

IMPLEMENTING WALRASIAN EQUILIBRIUM: THE LANGUAGES OF PRODUCT-MIX AUCTIONS*

ELIZABETH BALDWIN

Department of Economics, University of Oxford

PAUL KLEMPERER

Department of Economics, University of Oxford

EDWIN LOCK

Departments of Computer Science and Economics, University of Oxford

July 2024

Product-mix auctions are sealed-bid mechanisms for trading multiple divisible or indivisible units of multiple differentiated goods. They implement competitive-equilibrium allocations when these exist, based on the bids that participants make in a simple geometric language. All concave substitutes (respectively, strong-substitutes) valuations can be uniquely represented, and no other valuations can be represented, by bids in the corresponding version of this language. This provides new characterisations of ordinary substitutes, and of strong substitutes, when goods are indivisible. We discuss implementation of the auctions, and extensions and variants of the language, e.g., allowing for budget constraints.

KEYWORDS: product-mix auction, bidding language, competitive equilibrium, Walrasian equilibrium, arctic auction, substitutes, strong substitutes, indivisible goods, product mix auction.

1. INTRODUCTION

This paper develops the theoretical underpinnings, and further extensions, of the product-mix auction design that was originally developed for the Bank of England in 2007-8.

Product-mix auctions (PMAs) are easy-to-use, single-round, sealed-bid mechanisms to sell or procure multiple differentiated goods. Each good may be either divisible, or available in multiple indivisible units. The auctions allow both the bidders and the auctioneer to express rich preferences about how the allocations they receive depend on the auction prices.¹

PMAs implement competitive (i.e., Walrasian) equilibria when these exist, assuming participants express their preferences accurately.²

PMAs can thus obtain close to efficient results when bidders' strategic behaviour is not a first-order concern. Although most auction research emphasises game-theoretic issues of strategic interaction among bidders with asymmetric information, these issues are often not very important, in particular, when there are many bidders. If no bidder's demand is too large a fraction of

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

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

Edwin Lock: edwin.lock@cs.ox.ac.uk

* Acknowledgements to follow.

¹Free open-source software to run all our versions of the PMA is at <http://pma.nuff.ox.ac.uk/>.

²Alternative versions of the PMA implement the competitive equilibrium allocations based on the reported preferences, but then use a different pricing rule than charging the competitive equilibrium prices. (E.g., [Finster \(2021\)](#) experimentally tests discriminatory pricing in a PMA.) However, competitive pricing usually provides good incentives for truthful bidding.

Note that the auctioneer can make the PMA a profit-maximising auction by misrepresenting her own preferences, as we discuss below.

the market's, competitive equilibrium prices approximate Vickrey pricing which, in the right circumstances, induces truthful reporting of preferences, and so permits efficient allocations.³

However, inducing accurate reporting is not just a matter of incentives. Importantly, PMA “bidding languages” are easy to understand and use—they break down complex preferences into small pieces—thus allowing bidders (and the auctioneer) to accurately express their preferences in a wide range of contexts.

Because PMA bidding languages are geometric in basis, we can draw simple diagrams that help participants to visualise the auctions, and to understand how prices and allocations are determined.⁴ The geometric approach also facilitates the analysis of the auctions, and allows us to develop results about how different classes of preferences can be represented.

Practical bidding languages must balance simplicity that encourages participation against precision in what participants can express. So different PMA languages will be appropriate for different contexts.

Whichever PMA language is used, each bidder submits a set of bids that, taken together, express her preferences. The PMA then determines a price for each good, and allocates each bidder a bundle that maximises her utility at these prices, assuming she reported her preferences accurately.

Most naturally, the bidders also pay these prices for the goods in their bundles (so also every recipient of any quantity of any good pays the same per-unit price for that good)—and this incentivises (approximately) truthful behaviour by bidders whose demands are not too large relative to the auction's supply. However, alternative pricing rules are possible (see note 2 above).

An important feature of the PMA is that it gives the auctioneer flexibility about how the quantities sold depend on expressed demand. Moreover, the auctioneer need not behave competitively. It will do if (like the Bank of England, for whom the PMA was originally designed) its objective is social efficiency. But it can use its monopoly power and choose a “supply function” that does not reflect its true preferences, or even choose supply after seeing the bids, to pursue an alternative objective such as profit.⁵

Section 2 introduces the *substitutes PMA language* that is the main focus of our paper. We introduce and analyse it in its general form; practical implementations are likely to use simplifications of it. The *strong-substitutes PMA language* is a straightforward simplification. Further simplifications are easy to make, and have been made in, for example, the Bank of England's application.⁶

³Grace (2024a)'s empirical work finds that competitive behaviour is a better model than Nash equilibrium for the Bank of England's PMAs (which charge all bidders the prices that support competitive equilibrium, assuming bidders express preferences truthfully); and Milton Friedman assumes bidders bid truthfully in uniform-price Treasury auctions (which are a very simple special case of the Bank of England's PMA) in his (1959, 1963, 1991) arguments for uniform pricing. However, strategic behaviour is, of course, important in many contexts. Finster (2020) models strategic behaviour in a PMA; Holmberg et al. (2019)'s model can also be interpreted in this way.

⁴The set of bids that any individual bidder makes is simply a list of vectors, in which each vector expresses a component of her preferences.

⁵The allocation of goods among bidders will still be (approximately) efficient, assuming uniform pricing and that no bidder's demand is too large relative to supply. The Icelandic Government aimed to maximise profit in its proposed PMA (see Section 5.5).

⁶These languages were originally developed by Klemperer (2008), responding to the Governor of the Bank of England's 2007 request for a mechanism to allocate central-bank funds to bidders who would be permitted to offer different qualities of collateral. The context was that the UK suffered its first bank run for 140 years in September 2007 in an early sign of the financial crisis. Efficiency required charging different “prices” (i.e., different interest rates) for different loans; setting wrong prices would both mis-allocate funds in the current auction and incentivise undesirable activities such as commercial banks over-investing in “toxic assets”. The Bank of England implemented

Our languages can be used by buyers or sellers, or traders who may be on either or both sides of the market depending on prices.

Section 3 gives the representation theorem that is the main contribution of this paper: all concave⁷, quasilinear, substitutes preferences can be represented, and can be uniquely represented, by sets of bids in the substitutes PMA language. An immediate corollary is that all quasilinear, strong-substitutes preferences⁸ can be uniquely represented by sets of bids in the strong-substitutes PMA language. We explain the intuition for our results by describing the proof for multiple units of each of two goods in the context of a simple example. We sketch the additional arguments needed for the general case, but defer the full details to the Appendix.

Section 4 complements these results by identifying in a simple way which sets of substitutes PMA bids represent a (concave) valuation. These are the sets of bids that bidders are permitted to make in a PMA; we call them “valid”. We show that any valid set of substitutes PMA bids represents a substitutes valuation, and that any valid set of strong-substitutes PMA bids represents a strong-substitutes valuation. So there is a one-to-one correspondence between valid sets of substitutes PMA bids and (concave) substitutes valuations, and between valid sets of strong-substitutes PMA bids and strong-substitutes valuations.⁹

The PMA languages therefore provide new characterisations for indivisible goods of both ordinary substitutes and strong substitutes, and so—especially because our characterisations are as the sum of small, simple, pieces—gives us new ways to understand these classes of preferences. To our knowledge, no other language proposed for these preferences can represent the entirety of either class while representing no more than that class.¹⁰

Since our proof of our representation theorem is constructive, it leads naturally to straightforward algorithms that allow a bidder to create the set of bids that represent her preferences by responding to a list of questions that elicit her demand at different prices.

Section 5 discusses the implementation of PMAs using these bidding languages, including the Bank of England’s implementations of its PMAs.

Our language is “compact” in that many valuations can be expressed using only a small number of bids. However, we describe how predefining sets of bids—we call these sets of bids “words”—can allow to express natural preferences even more concisely.

In many auctions bidders are subject to income effects, or even hard budget constraints. For example, bidders offering to swap bonds for alternative financial assets cannot bid to swap

simplified versions of [Klemperer \(2008\)](#); these auctions are currently run weekly and have been used to auction approximately £240 billion in repos.

⁷Concavity keeps our language simple. (E.g., a bid that expresses a willingness to pay (up to) a per-unit price p for 100 units will imply that, if the per-unit price is exactly p , then the bidder is indifferent between buying any number of units between 0 and 100.) Moreover, we are particularly focused on contexts where we can find competitive equilibrium, for which concavity is important.

⁸Strong substitutes is the terminology coined by [Milgrom and Strulovici \(2009\)](#). It is, by definition, concave, and equivalent to M^\natural -concavity (see [Murota and Shioura \(1999\)](#), [Murota \(2003\)](#) and [Shioura and Tamura \(2015\)](#)).

⁹Alternative, longer proofs of the results of the preceding two paragraphs are developed in a technical note, [Baldwin and Klemperer \(2021\)](#). [Baldwin and Klemperer \(2016\)](#) presented a proof for the strong-substitutes case. [Lin and Tran \(2017\)](#) showed how any valuation associated with a “full positive basis” can be decomposed into a combination of simpler pieces, which gives the same decomposition as we do in the strong substitutes case, but the wider classes of valuations which we can decompose are different. [Klemperer \(2010\)](#) stated the result for strong-substitutes with multiple units of each of two goods.

¹⁰For example, neither [Hatfield and Milgrom \(2005\)](#)’s endowed assignment messages nor [Milgrom \(2009\)](#)’s (integer) assignment messages can express all strong-substitute valuations, see [Ostrovsky and Paes Leme \(2015\)](#), and [Fichtl \(2021\)](#), respectively. Furthermore, it is not possible to build up all strong-substitute valuations from weighted matroid rank functions, either by using positive linear combinations ([Balkanski and Paes Leme, 2020](#)); or by merging and endowing these rank functions, while restricting the ground sets of the matroids to the number of goods ([Tran, 2021](#)).

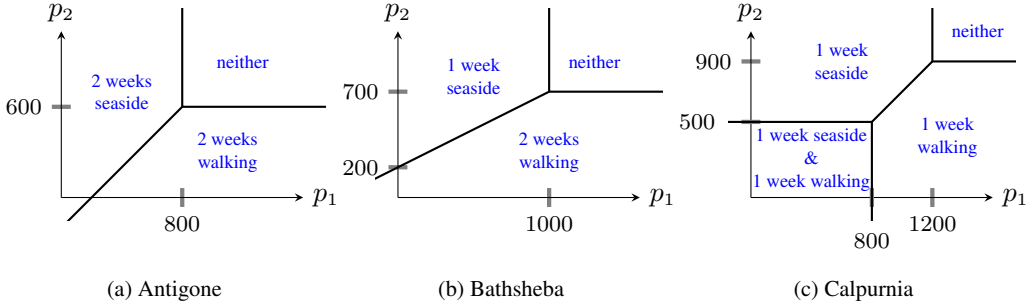


FIGURE 1.—Illustrations of the demands of Antigone, Bathsheba and Calpurnia in Example 1. Here p_1 denotes the price per week at the seaside, and p_2 is the price per week of walking.

more than their current holding. Since the Icelandic government asked one of us to develop a variant of the product-mix auction for this environment, we describe an appropriate language for this as “arctic” (see [Klemperer \(2018\)](#)). We note that an arctic bid can be understood as a “word” in the standard language.

Section 6 briefly compares PMAs and their languages to related auctions and languages.

Section 7 concludes. Omitted proofs are in Appendices. A less technical exposition of the PMA (more suitable for most users of the auction) can be found in [Klemperer \(2018\)](#).

2. THE SUBSTITUTES PRODUCT-MIX AUCTION LANGUAGE

The substitutes product-mix auction language allows each bidder to express (concave quasi-linear) substitutes preferences over arbitrary quantities of differentiated goods.

Our bidders can be buyers or sellers, or traders who may be on either or both sides of the market depending on prices. However, for simplicity, we will assume they are buyers throughout the paper except where stated otherwise. We introduce the language using Example 1 which describes the valuations of three potential buyers for two kinds of holiday. Figure 1 represents these buyers’ preferences by showing their demands as functions of the prices of goods 1 (the seaside) and 2 (walking).

EXAMPLE 1: Antigone has up to two weeks’ vacation time available, and is willing to pay up to £800 per week to be at the seaside, or £600 per week for a walking holiday.

Bathsheba would pay up to £1,000 for a week at the seaside; she would also spend up to £1,400 for *two* weeks of walking. (She is alternatively willing to pay £700 for *one* week of walking, but isn’t interested in more than one week at the seaside or in having both kinds of holiday.)

Calpurnia values a week at the seaside at £1,200, and a week’s walking at £900. She doesn’t want to spend more than one week on either holiday, but is interested in taking both—she values taking both at £1,700. So if she was already taking the walking holiday, she would pay up to an additional £800 ($=£1,700-£900$) to also take the seaside holiday; if she was anyway taking the seaside holiday, she would pay up to an additional £500 ($=£1,700-£1,200$) to also take the walking holiday.

We will see that, for example, Antigone can precisely describe her preferences using a single PMA bid that (i) states that she is indifferent among the three options of buying 1 unit of good 1, or 1 unit of good 2, or no purchase, at prices (800, 600), (ii) has “tradeoff” (1,1) between the goods (if she switched between the two kinds of holiday, she would still take the same amount

of vacation time), and (iii) has “multiplicity” 2 (since she is interested either 1 or 2 weeks’ vacation). Bathsheba can also precisely describe her preferences using a single bid (but has tradeoff (1, 2)), while buyers with more complex preferences, such as Calpurnia, can describe them with a set of bids of this kind.

2.1. Preliminaries

Our results apply to both indivisible and divisible goods. Because the divisible case follows from the indivisible one (see Section 2.4), we will mostly discuss the latter, providing the corresponding results where required.

There are $n + 1$ goods $[n]_0 := \{0, 1, \dots, n\}$, in which the *true* goods are $[n] := \{1, \dots, n\}$, and 0 corresponds to a notional *null* good. Receiving a quantity of the null good, 0, corresponds to receiving nothing. Our vectors are indexed by $[n]_0$. We write e^i for the coordinate vectors for $i \in [n]$ and notate $e^0 := \mathbf{0}$.¹¹

Prices \mathbf{p} are arbitrary for the true goods, and we set $p_0 = 0$, so the set of all possible prices is $\mathcal{P} := \{\mathbf{p} \in \mathbb{R}^{[n]_0} \mid p_0 = 0\}$. Similarly, *bundles* \mathbf{x} of goods are $n + 1$ -dimensional real vectors, and we fix $x_0 = 0$. Although the 0th entries are formally important, we write examples of prices and bundles as n -dimensional vectors for easier readability (because $p_0 = 0$ and $x_0 = 0$ are fixed and payoff-irrelevant), and describe the case of n true goods as the n -good case.

A *valuation* v maps a set X of bundles to real values. It gives rise to *quasi-linear demand* $D_v(\mathbf{p}) := \operatorname{argmax}_{\mathbf{x} \in X} (v(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x})$, defined for any price $\mathbf{p} \in \mathcal{P}$. We focus on the setting with indivisible goods in which the domain X of the valuation is finite and consists of integer bundles, so it is a subset of $\mathcal{X} := \{\mathbf{x} \in \mathbb{Z}^{[n]_0} \mid x_0 = 0\}$. For divisible goods, the domain is the set of all convex combinations of bundles in some such X and the valuation is continuous. Most of the valuations we consider are concave.¹²

A valuation v is *substitutes* if $\{\mathbf{x}\} = D_v(\mathbf{p})$ and $\{\mathbf{x}'\} = D_v(\mathbf{p} + \lambda e^i)$ implies $x_j \leq x'_j$ whenever $\lambda > 0$, $\mathbf{p} \in \mathcal{P}$ and $i, j \in [n]$ with $j \neq i$. A valuation for indivisible goods is *strong substitutes* if it is substitutes when we consider every unit of every good to be a separate good (see Milgrom and Strulovici (2009)). Strong substitutes guarantee existence of competitive equilibrium for indivisible goods, and we will see in Section 3.1 that they correspond to “one-to-one tradeoffs”.¹³ So Antigone, Bathsheba and Calpurnia all express substitutes preferences, but only Antigone and Calpurnia express strong-substitutes preferences.

2.2. The Substitutes PMA Language

A bidder in a substitutes PMA submits a set of bids that will depict a substitutes valuation (see Sections 3 and 4).

¹¹Note that e^0 is *not* the 0-th coordinate vector; this allows us to more conveniently notate bundles.

¹²A valuation for indivisible goods is *concave* (also known as *concave-extensible*) if it can be extended to a concave function. (That is, there exists a concave function $\hat{v} : \operatorname{conv}(X) \rightarrow \mathbb{R}$ such that $\hat{v}(\mathbf{x}) = v(\mathbf{x})$ for all $\mathbf{x} \in \operatorname{conv}(X) \cap \mathcal{X}$.) It is a standard consequence of the supporting hyperplane theorem that a valuation v is concave if and only if $D_v(\mathbf{p}) = \mathcal{X} \cap \operatorname{conv} D_v(\mathbf{p})$ at any $\mathbf{p} \in \mathcal{P}$.

¹³We use the same definition of substitutes for divisible and indivisible goods; it is clear that a valuation for indivisible goods is substitutes if and only if its concave extension is. When goods are indivisible, this definition is the most permissive extension of the standard definition of divisible gross substitutes to the indivisible case, and so these “substitutes” are the largest class of “substitutes” valuations for indivisible goods identified in the literature (see Shioura and Tamura (2015) for a survey). Baldwin and Klemperer (2019b) call this definition “ordinary substitutes”; it is strictly weaker than Milgrom and Strulovici (2009)’s “weak substitutes”, but the definitions coincide when there is only one unit of each good available, so our definition of “strong substitutes” is equivalent to Milgrom and Strulovici (2009)’s (see Danilov et al. (2003) and Shioura and Tamura (2015)). In the divisible case, the definitions based on “ordinary” and “weak” substitutes identify the same set of valuations.

Each *bid* is a triple $(\mathbf{r}; \mathbf{t}; m)$ consisting of a *root vector* \mathbf{r} whose coordinates are the values the bid places on the goods; a *tradeoff vector* \mathbf{t} describing the tradeoffs between the goods; and a *multiplicity* $m \in \mathbb{Z}$ scaling up how many units are demanded. We call a bid *positive* if $m > 0$ and *negative* if $m < 0$.

The root's entries (indexed by $[n]_0$) are drawn from \mathbb{R} together with $-\infty$. The 0th entry allows us to conveniently notate the case in which a bid regards an allocation of nothing (that is, of only the null good) *unacceptable* at any prices: in that case $r_0 = -\infty$. Otherwise we set $r_0 = 0$, indicating a zero value for an allocation of nothing. We also allow for the possibility that some true goods are completely unacceptable, that is, have value “ $-\infty$ ”, and we say that $I := \{i \in [n]_0 \mid r_i > -\infty\}$ is the set of goods in which the bid is *interested*. We say a bid is *regular* if it is interested in all the goods including the null good, so $I = [n]_0$. So a regular bid has a zero value for $\mathbf{0}$ and values every good above $-\infty$. We expect buyers to only need to use regular bids in any standard auction for substitutes (see Corollary 3.3 and the discussion following it). However, other bids can be useful (see, especially, Section 2.3).

The set \mathcal{T} of tradeoff vectors is also indexed by $[n]_0$, but we always set $t_0 = 1$. We also require that the vector $\mathbf{t}_{-0} = (t_1, \dots, t_n)$ is a non-negative *primitive* integer vector (the greatest common divisor of its entries is 1). So the set \mathcal{T} consists of the vectors in $\mathbb{Z}_{\geq 0}^{[n]_0}$ satisfying these properties. We set $t_i = 0$ for any true goods in which the bid is not interested (i.e., for $i \in [n] \setminus I$).

We now define the demand correspondence D_b of a positive bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ (i.e., a bid with $m > 0$). For any price \mathbf{p} , first define the bid's utility from mt_i units of a good i as $mt_i(r_i - p_i)$ for goods which it is interested in and $-\infty$ otherwise. At generic prices \mathbf{p} , the demand $D_b(\mathbf{p})$ of a positive bid is the unique bundle containing mt_i units of (only) the good i maximising its utility. (The bid demands only the zero bundle, $\mathbf{0}$, if the null good, 0 , uniquely yields the greatest utility.) At non-generic prices \mathbf{p} at which no bundle uniquely maximises utility, $D_b(\mathbf{p})$ is the set of all the integer bundles that are convex combinations of the utility-maximising bundles.

In particular, if $p_i = r_i$ for all goods in which the bid is interested (all $i \in I$), the bid is indifferent among all the bundles of mt_i items of any good i in which the bid is interested, including being indifferent about receiving nothing if $0 \in I$ (and the bid is also indifferent about receiving any integer bundle which is a convex combination of these bundles).

The demand D_b of a negative bid $(\mathbf{r}; \mathbf{t}; m)$ is the *negative* of the demand of bid $(\mathbf{r}; \mathbf{t}; |m|)$. That is, if such a bid is part of a set of bids, then the demand $D_b(\mathbf{p})$ of the bid $(\mathbf{r}; \mathbf{t}; |m|)$ will be subtracted from the demand of the remainder of the set.¹⁴

So $D_b(\mathbf{p}) = \mathcal{X} \cap \text{conv}\{mt_i e^i \mid i \in J\}$ where $J = \text{argmax}_{i \in I} t_i(r_i - p_i)$. We interpret D_b in terms of the demand of valuations, in Lemma 2.1 below.

Examples of the different possible types of bids for the two-good case, and their demands, $D_b(\mathbf{p})$, are illustrated in Figure 2. (The demand of an *unconditional* bid interested in only one good is constant, so not shown.)

In examples, we will often find it clearer to write bids in the format $(\mathbf{r}_{-0}, \mathbf{t}_{-0}, m)$; as $t_0 = 1$ always, this only loses the specification of r_0 , which is always implicitly 0 when we use this shorthand. Doing this only excludes bids for which an allocation of $\mathbf{0}$ is unacceptable. (All regular bids have $r_0 = 0$.)

In Example 1 (see also Figure 1), Antigone and Bathsheba can express their demands with a single bid each, namely $(\mathbf{r}_{-0}; \mathbf{t}_{-0}; m) = (800, 600; 1, 1; 2)$ for Antigone, and $(1000, 700; 1, 2; 1)$ for Bathsheba. We will explain below that Calpurnia can express her demand by making

¹⁴Note that demanding negative units is *not* the same as selling, since the negative units are only demanded at prices that are *low* enough, whereas a seller would only wish to sell at prices that are *high* enough. So a negative bid should be understood as a “cancellation” bid, *not* as an offer to sell.




FIGURE 2.—Demands of different possible types of bids for $n = 2$. (a) a regular bid (i.e., interested in all the goods, $\{0, 1, 2\}$) with root $\mathbf{r} = (0, r_1, r_2)$, tradeoffs $\mathbf{t} = (1, t_1, t_2)$, and multiplicity m ; (b) a bid interested in goods $\{0, 1\}$, with root $\mathbf{r} = (0, r_1, -\infty)$, tradeoffs $\mathbf{t} = (1, 1, 0)$ and multiplicity m ; (c) a bid interested in goods $\{0, 2\}$, with root $\mathbf{r} = (0, -\infty, r_2)$, tradeoffs $\mathbf{t} = (1, 0, 1)$ and multiplicity m ; (d) a bid interested in the “true” goods $\{1, 2\}$ but not in the null good, with root $\mathbf{r} = (-\infty, 0, r_2)$, tradeoffs $\mathbf{t} = (1, t_1, t_2)$, and multiplicity m ; (not shown) demand is constant if a bid is an unconditional bid only interested in one good. The demand in each region of price space is shown in the corresponding region. Demand on a boundary of a region is the set of integer bundles that are convex combinations of the bundles demanded in the adjacent regions. (The slopes of the diagonal lines are t_1/t_2 , and the circled numbers are “weights”, as explained in Section 3.1 below.)

three positive bids $(1200, 900; 1, 1; 1)$, $(800, -\infty; 1, 0; 1)$, $(-\infty, 500; 0, 1; 1)$, and one negative bid $(800, 500; 1, 1; -1)$.¹⁵ In particular Antigone and Calpurnia have one-to-one tradeoffs, but Bathsheba does not.

