e}^i$ is the j -th row in G , so $L_{ji} = 1$ if, and only if, there exists an arc from i to j .) In the following, let L^k be the k -th power of L (with $L^0 = I_n$). We recall the well-known result that L_{ji}^k denotes the number of directed walks of length k from vertex i to j . In particular, as D is a tree, all walks are paths. This implies that, for all $i, j \in [n]$, we have $L_{ij}^k \in \{0, 1\}$ and $L_{ij}^k = 1$ for at most one value of k . Moreover, $L^n = \mathbf{0}$, as no path can have length n . It follows that the inverse of G is

$$G^{-1} = \sum_{k=0}^{n-1} L^k, \quad (13)$$

and G^{-1} has only entries $\{0, 1\}$.

Fix $i \in [n]$ and let J_i be any subset of the in-neighbours of i . (The *in-neighbours* of vertex i in D are the vertices for which there exists an arc from i to j .) We claim that for $\mathbf{h} = \sum_{j \in J_i} \mathbf{e}^j$, there exists a subset S such that $i \in S$ and $G^{-1}\mathbf{h} = \mathbf{e}^S$, and for $\mathbf{h} = \mathbf{e}^i - \sum_{j \in J_i} \mathbf{e}^j$, there exists a subset S such that $i \in S$ and $G^{-1}\mathbf{h} = \mathbf{e}^S$. This claim immediately implies the main statement of Lemma 26, as H_i takes the form $H_i = \sum_{j \in J_i} \mathbf{e}^j$ or $H_i = \mathbf{e}^i - \sum_{j \in J_i} \mathbf{e}^j$ for some J_i .

It remains to prove our claim. Note first that $J_i = \{i-1, \dots, n\}$. For any $j \in [n]_0$, let $S(j)$ denote the vertices reachable from vertex j in D . As $L_{jl}^k = 1$ if there exists a path of length k from vertex l to j , and $L_{jl}^k = 0$ otherwise, (13) implies that $G^{-1}\mathbf{e}^i = \mathbf{e}^{S(i)}$ is the indicator vector denoting, for each $l \in [n]$, whether l is reachable from j (along a path of arbitrary length). Next, observe that J_i contains endpoints of arcs from i , so the arborescence property of D implies that the sets $S(j)$ for $j \in J_i$ are pairwise disjoint. Moreover, we see that $S(j) \subseteq S(i)$, as there is an arc from i to every $j \in J_i$. This allows us to say the following. If $\mathbf{h} = \sum_{j \in J_i} \mathbf{e}^j$, then we have $G^{-1}\mathbf{h} = \sum_{j \in J_i} \mathbf{e}^{S(j)} = \mathbf{e}^S$, where we define $S := \bigcup_{j \in J_i} S(j)$, and we note that $i \notin S$. On the other hand, if $\mathbf{h} = \mathbf{e}^i - \sum_{j \in J_i} \mathbf{e}^j$, then we have $G^{-1}\mathbf{h} = \mathbf{e}^{S(i)} - \sum_{j \in J_i} \mathbf{e}^{S(j)} = \mathbf{e}^S$, where we define $S := S(i) \cap \bigcup_{j \in J_i} S(j)$, and we note that $i \in S$. The last equality holds as $S(j) \subseteq S(i)$ for all $j \in J_i$. \square

Observation 27. *Suppose that G is an $n \times n$ matrix with strong-substitutes rows and that $S \subseteq [n]$. Then $G\mathbf{e}^S \in \{1, 0, 1\}^n$ and, in particular, $G\mathbf{e}^S = \mathbf{1}$.*

Lemma 28. *Suppose we have two strong-substitutes integer-valued valuations, u^1 and u^2 , and let $u^{\{1,2\}}$ denote the aggregate valuation. Fix $i \in [n]$ and $\varepsilon \in (0, 1)$, and write $u^{2[\varepsilon, i]}$ for the valuation given by $u^{2[\varepsilon, i]}(\mathbf{x}) = u^2(\mathbf{x}) + \varepsilon x_i$. For any $\mathbf{r}^1, \mathbf{r}^2 \in \mathbb{Z}^n$, suppose there exists a price \mathbf{p} such that $\mathbf{r}^1 \in D_{u^1}(\mathbf{p})$, and $\mathbf{r}^2 \in D_{u^{2[\varepsilon, i]}}(\mathbf{p})$. Then there exists a price \mathbf{p}^0 such that $\mathbf{r}^1 \in D_{u^1}(\mathbf{p}^0)$, and $\mathbf{r}^2 \in D_{u^2}(\mathbf{p}^0)$. Moreover, for any price \mathbf{p} such that $\mathbf{r} := \mathbf{r}^1 + \mathbf{r}^2 \in D_{u^{\{1,2\}}}(\mathbf{p})$, we have $\mathbf{r}^1 \in D_{u^1}(\mathbf{p})$, and $\mathbf{r}^2 \in D_{u^2}(\mathbf{p})$.*

Proof. Fix bundles \mathbf{r}^1 and \mathbf{r}^2 that are, by assumption, demanded at \mathbf{p} under u^1 and $u^{2[\varepsilon, i]}$. For $j \in \{1, 2\}$, we define $P^j := \{\mathbf{p} \mid \mathbf{r}^j \in D_{u^j}(\mathbf{p})\}$, the set of prices at which \mathbf{r}^j is demanded under u^j . It follows from the definition of $u^{2[\varepsilon, i]}$ that the set $P^{2[\varepsilon, i]}$ of prices at which \mathbf{r}^2 is demanded by \mathbf{p} under $u^{2[\varepsilon, i]}$ is equal to $P^2 + \varepsilon \mathbf{e}^i$. Moreover, by assumption we have $\mathbf{p} \in P^1 \setminus P^{2[\varepsilon, i]}$, so this intersection is non-empty.

As P^1 and P^2 are polyhedra in \mathbb{R}^n (cf. [Baldwin and Klemperer, 2019, Proposition 2.7]), we can describe each region P^j by the system of linear inequalities

$$H^j \mathbf{p} \leq \mathbf{A}^j, \quad (14)$$

where H^j is an $a_j \times n$ matrix for some $a_j \geq \mathbb{Z}_+$ and $A^j \geq \mathbb{R}^{a_j}$. Moreover, Proposition 3.10 in [Baldwin and Klemperer, 2019] tells us that we can choose H^j such that each row of H^j is a strong-substitutes vector. Additionally, as valuation w^j is integer-valued, we have $A^j \geq \mathbb{Z}^{a_j}$. It follows that $P^{2[\varepsilon, i]} = P^2 + \varepsilon e^i$ is defined by

$$H^2(\mathbf{p} - \varepsilon e^i) \leq A^2, \quad H^2 \mathbf{p} \leq A^2 + \varepsilon H_i^2, \quad (15)$$

where H_i^2 is the i -th column of H^2 . This allows us to express the intersection $P^1 \setminus P^{2[\varepsilon, i]}$ as

$$H \mathbf{p} \leq A + \varepsilon \hat{H}_i, \quad (16)$$

where we define $H = \begin{pmatrix} H^1 \\ H^2 \end{pmatrix}$, $A = \begin{pmatrix} A^1 \\ A^2 \end{pmatrix}$, and $\hat{H} = \begin{pmatrix} 0 \\ H^2 \end{pmatrix}$.

Now we argue similarly to Theorem 19.1 in [Schrijver, 1998]. We know that $P^1 \setminus P^{2[\varepsilon, i]} \neq \emptyset$. Let $F = \{ \mathbf{p} \mid H^0 \mathbf{p} = A^0 + \varepsilon \hat{H}_i^0 g \}$ be a minimal face of this intersection, where $H^0 \mathbf{p} \leq A^0 + \varepsilon \hat{H}_i^0$ is a subsystem of (16) with linearly independent rows. Without loss of generality, we can write $H^0 = (G \ G)$, where G is invertible (permuting the columns of H^0 if necessary, which is equivalent to relabelling goods). Moreover, we can choose this permutation such that the i -th coordinate remains within the first square matrix – that is, if G is $n^0 \times n^0$ then $i \leq n^0$. In addition, G inherits from H^1, H^2 the property that its rows are strong-substitutes vectors. Now define the two price vectors

$$\mathbf{p}^\varepsilon := \begin{pmatrix} G^{-1} (A^0 + \varepsilon \hat{H}_i^0) \\ 0 \end{pmatrix} \quad \text{and} \quad \mathbf{p}^0 := \begin{pmatrix} G^{-1} A^0 \\ 0 \end{pmatrix}.$$

Note that \mathbf{p}^0 is an integral vector, since A is integral and G is unimodular. Moreover, by Lemma 26, we know that either (i) $G^{-1} \hat{H}_i^0 = e^S$ for some $S \subseteq [n^0]$ with $i \in S$, or (ii) $G^{-1} \hat{H}_i^0 = e^S$ for some $S \subseteq [n^0]$ with $i \notin S$ (here n^0 is the dimension of G). In either case, $\mathbf{p}^\varepsilon - \varepsilon e^i \geq f \mathbf{p}^0 - \varepsilon e^{S^0} g$ for some $S^0 \subseteq [n]$. We also see immediately that \mathbf{p}^ε is a vector in F , and hence lies in $P^1 \setminus P^{2[\varepsilon, i]}$, so \mathbf{p}^ε satisfies (16). We now use this knowledge to show that $\mathbf{p}^0 \geq P^1 \setminus P^2$.

