

Learning Strong Substitutes Demand via Queries

PAUL W. GOLDBERG and EDWIN LOCK, University of Oxford, UK

FRANCISCO J. MARMOLEJO-COSSÍO, Harvard University, US

This paper addresses the computational challenges of learning strong substitutes demand when given access to a demand (or valuation) oracle. Strong substitutes demand generalises the well-studied gross substitutes demand to a multi-unit setting. Recent work by Baldwin and Klemperer shows that any such demand can be expressed in a natural way as a finite list of weighted bid vectors. A simplified version of this bidding language has been used by the Bank of England.

Assuming access to a demand oracle, we provide an algorithm that computes the unique list of weighted bid vectors corresponding to a bidder's demand preferences. In the special case where their demand can be expressed using positive bids only, we have an efficient algorithm that learns this list in linear time. We also show super-polynomial lower bounds on the query complexity of computing the list of bids in the general case where bids may be positive and negative. Our algorithms constitute the first systematic approach for bidders to construct a bid list corresponding to non-trivial demand, allowing them to participate in 'product-mix' auctions.

CCS Concepts: • **Theory of computation** → **Computational pricing and auctions**; *Algorithmic mechanism design*.

Additional Key Words and Phrases: learning demand, preference elicitation, bidding language, query protocol, product-mix auction, strong substitutes

ACM Reference Format:

Paul W. Goldberg, Edwin Lock, and Francisco J. Marmolejo-Cossio. 2022. Learning Strong Substitutes Demand via Queries. *ACM Trans. Econ. Comput.* 1, 1 (May 2022), 22 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

The product-mix auction [10–12] was devised by Klemperer as a means of providing liquidity to commercial banks and has been used regularly by the Bank of England since 2011. In it, there are a number of distinct goods available in multiple discrete units, and a set of buyers who express strong substitutes demands amongst these goods.¹ Given these strong substitutes constraints on the bidders' demand, it is possible to compute competitive equilibrium, which consists of market-clearing prices and allocations so that all buyers receive an allocation that they demand at those prices, and all goods are sold [1]. The strong substitutes property guarantees the existence of a competitive equilibrium.²

Importantly for the present paper, the auction introduces a novel bidding language in which buyers express their demands in terms of lists of bids, where each bid consists of a price vector (one price for each good) and a weight. Any

¹In the banking context, the goods correspond to liquidity secured against alternative kinds of collateral. Commercial banks pay for liquidity 'products' by committing to interest rates. The values of the bids submitted by the commercial banks correspond to the interest rates they are willing to pay. More information about how the Bank of England uses the product-mix auction can be found in [11, 12].

²In many markets, the fundamental theorems of welfare economics guarantee that competitive equilibrium coincides with social welfare maximising prices and allocations, and this also holds for the product-mix auction. Hence the product-mix auction is designed to find welfare-optimal outcomes.

Authors' addresses: Paul W. Goldberg, paul.goldberg@cs.ox.ac.uk; Edwin Lock, edwin.lock@cs.ox.ac.uk, University of Oxford, Department of Computer Science, Oxford, UK; Francisco J. Marmolejo-Cossio, fjmarmol@seas.harvard.edu, Harvard University, Department of Computer Science, Cambridge, MA, US.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

bid \mathbf{b} is understood as a willingness to buy some quantity of goods (the weight of \mathbf{b}), and for each good i a price b_i is offered. A bid is rejected if all prices offered are lower than the market-clearing prices of the corresponding goods, otherwise it is accepted on some good that maximises the price offered minus the market-clearing price.³ The auction currently run by the Bank of England only permits bidders to submit bids with one non-zero vector entry and positive weights, and it is straightforward to show that any such list of such positive bids has the strong substitutes property. Conversely, it has subsequently been shown that any strong-substitutes valuation with a discrete simplex domain can be uniquely represented as a list of bids with positive *and* negative weights. [3]. Moreover, for any $M > 0$, it is possible to compute an extension of the valuation to a discrete simplex domain while ensuring that demand remains the same at all prices in ‘bounding box’ $\{\mathbf{p} \in R^n \mid \|\mathbf{p}\|_\infty \leq M\}$, and the extension maintains the strong-substitutes property for these prices. In particular, by choosing M arbitrarily high, we can represent valuations using bids at any prices of interest.

While this gives the buyer a general-purpose means of communicating any strong substitutes demand, the buyer faces the problem of expressing her demand in this language. It may be easier for a buyer to answer queries of the form “What bundle would you demand, given the following per-unit prices of goods?”. In this paper, we develop query protocols that assist a buyer in constructing her demand function based on a sequence of such queries. Given an unknown demand correspondence, our algorithms are assumed to have access to a demand oracle: for any given prices for goods, the algorithms can learn a bundle of goods demanded at those prices. We are interested in minimising the number of queries to the demand oracle.

1.1 Our Contributions

This paper addresses the computational challenges of learning strong substitutes demand when given access to a demand oracle. Under the mild assumption that bidders are able to answer questions of the form “What bundle do you demand at the following per-unit prices?”, our algorithms constitute the first systematic approach for bidders to generate a bid list corresponding to their demand, allowing them to participate in product-mix auctions with non-trivial demand preferences. We provide upper and lower bounds on the query complexity of learning demand preferences and expressing these using the bidding language of the product-mix auction, which is able to encode any strong substitutes demand in a conceptually simple and natural fashion.

Section 2 outlines three complementary characterisations of the strong substitutes property and introduces the bidding language both algebraically and geometrically. A first result of this paper, given in Section 3, is to show that demand oracles are not unreasonably powerful: when given access instead to a valuation oracle, it is possible to simulate a demand oracle with $O(n^3 \log(W/n))$ valuation queries, where n is the number of goods and W is the maximum weight of a bid vector.

In Sections 4 and 5, we consider algorithms that learn the unique bid list corresponding to a bidder’s demand. The algorithm in Section 4 learns demands that can be represented by lists of positive bids, and has linear query complexity. In the setting where demand may require positive and negative bids to express, we provide an exponential-cost algorithm that proceeds by learning all hyperplanes that contain facets of the Locus of Indifference Prices (LIP), a geometric object introduced by Baldwin and Klemperer [2] to characterise demand.

Finally, in Section 6 we consider lower bounds on the query complexity of learning bid lists. We note briefly that $\Omega(B \log M)$ queries are required to learn a list of B positive bids, where M is the magnitude of the bid vectors w.r.t. the

³In this way, a single bid with vector \mathbf{b} and weight 1 is functionally equivalent to a unit-demand bidder with valuation \mathbf{b} and quasi-linear utility.

L_∞ norm. In order to identify the dependence on the number of goods n , we construct an adversarial game using a novel ‘island gadget’ consisting of bids with weight ± 1 . Crucially, the island gadget only changes demand in a local region. For fixed n , we identify the overall query complexity of learning bid lists corresponding to strong substitutes demand as $\Theta(B \log M + B^n)$.

1.2 Related Work

Our work relates to the theory of preference elicitation. In this setting, a centralised agent, such as an auctioneer, wishes to identify an optimal allocation of goods via queries to participants’ preferences. Queries typically take the form of value queries, where an agent reports a valuation for a given bundle of goods, or demand queries, where an agent reports a bundle that is demanded at given prices. This paper focuses on using demand queries to learn the bid list representation of strong substitutes demand preferences. This representation can then be used to compute an optimal allocation of goods to agents via the methods of [1].

Much early work in preference elicitation highlights the deep connections to exact learning via membership and equivalence queries from computational learning theory, and our results can also be viewed through this lens. Some notable examples include [4, 21]. The authors of [5] explore the use of ranking oracles to exploit the topological structure of bidder preferences to learn optimal allocations. This approach is extended in [6, 7], and verified empirically in [9]. The authors of [16] explore the communication complexity of preference elicitation in combinatorial auctions, where they show that for general valuations, finding a value-maximising allocation requires an exponential communication cost in the number of items. In [13], the authors explore connections between preference elicitation and exact learning, but they demonstrate that the representation length of the valuation is an important parameter in the query complexity of computing optimal allocations. This dependence on the representation length provides a way of side-stepping lower bounds from [16], and further justifies the need for succinct yet expressive bidding languages, as explored further in [15].

Our problem is conceptually similar to a problem studied recently in [20], in which the authors also consider algorithms with access to demand queries, and attempt to learn the underlying valuation that gives rise to the demand correspondence. The main difference between their work and ours however is that they consider a different class of value functions. In [20], there are n goods, each in unit supply (whereas we allow multiple copies of goods), and the buyer wants at most k goods, and has additive valuations (whereas our strong substitutes valuations are more general).

2 PRELIMINARIES

We denote $[n] := \{1, \dots, n\}$ and $[n]_0 := \{0, \dots, n\}$. In our auction model, there are n distinct goods numbered from 1 to n ; a single copy of a good is an *item*⁴. A *bundle* of goods, typically denoted \mathbf{x} or \mathbf{y} in this paper, is a vector in \mathbb{N}_0^n whose i -th entry denotes the number of items of good i . 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. Occasionally, it is convenient to work with a notional *reject good* 0 for which prices are always zero. The set of goods is then $[n]_0$ and we identify bundles and prices with the $n + 1$ -dimensional vectors obtained by adding a 0-th entry of value 0. 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.

⁴Items are also known as *units* in the literature.

For any vector $\mathbf{v} \in \mathbb{R}^n$, the L_1 and L_∞ norms are defined as $\|\mathbf{v}\|_1 = \sum_{i \in [n]} |v_i|$ and $\|\mathbf{v}\|_\infty = \max_{i \in [n]} |v_i|$. The L_∞ ball $B_\mathbf{p}^\varepsilon$ of radius ε at \mathbf{p} consists of all points $\mathbf{q} \in \mathbb{R}^n$ that satisfy $\|\mathbf{p} - \mathbf{q}\|_\infty \leq \varepsilon$; note that this is a hypercube with edge length 2ε . Any hypercube centred at \mathbf{p} can be partitioned into 2^n orthants, where every orthant O_a is described by some vector $\mathbf{a} \in \{-1, 1\}^n$ and consists of the set of points $O_a = \mathbf{p} + \{\mathbf{x} \in \mathbb{R}^n \mid a_i x_i \geq 0, \forall i \in [n]\}$.⁵ Every such orthant can be triangulated into $n!$ simplices as follows. For every ordering $[i_1, \dots, i_n]$ of the indices $[n]$, we define a simplex as the set of points in the orthant that satisfy $a_{i_1} x_{i_1} \leq a_{i_2} x_{i_2} \leq \dots \leq a_{i_n} x_{i_n}$.

2.1 Strong-Substitutes Demand Preferences

Throughout, we assume that bidders have quasi-linear *strong substitutes* (SS) demand. The SS property is appealing because it is a generalisation of *gross substitutes* (GS) from the single-unit setting that guarantees the existence of a competitive equilibrium in multi-unit auction markets. We first present a characterisation of SS by Shioura and Tamura [18] that elucidates the relationship between GS and SS before introducing the two equivalent characterisations that underpin our algorithmic results and draw from tropical geometry and discrete convex analysis, respectively. For a detailed survey on the relationship between GS and SS, we refer to Shioura and Tamura [18].

In this work, we assume that bidders have an implicit *valuation* $v : A \rightarrow \mathbb{R}$ for bundles of goods, where the domain A is a *discrete simplex domain* $A = \text{conv}\{\mathbf{m}\mathbf{e}^i \mid i \in [n]\}$ for some $m > \mathbb{N}_0$.⁶ This is equivalent to defining the valuation as $v : \mathbb{Z}^n \rightarrow \mathbb{R}$, where $\mathbb{R} := \mathbb{R} \cup \{-\infty\}$ denotes the partially extended reals, and we assume that the *effective domain* $\text{dom } v = \{\mathbf{x} \in \mathbb{Z}^n \mid v(\mathbf{x}) > -\infty\}$ of v is a finite simplex domain and non-negative in the sense that $\mathbf{x} \geq \mathbf{0}$ for all $\mathbf{x} \in \text{dom } v$. Moreover, bidders have quasi-linear utilities, i.e. the utility they derive from bundle \mathbf{x} at prices \mathbf{p} is $u(\mathbf{x}; \mathbf{p}) := v(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x}$. A bidder's *demand correspondence* D_v maps prices \mathbf{p} to the set of bundles $\mathbf{x} \in A$ that maximise $u(\mathbf{x}; \mathbf{p})$ for this \mathbf{p} .