A bid’s demand is associated with a valuation as follows.

LEMMA 2.1: *If bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ is positive (i.e., $m > 0$), then $D_{\mathbf{b}} = D_{v_{\mathbf{b}}}$, where the valuation $v_{\mathbf{b}} : \mathcal{X} \cap \text{conv}\{mt_i \mathbf{e}^i \mid i \in I\} \rightarrow \mathbb{R}$ is defined by $v_{\mathbf{b}}(\mathbf{x}) = \sum_{i \in I} r_i x_i$ and I is the set of goods in which \mathbf{b} is interested. If \mathbf{b} is negative (i.e., $m < 0$), then $D_{\mathbf{b}} = -D_{|\mathbf{b}|}$ in which $|\mathbf{b}| = (\mathbf{r}; \mathbf{t}; |m|)$.*

So positive bids represent simple valuations. Indeed, if the bid has tradeoffs $\mathbf{t} = \mathbf{1}$ and multiplicity 1, then it represents unit demand over goods I with values \mathbf{r} (cf. Gul and Stacchetti (1999)). More generally, the demand of a positive bid with tradeoffs $\mathbf{t} = \mathbf{1}$ is strong-substitutes (see the end of Section 3.1), so we call any bid with $\mathbf{t} = \mathbf{1}$ a *strong-substitutes bid*. Negative bids do not represent valuations, but rather the subtraction of the demand of an associated valuation.

We now address the uniqueness of a bid associated with a demand correspondence.

LEMMA 2.2: *For bids $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ and $\mathbf{b}' = (\mathbf{r}'; \mathbf{t}'; m')$, we have $D_{\mathbf{b}} = D_{\mathbf{b}'}$ if and only if $\mathbf{t} = \mathbf{t}'$, $m = m'$, they are interested in the same set of goods, I , and $t_i(r_i - r'_i) = t_j(r_j - r'_j)$ for all $i, j \in I$. In particular, if $0 \in I$, then $D_{\mathbf{b}} = D_{\mathbf{b}'}$ if and only if $\mathbf{b} = \mathbf{b}'$.*

That is, bids with value 0 for receiving nothing are uniquely identified by their demand correspondence, but bids for which receiving the null good is unacceptable can define the same demand even if their roots are distinct. For example, the demand shown in Figure 2d is defined by any bid with root $\mathbf{r} = (-\infty, r, r')$ such that (r, r') lies on the line shown in Figure 2d with the tradeoffs $\mathbf{t}_{-0} = (t_1, t_2)$, and multiplicity m .

¹⁵We will show in Section 3.2 that in any practical auction a buyer can express her demand at all relevant prices using only regular bids. For example, if the auctioneer’s reserve prices are (strictly) positive, Calpurnia can substitute her second and third positive bids by $(800, 0; 1, 0; 1)$ and $(0, 500; 0, 1; 1)$.




FIGURE 3.—The demand expressed by: (a) Calpurnia’s three positive bids; (b) Calpurnia’s single negative bid; and (c) all Calpurnia’s bids, \mathcal{B} ; in Example 1. The dashed lines in (c) show where demand changes in (a), but no longer changes after incorporating the negative bid shown in (b). For example, bundle $(1, 1)$ is the only bundle demanded in (c) in the quadrant of prices below $(800, 500)$ because both $(2, 1) + (-1, 0) = (1, 1)$ and $(1, 2) + (0, -1) = (1, 1)$.

So we normalise bid $(\mathbf{r}; \mathbf{t}; m)$ interested in goods I by replacing \mathbf{r} with the unique root \mathbf{r}' that satisfies $t_i(r_i - r'_i) = t_j(r_j - r'_j)$ for all $i, j \in I$ and $r'_i = 0$ for the smallest good $i \in I$. Then, after normalisation, any two distinct bids have distinct demand.

We can now see that Figure 2 illustrates *all* the possible bids in the two-good case, with the exception of bids that unconditionally demand only a single (true) good.

Each bidder gives the auctioneer a finite set, \mathcal{B} , of bids. Since demand $D_b(\mathbf{p})$ is single-valued at a dense set of prices in \mathbb{R}^n for each bid \mathbf{b} , we can define the demand $D_{\mathcal{B}}$ of a bid set \mathcal{B} as follows. For any $\mathbf{p} \in \mathbb{R}^n$, let $Q(\mathbf{p})$ be the set of all price vectors \mathbf{q} in a small local neighbourhood of \mathbf{p} at which every bid demands a unique bundle. Then:

$$D_{\mathcal{B}}(\mathbf{p}) := \mathcal{X} \cap \text{conv} \left\{ \sum_{\mathbf{b} \in \mathcal{B}} D_b(\mathbf{q}) \mid \mathbf{q} \in Q(\mathbf{p}) \right\}. \quad (2.1)$$

If \mathcal{B} is an empty set, we define $D_{\mathcal{B}}(\mathbf{p})$ to be $\{\mathbf{0}\}$ at all prices. Note that Equation (2.1) implies that any valuation v satisfying $D_v = D_{\mathcal{B}}$ is concave (see note 12).

A bid set is *parsimonious* if its bids are normalised and it contains at most one bid for any combination (\mathbf{r}, \mathbf{t}) of root and tradeoffs and no bids with multiplicity 0. Note that adding or removing bids with multiplicity 0, or substituting multiple bids with the same root and tradeoffs by a single bid with the sum of their multiplicities, does not change demand.¹⁶

DEFINITION 2.3: A *bid collection* \mathcal{B} is a parsimonious set of bids. A bid collection is *strong-substitutes* if it contains only strong-substitutes bids. It is *regular* if it contains only regular bids.

Figure 3 illustrates that Calpurnia’s bid collection corresponds to her demand (see also Figure 1).

2.3. Sellers and Traders

Because a supply of \mathbf{x} units is a demand of $-\mathbf{x}$ units, our language straightforwardly expresses preferences of sellers, and of traders.

¹⁶Baldwin et al. (2023a) used the alternative convention that bids could only have multiplicity ± 1 , and a bid set was parsimonious if there were no coincident bids with opposite signs. It is straightforward that a normalised parsimonious bid set under one convention corresponds to a unique normalised parsimonious bid set under the other.

Unconditionally selling m units of good i is represented by a bid only interested in good i , with multiplicity $-m$.¹⁷ Typically sellers in fact have a reservation price. And selling a good i if its price exceeds \tilde{r} is equivalent to selling it unconditionally *and* buying it back if (but only if) its price is less than or equal to \tilde{r} . Such “buyback bids” permit rich seller preferences across goods; see Section 5.2.

A trader, such as one described in Hatfield et al. (2013), who wishes, for example, to buy an apple (good 1) and sell an apple tart (good 2) if and only if the profit exceeds \tilde{r} (i.e., iff $p_2 - p_1 \geq \tilde{r}$) can do so by making one bid that sells an apple tart for sure, and a second bid that buys either an apple or an apple tart according as to whether or not $p_2 - p_1 \geq \tilde{r}$ (this latter bid is the bid shown in Figure 2d above, with $r_2 = \tilde{r}$ and $m = t_1 = t_2 = 1$).

So it is easy to use the PMA language to run a “product-mix market” for multiple agents who participate on either or both sides of the market.¹⁸

2.4. The Case of Divisible Goods

It is straightforward that we can approximate any divisible goods auction arbitrarily closely by an indivisible goods auction, by simply re-scaling quantities to arbitrarily tiny units. But our language can also be used to express demand for divisible goods directly: the demand of a bid collection \mathcal{B} of PMA substitutes bids for divisible goods at prices \mathbf{p} is then defined as $\hat{D}_{\mathcal{B}}(\mathbf{p}) := \text{conv } D_{\mathcal{B}}(\mathbf{p})$.¹⁹ In this divisible setting, we also allow bids to have rational multiplicities and tradeoffs. This is mathematically equivalent to integer multiplicities and tradeoffs, since we can re-scale quantities of goods appropriately.²⁰

PROPOSITION 2.4: *Suppose valuation \hat{v} for divisible goods is the concave envelope of concave valuation v for indivisible goods. A bid collection \mathcal{B} satisfies $\hat{D}_{\mathcal{B}} = D_{\hat{v}}$ if and only if $D_{\mathcal{B}} = D_v$.*

Proposition 2.4 allows us to handle all concave substitutes valuations, \hat{v} , for divisible goods that are piecewise-linear with rational vertices (because \hat{v} is the concave extension of an indivisible-goods valuation v after re-scaling quantities). Since any concave substitutes valuation for divisible goods can be approximated arbitrarily well with such piecewise functions, Proposition 2.4 seems unlikely to imply any restriction on a PMA in practice.²¹

So all our results on the expressivity of our language for indivisible goods can equivalently be expressed for appropriate valuations on divisible goods.

¹⁷Note that a negative bid interested in more than one good *cannot* be understood as a sale, since it selects good(s) with a *low* price relative to the root (whereas a seller prefers to sell *high* price good(s)). So such a bid only makes sense in combination with positive bid(s) which it will cancel at appropriate prices.

¹⁸Indeed one of us (Lock) recently worked with Bellus Ventures to implement a substitutes PMA as a platform for developers of renewable energy projects to sell tax credits, and hence allocate subsidies for clean energy more efficiently; software for this application is at <http://pma.nuff.ox.ac.uk/>.

¹⁹This, of course, also holds for the demand of an individual bid. While in the indivisible case we defined $D_b(\mathbf{p}) = \mathcal{X} \cap \text{conv}\{mt_i e^i \mid i \in J\}$, in the divisible case the demand is just $\hat{D}_b(\mathbf{p}) := \text{conv}\{mt_i e^i \mid i \in J\}$, where $J = \text{argmax}_{i \in I} t_i(r_i - p_i)$ and I is the set of goods.

²⁰We can, if desired, achieve a unique normalisation by scaling tradeoffs and multiplicities so that the tradeoff is a primitive integer vector.

²¹Moreover, after appropriately re-scaling quantities, Proposition 2.4 tells us that the unique bid collection \mathcal{B} expressing the same demand as \hat{v} is the unique bid collection satisfying $D_{\mathcal{B}} = D_v$ for indivisible goods.

3. OUR REPRESENTATION THEOREM

We now state the main results of this paper: any concave substitutes valuation can be expressed in the substitutes PMA language, and any strong-substitutes valuation can be expressed in the substitutes PMA language using strong-substitutes bids. Moreover, these expressions are unique (up to normalisation and removal of redundancies as incorporated in the definition of bid collections). We state our results for indivisible goods, but they can be extended to the divisible case by applying Proposition 2.4.

THEOREM 3.1: *For any substitutes concave valuation v , there exists a unique bid collection \mathcal{B} such that $D_{\mathcal{B}} = D_v$.*

COROLLARY 3.2: *For any strong-substitutes valuation v , there exists a unique bid collection \mathcal{B} such that $D_{\mathcal{B}} = D_v$, and all bids in \mathcal{B} are strong-substitutes bids.*

Conversely, we will see in Section 4 that if we have a strong-substitutes bid collection whose demand is that of a valuation, then that valuation must be strong substitutes. So it is natural to call the restriction of the substitutes PMA language to strong-substitutes bids the “strong-substitutes product-mix auction language”. We outline the proofs of Theorem 3.1 and Corollary 3.2 in Section 3.3 and give the full proofs in Appendix G.

Finally, in any auction in which sufficiently low bids are always rejected, a bidder can express her demand precisely at all relevant prices using only regular bids (those that are interested in all goods, as illustrated in Figure 2a for the two-good case), assuming she is not prepared to pay an arbitrarily large price.

COROLLARY 3.3: *For any substitutes concave valuation v , there exists a regular bid collection \mathcal{B} for any $\underline{\mathbf{p}} \in \mathcal{P}$, such that $D_{\mathcal{B}}(\mathbf{p}) = D_v(\mathbf{p})$ for all $\mathbf{p} \geq \underline{\mathbf{p}}$, if there exists $\bar{\mathbf{p}}$ such that $D_v(\mathbf{p}) = \{\mathbf{0}\}$ for all $\mathbf{p} \geq \bar{\mathbf{p}}$.*

For example, if the auctioneer’s reserve prices are (strictly) positive, then, as we noted in Section 2.2 above, substituting the regular bids $(800, 0; 1, 0; 1)$ and $(0, 500; 0, 1; 1)$ for the bids $(800, -\infty; 1, 0; 1)$, $(-\infty, 500; 0, 1; 1)$ in the bid collection that expresses Calpurnia’s true demand will change her demand only at (sufficiently) negative prices which will never arise.

3.1. The Geometry of Demand

The proofs of our main results rest on a geometric characterisation of indivisible demand introduced in Baldwin and Klemperer (2019b). We outline the key ideas here; for precise statements and further details we refer to Appendix C.

When goods are indivisible, demand is unique at generic prices and can only change by passing through a price at which there is indifference between two or more bundles. So, for any demand correspondence D arising from a valuation v , or an individual bid \mathbf{b} , or a bid collection \mathcal{B} , we define the *Locus of Indifference Prices (LIP)* $\mathcal{L} := \{\mathbf{p} \in \mathcal{P} \mid |D(\mathbf{p})| \geq 2\}$. We call the connected components of the complement of \mathcal{L} the *unique demand regions (UDRs)* of \mathcal{L} ; demand is constant within a UDR.

LIPs of valuations and bid collections can be decomposed into $(n - 1)$ -dimensional linear pieces which we call *facets*. Baldwin and Klemperer (2019b) show this for valuations, and by Lemma 2.1 it also holds for the LIP of a single bid; the LIP of a bid collection then also inherits this structure from the LIPs of its individual bids. In each panel of Figure 1, the line segments




FIGURE 4.—The UDRs and weighted LIPs associated with Antigone, Bathsheba and Calpurnia in Example 1. Antigone and Bathsheba both have three UDRs, while Calpurnia has four. The bundles demanded are shown in their corresponding UDRs. The facet weights are circled, and arrows show the normal vectors for the facets of Antigone and Bathsheba’s LIPs defined by a fixed rotational direction about the point at which the three facets meet.

are the facets, and the union of the line segments form the LIP. Antigone and Bathsheba have three UDRs, while Calpurnia has four.

We extend a LIP to a *weighted LIP* (\mathcal{L}, w) by associating each facet F of \mathcal{L} with a weight $w(F)$. For valuations v , weight $w_v(F)$ is the (positive) greatest common divisor of the coordinate entries of $\mathbf{x} - \mathbf{y}$, where \mathbf{x} and \mathbf{y} are the bundles demanded in the UDRs on either side of F . The change in demand as we cross a facet F of \mathcal{L}_v is then specified by $w_v(F)\mathbf{n}$, where \mathbf{n} is the normal to F expressed as a primitive integer vector and pointing in the opposite direction to the price change. Figure 4 illustrates this for Example 1; the figure shows the bundles in the UDRs, the facet weights, and the facet normals. For individual bids \mathbf{b} , we define $w_{\mathbf{b}} = w_{v_{\mathbf{b}}}$ if $m > 0$ and $w_{\mathbf{b}} = -w_{v_{|\mathbf{b}|}}$ if $m < 0$. It follows from Lemma 2.1 that the change in demand is similarly expressed by the facet normals and weights. This demand change property also extends to the weighted LIPs of bid collections \mathcal{B} if we define the weight $w_{\mathcal{B}}(F)$ of each facet F in $\mathcal{L}_{\mathcal{B}}$ as the sum, over all bids $\mathbf{b} \in \mathcal{B}$, of the weights $w_{\mathbf{b}}(F')$ of all facets F' in $\mathcal{L}_{\mathbf{b}}$ containing F .

The fact that demand change is governed by the facet weights and normals for all our weighted LIPs allows us to establish the following relation between concave valuations and bid collections. This plays a key role in our proofs of Theorem 3.1 and Proposition 4.3.

PROPOSITION 3.4: *Concave valuation v and bid collection \mathcal{B} satisfy $(\mathcal{L}_v, w_v) = (\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ and $D_v(\tilde{\mathbf{p}}) = D_{\mathcal{B}}(\tilde{\mathbf{p}})$ for some specific price $\tilde{\mathbf{p}} \in \mathcal{P}$ if and only if $D_v = D_{\mathcal{B}}$.*

Moreover, the net change in demand along a price path that ends where it started must be zero. So the weighted LIPs of valuations and bid collections both satisfy the following important ‘balancing condition’.

FACT 3.5—The “Balancing Condition” (Mikhalkin (2004)): Let (\mathcal{L}, w) be a weighted LIP. For any $(n - 2)$ -dimensional intersection G of two or more facets of \mathcal{L} , the weights $w(F^k)$ of the facets F^1, \dots, F^l that contain G , and primitive integer normal vectors $\mathbf{n}^1, \dots, \mathbf{n}^l$ for these facets defined by a fixed rotational direction about G , satisfy $\sum_{k=1}^l w(F^k)\mathbf{n}^k = 0$.

Figure 4 illustrates the balancing condition for Antigone and Bathsheba, with normal vectors shown as arrows. In Bathsheba’s LIP, for example, the facets meeting the LIP’s vertex at point $(1000, 700)$ are normal to $(-1, 0)$, $(1, -2)$ and $(0, 1)$, so the balancing condition states $1 \cdot (-1, 0) + 1 \cdot (1, -2) + 2 \cdot (0, 1) = 0$.

It is not hard to see that for a valuation to be substitutes, every facet’s normal vector must have at most one positive coordinate entry and at most one negative coordinate entry (with all others being 0)—if not we would be able to engineer a price change violating the substitutes condition. The converse is also clear. A valuation is strong-substitutes if and only if all the facet normal vectors have coordinates in $\{0, \pm 1\}$.²² So strong substitutes means that the tradeoffs between units of goods across facets are one-to-one.

It is straightforward that the LIP \mathcal{L}_b of a bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ interested in goods I consists of one facet normal to $t_i \mathbf{e}^i - t_j \mathbf{e}^j$ for each pair (i, j) of goods in I , so $v_{|\mathbf{b}|}$ is substitutes. Moreover, $v_{|\mathbf{b}|}$ is strong-substitutes if $\mathbf{t} = \mathbf{1}$.

We will refer to facets with normal $a\mathbf{e}^i - b\mathbf{e}^j$ for some $a, b \in \mathbb{Z}_{>0}$ as (i, j) -facets. All facets of \mathcal{L}_b and \mathcal{L}_B are (i, j) -facets for some $i, j \in [n]_0$.

3.2. Regular Valuations

Theorem 3.1 will be easiest to see for valuations which we call “regular”, which are those which can be expressed using only regular bids (see Corollary 3.7 below). Regular valuations are those for which the goods are “strictly substitutes” in the sense that the bidder will switch completely towards or away from a good if its price is sufficiently extreme. That is, for every good i , both (i) good i is not demanded if its price is high enough, and (ii) fixing the prices of all other goods, good i is the only good demanded if its price is low enough (perhaps very negative):

DEFINITION 3.6: A valuation v is *regular* if it is concave and substitutes, and if for all $i \in [n]$, both:

- (i) there exists $\bar{p}_i \in \mathbb{R}$ such that if $p_i > \bar{p}_i$ and $\mathbf{x} \in D_v(\mathbf{p})$ then $x_i = 0$,
- (ii) for all $\mathbf{p} \in \mathcal{P}$ there exists $\bar{\lambda} > 0$ such that if $\lambda > \bar{\lambda}$ then $x_j = 0$ for all $\mathbf{x} \in D_v(\mathbf{p} - \lambda \mathbf{e}^i)$ and $j \neq i$.

So Antigone and Bathsheba have regular valuations, but Calpurnia does not: if the price of a seaside holiday is below £800 then Calpurnia always demands a week at the seaside, even if the price of taking two weeks walking is arbitrarily *negative*.²³

COROLLARY 3.7: For any regular valuation v , there exists a unique bid collection \mathcal{B} such that $D_{\mathcal{B}} = D_v$, and all bids in \mathcal{B} are regular bids.

Corollary 3.7 follows largely from Theorem 3.1 and is proved in Appendix G.

We will use the following geometric characterisation of the LIPs of regular valuations.

²²For this material, see Baldwin and Klemperer (2019b, Proposition 3.6), Baldwin and Klemperer (2014, Corollary 5.20), and Shioura and Tamura (2015, Theorem 4.1(i)). The “demand types” of Baldwin and Klemperer (2019b) generalise these results by providing a taxonomy of valuations according to the facet normals of their LIPs. Baldwin et al. (2020) extends the “demand types” classification of valuations to settings with income effects, and Baldwin et al. (2021) provides further analysis (and an alternative definition) of “demand types”.

²³However, as Corollary 3.3 suggests, it is trivial to find a regular valuation whose demand precisely matches Calpurnia’s demand at all economically-relevant prices. For example, if only positive prices are economically-relevant, we can introduce the possibility that if she accepts two weeks’ holiday of either type then she has free disposal of the unwanted second week—but she cannot buy more than two weeks’ holiday in total. In this case, she would prefer getting two weeks’ walking (but just using one of the weeks) to her other options, if the price of walking were sufficiently negative, but her demand is unchanged at all positive prices.




FIGURE 5.—A price vector $\mathbf{p} \in \mathbb{R}^2$ can be adjacent to at most two horizontal and two vertical facets, and any number of diagonal facets, of a LIP. The weight labels are those defined in Section 3.3.1. The arrows illustrate a consistent fixed rotational direction for computing the balancing property.

LEMMA 3.8: *Suppose \mathcal{L}_v is the LIP of a substitutes valuation. Then it is the LIP of a regular valuation if and only if, for all $i, j \in [n]$,*

- (i) every $(i, 0)$ -facet of \mathcal{L}_v is bounded below in all coordinates;
- (ii) every (i, j) -facet of \mathcal{L}_v is bounded above in coordinates i and j .

It is not hard to show that the LIP of a regular bid satisfies these properties, and therefore that the LIP of a regular bid collection also does (see Corollary D.1).

3.3. Proving the Representation Theorem

To prove Theorem 3.1, we construct a bid collection \mathcal{B} such that $(\mathcal{L}_v, w_v) = (\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ and $D_v(\tilde{\mathbf{p}}) = D_{\mathcal{B}}(\tilde{\mathbf{p}})$ for some specific price $\tilde{\mathbf{p}} \in \mathcal{P}$ in a UDR of $\mathcal{L}_v = \mathcal{L}_{\mathcal{B}}$, for any given concave substitutes valuation v . By Proposition 3.4, $D_v(\mathbf{p})$ and $D_{\mathcal{B}}(\mathbf{p})$ then agree at all prices.

The substance of the proof is therefore the construction of the bid collection such that the weighted LIPs of v and \mathcal{B} agree. To do this, we choose bids so that the weights of specific facets of $\mathcal{L}_{\mathcal{B}}$ match the weights of the corresponding facets of \mathcal{L}_v . We then show that the balancing condition implies that the weights of all $\mathcal{L}_{\mathcal{B}}$'s other facets also match the weights of their corresponding facets in \mathcal{L}_v .


We first show how to do this for the regular two-good case, and illustrate the method in the context of a simple example, before sketching the additional arguments needed for the general case. The full proof is in Appendix G.