As $\mathbf{p}^\varepsilon \geq P^1$, it satisfies (14) for $j = 1$. Note also that the rows of H^1 are strong-substitutes and $(\mathbf{p}^0 - \mathbf{p}^\varepsilon) \geq f - \varepsilon e^{S^0} g$ for some $S^0 \subseteq [n]$, so (14) and Observation 27 imply

$$H^1 \mathbf{p}^0 = H^1 \mathbf{p}^\varepsilon + H^1 (\mathbf{p}^0 - \mathbf{p}^\varepsilon) \leq A^1 - \varepsilon \mathbf{1}.$$

The integrality of H^1, \mathbf{p}^0 and A^1 , together with $\varepsilon \in (0, 1)$, imply $H^1 \mathbf{p}^0 \leq A^1$. Similarly, we argue that $H^2 \mathbf{p}^0 \leq A^2$. Indeed, \mathbf{p}^ε satisfies (15) and we have $(\mathbf{p}^0 - \mathbf{p}^\varepsilon + \varepsilon e^i) \geq -\varepsilon e^S$ for some $S \subseteq [n]$, so (15) and Observation 27 imply

$$H^2 \mathbf{p}^0 = H^2 (\mathbf{p}^\varepsilon - \varepsilon e^i) + H^2 (\mathbf{p}^0 - \mathbf{p}^\varepsilon + \varepsilon e^i) \leq A^2 - \varepsilon \mathbf{1} \leq A^2.$$

The last inequality again follows due to the integrality of H^2, \mathbf{p}^0 and A^2 . Thus we have established that $\mathbf{p}^0 \geq P^1 \setminus P^2$, which is equivalent to $\mathbf{r}^j \geq D_{w^j}(\mathbf{p}^0)$ for $j = 1, 2$, our first claim.

Now let $\mathbf{r} := \mathbf{r}^1 + \mathbf{r}^2$ and let \mathbf{p} be a price at which \mathbf{r} is aggregately demanded, so $\mathbf{r} \geq D_{w^1, w^2}(\mathbf{p})$. Suppose $\mathbf{r} = \mathbf{s}^1 + \mathbf{s}^2$ is a potentially different equilibrium allocation of bundles across the two agents, so that $\mathbf{s}^j \geq D_{w^j}$ holds. By Proposition 2.F.1 in [Mas-Colell et al., 1995], we know that $(\mathbf{s}^j - \mathbf{r}^j) \cdot (\mathbf{p} - \mathbf{p}^0) = 0$ for $j = 1, 2$, with equality holding only if the agent is indifferent between both bundles at both prices. Hence, adding across $j = 1, 2$, we obtain $((\mathbf{s}^1 + \mathbf{s}^2) - \mathbf{r}) \cdot (\mathbf{p} - \mathbf{p}^0) = 0$, with equality holding only if both agents are indifferent between their two bundles at both prices. But as $\mathbf{s}^1 + \mathbf{s}^2 = \mathbf{r}$, equality holds, so we arrive at $\mathbf{r}^j \geq D_{w^j}(\mathbf{p})$ for $i = 1, 2$. \square

Proof of Proposition 14. By Lemma 25, every bid list that is η -valid at \mathbf{p} with $\eta < \frac{1}{2}$ can be associated with a strong-substitutes integral-valued valuation such that their indirect utility

function and demand correspondence coincide in $B(\mathbf{p}, \eta)$. The bid lists B^j and $B \cap B^j$ are $1/2$ -valid at \mathbf{p} by Lemma 9 and the fact that integral bids have a surplus of at least 1 at integral prices. Moreover, the shifted bid list of bidder j and the bid list B^θ are $(1/2 - j\varepsilon)$ -valid at \mathbf{p} (this follows from the definition of local validity). Recalling that $j\varepsilon < \frac{1}{4}$, we fix η with $2j\varepsilon < \eta < \frac{1}{2}$ and let u^2 denote a strong-substitutes integral-valued valuation of bidder j corresponding in $B(\mathbf{p}, \eta)$ to B^j ; and let u^1 denote a (strong-substitutes integral-valued) valuation corresponding in $B(\mathbf{p}, \eta)$ to the bids $B \cap B^j$ of the other bidders. As above, $u^{\ell_1, 2g}$ denotes the aggregate of these valuations, and it corresponds in $B(\mathbf{p}, \eta)$ to B . By assumption, we have $\mathbf{r} \geq D_{u^{\ell_1, 2g}}(\mathbf{p})$.

As in Lemma 28, we write $u^{2[\varepsilon, i]}$ for the valuation given by $u^{2[\varepsilon, i]}(\mathbf{x}) = u^2(\mathbf{x}) + \varepsilon x_i$. Note that the polyhedral complex of $u^{2[\varepsilon, i]}$ is the polyhedral complex of u^2 shifted by $\varepsilon < \eta/2$ in direction \mathbf{e}^i . Hence, in $B(\mathbf{p}; \frac{\eta}{2})$ the demand $D_{u^{2[\varepsilon, i]}}$ coincides with the demand of bidder j 's bids after shifting, and $D_{u^{1, 2[\varepsilon, i]}}$ coincides with D_{B^θ} in $B(\mathbf{p}; \frac{\eta}{2})$. Since both u^1 and $u^{2[\varepsilon, i]}$ are strong-substitutes valuations, there exists a price $\mathbf{p} \geq \mathbb{R}^n$ such that $\mathbf{r} \geq D_{u^{1, 2[\varepsilon, i]}}(\mathbf{p})$, and thus there exist $\mathbf{r}^1, \mathbf{r}^2 \geq \mathbb{Z}^n$ with $\mathbf{r}^1 + \mathbf{r}^2 = \mathbf{r}$ such that $\mathbf{r}^2 \geq D_{u^1}(\mathbf{p})$ and $\mathbf{r}^1 \geq D_{u^{2[\varepsilon, i]}}(\mathbf{p})$. By Lemma 28, it follows that $\mathbf{r}^j \geq D_{u^j}(\mathbf{p})$ for $j = 1, 2$.

Like in the proof of Lemma 28, we let $P^j := \{\mathbf{p}^\theta \mid 0 \leq \mathbf{r}^j \geq D_{u^j}(\mathbf{p}^\theta)g\}$ for $j \geq \ell_1, 2g$ and observe that $P^{2[\varepsilon, i]} := \{\mathbf{p}^\theta \mid 0 \leq \mathbf{r}^2 \geq D_{u^{2[\varepsilon, i]}}(\mathbf{p}^\theta)g\} = P^2 + \varepsilon \mathbf{e}^i$. We know that $\mathbf{p} \geq P^1 \cap P^2$ (so this intersection is non-empty) and $\mathbf{p} \geq P^1 \cap P^{2[\varepsilon, i]}$ (so this intersection is also non-empty). We seek $\mathbf{p}^\varepsilon \geq P^1 \cap P^{2[\varepsilon, i]}$ with $\mathbf{p}^\varepsilon \geq \mathbf{p} - \varepsilon \mathbf{e}^S$ for some $S \subseteq [n]$; this is sufficient to show $\mathbf{r} \geq D_{u^{1, 2[\varepsilon, i]}}(\mathbf{p}^\varepsilon)$. Since $j\varepsilon < \frac{\eta}{2}$, we know $\mathbf{p}^\varepsilon \geq B(\mathbf{p}, \frac{\eta}{2})$, the region in which demand from the valuations coincides with demand from the bids. Our first claim, that $\mathbf{r} \geq D_{B^\theta}(\mathbf{p}^\varepsilon)$ follows.

Analogous to line (14) above, we express the polyhedron P^j as $H^j \mathbf{p}^\theta \leq A^j$, for $j \geq \ell_1, 2g$, where the rows of H^j are strong-substitutes vectors. Therefore $P^1 \cap P^{2[\varepsilon, i]} = P^1 \cap (P^2 + \varepsilon \mathbf{e}^i)$ is defined by the set of inequalities

$$H^1 \mathbf{p}^\theta \leq A^1 \text{ and } H^2(\mathbf{p}^\theta - \varepsilon \mathbf{e}^i) \leq A^2. \quad (17)$$

Now, for $j = 1, 2$, let

$$(H^j)^\theta \mathbf{p}^\theta \leq (A^j)^\theta \quad (18)$$

be the subsystem of $H^j \mathbf{p}^\theta \leq A^j$ consisting of all rows which hold with equality at \mathbf{p} . (If both of these subsystems are empty then subsequent arguments will simply set $\mathbf{p}^\varepsilon = \mathbf{p}$). Write $H = \begin{pmatrix} (H^1)^\theta \\ (H^2)^\theta \end{pmatrix}$, $A = \begin{pmatrix} (A^1)^\theta \\ (A^2)^\theta \end{pmatrix}$ and $\hat{H} = \begin{pmatrix} 0 \\ (H^2)^\theta \end{pmatrix}$. Thus when we stack the corresponding rows of (17), we obtain the new system

$$H \mathbf{p}^\theta \leq A + \varepsilon \hat{H}_i. \quad (19)$$

This defines a superset of $P^1 \cap (P^2 + \varepsilon)$, since we have removed some rows from $H^j \mathbf{p}^\theta \leq A^j$ for $j = 1, 2$. But we have seen that $P^1 \cap (P^2 + \varepsilon)$ contains \mathbf{p} , so (19) defines a non-empty polytope. So let $F = \{\mathbf{p}^\theta \mid 0 \leq H \mathbf{p}^\theta \leq A + \varepsilon \hat{H}_i^g\}$ be a minimal face of this polytope, where

$$H^\theta \mathbf{p}^\theta \leq A^\theta + \varepsilon \hat{H}_i^\theta \quad (20)$$

is a subsystem of (19) with linearly independent rows.