Definition 2.1 captures how GS and SS demand changes when we (weakly) increase prices. Intuitively, the GS property states that the bidder's demand for those goods with unchanged prices does not decrease, while the law of aggregate demand (LAD) guarantees that the overall number of items that are demanded does not increase. The SS property combines the GS property with LAD.

Definition 2.1 (cf. [18]). A demand correspondence D_v is *gross substitutes* (GS) if, for any prices $\mathbf{p}' \geq \mathbf{p}$ with $D_v(\mathbf{p}) = \{\mathbf{x}\}$ and $D_v(\mathbf{p}') = \{\mathbf{x}'\}$, we have $x'_k \geq x_k$ for all k such that $p_k = p'_k$. D_v is *strong substitutes* (SS) if \mathbf{x} and \mathbf{x}' additionally satisfy $\|\mathbf{x}'\|_1 \leq \|\mathbf{x}\|_1$ (the law of aggregate demand).

2.1.1 Geometric Approach. We give some geometric intuition for strong substitutes demand correspondences that underpins the algorithmic ideas in this paper. It is well-known that any quasi-linear demand divides price space into piecewise-linear convex regions corresponding to bundles. When demand is SS, each such region is a convex lattice [14]. Figure 1 illustrates this.

Recently, Baldwin and Klemperer [2] proposed a new way of characterising demand types. Borrowing from the tropical geometry literature, they introduce the *Locus of Indifference Prices* (LIP), a piecewise-linear geometric object consisting of the set of all prices at which the bidder is indifferent between two or more bundles. They show that the LIP corresponds in a natural way to a polyhedral complex with $n - 1$ -dimensional facets. In Figure 1, the LIP is drawn using dashed lines. Noting that the orientation of the separating facet between two adjacent demand regions characterises how demand changes when moving from one region to the other, Baldwin and Klemperer [2] propose a new way

⁵Orthants in n -dimensional space generalise the notion of quadrants and octants in two- and three-dimensional space, respectively.

⁶If the domain of v is not a discrete simplex domain, we can, for any arbitrary bounding box $\{\mathbf{p} \in \mathbb{R}^n \mid \|\mathbf{p}\|_\infty \leq 1\}$, compute an extension of v to a discrete simplex domain that remains strong-substitutes on this bounding box, and such that demand coincides with the original valuation at any prices in the box. All results in this paper follow through for such an extension, if we restrict our algorithms, definitions and bids to this bounding box.




Fig. 1. Left: An illustration of a strong-substitutes demand correspondence with two goods, partitioning price space into piecewise-linear convex regions. Each region is labelled with the bundle demanded at prices in the region. The dashed lines comprise the Locus of Indifference Prices (LIP). Right: Six positive (solid) and two negative (hollow) bids of unit weight are required to express this demand. (Note the two positive unit-weight bids at $(2, 4)$, which can alternatively be normalised to a single bid of weight 2.)

of defining demand types by the set of facet-normal vectors of the LIP. In this new paradigm, the strong substitutes demand type is defined as the family of demand correspondences whose LIP facets are normal to \mathbf{e}^i or $\mathbf{e}^i - \mathbf{e}^j$ for some $i, j \in [n]$. In two dimensions, facets of SS LIPs are either horizontal, vertical or normal to $(1, -1)$. Hence it follows directly from this definition that the demand correspondence in Figure 1 enjoys the strong substitutes property.

2.1.2 Discrete Convex Analysis. A ‘price-free’ characterisation of the strong substitutes property using the language of discrete convex optimisation is given by Shioura and Tamura [18]. A function $f : \mathbb{Z}^n \rightarrow \mathbb{R}$ is called M^{\natural} -concave if it satisfies the following *exchange property*. For any $\mathbf{x}, \mathbf{y} \in \text{dom } f$ and $i \in \text{supp}^+(\mathbf{x} - \mathbf{y})$, there exists $j \in \text{supp}^-(\mathbf{x} - \mathbf{y}) \cup \{0\}$ such that

$$f(\mathbf{x}) + f(\mathbf{y}) \leq f(\mathbf{x} - \mathbf{e}^i + \mathbf{e}^j) + f(\mathbf{y} + \mathbf{e}^i - \mathbf{e}^j). \quad (1)$$

Here we define $\mathbf{e}^0 = 0$, and the positive and negative support of a vector $\mathbf{z} \in \mathbb{Z}^n$ as $\text{supp}^+(\mathbf{z}) = \{i \in [n] \mid z_i > 0\}$ and $\text{supp}^-(\mathbf{z}) = \{i \in [n] \mid z_i < 0\}$.

THEOREM 2.2 ([18, THEOREM 4.1]). *A quasi-linear demand correspondence D_u is strong substitutes if and only if its valuation u is M^{\natural} -concave.*

M^{\natural} -concave functions are closely related to M -concave functions, which satisfy the exchange property (1) for some non-zero $j \in \text{supp}^-(\mathbf{x} - \mathbf{y})$. Every n -dimensional M^{\natural} -concave function can be obtained as the projection of an $n + 1$ -dimensional M -concave function onto an n -dimensional hyperplane. Conversely, we can obtain the corresponding M -concave function \hat{f} of an M^{\natural} -concave function f as

$$\hat{f}(x_0, \mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } x_0 = -\sum_{i \in [n]} x_i, \\ -\infty & \text{otherwise,} \end{cases} \quad (2)$$

where $(x_0, \mathbf{x}) \in \mathbb{N}_0^{n+1}$ is an $n + 1$ -dimensional vector. For details on M^{\natural} - and M -concave functions, we refer to Murota [14].

2.2 The Bidding Language

The product-mix auction introduces a novel bidding language that allows us to express every strong substitutes demand as a finite list \mathcal{B} of positive and negative bids. A *bid* consists of an n -dimensional integral⁷ vector $\mathbf{b} \in \mathbb{N}_0^n$ and a weight $w(\mathbf{b}) \in \mathbb{Z} \setminus \{0\}$. When working with the notional reject good 0 introduced above, we identify a bid vector \mathbf{b} with the $n+1$ -dimensional vector obtained by adding a 0-th entry of value 0. As we will see, any bid with a weight of $w(\mathbf{b}) \in \mathbb{Z}$ is equivalent to $|w(\mathbf{b})|$ unit bids with the same vector and sign. This allows us to normalise bid lists to their most succinct form, where no two bids share the same vector. In this paper, we wish to learn the unique normalised bid list that represents the bidder's demand correspondence.

For each bid \mathbf{b} , we can understand b_i as the amount that \mathbf{b} is willing to spend on good i . It is also convenient to understand the vector \mathbf{b} as the location of the bid in price-space, as illustrated in Figure 2. Suppose the auctioneer sets prices \mathbf{p} . The bid is *rejected* at \mathbf{p} if $b_i < p_i$ for all goods i . On the other hand, if $\max_{i \in [n]} (b_i - p_i) > 0$, then the bid demands a bundle containing a total of $w(\mathbf{b})$ items of the goods $i \in [n]$ maximising $b_i - p_i$. Finally, if $\max_{i \in [n]} (b_i - p_i) = 0$, then the bid demands any bundle containing *up to* $w(\mathbf{b})$ items of goods $\arg \max_{i \in [n]} (b_i - p_i)$ (including the empty bundle 0). The notional 'reject' good 0 simplifies notation: recalling that we defined $b_0 = 0 = p_0$, we say that \mathbf{b} demands good $i \in [n]_0$ if $i \in \arg \max_{i \in [n]_0} (b_i - p_i)$, and receiving an item of the 'reject' good 0 is equivalent to receiving an item of 'nothing'. In particular, receiving $w(\mathbf{b})$ items of 0 is equivalent to the bid being rejected. If the set of demanded goods $\arg \max_{i \in [n]_0} (b_i - p_i)$ at \mathbf{p} contains more than one good, we say that \mathbf{b} is *indifferent* between these goods at \mathbf{p} . (In particular, a bid may be indifferent between demanding goods and being rejected when $\max_{i \in [n]_0} (b_i - p_i) = 0$). A price \mathbf{p} is *marginal* if there are bids indifferent between goods (which may include the reject good 0) at \mathbf{p} , and non-marginal otherwise.

We can now introduce the *demand correspondence* $D_{\mathcal{B}}(\mathbf{p})$ for a bid list \mathcal{B} as follows. If \mathbf{p} is non-marginal, the unique bundle demanded at \mathbf{p} is obtained by adding $w(\mathbf{b})$ items of $i(\mathbf{b})$ to the bundle for each $\mathbf{b} \in \mathcal{B}$, where $i(\mathbf{b})$ is the unique good that \mathbf{b} demands at \mathbf{p} .⁸ If \mathbf{p} is marginal, $D_{\mathcal{B}}(\mathbf{p})$ 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 $\text{conv}(X) \cap \mathbb{N}_0^n$. In general, this implies that when \mathcal{B} contains negative bids, we cannot independently allocate to each bid one of the goods it demands, as this may result in bundles that are not in $D_{\mathcal{B}}(\mathbf{p})$. While a single bid with positive weight $w(\mathbf{b}) = 1$ corresponds to a unit-demand bidder, a negatively-weighted bid has no immediate economic interpretation in isolation. However, we can still express the demand correspondence $D_{\mathbf{b}}$ of a single bid \mathbf{b} with negative weight $w(\mathbf{b})$ as the negation of the correspondence $D_{\mathbf{b}'}$ of the 'negated' positive bid \mathbf{b}' with $\mathbf{b}' = \mathbf{b}$ and $w(\mathbf{b}') = |w(\mathbf{b})|$, so that $D_{\mathbf{b}} = -D_{\mathbf{b}'}$. Moreover, it is immediately apparent from the definition of $D_{\mathcal{B}}$ that a bid \mathbf{b} with weight $w(\mathbf{b}) \in \mathbb{Z} \setminus \{0\}$ is equivalent to $|w(\mathbf{b})|$ unit-weight bids of the same sign and at the same location.

Given a 'bounding box' $\{\mathbf{p} \in \mathbb{R}^n \mid \|\mathbf{p}\|_{\infty} \leq M\}$ for arbitrarily large M , Baldwin and Klemperer [3] show that any strong substitutes demand correspondence D_v can be represented as a finite list \mathcal{B} of positive and negative bids contained in the box, such that $D_v(\mathbf{p}) = D_{\mathcal{B}}(\mathbf{p})$ at all prices \mathbf{p} in this box, and the bid representation is essentially unique (if we restrict ourselves to normalised bid lists as described above). Moreover, if the valuation v has a discrete simplex domain, as we assume throughout, demand coincides at all prices $\mathbf{p} \in \mathbb{R}^n$. The bids in Figure 1 (right) represent the strong substitutes demand shown in Figure 1 (left). Conversely, however, not all lists of positive and negative bids

⁷While it is unproblematic technically to allow rational bid entries, we follow the existing literature on product-mix auctions in only considering integral bids. We note that allowing rational bid entries with denominators bounded by some constant does not affect the correctness of the algorithms or lower bound construction, as we can conceptually scale price space (and the bounding box) to make all bids integral. Since the size of the bounding box enters the running time of our algorithms logarithmically, this does not affect the complexity.

⁸In particular, we emphasise that if a bid \mathbf{b} has negative weight $w(\mathbf{b}) < 0$, then we *subtract* $|w(\mathbf{b})|$ items of good $i(\mathbf{b})$ from the aggregate bundle.

induce a strong substitutes demand correspondence; we call a bid list *valid* if it does. Theorem 2.3, taken from [1], gives a criterion that allows us to check validity. It is known that the problem of deciding the validity of a bid list is coNP-complete [1].

THEOREM 2.3. *A bid list is valid if and only if the weights of the bids indifferent between i and i' at \mathbf{p} sum to a non-negative number, for all $\mathbf{p} \in \mathbb{R}^n$ and $i, i' \in [n]_0$.*

A special subclass of the strong substitutes demand type is the family of demand correspondences that can be expressed using only positive bids. This family is of particular practical interest, as the Bank of England currently runs the product-mix auction with positive bids only. Moreover, as a single positive bid corresponds to the demand of a unit-demand consumer, learning a list of positive bids is equivalent to learning the demands of a collection of unit-demand consumers; the latter have been studied in a variety of settings such as profit-maximising envy-free pricing [8]. Note that any list of positive bids is valid, as it trivially satisfies Theorem 2.3.