Corollary 3.2 follows immediately from our proof of Theorem 3.1: if v is strong-substitutes, then all facets of \mathcal{L}_v represent one-to-one tradeoffs between goods, and so the bids we construct all incorporate tradeoffs $\mathbf{t} = \mathbf{1}$.

3.3.1. The Representation Theorem for the Regular Two-Good Case

We first show that we can represent regular valuations for two goods using only regular bids, that is, only the bids illustrated in Figure 2a. For this case we can identify bids $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ using $(\mathbf{r}_{-0}; \mathbf{t}_{-0}; m)$.

To explain our procedure, we notate facet weights as follows (see Figure 5): For any $\mathbf{p} \in \mathbb{R}^2$, we let $w^l(\mathbf{p})$ and $w^r(\mathbf{p})$ be the weights of the horizontal facets containing \mathbf{p} to the left and right of \mathbf{p} ; let $w^u(\mathbf{p})$ and $w^d(\mathbf{p})$ be the weights of the upper and lower vertical facets containing \mathbf{p} ;



Bid	r_{-0}	t_{-0}	m
1	(9, 5)	(1, 2)	1
2	(4, 4)	(2, 1)	1
3	(3, 2)	(1, 2)	-1
4	(3, 2)	(2, 1)	-1
5	(3, 2)	(1, 1)	3

(a)

(b)

FIGURE 6.—(a) The LIP \mathcal{L}_v of a regular valuation v . The facets are labelled with their weights, and the bundle of goods demanded in each UDR is shown. (b) The bid collection representing v .

and let $w^{dl}(\mathbf{p}; \frac{t_1}{t_2})$ and $w^{ur}(\mathbf{p}; \frac{t_1}{t_2})$ be the weights of the diagonal facets with slope $\frac{t_1}{t_2}$ below (“down-left”) and above (“up-right”) \mathbf{p} , where we write $\frac{t_1}{t_2}$ as a fraction in lowest terms. Where a corresponding facet does not exist, we say its weight is 0. We use subscripts v , \mathbf{b} and \mathcal{B} for the weights, according to which LIP we are referring to.

Consider any regular valuation, v , on two goods 1 and 2. Figure 6 shows the LIP, \mathcal{L}_v , of an example of such a valuation. We can create the bid collection \mathcal{B} that expresses such a valuation from regular bids whose roots r_{-0} are vertices of \mathcal{L}_v . For every vertex $\mathbf{p} = (p_1, p_2)$, and for every $\frac{t_1}{t_2}$ such that a diagonal facet with slope $\frac{t_1}{t_2}$ starts or ends at \mathbf{p} , include a bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ with $r_{-0} = (p_1, p_2)$, tradeoff $\mathbf{t}_{-0} = (t_1, t_2)$, and multiplicity $m = w_v^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_v^{ur}(\mathbf{p}, \frac{t_1}{t_2})$, except if $m = 0$. (An $m = 0$ bid would be vacuous.) So $w_b^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_b^{ur}(\mathbf{p}, \frac{t_1}{t_2}) = m$ (see Figure 2), and therefore also $w_{\mathcal{B}}^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_{\mathcal{B}}^{ur}(\mathbf{p}, \frac{t_1}{t_2}) = m$, since \mathbf{b} is the only bid (if any) which influences the change in weight at \mathbf{p} along the diagonal with slope $\frac{t_1}{t_2}$. The bids for our example are those given on the right in Figure 6.

We can now show that $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}}) = (\mathcal{L}_v, w_v)$. First consider diagonal facets. Trace down any (doubly-infinite) diagonal line that contains a facet in either \mathcal{L}_v or $\mathcal{L}_{\mathcal{B}}$. At high enough prices, the weights of both LIPs’ facets along this line are 0; that is, neither LIP has a facet because 0 is demanded everywhere at such prices, as valuation v is regular (recall Definition 3.6 (i)). And by construction, there is a bid in \mathcal{B} whose tradeoff matches the slope of this line wherever the weights of the facets in this line changes; moreover this bid is such that the weights change in both LIPs in the same way. So the weights of the facets in $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ match those in (\mathcal{L}_v, w_v) all along the line.

For example, in Figure 6, as we trace down the line with slope $\frac{1}{2}$ through $(7, 4)$, we first have no facet (above $(9, 5)$), then two consecutive facets between $(9, 5)$ and $(3, 2)$, both with weight 1 due to bid 1, and then no facet again (below $(3, 2)$), as bid 3 with multiplicity -1 cancels out the facet of bid 1.

The key to our proof is that the balancing condition (Fact 3.5) implies that any regular LIP is fully determined by its diagonal facets. (Figure 2a provides an elementary illustration.) So, because our bids’ diagonal facets were specified to match the valuation’s diagonal facets, (\mathcal{L}_v, w_v) and $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ must match everywhere.

In detail, for the facet weights and normals (in counterclockwise direction) of any LIP around \mathbf{p} , as illustrated in Figure 5, the balancing condition tells us that

$$\begin{aligned} w^u(\mathbf{p}) \begin{pmatrix} -1 \\ 0 \end{pmatrix} + w^l(\mathbf{p}) \begin{pmatrix} 0 \\ -1 \end{pmatrix} + \sum_{\frac{t_1}{t_2} \in \mathcal{S}} w^{dl}(\mathbf{p}, \frac{t_1}{t_2}) \begin{pmatrix} t_1 \\ -t_2 \end{pmatrix} \\ + w^d(\mathbf{p}) \begin{pmatrix} 1 \\ 0 \end{pmatrix} + w^r(\mathbf{p}) \begin{pmatrix} 0 \\ 1 \end{pmatrix} + \sum_{\frac{t_1}{t_2} \in \mathcal{S}} w^{ur}(\mathbf{p}, \frac{t_1}{t_2}) \begin{pmatrix} -t_1 \\ t_2 \end{pmatrix} = 0, \end{aligned}$$

where \mathcal{S} is the set of all slopes of diagonal facets at \mathbf{p} . The first coordinate of this equation tells us $w^u(\mathbf{p}) - w^d(\mathbf{p}) = \sum_{\frac{t_1}{t_2} \in \mathcal{S}} (w^{dl}(\mathbf{p}; \frac{t_1}{t_2}) - w^{ur}(\mathbf{p}; \frac{t_1}{t_2})) t_1$, for the weight functions of v and \mathcal{B} .

So since, for every slope $\frac{t_1}{t_2}$, we already know that $w_{\mathcal{B}}^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_{\mathcal{B}}^{ur}(\mathbf{p}, \frac{t_1}{t_2}) = w_v^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_v^{ur}(\mathbf{p}, \frac{t_1}{t_2})$, it follows that $w_{\mathcal{B}}^u(\mathbf{p}) - w_{\mathcal{B}}^d(\mathbf{p}) = w_v^u(\mathbf{p}) - w_v^d(\mathbf{p})$.

Consider, for example, the three bids at $(3, 2)$ in Figure 6. As illustrated in Figure 2a, bid 3 gives a vertical facet of weight $t_1 m = -1$ above $(3, 2)$, and bid 4 gives a vertical facet of weight $t_1 m = -2$ here. But bid 5's vertical facet weight is $t_1 m = 3$. So these facets all cancel in $\mathcal{L}_{\mathcal{B}}$ and so $w_{\mathcal{B}}^u(\mathbf{p}) = w_{\mathcal{B}}^d(\mathbf{p}) = 0$, giving the same 0 difference in weight here as on the LIP \mathcal{L}_v .

So tracing up from the bottom of any vertical line containing a facet of either \mathcal{L}_v or $\mathcal{L}_{\mathcal{B}}$, the weight changes are the same in both LIPs. Since (by our regularity assumption and Lemma 3.8) neither LIP has a vertical facet at low enough prices, it follows that their weighted vertical facets are identical.

The argument for the horizontal facets of $\mathcal{L}_{\mathcal{B}}$ and \mathcal{L}_v is the same as for the vertical ones; the fact that any changes in the horizontal facets' weights match at any \mathbf{p} , i.e., $w_{\mathcal{B}}^r(\mathbf{p}) - w_{\mathcal{B}}^l(\mathbf{p}) = w_v^r(\mathbf{p}) - w_v^l(\mathbf{p})$, follows from the second coordinate of the balancing equation.

So $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}}) = (\mathcal{L}_v, w_v)$. And since both our bid collection and the valuation are regular, both demand $\mathbf{0}$ at high enough prices, so $D_{\mathcal{B}}(\mathbf{p}) = D_v(\mathbf{p})$ at all prices by Proposition 3.4.²⁴

As \mathcal{B} is parsimonious and normalised by construction, it is a bid collection. We now show that \mathcal{B} is the unique bid collection such that $D_{\mathcal{B}}(\mathbf{p}) = D_v(\mathbf{p})$ at all prices: suppose that $D_{\mathcal{B}'} = D_v$ for some bid collection \mathcal{B}' . Then $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}}) = (\mathcal{L}_{\mathcal{B}'}, w_{\mathcal{B}'})$ by Proposition 3.4, and so $w_{\mathcal{B}}^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_{\mathcal{B}}^{ur}(\mathbf{p}, \frac{t_1}{t_2}) = w_{\mathcal{B}'}^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_{\mathcal{B}'}^{ur}(\mathbf{p}, \frac{t_1}{t_2})$ for all possible \mathbf{p} and \mathbf{t} . But we saw above that $w_{\mathcal{B}}^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_{\mathcal{B}}^{ur}(\mathbf{p}, \frac{t_1}{t_2}) = w_{\mathbf{b}}^{dl}(\mathbf{p}, \frac{t_1}{t_2}) - w_{\mathbf{b}}^{ur}(\mathbf{p}, \frac{t_1}{t_2})$ where \mathbf{b} is the unique bid in \mathcal{B} with root $\mathbf{r} = (0, p_1, p_2)$, tradeoff \mathbf{t} , and multiplicity given by this difference; if no such bid exists, then this difference must be zero. So the existence and multiplicity of such bids must match across \mathcal{B} and \mathcal{B}' , so $\mathcal{B}' = \mathcal{B}$.

3.3.2. The Representation Theorem for the General Case

We begin with the case of a regular valuation, v . The LIP of a valuation for n goods potentially has many more facets than that of a two-good valuation. (Figure 7 illustrates a bid in the $n = 3$ case.) However, the complexity of our problem is greatly reduced by the fact that the LIPs of substitutes valuations and bid collections only have facets that are normal to vectors

²⁴Note that it was important to start by considering diagonal facets, because we needed to know the tradeoff of each bid, and only the diagonal facets tell us this. (Starting by looking at horizontal or vertical facets would not have worked: doing so would have suggested that there is no bid at $(3, 2)$, for example.) An exception is if the valuation is strong-substitutes, in which case every bid has $\mathbf{t}_{-0} = (1, 1)$, so either the horizontal facets, or the vertical ones, will tell us all we need to know.




FIGURE 7.—The LIP of a bid with root $\mathbf{r}_{-0} = (3, 6, 2)$, tradeoffs $\mathbf{t}_{-0} = (5, 6, 4)$ and multiplicity $m = 1$ divides price-space into four 3-dimensional UDRs. Demand is $(5, 0, 0)$, $(0, 6, 0)$, $(0, 0, 4)$, and $(0, 0, 0)$, in the UDRs that are to the left of the left vertical facet, in front of the right vertical facet, below the horizontal facet, and at higher prices than any of these three facets, respectively. The facets' weights are shown in the circles. Their normal vectors are shown as arrows; they are $(0, 0, 1)$, $(1, 0, 0)$, and $(0, -1, 0)$, for the *non*-diagonal facets (with weights 4, 5, and 6, respectively), and $(-5, 0, 4)$, $(-5, 6, 0)$, and $(0, -6, 4)$ for the diagonal facets (from left to right in the view shown).

with at most two non-zero entries (see Section 3.1). This, together with the balancing condition, will mean that we only need to find bids matching the facets in a *single* diagonal (i, j) orientation in order to find bids that match the facets in all the other orientations.

We fix two goods $i, j \in [n]$ and construct a bid collection \mathcal{B} so that the weighted LIPs of v and \mathcal{B} agree on all diagonal (i, j) -facets for this particular choice of i and j . In the two-good case, the multiplicity of a bid at \mathbf{p} with tradeoffs (t_1, t_2) is the difference in weight between the upper and lower facet along the line through \mathbf{p} with slope $\frac{t_1}{t_2}$, i.e., in direction $(\frac{1}{t_1}, \frac{1}{t_2})$ (this is the diagonal line in Figure 2a). In higher dimensions, the area of interest becomes a hyperplane neighbourhood around \mathbf{p} , which can contain many facets; the multiplicity of a bid at \mathbf{p} with tradeoffs \mathbf{t} is the weighted sum of weights of the facets meeting at \mathbf{p} that contain part of the line through \mathbf{p} in direction $(\frac{1}{t_1}, \dots, \frac{1}{t_n})$. (In Figure 7, this is the diagonal line along which all three diagonal facets meet.)

In the $n = 2$ case, we picked a specific diagonal line and moved along it to show that the weighted LIPs of v and \mathcal{B} agree on all the facets of the line. In our more general setting, we now pick a specific (i, j) -hyperplane and sweep across it, visiting vertices in an appropriate order, to show that the weighted LIPs of v and \mathcal{B} are identical on all (i, j) -facets in this hyperplane, and hence all hyperplanes for our fixed i, j .

We then use the balancing condition to show that we would have got the same bid collection if we had started with a different diagonal orientation than (i, j) . So the weighted LIPs of v and \mathcal{B} have the same weighted facets for all diagonal orientations. Moreover, the balancing condition also tells us that the $(i, 0)$ -facets of both weighted LIPs match for all $i \in [n]$.

When valuations are not regular, we start in exactly the same way by creating regular bids at each vertex of \mathcal{L}_v . However, in this case, the construction fails to account for the fact that \mathcal{L}_v has diagonal (i, j) -facets which are unbounded *above* in coordinates i and j and/or $(i, 0)$ -facets which are unbounded *below* (cf. Lemma 3.8). Because the facets of regular bids *are* bounded

in these ways, they cannot depict such facets. For example, a valuation that is interested only in good i and the null good has a LIP consisting of a single (unbounded) $(i, 0)$ -hyperplane (as illustrated in Figure 2b or 2c for the two-good case) but has no vertex, so we have thus far have generated no bid(s) to match it.

We therefore need to add some non-regular bids to our construction. To identify these, we draw a “bounding box” which is a hyperrectangle large enough that its interior contains at least part of the interior of every facet of \mathcal{L}_v . The $(i, 0)$ -facets which are unbounded below, and (i, j) -facets which are unbounded above in coordinates i and j , must intersect the boundary of this box. So at every vertex at which \mathcal{L}_v and the bounding box intersect, we add non-regular bids so that the weighted LIPs also agree on all the unbounded facets. (In the two-good case, these are bids of the kind illustrated in Figure 2b for $(1, 0)$ -facets, Figure 2c for $(2, 0)$ -facets, and Figure 2d for $(1, 2)$ -facets.) Appendix A.1 gives a detailed example of adding non-regular bids.

The weights of the LIPs of the valuation and the bid collection now match everywhere. It remains to adjust demand globally to the correct levels. To do this, we fix any price in a unique demand region of \mathcal{L}_v . Then, for each good i , add a bid interested only in i , with multiplicity equal the number of units of i demanded by v , minus the number of units of i demanded by our current bid collection, at that price.

The final steps of the proof are as for the regular two-good case. Appendix A.2 provides a more detailed sketch of the proof of the general case; the full details are in Appendix G.

4. VALUATIONS FROM BIDS

Theorem 3.1 showed that any substitutes valuation corresponds to a collection of bids. Conversely, we call a collection \mathcal{B} *valid* if there is a valuation v with $D_v = D_{\mathcal{B}}$. We will see that all bid collections of positive bids are valid, but not all bid collections that include negative bids are valid. (Calpurnia’s bid collection includes negative bids and *is* valid.) We develop multiple characterisations of validity in Proposition 4.3.

We now show that the class of valid bid collections in the substitutes PMA bidding language corresponds *exactly* to the class of concave substitutes valuations, and analogous equivalences hold for strong-substitutes and regular valuations.

Note that adding a constant to the valuation v (including to the value of the zero bundle, $v(\mathbf{0})$, if $\mathbf{0}$ is in v ’s domain) does not change its demand, D_v . Conversely, any two valuations have identical demands only if their normalisations are the same (Mikhalkin 2004, Proposition 2.1). So we normalise valuations by fixing the value of any consistently chosen bundle from their domains. We already fixed a normalisation of bids, but we could have made any other choice consistent with Lemma 2.2. Now:

COROLLARY 4.1: *Fix any normalisations for valuations and bids.*

- (i) *Valid bid collections \mathcal{B} correspond one-to-one to concave substitutes valuations v , such that $D_{\mathcal{B}} = D_v$.*
- (ii) *Valid strong-substitutes bid collections \mathcal{B} correspond one-to-one to strong-substitutes valuations v , such that $D_{\mathcal{B}} = D_v$.*
- (iii) *Valid regular bid collections \mathcal{B} correspond one-to-one to regular valuations v , such that $D_{\mathcal{B}} = D_v$.*

In each case, given a valuation, we obtain a bid collection satisfying the specified properties and unique up to normalisation, by Theorem 3.1 and Corollaries 3.2 and 3.7, respectively. Conversely, validity of a bid collection means that it defines a valuation, and this is concave (as we noted below Equation (2.1)) and unique up to normalisation (see the previous paragraph). Seeing that this valuation has the specified properties is easy using the geometric tools

of Section 3.1. For (i), every valid bid collection \mathcal{B} is associated with the valuation v satisfying $D_{\mathcal{B}} = D_v$ by definition, so $\mathcal{L}_{\mathcal{B}} = \mathcal{L}_v$. The facets of each bid's LIP are (i, j) -facets, so the same holds for $\mathcal{L}_{\mathcal{B}}$, implying that v is substitutes. For (ii), we add that if \mathcal{B} is a strong-substitutes bid collection, then the normals of each bid's LIP facets are normal to $e^i - e^j$ for some $i, j \in [n]_0$, so the same holds for $\mathcal{L}_{\mathcal{B}} = \mathcal{L}_v$ and v is strong-substitutes. For (iii), if \mathcal{B} is a regular bid collection, its LIP satisfies the properties of Lemma 3.8 (see discussion directly below that lemma), so v is regular by Lemma 3.8. This proves Corollary 4.1.

4.1. Explicit Valuations

We can specify the valuation $v_{\mathcal{B}}$ of a valid bid collection \mathcal{B} as follows. For any bundle \mathbf{x} that is uniquely demanded by \mathcal{B} at some prices, pick any such specific price \mathbf{p} at which each bid $\mathbf{b} = (r^{\mathbf{b}}, t^{\mathbf{b}}, m^{\mathbf{b}}) \in \mathcal{B}$ uniquely demands some good $i^{\mathbf{b}}$ and define

$$v_{\mathcal{B}}(\mathbf{x}) := \sum_{\mathbf{b} \in \mathcal{B}} m^{\mathbf{b}} t_{i^{\mathbf{b}}}^{\mathbf{b}} r_{i^{\mathbf{b}}}^{\mathbf{b}}.$$

We can extend this definition to a (concave) valuation on the set of all bundles demanded by \mathcal{B} (at any price): if \mathbf{x} is (non-uniquely) demanded by \mathcal{B} , fix prices \mathbf{p} at which \mathbf{x} is demanded and perturb them generically and infinitesimally to get any prices \mathbf{p}' at which every bid's demand is unique. The bundle \mathbf{x}' that \mathcal{B} demands at \mathbf{p}' is also demanded at \mathbf{p} , and we define $v_{\mathcal{B}}(\mathbf{x}) := v_{\mathcal{B}}(\mathbf{x}') + \mathbf{p} \cdot (\mathbf{x} - \mathbf{x}')$. The fact that $v_{\mathcal{B}}$ is well-defined and satisfies $D_{\mathcal{B}} = D_{v_{\mathcal{B}}}$ is part (iii) of Proposition 4.3.

Alternatively, Lemma F.3 and Proposition F.6 show that we can write $v_{\mathcal{B}}$ in terms of the valuation functions $v_{\mathbf{b}}$ associated with each bid $\mathbf{b} \in \mathcal{B}$ in Lemma 2.1.

4.2. Other Characterisations of Validity

Proposition 4.3 shows that \mathcal{B} is valid if and only if $D_{\mathcal{B}} = D_{v_{\mathcal{B}}}$ for the valuation $v_{\mathcal{B}}$ from Section 4.1. It also gathers further equivalent characterisations of validity of bid collections.

In standard economic contexts, without Giffen goods, a “law of demand” holds; that is, an increase in price for any one good will (weakly) reduce demand for that good (see Definition 4.2). In particular, this holds for the demand correspondence of any valuation. So if $D_{\mathcal{B}} = D_v$ holds for some valuation v , then $D_{\mathcal{B}}$ also satisfies the law of demand. We show in Proposition 4.3 that this is in fact the *only* condition \mathcal{B} needs to satisfy: \mathcal{B} is valid if and only if $D_{\mathcal{B}}$ satisfies the law of demand.

DEFINITION 4.2: A demand D satisfies the *law of demand* if, for any two prices \mathbf{p} and $\mathbf{p}' := \mathbf{p} + \lambda e^i$ with $\lambda \geq 0$ and $i \in [n]$ that satisfy $D(\mathbf{p}) = \{\mathbf{x}\}$ and $D(\mathbf{p}') = \{\mathbf{x}'\}$, we have $x'_i \leq x_i$ and equality holds if and only if $\mathbf{x} = \mathbf{x}'$.

Note that it follows that not all bid collections that include negative bids are valid. For example, a single negative bid on its own violates the law of demand—see any panel of Figure 2 with $m < 0$ —so corresponds to no valuation. However, all positive bid collections are valid, as will become clear shortly.

We can also characterise validity using the indirect utility functions of v and \mathcal{B} defined, respectively, by $\pi_v(\mathbf{p}) := \max_{\mathbf{x} \in X} (v(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x})$ and

$$\pi_{\mathcal{B}}(\mathbf{p}) := \sum_{(r; t; m) \in \mathcal{B}} m \max_{i \in [n]_0} t_i (r_i - p_i). \quad (4.1)$$

We show that \mathcal{B} is valid if and only if $\pi_{v_{\mathcal{B}}} = \pi_{\mathcal{B}}$. Indeed, we need only impose a weaker condition on $\pi_{\mathcal{B}}$. We show that \mathcal{B} is valid if and only if $\pi_{\mathcal{B}}$ is convex—a property necessarily satisfied by the indirect utility function of any valuation.

We also show that the demand of a valid bid collection \mathcal{B} is the Minkowski difference between the demand of \mathcal{B} 's positive bids, \mathcal{B}^+ , and the demand of \mathcal{B} 's negative bids with the absolute value taken for their multiplicities, $|\mathcal{B}^-|$. The converse also holds, which provides an additional characterisation of validity. (The Minkowski difference $X - Y$ of two sets $X, Y \subseteq \mathcal{P}$ is defined as $\{\mathbf{p} \in \mathcal{P} \mid \mathbf{p} + Y \subseteq X\}$.)

Finally, we show that \mathcal{B} is valid if and only if every facet of its weighted LIP ($\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}}$) has a positive weight.