Without loss of generality, we can write $H^\theta = (G \ G)$, where G is invertible (permuting the columns of H^θ if necessary, which is equivalent to relabelling goods). Moreover we can choose this permutation such that the i -th coordinate remains within the first square matrix – that is, if G is $n^\theta \times n^\theta$ then assume when we re-order that $i \leq n^\theta$. Correspondingly write $\mathbf{p}^\theta = \begin{pmatrix} \mathbf{p}_1^\theta \\ \mathbf{p}_2^\theta \end{pmatrix}$. In summary, our system (19) is now written

$$(G \ G) \begin{pmatrix} \mathbf{p}_1^\theta \\ \mathbf{p}_2^\theta \end{pmatrix} \leq A^\theta + \varepsilon \hat{H}_i^\theta. \quad (21)$$

By definition of the rows identified at (18), we have $(G \quad \mathcal{G}) \begin{pmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \end{pmatrix} = A^\theta$ and so

$$\mathbf{p}_1 = G^{-1}A^\theta - G^{-1}\mathcal{G}\mathbf{p}_2. \quad (22)$$

We now define \mathbf{p}^ε as follows

$$\mathbf{p}^\varepsilon := \begin{pmatrix} G^{-1}(A^\theta + \varepsilon\hat{H}_i^\theta) & G^{-1}\mathcal{G}\mathbf{p}_2 \\ & \mathbf{p}_2 \end{pmatrix} \quad (23)$$

in order to obtain

$$H^\theta \mathbf{p}^\varepsilon = (G \quad \mathcal{G}) \mathbf{p}^\varepsilon = A^\theta + \varepsilon\hat{H}_i^\theta. \quad (24)$$

We wish to show that (17) holds for $\mathbf{p}^\theta = \mathbf{p}^\varepsilon$. From (22) we observe that $\mathbf{p}^\varepsilon - \mathbf{p} = \begin{pmatrix} \varepsilon G^{-1}\hat{H}_i^\theta \\ 0 \end{pmatrix}$. So, by Lemma 26, either $\mathbf{p}^\varepsilon = \mathbf{p} + \varepsilon\mathbf{e}^S$ for some $S \subseteq [n]$ with $i \in S$, or $\mathbf{p}^\varepsilon = \mathbf{p} - \varepsilon\mathbf{e}^S$ where $S \subseteq [n]$ with $i \notin S$. In either case, $\mathbf{p}^\varepsilon - \varepsilon\mathbf{e}^i \geq \mathbf{p} - \varepsilon\mathbf{e}^{S^c}g$ for some $S^c \subseteq [n]$.

We also immediately observe that \mathbf{p}^ε satisfies (20) with equality, and so $\mathbf{p}^\varepsilon \geq F$. Thus, by definition of F , \mathbf{p}^ε satisfies (19). It follows by definition of (18) that, for every row in the original systems $H^j \mathbf{p}^\theta = A^j$ that holds with equality at \mathbf{p} , the corresponding shifted equation in (17) holds at \mathbf{p}^ε . It remains to show that this is also true for rows that are slack at \mathbf{p} . Let $(H^j)^\theta \mathbf{p}^\theta = (A^j)^\theta$ be the subsystems of such rows, for $j = 1, 2$, and observe that by integrality it follows that

$$(H^j)^\theta \mathbf{p} = (A^j)^\theta + \mathbf{1}. \quad (25)$$

Now, applying Observation 27 to $\mathbf{p}^\varepsilon - \mathbf{p}$ and (25) for $j = 1$, we obtain

$$(H^1)^\theta \mathbf{p}^\varepsilon = (H^1)^\theta \mathbf{p} + (H^1)^\theta (\mathbf{p}^\varepsilon - \mathbf{p}) = (A^1)^\theta + (1 - \varepsilon)\mathbf{1} = (A^1)^\theta.$$

And, applying Observation 27 to $\mathbf{p}^\varepsilon - \mathbf{p} - \varepsilon\mathbf{e}^i$, and (25) for $j = 2$,

$$(H^2)^\theta (\mathbf{p}^\varepsilon - \varepsilon\mathbf{e}^i) = (H^2)^\theta \mathbf{p} + (H^2)^\theta (\mathbf{p}^\varepsilon - \mathbf{p} - \varepsilon\mathbf{e}^i) = (A^2)^\theta + (1 - \varepsilon)\mathbf{1} = (A^2)^\theta.$$

So, for every row in the original systems $H^j \mathbf{p}^\theta = A^j$ that is slack at \mathbf{p} , the corresponding shifted equation in (17) holds for \mathbf{p}^ε . This completes the proof that $\mathbf{p}^\varepsilon \geq P^1 \setminus (P^2 + \varepsilon\mathbf{e}^i g)$. It follows that $\mathbf{r} \geq D_{u,1,2,[1:n]}(\mathbf{p}^\varepsilon)$, which concludes the proof of our first claim.

To see that we can compute \mathbf{p}^ε in polynomial time, recall that the prices at which \mathbf{r} is demanded are equivalent to the prices that minimise the Lyapunov function $g(\mathbf{p}) = f_{B^\theta}(\mathbf{p}) + \mathbf{r} \cdot \mathbf{p}$. It is known (see [Shioura and Tamura, 2015, Theorem 7.3]) that f_{B^θ} (and thus g) satisfies the *generalised submodularity* property

$$f_{B^\theta}(\mathbf{p}) + f_{B^\theta}(\mathbf{q}) = f_{B^\theta}(\mathbf{p} _ \mathbf{q}) + f_{B^\theta}(\mathbf{p} \wedge \mathbf{q}) \quad (\mathcal{S}\mathbf{p}, \mathbf{q} \geq \mathbb{R}^n),$$

where $_$ and \wedge denote the component-wise maximum and minimum, respectively. Hence

$$\begin{aligned} h^+(S) &:= g(\mathbf{p} + \varepsilon\mathbf{e}^S) - g(\mathbf{p}), \\ h^-(S) &:= g(\mathbf{p} - \varepsilon\mathbf{e}^S) - g(\mathbf{p}), \end{aligned}$$

are submodular set functions. Hence, in order to determine a price vector $\mathbf{p}^\varepsilon \geq \{\mathbf{p} + \varepsilon\mathbf{e}^S \mid S \subseteq [n]\}$ at which \mathbf{r} is demanded, we find minimisers S^+ and S^- of h^+ and h^- using SFM, then let

$$\mathbf{p}^\varepsilon = \arg \min \left\{ g(\mathbf{p} + \varepsilon\mathbf{e}^{S^+}), g(\mathbf{p} - \varepsilon\mathbf{e}^{S^-}) \right\}.$$

Note that we do not require *minimal* submodular minimisers to find \mathbf{p}^ε , and so this step takes $2T(n)$ time, where $T(n)$ is the time it takes to perform submodular minimisation on h^- . We also note that finding \mathbf{p}^ε is analogous to finding a steepest descent direction in Appendix C. \square

B.5 Proof of Lemma 17

Here we prove that applying ShiftProjectUnshift strictly reduces the number of edges in the marginal bids graph G_A .

Proof of Lemma 17. First we show that every edge in G_{A^θ} is present in G_A . To see this, fix some bid \mathbf{b} and note by Observation 7 and Lemma 16 that Steps 1, 2 and 3 do not make \mathbf{b} marginal on any new goods.