2.3 The Geometry of Bids

In the previous section, we explained the algebraic relationship between a bid list and its resulting demand correspondence. Here we highlight the geometry of such a demand correspondence, as this forms the basis of our algorithms.

Fix a bid $\mathbf{b} \in \mathbb{N}_0^n$ and an item $i \in [n]_0$. We let R_i denote the set of prices at which \mathbf{b} demands good i . Each R_i is an unbounded convex polytope in \mathbb{R}^n , which can be expressed succinctly as $R_i = \mathbf{b} + H_i$, where H_i is the conic hull of the vectors \mathbf{e}^j , $j \in [n]_0 \setminus \{i\}$, if we define $\mathbf{e}^0 = -\mathbf{e}^{[n]}$ (see Figure 2). Note that each R_i is of full affine dimension and the R_i together cover the entirety of \mathbb{R}^n . In line with the previous section, \mathbf{b} is indifferent between two goods $i, j \in [n]_0$ at \mathbf{p} if and only if $\mathbf{p} \in R_i \cap R_j$. Moreover, if a price lies in the interior of any given R_i , good i is the unique good demanded by \mathbf{b} . For a given list of bids \mathcal{B} , we recall that a price \mathbf{p} is non-marginal if each bid in \mathcal{B} demands a unique item at \mathbf{p} . If we define $R_i^j := \mathbf{b}^j + H_i$ for each $\mathbf{b}^j \in \mathcal{B}$, then it is straightforward to see that \mathbf{p} is non-marginal if and only if it does not lie on the boundary of any R_i^j . This allows us to describe the geometry of the demand correspondence $D_{\mathcal{B}}$. Suppose that \mathbf{p} is non-marginal, and that for each $\mathbf{b} \in \mathcal{B}$, $i(\mathbf{b})$ is the unique good demanded by \mathbf{b} at \mathbf{p} . We recall that the unique bundle demanded at \mathbf{p} is obtained by adding $w(\mathbf{b})$ items of good $i(\mathbf{b})$ for each $\mathbf{b} \in \mathcal{B}$. We can express this as

$$D_{\mathcal{B}}(\mathbf{p}) = \sum_{\substack{(i,j): \\ \mathbf{p} \in R_i^j}} w(\mathbf{b}^j) \mathbf{e}^i.$$

This is illustrated in Figure 2. We extend the definition of $D_{\mathcal{B}}$ to marginal prices as above: if \mathbf{p} is marginal, $D_{\mathcal{B}}(\mathbf{p})$ is the discrete convex hull of the bundles demanded at non-marginal prices arbitrarily close to \mathbf{p} .

2.3.1 Valid Bid Lists. As mentioned above, not all demand correspondences $D_{\mathcal{B}}$ arising from bid lists encode a quasi-linear, strong substitutes demand correspondence D_{σ} as defined in Section 2.1. To obtain SS demand, the bid list must satisfy the property given in Theorem 2.3. Here we explore what validity means geometrically. Suppose that \mathcal{B} is a bid list, and for any given price \mathbf{p} , define the (i, ℓ) -support $\text{supp}_{i,\ell}(\mathbf{p})$ of \mathbf{p} to be the set of all bids $\mathbf{b}^j \in \mathcal{B}$ such that $\mathbf{p} \in R_i^j \cap R_{\ell}^j$. In other words, all bids that are indifferent between goods i and ℓ at price \mathbf{p} . This allows us to give a condition for non-validity of our bid list \mathcal{B} . In particular, \mathcal{B} is valid if and only if $\sum_{\mathbf{b} \in \text{supp}_{i,\ell}(\mathbf{p})} w(\mathbf{b}) \geq 0$ for all prices \mathbf{p} and pairs of goods $i, \ell \in [n]_0$.




Fig. 2. Left: The demand regions R_i of a single positive bid b in the case of two goods 1, 2 and the reject good 0. Right: a three-element bid list $\mathcal{B} = \{b^1, b^2, b^3\}$. The price vector p (marked by a red cross) lies in R_1^1, R_0^2 , and R_2^3 . Hence, at prices p , the bid b^1 demands good 1, b^2 is rejected, and b^3 demands good 2. Putting this together, $D_{\mathcal{B}}(p) = (1, 1)$.

2.4 The Computational Challenges

Consider a bidder who has an (unknown) strong substitutes demand correspondence D_v on n goods. We study the problem of learning the unique list \mathcal{B} of positive and negative bids of unit weight that represent a bidder's demand correspondence, i.e. such that $D_v = D_{\mathcal{B}}$. We consider algorithms that learn \mathcal{B} by querying the demand correspondence D_v at different price vectors. More specifically, our algorithms have access to an *adversarial demand oracle* $Q_{\mathcal{B}}$; given any price vector p , $Q_{\mathcal{B}}(p)$ returns a bundle from $D_{\mathcal{B}}(p)$. A bidder may demand multiple bundles at some price (i.e. when $|D_{\mathcal{B}}(p)| > 1$), in which case the adversarial oracle simply returns a single demanded bundle at that price, and we have no control over which such bundle is returned. Another related setting we address in Section 3 is the complexity of learning \mathcal{B} given access to a *valuation oracle*, i.e. given a bundle x , the bidder reports their valuation $v(x)$ for this bundle.

Let $B := |\mathcal{B}|$ be the number of bids we wish to learn. Moreover, let $M := \max_{b \in \mathcal{B}} \|b\|_{\infty}$ be the *magnitude* of the bids w.r.t to the L_{∞} norm and $W := \max_{b \in \mathcal{B}} w(b)$ be the maximum bid weight. Note that for any unknown bid list \mathcal{B} , we can determine the value of M with $O(\log M)$ demand queries, as M corresponds to the smallest value m such that the bidder demands the empty bundle at price vector $p = m\mathbf{e}^{[n]}$, which can be found using binary search. We are interested in the query complexity of learning \mathcal{B} , measured in terms of $n, B, \log M$ and $\log W$. Note that $nB \log M + B \log W$ bits are required to store the bid list \mathcal{B} , under the natural assumption that bid vectors and weights are encoded in binary.

3 SIMULATING $Q_{\mathcal{B}}$ WITH A VALUATION ORACLE

In this section we show that demand oracles are not unreasonably powerful, in the sense that we can use a valuation oracle to simulate a demand oracle with polynomial overhead. Consider the setting where we are given query access to a bidder's valuation function v . We show that a single query to $Q_{\mathcal{B}}$ can be simulated with a polynomial number of queries to a valuation oracle. This result utilises the equivalence of the strong substitutes property and M^h -convexity from the discrete convex analysis literature.

Recall that the utility of bundle x at prices p is given by $u(x; p) = v(x) - p \cdot x$. We define $u_p := u(\cdot; p)$ for convenience. In order to simulate a demand oracle on input p , we wish to compute a bundle $x \in D_v$ that maximises $u_p(x)$. Note that we can compute $u_p(x)$ for any bundle x using a single query to the valuation oracle. In order to compute a maximiser of $u_p(\cdot)$, we draw from the discrete convex analysis literature. Firstly, we see that u_p is M^h -concave. Indeed,

it is well-known that strong substitutes valuations are M^{\natural} -concave [18] and subtracting a linear term preserves this property. Secondly, let \hat{u} be the corresponding M -concave function to $u_{\mathbf{p}}$ as defined in (2). We see that maximising $u_{\mathbf{p}}$ is equivalent to maximising \hat{u} . Moreover, we can compute $\hat{u}(x_0, \mathbf{x})$ using at most one query to the valuation oracle. Thirdly, note that we have $\|\mathbf{x}\|_1 \leq BW$ for any bundle \mathbf{x} that the bidder demands, as every bid $\mathbf{b} \in \mathcal{B}$ contributes at most W items to \mathbf{x} .

Murota [14, Chapter 10] provides multiple algorithms for maximising M -concave functions f with bounded effective domains $\text{dom } f$. The simplest such algorithm, a straightforward steepest descent method, finds a maximiser with $O(n^2 L)$ queries, where $L := \max\{\|\mathbf{x} - \mathbf{y}\|_1 \mid \mathbf{x}, \mathbf{y} \in \text{dom } f\}$. In our setting, we have $L = BW$, yielding a query complexity of $O(n^2 BW)$. This query complexity can be improved to $O(n^3 \log(BW/n))$ by applying the more involved algorithms for maximising M -concave functions given in [17] and [19]. We note that this query complexity is polynomial in n, B and $\log W$.

4 LEARNING POSITIVE-WEIGHTED BIDS

In this section we assume that the bidder's demand correspondence can be expressed by a list of positive bids. This is the case, for instance, when we wish to learn the individual demands of a collection of unit-demand consumers.⁹ Our algorithm learns a list of B positive bids using $O(nB \log M)$ demand queries. This is close to our lower bound of $\Omega(B \log M)$ given in Theorem 6.1 below.

We proceed by repeatedly finding a bid and 'removing' it, thereby reducing the size of the remaining demand correspondence until all bids have been found. Let \mathcal{L} denote the subset of bids from \mathcal{B} that have already been learnt, and let $\mathcal{B}' := \mathcal{B} - \mathcal{L}$ be the list of remaining bids. We can simulate a demand oracle $Q_{\mathcal{B}'}(\mathbf{p})$ for the demand correspondence associated with \mathcal{B}' as follows. At price vector \mathbf{p} , first determine a bundle \mathbf{x} demanded by all bids in \mathcal{B} with a single query $Q_{\mathcal{B}}(\mathbf{p})$, and then subtract from \mathbf{x} a bundle \mathbf{y} demanded at \mathbf{p} by the bids in \mathcal{L} .¹⁰

In this way, the problem of learning a list of positive bids reduces to repeatedly identifying a single bid. In the next section, we describe a subroutine that learns the location of a single bid in \mathcal{B}' using $O(n \log M)$ queries. As this subroutine is called B times, this yields an overall query complexity of $O(nB \log M)$ for learning all bids in \mathcal{B} . Recall that we can compute M with $O(\log M)$ queries.

4.1 Finding a Single Positive Bid

We present an algorithm that performs binary searches using *delta queries* to successively learn the coordinates x_1, \dots, x_n of a bid's location \mathbf{x} together with its weight. We begin by defining delta queries and establishing some fundamental facts about the results returned by these queries.

Definition 4.1. A *delta query* $\Delta(\mathbf{q})$ at $\mathbf{q} \in \mathbb{Z}^n$ consists of two queries \mathbf{q}^+ and \mathbf{q}^- defined by $\mathbf{q}^+ = \mathbf{q}' + \frac{1}{2n}\mathbf{e}^1$ and $\mathbf{q}^- = \mathbf{q}' - \frac{1}{2n}\mathbf{e}^1$, where we let

$$\mathbf{q}' := \mathbf{q} + \sum_{i \in \{2, \dots, n\}} \frac{1}{2(n-i+1)} \mathbf{e}^i.$$

The return value of the delta query is defined as $\Delta(\mathbf{q}) = \mathbf{x}_1^- - \mathbf{x}_1^+$, where \mathbf{x}^+ and \mathbf{x}^- are the bundles of goods uniquely demanded at \mathbf{q}^+ and \mathbf{q}^- .

⁹In this setting we would have access only to an oracle that returns an aggregately demanded bundle at given prices.

¹⁰Note that \mathcal{B}' is valid, as lists of positive bids are always valid. If \mathcal{B} consisted of positive and negative bids, removing a single positive bid might result in a bid list that is no longer valid. In this case, the algorithm described in this section may fail and return points not corresponding to bid locations.

Algorithm 1 Learning Positive Bids

-
- 1: Perform binary search to find largest price $\mathbf{p} \in \{(k, M, \dots, M) \mid k \in [M]_0\}$ at which the demand of good 1 is positive, and fix $x_1 = p_1$.
 - 2: **for** $i = 2 \dots n$ **do**
 - 3: Binary search to find smallest price $\mathbf{p} \in \{(x_1, \dots, x_{i-1}, k, M, \dots, M) \mid k \in [M]_0\}$ at which $\Delta(\mathbf{p}) > 0$, and fix $x_i = p_i$.
 - 4: **return** bid vector $\mathbf{x} = (x_1, \dots, x_n)$ and weight $\Delta(\mathbf{x})$.
-

Note that $q_1^+ = q_1 + \frac{1}{2n}$ and $q_1^- = q_1 - \frac{1}{2n}$, and the two query points \mathbf{q}^+ and \mathbf{q}^- agree on all other coordinates $i \geq 2$. Secondly, \mathbf{q}^\pm is non-marginal by construction, so any bid $\mathbf{b} \in \mathbb{N}_0^n$ uniquely demands some good i at \mathbf{q}^\pm . The intuition behind delta queries is as follows. Consider the hyperplane normal to \mathbf{e}^1 that contains \mathbf{q} . In a first step, we carefully perturb \mathbf{q} such that the resulting point \mathbf{q}' remains on the hyperplane and no bid is indifferent between any two goods in $\{2, \dots, n\}$. The points \mathbf{q}^- and \mathbf{q}^+ are then obtained by perturbing \mathbf{q}' in directions $\pm \mathbf{e}^1$ such that the prices become non-marginal.