PROPOSITION 4.3: *For any bid collection \mathcal{B} , the following statements are equivalent.*

- (i) *Bid collection \mathcal{B} is valid, i.e., there exists a valuation v such that $D_v = D_{\mathcal{B}}$.*
- (ii) *There exists a concave substitutes valuation v such that $D_v = D_{\mathcal{B}}$.*
- (iii) *The valuation $v_{\mathcal{B}}$ is well-defined and satisfies $D_{v_{\mathcal{B}}} = D_{\mathcal{B}}$.*
- (iv) *$D_{\mathcal{B}}$ satisfies the law of demand.*
- (v) *The valuation $v_{\mathcal{B}}$ satisfies $\pi_{v_{\mathcal{B}}} = \pi_{\mathcal{B}}$.*
- (vi) *The indirect utility function $\pi_{\mathcal{B}}$ of \mathcal{B} is convex.*
- (vii) *For all $\mathbf{p} \in \mathcal{P}$, we have $D_{\mathcal{B}}(\mathbf{p}) = D_{\mathcal{B}^+}(\mathbf{p}) - D_{|\mathcal{B}^-|}(\mathbf{p})$.*
- (viii) *The weight $w_{\mathcal{B}}(F)$ of every facet of $\mathcal{L}_{\mathcal{B}}$ is positive.*

(Statements (v) and (vi) are generalisations of [Baldwin et al. 2023a](#), Theorem 1 and (vii) generalises a result in [Baldwin et al. 2024](#), all originally just in the strong substitutes case.)

It is now immediate that *any* positive bid collection \mathcal{B} is valid: (vii) trivially holds when $\mathcal{B} = \mathcal{B}^+$. (An easy alternative argument is that the weight $w_{\mathcal{B}}(F)$ of any facet F of $\mathcal{L}_{\mathcal{B}}$ equals the sum of the weights of the facets of the bids in $\mathcal{L}_{\mathcal{B}}$ that contain F (Section 3.1). Since any positive bid's facets' weights are all positive (see Section 3.1), $w_{\mathcal{B}}(F)$ is also positive, so \mathcal{B} is valid by (viii) of Proposition 4.3.)

5. IMPLEMENTING PRODUCT-MIX AUCTIONS

5.1. Constructing Bid Collections

A bidder who knows the weighted LIP of her valuation can find the unique bid collection representing her valuation by applying the constructive procedure that underpins the proof of Theorem 3.1 (see Algorithm 1 in Appendix G). And the mathematical software tool Gfan ([Jensen, 2024](#)) can find the weighted LIP if the bidder knows her valuation function.

In practice a bidder may find it hard to articulate her entire valuation function. However, PMA languages allow bidders to construct complex valuations by combining smaller structures. This can be especially helpful for, for example, an organisation such as a commercial bank whose demand in the Bank of England's PMA for loans (see Section 5.4) may be the combination of the demands of several separate divisions. Moreover, even a simpler bid structure (such as the demand of a single division) can be pieced together, bit by bit, from the answers to straightforward questions.

For example, a bidder might know that she wants up to two units in total of two goods on offer. So she knows the configuration of her bids is of the form shown in Figure 8, with three positive bids, A, B , and C , and one negative bid, D . So she could proceed by identifying the four separate price vectors at each of which she is indifferent between three alternative bundles. But she could instead ask: if p_2 were high enough that I would never buy good 2, what is the maximum price, p_1' , that I would pay for a single unit of good 1, and what maximum per-unit




FIGURE 8.—Example of strong-substitutes preferences expressed by three positive and one negative bid. The roots of the positive bids A , B and C are drawn solid, and the root of the negative bid D is hollow.

price, p_1'' , would I pay for two units of good 1? (thus identifying the first coordinates of bids A and B). Also, if p_1 were very high, what maximum per-unit prices, p_2' , and p_2'' , would I pay for one unit, and two units, of good 2, respectively? (thus identifying the second coordinates of bids A and C). And finally, if both goods' prices were low, at what relative prices $p_1 - p_2 = \tilde{p}_1$ would I switch from demanding $(0, 2)$ to $(1, 1)$?, and at what $p_1 - p_2 = \tilde{p}_1$ would I switch from demanding $(1, 1)$ to $(2, 0)$? (thus pinning down the remaining coordinates of the bids²⁵). This elicits the required information in a way that may be much easier for the bidder.

Quite generally, it suffices for a bidder to be able to answer simple questions of the type “At the following prices, what bundle of goods do you demand?”. For this “demand oracle” case, [Goldberg et al. \(2022\)](#) present algorithms to generate strong-substitutes bid collections by making a series of demand queries. Moreover, [Baldwin et al. \(2023a, Appendix C2\)](#) presents a simple algorithm to test for validity of a strong-substitutes bid collection.²⁶ We expect that [Theorem 3.1](#) and [Proposition 4.3](#) can be used to extend these algorithms to substitutes preferences more generally.

5.2. Expressing The Auctioneer's Preferences

An important feature of the PMA is that it allows the auctioneer to choose how the quantities sold depend upon expressed demand. The auctioneer can express her preferences in the same way as any other bidder who is a seller (see [Section 2.3](#)).

We can write $S(\mathbf{p})$ for the auctioneer's supply, and choose a vector \mathbf{s} such that s_i is the maximum number of units of good i that the auctioneer would sell at any price \mathbf{p} (i.e., $s_i = \max_{\mathbf{p}} S_i(\mathbf{p})$). Then offering supply $S(\mathbf{p})$ is equivalent to always selling a fixed supply \mathbf{s} , but *buying back* the demand of a bidder whose demand is $D(\mathbf{p}) := -S(\mathbf{p}) + \mathbf{s}$. In the next subsections we will think of the auctioneer in just this way, as offering a fixed supply bundle, \mathbf{s} , and also acting as an additional bidder who buys back goods to ensure that the “correct quantity” is sold to the remaining bidders at any price vector.²⁷

²⁵See [Figure 8](#): writing r_i^J for the i th coordinate of bid J , we have $r_1^A = p_1'$, $r_1^B = p_1''$, $r_2^A = p_2'$, $r_2^C = p_2''$, $r_1^C - r_2^C = \tilde{p}_1$, and $r_1^B - r_2^B = \tilde{p}_1$, which allows us to solve for $r_1^C (= r_1^D)$ and $r_2^B (= r_2^D)$. In non-generic cases D can coincide with one of A , B and C , in which case there are just two positive bids, and no negative bid.

²⁶Given a fixed number of negative bids, the algorithm tests for validity in time polynomial in the total number of bids and goods. An implementation of this algorithm is at <http://pma.nuff.ox.ac.uk>.

²⁷For example, an auctioneer who wishes to sell Q units in total of two goods, and to sell good 2 if and only if $p_2 \geq p_1 + \tilde{r}$, can do this by choosing $s_1 = s_2 = Q$ and also making the bid illustrated in [Figure 2d](#) (with $r_2 = \tilde{r}$,

An auctioneer may also wish to set a reserve price for each good. This is usually most easily implemented via restrictions on bids and clearing prices, but can alternatively be achieved by making appropriate buyback bids.

5.3. Determining the auction allocation and price vector

The auction's *allocations* are chosen to be consistent with competitive equilibrium, given the supplies and demands expressed by the bids, where possible. (If the auctioneer's objective is profit maximisation, it can achieve this by choosing a supply that does not reflect its actual preferences.²⁸)

We also assume the auction's *prices* are the competitive equilibrium prices, given the supplies and demands expressed by the bids, where these exist. ("Pay-your-bid" or other pricing rules can be used instead if preferred, though this would give even "small" bidders incentives to deviate from truthful reporting of preferences.) If multiple price vectors support competitive equilibrium, we prefer to choose the lowest one to minimise bidders' incentives to distort their preferences.²⁹

Competitive equilibrium always exists in a substitutes PMA with divisible goods. Moreover, a merit of the strong-substitutes PMA is that equilibrium continues to exist even when goods are indivisible (under mild assumptions, such as those we make in the next paragraph)—see Proposition F.2. A fortiori, competitive equilibrium also exists in simpler cases such as the Bank of England's PMA described below. Beyond strong-substitutes, however, it is well-known that equilibrium may fail to exist.³⁰ So, if there is no equilibrium for a general substitutes PMA with indivisible goods, we choose an outcome close to a pseudo-equilibrium (that is, what would be an equilibrium if we treated goods as divisible, cf. Milgrom and Strulovici (2009)). This may result in some bidders getting less than they want at the auction prices, or alternatively a change in the auctioneer's supply.

Solving the PMA is straightforward when all bids are positive. Suppose the auctioneer sells a fixed supply bundle s and makes a collection of (positive) buyback bids, \mathcal{B}^0 , that demands s at some sufficiently low price. This allows the auctioneer to express reserve prices for each good and any further preferences (cf. Section 5.2). Let $\mathcal{B}^1, \dots, \mathcal{B}^m$ be the collections of bids made by the bidders, $[m]$, and \mathcal{B} be the set of all the bids made in the PMA, that is, $\mathcal{B}^0 \cup \dots \cup \mathcal{B}^m$. We assume all bids demand $\mathbf{0}$ at sufficiently high prices. In our setting, a competitive equilibrium allocation of bundles x^0, \dots, x^m to the bid collections $\mathcal{B}^0, \mathcal{B}^1, \dots, \mathcal{B}^m$ is exactly one which maximises welfare, that is, the sum of the utilities of the auctioneer and the bidders, subject to these allocations summing to supply s . So we can use mathematical programming to find this allocation.

$t_1 = t_2 = 1$, and $m = Q$) to buy back Q units of good 1 if $p_2 \geq p_1 + \tilde{r}$ and of good 2 otherwise. (Cf. the related bid in Section 2.3; bids of this kind are used in the Bank of England's PMA, as discussed in Section 5.4.)

Note that $D(p)$ is (strong) substitutes if $-S$ is.

²⁸In particular, the auctioneer can choose its supply after seeing the other bids (assuming it chooses rules to permit itself to do this). Assuming the auction implements the resulting equilibrium prices, this will not affect the fact that "small" bidders, at least, have little incentive to distort their expressed preferences.

²⁹Since preferences are substitutes, there is a unique price vector, all of whose entries are the lowest among any equilibrium price vectors (see Milgrom and Strulovici 2009, Proposition 3); it is not hard to find this price vector in practical applications.

³⁰See Milgrom and Strulovici (2009) and Baldwin and Klemperer (2019b) for examples with substitutes bidders without an equilibrium. Baldwin and Klemperer (2019b) use the geometric characterisation of demand described in Section 3.1 to develop conditions under which equilibrium is guaranteed. Baldwin et al. (2020, 2021, 2023b) offer extensions of equilibrium existence guarantees to non-transferable utility.

If goods are divisible, the PMA with positive bids is solved by the following linear program (LP) in which \mathbf{y}^b represents the allocation to bid b .³¹ Writing $\mathbf{b} = (r^b; t^b; m^b)$, and I_b for the set of goods in which b is interested, the program is:

$$\begin{aligned}
\max \quad & \sum_{j \in [m]_0} \sum_{b \in \mathcal{B}^j} r^b \cdot \mathbf{y}^b \\
\text{s.t.} \quad & \sum_{i \in I_b} \frac{1}{t_i^b} y_i^b = m^b, \quad \forall b \in \mathcal{B} && \text{(bid demand constraints),} \\
& \sum_{b \in \mathcal{B}} y_i^b = s_i, \quad \forall i \in [n] && \text{(supply constraints),} \\
& \mathbf{x}^j = \sum_{b \in \mathcal{B}^j} \mathbf{y}^b, \quad \forall j \in [m]_0 && \text{(bidder allocations),} \\
& \mathbf{y}^b \geq \mathbf{0}, \quad \forall b \in \mathcal{B}.
\end{aligned}$$

For divisible goods, the solution(s) of this LP and the corresponding shadow prices in its dual are guaranteed to exist and are competitive equilibrium allocations and prices. See Corollary F.4 for details.

With indivisible goods (and positive bids), the LP solution is a solution of the PMA only if the allocations to bidders are integer bundles. However, we can find such a solution if it exists (as it always does for strong substitutes, even though not in general) by also solving the mixed-integer linear program (MILP) obtained by restricting the x^j in the LP to integer values. A solution to the MILP is an equilibrium of the indivisible PMA if and only if it achieves the same welfare as the solution of the original LP (cf. Bikhchandani and Mamer (1997)). If no equilibrium exists, the LP solution forms a pseudo-equilibrium which we can use to find a close outcome with integer bundles. In practice, we find that reasonably-sized auctions can easily be solved using commercial MILP solvers. The strong-substitutes PMA can be solved even more efficiently by formulating the LP as a minimum-cost network flow problem and applying the network simplex algorithm. (Corollary F.5 gives details.)

When all the bids are positive, the LP's solution automatically allocates goods to a set of bids that would choose them at the shadow price vector in the solution. But if there are also negative bids, our simple program will not in general handle them correctly.³² So we need to proceed a little differently: if the auction wants to allocate \mathbf{s} in total, as above, then it must allocate $\mathbf{s} + \mathbf{z}$ to positive bids, and $-\mathbf{z}$ to negative bids, for some \mathbf{z} . The value of the solution to the LP above for supply $\mathbf{s} + \mathbf{z}$ when only the positive bids are included is $v_{\mathcal{B}^+}(\mathbf{s} + \mathbf{z})$, and the value of the solution to the LP above for supply \mathbf{z} when only the negative bids are included, but using the absolute values of their multiplicities, is $v_{|\mathcal{B}^-|}(\mathbf{z})$ (see Lemma F.3). So solving our problem requires finding the correct value of \mathbf{z} , and the total welfare for our problem will be $v_{\mathcal{B}^+}(\mathbf{s} + \mathbf{z}) - v_{|\mathcal{B}^-|}(\mathbf{z})$ for this \mathbf{z} . But we know that the shadow price vectors in the two programs must be the same at the auction solution (so that the two programs are selecting consistent sets of bids). And this in fact means that the solution must be a stationary point of $v_{\mathcal{B}^+}(\mathbf{s} + \mathbf{z}) - v_{|\mathcal{B}^-|}(\mathbf{z})$ with respect to \mathbf{z} (and the corresponding shadow prices are

³¹See Corollary F.4, Baldwin and Klempner (2019a) and Baldwin et al. (2024). In implementations we replace any $-\infty$ entries in bids' roots by large negative values. We also allow y_0^b to take arbitrary positive values (unlike the convention for bundles in Section 2.1)—the quantities of the null good allocated clearly have no significance.

³²Allocations to negative bids reduce the objective, so the LP would wrongly allocate (true) goods to negative bids with low value roots, while not allocating them to negative bids with higher value roots.

the competitive equilibrium prices). Proposition F.6 shows this for the general substitutes case, and Baldwin et al. (2024) present a method for finding the correct z , and thereby solving, the strong-substitutes case.

Baldwin et al. (2023a) provides an alternative method of finding the equilibrium prices in the strong-substitutes case. This method takes advantage of the special structure of strong-substitutes preferences to use a variant of standard steepest descent methods for discrete functions (following Kelso and Crawford (1982), Milgrom (2000), Gul and Stacchetti (2000), Ausubel (2006) and Murota et al. (2013, 2016)). However, it achieves greater efficiency by ‘leaping’ between boundaries of the UDRs of the aggregate demand’s LIP, instead of taking unit steps. This method is proved to be polynomial-time.

Baldwin et al. (2024) tests a range of examples and shows that neither of these two methods for finding equilibrium prices (in the strong substitutes case) consistently dominates the other in practice. We expect both methods can be extended to general substitutes PMAs.

Finding equilibrium entails finding allocations as well as prices, of course. If many bids are interested in more than one good (including the null good) at the equilibrium prices, this creates a possibly-complicated tie-breaking problem. When all bids are positive, a solution arises automatically as a solution to the LP. Baldwin et al. (2023a) offers a polynomial-time algorithm that solves the general (positive and negative bids) strong-substitutes case.

5.4. The Bank of England’s Product-Mix Auction

The “Indexed Long-Term Repo” auction that the Bank of England (“the Bank”) uses to sell loans to financial institutions was originally a special case of a strong-substitutes PMA, and it remains close to a special case. Its implementation is now based on a linear program that is related to the one described in the previous subsection.

The Bank auctions n “vertically differentiated” goods, so good j is always valued more highly than good $j - 1$.³³ Bidders were initially permitted to make collections of positive strong-substitutes bids. Although Klemperer (2008) gave an example of the use of negative bids, and also mentioned the possibility of allowing general substitutes bids, he did not recommend implementing these initially, since preferences requiring them seemed unlikely to be common.³⁴

The Bank also followed Klemperer (2008)’s proposal to describe its own preferences for the quantities, $q_j(\mathbf{p})$, it wishes to supply in terms of a set of “supply functions” that are easy to represent graphically. This helps both the Bank and its bidders understand the auction. The Bank initially chose supply functions that specified that $\sum_{i \geq j} q_i(\mathbf{p})$ was a function only of the price difference $p_j - p_{j-1}$ (for $2 \leq j \leq n$). It also initially specified a maximum total quantity Q (i.e., $\sum_i q_i(\mathbf{p}) \leq Q$), and that all bids should meet or exceed reserve prices on the respective goods.

This “supply functions” approach is equivalent to Section 5.2’s approach of representing preferences as a vector of fixed supplies, \mathbf{s} , plus buyback bids. For example, if $n = 2$, the (only) supply function is $q_2(\mathbf{p}) = f(p_2 - p_1)$ for some increasing function f , and we also have $q_1(\mathbf{p}) + q_2(\mathbf{p}) \leq Q$. That is, the Bank would sell up to Q units in total of two goods, and (if prices exceed the reserves) sell $\geq q_2$ units of good 2 if $p_2 \geq p_1 + f^{-1}(q_2)$. This corresponds, in Section 5.2’s description, to choosing $s_1 = s_2 = Q$ and making Q separate buyback bids, in

³³Good j is a loan which requires collateral that is less liquid than that required for the otherwise-identical loan that is good $j - 1$, so good j commands the higher interest rate (i.e., higher price). Fisher et al. (2011), Frost et al. (2015) and Fulmer (2022) describe the context in which the auction was introduced, and the Bank’s objectives.

³⁴The usefulness of negative bids in the Bank of England’s context may increase as technology develops to allow banks to better coordinate their different operations (see Baldwin et al. (2024)).

which the k th buyback bid buys 1 unit of good 1 if $p_2 \geq p_1 + f^{-1}(k)$ and 1 unit of good 2 otherwise, with additional buyback bids at the reserves.³⁵

In fact, the Bank did choose to sell just two goods initially, making the auction simple to solve: just as the Bank’s “supply function” preferences could equivalently be represented as a set of positive bids so, conversely, the aggregation of all bidder’s bids can equivalently be represented as (relative) “demand curves”. With just two goods, therefore, competitive equilibrium—and so the auction’s prices, and hence its allocations—is defined by the intersection of the single supply function with a single demand curve (see [Klemperer \(2008\)](#)). So the auction could easily be solved graphically, further aiding participants’ understanding and increasing their comfort with the procedure.

Once the Bank (and its bidders) had developed experience with its relatively simple initial case, it felt comfortable auctioning more goods and permitting the expression of more sophisticated preferences. However, it also became clear during the initial auctions that the Bank was much keener than the bidders to be able express a rich set of preferences.³⁶ So the current version of the Bank’s PMA takes advantage of the fact that the goods it is auctioning (loans) are in effect continuously divisible, which permits much greater flexibility in preferences without risking the failure of equilibrium. The total quantity of loans allocated, $\sum_i q_i(\mathbf{p})$, is now an increasing function of all the prices (and $\sum_{i \geq j} q_i(\mathbf{p}) / \sum_i q_i(\mathbf{p})$ is now a function of $p_j - p_{j-1}$).³⁷

The Bank has run its PMA at least monthly since 2010, and more often when institutions are likely to be under stress. It has been run weekly for the last several years. Most recently the Bank has been auctioning three varieties of 6 month repos in each auction, though its current implementation allows it to auction many more varieties simultaneously. It allocates a variable, and in principle unlimited, total quantity—as much as £7.2 billion was allocated in one auction, and around £240 billion has been allocated in all.³⁸

5.5. “Arctic” (Budget-Constrained) Product-Mix Auctions

Our standard PMA substitutes auction does not allow bidders to express the budget constraints that are important in many real-world auctions. So in 2015 the Icelandic government asked one of us to design a version of the PMA in which holders of blocked “offshore” funds could use their funds to bid for alternative financial instruments, and each bidder would be

³⁵So the k th buyback bid is as illustrated in [Figure 2d](#) with $r_2 = f^{-1}(k)$, $t_1 = t_2 = 1$ and $m = 1$, cf. the example in [note 27](#) in [Section 5.2](#). See [Klemperer \(2018, Appendix IE\)](#) and [Baldwin and Klemperer \(2019a\)](#) for further discussion. Note that it is now clear that the Bank’s preferences are strong substitutes, if we treat the units as indivisible. (The goods (loans) are in effect continuously divisible, but the Bank discretises the quantity available for the purpose of creating supply functions.)

³⁶So, as of this writing, the Bank is permitting bidders to bid on only a single good in each bid in their bid collection. However, the Bank’s current implementation allows it, if it wishes, to permit bidders to make any positive strong-substitutes bids (and many bidders bid on multiple varieties of good in the same auction).

It may seem unsurprising that the auctioneer (the Bank) wants to express more sophisticated preferences than the bidders, since it is interested in the profit from, and the composition of, the entire volume traded, whereas an individual bidder cares only about the price and composition of her own bundle. Note, however, that simultaneous multiple round auctions allow the auctioneer to express *less* sophisticated preferences than the bidders.

³⁷This aspect of the Bank’s preferences is implemented using a “Total Quantity Supply Schedule”. Precise details of how the Bank now specifies its preferences are confidential, but [Frost et al. \(2015\)](#) give some details. The current solution method was developed by Baldwin and Klemperer and introduced in 2014; it is described in [Baldwin and Klemperer \(2019a\)](#). See also the implementation at <http://pma.nuff.ox.ac.uk>.

³⁸[Giese and Grace \(2023\)](#) find that the Bank of England’s PMA increased welfare (as conventionally measured by the sum of the bidders’ and the Bank’s surpluses) by around 50% in 2010 to 2014, relative to if it had either auctioned each good in a separate auction with the quantities allocated to each good’s auction held constant over time, or run a “reference price auction” (see [note 49](#)) with constant price differences between goods.

constrained to a budget corresponding to the quantity of blocked funds she owned. We call this version of the PMA the “arctic” PMA.³⁹

More recently, IMF staff proposed using a version of the arctic PMA to restructure sovereign debt: creditors would bid to exchange their claims for alternative debt instruments, while the auctioneer (the debtor country), and potentially outsiders offering new funds, would choose supply schedules, as discussed in Section 5.2.⁴⁰ Another potential application of an arctic PMA is to allow shareholders to bid for choices of alternative securities for their shares in a firm that is being acquired or otherwise restructured. See Klemperer (2018, Appendix II) for further discussion.

An *arctic bid* $(\mathbf{r}; \beta)$ consists of a root vector \mathbf{r} whose coordinates $r_i \in \mathbb{R}$ are the values the bid places on the true goods $i \in [n]$, as in the standard substitutes PMA, and a budget $\beta \in \mathbb{R}_{>0}$.⁴¹ The bid can be interpreted as an agent with linear valuation $v(\mathbf{x}) = \mathbf{r} \cdot \mathbf{x}$ and budget β . So at prices $\mathbf{p} > \mathbf{0}$, the bidder has quasilinear utility $\mathbf{r} \cdot \mathbf{x} - \mathbf{p} \cdot \mathbf{x}$ for all non-negative bundles such that $\mathbf{p} \cdot \mathbf{x} \leq \beta$ and utility $-\infty$ for all other bundles. We assume divisible goods in this subsection, for simplicity.

The bid’s demand $D(\mathbf{p})$ consists of the utility-maximising bundles, so we have $D(\mathbf{p}) := \operatorname{argmax}_{\mathbf{x} \in \mathbb{R}_{\geq 0}, \mathbf{p} \cdot \mathbf{x} \leq \beta} (\mathbf{r} - \mathbf{p}) \cdot \mathbf{x}$; spending the bid’s budget β only on good i obtains $\frac{\beta}{p_i}$ units of that good, so utility equal to $(\frac{r_i}{p_i} - 1)\beta$.