Secondly we prove that for any multi-bidder cycle C with cycle-link good i and incident label j , ShiftProjectUnshift removes at least one of the edges of C from the marginal bids graph. Re-label the goods going around cycle C as $1, \dots, k$, so that $1 = i$ and so that bidder j placed the bid marginal between goods 1 and 2. Also, for convenience, index the marginal bid that is marginal between goods i and $i + 1$ as ‘ i ’, and index the marginal bid between goods k and 1 as ‘ k ’. Thus $\mathbf{b}^1 \succeq B^j$ and $\mathbf{b}^k \not\succeq B^j$. (In general the differently labelled bids need not be from different bidders). The existence of a marginal bid between goods i and $i + 1$ means that we have an equality

$$b_i^i - p_i = b_{i+1}^i - p_{i+1} \quad (26)$$

for $i = 1, \dots, k - 1$, and also

$$b_k^k - p_k = b_1^k - p_1 \quad (27)$$

But if we take the sum of the first $k - 1$ equations, and cancel, we find

$$\sum_{i=1}^{k-1} b_i^i - \sum_{i=1}^{k-1} b_{i+1}^i = \sum_{i=1}^{k-1} (p_i - p_{i+1}) = p_1 - p_k$$

So it must hold that

$$\sum_{i=1}^{k-1} (b_i^i - b_{i+1}^i) = b_1^k - b_k^k. \quad (28)$$

For any bid \mathbf{b} in the bid lists $(B^j)_{j \in J}$, fixed $j \in J$ and $i \in [n]$, let

$$\mathbf{b} := \begin{cases} \mathbf{b} + \frac{1}{10} \mathbf{e}^i & \mathbf{b} \succeq B^j \\ \mathbf{b} & \text{otherwise} \end{cases}$$

Suppose we replace each bid \mathbf{b} by \mathbf{b} as above and recompute the prices. Since prices shift by at most $\frac{1}{10}$, a bid that is *not* marginal on a pair of goods cannot *become* marginal on them. If we assume for a contradiction that all edges in C are all still present in G^θ then we can write down similar expressions to (26) and (27) (w.r.t. new prices), then eliminate those prices obtaining a version of (28), namely $\sum_{i=1}^{k-1} (\bar{b}_i^i - \bar{b}_{i+1}^i) = \bar{b}_1^k - \bar{b}_k^k$. But by definition of \mathbf{b} , we know that

$$\sum_{i=1}^{k-1} (\bar{b}_i^i - \bar{b}_{i+1}^i) = \sum_{i=1}^{k-1} (b_i^i - b_{i+1}^i) + \frac{1}{10},$$

and $\bar{b}_1^k - \bar{b}_k^k = b_1^k - b_k^k$.

This is inconsistent with (28), so we have the required contradiction. \square

C Finding market-clearing prices

We discuss two iterative steepest descent algorithms, MinUp and LongStepMinUp, from [Shioura, 2017] that determine a minimal minimiser of an L^1 -convex function. Both algorithms use the SFM subroutine described in Section 2 to find the component-wise minimal discrete

steepest descent direction. For the second algorithm, we present two methods of computing step lengths and show that both methods yield a polynomial running time in our bidding-language setting.

In order to apply the steepest descent method to our price-finding problem, we define a Lyapunov function g and note in Proposition 29 that its restriction to Z_+^n is L^{\natural} -convex. Moreover, Lemma 30 states that the lowest market-clearing price is integral and finding it reduces to determining the minimal minimiser of g . This approach generalises an algorithm used by Ausubel's ascending auction design [Ausubel, 2006] and Gul and Stacchetti [2000] for the task of finding equilibrium prices in single-unit markets.

Let B be a valid bid list. By Theorem 1, the function f_B defined in (3) is the indirect utility function f_u of some strong-substitutes valuation u . For any strong-substitutes valuation u , the *Lyapunov function* with regard to indirect utility function f_u and target bundle \mathbf{t} is defined as $g_{\mathbf{t}}(\mathbf{p}) := f_u(\mathbf{p}) + \mathbf{t} \cdot \mathbf{p}$. We suppress the subscript \mathbf{t} if it is clear from context. We can use our knowledge of the bids in B and (3) to express $g_{\mathbf{t}}$ for the list B as

$$g_{\mathbf{t}}(\mathbf{p}) := f_B(\mathbf{p}) + \mathbf{t} \cdot \mathbf{p} = \sum_{\mathbf{b} \in B} w(\mathbf{b}) \max_{i \in [n]_0} (b_i - p_i) + \mathbf{t} \cdot \mathbf{p}. \quad (29)$$

From (29) it is clear that we can evaluate g at any price \mathbf{p} in time $O(n|B|)$.

A function $f : Z^n \rightarrow \mathbb{R}$ is L^{\natural} -convex if it satisfies the *translation submodularity* property,

$$f(\mathbf{p}) + f(\mathbf{q}) \leq f((\mathbf{p} \vee \alpha \mathbf{1}) \wedge \mathbf{q}) + f(\mathbf{p} \wedge (\mathbf{q} + \alpha \mathbf{1})) \quad (\delta \mathbf{p}, \mathbf{q} \geq Z_+^n, \delta \alpha \geq Z_+).$$
 (30)

Here, \vee and \wedge denote the component-wise maximum and minimum, respectively.

The following proposition and lemma demonstrate that we can use algorithms for minimising L^{\natural} -convex functions to find the minimal price \mathbf{p} at which \mathbf{t} is demanded.

Proposition 29 ([Murota et al., 2013, 2016]). *The Lyapunov function $g_{\mathbf{t}}$ restricted to Z_+^n is L^{\natural} -convex.*

Proof. It is known that an indirect utility function f_u restricted to Z^n is L^{\natural} -convex if and only if the valuation function u is strong-substitutes (cf. [Shioura and Tamura, 2015, Theorem 7.1]). Secondly, it is easy to verify that adding a linear term to an L^{\natural} -convex function preserves L^{\natural} -convexity. As the bid list B is valid, we have $f_B = f_u$ for some strong-substitutes valuation u and it follows that the function $f_B(\mathbf{p}) + \mathbf{t} \cdot \mathbf{p}$ is L^{\natural} -convex. \square

Lemma 30. *The Lyapunov function $g_{\mathbf{t}}$ with regard to a valid (integral) bid list and any target bundle \mathbf{t} is convex, and \mathbf{p} is a minimiser of g if and only if it is a market-clearing price for target bundle \mathbf{t} . Moreover, the minimal minimiser of g is integral.*

Proof. The first statement is immediate from the fact that f_B is convex and $\mathbf{t} \cdot \mathbf{p}$ is a linear term. To see the second statement, note that for any market-clearing price \mathbf{p} of target bundle \mathbf{t} , we have $g(\mathbf{p}) = u(\mathbf{t})$, whereas for any price \mathbf{p} at which \mathbf{t} is not demanded, we have $g(\mathbf{p}) = \max_{\mathbf{x} \geq D(\mathbf{p})} (u(\mathbf{p}) - \mathbf{x} \cdot \mathbf{p}) + \mathbf{t} \cdot \mathbf{p} > u(\mathbf{t})$. Here u is the valuation function as defined in Section 2 and we use (2). Integrality of the minimal minimiser of g was shown in [Murota et al., 2016, Corollary 4.4]. Alternatively, we can see that this follows from the fact that if \mathbf{t} is demanded at \mathbf{p} , then \mathbf{t} is also demanded at $b\mathbf{p}c$. To see this, fix a bid \mathbf{b} and note that if \mathbf{b} demands good i at \mathbf{p} , it still demands i at $b\mathbf{p}c$. Indeed, as \mathbf{b} demands i at \mathbf{p} , we have $b_j - p_j \leq b_i - p_i$ for all goods j , which implies

$$b_j - b p_j c < b_j - p_j + 1 \leq b_i - p_i + 1 \leq b_i - b p_i c + 1.$$

Due to the integrality of the bids and prices in $b\mathbf{p}c$, this implies $b_j - b p_j c \leq b_i - b p_i c$. \square

Let \mathbf{p} be a point that is dominated by some minimiser of g . Shioura [2017] noted that \mathbf{p} minimises g if and only if $g(\mathbf{p}) = g(\mathbf{p} + \mathbf{e}^S)$ for every $S \subseteq [n]$. For any integral point $\mathbf{p} \in \mathbb{Z}_+^n$ and $S \subseteq [n]$, let the *slope function* $g^\theta(\mathbf{p}; S) := g(\mathbf{p} + \mathbf{e}^S) - g(\mathbf{p})$ denote the amount by which \mathbf{p} decreases when moving in the direction of \mathbf{e}^S . If S minimises $g^\theta(\mathbf{p}; S)$, we call \mathbf{e}^S a *steepest descent direction*. For any integral vector \mathbf{p} , the L^1 -convexity of g implies that $g^\theta(\mathbf{p}; S)$ is an integral submodular function [Shioura and Tamura, 2015, Theorem 7.2] and hence there exists a unique component-wise *minimal steepest descent direction* \mathbf{e}^{S_0} .

Let \mathbf{p} be the minimal minimiser of g . If \mathbf{p} is dominated by \mathbf{p} and we move in the minimal steepest descent direction, the point $\mathbf{p} + \mathbf{e}^{S_0}$ is also dominated by \mathbf{p} . This suggests Algorithm 8 (called GreedyUpMinimal in [Shioura, 2017]), which iterates a point \mathbf{p} by moving by some step \mathbf{e}^{S_0} , all the while remaining dominated by \mathbf{p} , until $\mathbf{p} = \mathbf{p}$.

Algorithm 8 MinUp

- 1: Pick a point $\mathbf{p} \succeq \mathbf{p}$ (e.g. $\mathbf{p} = \mathbf{0}$).
 - 2: Find the inclusion-wise minimal set $S_0 \subseteq [n]$ minimising $g^\theta(\mathbf{p}; S)$.
 - 3: **if** $S_0 = \emptyset$ **then**
 - 4: **return** \mathbf{p} .
 - 5: **else**
 - 6: Set $\mathbf{p} = \mathbf{p} + \mathbf{e}^{S_0}$ and go to line 2.
-

By the existence of the auctioneer's reserve bids, we can initialise \mathbf{p} to $\mathbf{0}$. For this starting point, running time analysis by Murota and Shioura [2014] implies that Algorithm 8 iterates exactly $k\mathbf{p} \cdot k_1$ times. As we know that \mathbf{p} is bounded from above by the component-wise maximum over all bids B , the number of iterations is at most $M := \max_{\mathbf{b} \in B} k\mathbf{b} \cdot k_1$. In each iteration of Step 2, we can determine S_0 using the SFM subroutine described in Section 2 that finds a minimal minimiser of a given function in time $T(n)$. This leads to the following running time for MinUp.