Lemma 4.2 makes the observation that bids \mathbf{b} satisfying $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}$ demand good 1 at \mathbf{q}^- and are rejected at \mathbf{q}^+ , while all other bids demand the same good at both prices \mathbf{q}^\pm . Hence demand changes only in terms of good 1, and $\Delta(\mathbf{q})$ captures the magnitude of this change.¹¹ In our current setting where all bids have positive weights, Corollary 4.3 notes that this is equivalent to summing the weights of the bids \mathbf{b} that satisfy $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}$. Our algorithm exploits this fact in order to learn the coordinates of a bid location as well as the bid weight.

LEMMA 4.2. *Suppose we place a delta query at $\mathbf{q} \in \mathbb{Z}^n$. Then any bid $\mathbf{b} \in \mathbb{N}_0^n$ demands different goods at \mathbf{q}^- and \mathbf{q}^+ if and only if $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}'$. Moreover, any such bid demands an item of good 1 at \mathbf{q}^- and an item of the reject good 0 at \mathbf{q}^+ .*

PROOF. Suppose \mathbf{b} is a bid with $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}'$. As \mathbf{b} is integral, we have $b_i < q_i^+ = q_i^-$ for all goods $i \geq 2$, which implies that \mathbf{b} can only demand goods 0 or 1 at \mathbf{q}^- and \mathbf{q}^+ . At \mathbf{q}^- , bid \mathbf{b} uniquely demands good 1, as we have $b_1 - q_1^- = \frac{1}{2n} > 0$. Similarly, we see that \mathbf{b} uniquely demands the reject good 0 at \mathbf{q}^+ , as $b_1 - q_1^+ = -\frac{1}{2n} < 0$.

Conversely, suppose \mathbf{b} demands distinct goods i and j at \mathbf{q}^- and \mathbf{q}^+ , respectively. This implies the two fundamental inequalities (a) $b_i - q_i^- > b_j - q_j^-$ and (b) $b_j - q_j^+ > b_i - q_i^+$. We first show that we must have $i = 1$ and $j = 0$ by excluding all other possibilities. (Recall that prices p_0 and bid values b_0 are 0 for the reject good 0, by definition.) Suppose that $i = 0$ or $i \geq 2$. Then $q_i^- = q_i^+$ by construction of \mathbf{q} , so (a) and (b) imply $b_j - q_j^- < b_j - q_j^+$ in contradiction to $q_j^- \leq q_j^+$, which holds by construction of \mathbf{q} . Next suppose $i = 1$ and $j \geq 2$. Then $q_j^- = q_j^+ = q_j + \frac{1}{2(n-j+1)}$. The inequalities (a) and (b) imply $b_1 - q_1 - \frac{1}{2n} < b_j - q_j^- < b_1 - q_1 + \frac{1}{2n}$. As \mathbf{b} and \mathbf{q} are integral, it follows that $b_j - q_j^-$ must lie within $\frac{1}{2n}$ of an integral point, in contradiction to $q_j^- = q_j + \frac{1}{2(n-j+1)}$.

Finally, we show that a bid \mathbf{b} that demands good $i = 1$ at \mathbf{q}^- and $j = 0$ at \mathbf{q}^+ satisfies $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}'$, again using the integrality of \mathbf{b} and \mathbf{q} . Firstly, (a) and (b) imply $q_1 - \frac{1}{2n} < b_1 < q_1 + \frac{1}{2n}$ and it follows that $b_1 = q_1$. Secondly, we have $b_k < q_k^+$ for all $k \geq 1$, as \mathbf{b} uniquely demands the reject good 0 at \mathbf{q}^+ , which implies $\mathbf{b} \leq \mathbf{q}$ by integrality. \square

COROLLARY 4.3. $\Delta(\mathbf{q})$ is the sum of the weights of all bids $\mathbf{b} \in \mathcal{B}$ satisfying $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}$.

4.1.1 The Algorithm. Algorithm 1 learns the vector \mathbf{x} and weight of a single positive bid with $O(n \log M)$ queries. It determines the value of x_i , $i \in [n]$, by performing a binary search on line segment $L_i := \{(x_1, \dots, x_{i-1}, z, M, \dots, M) \mid$

¹¹Lemma 4.2 still holds if the bid list contains positive and negative bids. Indeed, a generalised version of delta queries is used in our algorithm for learning positive and negative bids in Section 5.

$0 \leq z \leq M\}$, where the values of x_1, \dots, x_{i-1} have already been determined and are fixed. As L_i is well-ordered, we can define the ‘smallest’ and ‘largest’ points on L_i as $\mathbf{s}^i := (x_1, \dots, x_{i-1}, 0, M, \dots, M)$ and $\mathbf{l}^i := (x_1, \dots, x_{i-1}, M, \dots, M)$.

In a first step, Algorithm 1 performs binary search on L_1 in order to find the largest point at which demand for good 1 is positive. Note that at any $\mathbf{p} \in L_1$, no bid demands items of good $i \geq 2$, i.e. every bid demands an item of good 0 or 1 (or is indifferent between the two). Moreover, the function mapping prices \mathbf{p} on L_1 to the demand of good 1 at \mathbf{p} is monotonically decreasing and changes only at integral points, as the bids are integral. As B items of good 1 are demanded at \mathbf{s}^1 , there is a largest price $\mathbf{p}^* \in L_1$ at which demand is positive. Hence, we can find \mathbf{p}^* using binary search on L_1 by querying demand at $O(\log M)$ prices of the form (k, M, \dots, M) with $k \in [M]_0$.

The second kind of binary search uses delta queries to find the smallest point \mathbf{p}^* for each line segment L_i , $i \geq 2$, at which $\Delta(\mathbf{p}^*)$ is positive. Suppose $i \geq 2$. Corollary 4.3 implies that $\Delta(\mathbf{q})$ restricted to the line L_i is monotonically increasing. We show in the proof of Theorem 4.4 that the invariant $\Delta(\mathbf{l}^i) > 0$ holds when we perform binary search on L_i . Moreover, Δ only changes in value at integral points along L_i , so we can perform binary search to find \mathbf{p}^* with $O(\log M)$ delta queries at prices $(x_1, \dots, x_{i-1}, k, M, \dots, M)$, where $k \in [M]_0$.

Theorem 4.4 establishes the correctness and running time of Algorithm 1. The proof proceeds by induction and makes use of Observation 1.

OBSERVATION 1. *Let $i \geq 2$. Suppose Algorithm 1 has successfully determined the first $i - 1$ coordinates x_1, \dots, x_{i-1} and binary search (using delta queries) finds the smallest point $\mathbf{p} = (x_1, \dots, x_i, M, \dots, M)$ on L_i at which $\Delta(\mathbf{p}) > 0$. Corollary 4.3 implies the following.*

- (a) *None of the bids $\mathbf{b} \in \mathcal{B}$ satisfy $\mathbf{b} \leq \mathbf{p}$ and $b_i < p_i$.*
- (b) *There is at least one bid \mathbf{b} with $\mathbf{b} \leq \mathbf{p}$ and $b_i = p_i$.*

THEOREM 4.4. *Algorithm 1 returns the vector and weight of a bid using $O(n \log M)$ queries.*

PROOF. First we show that the algorithm is well-defined. Clearly, the first coordinate x_1 can be found, as outlined above. To show that the binary search on L_i using delta queries to find x_i is well-defined for all $i \geq 2$, it suffices to show that $\Delta(\mathbf{l}^i)$ is positive. Fix $i \leq 2$ and suppose the algorithm has found x_1, \dots, x_{i-1} . If \mathbf{p} is the point found by binary searching on line L_{i-1} , then Observation 1 (b) tells us that there is at least one bid \mathbf{b} satisfying $b_1 = p_1$ and $\mathbf{b} \leq \mathbf{p}$. Moreover, we have $p_1 = \mathbf{l}_1^i$ and $\mathbf{p} \leq \mathbf{l}^i$ by construction. By transitivity and Corollary 4.3, it follows that $\Delta(\mathbf{l}^i) > 0$.

Next, we prove that $\mathbf{x} = (x_1, \dots, x_n)$ returned by the algorithm corresponds to the vector of a bid and that $\Delta(\mathbf{x})$ is its weight. By Observation 1 (b) and construction of x_n , we know that there is a bid \mathbf{b}^* that satisfies $\mathbf{b}^* \leq \mathbf{x}$ and $b_n^* = x_n$. On the other hand, for any $1 \leq i \leq n - 1$, the construction of x_i together with Observation 1 (a) implies that no bid \mathbf{b} satisfies $\mathbf{b} \leq \mathbf{x}$ and $b_i < x_i$. As a result, we get $b_i^* = x_i$ for all $1 \leq i \leq n$ and it follows that \mathbf{b}^* lies at point \mathbf{x} . Moreover, we see that \mathbf{b}^* is the only bid in \mathcal{B} that satisfies $b_1 = x_1$ and $\mathbf{b} \leq \mathbf{x}$. By Corollary 4.3, $\Delta(\mathbf{x})$ is the weight of \mathbf{b}^* .

Finally, to see that the algorithm has query complexity $O(n \log M)$, note that it performs n binary searches along lines L_i , each of which incurring $O(\log M)$ queries. \square

5 LEARNING POSITIVE AND NEGATIVE BIDS

In this section, we provide an algorithm for learning valid bid lists that may contain both positive and negative bids. Introducing negative bids significantly complicates learning the location of bids, as negative bids are able to ‘cancel out’ facets. Consider, for instance, a bid \mathbf{b} of weight 2 at position $(2, 4)$ in the setting with two goods. Without negative bids, its horizontal and vertical facets extend indefinitely from \mathbf{b} in direction \mathbf{e}^1 and \mathbf{e}^2 . In contrast, Figure 1 demonstrates

how two negative bids of weight 1 cancel out \mathbf{b} 's horizontal facet from $(7, 4)$ onwards. Similarly, in the same figure the vertical facet of the bid at $(4, 2)$ is cancelled out by the negative bid at $(4, 4)$. As a result, some bids may only be detectable if we query in their local neighbourhood; for more details we refer to Section 6.2, where we exploit this phenomenon to show super-polynomial lower bound on learning positive and negative bids.

Recall from Section 2.1.1 that the demand correspondence of a SS demand (and hence of a valid bid list) corresponds to a polyhedral complex over price space, the LIP, where the boundaries between unique demand regions are $(n - 1)$ -dimensional facets. Our algorithm learns the collection of all hyperplanes that contain these facets, as well as each vertex arising from the intersection of n such hyperplanes. We note that every bid must lie at a vertex but, conversely, not every such vertex contains a bid. We introduce *bid existence queries* at integral price vectors, which use a set of 2^n demand queries in the local neighbourhood of the vector in order to check for existence of a bid there. In addition, we define (more powerful and costly) *super queries* which for integral points \mathbf{p} , use a collection of demand queries to provide complete information about the demand correspondence in the local neighbourhood of \mathbf{p} . This allows us to simulate a bid existence query with a super query and also to perform a principled search for new hyperplanes: at each iteration of the algorithm, either the local information around two vertices points us in the direction of a new hyperplane, or we have succeeded in learning all hyperplanes, and thus all bids.

5.1 Bid Existence Queries

Our algorithm computes a set of integral candidate price vectors at which bids may lie. In order to check for any such candidate vector $\mathbf{p} \in \mathbb{Z}^n$ whether there is a bid at this location and, if so, what its weight is, we first introduce a generalisation of the delta queries that were first introduced in Section 4.1. This in turn allows us to define the *bid existence query*, which consists of 2^{n-1} generalised delta queries. In total, we see that as each generalised delta query consists of two elementary demand queries, a total of 2^n queries to the demand oracle is sufficient to determine the existence and weight of a bid at a given integral price vector.