So the arctic bid demands good i at $\mathbf{p} (> \mathbf{0})$ if i maximises the “bang per buck”, $\frac{r_i}{p_i}$, among all the goods $i \in [n]$ and $r_i \geq p_i$, see Figure 9a. By contrast, the standard substitutes PMA bid $(\mathbf{r}; \mathbf{t}; m)$ demands good i at \mathbf{p} if i maximises $t_i(r_i - p_i)$ among the goods $i \in [n]$ and $r_i \geq p_i$. The standard substitutes and arctic bids therefore demand the *same* good(s) at all \mathbf{p} if $t_i = \frac{1}{r_i}$, for $i \in [n]$, see Figure 9b. (Indeed, if $m = \beta$ and prices are such that the bids are indifferent about buying the zero bundle, then the bids demand identical bundles.) However, the arctic bid spends its entire budget, so demands $\frac{\beta}{p_i}$ of good i at \mathbf{p} (whenever good i is uniquely demanded), while the standard substitutes bid demands $mt_i = \frac{m}{r_i}$ of good i , independent of p_i .

Observe that any hyperplane along which an arctic bid is indifferent between spending its budget on either of two (true) goods i and j is normal to $\frac{1}{r_i} \mathbf{e}^i - \frac{1}{r_j} \mathbf{e}^j$, and always passes through the origin (whereas the diagonal facets of standard substitutes bids generally do not).

A bidder who prefers a portfolio of different goods can divide her budget across multiple separate bids; the demand of a set of arctic bids is simply the Minkowski sum of each bid’s demand.⁴²

³⁹The government hired a consultancy (dotEcon) who programmed and tested Klemperer’s design, and publicly announced it would run a PMA, but later abandoned it after the April 2016 political crisis (see Klemperer (2018), Appendix II). Fichtl (2022) describes and analyses the algorithm used to solve this PMA. Software to run it is at <http://pma.nuff.ox.ac.uk/>. Finding a competitive equilibrium in the special case that all bids are “Fisher bids” (see below) has been extensively studied (see, e.g., Vazirani (2007)).


The nomenclature “arctic” reflects both the auction’s Icelandic origin and also the fact that budget constraints generate preferences that cannot be described “tropically”. By contrast, tropical-geometric techniques can obtain results about equilibrium existence for agents with the quasilinear preferences for indivisible goods that the standard substitutes PMA expresses (see Baldwin and Klemperer (2019b)).

⁴⁰Klemperer suggested the design for this auction. See Willems (2021) and Lock’s implementation at <http://pma.nuff.ox.ac.uk/> for details.

⁴¹We do not include a “null” good, 0, so the vectors \mathbf{r} , \mathbf{p} , and \mathbf{x} are n -dimensional in this subsection, but we discuss below how a bidder who is not interested in the zero bundle can express this by “scaling up” \mathbf{r} while retaining the relative values of the r_i .

⁴²In principle, we could allow—and aggregate—standard substitutes PMA and arctic bids in the same auction.

Finster et al. (2023) analyse the case corresponding to, for example, the sale of internet adverts, where the auctioneer has no cost (i.e., zero valuation) up to a fixed supply. In particular they show that welfare maximisation coincides with the auctioneer’s revenue maximisation in this case.



(a) The demand of an arctic bid with root $r = (r_1, r_2)$ and budget β . (b) The demand of a standard substitutes PMA bid (for divisible goods) with root $r_{-0} = (r_1, r_2)$, tradeoffs $t_{-0} = (\frac{1}{r_1}, \frac{1}{r_2})$ and multiplicity m .

FIGURE 9.—Comparison of arctic and standard bids for two goods. Both bids demand the same goods at the same prices, but the quantities of these goods demanded by the arctic bid in the top left and bottom right regions decrease continuously with prices p_1 and p_2 , respectively. (If $m = \beta$ they demand identical bundles on the horizontal and vertical line segments where they are indifferent between two or more bundles.)

“Fisher markets” (see, e.g., [Gale \(1989\)](#)) can be understood as a special case of arctic PMAs: a *Fisher bid* $(r; \beta)$ can be interpreted as an agent with linear valuation $v(x) = r \cdot x$ and budget β who has no value for any unused budget. (For example, a bidder’s budget is “points” or “tokens” that can only be used in the current auction.) So at $p > \mathbf{0}$, the bidder’s utility is $r \cdot x$ for all bundles such that $p \cdot x = \beta$, and $-\infty$ for all other bundles. Observe that only the relative values of the r_i matter in a Fisher bid. So if we multiply the Fisher bid’s root by a sufficiently large constant,⁴³ we can then treat the bid as an arctic PMA bid.

[Huang et al. \(2024\)](#) show how bidders can construct arctic bid collections by answering a series of demand queries, analogously to [Goldberg et al. \(2022\)](#)’s methods for the standard substitutes PMA bidding language (see Section 5.1).

5.6. Special Words

In realistic settings of which we are aware, our languages are “compact” in that they require fewer bids than, say, listing the valuations of all bundles explicitly (see [Baldwin et al. \(2024, Section 2.3\)](#)). Nevertheless, especially with more than two goods, some natural preferences may be cumbersome to express with bids. So we further simplify our languages by introducing “words”. A *word* is a valid collection of bids that captures an economically intuitive preference, and can be communicated more concisely. Words thus simplify the expression of preferences and also make their representation easier to understand. Moreover, allowing bidders to use a few special words together with all positive bids may obviate the need to use any negative bids, and so also obviate any need for testing bid collections for validity.⁴⁴

For example, a bidder who wants to select up to *two* pieces of fruit among an apple, a banana, and a clementine, but at most one piece of any kind, needs to use four standard substitutes PMA

⁴³The constant is “sufficiently large” if there is at least one good, i , for which $r_i > p_i$ for every “relevant” p ; a price p is not “relevant” if every p_i is high enough that supply exceeds the sum of quantities $\frac{\beta}{p_i}$ that all bidders might demand of that good.

⁴⁴Using Proposition 4.3 (viii), it is not hard to see that the combination of any words and collections of positive bids will be valid, since both words (by construction) and collections of positive bids are always valid.

bids (three positive and one negative) to express these preferences.⁴⁵ So we can define a “word” that means “(up to) best k_1 units from k_2 specified goods (but no more than one unit of any good)”.⁴⁶

Which words are useful in a practical implementation will depend on the context. For example, it might be helpful for a central bank running liquidity auctions to offer potentially-distressed bidders a word signifying: “a need to borrow at least $\text{£}x$ whatever the interest rate, however high, preferring a type-1 loan to a type-2 loan if the interest-rate premium for type 1 is less than $r\%$, but a wish to borrow $\text{£}x_1 \geq x$ (in total) of type 1 at any interest rate less than r_1 and also $\text{£}x_2 \geq x$ (in total) of type 2 at any interest rate less than $r_2 = r_1 - r$ ”.⁴⁷

It is not hard to see that any arctic bid $(r; \beta)$ can be approximated arbitrarily closely by a collection of standard substitutes PMA bids arrayed along the diagonal line segment from the root of that arctic bid down to the origin.⁴⁸ Thus an arctic bid can be understood as another special “word” in a standard substitutes PMA.

6. RELATED LANGUAGES AND AUCTIONS

The alternative mechanism that is most commonly used in practice instead of a PMA is simply running a set of independent auctions, one for each good, with a fixed quantity available in each. This is clearly less efficient than running a PMA, and [Grace \(2024b\)](#) shows running separate simultaneous auctions gives lower expected surplus to both the bidders (in aggregate) and the auctioneer, under broad conditions.⁴⁹

The mechanism that corresponds most closely to the PMA is [Milgrom \(2009\)](#)’s proposed “assignment auction”. As with the PMA, this would be a static auction which finds competitive equilibrium prices and allocations assuming bidders express substitutes preferences. The crucial distinction is that Milgrom’s proposal requires bidders to express their preferences using constraints that satisfy a set of partly-overlapping algebraic tree-structures. This seems harder for non-economists to understand than the PMA’s language, which links directly to the geometry of what bundles of goods a bidder demands at which prices in price space. Also

⁴⁵The three positive bids have roots, r_{-0} , equal to $(v_a, v_b, -\infty)$, $(v_a, -\infty, v_c)$ and $(-\infty, v_b, v_c)$, and multiplicity 1. The negative bid has root $r_{-0} = (v_a, v_b, v_c)$ and multiplicity -1 . All the bids have tradeoffs $t = \mathbf{1}$.

⁴⁶As an example of how words make our representation of preferences easier to understand, 5 standard positive strong-substitutes PMA bids plus 4 “(up to) best 2 from 3” bids is much more easily understood than its equivalent representation as $5 + (4 \cdot 3) = 17$ positive bids plus 4 negative bids. The latter representation both contains more vectors and has its meaning obscured by the entangling of the 16 bids that are 4 groups of 4 connected bids with each other and with the 5 “individual” bids.

We implement “words” including “(up to) best k_1 from k_2 goods” in our software at <http://pma.nuff.ox.ac.uk>.

⁴⁷So the bidder wants to borrow a total of $\text{£}x_1$ of type 1 and $\text{£}x_2$ of type 2 if both interest rates are sufficiently low. In the standard substitutes PMA language, these preferences are expressed by the bid collection $(-\infty, 0, -r; 1, 1, 1; x)$, $(0, r_1, -\infty; 1, 1, 1; x_1)$, $(0, -\infty, r_2; 1, 1, 1; x_2)$, and $(0, r_1, r_2; 1, 1, 1; -x)$. (See [Klemperer 2008, 2010](#).)

⁴⁸The required bids are a standard substitutes PMA bid for the maximum quantity of each good that the arctic bid could afford at prices r (this bid is $(r; t; \beta)$ with $t_i := \frac{1}{r_i}$), plus a list of “tiny” standard substitutes PMA bids at decreasing (but sufficiently close) price vectors, each bid of which is for the additional quantity of each good that the arctic bid could afford if the actual prices fell the additional distance below the prices of the previous bid (these bids are $b^j := (\lambda_j r; t; m_j)$ with $m_j := \frac{1}{\lambda_j} - \frac{1}{\lambda_{j-1}}$ for $j \geq 1$ and $1 = \lambda_0 > \lambda_1 > \lambda_2 > \dots > 0$). See [Lemma F.12](#).

⁴⁹Another inefficient alternative to a PMA that is sometimes used in government securities and energy markets is a “reference price auction” in which all goods are sold in a single auction, but the auctioneer specifies fixed price differences between goods (see, e.g., [Armantier and Holt \(2023\)](#) and [Fabra and Montero \(2023\)](#)). [Grace \(2024b\)](#) gives conditions under which this gives less expected surplus than a PMA to both the bidders and the auctioneer. [Giese and Grace \(2023\)](#) analyses the benefits of the Bank of England’s PMA relative to running either separate simultaneous auctions or reference price auctions (see note 38).

important, the assignment-auction language neither permits the representation of all possible strong-substitutes preferences (see [Fichtl \(2021\)](#)) nor, to our knowledge, exactly represents any other standard preference class.

Both the PMA and the assignment auction can be understood as static versions of the (dynamic) simultaneous multiple round auction (SMRA) ([Milgrom, 2000](#)). All three auctions find competitive equilibrium if bidders with strong-substitutes valuations bid their true preferences, but the PMA and assignment auctions also find exact competitive equilibria for broader classes of private-value preferences, whereas the SMRA in general does not. Importantly, the PMA and assignment auctions allow the quantities traded, as well as the prices, to depend on participants' preferences, by contrast with the SMRA which sells a fixed bundle of goods.⁵⁰ Moreover, the PMA and assignment auctions can also implement objectives other than efficiency, which the SMRA cannot do. And, because they are static auctions, the PMA and assignment auctions also have the advantages of speed and/or not requiring bidders to bid in real time. Finally, although bidders' ability to observe others' behaviour in a dynamic auction may make the SMRA more efficient if preferences have significant "common-value" components and/or exhibit some forms of complements,⁵¹ it can also facilitate anti-competitive practices (collusion, etc. ([Klemperer, 2002, 2004](#))).

The PMA language is also related to the 'logical bidding languages', introduced by [Sandholm \(1999\)](#) and [Fujishima et al. \(1999\)](#), whose bids are constructed recursively from valuations for individual bundles of goods using 'OR' and 'XOR' operators.⁵² A single positive PMA bid with multiplicity 1 corresponds to multiple 'XOR'ed unit-demand valuations, while a collection of positive PMA bids corresponds to the 'OR' of the valuations corresponding to the bids in the collection. However, the logical languages offer no straightforward analogues to the PMA's negative bids. Also importantly, there is, to our knowledge, no way to restrict the syntax of these languages to express either the set of all, but only, concave substitutes valuations, or the set of all, but only, strong-substitutes valuations.⁵³ Nor do the logical languages facilitate the geometric representation and analysis of preferences.

7. CONCLUSION

It is now commonly said that the (only) analytical framework for auctions is game theory. By contrast, we believe competitive-equilibrium analysis can often provide a good approximation of behaviour, so we have instead focused attention on an important aspect of auction design that is often overlooked—making it straightforward for participants to communicate their preferences.

⁵⁰However, an SMR auctioneer could sell a variable supply by offering the fixed supply which comprises the maximum number of units of each good that could be sold of that good, and then buying back goods to ensure that the "correct quantity" is sold (cf. our discussion in Section 5.2).

⁵¹However, the Bank of England's current implementation of the PMA permits some limited complements preferences, and [Klemperer \(2010, note 20\)](#) suggested a method of extending the PMA to allow bidders to update their bids based on reported "interim" auction prices to improve efficiency if there are significant "common-value" components. Moreover, we know of no compelling theoretical results about the efficiency of SMRAs or related dynamic auctions in the context of common values or complements.

⁵²See also [Boutlier and Hoos \(2001\)](#), [Nisan \(2000\)](#), [Lehmann et al. \(2006\)](#), and [Nisan \(2005\)](#)'s survey. As defined, these languages only express strong-substitutes valuations for up to one unit of each indivisible good, but they can easily be extended to allow for divisible goods (see [Kaleta \(2013\)](#)) and multiple units of goods.

⁵³Existing logical languages can represent either a strict subset of substitutes valuations or (e.g., in the case of the OXS and XOS languages ([Lehmann et al., 2006](#))) a strict superset. For example, [Lehmann et al.](#)'s XS language is very slightly less general than the restriction to only positive bids of the strong-substitutes PMA language.

Importantly, an institution can begin by using even simpler versions of the product-mix auction than those described here, and then allow participants to express more sophisticated preferences in later auctions when they are comfortable with the mechanism. For example, although no other central bank has yet copied the Bank of England’s original PMA, other central bank auctions have introduced some PMA features and we hope more will follow.

REFERENCES

- AHUJA, RAVINDRA, THOMAS MAGNANTI, AND JAMES ORLIN (1993): *Network Flows: Theory, Algorithms, and Applications*, USA: Prentice-Hall, Inc. [49]
- ARMANTIER, OLIVIER, AND CHARLES HOLT (2023): “Endogenous Reference Price Auctions for a Diverse Set of Commodities: an Experimental Analysis,” *Experimental Economics*, 27, 1–27. [27]
- AUSUBEL, LAWRENCE (2006): “An Efficient Dynamic Auction for Heterogeneous Commodities,” *American Economic Review*, 96, 602–629. [23]
- BALDWIN, ELIZABETH, MARTIN BICHLER, MAXIMILIAN FICHTL, AND PAUL KLEMPERER (2024): “Strong Substitutes: Structural Properties, and a New Algorithm for Competitive Equilibrium Prices,” *Mathematical Programming*, 203, 611–643. [19, 22, 23, 26, 43, 49, 50]
- BALDWIN, ELIZABETH, PAUL GOLDBERG, PAUL KLEMPERER, AND EDWIN LOCK (2023a): “Solving Strong-Substitutes Product-Mix Auctions,” *Mathematics of Operations Research*, in press; published online 2023. [8, 19, 20, 23]
- BALDWIN, ELIZABETH, RAVI JAGADEESAN, PAUL KLEMPERER, OMER EDHAN, AND ALEXANDER TEYTELBOYM (2020): “The Equilibrium Existence Duality: Equilibrium with Indivisibilities & Income Effects,” CEPR Discussion Paper No. 14926. CEPR Press, Paris & London. Available at <https://cepr.org/publications/dp14926>. [12, 21]
- BALDWIN, ELIZABETH, RAVI JAGADEESAN, PAUL KLEMPERER, AND ALEXANDER TEYTELBOYM (2021): “On Consumer Theory with Indivisible Goods,” Working Paper. Available at www.nuffield.ox.ac.uk/users/klemperer/consumer-theory-indivisible.pdf. [12, 21]
- (2023b): “The Equilibrium Existence Duality,” *Journal of Political Economy*, 131, 1440–1476. [21]
- BALDWIN, ELIZABETH, AND PAUL KLEMPERER (2014): “Tropical Geometry to Analyse Demand,” Working Paper, Nuffield College. Available at <https://www.nuff.ox.ac.uk/users/klemperer/Tropical2014.pdf>. [12]
- (2016): “Proof that any Strong Substitutes Preferences can be Represented by the Strong Substitutes Product-Mix Auction Bidding Language. Tech. rep.. (day 2 slides on Tropical intersections and equilibrium, Baldwin, E., Goldberg, P. and Klemperer, P.)” Hausdorff School on Tropical Geometry and Economics. Available at <http://people.math.gatech.edu/~jyu67/HCM/Baldwin2.pdf>. [3]
- (2019a): “Notes on Linear Programming to Solve the Product-Mix Auction,” Working Paper, Nuffield College. Available at <https://www.nuffield.ox.ac.uk/users/klemperer/linprog.pdf>. [22, 24]
- (2019b): “Understanding Preferences: “Demand Types”, and the Existence of Equilibrium with Indivisibilities,” *Econometrica*, 87, 867–932. [5, 10, 12, 21, 25, 35, 36, 50]
- (2021): “Proof that the Product-Mix Auction Bidding Language Can Represent Any Substitutes Preferences,” Nuffield College Economics Discussion Paper 2021-W05. Available at https://www.nuffield.ox.ac.uk/economics/Papers/2021/2021-W05_subproof.pdf. [3]
- BALKANSKI, ERIC, AND RENATO PAES LEME (2020): “On the Construction of Substitutes,” *Mathematics of Operations Research*, 45, 272–291. [3]
- BIKHCHANDANI, SUSHIL, AND JOHN W. MAMER (1997): “Competitive Equilibrium in an Exchange Economy with Indivisibilities,” *Journal of Economic Theory*, 74, 385–413. [22]
- BOUTILIER, CRAIG, AND HOLGER HOOS (2001): “Bidding Languages for Combinatorial Auctions,” *International Joint Conference on Artificial Intelligence*, 17, 1211–1217. [28]
- DANILOV, VLADIMIR, GLEB KOSHEVOY, AND CHRISTINE LANG (2003): “Gross Substitution, Discrete Convexity, and Submodularity,” *Discrete Applied Mathematics*, 131, 283–298. [5]
- DANILOV, VLADIMIR, GLEB KOSHEVOY, AND KAZUO MUROTA (2001): “Discrete Convexity and Equilibria in Economies with Indivisible Goods and Money,” *Mathematical Social Sciences*, 41, 251–273. [48]
- FABRA, NATALIA AND JUAN-PABLO MONTERO (2023): “Technology-Neutral Versus Technology-Specific Procurement,” *Economic Journal*, 133, 669–705. [27]
- FICHTL, MAXIMILIAN (2021): “On the Expressiveness of Assignment Messages,” *Economics Letters*, 208, 1–4. [3, 28]
- (2022): “Computing Candidate Prices in Budget-Constrained Product-Mix Auctions,” *arXiv preprint arXiv:2208.02633*. [25]

- FINSTER, SIMON (2020): “Strategic Bidding in Product-Mix, Sequential, and Simultaneous Auctions,” Nuffield College Discussion Paper 2020-W03. Available at https://www.nuffield.ox.ac.uk/economics/Papers/2020/2020W03_pma_strategic_5.pdf. [2]
- (2021): “Strategic Behaviour in Multi-Object Auctions: Theory and Experiment,” Working Paper. Available at https://drive.google.com/file/d/1ayPs1_EqjzG81ommXMWrWyB1OpvrpYtc/view. [1]
- FINSTER, SIMON, PAUL GOLDBERG, AND EDWIN LOCK (2023): “Substitutes Markets with Budget Constraints: Solving for Competitive and Optimal Prices,” in *The 19th Conference On Web And InterNet Economics*. [25]
- FISHER, PAUL, TARKUS FROST, AND OLAF WEEKEN (2011): “Pricing Central Bank Liquidity through Product-Mix Auctions—the First Anniversary of the Bank of England’s Indexed Long-Term Repo Operations,” Bank of England Working Paper. Available at <https://elischolar.library.yale.edu/ypfs-documents/10672>. [23]
- FRIEDMAN, MILTON (1959): “Testimony in Employment, Growth, and Price Levels. Hearings before the Joint Economic Committee,” in *86th Congress, 1st Session, October, 30, 3032–3026*. [2]
- (1963): “Price Determination in the United States Treasury Bill Market: a Comment,” *Review of Economics and Statistics*, 45, 318–320. [2]
- (1991): “How to Sell Government Securities,” *Wall Street Journal*, August 28. [2]
- FROST, TARKUS, NICK GOVIER, AND TOM HORN (2015): “Innovations in the Bank’s Provision of Liquidity Insurance via Indexed Long-Term Repo (ILTR) Operations,” *Bank of England Quarterly Bulletin*, 2015, 181–188, available at <https://www.bankofengland.co.uk/quarterly-bulletin/2015/q2/innovations-in-the-banks-provision-of-liquidity-insurance-via-indexed-long-term-repo-operations>. [23, 24]
- FUJISHIMA, YUZO, KEVIN LEYTON-BROWN, AND YOAV SHOHAM (1999): “Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches,” *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, 1, 548–553. [28]
- FULMER, SEAN (2022): “United Kingdom: Indexed Long-Term Repo Operations,” *Journal of Financial Crises*, 4, 1126–1155. [23]
- GALE, DAVID (1989): *The Theory of Linear Economic Models*, University of Chicago Press. [26]
- GIESE, JULIA, AND CHARLOTTE GRACE (2023): “An Evaluation of the Bank of England’s ILTR Operations: Comparing the Product-Mix Auction to Alternatives,” Bank of England Working Paper 1044. [24, 27]
- GOLDBERG, PAUL, EDWIN LOCK, AND FRANCISCO MARMOLEJO-COSSÍO (2022): “Learning Strong Substitutes Demand via Queries,” *ACM Transactions on Economics and Computation*, 10, 1–22. [20, 26]
- GRACE, CHARLOTTE (2024a): “Competing Models of the Bank of England’s Liquidity Auctions: Truthful Bidding is a Good Approximation,” Bank of England Working Paper 1061. Available at <https://www.bankofengland.co.uk/working-paper/2024/competing-models-of-the-bank-of-englands-liquidity-auctions>. [2]
- (2024b): “Who Gains from Efficient Auctions?” Nuffield College Economics Discussion Paper 2024-W01. Available at https://www.nuff.ox.ac.uk/economics/papers/2024/2024-W01grace_2024_who_gains.pdf. [27]
- GRÜNBAUM, BRANKO (1967): *Convex Polytopes*, John Wiley & Sons, 1 ed. [43]
- GUL, FARUK, AND ENNIO STACCHETTI (1999): “Walrasian Equilibrium with Gross Substitutes,” *Journal of Economic Theory*, 87, 95–124. [7]
- (2000): “The English Auction with Differentiated Commodities,” *Journal of Economic Theory*, 92, 66–95. [23]
- HATFIELD, JOHN WILLIAM, SCOTT DUKE KOMINERS, ALEXANDRU NICHIFOR, MICHAEL OSTROVSKY, AND ALEXANDER WESTKAMP (2013): “Stability and Competitive Equilibrium in Trading Networks,” *Journal of Political Economy*, 121, 966–1005. [9]
- HATFIELD, JOHN WILLIAM AND PAUL MILGROM (2005): “Matching with Contracts,” *American Economic Review*, 95, 913–935. [3]
- HOLMBERG, PÄR, KEITH RUDELL, AND BERT WILLEMS (2019): “Multi-Product Supply Function Equilibria,” Working paper. Research Institute of Industrial Economics (IFN), Stockholm. [2]
- HUANG, DAVID, EDWIN LOCK, AND FRANCISCO MARMOLEJO-COSSÍO (2024): “Learning Demand with LLM Proxies,” In preparation. [26]
- JENSEN, ANDERS (2024): “Gfan, a Software System for Gröbner Fans,” available at <https://users-math.au.dk/~jensen/software/gfan/gfan.html>. [19]
- KALETA, MARIUSZ (2013): “Bidding Languages for Continuous Auctions,” in *New Trends in Databases and Information Systems*, Springer, 211–220. [28]
- KELSO, ALEXANDER, AND VINCENT CRAWFORD (1982): “Job Matching, Coalition Formation, and Gross Substitutes,” *Econometrica: Journal of the Econometric Society*, 1483–1504. [23]
- KLEMPERER, PAUL (2002): “What Really Matters in Auction Design,” *Journal of Economic Perspectives*, 16, 169–189. [28]
- (2004): *Auctions: Theory and Practice*, Princeton University Press. [28]
- (2008): “A New Auction for Substitutes: Central Bank Liquidity Auctions, the U.S. TARP, and Variable Product-Mix Auctions,” Nuffield College Working Paper, available at <https://www.nuffield.ox.ac.uk/economics/Papers/2008/substsauc.pdf>. [2, 3, 23, 24, 27]