Theorem 31 (cf. [Murota and Shioura, 2014]). *The algorithm MinUp finds the component-wise minimal market-clearing price in time $O(MnT(n))$, where $T(n)$ denotes the time it takes to find a minimiser of g^θ .*

We note that the running times of the two practical SFM algorithms mentioned in Section 2 are given with respect to an upper bound on the absolute value of the objective function. In order to provide such an upper bound for our slope function g^θ , observe that the two points \mathbf{p} and $\mathbf{p} + \mathbf{e}^S$ share a demanded bundle \mathbf{x} , for any \mathbf{p} and $S \subseteq [n]$. As every bid contributes at most one item to \mathbf{t} and \mathbf{x} , this implies $g^\theta(\mathbf{p}; S) = (\mathbf{t} - \mathbf{x}) \cdot \mathbf{e}^S = \sum_j B_j$.

C.1 Longer step sizes

The analysis by Murota and Shioura [2014] shows that MinUp performs optimally for an iterative algorithm that is constrained to steps with an L^1 -size of at most 1. As described by Shioura [2017], we can, however, exploit monotonicity properties of the function $g^\theta(\mathbf{p}; S_0)$ in order to increase step sizes without changing the trajectory of \mathbf{p} as the algorithm runs. This reduces the number of SFM subroutine calls, the most expensive part of the algorithm. In particular, if \mathbf{e}^{S_0} denotes the minimal steepest descent direction at \mathbf{p} , we take a single long step $\lambda \mathbf{e}^{S_0}$ for some $\lambda \in \mathbb{Z}_+$ that is equivalent to several consecutive steps of MinUp in the same direction \mathbf{e}^{S_0} .

We follow [Shioura, 2017] in choosing step length

$$\lambda(\mathbf{p}, S_0) = \max_{\lambda \in \mathbb{Z}_+} \lambda g^\theta(\mathbf{p}; S_0) = g^\theta(\mathbf{p} + (\lambda - 1)\mathbf{e}^{S_0}; S_0), \quad (31)$$

that is, the farthest distance we can move before the slope

$$g^\theta(\mathbf{p} + (\lambda - 1)\mathbf{e}^{S_0}; S_0) = g(\mathbf{p} + \lambda \mathbf{e}^{S_0}) - g(\mathbf{p} + (\lambda - 1)\mathbf{e}^{S_0})$$

in the direction of \mathbf{e}^{S_0} changes. We give two methods to compute (31) in Section C.1.1. This leads to the following algorithm (referred to as GreedyUp-LS in [Shioura, 2017]).

Algorithm 9 LongStepMinUp

- 1: Pick a point \mathbf{p} \mathbf{p} (e.g. $\mathbf{p} = \mathbf{0}$).
 - 2: Compute inclusion-wise minimal minimiser $S_0 \subseteq [n]$ of $g^\theta(\mathbf{p}; S_0)$ using SFM.
 - 3: Determine $\lambda(\mathbf{p}, S_0)$, as defined by (31), using a method from Section C.1.1.
 - 4: **if** $S_0 \neq \emptyset$; **then**
 - 5: Set $\mathbf{p} = \mathbf{p} + \lambda(\mathbf{p}, S_0)\mathbf{e}^{S_0}$ and go to line 2.
 - 6: **return** \mathbf{p} .
-

Note that the values of \mathbf{p} follow the same trajectory for MinUp and LongStepMinUp. This is an immediate consequence of Lemma 33, which rests on the monotonicity properties of $g^\theta(\mathbf{p}, S_0)$ stated in Proposition 32.

Proposition 32 ([Shioura, 2017], Theorem 4.16). *Let S_0 and S_0^θ be minimal steepest descent directions at \mathbf{p} and $\mathbf{p} + \mathbf{e}^{S_0}$, respectively. Then we have*

1. $g^\theta(\mathbf{p} + \mathbf{e}^{S_0}; S_0^\theta) > g^\theta(\mathbf{p}; S_0)$ or
2. $g^\theta(\mathbf{p} + \mathbf{e}^{S_0}; S_0^\theta) = g^\theta(\mathbf{p}; S_0)$ and $S_0 \subseteq S_0^\theta$.

Lemma 33. *Let S_0 denote the minimal steepest descent at \mathbf{p} and let $\lambda(\mathbf{p}, S_0)$ be defined by (31). Then \mathbf{e}^{S_0} is the minimal steepest descent at $\mathbf{p} + (\lambda - 1)\mathbf{e}^{S_0}$ for any $1 \leq \lambda \leq \lambda(\mathbf{p}, S_0)$.*

Shioura [2017] bounds the number of iterations of LongStepMinUp as follows.

Theorem 34 ([Shioura, 2017], Theorem 4.17). *The number of iterations of LongStepMinUp is at most $n \max_{f \in \mathcal{F}} g^\theta(\mathbf{0}; S) / j S \subseteq [n]g$.*

Note that $\mathbf{0}$ and \mathbf{e}^{S_0} share a demanded bundle \mathbf{x} , so for any $S \subseteq [n]$ we have

$$g^\theta(\mathbf{0}; S) = g(\mathbf{e}^S) \quad g(\mathbf{0}) = f_u(\mathbf{e}^S) + \mathbf{t} \cdot \mathbf{e}^S \quad f_u(\mathbf{0}) = (\mathbf{t} \cdot \mathbf{x}) \cdot \mathbf{e}^S \quad \sum_{i \in 2[n]} x_i \leq jBj,$$

as each bid contributes at most one unit to the demanded bundle. This implies that the LongStepMinUp algorithm takes at most n/jBj iterations to find component-wise minimal equilibrium prices.

C.1.1 Computing the step length

Fix a price \mathbf{p} and let \mathbf{e}^{S_0} be the component-wise minimal steepest descent direction at \mathbf{p} . We describe two methods to compute the step length defined by (31). The first method uses binary search and is also suggested in [Shioura, 2017], while second method exploits our knowledge of the bids to determine $\lambda(\mathbf{p}; S_0)$. Note that we can evaluate $g^\theta(\mathbf{p}; S_0)$ in time $O(njBj)$, as $g^\theta(\mathbf{p}; S_0) = g(\mathbf{p} + \mathbf{e}^{S_0}) - g(\mathbf{p})$.

Theorem 35. *The LongStepMinUp algorithm in combination with the binary search and demand change methods has a respective running time of $O(n^2jBj^2 \log M + njBjT(n))$ and $O(n^2jBj^3 + njBjT(n))$.*

A description of the two methods, as well as a proof of this theorem, is provided below. Note that as the two methods have different running time guarantees, the best method in practice is context-specific.

Binary search. Note that $\lambda(\mathbf{p}; S_0)$ can be bounded by the total number of unit steps in MinUp, which in turn is bounded by M . By Proposition 32, we have that $g^\theta(\mathbf{p}; S_0) < g^\theta(\mathbf{p} + \lambda \mathbf{e}^{S_0}; S_0)$ implies $g^\theta(\mathbf{p}; S_0) < g^\theta(\mathbf{p} + \lambda^\theta \mathbf{e}^{S_0}; S_0)$ for all $\lambda^\theta > \lambda$. Hence we can apply binary search to find $\lambda(\mathbf{p}; S_0)$ in time $O(njBj \log M)$.

Demand change. Alternatively, we can exploit our knowledge of the individual bids to determine $\lambda(\mathbf{p}; S_0)$. We proceed by performing a demand-change procedure, which repeatedly determines the highest value μ for which every bid $\mathbf{b} \succeq B$ demands the same goods (or a superset thereof) at prices $\mathbf{p} + \mu \mathbf{e}^{S_0}$ that it demands at $\mathbf{p} + \mathbf{e}^{S_0}$ and updates \mathbf{p} to $\mathbf{p} + \mu \mathbf{e}^{S_0}$.

Fix a bid \mathbf{b} and let I denote the goods it demands at \mathbf{p} . If $I \subseteq S_0$, then \mathbf{b} demands the same set of goods $I_{\mathbf{b}} = I \cap S_0$ at all prices $\mathbf{p} + \mu^\theta \mathbf{e}^{S_0}$ with $\mu^\theta \leq 1$. On the other hand, suppose $I \not\subseteq S_0$. We define $\mu_{\mathbf{b}} := \min_{j \in [n] \cap S_0} ((b_i - p_i) - (b_j - p_j))$, where i is any good in I , and consider the set of goods that \mathbf{b} demands at prices $\mathbf{p} + \mu^\theta \mathbf{e}^{S_0}$ with $\mu^\theta \leq 1$. If $1 - \mu^\theta < \mu_{\mathbf{b}}$, then \mathbf{b} demands I , if $\mu^\theta = \mu_{\mathbf{b}}$ then \mathbf{b} demands a superset of I and if $\mu^\theta > \mu_{\mathbf{b}}$, then \mathbf{b} demands none of the goods in I . We define C to be the list of bids that demand a subset of S_0 at \mathbf{p} , and let $\mu(\mathbf{p}, S_0) := \min_{\mathbf{b} \succeq C} \mu_{\mathbf{b}}$. Note that we can determine the value of $\mu(\mathbf{p}, S_0)$ in time $O(njBj)$.