5.1.1 Generalised delta queries. Recall that a delta query at $\mathbf{q} \in \mathbb{Z}^n$ works by first perturbing the entries $i \geq 2$ of \mathbf{q} in a specific way. We now generalise the directions in which these perturbations happen. For any point $\mathbf{q} \in \mathbb{Z}^n$ and every $S \subseteq \{2, \dots, n\}$, we define $\mathbf{q}(S)$ by $q(S)_1 := q_1$ and $q(S)_i := q_i - \frac{1}{2(n-i+1)}$ if $i \in S$ and $q(S)_i := q_i + \frac{1}{2(n-i+1)}$ otherwise for all $i \geq 2$. This allows us to define a *generalised delta query* at \mathbf{q} w.r.t. S as follows.

Definition 5.1. Let $\mathbf{q} \in \mathbb{Z}^n$ and $S \subseteq \{2, \dots, n\}$. A (generalised) delta query consists of two queries $\mathbf{q}^+(S) := \mathbf{q}(S) + \frac{1}{2n}\mathbf{e}^1$ and $\mathbf{q}^-(S) := \mathbf{q}(S) - \frac{1}{2n}\mathbf{e}^1$, where we let

$$\mathbf{q}(S) := \mathbf{q} - \sum_{i \in S} \frac{1}{2(n-i+1)} \mathbf{e}^i + \sum_{i \notin S \cup \{1\}} \frac{1}{2(n-i+1)} \mathbf{e}^i.$$

The return value of the delta query is defined as $\Delta(\mathbf{q}; S) = \mathbf{x}_1^- - \mathbf{x}_1^+$, where \mathbf{x}^+ and \mathbf{x}^- are the bundles of goods uniquely demanded at \mathbf{q}^+ and \mathbf{q}^- .

Note that $\mathbf{q}^+(S)_1 = q(S)_1 + \frac{1}{2n}$ and $\mathbf{q}^-(S)_1 = q(S)_1 - \frac{1}{2n}$, and the two query points $\mathbf{q}^+(S)$ and $\mathbf{q}^-(S)$ agree on all other coordinates $i \geq 2$. Secondly, \mathbf{q}^\pm are non-marginal prices by construction, so every bid uniquely demands some good i .

Lemma 5.2 makes the observation that bids satisfying $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}(S)$ demand good 1 at \mathbf{q}^- and are rejected at \mathbf{q}^+ , while all other bids demand the same good at both prices \mathbf{q}^\pm . Hence demand changes only in terms of good 1, and $\Delta(\mathbf{q}; S)$ captures this change. Another useful interpretation of $\Delta(\mathbf{q}; S)$ is that it returns the sum of the weights of all bids that satisfy $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}$.

LEMMA 5.2. Suppose we make a delta query at $\mathbf{q}(S) \in \mathbb{Z}^n$ w.r.t. $S \subseteq \{2, \dots, n\}$. Then any bid $\mathbf{b} \in \mathbb{Z}^n$ demands different goods at \mathbf{q}^- and \mathbf{q}^+ if and only if $b_1 = q_1$ and $\mathbf{b} \leq \mathbf{q}(S)$. Moreover, any such bid demands an item of good 1 at \mathbf{q}^- and an item of reject good 0 at \mathbf{q}^+ .

PROOF. Analogous to the proof of Lemma 4.2. \square

5.1.2 *Defining Bid Existence Queries.* In order to check for the existence (and weight) of a bid at a given price vector $\mathbf{p} \in \mathbb{Z}^n$, we introduce the existence query consisting of 2^{n-1} generalised delta queries. Indeed, an existence query performed at $\mathbf{p} \in \mathbb{Z}^n$ consists of the set of delta queries $\{\Delta(\mathbf{p}; S) \mid S \subseteq \{2, \dots, n\}\}$. We now show how the existence and weight of a bid at \mathbf{p} can be inferred from the return values of these 2^{n-1} delta queries. For convenience, we introduce the notation $w(C) = \sum_{\mathbf{b} \in C} w(\mathbf{b})$ for any list of bids C . Our approach is based on the following straightforward observation.

OBSERVATION 2. For any bid list \mathcal{B} , define the two sub-lists $\mathcal{B}' = \{\mathbf{b} \in \mathcal{B} \mid b_1 = p_1, \mathbf{b} \leq \mathbf{p}\}$ and $\mathcal{B}'' = \{\mathbf{b} \in \mathcal{B} \mid b_1 = p_1, \mathbf{b} \leq \mathbf{p}, \mathbf{b} \neq \mathbf{p}\}$. There exists a bid at \mathbf{p} in \mathcal{B} if and only if the term $w := w(\mathcal{B}') - w(\mathcal{B}'')$ is non-zero. Moreover, if the bid exists, then its weight is w .

Note that \mathcal{B}' contains all bids dominated by \mathbf{p} and we can compute $w(\mathcal{B}')$ with a single delta query $\Delta(\mathbf{p}, \emptyset)$. Moreover, the set \mathcal{B}'' contains all bids that are dominated by \mathbf{p} and do not lie at \mathbf{p} . In order to determine $w(\mathcal{B}'')$, we first partition \mathcal{B}'' into $2^{n-1} - 1$ sub-lists as follows. For each non-empty set of goods $T \subseteq \{2, \dots, n\}$, we let $\mathcal{B}_T := \{\mathbf{b} \in \mathcal{B} \mid b_i < p_i, \forall i \in T \text{ and } b_j = p_j, \forall j \notin T\}$. Lemma 5.3 now provides us with a method of computing $w(\mathcal{B}'')$ using the $2^{n-1} - 1$ delta queries $\Delta(\mathbf{p}, S)$ with $\emptyset \subsetneq S \subseteq \{2, \dots, n\}$.

LEMMA 5.3. We have

$$w(\mathcal{B}'') = \sum_{\emptyset \subsetneq T \subseteq \{2, \dots, n\}} w(\mathcal{B}_T) = \sum_{\emptyset \subsetneq S \subseteq \{2, \dots, n\}} (-1)^{|S|+1} \Delta(\mathbf{p}; S).$$

PROOF. Firstly, observe how the result of the delta query $\Delta(\mathbf{p}; S)$ for a given $S \subseteq \{2, \dots, n\}$ satisfies $\Delta(\mathbf{p}; S) = \sum_{T \supseteq S} w(\mathcal{B}_T)$. Indeed, recall that $\Delta(S)$ sums the weights of the bids $\mathbf{b} \in \mathcal{B}$ satisfying $b_1 = p_1$ and $\mathbf{b} \leq \mathbf{p}(S)$. As we have $\mathbf{p}(S)_i < p_i$ for all $i \in S$ and the bids are integral, the set of bids that $\Delta(\mathbf{p}; S)$ considers is identical to $\bigcup_{T \supseteq S} \mathcal{B}_T$. Hence

$$\begin{aligned} \sum_{\emptyset \subsetneq S \subseteq \{2, \dots, n\}} (-1)^{|S|+1} \Delta(\mathbf{p}; S) &= \sum_{\emptyset \subsetneq S \subseteq \{2, \dots, n\}} (-1)^{|S|+1} \sum_{T \supseteq S} w(\mathcal{B}_T) \\ &= \sum_{\emptyset \subsetneq T \subseteq \{2, \dots, n\}} w(\mathcal{B}_T) \sum_{S \subseteq T} (-1)^{|S|+1} \\ &= \sum_{\emptyset \subsetneq T \subseteq \{2, \dots, n\}} w(\mathcal{B}_T). \end{aligned}$$

The last equality follows from the binomial identity in Lemma 5.4 below. \square

LEMMA 5.4. For any $n \geq 1$ we have $\sum_{i=1}^n (-1)^{i+1} \binom{n}{i} = 1$.

PROOF. Recalling Pascal's Rule, $\binom{n}{i} = \binom{n-1}{i} + \binom{n-1}{i-1}$, we have

$$\begin{aligned} \sum_{i=1}^n (-1)^{i+1} \binom{n}{i} &= \sum_{i=1}^n (-1)^{i+1} \binom{n-1}{i} + \sum_{i=1}^n (-1)^{i+1} \binom{n-1}{i-1} \\ &= \sum_{i=1}^{n-1} (-1)^{i+1} \binom{n-1}{i} + \sum_{i=1}^{n-1} (-1)^i \binom{n-1}{i} + (-1)^2 \binom{n-1}{0} = 1. \end{aligned}$$

\square

5.2 Super Queries

Suppose $\mathbf{p} \in \mathbb{Z}^n$ is an integral price vector. We show that it is possible to obtain complete knowledge of $D_{\mathcal{B}}(\mathbf{p}')$ for all prices \mathbf{p}' with $\|\mathbf{p} - \mathbf{p}'\|_{\infty} < 1$ using a *super query*, which consists of a specific set of demand queries at non-marginal query points in the vicinity of \mathbf{p} . Intuitively, this works because bid vectors are integral and facets of an SS LIP can only have specific orientations. Super queries are used by our algorithm in two ways: firstly to determine the existence and weight of a bid at a given integral point \mathbf{p} , and secondly to provide information that leads to a new separating hyperplane.

Let $U_1(\mathbf{p}), \dots, U_{2^n}(\mathbf{p})$ denote the 2^n orthants of the unit L_{∞} -ball around \mathbf{p} . Each orthant is a hypercube that can be triangulated into $n!$ simplices (one for each permutation of the coordinates $[n]$), as described in Section 2. We denote these simplices for the i -th orthant by $U_i^1(\mathbf{p}), \dots, U_i^{n!}(\mathbf{p})$.

Definition 5.5. A *super query* at $\mathbf{p} \in \mathbb{Z}^n$ is a collection $SQ(\mathbf{p})$ of representative prices from the interior of each $U_i^j(\mathbf{p})$, where $i \in [2^n]$ and $j \in [n!]$.

With a slight abuse of notation, we say that we ‘super query’ a price vector \mathbf{p} if we query all price vectors in $SQ(\mathbf{p})$. Figure 3 illustrates super queries in the case of two goods ($n = 2$). The following lemma demonstrates the use of super queries.

LEMMA 5.6. *Querying the points in $SQ(\mathbf{p})$ once is sufficient to ascertain $D_{\mathcal{B}}(\mathbf{p}')$ for any \mathbf{p}' with $\|\mathbf{p} - \mathbf{p}'\|_{\infty} < 1$. With this information, we can learn all facets of the LIP containing \mathbf{p} and establish the existence and weight of a bid at \mathbf{p} .*

PROOF. Suppose that \mathbf{p}' is such that $\|\mathbf{p} - \mathbf{p}'\|_{\infty} < 1$. From construction, it follows that $\mathbf{p}' \in U_i^j(\mathbf{p})$ for some $U_i^j(\mathbf{p})$. However, we also know that facets in the polyhedral complex resulting from $D_{\mathcal{B}}$ can only be normal to a unit vector \mathbf{e}^i , for some $i \in [n]$, or a vector of the form $\mathbf{e}^i - \mathbf{e}^j$, for some $i, j \in [n]$. As a consequence, any price in the interior of $U_i^j(\mathbf{p})$ demands the same unique bundle. If \mathbf{p}' also lies in the interior of $U_i^j(\mathbf{p})$, then we know the bundle it demands and we are done. On the other hand, if \mathbf{p}' is on the boundary $U_i^j(\mathbf{p})$, then it could be marginal, but if this is the case, from construction, we know all bundles demanded at non-marginal prices neighbouring \mathbf{p}' , as they must lie in the interior of other $U_{i'}^{j'}(\mathbf{p})$. This in turn implies that we can infer all bundles demanded at \mathbf{p}' as desired. Moreover, it follows immediately that $SQ(\mathbf{p})$ holds all information regarding facets containing \mathbf{p} , as facets can only take specific orientations in a SS demand correspondence. Finally, as we have complete information about demand in the unit ball around \mathbf{p} , we can simulate a bid existence query as defined in Section 5.1 at \mathbf{p} without making any further queries to the demand oracle. This tells us whether there exists a bid at \mathbf{p} and, if so, what its weight is. \square

5.3 Finding a Separating Hyperplane

Suppose $\mathbf{0} \leq \mathbf{q}, \mathbf{q}' \leq (M+1)\mathbf{e}^{[n]}$ are distinct price vectors that lie in the interiors of different demand regions. Note that $\|\mathbf{q} - \mathbf{q}'\|_{\infty} \leq M+1$. As demand regions are convex and have piecewise-linear boundaries, there exists some facet F of the LIP separating \mathbf{q} and \mathbf{q}' . In order to find the hyperplane containing F , we first perform $O(\log M)$ steps of binary search on $\text{conv}(\mathbf{q}, \mathbf{q}')$ to obtain a pair of points \mathbf{s} and \mathbf{s}' on either side of F with $\|\mathbf{s} - \mathbf{s}'\|_{\infty} \leq 1/4$. By the geometry of SS demand, we know that F must be normal to a vector $\mathbf{e}^i - \mathbf{e}^j$, where $i, j \in [n]_0$ and we denote $\mathbf{e}^0 = \mathbf{0}$ for convenience.