- (2010): “The Product-Mix Auction: a New Auction Design for Differentiated Goods,” *Journal of the European Economic Association*, 8, 526–536. [3, 27, 28]
- (2018): “Product-Mix Auctions,” Nuffield College Economics Discussion Paper 2018-W07. Available at <https://www.nuffield.ox.ac.uk/economics/Papers/2018/2018W07productmix.pdf>. [4, 24, 25]
- LEHMANN, BENNY, DANIEL LEHMANN, AND NOAM NISAN (2006): “Combinatorial Auctions with Decreasing Marginal Utilities,” *Games and Economic Behavior*, 55, 270–296. [28]
- LIN, BO, AND NGOC MAI TRAN (2017): “Linear and Rational Factorization of Tropical Polynomials,” *arXiv preprint arXiv:1707.03332*. [3]
- MAS-COLELL, ANDREU, MICHAEL D. WHINSTON, AND JERRY R. GREEN (1995): *Microeconomic Theory*, Oxford University Press. [47]
- MIKHALKIN, GRIGORY (2004): “Decomposition into Pairs-of-Pants for Complex Algebraic Hypersurfaces,” *Topology*, 43, 1035–1065. [11, 17, 36, 42]
- MILGROM, PAUL (2000): “Putting Auction Theory to Work: the Simultaneous Ascending Auction,” *Journal of Political Economy*, 108, 245–272. [23, 28]
- (2009): “Assignment Messages and Exchanges,” *American Economic Journal: Microeconomics*, 1, 95–113. [3, 27]
- MILGROM, PAUL, AND CHRIS SHANNON (1994): “Monotone Comparative Statics,” *Econometrica*, 62, 157–180. [68]
- MILGROM, PAUL, AND BRUNO STRULOVICI (2009): “Substitute Goods, Auctions, and Equilibrium,” *Journal of Economic Theory*, 144, 212–247. [3, 5, 21]
- MUROTA, KAZUO (2003): *Discrete Convex Analysis*, Society for Industrial and Applied Mathematics. [3]
- MUROTA, KAZUO, AND AKIYOSHI SHIOURA (1999): “M-convex Function on Generalized Polymatroid,” *Mathematics of Operations Research*, 24, 95–105. [3]
- MUROTA, KAZUO, AKIYOSHI SHIOURA, AND ZAIFU YANG (2013): “Computing a Walrasian Equilibrium in Iterative Auctions with Multiple Differentiated Items,” in *Algorithms and Computation: 24th International Symposium, ISAAC 2013, Hong Kong, China, December 16-18, 2013, Proceedings 24*, Springer, 468–478. [23]
- (2016): “Time Bounds for Iterative Auctions: a Unified Approach by Discrete Convex Analysis,” *Discrete Optimization*, 19, 36–62. [23]
- NISAN, NOAM (2000): “Bidding and Allocation in Combinatorial Auctions,” in *Proceedings of the 2nd ACM Conference on Electronic Commerce*, Association for Computing Machinery, 1–12. [28]
- (2005): “Bidding Languages for Combinatorial Auctions,” in *Combinatorial Auctions*, The MIT Press, 215–232. [28]
- OSTROVSKY, MICHAEL, AND RENATO PAES LEME (2015): “Gross Substitutes and Endowed Assignment Valuations,” *Theoretical Economics*, 10, 853–865. [3]
- ROCKAFELLAR, RALPH (1970): *Convex Analysis*, Princeton: Princeton University Press. [50]
- SANDHOLM, TUOMAS (1999): “An Algorithm for Optimal Winner Determination in Combinatorial Auctions,” in *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, Morgan Kaufmann Publishers Inc., 542–547. [28]
- SCHNEIDER, ROLF (2013): *Convex Bodies: the Brunn-Minkowski Theory*, Cambridge University Press, 2 ed. [43]
- SHIOURA, AKIYOSHI, 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, 61–103. [3, 5, 12]
- TAO, PHAM DINH, AND LE THI HOAI AN (1997): “Convex Analysis Approach to DC Programming: Theory, Algorithms and Applications,” *Acta Mathematica Vietnamica*, 22, 289–355. [50]
- TRAN, NGOC MAI (2021): “The Finite Matroid-Based Valuation Conjecture is False,” *SIAM Journal on Applied Algebra and Geometry*, 5, 506–525. [3]
- VAZIRANI, VIJAY (2007): “Combinatorial Algorithms for Market Equilibria,” in *Algorithmic Game Theory*, ed. by Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, Cambridge University Press, 103–134. [25]
- WILLEMS, TIM (2021): “An Auction-Based Sovereign Debt Restructuring Mechanism,” Working Paper. Available at https://drive.google.com/file/d/1fJUrFM860-n4xVDaivIz_t2IjkdAn9_/view. [25]

APPENDIX

Appendix **A** provides additional proof sketches for the Representation Theorem. Appendix **B** contains proofs for Section 2.2. Appendix **C** present additional details on the geometry of demand used in the remaining appendices. Appendices **D** to **F** accompany Sections 3 to 5, respectively, with the exception of our main results, namely Theorem 3.1 and Corollaries 3.2, 3.3 and 3.7. Finally, Appendix **G** contains the full proofs of our main results.

APPENDIX A: ADDITIONAL PROOF SKETCHES FOR THE REPRESENTATION THEOREM


A.1. *The Representation Theorem for the General Two-Good Case*

FIGURE A.1.—(a) LIP \mathcal{L}_v of a general concave substitutes valuation v . The facets are labelled with their weights, and the bundle of goods demanded in each UDR is shown. The dashed box depicts a possible bounding box. (b) The bid collection representing v .

Figure A.1 shows an example of a concave substitutes valuation v on two goods that is not regular. It also shows an example of a “bounding box” with horizontal and vertical sides whose corners (in this example including $(1, 1)$ and $(6, 5)$) are chosen so that it is large enough to contain at least part of the interior of every facet of \mathcal{L}_v .

In order to determine the bid collection corresponding to v , we first create regular bids at each vertex of \mathcal{L}_v exactly as in Section 3.3.1. As discussed in Section 3.1, these bids (bids 1, 2, and 3 in the table in Figure A.1) only generate vertical and horizontal facets that go up, and diagonal facets that go down, so we also need to include non-regular bids.

At each intersection, $\mathbf{p} = (p_1, p_2)$, of a vertical facet of \mathcal{L}_v with the bottom boundary of the box, we create a bid of the type shown in Figure 2b whose LIP matches this facet around \mathbf{p} . Each such bid (bids 4 and 5 in our example) has root $\mathbf{r} = (0, p_1, -\infty)$, tradeoff $\mathbf{t}_{-0} = (1, 0)$ and multiplicity $m = w_v^u(\mathbf{p})$ (see Figure 5). These bids mean that the weights of the vertical facets of our bid collection now all match those of \mathcal{L}_v . For example, the newly-created bid 4 creates a facet of weight 1 at low prices along the line $p_1 = 2$, and then cancels with bid 1 to give weight 0 above $(2, 3)$.

Similarly, at each intersection, $\mathbf{p} = (p_1, p_2)$, of a horizontal facet of \mathcal{L}_v with the left-hand boundary of the box, we create a bid of the type shown in Figure 2c with root $\mathbf{r} = (0, -\infty, p_2)$, tradeoff $\mathbf{t}_{-0} = (0, 1)$, and multiplicity $m = w_v^r(\mathbf{p})$. Including these bids (in our example this is just bid 6) ensures that the weights of the horizontal facets of our bid collection now all match those of \mathcal{L}_v , by the same argument as for the vertical facets. And since these bids’ LIPs all consist of a single horizontal facet, the matching of the vertical facets’ weights of the LIPs of the bid collection and of the valuation is unaffected.

Next, at each intersection, $\mathbf{p} = (p_1, p_2)$, of a diagonal facet of \mathcal{L}_v with slope $\frac{t_1}{t_2}$ (expressed in lowest terms) on either the right-hand or upper boundaries of the box, we create a bid of the type shown in Figure 2d with root $\mathbf{r} = (-\infty, 0, p_2 - \frac{t_1}{t_2}p_1)$,⁵⁴ tradeoff $\mathbf{t}_{-0} = (t_1, t_2)$, and

⁵⁴This is equivalent to choosing $\mathbf{r} = (-\infty, p_1, p_2)$ and then normalising the bid as described in Section 2.2.

multiplicity $m = w_v^{dl}(\mathbf{p}, \frac{t_1}{t_2})$. Including these bids (just bid 7 in our example) means the weights of these facets of our bid collection match the weights of the corresponding facets of \mathcal{L}_v . And, since the LIPs of these bids each consist of a single diagonal facet, their inclusion does not alter the fact that the vertical and horizontal facet weights, and therefore *all* facets and weights of the LIPs of the bid collection and of the valuation also match.

Finally, we need to adjust demand globally to the correct levels by adding bids that unconditionally demand a single good. To do this, we fix a price in a unique demand region of \mathcal{L}_v , for example $\mathbf{p} = (6, 5)$. We see $D_v(\mathbf{p}) = \{(1, 2)\}$ in our example, but that the uniquely demanded bundle for our current bid collection at the same price is $(0, 2)$. (This demand comes from bid 7 only, since all other bids demand $\mathbf{0}$ at \mathbf{p} because \mathbf{p} is sufficiently high (see Figures 2a to 2c).) So for our example we add a bid that unconditionally demands good 1 with multiplicity 1 (bid 8 in the table) so that the demand of our bid collection and of the valuation are the same at \mathbf{p} . This does not affect the LIP of the bid collection, as the LIP of an unconditional bid is “empty”. So our final bid collection, \mathcal{B} , satisfies both $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}}) = (\mathcal{L}_v, w_v)$ and $D_{\mathcal{B}}(\mathbf{p}) = D_v(\mathbf{p})$ at all prices \mathbf{p} , as required (cf. Proposition 3.4).

A.2. More Details for the Representation Theorem for the General Case

In n dimensions, facets are $(n-1)$ -dimensional objects, so we refer to them using the vectors normal to them. The LIP of a regular bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ has $n(n-1)/2$ “diagonal” facets, one with normal $t_i e^i - t_j e^j$ for each distinct pair $i, j \in [n]$, which we correspondingly label $F_{\mathbf{b}}^{ij}$. It also has n facets whose normal vectors are parallel to the axes; we label the facet normal to the i -axis as $F_{\mathbf{b}}^{i0}$. (Recall Figure 7, which illustrates a bid for $n = 3$ goods.)

We begin with the case in which v is a regular valuation. For every vertex \mathbf{p} of \mathcal{L}_v , we include a bid $(\mathbf{r}; \mathbf{t}; m)$ with root $\mathbf{r} = (0, \mathbf{p})$ for every tradeoff \mathbf{t} such that $t_i e^i - t_j e^j$ is normal to a facet of \mathcal{L}_v at \mathbf{p} , for all $i, j \in [n]$, just as in the $n = 2$ case. But identifying the multiplicity for a given (\mathbf{r}, \mathbf{t}) pair is more intricate than before. We first fix an orientation (i, j) . Let $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$ be the hyperplane through \mathbf{p} with normal $t_i e^i - t_j e^j$, and refer to facets lying in $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$ as $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$ -facets. When $n = 2$, $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$ is a line, and so there are at most two $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$ -facets in \mathcal{L}_v at any \mathbf{p} : “up-right” and “down-left”. But for $n > 2$, there may be many more.

We want to use each bid $\mathbf{b} = ((0, \mathbf{p}); \mathbf{t}; m)$ to set the weight in $\mathcal{L}_{\mathcal{B}}$ of a specific facet $F \in F_{\mathbf{b}}^{ij}$, so that $w_{\mathcal{B}}(F) = w_v(F)$.⁵⁵ But, as for $n = 2$, we cannot simply set m equal to $w_v(F)$, because F is potentially contained in $F_{\mathbf{b}'}^{ij}$ for other bids, \mathbf{b}' , that we *also* need to include in order to depict the LIP locally in the neighbourhood of their own roots. For $n = 2$, we only have to account for any diagonal facet on the other side of the bid’s root, \mathbf{r} , but for $n > 2$ there are generally multiple other bids.

By accounting for the weight in \mathcal{L}_v of other $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$ -facets containing \mathbf{p} , we account for the effects of these other bids. So we define a function $m^{ij}(\mathbf{p}, \mathbf{t})$, which is a signed sum of the weights of these facets. When $n = 2$, for example, $m^{ij}(\mathbf{p}, \mathbf{t}) = w^{dl}(\mathbf{p}; \mathbf{t}) - w^{ur}(\mathbf{p}; \mathbf{t})$.⁵⁶ We

⁵⁵The facet $F \subseteq F_{\mathbf{b}}^{ij}$ we target is the one which contains \mathbf{p} and part of the line $\mathbf{p} - \lambda(0, \frac{1}{t_1}, \dots, \frac{1}{t_n})$ at which the bid’s diagonal facets all meet. This facet may not exist in \mathcal{L}_v . So our proof extends each LIP to a “hyperplane of indifference prices” (HIP) which is the union of all hyperplanes in \mathcal{P} containing a facet of the LIP. The HIP decomposes into facets in a natural way, and the facet we require does exist in the HIP.

Facets of \mathcal{L}_v contained in $F_{\mathbf{b}}^{ij}$ and containing \mathbf{p} , but not containing part of this line will also be contained in $F_{\mathbf{b}'}^{ij}$ for other bids \mathbf{b}' with roots at \mathbf{r} but with tradeoffs $\mathbf{t}' \neq \mathbf{t}$.

⁵⁶When $n > 2$, we need to account for the fact that the weight of the facet of a bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ with normal $t_i e^i - t_j e^j$ is $m \gcd(t_i, t_j)$ (and not simply m). So we need to divide the sum of the weights of the facets with this normal that meet at \mathbf{r} by $\gcd(t_i, t_j)$, and show that this yields an integer weight.

can apply this function to the weights of either \mathcal{L}_v or \mathcal{L}_B (using subscripts to denote the case, as before). Now set the multiplicity of every bid via $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m_v^{ij}(\mathbf{r}_{-0}; \mathbf{t}))$. Having done so, we can show that it follows that $m_v^{ij}(\mathbf{p}; \mathbf{t}) = m_B^{ij}(\mathbf{p}; \mathbf{t})$, as when $n = 2$. When v is regular, this equality implies that the weights of all $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$ -facets in \mathcal{L}_B match those in \mathcal{L}_v ; we prove this by induction across $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$, just as we did for $n = 2$.⁵⁷ We apply this argument for every such hyperplane $H(\mathbf{p}; i, j; \frac{t_i}{t_j})$. We now know that all the weighted facets with normal $t_i e^i - t_j e^j$ match across \mathcal{L}_v and \mathcal{L}_B , for any t_i and t_j (for our fixed i and j).

Now, as in the $n = 2$ case, the balancing condition tells us that finding the weights for the diagonal (i, j) -facets automatically finds them for all the other diagonal facets. We confirm this by observing that the choice of i and j was arbitrary: for any $k \neq l$ in $[n]$ we can match the facets with normal $t_k e^k - t_l e^l$ across \mathcal{L}_v and \mathcal{L}_B , using a multiplicity function $m^{kl}(\mathbf{p}; \mathbf{t})$ on the weights of facets in $H(\mathbf{p}; k, l; \frac{t_k}{t_l})$, and the balancing condition shows us that $m^{kl}(\mathbf{p}; \mathbf{t}) = m^{ij}(\mathbf{p}; \mathbf{t})$ for balanced LIPs. So since both \mathcal{L}_v and \mathcal{L}_B are balanced, we have $m_v^{kl}(\mathbf{p}; \mathbf{t}) = m_B^{kl}(\mathbf{p}; \mathbf{t})$ for all $k \neq l$ in $[n]$.

And we can similarly define a multiplicity function, $m^{i0}(\mathbf{p})$, which is a signed sum of the weights of facets meeting at \mathbf{p} and with normal e^i . (For example, when $n = 2$, we have $m^{10}(\mathbf{p}) = w_v^u(\mathbf{p}) - w_v^d(\mathbf{p})$.) The balancing condition shows that $\sum_{\mathbf{t}} t_i m^{ij}(\mathbf{p}, \mathbf{t}) = m^{i0}(\mathbf{p})$, where the sum is taken over all the tradeoffs \mathbf{t} used at root $(0, \mathbf{p})$. So, again since \mathcal{L}_v and \mathcal{L}_B are balanced, it follows that $m_v^{i0}(\mathbf{p}) = m_B^{i0}(\mathbf{p})$, and our same inductive proof again shows that, in the regular case, these weighted facets match everywhere. So now $(\mathcal{L}_v, w_v) = (\mathcal{L}_B, w_B)$, and $D_v = D_B$ follows as in the $n = 2$ case.

When v is not regular we proceed as above but, as explained in Section 3.3.2, we also need to add bids that are not regular. We do this by introducing a bounding box that is a hyperrectangle which is large enough that its interior contains at least part of the interior of every facet of \mathcal{L}_v .

For every vertex \mathbf{p} where \mathcal{L}_v intersects the boundary of the bounding box, we create bids whose roots' coordinates are $r_0 = -\infty$ if \mathbf{p} lies on any upper boundary, and $r_0 = 0$ otherwise; $r_i = -\infty$ for any i for which \mathbf{p} lies on the lower i -boundary; and $r_i = p_i$ otherwise. We consider the same tradeoffs for these bids as in the regular case, except that we restrict attention to the (true) goods in which the bid is interested (the $i \in [n]$ for which $r_i > -\infty$). We then normalise the roots of the bids as described in Section 2.2. We compute the multiplicities of our new bids in the same way as when v is regular, except the definitions of the $m^{ij}(\mathbf{p}; \mathbf{t})$ and $m^{i0}(\mathbf{p})$ functions consider fewer facets than before. (For example, when $n = 2$, we have $m^{10}(\mathbf{p}) = w_v^u(\mathbf{p})$ for vertices \mathbf{p} on the boundary of the box, by contrast with (recall from above) $m^{10}(\mathbf{p}) = w_v^u(\mathbf{p}) - w_v^d(\mathbf{p})$ if \mathbf{p} is in the interior of the box.)

Appendix A.1 illustrates the procedure for finding the non-regular bids when $n = 2$.

Finally, we show uniqueness of \mathcal{B} in the same way as for $n = 2$. That is, if $D_B = D_{B'}$ then $m_B^{ij}(\mathbf{p}; \mathbf{t}) = m_{B'}^{ij}(\mathbf{p}; \mathbf{t})$ for all goods i, j , prices \mathbf{p} and tradeoffs \mathbf{t} . As \mathcal{B} and \mathcal{B}' are parsimonious, there are unique bids $\mathbf{b} \in \mathcal{B}$ and $\mathbf{b}' \in \mathcal{B}'$ with any one pair (\mathbf{r}, \mathbf{t}) of root and tradeoff, where \mathbf{r} is the unique root corresponding to \mathbf{p} . Moreover, we show that $m_B^{ij}(\mathbf{p}; \mathbf{t}) = m_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{t})$ must be the multiplicity of \mathbf{b} and, similarly, $m_{B'}^{ij}(\mathbf{p}; \mathbf{t}) = m_{\mathbf{b}'}^{ij}(\mathbf{p}; \mathbf{t})$ must be the multiplicity of \mathbf{b}' . So $\mathbf{b} = \mathbf{b}'$, that is, the bid in \mathcal{B} with root \mathbf{r} and tradeoff \mathbf{t} has an identical counterpart in \mathcal{B}' . So $\mathcal{B} = \mathcal{B}'$.

APPENDIX B: PROOFS FOR SECTION 2.2

LEMMA 2.1: *If bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ is positive (i.e., $m > 0$), then $D_{\mathbf{b}} = D_{v_{\mathbf{b}}}$, where the valuation $v_{\mathbf{b}} : \mathcal{X} \cap \text{conv}\{mt_i e^i \mid i \in I\} \rightarrow \mathbb{R}$ is defined by $v_{\mathbf{b}}(\mathbf{x}) = \sum_{i \in I} r_i x_i$ and I is the set*

⁵⁷For $n > 2$ we have to use a partial ordering: we consider the effect of a bid \mathbf{b} after the effect of a bid \mathbf{b}' if $F_{\mathbf{b}}^{ij}$ is contained in $F_{\mathbf{b}'}^{ij}$ in the neighbourhood of the root of \mathbf{b} .

of goods in which \mathbf{b} is interested. If \mathbf{b} is negative (i.e., $m < 0$), then $D_{\mathbf{b}} = -D_{|\mathbf{b}|}$ in which $|\mathbf{b}| = (\mathbf{r}; \mathbf{t}; |m|)$.

PROOF: Fix prices $\mathbf{p} \in \mathcal{P}$. Suppose first that $m > 0$. Let $X_{\mathbf{b}} := \mathcal{X} \cap \text{conv}\{mt_i \mathbf{e}^i \mid i \in I\}$ be the domain of $v_{\mathbf{b}}$. Recall $D_{v_{\mathbf{b}}}(\mathbf{p}) = \text{argmax}_{\mathbf{x} \in X_{\mathbf{b}}} (v_{\mathbf{b}}(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x})$ and $D_{\mathbf{b}}(\mathbf{p}) = \mathcal{X} \cap \text{conv}\{mt_i \mathbf{e}^i \mid i \in J\}$, where $J = \text{argmax}_{i \in I} t_i(r_i - p_i)$. Every bundle \mathbf{x} in the domain $X_{\mathbf{b}}$ of $v_{\mathbf{b}}$ is associated with the (unique) coefficients $\lambda(\mathbf{x}) \in [0, 1]^I$ satisfying $\sum_{i \in I} \lambda_i(\mathbf{x}) = 1$ for which $\mathbf{x} = \sum_{i \in I} \lambda_i(\mathbf{x}) mt_i \mathbf{e}^i$. The utility of bundle \mathbf{x} at \mathbf{p} is then $v_{\mathbf{b}}(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x} = m \sum_{i \in I} \lambda_i(\mathbf{x}) t_i(r_i - p_i)$. As $t_i(r_i - p_i)$ is maximised for exactly the goods in J , bundle \mathbf{x} achieves maximal utility if and only if all goods i with $\lambda_i(\mathbf{x}) > 0$ lie in J . Hence, $D_{v_{\mathbf{b}}}(\mathbf{p}) = \mathcal{X} \cap \text{conv}\{mt_i \mathbf{e}^i \mid i \in J\} = D_{\mathbf{b}}(\mathbf{p})$.