The demand-change procedure takes as input a price $\mathbf{p}^0 := \mathbf{p}$ and direction \mathbf{e}^{S_0} , and consists of the following steps. Initially, we set λ to 0. Compute $\mu(\mathbf{p}, S_0)$, and increment λ by $\mu(\mathbf{p}, S_0)$. If we have $\lambda = \lambda(\mathbf{p}^0, S_0)$, return λ . Otherwise, increment \mathbf{p} by $\mu \mathbf{e}^{S_0}$ and repeat the above with the same value for \mathbf{e}^{S_0} .

Note that in order to check whether $\lambda = \lambda(\mathbf{p}^0; S_0)$, we can compute $g^\theta(\mathbf{p}^0 + \mu \mathbf{e}^{S_0}; S_0)$ and verify that $g^\theta(\mathbf{p}^0 + \mu \mathbf{e}^{S_0}; S_0) = g^\theta(\mathbf{p}^0; S_0)$, which takes time $O(njBj)$. Lemma 36 proves that the demand-change procedure correctly computes the value of $\lambda(\mathbf{p}; S_0)$ in time $O(njBj^2)$.

Lemma 36. *The demand-change procedure returns $\lambda(\mathbf{p}; S_0)$ in at most jBj iterations.*

Proof. Let K denote the number of iterations in the demand-change procedure and let \mathbf{p}^k, μ^k and C^k be the values of $\mathbf{p}, \mu(\mathbf{p}, S_0)$ and C after the k -th iteration. For notational convenience, let $\mathbf{p}^0 = \mathbf{p}$ and $\mu^0 = 0$. Proposition 32 implies that it suffices to show

$$g^\theta(\mathbf{p}; S_0) = g^\theta(\mathbf{p}^K - \mathbf{e}^S; S_0) \text{ and } g^\theta(\mathbf{p}; S_0) < g^\theta(\mathbf{p}^K; S_0) \quad (32)$$

in order to prove the first claim that $\lambda = \lambda(\mathbf{p}, S_0)$.

Firstly, note that there is a bundle \mathbf{x} that is demanded at \mathbf{p}^{k-1} and \mathbf{p}^k , as well as all prices in between these two points. Indeed, if a bid \mathbf{b} demands good i at $\mathbf{p}^{k-1} + \mathbf{e}^{S_0}$, then \mathbf{b} also demands i at the two prices \mathbf{p}^{k-1} and $\mathbf{p}^k - \mathbf{e}^{S_0} = \mathbf{p}^{k-1} + (\mu^{k-1} - 1) \mathbf{e}^{S_0}$, by construction of μ^k . Hence, making use of (2), we get

$$g^\theta(\mathbf{p}^{k-1}; S_0) = (\mathbf{t} - \mathbf{x}) \cdot \mathbf{e}^{S_0} = g^\theta(\mathbf{p}^k - \mathbf{e}^{S_0}; S_0).$$

Secondly, note that for every $0 \leq k < K$, we have $g^\theta(\mathbf{p}^k - \mathbf{e}^{S_0}; S_0) = g^\theta(\mathbf{p}^k; S_0)$, and for the last iteration K we have $g^\theta(\mathbf{p}^K - \mathbf{e}^{S_0}; S_0) < g^\theta(\mathbf{p}^K; S_0)$. This implies (32).

Now we turn to the second claim, that $K \leq jBj$. This follows from the fact that $C^{k-1} \subseteq C^k$ and $|C^0| \leq jBj$. Indeed, if $\mathbf{b} \notin C^{k-1}$, then \mathbf{b} demands goods $I \subseteq S_0$ at \mathbf{p}^{k-1} and $I_{\mathbf{b}} = I \cap S_0$ at any prices $\mathbf{p}^{k-1} + \mu \mathbf{e}^{S_0}$ for any $\mu \leq 1$, so $\mathbf{b} \in C^k$. Now suppose $\mathbf{b} \notin C^k$ is a bid for which $\mu^{k-1} = \mu_{\mathbf{b}}$ (at \mathbf{p}^{k-1}). Then we claim that $\mathbf{b} \notin C^k$. Indeed, the demanded goods I^θ of \mathbf{b} at $\mathbf{p} + (\mu_{\mathbf{b}} + 1) \mathbf{e}^{S_0}$ satisfy $I^\theta \subseteq S_0$ by construction of $\mu_{\mathbf{b}}$. By the same argument as above, \mathbf{b} demands goods not in S_0 at all prices $\mathbf{p} + \mu \mathbf{e}^{S_0}$ with $\mu \leq \mu_{\mathbf{b}} + 1$. \square

C.2 Some practical improvements

The computation time of MinUp and LongStepMinUp is dominated by the task of finding a minimal set S_0 minimising $g^\theta(\mathbf{p}; S)$ using SFM. In practice, we can exploit our direct access to

the bids to speed up the computation of S_0 by decreasing the dimensionality of the submodular function to be minimised. The following observations can be seen as a special case of observations by Ausubel [2006]. Fix $\mathbf{p} \geq \mathbf{Z}_+^n$ and let \mathbf{e}^{S_0} be the minimal steepest descent direction at \mathbf{p} . Note that \mathbf{p} and $\mathbf{p} + \mathbf{e}^{S_0}$ share a demanded bundle \mathbf{x} due to the structure of our price space. Hence

$$g^\theta(\mathbf{p}; S_0) = g(\mathbf{p} + \mathbf{e}^{S_0}) \quad g(\mathbf{p}) = (\mathbf{t} \quad \mathbf{x}) \cdot \mathbf{e}^{S_0},$$

and S_0 minimises this term if and only if S_0 contains all indices $i \in [n]$ with $t_i < x_i$ and no indices $j \in [n]$ with $t_j > x_j$. In particular, we have $S_0 := \{i \in [n] \mid t_i < x_i\}$ due to the minimal minimiser property of S_0 .

If \mathbf{p} has a unique demanded bundle \mathbf{x} , which is easy to verify, it is straightforward to compute \mathbf{x} . Hence, in this case we can determine the minimal steepest descent direction \mathbf{e}^{S_0} without performing SFM. In the case that \mathbf{p} has at least two demanded bundles, the demanded bundle \mathbf{x} shared by \mathbf{p} and $\mathbf{p} + \mathbf{e}^{S_0}$ is unknown. However, if we define index sets $I_{\mathbf{p}}$ and $J_{\mathbf{p}}$ as

$$\begin{aligned} I_{\mathbf{p}} &:= \{i \in [n] \mid x_i > t_i, \exists \mathbf{x} \in D(\mathbf{p})\}, \\ J_{\mathbf{p}} &:= \{i \in [n] \mid x_i < t_i, \exists \mathbf{x} \in D(\mathbf{p})\}, \end{aligned}$$

we have $I_{\mathbf{p}} \cap S_0 = \emptyset$ and $J_{\mathbf{p}} \cap S_0 = \emptyset$. (Note that the inequality in the definition of $I_{\mathbf{p}}$ and $J_{\mathbf{p}}$ is strict and not strict, respectively.) Hence we can restrict ourselves to minimising the submodular function $h: [n] \setminus (I_{\mathbf{p}} \cup J_{\mathbf{p}}) \rightarrow \mathbb{Z}$ defined by $h(T) = g^\theta(\mathbf{p}; T \cup I_{\mathbf{p}})$ and reduce the dimensionality of the SFM problem from n to $n - |I_{\mathbf{p}} \cup J_{\mathbf{p}}|$. The following lemma shows that computing $I_{\mathbf{p}}$ and $J_{\mathbf{p}}$ is cheap.

Lemma 37. *$I_{\mathbf{p}}$ and $J_{\mathbf{p}}$ can be computed in time $O(n|B|)$.*

Proof. Note that $I_{\mathbf{p}} = \{i \in [n] \mid \min_{\mathbf{x} \in D(\mathbf{p})} x_i > t_i\}$ and $J_{\mathbf{p}} = \{i \in [n] \mid \max_{\mathbf{x} \in D(\mathbf{p})} x_i < t_i\}$, where we compute the minimum and maximum component-wise. Fix $i \in [n]$. The minimum $\min_{\mathbf{x} \in D(\mathbf{p})} x_i$ is attained if no marginal bid selects good i . Hence

$$\min_{\mathbf{x} \in D(\mathbf{p})} x_i = \sum_{b \in B} w(b) \cdot \mathbb{1}_{b \text{ is non-marginal and demands } i}.$$

Similarly, we claim that $\max_{\mathbf{x} \in D(\mathbf{p})} x_i$ is attained if good i is selected whenever possible and thus $\max_{\mathbf{x} \in D(\mathbf{p})} x_i$ is the sum of the weights of the bids for which i is demanded. Consider the price perturbation \mathbf{p}^θ defined by $p_i^\theta = p_i$ and $p_j^\theta = p_j + \varepsilon$ for $j \neq i$. Here $\varepsilon > 0$ is chosen sufficiently small so that every bid that is marginal on i at \mathbf{p} is non-marginal and demands i at \mathbf{p}^θ , while every non-marginal bid demanding good i at \mathbf{p} is also non-marginal and demands i at \mathbf{p}^θ . This perturbation corresponds to our proposed rule to compute $\max_{\mathbf{x} \in D(\mathbf{p})} x_i$ by selecting good i whenever possible.