Suppose we know that F is normal to $\mathbf{e}^i - \mathbf{e}^j$ for some fixed i and j . Then we can determine the point $\mathbf{p} = \lambda \mathbf{s} + (1-\lambda)\mathbf{s}'$ ($0 \leq \lambda \leq 1$) at which the line segment $\overline{\mathbf{s}\mathbf{s}'}$ intersects F , as this is obtained for the unique value $0 \leq \lambda \leq 1$ such that \mathbf{p}

Algorithm 2 Finding a Separating Hyperplane

Input: Two hyperplane witnesses \mathbf{q} and \mathbf{q}' that are separated by a hyperplane.

- 1: Binary search between \mathbf{q} and \mathbf{q}' to find a pair of points \mathbf{s}, \mathbf{s}' with distance $\|\mathbf{s} - \mathbf{s}'\|_\infty < 1/4$ and distinct demands.
- 2: **for all** $(i, j) \in [n]_0^2$ **do**
- 3: Suppose \mathbf{s} and \mathbf{s}' are separated by a facet F normal to $\mathbf{e}^i - \mathbf{e}^j$.
- 4: Find the unique intersection point \mathbf{p} of F with the line segment from \mathbf{s} to \mathbf{s}' .
- 5: Super query at $\lceil \mathbf{p} \rceil$.
- 6: **if** the super query learns a facet containing $\lceil \mathbf{p} \rceil$ **then**
- 7: **return** the hyperplane containing this facet.

satisfies $p_i - p_j \in \mathbb{Z}$. Next, we round up \mathbf{p} to the nearest integer and claim that $\lceil \mathbf{p} \rceil$ also lies in F . Indeed, we see that if a bid \mathbf{b} demands good i at \mathbf{p} , then it also demands i at $\lceil \mathbf{p} \rceil$:

$$b_j - \lceil p_j \rceil - 1 \leq b_j - p_j - 1 \leq b_i - p_i - 1 < b_i - \lceil p_i \rceil.$$

Here the second inequality follows from the fact that \mathbf{b} demands good i at \mathbf{p} . As \mathbf{b} and $\lceil \mathbf{p} \rceil$ are integral, this chain of inequalities implies $b_i - \lceil p_i \rceil \geq b_j - \lceil p_j \rceil$. Hence, as this holds for every bid, we see that \mathbf{p} and $\lceil \mathbf{p} \rceil$ lie in the same facet. We perform a super query at $\lceil \mathbf{p} \rceil$ in order to find F (as well any other adjacent facets that \mathbf{p} might lie in).

This suggests the following algorithm for finding a separating hyperplane, specified in Algorithm 2. First we perform binary search between \mathbf{q} and \mathbf{q}' in order to find points \mathbf{s} and \mathbf{s}' that are $1/4$ apart (w.r.t. the L_∞ -norm). Then, for every possible pair of goods $(i, j) \in [n]_0^2$, we suppose that \mathbf{s} and \mathbf{s}' are separated by a facet that is normal to $\mathbf{e}^i - \mathbf{e}^j$ and perform a super query at prices $\lceil \mathbf{p} \rceil$ obtained as described above. As soon as such a super query finds a facet containing $\lceil \mathbf{p} \rceil$, we return the hyperplane containing this facet. Overall, finding a separating hyperplane between \mathbf{q} and \mathbf{q}' in this way costs $O(\log M + n^2 2^n n!)$ queries, as each super query costs $O(2^n n!)$. Moreover, our procedure is guaranteed to find a separating hyperplane, as it exhaustively checks all possible orientations for the separating facet F between \mathbf{s} and \mathbf{s}' .

5.4 The Main Algorithm

Algorithm 3 learns bid lists that may comprise positive and negative bids. The algorithm maintains a set of hyperplanes \mathcal{H} that it has learnt. We initialise \mathcal{H} with the axis-aligned hyperplanes $\mathbf{e}^i = 0$ and $\mathbf{e}^i = M$ for all $i \in [n]$. The algorithm also keeps track of the corresponding set of vertices \mathcal{V} arising from intersections of hyperplanes in \mathcal{H} , the set of query points $\mathcal{Q} = \bigcup_{\mathbf{v} \in \mathcal{V}} \text{SQ}(\mathbf{v})$, as well as the set of polytopes \mathcal{P} of the subdivision of \mathbb{R}^n by the hyperplanes in \mathcal{H} . Finally, $\hat{\mathcal{B}}$ denotes the set of bids that the algorithm has learnt. Two query points $\mathbf{q}, \mathbf{q}' \in \mathcal{Q}$ are *hyperplane witnesses* if they lie in the same polytope $P \in \mathcal{P}$ but have different demand. Figure 3 shows a snapshot of Algorithm 3 after learning a single hyperplane.

We now argue that Algorithm 3 is well-defined and learns a bid list in $O(Bn^2)$ iterations by identifying all the hyperplanes containing a facet of the LIP. As each bid gives rise to at most $O(n^2)$ facets, the total number of such hyperplanes is $O(Bn^2)$. The algorithm learns a new hyperplane in each iteration, as hyperplane witnesses lie in the same polytope by definition, which implies that Step 5 of the algorithm finds a new hyperplane that is not in \mathcal{H} . Moreover, every bid must lie at the intersection of n such hyperplanes, and we perform a super query at each intersection to check for the existence of a bid at that point. All that remains is to show that Algorithm 3 does not terminate until all $O(Bn^2)$ hyperplanes containing facets of the LIP have been learnt, as this immediately implies that the algorithm identifies the locations and weights of all bids.




Fig. 3. A snapshot of Algorithm 3 learning the bid list consisting of one negative bid of weight -1 at $(3, 3)$ and three positive bids of weight 1 at $(1, 3)$, $(3, 1)$ and $(5, 5)$. The initial axis-aligned hyperplanes are drawn in orange. Moreover, the algorithm has learnt one additional vertical hyperplane ($x = 1$) drawn in purple. The super queries made at the hyperplane intersections are represented by blue query points. Two super queries are highlighted to illustrate how the existence of a bid is determined.

LEMMA 5.7. *Algorithm 3 learns all $O(Bn^2)$ hyperplanes containing a facet of the LIP.*

PROOF. It suffices to show that there always exists a pair of hyperplane witnesses if we have not learnt all hyperplanes. Suppose F is a facet of the LIP that is not contained in any hyperplane in \mathcal{H} . Then F separates two neighbouring demand regions. Moreover, by assumption there is a polytope $P \in \mathcal{P}$ that intersects both these regions, hence there exist two non-marginal points $p, p' \in P$ at which demand differs.

Next, recall that we perform a super query at every vertex of the subdivision of \mathbb{R}^n by the hyperplanes in \mathcal{H} . Hence, for every polytope $P \in \mathcal{P}$ we have a query point close to its vertices. Suppose all these query points in P demand the same bundle x . This implies that x is also demanded at all vertices of P . By convexity of demand, x is then uniquely demanded at any non-marginal point of the polytope. But this contradicts the fact that demand differs at p and p' . It follows that at least two of the query points close the vertices of P have distinct demand. \square

Theorem 5.8 gives the overall query complexity of learning a bid list with Algorithm 3.

Algorithm 3 Learning Positive and Negative Bids

Initialisation:

- 1: Let \mathcal{H} be hyperplanes of the form $e^i \cdot x = 0$ and $e^i \cdot x = M$ for all $i \in [n]$.
- 2: Update \mathcal{V}, \mathcal{Q} and \mathcal{P} according to \mathcal{H} .
- 3: Super query at each $v \in \mathcal{V}$ to check for bid at v and add each new-found bid to $\hat{\mathcal{B}}$.

Main loop:

- 4: **while** \exists hyperplane witnesses $q, q' \in \mathcal{Q}$ **do**
 - 5: Use Algorithm 2 to find a hyperplane separating q and q' and add it to \mathcal{H} .
 - 6: Update \mathcal{V}, \mathcal{Q} and \mathcal{P} accordingly.
 - 7: **for all** new vertices v in \mathcal{V} **do**
 - 8: Super query at v to check for bid at v and add a newly found bid to $\hat{\mathcal{B}}$.
 - 9: **return** $\hat{\mathcal{B}}$
-

THEOREM 5.8. *Algorithm 3 requires $O\left(Bn^2 \log M + 2^n n! \binom{Bn^2}{n}\right)$ queries to learn a bid list that may consist of positive and negative bids. For n constant, this is $O(B \log M + B^n)$.*

PROOF. Let H, V and Q be the number of hyperplanes, vertices and query points in the sets \mathcal{H}, \mathcal{V} and \mathcal{Q} at the conclusion of Algorithm 3. Note that Algorithm 3 only makes demand queries in Steps 3, 5 and 8. We first count the number of queries that arise from checking for existence of a bid at vertex locations (Steps 3 and 8). \mathcal{H} is initialised with $2n$ axis-aligned hyperplanes and learns an additional $O(Bn^2)$ hyperplanes (cf. Lemma 5.7). Hence $H = O(Bn^2)$. As every vertex in \mathcal{V} arises as the intersection of n distinct hyperplanes, we have $V = O\left(\binom{Bn^2}{n}\right)$. The algorithm performs a super query of cost $2^n n!$ at each vertex in \mathcal{V} , leading to a total cost of $O\left(2^n n! \binom{Bn^2}{n}\right)$. Next we count the number of queries required to find new hyperplanes (Step 5). Section 5.3 tells us that finding a single hyperplane costs $O(\log M + n^2 2^n n!)$ queries. As the algorithm learns $O(Bn^2)$ hyperplanes, the aggregate number of queries performed by Step 5 is $O(Bn^2 (\log M + n^2 2^n n!))$. Summing these two cost terms for super queries and hyperplane finding gives the desired query complexity for Algorithm 3. \square

6 LOWER BOUNDS

We describe lower bounds for the complexity of learning bid lists. First we show that learning B positive bids requires $\Omega(B \log M)$ queries. Our other lower bounds apply in the setting where the bid list may comprise positive and negative bids and make use of a carefully constructed ‘island gadget’ that consists of 2^n positive and 2^n negative bids of unit weight. For bid lists \mathcal{B} that need not be valid, the island gadget immediately implies that $\Omega\left(\left(\frac{M}{4}\right)^n\right)$ queries are required to learn \mathcal{B} (roughly, price-space has to be queried exhaustively). For the case with valid bid lists, we construct an adversarial game to obtain a lower bound of $\Omega\left(\left((B - 2^{n+1})/8n^2\right)^n\right)$ on the query complexity. We see that n must be held constant for the query complexity to be polynomial. In this regime, our lower bounds and the upper bounds from Section 5 imply a query complexity of $\Theta(B \log M + B^n)$ for constant n .

THEOREM 6.1. *Any algorithm for learning bid lists requires $\Omega(B \log M)$ queries.*

PROOF. Suppose an adversary places B positive bids at B integral points in the space $\{(x_1, 0, \dots, 0) \mid x_1 \in [M]_0\}$. By a standard decision tree argument, any algorithm must make $\Omega(B \log M)$ queries to learn the location of the bids. \square

6.1 The Island Gadget

We now introduce the *island gadget*, which allows us to locally change the demand correspondence without affecting demand outside the convex hull of the gadget bids. This is illustrated in Figure 4. Lemma 6.3 establishes that the gadget only influences demand locally. Let $\rho(\mathbf{x})$ denote the weight function that assigns positive weight 1 to bids with an even number of odd entries and negative weight -1 otherwise.

Definition 6.2. The gadget G at position $\mathbf{0}$ consists of the following 2^{n+1} positive and negative unit bids that sit on the vertices of two unit hyper-cubes. The bids on the first hyper-cube lie at $\mathbf{b} \in \{0, 1\}^n$ and have weight $\rho(\mathbf{b})$. The bids on the second hyper-cube lie at $\mathbf{b} \in \{2, 3\}^n$ and have weight $-\rho(\mathbf{b})$. In order to place G at position \mathbf{x} , we add \mathbf{x} to the position of each bid (without changing the weights).