For bids $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ with negative multiplicity $m < 0$, we have $D_{\mathbf{b}}(\mathbf{p}) = -D_{|\mathbf{b}|}(\mathbf{p}) = -D_{v_{|\mathbf{b}|}}(\mathbf{p})$, where the first equality holds by definition. Q.E.D.

LEMMA 2.2: For bids $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ and $\mathbf{b}' = (\mathbf{r}'; \mathbf{t}', m')$, we have $D_{\mathbf{b}} = D_{\mathbf{b}'}$ if and only if $\mathbf{t} = \mathbf{t}'$, $m = m'$, they are interested in the same set of goods, I , and $t_i(r_i - r'_i) = t'_j(r_j - r'_j)$ for all $i, j \in I$. In particular, if $0 \in I$, then $D_{\mathbf{b}} = D_{\mathbf{b}'}$ if and only if $\mathbf{b} = \mathbf{b}'$.

PROOF: If \mathbf{b} and \mathbf{b}' satisfy the conditions given, then it is easy to check that $D_{\mathbf{b}} = D_{\mathbf{b}'}$. We now prove the converse, so suppose $D_{\mathbf{b}} = D_{\mathbf{b}'}$. For price \mathbf{p} , let $J_{\mathbf{b}}(\mathbf{p}) := \text{argmax}_{i \in [n]_0} t_i(r_i - p_i)$ be the goods that \mathbf{b} demands at \mathbf{p} . Note that $J_{\mathbf{b}}(\mathbf{p}) \subseteq I$. Fix $\tilde{\mathbf{p}}$ such that $\tilde{p}_i = r_i$ for all $i \in I$. Then $J_{\mathbf{b}}(\tilde{\mathbf{p}}) = I$, and this is the maximal set that $J_{\mathbf{b}}(\mathbf{p})$ can be. Analogously define $\tilde{\mathbf{p}}'$ so that $\tilde{p}'_i = r'_i$ for all $i \in I'$, meaning $J_{\mathbf{b}'}(\tilde{\mathbf{p}}') = I'$ is maximal for $J_{\mathbf{b}'}(\mathbf{p})$. But $D_{\mathbf{b}} = D_{\mathbf{b}'}$ implies $J_{\mathbf{b}} = J_{\mathbf{b}'}$ and so these maximal sets must be equal: $I = I'$. Moreover, now $\mathcal{X} \cap \text{conv}\{mt_i \mathbf{e}^i \mid i \in I\} = D_{\mathbf{b}}(\tilde{\mathbf{p}}) = D_{\mathbf{b}'}(\tilde{\mathbf{p}}) = \mathcal{X} \cap \text{conv}\{m't'_i \mathbf{e}^i \mid i \in I\}$, so that $mt_i = m't'_i$ for all $i \in I$. Since $\mathbf{t}, \mathbf{t}' \in \mathcal{T}$, and in particular are primitive integer vectors, it follows that $\mathbf{t} = \mathbf{t}'$ and $m = m'$. Moreover, $J_{\mathbf{b}}(\tilde{\mathbf{p}}') = I$ implies $t_i(r_i - r'_i) = t_j(r_j - r'_j)$ for all $i, j \in I$. Finally, if $0 \in I$ then $r_0 = r'_0 = 0$ and so $t_i(r_i - r'_i) = 0$ for all $i \in I$, that is, $\mathbf{r} = \mathbf{r}'$. Q.E.D.

PROPOSITION 2.4: Suppose valuation \hat{v} for divisible goods is the concave envelope of concave valuation v for indivisible goods. A bid collection \mathcal{B} satisfies $\hat{D}_{\mathcal{B}} = D_{\hat{v}}$ if and only if $D_{\mathcal{B}} = D_v$.

PROOF: First we recall that Lemma 2.17 of Baldwin and Klemperer (2019b) uses the supporting hyperplane theorem to show that, if \hat{v} is the concave envelope of v , then $\text{conv } D_v(\mathbf{p}) = D_{\hat{v}}(\mathbf{p})$ for all $\mathbf{p} \in \mathcal{P}$. Now suppose $D_{\mathcal{B}} = D_v$; it follows that, for all $\mathbf{p} \in \mathcal{P}$, we have $\hat{D}_{\mathcal{B}}(\mathbf{p}) = \text{conv } D_{\mathcal{B}}(\mathbf{p}) = \text{conv } D_v(\mathbf{p}) = D_{\hat{v}}(\mathbf{p})$. Conversely, suppose that $\hat{D}_{\mathcal{B}} = D_{\hat{v}}$, that is, that $\text{conv } D_{\mathcal{B}}(\mathbf{p}) = \text{conv } D_v(\mathbf{p})$ for all $\mathbf{p} \in \mathcal{P}$. So, for all $\mathbf{p} \in \mathcal{P}$, we have $D_{\mathcal{B}}(\mathbf{p}) = \mathcal{X} \cap \text{conv } D_{\mathcal{B}}(\mathbf{p}) = \mathcal{X} \cap \text{conv } D_v(\mathbf{p}) = D_v(\mathbf{p})$, where the first equality follows by Equation (2.1) and the third by concavity of v and note 12. Q.E.D.

APPENDIX C: DETAILS OF THE GEOMETRY OF DEMAND

Here we provide the necessary mathematical detail on the geometry of demand that is required for the proofs in these appendices. This material is developed with more detail and discussion in Section 2 of Baldwin and Klemperer (2019b). We begin by showing that the LIPs of valuations, bids, and bid collections are the union of $(n-1)$ -dimensional rational polyhedral complexes. Recall that \mathcal{X} consists of all the bundles $\mathbf{x} \in \mathbb{Z}^{[n]_0}$ with $x_0 = 0$.

DEFINITION C.1—Weighted rational polyhedral complexes, (cf. [Mikhalkin 2004](#), Definitions 1,2):

- (i) A *rational polyhedron* is the intersection of one or more half-spaces $\{\mathbf{p} \in \mathcal{P} \mid \mathbf{a} \cdot \mathbf{p} \geq b\}$ for some $\mathbf{a} \in \mathcal{X}$ and $b \in \mathbb{R}$.
- (ii) A *face* of a polyhedron P consists of the price set $\operatorname{argmax}_{\mathbf{p} \in P} \mathbf{a} \cdot \mathbf{p}$ for some fixed $\mathbf{a} \in \mathcal{X}$.
- (iii) A *rational polyhedral complex* Π is a finite collection of *cells* $C \subseteq \mathcal{P}$ such that
 - (a) Every $C \in \Pi$ is a polyhedron, and every face of C is in Π .
 - (b) For any two cells C, C' of Π with non-empty intersection, the intersection $C \cap C'$ is a face of both C and C' .
- (iv) A *k-cell* is a cell of dimension k and a *facet* is a cell of dimension $n - 1$.
- (v) A polyhedral complex is *k-dimensional* if all its cells are contained in its k -cells.
- (vi) A *weighted polyhedral complex* is a pair (Π, w) with a polyhedral complex Π and a function w that assigns a weight $w(F) \in \mathbb{Z}$ to each facet F of Π .

Note that, unlike [Mikhalkin \(2004\)](#), we allow the weights of a polyhedral complex to be negative as well as positive.

DEFINITION C.2—cf. [Mikhalkin \(2004, Definition 3\)](#): An $(n - 1)$ -dimensional rational polyhedral complex (Π, w) is *balanced* if, for every $(n - 2)$ -cell G of Π , the weights $w(F^k)$ of the facets F^1, \dots, F^l that contain G , and primitive integer normal vectors $\mathbf{n}^1, \dots, \mathbf{n}^l$ for these facets defined by a fixed rotational direction about G , satisfy $\sum_{k=1}^l w(F^k) \mathbf{n}^k = 0$.

It is known ([Mikhalkin 2004, Proposition 2.1](#), and [Baldwin and Klemperer 2019b, Proposition 2.7](#) in the economic setting) that \mathcal{L}_v is the union of an $(n - 1)$ -dimensional rational polyhedral complex Π_v (see also [Fact C.4](#) below). We call the facets of Π_v the *facets of \mathcal{L}_v* and associate the following weight function w_v with Π_v and \mathcal{L}_v . Fix some facet F of Π_v , and let \mathbf{x} and \mathbf{x}' be the two bundles (necessarily distinct) that are uniquely demanded on either side of F . We define $w_v(F)$ as the greatest common divisor of the coordinates entries of $\mathbf{x} - \mathbf{x}'$. It is immediate that $w_v(F) > 0$ for all facets F of Π_v . We call (Π_v, w_v) and (\mathcal{L}_v, w_v) the weighted polyhedral complex and weighted LIP of v . Each facet's weight and normal vector jointly specify the change in demand as we cross the facet.

PROPOSITION C.3—[Baldwin and Klemperer \(2019b, Proposition 2.4\)](#): *The change in demand as we change prices to cross a facet F of \mathcal{L}_v is $w_v(F)\mathbf{n}$, where \mathbf{n} is the primitive integer vector that is normal to F and points in the opposite direction to the price change.*

As the net change in demand along any price path that begins and ends at the same point is zero, it follows that (Π_v, w_v) satisfies [Definition C.2](#). Balancing and this positivity are, in fact, the only conditions necessary for a weighted rational polyhedral complex to lead to a weighted LIP (\mathcal{L}_v, w_v) of a valuation.

FACT C.4—The “Valuation-Complex Equivalence Theorem” (see [Mikhalkin 2004, Remark 2.3](#) and [Proposition 2.4](#), and [Baldwin and Klemperer 2019b, Theorem 2.14](#), in the economic context):

- (i) If v is a valuation, then \mathcal{L}_v is the union of an $(n - 1)$ -dimensional rational polyhedral complex Π_v , and (Π_v, w_v) is balanced.
- (ii) A polyhedral complex (Π, w) satisfies both the balancing condition and $w > 0$ if and only if it is the weighted polyhedral complex (Π_v, w_v) of some valuation v . Moreover, in this case, for any combination of price $\tilde{\mathbf{p}} \in \mathcal{P} \setminus \Pi$, bundle $\tilde{\mathbf{x}} \in \mathcal{X}$ and $r \in \mathbb{R}$, there exists a unique concave valuation v satisfying $D_v(\tilde{\mathbf{p}}) = \{\tilde{\mathbf{x}}\}$ and $v(\tilde{\mathbf{x}}) = r$.

We will now see that the weighted LIPs of bids and bid collections also decompose into weighted polyhedral complexes, and satisfy an analogous demand change property. Lemma 2.1 defines a valuation v_b whose demand is the same as that of a single positive bid b so the valuation and bid have the same weighted LIP, which decomposes into polyhedral complex $\Pi_{v_b} =: \Pi_b$. As in Section 3.1, we write w_b for the weight function w_{v_b} of v_b , and we see immediately that weighted polyhedral complex (Π_b, w_b) is balanced and satisfies the demand change described in Proposition C.3. Moreover, a negative bid b has the same LIP as the associated valuation $v_{|b|}$ but now $(\Pi_b, w_b) = (\Pi_{v_{|b|}}, -w_{|b|})$. It is easy to verify that the balancing and demand change properties still hold.

In order to formally define and understand the weighted LIP (\mathcal{L}_B, w_B) of a bid collection B , we first consider an ‘extended’ weighted LIP that includes 0-weighted facets. Let $\mathcal{L}'_B = \bigcup_{b \in B} \mathcal{L}_b$. Note that \mathcal{L}'_B is the union of an $(n-1)$ -dimensional rational polyhedral complex Π'_B , a property that it inherits from the structure of the individual Π_b . By construction, $\mathcal{L}_B \subseteq \mathcal{L}'_B$. We call the connected components of $\mathcal{P} \setminus \mathcal{L}'_B$ the *unique bid demand regions (UBDRs)* of \mathcal{L}'_B . By construction, these UBDRs consists of all the n -dimensional polyhedra that arise as the intersection of a UDR from \mathcal{L}_b for each $b \in B$:

$$\left\{ \bigcap_{b \in B} U^b \mid U^b \text{ is a UDR of } \mathcal{L}_b \text{ and } \bigcap_{b \in B} U^b \text{ is } n\text{-dimensional} \right\}.$$

Every facet of \mathcal{L}'_B is contained in a facet of \mathcal{L}_b for at least one $b \in B$. We define a weight function $w'_B(F) = \sum_{b \in B} w_b(F')$ for every facet F of \mathcal{L}'_B . Here we have extended the weight function w_b of each bid LIP such that $w_b(F) := w_b(F')$ for any $(n-1)$ -dimensional subset F of a facet F' , and $w_b(F) = 0$ for all other subsets. (Note that it is possible for $w'_B(F)$ to be 0 if F is contained in \mathcal{L}_b for multiple bids b with positive and negative multiplicities.)

PROPOSITION C.5: *At all prices in a given UBDR of (\mathcal{L}'_B, w'_B) , the same unique bundle is demanded. The change in demand as we change prices to cross a facet F of \mathcal{L}'_B is $w'_B(F)\mathbf{n}$, where \mathbf{n} is the primitive integer vector that is normal to F and points in the opposite direction to the price change.*

PROOF: Suppose \mathbf{p} and \mathbf{q} are two prices in the same UBDR of \mathcal{L}'_B . By construction, both lie in the same UDR of \mathcal{L}_b for each bid $b \in B$, so every bid uniquely demands the same bundle at \mathbf{p} and \mathbf{q} . By definition of D_B , aggregate demand is thus unique and identical at \mathbf{p} and \mathbf{q} .

Now suppose R and R' are two UBDRs of (\mathcal{L}'_B, w'_B) separated by facet F of \mathcal{L}'_B , and let \mathbf{x} and \mathbf{y} be the unique bundles demanded in R and R' . We will show that $\mathbf{x} - \mathbf{y} = w_B(F)\mathbf{n}$ for the unique primitive vector \mathbf{n} normal to F pointing from R to R' .

Suppose first that B consists of a single bid b . Recall that $(\mathcal{L}_b, w_b) = (\mathcal{L}_{v_b}, w_{v_b})$ if the bid is positive and $(\mathcal{L}_b, w_b) = (\mathcal{L}_{v_{|b|}}, -w_{v_{|b|}})$ otherwise, so the statement follows from Lemma 2.1 and Proposition C.3. Now suppose B contains two or more bids. Let U and V be the two UBDRs of \mathcal{L}'_B on either side of F . For each bid $b \in B$, these UBDRs U and V are either contained in the same UDR of \mathcal{L}_b , or they are contained in neighbouring UDRs of \mathcal{L}_b that are separated only by a facet F^b that contains F . Let \mathbf{x}^b and \mathbf{y}^b be the bundles uniquely demanded by each $b \in B$ at prices in U and V , so that $\mathbf{x} = \sum_{b \in B} \mathbf{x}^b$ and $\mathbf{y} = \sum_{b \in B} \mathbf{y}^b$ (by definition of D_B). Fix some bid $b \in B$. If U and V lie in distinct neighbouring UDRs of \mathcal{L}_b , then we already know that $\mathbf{x}^b - \mathbf{y}^b = w_b(F^b)\mathbf{n} = w_b(F)\mathbf{n}$. Otherwise, i.e., U and V lie in the same UDR of \mathcal{L}_b , it is immediate that $\mathbf{x}^b = \mathbf{y}^b$ and $w_b(F) = 0$ imply the same. Hence $\mathbf{x} - \mathbf{y} = \sum_{b \in B} (\mathbf{x}^b - \mathbf{y}^b) = \sum_{b \in B} w_b(F)\mathbf{n} = w_B(F)\mathbf{n}$, where the last equality holds by definition of $w_B = \sum_{b \in B} w_b$. Q.E.D.

Now let $(\Pi_{\mathcal{B}}, w_{\mathcal{B}})$ be the weighted $(n-1)$ -dimensional rational polyhedral complex obtained by removing all zero-weighted facets from $(\Pi'_{\mathcal{B}}, w'_{\mathcal{B}})$, as well as all lower-dimensional cells that are faces only of zero-weighted facets. Note that removing these cells results in another $(n-1)$ -dimensional complex, so $\Pi_{\mathcal{B}}$ is well-defined. Define $\mathcal{L}_{\mathcal{B}}$ as the union of the polyhedra in $\Pi_{\mathcal{B}}$ to get the weighted LIP $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$. In Section 3.1, $\mathcal{L}_{\mathcal{B}}$ is defined as the set $\{\mathbf{p} \in \mathcal{P} \mid |D_{\mathcal{B}}(\mathbf{p})| \geq 2\}$ of prices at which \mathcal{B} is indifferent between two or more bundles. These two definitions are equivalent: by Proposition C.5, $\mathbf{p} \in \mathcal{P}$ satisfies $|D_{\mathcal{B}}(\mathbf{p})| \geq 2$ if and only if it lies in a facet of $\mathcal{L}'_{\mathcal{B}}$ with non-zero weight; and the latter holds if and only if $\mathbf{p} \in \mathcal{L}_{\mathcal{B}}$.

Note that the UDRs of $\mathcal{L}_{\mathcal{B}}$ may not be convex. However, each UDR of $\mathcal{L}_{\mathcal{B}}$ is the union of one or more convex UDRs of $\mathcal{L}'_{\mathcal{B}}$ with zero-weighted facets separating any two neighbouring UDRs of this union. This allows us to prove an analogue to Proposition C.5 for $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$.

PROPOSITION C.6: *At all prices in a given UDR of $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$, the same unique bundle is demanded. The change in demand as we change prices to cross a facet F of $\mathcal{L}_{\mathcal{B}}$ is $w_{\mathcal{B}}(F)\mathbf{n}$, where \mathbf{n} is the primitive integer vector that is normal to F and points in the opposite direction to the price change.*

PROOF: Suppose \mathbf{p} and \mathbf{q} are two prices in the same UDR of $\mathcal{L}_{\mathcal{B}}$. If \mathbf{p} and \mathbf{q} lie in the same UDR of $(\mathcal{L}'_{\mathcal{B}}, w'_{\mathcal{B}})$, or in neighbouring UDRs separated by a zero-weighted facet, then Proposition C.5 tells us that the same unique bundle is demanded at \mathbf{p} and \mathbf{q} . Otherwise, there exists a finite sequence of prices starting with \mathbf{p} and ending with \mathbf{q} so that every consecutive pair of prices lies in neighbouring UDRs of $(\mathcal{L}'_{\mathcal{B}}, w'_{\mathcal{B}})$ separated by a zero-weighted facet, so the same result holds. For the second statement, observe that facet F of $\mathcal{L}_{\mathcal{B}}$ contains at least one facet F' of $\mathcal{L}'_{\mathcal{B}}$, and these two facets share the same weight and normal vector, so Proposition C.5 implies the statement. *Q.E.D.*

APPENDIX D: PROOFS FOR SECTION 3

We defer the proofs of Theorem 3.1 and Corollaries 3.2, 3.3 and 3.7 to Appendix G. This appendix proves the remaining results from Section 3.

PROPOSITION 3.4: *Concave valuation v and bid collection \mathcal{B} satisfy $(\mathcal{L}_v, w_v) = (\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ and $D_v(\tilde{\mathbf{p}}) = D_{\mathcal{B}}(\tilde{\mathbf{p}})$ for some specific price $\tilde{\mathbf{p}} \in \mathcal{P}$ if and only if $D_v = D_{\mathcal{B}}$.*

PROOF: If $D_v = D_{\mathcal{B}}$, then $\mathcal{L}_v = \mathcal{L}_{\mathcal{B}}$ and $D_v(\tilde{\mathbf{p}}) = D_{\mathcal{B}}(\tilde{\mathbf{p}})$ for any $\tilde{\mathbf{p}}$ are immediate. Moreover, by Propositions C.3 and C.6, $(\mathcal{L}_v, w_v) = (\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$.

Conversely, suppose $(\mathcal{L}_v, w_v) = (\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ and $D_v(\tilde{\mathbf{p}}) = D_{\mathcal{B}}(\tilde{\mathbf{p}})$ for some specific $\tilde{\mathbf{p}} \in \mathcal{P}$. Fix any $\hat{\mathbf{p}} \in \mathcal{P} \setminus \mathcal{L}_v$, let $D_v(\hat{\mathbf{p}}) = \{\mathbf{x}\}$ and $D_{\mathcal{B}}(\hat{\mathbf{p}}) = \{\mathbf{y}\}$, and define $\Delta = \mathbf{x} - \mathbf{y}$. Propositions C.3 and C.6 tell us that the change in demand as we cross facets of $(\mathcal{L}_v, w_v) = (\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ is identical for v and \mathcal{B} , so $D_{\mathcal{B}}(\mathbf{p}) = \Delta + D_v(\mathbf{p})$ for all $\mathbf{p} \in \mathcal{P} \setminus \mathcal{L}_v$. For any price $\mathbf{p} \in \mathcal{L}_v$, both $D_v(\mathbf{p})$ and $D_{\mathcal{B}}(\mathbf{p})$ consist of the integer bundles in the convex hull of demand at prices in the UDRs infinitesimally close to \mathbf{p} (see note 12 and Equation (2.1)), so $D_v(\mathbf{p}) = \Delta + D_{\mathcal{B}}(\mathbf{p})$ holds for all $\mathbf{p} \in \mathcal{P}$. Finally, $D_v(\tilde{\mathbf{p}}) = D_{\mathcal{B}}(\tilde{\mathbf{p}})$ implies $\Delta = \mathbf{0}$. *Q.E.D.*

LEMMA 3.8: *Suppose \mathcal{L}_v is the LIP of a substitutes valuation. Then it is the LIP of a regular valuation if and only if, for all $i, j \in [n]$,*

- (i) every $(i, 0)$ -facet of \mathcal{L}_v is bounded below in all coordinates;
- (ii) every (i, j) -facet of \mathcal{L}_v is bounded above in coordinates i and j .

PROOF: Suppose first that (\mathcal{L}_v, w_v) is the weighted LIP of a regular valuation v . We will show in turn that the two conditions (ii) and (i) hold.

Assume for a contradiction that F is an (i, j) -facet of \mathcal{L}_v with $i, j \in [n]$ that is not upper-bounded in coordinates i and j . By the first regularity condition, there exists $\bar{p}_i \in \mathbb{R}$ such that $x_i = 0$ for all $\mathbf{x} \in D_v(\mathbf{p})$ at prices \mathbf{p} satisfying $p_i > \bar{p}_i$. As F is not upper-bounded in coordinates i and j , its relative interior contains a price \mathbf{p}' with $p'_i > \bar{p}_i$. F has positive weight $w_v(F) > 0$ and is normal to $ae^i - be^j$ for some $a, b \in \mathbb{Z}_{>0}$. So $x_i \neq 0$ for some bundle $\mathbf{x} \in D_v(\mathbf{p}')$ by Proposition C.3, a contradiction.