Note that all bundles \mathbf{x}^θ at price \mathbf{p}^θ have the same number of items of i and are also demanded at \mathbf{p} . As $\mathbf{p} \preceq \mathbf{p}^\theta$ and $p_i = p_i^\theta$, the strong-substitutes property implies that there exists $\mathbf{x}^\theta \in D(\mathbf{p}^\theta)$ such that $x_i = x_i^\theta$ for all $\mathbf{x} \in D(\mathbf{p})$. In other words, any demanded bundle \mathbf{x}^θ at \mathbf{p}^θ maximises x_i . \square

D Hardness of testing validity of unrestricted bid lists

Here we show that the question of whether a list of bids is valid is coNP-complete, even when we restrict the problem to unit bids (bids with weights = 1). We also present a simple algorithm verifying the validity of a given list of positive and negative unit bids that runs in polynomial time if the number of goods, or the number of negative bids, is bounded by a constant. Let B denote a list of unit bids, and let B^+ and B^- denote the positive and negative bids.

D.1 Checking validity is coNP-complete

Definition 8. Given a list of positive and negative bids B , the problem Valid Bids is to decide whether B is valid.

Theorem 38 (Theorem 2). Valid Bids is coNP-complete.

The proof of Theorem 38 uses an equivalent definition of validity for the list of bids. Define regions in \mathbb{R}_+^n that are generated by a point \boldsymbol{p} and coordinates $i, j \in [n]$ as follows.

$$H_i^{\boldsymbol{p}} := \{ \boldsymbol{x} \in \mathbb{R}_+^n \mid \boldsymbol{p}, x_i = p_i \} \quad (33)$$

$$F_{ij}^{\boldsymbol{p}} := \{ \boldsymbol{x} \in \mathbb{R}_+^n \mid \boldsymbol{p}, x_i = p_i \text{ and } x_j = p_j \} \quad (34)$$

A bid $\boldsymbol{b} = (b_1, \dots, b_n; b_{n+1})$ is *contained* in a region $H_i^{\boldsymbol{p}}$ (or $F_{ij}^{\boldsymbol{p}}$) if the ‘valuation’ vector (b_1, \dots, b_n) consisting of the first n components lies in it. Moreover, we say that $H_i^{\boldsymbol{p}}$ or $F_{ij}^{\boldsymbol{p}}$ is *negative* for a list of bids if it contains more negative than positive bids, otherwise it is *non-negative*.

Observation 39. For any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}_+^n$, we have $\boldsymbol{x} \in F_{ij}^{\boldsymbol{y}}$ if and only if $x_i - y_i = x_j - y_j$ and $\boldsymbol{x} \leq (x_i - y_i)\mathbf{1} - \boldsymbol{y}$.

Observation 40. Fix $i, j \in [n], i \neq j$. Then containment in the regions given in (33) and (34) is transitive in the sense that, for any $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathbb{R}_+^n$,

- $\boldsymbol{x} \in H_i^{\boldsymbol{y}}$ and $\boldsymbol{y} \in H_i^{\boldsymbol{z}}$, implies $\boldsymbol{x} \in H_i^{\boldsymbol{z}}$.
- $\boldsymbol{x} \in F_{ij}^{\boldsymbol{y}}$ and $\boldsymbol{y} \in F_{ij}^{\boldsymbol{z}}$, implies $\boldsymbol{x} \in F_{ij}^{\boldsymbol{z}}$.

The following definition of valid bids restates the second part of Theorem 1 in terms of regions (33) and (34).

Definition 9 (Valid bids). A list of positive and negative bids is *valid* if, for any point $\boldsymbol{p} \in \mathbb{R}_+^n$ and two coordinates $i, j \in [n]$, $H_i^{\boldsymbol{p}}$ and $F_{ij}^{\boldsymbol{p}}$ are non-negative.

On the basis of Definition 9, the proof of Theorem 2 works by reducing the well-known NP-complete problem 3-CNF Satisfiability to Valid Bids by means of an intermediate NP-complete decision problem More Negatives Dominated, which we define in Definition 10. The NP-completeness of More Negatives Dominated is established in Theorem 41 and the reduction from More Negatives Dominated to Valid Bids is given below as the proof of Theorem 2. For any two points $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, we say that \boldsymbol{x} *dominates* \boldsymbol{y} if $x_i - y_i$ for all $i \in \{1, \dots, n\}$.

Definition 10. Given two lists P and N of vectors in \mathbb{R}_+^n , More Negatives Dominated is the problem of deciding whether there exists a vector $\boldsymbol{z} \in \mathbb{R}_+^n$ that dominates more vectors in N than in P .

Theorem 41. The problem More Negatives Dominated is NP-complete.

Proof. Note that verifying whether a given point $\boldsymbol{z} \in \mathbb{R}_+^n$ dominates more vectors in N than in P can be done in polynomial time. This establishes membership of More Negatives Dominated in the class NP. We show completeness by reducing from 3-CNF Satisfiability. Recall that a boolean formula ϕ defined on n variables x_1, \dots, x_n is 3-CNF if

$$\phi = \bigwedge_{i=1}^m C_i,$$

where $C_i = L_1^i _ L_2^i _ L_3^i$ are its clauses and $L_j^i \in \{f x_1, \dots, x_n, \overline{x_1}, \dots, \overline{x_n}\}g$ are its literals. Here $\overline{x_i}$ denotes the negation of x_i . Without loss of generality, we assume that the three variables appearing in each clause are distinct. Given such a 3-CNF formula ϕ , first construct a pair of lists (P_i, N_i) of vectors in \mathbb{R}_+^n for each clause C_i as follows. For notational convenience, $\mathbf{e}^{a\dots z}$ denotes the characteristic vector $\mathbf{e}^{f a, \dots, z g}$ of $f a, \dots, z g$.

- If $C_i = x_a _ x_b _ x_c$, let $N_i = f \mathbf{e}^i g$ and $P_i = f \mathbf{e}^{abc} g$.
- If $C_i = x_a _ x_b _ \overline{x_c}$, let $N_i = f \mathbf{e}^i, \mathbf{e}^{abc} g$ and $P_i = f \mathbf{e}^{ab} g$.
- If $C_i = x_a _ \overline{x_b} _ \overline{x_c}$, let $N_i = f \mathbf{e}^i, \mathbf{e}^{ab}, \mathbf{e}^{ac} g$ and $P_i = f \mathbf{e}^a, \mathbf{e}^{bc} g$.
- If $C_i = \overline{x_a} _ \overline{x_b} _ \overline{x_c}$, let $N_i = f \mathbf{e}^a, \mathbf{e}^b, \mathbf{e}^c, \mathbf{e}^{abc} g$ and $P_i = f \mathbf{e}^{ab}, \mathbf{e}^{ac}, \mathbf{e}^{bc} g$.

Finally, let N be the sum of the N_i and let P be the sum of the P_i together with $m - 1$ copies of \mathbf{e}^i . This completes our reduction from 3-CNF Satisfiability to More Negatives Dominated. Clearly, P and N can be constructed from ϕ in polynomial time. To show correctness of the reduction, associate with every vector $\mathbf{z} \in \mathbb{R}_+^n$ a truth assignment β_z (on boolean variables x_1, \dots, x_n) that sets x_i to true if $z_i < 1$ and false otherwise.

Observation 42. *Let $\mathbf{z} \in \mathbb{R}_+^n$ be a point and β_z be its associated truth assignment. For any $i \in \{1, \dots, mg\}$, \mathbf{z} dominates one more point in N_i than in P_i if β_z satisfies C_i and an equal number of points in N_i and P_i if β_z does not satisfy C_i .*

Suppose ϕ is a satisfiable 3-CNF formula with satisfying truth assignment β . Then define $\mathbf{z} \in \mathbb{R}_+^n$ by

$$z_i = \begin{cases} 0 & \text{if } \beta[x_i] = \text{true,} \\ 1 & \text{else.} \end{cases}$$

Hence β is the truth assignment associated with \mathbf{z} . By Observation 42, \mathbf{z} dominates one more point in N_i than in P_i for each i , so due to the additional $m - 1$ origin points \mathbf{e}^i added to P , \mathbf{z} dominates exactly one more point in N than in P . Conversely, suppose there exists a point $\mathbf{z} \in \mathbb{R}_+^n$ that dominates more points in N than in P , and let β_z be its corresponding truth assignment. Then \mathbf{z} dominates at least m more points in N than in $\sum_{i=1}^m P_i$. By Observation 42, we know that \mathbf{z} dominates at most one more point in N_i than in P_i , which implies that this is the case for each i . Hence by the same observation, β_z satisfies all clauses of ϕ and thus ϕ itself. \square

Proof of Theorem 2. In order to show that a list of bids B is not valid, it suffices to provide a certificate in the form of a point \mathbf{p} and coordinates $i, j \in [n]$. We can verify in polynomial time whether $H_i^{\mathbf{p}}$ or $F_{ij}^{\mathbf{p}}$ is negative. This establishes membership of Valid Bids in coNP.