LEMMA 6.3. *The bids of gadget G placed at position \mathbf{x} demand nothing at prices $\mathbf{p} \notin \mathbf{x} + [3]_0^n$.*

The proof of Lemma 6.3 makes use of the following technical result.

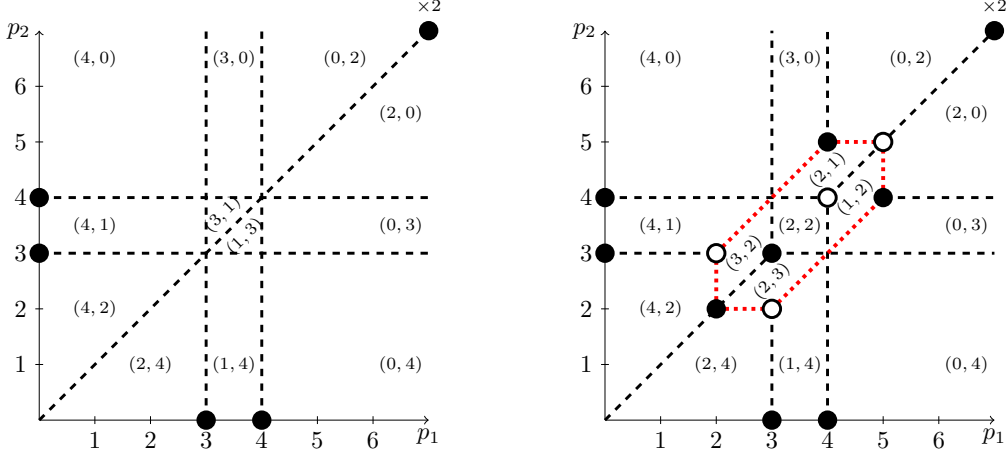


Fig. 4. A demand correspondence with (right) and without (left) an *island gadget* placed at $(2, 2)$. The gadget consists of the eight bids contained in the convex region denoted by the red dotted lines. Note that demand differs only at prices *inside* inside this region.

LEMMA 6.4. Let $\mathbf{c} \in \mathbb{R}^n$ be a constraint vector. If $c_i \geq 1$ for some $i \in [n]$, then the number of vectors with an even number of 1 entries in $\{\mathbf{x} \in \{0, 1\}^n \mid \mathbf{x} \leq \mathbf{c}\}$ is equal to the number of vectors with an odd number of 1 entries.

PROOF. Induction on n . The case $n = 1$ is immediate. For the inductive step, suppose, without loss of generality, that $c_1 \geq 1$. By the induction hypothesis, the set $X^{n-1} := \{\mathbf{x} \in \{0, 1\}^{n-1} \mid x_i \leq c_i, \forall i \in [n-1]\}$ has an equal number of vectors with even and odd numbers of 1 entries. If $c_n = 0$, then the set $X^n := \{\mathbf{x} \in \{0, 1\}^n \mid \mathbf{x} \leq \mathbf{c}\}$ can be constructed by appending an n -th entry 0 to each vector in X^{n-1} , so the number of entries with even and odd 1 entries remains unchanged. If $c_n \geq 1$, X^n consists of the set of vectors obtained by appending n -th entries 0 or 1 to each vector in X^{n-1} , so the number of odd and even vectors doubles. In both cases, the claim holds. \square

PROOF OF LEMMA 6.3. Recall that gadget G at position \mathbf{x} consists of the bids $\mathbf{b} \in \mathbf{x} + \{0, 1\}^n$ of weight $\rho(\mathbf{b} - \mathbf{x})$ and the bids $\mathbf{b} \in \mathbf{x} + \{2, 3\}^n$ of weight $-\rho(\mathbf{b} - \mathbf{x})$. In order to prove that the bids of G do not aggregately demand items of any goods at prices $\mathbf{p} \notin \mathbf{x} + [3]_0^n$, we show for every good i and every such price that the number of negative and positive gadget bids demanding i at \mathbf{p} is the same.

Fix a good $i \in [n]$ and prices $\mathbf{p} \notin \mathbf{x} + [0, 3]^n$. Without loss of generality, we assume that \mathbf{p} is non-marginal and $p_i \neq p_j$ for all $i, j \in [n]$. Recall that an integral bid \mathbf{b} uniquely demands good i at \mathbf{p} if and only if we have $b_i > p_i$ and $b_i - p_i > b_j - p_j$ for all $j \in [n] \setminus \{i\}$. The last condition can be rewritten as $b_i - b_j > p_i - p_j$. Hence we can express the set of bids in the bottom and top cube that (uniquely) demand a positive or negative item of good i as

$$\{\mathbf{b} \in \mathbf{x} + \{0, 1\}^n \mid b_i > p_i \text{ and } b_i - b_j > p_i - p_j, \forall j \in [n] \setminus \{i\}\} \quad (3)$$

$$\text{and } \{\mathbf{b} \in \mathbf{x} + \{2, 3\}^n \mid b_i > p_i \text{ and } b_i - b_j > p_i - p_j, \forall j \in [n] \setminus \{i\}\}, \quad (4)$$

respectively. We show by case distinction on the possible values of p_i that the number of positive and negative bids in (3) and (4) is equal.

Case I: Suppose first that $p_i \notin \mathbf{x} + [3]_0^n$. If $p_i > x_i + 3$, the two sets (3) and (4) are empty and we are done. Hence assume that $p_i < x_i$. Then we can express (3) and (4) as

$$\mathbf{x} + \{\mathbf{v} \in \{0, 1\}^n \mid v_i - v_j > p_i - p_j, \forall j \in [n] \setminus \{i\}\}$$

$$\text{and } \mathbf{x} + 2\mathbf{e}^{[n]} + \{\mathbf{v} \in \{0, 1\}^n \mid v_i - v_j > p_i - p_j, \forall j \in [n] \setminus \{i\}\}.$$

We see that there is a one-to-one correspondence $\mathbf{b} \leftrightarrow \mathbf{b} + 2\mathbf{e}^{[n]}$ between bids in (3) and (4) that preserves the number of odd entries in the bid vectors. A bid \mathbf{b} in the bottom cube is positive if and only if $\mathbf{b} - \mathbf{x}$ has an even number of odd entries and the reverse is true for the top cube. Hence (3) and (4) contain the same number of positive and negative bids.

Case II: Suppose that $p_i \in [x_i, x_i + 3]$. Then there exists a good $k \in [n] \setminus \{i\}$ with $p_k < x_k$ or $p_k > x_k + 3$. We proceed by further case distinction. Note that $b_i - b_j \in \{-1, 0, 1\}$ for all gadget bids \mathbf{b} .

Case II.1: Suppose $p_i - p_j \geq 1$ for some $j \in [n] \setminus \{i\}$. Then the sets (3) and (4) are empty, as the constraint $b_i - b_j > p_i - p_j$ is violated.

Case II.2: Now suppose $p_i - p_j < 1$ for every $j \in [n] \setminus \{i\}$ and we have a non-empty set K of indices k for which $0 < p_i - p_k < 1$. Then for all bids \mathbf{b} in (3) and (4) we have $b_i - b_k = 1$ for all $k \in K$. Hence b_i and b_k can only take one value. If $K = [n] \setminus \{i\}$, this implies that (3) and (4) both contain a single bid of odd parity and we are done. If $K \subset [n] \setminus \{i\}$, then for $j \notin K$ we have $p_i - p_j < 0$ and the constraint $b_i - b_j > p_i - p_j$ holds vacuously, so b_j can take any value. We apply Lemma 6.4 to see that (3) and (4) each contain an equal number of positive and negative bids.

Case II.3: Suppose $p_j > p_i$ for all $j \in [n] \setminus \{i\}$. Then there exists a good $k \in [n] \setminus \{i\}$ such that $p_k > x + 3$. We consider the two subsets $\{b \in (3) \mid b_i = x_i\}$ and $\{b \in (3) \mid b_i = x_i + 1\}$ of (3) separately. If \mathbf{b} is a bid in the first subset, then $b_i = x_i$ and $b_i > p_i$ imply that $p_i - p_k < -1$, so there is no constraint on the value of b_k . If \mathbf{b} is a bid in the second subset, then there is no constraint on the value of b_k . Hence in both cases each subset is either empty or contains an equal number of positive and negative bids by Lemma 6.4. \square

6.2 A Lower Bound for Positive and Negative Bids

From now on, we assume that the unknown bid list may comprise positive and negative bids. We primarily consider the setting where the bid list is assumed to be valid and specify an adversarial game where the adversary can force the player to make $\Omega\left(\left(\frac{B-2^{n+1}}{8n^2}\right)^n\right)$ queries. Note that this yields a query complexity of $\Omega(B^n)$ if n is constant. A similar but simpler adversarial game is then applied to obtain a lower bound for the setting where the bid list may be invalid. The resulting lower bounds for bid lists comprising positive and negative bids are stated in Theorem 6.8.

Fix a parameter $k \in \mathbb{N}$ and let $M = 4k$. The adversary positions the island gadget consisting of 2^{n+1} bids at exactly one of k^n possible points \mathbf{x} of the lattice $4[k-1]_0^n$. He also places *boundary bids* of weight 1 on the boundary of the M -cube as follows. For every good $i \in [n]$, there are M positive bids at all points $m\mathbf{e}^i$ with $m \in [M-1]_0$, and for every pair of goods $(i, j) \in [n]^2$ with $i \neq j$, there are M positive bids at all points $m\mathbf{e}^i + M\mathbf{e}^j$ with $m \in [M-1]_0$. Hence, in total, the adversary creates a bid list with $B = 2^{n+1} + 4k\left(n + \binom{n}{2}\right)$ positive and negative bids. The player wishes to identify where the adversary placed the gadget. Lemma 6.3 shows that placing the gadget at \mathbf{x} only influences demand at prices inside the cube $\mathbf{x} + [3]_0^n$. Hence the player must make queries inside at least $k^n - 1$ cubes to determine where the gadget was placed. Lemma 6.5 shows that the bid list created by the adversary is valid. This leads to the lower bounds stated in Theorem 6.8.

LEMMA 6.5. *The bids placed by the adversary are valid.*

Recall that \mathcal{B}^G denotes the bids of the island gadget. In order to prove Lemma 6.5, we define the subset \mathcal{B}_i^G of \mathcal{B}^G that contains all bids indifferent between a good i and the reject good at some fixed price vector \mathbf{p} . The subset of bids indifferent between two non-reject goods i and j is defined similarly, albeit separately for the ‘lower’ and the ‘upper’ cube of the gadget. Specifically, we define

$$\mathcal{B}_i^G := \{\mathbf{b} \in \{0, 1\}^n \cup \{2, 3\}^n \mid b_i = p_i \text{ and } b_k \leq p_k, \forall k \in [n]\} \quad (5)$$

as well as

$$\mathcal{B}_{ij}^L := \{\mathbf{b} \in \{0, 1\}^n \mid b_i \geq p_i, b_j = b_i - p_i + p_j \text{ and } b_k \leq b_i - p_i + p_k \forall k \notin \{i, j\}\} \quad (6)$$

$$\mathcal{B}_{ij}^U := \{\mathbf{b} \in \{2, 3\}^n \mid b_i \geq p_i, b_j = b_i - p_i + p_j \text{ and } b_k \leq b_i - p_i + p_k \forall k \notin \{i, j\}\}. \quad (7)$$

LEMMA 6.6. *Let $n \geq 3$ and $\mathbf{p} \in \{0, 1\}^n \cup \{2, 3\}^n$. The sum of the weights of the gadget bids indifferent between i and the reject good 0 \mathbf{p} is $-1, 0$ or 1 .*

PROOF. Suppose first that $\mathbf{p} \in \{0, 1\}^n$. We show that the sum of the weights of bids in \mathcal{B}_i^G is 0 or 1. Recalling the definition of \mathcal{B}_i^G from (5), note that if $p_k = 0$ for all $k \neq i$, then $\mathbf{b} = \mathbf{0}$ is the only bid in this set and its weight is positive. If we have $p_k = 1$ for some $k \neq i$, then the set contains an equal number of positive and negative bids by Lemma 6.4. Similarly, we can show that the sum of the weights of bids in \mathcal{B}_i^G is -1 or 0 when $\mathbf{p} \in \{2, 3\}^n$. \square

LEMMA 6.7. *Let $n \geq 3$ and $\mathbf{p} \in \{0, 1\}^n \cup \{2, 3\}^n$. If \mathcal{B}_{ij}^L has more than one bid, the sum of the weights of the bids indifferent between i and j is 0 or 1. Similarly, if \mathcal{B}_{ij}^U has more than one bid, the sum of the respective bid weights is -1 or 0 .*

PROOF. We prove the claim for \mathcal{B}_{ij}^L ; the proof for \mathcal{B}_{ij}^U is identical. Note that in order for \mathcal{B}_{ij}^L to contain any bids, we must have $\mathbf{p} \in \{0, 1\}^n$. Suppose $p_i = 1$. Then for any bid $\mathbf{b} \in \mathcal{B}_{ij}^L$, the values of b_i and b_j are constrained to $b_i = 1$ and $b_j = p_j$. Moreover, the values of b_k are constrained by $b_k \leq p_k$. Hence for \mathcal{B}_{ij}^L to contain at least two bids, we must have $p_k = 1$ for some $k \notin \{i, j\}$, allowing b_k to take values 0 and 1. Lemma 6.4 implies that \mathcal{B}_{ij}^L contains an equal number of positive and negative bids. Now suppose $p_i = 0$. Then b_i can take value 0 or 1 and we consider the two subsets of \mathcal{B}_{ij}^L where b_i takes each value separately. Let $S := \{\mathbf{b} \in \mathcal{B}_{ij}^L \mid b_i = 0\}$ and $T := \{\mathbf{b} \in \mathcal{B}_{ij}^L \mid b_i = 1\}$. By definition, \mathbf{b} is a bid in S if $b_i = 0$, $b_j = p_j$ and $b_k \leq p_k$. By the same argument as above, S either contains a single positive bid at $b_j = p_j$ and $b_k = 0$ for all $k \neq i$, or it contains an equal number of positive and negative bids. Finally, note that \mathbf{b} is a bid in T if $b_i = 1$, $b_j = 1 + p_j$ and $b_k \leq 1 + p_k$ for all $k \in [n]$. The last constraint is vacuous as it holds for all possible values of b_k and p_k . Hence T is empty if $p_j = 1$ and contains an equal number of positive and negative bids otherwise. Together, these properties of S and T imply the claim. \square

We are now ready to prove Lemma 6.5.

PROOF OF LEMMA 6.5. Suppose that gadget G is placed at \mathbf{x} . We apply Theorem 2.3 to prove the claim. Moreover, we can use a result given by Lemma B.10 in [1] to restrict ourselves to checking the validity criterion in Theorem 2.3 only at a finite number of price points (see [1], Appendix B for details). In our case, it suffices to verify the conditions for each vertex $\mathbf{p} \in \mathbf{x} + \{0, 1\}^n$ and $\mathbf{p} \in \mathbf{x} + \{2, 3\}^n$.

Fix $\mathbf{p} \in \mathbf{x} + \{0, 1\}^n \cup \{2, 3\}^n$ and distinct goods $i, j \in [n]_0$. We consider first the case where $i = 0$ and $j \in [n]$. By Lemma 6.6, the sum of weights of the gadget bids indifferent between 0 and j is at least -1 . Secondly, by construction of the adversarial game, there is a gadget bid at $p_i \mathbf{e}^i$ and this bid is indifferent between i and the reject good. Hence the total sum of weights of the bids in the game indifferent between i and rejecting is at least 0.

Now let i, j be distinct (non-reject) goods in $[n]$. Note that $\mathbf{b} \in \{0, 1\}^n$ is indifferent between i and j at \mathbf{p} , then $\mathbf{b}' = \mathbf{b} + 2 \in \mathbf{x} + \{2, 3\}^n$ is also indifferent between i and j . Moreover, \mathbf{b} and \mathbf{b}' have opposite weights. This fact, together with Lemma 6.7, implies that the sum of the weights of the gadget bids is at least -1 . Again, note that for any \mathbf{p} there is a boundary bid that is indifferent between i and j . Indeed, recall that a bid is indifferent between i and j if and only if $b_i - b_j = p_i - p_j$, $b_i \geq p_i$ and $b_k - p_k \leq b_i - p_i$. If $p_i \geq p_j$, then the bid $Me^i + (M - (p_i - p_j))e^j$ satisfies these conditions. Similarly, if $p_i < p_j$, then the bid $(M - (p_i - p_j))e^i - Me^j$ satisfies this condition. As this bid is present by construction of the adversarial game, we see that the total sum of weights of all bids that are indifferent between i and j in the game is at least 0 in all cases. \square

THEOREM 6.8. *Let \mathcal{B} be a bid list that may comprise positive and negative bids. If \mathcal{B} is valid, any algorithm requires $\Omega\left(\left(\frac{B-2^{n+1}}{8n^2}\right)^n\right)$ queries to learn \mathcal{B} . If \mathcal{B} is allowed to be invalid, $\Omega\left(\left(\frac{M}{4}\right)^n\right)$ queries are required.*

PROOF. Consider first the case where \mathcal{B} is valid. Let \mathcal{B} be the adversary's bid list, which is valid by Lemma 6.5. By Lemma 6.3, we know that the gadget placed at \mathbf{x} does not affect demand outside the cube $\mathbf{x} + [3]_0^n$. Hence by a standard decision tree argument, any algorithm requires $k^n - 1$ queries to learn the location of the gadget G . Note that the adversary places a total of $B \leq 2^{n+1} + 8kn^2$ gadget and boundary bids. Solving for k gives $k \geq \frac{B-2^{n+1}}{8n^2}$. Hence expressing the query complexity of $k^n - 1$ in terms of B and n yields the desired expression.

For the case where \mathcal{B} may be invalid, we construct an even simpler adversarial game, where the adversary positions the island gadget as above but places no boundary bids. By the same argument, learning the location of the gadget incurs at least $\Omega\left(\left(\frac{M}{4}\right)^n\right)$ queries. \square

Using the fact that $\max\{f, g\} = \Omega(f + g)$ for any positive functions f and g , we can combine the lower bounds from Theorem 6.1 and Theorem 6.8, along with the upper bound from Algorithm 3 in Section 5, to identify the query complexity of expressing SS demand by means of positive and negative bids.

COROLLARY 6.9. *For constant n , learning SS demand requires $\Theta(B^n + B \log M)$ queries.*

7 CONCLUSIONS

Our algorithms for learning demand are conceptually simple and provide the first systematic approach for bidders to express their preferences in the bidding language used by the product-mix auction. This allows bidders with non-technical backgrounds to participate in these auctions under the mild assumption that they are able to answer demand oracle queries. In the setting where demand can be expressed using positive bids only, our algorithm achieves linear query complexity. When demand may only be expressible using positive *and* negative bids, our hyperplane finding algorithm performs well if the number of goods is not too large. Further work could address extending our positive-bid algorithm to allow a small number of negative bids, approximate learning of the demand function, and dealing with errors in answers to queries.

ACKNOWLEDGMENTS

The authors thank their anonymous reviewers for insightful comments and suggestions. We are also grateful for valuable comments from Elizabeth Baldwin.

REFERENCES

- [1] Elizabeth Baldwin, Paul W. Goldberg, Paul Klemperer, and Edwin Lock. 2019. Solving Strong-Substitutes Product-Mix Auctions. *ArXiv preprint abs/1909.07313* (2019). <http://arxiv.org/abs/1909.07313>
- [2] Elizabeth Baldwin and Paul Klemperer. 2019. Understanding Preferences: “Demand Types”, and the Existence of Equilibrium with Indivisibilities. *Econometrica* 87, 3 (May 2019), 867–932. <https://doi.org/10.3982/ECTA13693>
- [3] Elizabeth Baldwin and Paul Klemperer. 2021. Proof that the Product-Mix Bidding Language can represent any Substitutes Preferences. (2021). https://www.nuffield.ox.ac.uk/economics/Papers/2021/2021-W05_subproof.pdf Nuffield Discussion Papers (working paper).
- [4] Avrim Blum, Jeffrey C. Jackson, Tuomas Sandholm, and Martin Zinkevich. 2004. Preference Elicitation and Query Learning. *J. Mach. Learn. Res.* 5 (2004), 649–667. <http://jmlr.org/papers/volume5/blum04a/blum04a.pdf>
- [5] Wolfram Conen and Tuomas Sandholm. 2001. Preference elicitation in combinatorial auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*. 256–259.
- [6] Wolfram Conen and Tuomas Sandholm. 2002. Differential-Revelation VCG Mechanisms for Combinatorial Auctions. In *Agent-Mediated Electronic Commerce IV, Designing Mechanisms and Systems, AAMAS 2002 Workshop on Agent Mediated Electronic Commerce, Bologna, Italy, July 16, 2002, Revised Papers (Lecture Notes in Computer Science, Vol. 2531)*, Julian A. Padget, Onn Shehory, David C. Parkes, Norman M. Sadeh, and William E. Walsh (Eds.). Springer, 34–51. https://doi.org/10.1007/3-540-36378-5_3
- [7] Wolfram Conen and Tuomas Sandholm. 2002. Partial-Revelation VCG Mechanism for Combinatorial Auctions. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence and Fourteenth Conference on Innovative Applications of Artificial Intelligence, July 28 - August 1, 2002, Edmonton, Alberta, Canada*, Rina Dechter, Michael J. Kearns, and Richard S. Sutton (Eds.). AAAI Press / The MIT Press, 367–372. <http://www.aaai.org/Library/AAAI/2002/aaai02-056.php>
- [8] Venkatesan Guruswami, Jason D Hartline, Anna R Karlin, David Kempe, Claire Kenyon, and Frank McSherry. 2005. On profit-maximizing envy-free pricing. In *SODA*, Vol. 5. 1164–1173.
- [9] Benoit Hudson and Tuomas Sandholm. 2004. Effectiveness of query types and policies for preference elicitation in combinatorial auctions. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 1*. IEEE Computer Society, 386–393.
- [10] Paul Klemperer. 2008. A New Auction for Substitutes: Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions. (2008). <https://www.nuffield.ox.ac.uk/economics/Papers/2008/substsauc.pdf> working paper, Nuffield College.
- [11] Paul Klemperer. 2010. The Product-Mix Auction: A New Auction Design for Differentiated Goods. *Journal of the European Economic Association* 8, 2–3 (2010), 526–536.
- [12] Paul Klemperer. 2018. Product-Mix Auctions. (2018). <https://www.nuffield.ox.ac.uk/economics/Papers/2018/2018W07productmix.pdf> working paper 2018-W07, Nuffield College.
- [13] Sebastien M Lahaie and David C Parkes. 2004. Applying learning algorithms to preference elicitation. In *Proceedings of the 5th ACM conference on Electronic commerce*. 180–188.
- [14] Kazuo Murota. 2013. *Discrete Convex Analysis*. Society for Industrial and Applied Mathematics.
- [15] Noam Nisan. 2000. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM conference on Electronic commerce*. 1–12.
- [16] Noam Nisan and Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129, 1 (2006), 192–224. <https://doi.org/10.1016/j.jet.2004.10.007>
- [17] Akiyoshi Shioura. 2004. Fast scaling algorithms for M-convex function minimization with application to the resource allocation problem. *Discrete Applied Mathematics* 134, 1 (2004), 303 – 316. [https://doi.org/10.1016/S0166-218X\(03\)00255-5](https://doi.org/10.1016/S0166-218X(03)00255-5)
- [18] Akiyoshi Shioura and Akihisa Tamura. 2015. Gross Substitutes Condition and Discrete Concavity for Multi-Unit Valuations: a Survey. *Journal of the Operations Research Society of Japan* 58, 1 (2015), 61–103. <https://doi.org/10.15807/jorsj.58.61>
- [19] Akihisa Tamura. 2004. Coordinatewise domain scaling algorithm for M-convex function minimization. *Mathematical Programming* 102, 2 (2004), 339. <https://doi.org/10.1007/s10107-004-0522-y>
- [20] Hanrui Zhang and Vincent Conitzer. 2020. Learning the Valuations of a k -demand Agent. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*, Hal Daumé III and Aarti Singh (Eds.). PMLR, 11066–11075.
- [21] Martin A Zinkevich, Avrim Blum, and Tuomas Sandholm. 2003. On polynomial-time preference elicitation with value queries. In *Proceedings of the 4th ACM Conference on Electronic Commerce*. 176–185.