Let (N, P) be an instance of More Negatives Dominated and n be the dimension of its vectors. We construct an $(n + 1)$ -dimensional instance B of Valid Bids such that $(N, P) \in$ More Negatives Dominated if and only if $B \notin$ Valid Bids as follows. Let the negative and positive bids of B be given by prepending a component to the vectors as follows.

$$B^- := f(1, \mathbf{v}) \mid \mathbf{v} \in Ng$$

$$\text{and } B^+ := f(1, \mathbf{w}) \mid \mathbf{w} \in Pg + f(0, \mathbf{v}) \mid \mathbf{v} \in Ng + f(1, \mathbf{v}) + \mathbf{e}^{[n+1]} \mid \mathbf{v} \in Ng.$$

Clearly, $B = B^+ + B^-$ can be constructed efficiently. Note that all $F_{ij}^{\mathbf{x}}$ and all $H_i^{\mathbf{x}}$ with $i \neq 0$ are non-negative by construction. Indeed, fix $F_{ij}^{\mathbf{x}}$ and suppose it contains r negative bids. Then for each such bid $(1, \mathbf{v})$, where $\mathbf{v} \in N$, there exists a positive bid $(1, \mathbf{v}) + \mathbf{e}^{[n+1]}$ that is contained in $F_{ij}^{\mathbf{x}}$, so that the region is non-negative. Analogously, for any $H_i^{\mathbf{x}}$ with $i \neq 0$, each negative bid $(1, \mathbf{v})$ has a corresponding positive bid $(0, \mathbf{v})$. Further, as $H_0^{\mathbf{x}}$ contain no negative points unless $x_0 = 1$, the correctness of our reduction thus follows from the fact that \mathbf{z} dominates r points from N and s points from P if and only if $H_0^{\mathbf{x}}$ with $\mathbf{x} = (1, \mathbf{z})$ contains r negative and s positive bids. \square

D.2 A simple algorithm for testing validity

In Section D.1, we show that checking the validity of a bid list in the general case is coNP-complete. Hence we cannot expect an efficient algorithm for the task of checking validity of a list of bids in the general case. Here we present a simple algorithm that tests validity of a given list of bids in polynomial time if the number of goods or the number of negative bids is bounded by a constant. Such constraints may be reasonable in certain economic settings.

D.2.1 The algorithm

Our procedure rests on Theorem 43, which reduces the validity condition from Definition 9 to a finite number of checks. Lemma 44 reduces the number of checks further. Together, the two results immediately yield Algorithm 10, whose running time is given in Theorem 45. For any list of bids U , define the minimal dominating vector $md(U)$ of U as the component-wise maximum over the valuation vectors of the bids in U , so that $md(U)_i = \max_{\mathbf{b} \in U} b_i$ for all $i \in [n]$. We note that if all bids of U agree on some coordinate i (that is, if $b_i = b_i^\theta$ for all $\mathbf{b}, \mathbf{b}^\theta \in U$), they lie in H_i^x for large enough \mathbf{x} , and $md(U)$ is the minimal such \mathbf{x} so that H_i^x contains all points in U . Similarly, for any list of bids U and $i \in [n]$, define $mdF(i, U)$ as $mdF(i, U) := \min_{\mathbf{b} \in U} b_i \mathbf{1} + md(\mathbf{b} - b_i \mathbf{1} \mid \mathbf{b} \in U)$. We note that if, for some $i \neq j$, the set U satisfies $b_i - b_j = b_i^\theta - b_j^\theta$ for all $\mathbf{b}, \mathbf{b}^\theta \in U$, then the bids in U lie in F_{ij}^x for small enough \mathbf{x} , and it follows from Observation 39 that $mdF(i, U) = mdF(j, U)$ is the maximal such \mathbf{x} .

Theorem 43. *A list of bids B is valid if and only if the following two conditions hold.*

1. *For every set $U \subseteq B$ of negative bids that agree on the i -th coordinate, that is $b_i = b_i^\theta$ $\forall \mathbf{b}, \mathbf{b}^\theta \in U$, the region $H_i^{md(U)}$ is non-negative.*
2. *For every set $U \subseteq B$ of negative bids that satisfies $b_i - b_i^\theta = b_j - b_j^\theta$ for all $\mathbf{b}, \mathbf{b}^\theta \in U$ and some $i \neq j$, the region $F_{ij}^{mdF(i, U)}$ is non-negative.*

Proof. The implication is immediate by Definition 9. Conversely, suppose B satisfies the second condition. First we show that condition 1 of Definition 9 is satisfied. Fix H_i^x , let $U = B \cap H_i^x$ be the set of negative bids in H_i^x and $\mathbf{y} = md(U)$. Then by construction, we have $\mathbf{b} \in H_i^y$ for all $\mathbf{b} \in U$ and $\mathbf{y} \in H_i^x$. As H_i^y has $|U|$ negative bids, it also contains at least $|U|$ positive bids by assumption and by Observation 40, these positive bids are also in H_i^x . Condition 2 is shown analogously. Fix F_{ij}^z , let U be the negative bids in this region and $\mathbf{z} = mdF(U)$. Suppose $\mathbf{b} \in F_{ij}^z$ for every $\mathbf{b} \in U$ and $\mathbf{z} \in F_{ij}^z$. Then Observation 40 implies that F_{ij}^z is non-negative. To see that $\mathbf{b} \in F_{ij}^z$ for every $\mathbf{b} \in U$, we verify the conditions of Observation 39. Firstly,

$$z_i - z_j = \min_{\mathbf{b} \in U} b_i - (\min_{\mathbf{b} \in U} b_i + \max_{\mathbf{b} \in U} (b_j - b_i)) = \max_{\mathbf{b} \in U} (b_j - b_i) = b_j - b_i$$

for all $\mathbf{b} \in U$, as the difference $b_j - b_i$ is the same for all $\mathbf{b} \in F_{ij}^z$. Secondly, for any $k \in [n]$ and $\mathbf{b} \in U$, we have

$$(\mathbf{b} - (b_i - z_i)\mathbf{1})_k = z_i + b_k - b_i - z_i + \max_{\mathbf{b} \in U} (b_k - b_i) = z_k.$$

We can verify $\mathbf{z} \in F_{ij}^z$ similarly. □

Lemma 44. *For any list of negative bids $U \subseteq B$ there exists a list $U^\theta \subseteq U$ with $|U^\theta| = n$ so that $md(U^\theta) = md(U)$. Moreover, for each coordinate $i \in [n]$, there exists a list $U^{\theta\theta} \subseteq U$ with $|U^{\theta\theta}| = n + 1$ so that $mdF(i, U^{\theta\theta}) = mdF(i, U)$.*

Algorithm 10 Checking the validity of a bid list B

- 1: **for all** subsets $U \subseteq B$ of negative bids with $|U| = n + 1$ **do**
 - 2: **for all** $i, j \in [n]$ **do**
 - 3: Verify the two conditions given in Theorem 43.
-

Proof. Let U^θ be a list of bids $\mathbf{u}^1, \dots, \mathbf{u}^n$, where \mathbf{u}^k is a bid from U that maximises the k -th component, i.e. $\mathbf{u}^k \in \arg \max_{\mathbf{v} \in U} v_k$. Then $md(U^\theta) = md(U)$ by construction. Secondly, fix $i \in [n]$ and let $U^{\theta\theta}$ be a list of bids $\mathbf{u}^1, \dots, \mathbf{u}^{n+1}$, where $\mathbf{u}^1 \in \arg \min_{\mathbf{b} \in U} b_i$ and $\mathbf{u}^k \in \arg \max_{\mathbf{b} \in U} (b_k - b_i)$ for $k \in \{2, \dots, n+1\}$. Then $mdF(U^{\theta\theta}) = mdF(U)$ by construction and we are done. \square

Theorem 45. *Algorithm 10 runs in time $O\left(\binom{jB}{n} n^3 jBj\right)$. Moreover, if $jB - j = k$ or $n - k$, the algorithm's running time is polynomial in the input size $njBj$.*

Proof. There are $\binom{jB}{n+1}$ subsets U of B for the algorithm to iterate over, while there are n^2 possibilities for the index pair i, j . For each U and i, j , checking the conditions in Theorem 43 takes $njBj$ time. This implies the running time. Suppose the number of negative bids $jB - j = k$ is bounded by some $k \in \mathbb{N}$. Then $\binom{jB}{n+1}$ is constant and the running time of our algorithm is $O(n^3(jVj + jWj))$. Secondly, if the number of goods $n - k$ is bounded, we have $\binom{jB}{n+1} = jB - j^{k+1}$ and hence a running time of $O(jB - j^{k+1} jBj)$. \square