Now suppose for a contradiction that F is an $(i, 0)$ -facet of \mathcal{L}_v not lower-bounded in coordinate $j \in [n] \setminus \{i\}$, and let H be the hyperplane containing F . Our goal is to show that the existence of F violates the second regularity condition, which stipulates that, for any $\mathbf{p} \in \mathcal{P}$, we have $x_i = 0$ for all bundles $\mathbf{x} \in D_v(\mathbf{p} - \lambda e^j)$ when λ is sufficiently large. To do this, we now find prices $\mathbf{p} \in F$ so that the half-ray $\mathbf{p} - \lambda e^j$ with $\lambda > 0$ lies entirely in facets contained in H . By Proposition C.3, this then implies, for every $\lambda > 0$, the existence of a bundle \mathbf{x} with $x_i \neq 0$ that is demanded at $\mathbf{p} - \lambda e^j$: a contradiction.

Firstly, every $(j, 0)$ -facet and (i, j) -facet which meets F does so in a face for which j is fixed; as \mathcal{L}_v has finitely many faces, and F is not bounded below in coordinate j , we can choose \mathbf{p} in the relative interior of F with p_j below any such intersections. Secondly, \mathcal{L}_v has finitely many (i, k) -, (k, l) - and $(k, 0)$ -facets with $k, l \in [n]_0 \setminus \{i, j\}$ and we can make small changes in p_k (without altering the fact that $\mathbf{p} \in F$) to ensure that $\mathbf{p} - \lambda e^j$ is disjoint from all such facets. So the only facets of \mathcal{L}_v not in H which might meet our half-ray $\mathbf{p} - \lambda e^i$ are (j, k) -facets with $k \in [n] \setminus \{i, j\}$.

If the half-ray is contained entirely in F , we are done. So suppose not. Let \mathbf{p}' be the point of intersection of the half-ray with the boundary of F , and let G be an $(n-2)$ -dimensional face of F containing \mathbf{p}' . By construction, G is the intersection of H and a (j, k) -facet F' of \mathcal{L}_v with $k \notin \{i, 0\}$ and normal vector $t_j e^j - t_k e^k$, so the affine span of G is $\left\{ e^l, \frac{1}{t_j} e^k + \frac{1}{t_k} e^k \mid l \in [n] \setminus \{i, j, k\} \right\}$. Observe that G therefore does not lie in any other (j, k) -facet with different slope $\sigma \neq \frac{t_j}{t_k}$, or any (j, l) -facet with $l \in [n] \setminus \{i, j, k\}$. As the normal vectors of F and H are linearly independent and (\mathcal{L}_v, w_v) is balanced, H contains a facet of the same weight on either side of G . So, after crossing G , our half-ray $\mathbf{p} - \lambda e^j$ lies in another facet in H of the same weight as F . Apply this argument repeatedly to see that $\mathbf{p} - \lambda e^j$ is contained in such a facet for all $\lambda > 0$, so regularity is violated.

Now suppose that \mathcal{L}_v satisfies conditions (i) and (ii). There are in fact many valuations giving rise to the same LIP: we can add a constant bundle to that which is demanded in any UDR, and we can add a constant scalar to the valuation. We need to show that one of these valuations v is regular. In order to show that v satisfies the first regularity condition, we first fix $\bar{\mathbf{p}} = \lambda \mathbf{1}$ and argue that all prices $\mathbf{p} \geq \bar{\mathbf{p}}$ lie in the same UDR of \mathcal{L}_v if we choose $\lambda > 0$ sufficiently large. As \mathcal{L}_v has a finite number of $(i, 0)$ -facets, we can choose λ greater than the i -coordinate of each such facet. And by condition (ii), for every $i, j \in [n]_0$ we can choose λ greater than the upper bound (in coordinate i or j) for every (i, j) -facet. Now there are no facets in the price region $\mathbf{p} \geq \bar{\mathbf{p}}$, and so this price region must be contained within a single UDR. Then, by Fact C.4, there exists a (concave) valuation—call it v without loss of generality—with LIP \mathcal{L}_v that demands 0 in this UDR. Suppose also that $v(\mathbf{0}) = 0$.

Now, fix a good i and prices \mathbf{p} satisfying $p_i \geq \bar{p}_i$, and suppose $\mathbf{x} \in D_v(\mathbf{p})$. If $\mathbf{p} \geq \bar{\mathbf{p}}$ then $x_i = 0$ because $\mathbf{x} = \mathbf{0}$. So suppose $\mathbf{p} \not\geq \bar{\mathbf{p}}$, and let \mathbf{p}' be the prices obtained by setting $p'_j = \max\{p_j, \bar{p}_j\}$ for all goods $j \in [n] \setminus \{i\}$. We can consider the change in price from \mathbf{p} to \mathbf{p}' as an increase in price in each of the goods $j \neq i$ in turn. It is possible that demand is not unique at \mathbf{p} , at \mathbf{p}' , or at intermediate prices in this sequence of price increases, but we can make a small

generic perturbation of all the prices in this chain so that we have a sequence of increases in the price of one good $j \neq i$ between two prices at which demand is unique. Now the substitutes property tells us that this overall price change from prices near \mathbf{p} to prices near \mathbf{p}' cannot decrease the demand of good i . As we know that the demand of good i is 0 at or near \mathbf{p}' , it is thus also 0 near \mathbf{p} . By taking small generic perturbations in enough directions, we can infer that this holds for all bundles demanded at prices close to \mathbf{p} , and so that $x_i = 0$ itself. Thus v satisfies the first regularity condition (with threshold \bar{p}).

In particular, this also implies that this valuation v demands only non-negative bundles at all prices, as increasing the price of any one good sufficiently leads to a demand of 0 for that good, and as D_v satisfies the law of demand (cf. Definition 4.2).

Now we argue that v also satisfies the second regularity condition. Fix a good i and prices \mathbf{p} . We want to show that there exists large $\bar{\lambda} > 0$ such that demand of goods $j \in [n] \setminus \{i\}$ is 0 at prices $\mathbf{p} - \lambda \mathbf{e}^i$ for all $\lambda \geq \bar{\lambda}$. We will argue this by first identifying sufficiently large $\bar{\lambda}$ and then considering prices $\mathbf{p}', \mathbf{p}''$, where $\mathbf{p}' = \mathbf{p} - \lambda \mathbf{e}^i$ with $\lambda \geq \bar{\lambda}$, and \mathbf{p}'' satisfies $p'_i = p'_i$ and $p''_j \geq p'_j$.

First, we know from condition (i) that all $(j, 0)$ -facets are bounded below in coordinate i , as (by their nature) are all $(i, 0)$ -facets, so for λ exceeding large enough $\bar{\lambda}$ we know that every $\mathbf{p}', \mathbf{p}''$ with $p'_i = p''_i = p_i - \lambda$ will be below all such facets. Second, consider (i, j) -facets where $j \in [n] \setminus \{i\}$. Such a facet lies in a hyperplane $H = \{\mathbf{p} \in \mathbb{R}^n \mid t_i p_i - t_j p_j = \alpha\}$ for some $\alpha > 0$. Observe that for sufficiently large λ , prices $\mathbf{p}' = \mathbf{p} - \lambda \mathbf{e}^i$ always satisfy $t_i p'_i - t_j p'_j < \alpha$ and moreover if $p''_i = p'_i$ and $p''_j \geq p'_j$ then also $t_i p''_i - t_j p''_j < \alpha$. As every LIP has finitely many facets, we can pick $\bar{\lambda}$ large enough that these conditions are satisfied for all such facets and all $\lambda > \bar{\lambda}$. So these facets do not separate such \mathbf{p}' and \mathbf{p}'' . Now, for convenience, fix $\mathbf{p}' = \mathbf{p} - \lambda \mathbf{e}^i$ for one such λ , and let \mathbf{x}' be a demanded bundle at \mathbf{p}' . We wish to show that $x'_j = 0$ for all $j \in [n] \setminus \{i\}$.

Let \mathbf{p}'' be the prices obtained from \mathbf{p}' by setting $p''_i = p'_i$ and $p''_j = \max\{p'_j, \bar{p}_j\}$ for every $j \in [n] \setminus \{i\}$, and let \mathbf{x}'' be demanded at \mathbf{p}'' . Observe $p''_j \geq p'_j$ for all $j \neq i$. As v satisfies the first regularity condition with thresholds \bar{p} , we see that $x''_j = 0$ for $j \in [n] \setminus \{i\}$. By our choice of $\bar{\lambda}$, by changing prices from \mathbf{p}' to \mathbf{p}'' we do not cross any $(j, 0)$ or (i, j) -facets, for any $j \in [n] \setminus \{i\}$. Now we see that we also do not cross (j, k) -facets with goods $j, k \in [n] \setminus \{i\}$. If we were to do so, then demand would transfer between goods j and k (at tradeoffs specified by the facet's normal vector). But as $x''_j = x''_k = 0$, this would imply that demand for either good j or good k were negative on the other side of such a facet. We have already seen that there are no such bundles in the domain of v . So we cross no facets between \mathbf{p}' and \mathbf{p}'' , and so $x'_j = 0$ for all $j \neq i$ in $[n]$. *Q.E.D.*

COROLLARY D.1: *The LIP $\mathcal{L}_{\mathcal{B}}$ of a regular bid collection satisfies the conditions of Lemma 3.8.*

PROOF: First suppose that \mathcal{B} contains only one bid, \mathbf{b} . If \mathbf{b} has positive multiplicity then $D_{\mathbf{b}} = D_{v_{\mathbf{b}}}$ by Lemma 2.1, and $v_{\mathbf{b}}$ clearly satisfies Definition 3.6, and so $\mathcal{L}_{\mathbf{b}}$ satisfies the conditions of Lemma 3.8. The case of negative multiplicity \mathbf{b} follows because $\mathcal{L}_{\mathbf{b}} = \mathcal{L}_{|\mathbf{b}|}$. The result follows for arbitrary \mathcal{B} because $\mathcal{L}_{\mathcal{B}} \subseteq \bigcup_{\mathbf{b} \in \mathcal{B}} \mathcal{L}_{\mathbf{b}}$ and all facets of all LIPs $\mathcal{L}_{\mathbf{b}}$ are bounded. *Q.E.D.*

The following corollary is also useful to prove Corollaries 3.3 and 3.7.

COROLLARY D.2:

(i) *If there exists $\mathbf{p} \in \mathcal{P}$, $J \subseteq [n]$ and $\mathbf{t} \in \mathcal{T}$ such that $\mathbf{p} + \lambda \sum_{i \in J} \frac{\mathbf{e}^i}{t_i} \in \mathcal{L}_v$ for every $\lambda > 0$, then part (i) of Definition 3.6 (regular valuations) fails.*

(ii) If there exists $\mathbf{p} \in \mathcal{P}$ and $i \in [n]$ such that $\mathbf{p} - \lambda \mathbf{e}^i \in \mathcal{L}_v$ for every $\lambda > 0$, then part (ii) of Definition 3.6 (regular valuations) fails.

PROOF: Suppose (i) holds. Then, because it is contained in the (finite) set of facets of \mathcal{L}_v , the half-ray $\mathbf{p} + \lambda \sum_{i \in J} \frac{\mathbf{e}^i}{t_i} \in \mathcal{L}_v$ must at all points be contained in some (j, k) -facet for some distinct $j, k \in J \subseteq [n]$. As there are finitely many such facets, there must be at least one such that is unbounded above, and in particular unbounded above in coordinate j . But this facet has positive weight and so $x_j \neq 0$ for some bundle demanded everywhere in this facet, by Proposition C.3. This implies failure of part (i) of Definition 3.6.

Suppose (ii) holds, so $\mathbf{p} - \lambda \mathbf{e}^i \in \mathcal{L}_v$ for all $\lambda > 0$. As \mathcal{L}_v is a finite union of facets, it follows that any point in $\mathbf{p} - \lambda \mathbf{e}^i$ must be contained in a (j, k) -facet F for distinct $j, k \in [n]_0 \setminus \{i\}$ where $j \neq 0$. But each such facet has positive weight and so $x_j \neq 0$ for some bundle $\mathbf{x} \in D_v(\mathbf{p} - \lambda \mathbf{e}^i)$ by Proposition C.3, so that part (ii) of Definition 3.6 fails. Q.E.D.

APPENDIX E: PROOFS FOR SECTION 4

PROPOSITION 4.3: For any bid collection \mathcal{B} , the following statements are equivalent.

- (i) Bid collection \mathcal{B} is valid, i.e., there exists a valuation v such that $D_v = D_{\mathcal{B}}$.
- (ii) There exists a concave substitutes valuation v such that $D_v = D_{\mathcal{B}}$.
- (iii) The valuation $v_{\mathcal{B}}$ is well-defined and satisfies $D_{v_{\mathcal{B}}} = D_{\mathcal{B}}$.
- (iv) $D_{\mathcal{B}}$ satisfies the law of demand.
- (v) The valuation $v_{\mathcal{B}}$ satisfies $\pi_{v_{\mathcal{B}}} = \pi_{\mathcal{B}}$.
- (vi) The indirect utility function $\pi_{\mathcal{B}}$ of \mathcal{B} is convex.
- (vii) For all $\mathbf{p} \in \mathcal{P}$, we have $D_{\mathcal{B}}(\mathbf{p}) = D_{\mathcal{B}^+}(\mathbf{p}) - D_{|\mathcal{B}^-|}(\mathbf{p})$.
- (viii) The weight $w_{\mathcal{B}}(F)$ of every facet of $\mathcal{L}_{\mathcal{B}}$ is positive.

PROOF: We first show (iv) \implies (viii). For the purpose of contradiction, suppose F is a facet of $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ with $w_{\mathcal{B}}(F) < 0$. By construction, F is normal to a vector $a\mathbf{e}^i - b\mathbf{e}^j$ for some $i, j \in [n]_0$ with $i > j$ and $a, b \in \mathbb{Z}_{>0}$. Pick a point \mathbf{r} in the relative interior of F , and let $\mathbf{p} := \mathbf{r} - \varepsilon \mathbf{e}^i$ and $\mathbf{q} := \mathbf{r} + \varepsilon \mathbf{e}^i$ for some infinitesimal $\varepsilon > 0$. We will show that the law of demand (cf. Definition 4.2) is violated for \mathbf{p} and \mathbf{q} . By construction, \mathbf{p} and \mathbf{q} lie in the UDRs of $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ on either side of F . By Proposition C.6, we have $\mathbf{x} - \mathbf{y} = w_{\mathcal{B}}(F)(a\mathbf{e}^i - b\mathbf{e}^j)$ for the bundles \mathbf{x} and \mathbf{y} uniquely demanded at \mathbf{p} and \mathbf{q} . As $x_i - y_i = w_{\mathcal{B}}(F)a > 0$ and so $x_i > y_i$, the law of demand is violated.

Next we argue (viii) \implies (ii). Recall that $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ can be decomposed into an $(n - 1)$ -dimensional rational polyhedral complex $\Pi_{\mathcal{B}}$, and that $w_{\mathcal{B}}(F) > 0$ for every facet F of $\mathcal{L}_{\mathcal{B}}$ (which is, by definition, also a facet of $\Pi_{\mathcal{B}}$). Moreover, $(\Pi_{\mathcal{B}}, w_{\mathcal{B}})$ is balanced for any bid collection \mathcal{B} . Fix some $\tilde{\mathbf{p}} \in \mathcal{P} \setminus \mathcal{L}_{\mathcal{B}}$ and $\{\tilde{\mathbf{x}}\} = D_{\mathcal{B}}(\tilde{\mathbf{p}})$. By Fact C.4, there exists a concave valuation v such that $(\mathcal{L}_v, w_v) = (\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$, $v(\tilde{\mathbf{x}}) = 0$, and $D_v(\tilde{\mathbf{p}}) = \{\tilde{\mathbf{x}}\}$, so Proposition 3.4 tells us that $D_v = D_{\mathcal{B}}$. Finally, as the facets of $\mathcal{L}_{\mathcal{B}}$ (and so \mathcal{L}_v) are normal to vectors of type $a\mathbf{e}^i - b\mathbf{e}^j$ for some $i, j \in [n]_0$ and $a, b \in \mathbb{N}$, we have that v is a substitutes valuation (cf. Section 3.1).

It is immediate that (ii) implies (iv) and (i). Next, we show (i) \implies (iii). Recall (defined above Proposition C.5) that $(\mathcal{L}'_{\mathcal{B}}, w'_{\mathcal{B}})$ is the ‘extended’ weighted LIP that contains zero-weighted facets, and that $\mathcal{L}_{\mathcal{B}} \subseteq \mathcal{L}'_{\mathcal{B}} = \bigcup_{b \in \mathcal{B}} \mathcal{L}_b$. We modify the notation for the definition of $v_{\mathcal{B}}$ used in the text to incorporate the role of prices: for any price $\mathbf{p} \in \mathcal{P} \setminus \mathcal{L}'_{\mathcal{B}}$, let $i(\mathbf{b}, \mathbf{p})$ denote the unique good that each $\mathbf{b} \in \mathcal{B}$ demands at \mathbf{p} . For the valuation v with $D_v = D_{\mathcal{B}}$, which exists by assumption, we start by showing that for every bundle \mathbf{x} demanded in a UDBR of $\mathcal{L}'_{\mathcal{B}}$, and

every price \mathbf{p} within these UBDRs at which \mathbf{x} is demanded ($\mathbf{x} \in D_v(\mathbf{p})$), we have

$$v(\mathbf{x}) = v_{\mathcal{B}}(\mathbf{x}) = \sum_{(r;t;m)=b \in \mathcal{B}} mt_{i(b,\mathbf{p})} r_{i(b,\mathbf{p})}. \quad (\text{E.1})$$

To see this, first fix a specific bundle $\tilde{\mathbf{x}}$ demanded in a UBDR of $\mathcal{L}'_{\mathcal{B}}$, and a price $\tilde{\mathbf{p}}$ in this UBDR. We can assume that Equation (E.1) holds for $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{p}}$, by adding a constant term to v if necessary (which does not change the demand of v , cf. [Mikhalkin 2004](#), Remark 2.3). Secondly, $i(\mathbf{b}, \mathbf{p}) = i(\mathbf{b}, \mathbf{p}')$ for any two prices \mathbf{p}, \mathbf{p}' in the same UBDR of $\mathcal{L}'_{\mathcal{B}}$, for any bid $\mathbf{b} \in \mathcal{B}$, so if Equation (E.1) holds for \mathbf{x} and \mathbf{p} , then it also holds for \mathbf{x} and \mathbf{p}' . Now suppose \mathbf{p}, \mathbf{p}' are in neighbouring UBDRs of $(\mathcal{L}'_{\mathcal{B}}, w'_{\mathcal{B}})$, and let \mathbf{x} and \mathbf{x}' be the (not necessarily distinct) two bundles demanded in these UBDRs. If price vector \mathbf{q} is on the facet separating the two UBDRs, then each bid $\mathbf{b} = (r; t; m) \in \mathcal{B}$ demands both goods $i(\mathbf{b}, \mathbf{p})$ and $i(\mathbf{b}, \mathbf{p}')$ at \mathbf{q} , so

$$m \max_{i \in [n]_0} t_i (r_i - q_i) = mt_{i(\mathbf{b}, \mathbf{p})} (r_{i(\mathbf{b}, \mathbf{p})} - q_{i(\mathbf{b}, \mathbf{p})}) = mt_{i(\mathbf{b}, \mathbf{p}')} (r_{i(\mathbf{b}, \mathbf{p}')} - q_{i(\mathbf{b}, \mathbf{p}')}).$$

Taking the sum over all bids and recalling that $mt_{i(\mathbf{b}, \mathbf{p})}$ is the number of units of good $i(\mathbf{b}, \mathbf{p})$ that bid $\mathbf{b} = (r, \mathbf{p}, m)$ demands at \mathbf{p} , we see

$$\sum_{(r;t;m) \in \mathcal{B}} mt_{i(\mathbf{b}, \mathbf{p})} r_{i(\mathbf{b}, \mathbf{p})} - \mathbf{q} \cdot \mathbf{x} = \sum_{(r;t;m) \in \mathcal{B}} mt_{i(\mathbf{b}, \mathbf{p}')} r_{i(\mathbf{b}, \mathbf{p}')} - \mathbf{q} \cdot \mathbf{x}'. \quad (\text{E.2})$$

But as \mathbf{x} and \mathbf{x}' are bundles respectively demanded by \mathcal{B} , and hence by v , at \mathbf{p} and \mathbf{p}' , both are demanded by v at prices \mathbf{q} on the facet separating the two UBDRs, so $v(\mathbf{x}) - \mathbf{q} \cdot \mathbf{x} = v(\mathbf{x}') - \mathbf{q} \cdot \mathbf{x}'$. Subtracting this expression from Equation (E.2) demonstrates that if Equation (E.1) holds for \mathbf{x} and \mathbf{p} , then it also holds for \mathbf{x}' and \mathbf{p}' . Inductively apply this result to see that Equation (E.1) holds for all bundles demanded within UBDRs, and all prices in these respective UBDRs.

For any bundle \mathbf{x} that is demanded at $\mathbf{p} \in \mathcal{L}'_{\mathcal{B}}$, we can perturb generically to get prices \mathbf{p}' and note that the unique bundle \mathbf{x}' demanded at \mathbf{p}' is also demanded at \mathbf{p} . This implies $v(\mathbf{x}) - v(\mathbf{x}') = \mathbf{p} \cdot (\mathbf{x} - \mathbf{x}')$. But this corresponds to the definition of $v_{\mathcal{B}}(\mathbf{x})$, so $v_{\mathcal{B}}(\mathbf{x}) = v(\mathbf{x})$ for all bundles demanded by \mathcal{B} .

We now show that (iii) \implies (v). For prices \mathbf{p} not in $\mathcal{L}'_{\mathcal{B}}$, the uniquely demanded bundle at \mathbf{p} is $\mathbf{x} = \sum_{(r;t;m)=b \in \mathcal{B}} mt_{i(b)} e^{i(b)}$. So, applying the definition of $\pi_{\mathcal{B}}$ and Equation (E.1), we see that $\pi_{\mathcal{B}}(\mathbf{p}) = v_{\mathcal{B}}(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x} = \pi_{v_{\mathcal{B}}}(\mathbf{p})$. Moreover, as $\mathcal{P} \setminus \mathcal{L}'_{\mathcal{B}}$ is dense, it then follows by continuity of $\pi_{\mathcal{B}}$ and $\pi_{v_{\mathcal{B}}}$ that $\pi_{\mathcal{B}}(\mathbf{p}) = \pi_{v_{\mathcal{B}}}(\mathbf{p})$ for every $\mathbf{p} \in \mathcal{P}$.

It is well-known that the indirect utility function of a valuation is convex, so (v) \implies (vi).

We show that (vi) \implies (viii), so suppose π_b is convex. Fix a facet F of $(\mathcal{L}_{\mathcal{B}}, w_{\mathcal{B}})$ normal to $ae^i - be^j$ for some $i, j \in [n]_0$ with $i > j$ and $a, b \in \mathbb{N}$. Let \mathbf{p} be a point in the relative interior of F and define $\mathbf{q}^{\pm} := \mathbf{p} \pm \varepsilon e^i$ for some infinitesimal $\varepsilon > 0$. Note that \mathbf{q}^+ and \mathbf{q}^- lie in the UDRs of $\mathcal{L}_{\mathcal{B}}$ on either side of F . Let \mathbf{x}^+ and \mathbf{x}^- be the bundles demanded at \mathbf{q}^+ and \mathbf{q}^- . We define $\hat{\mathcal{B}}^-$ as the bids in \mathcal{B} that demand good i at \mathbf{q}^- , and $\hat{\mathcal{B}}^+$ likewise, so that $x_i^- = \sum_{(r;t;m) \in \hat{\mathcal{B}}^-} mt_i$ and $x_i^+ = \sum_{(r;t;m) \in \hat{\mathcal{B}}^+} mt_i$. Proposition C.6 tells us that the demand change between prices \mathbf{q}^{\pm} is given by $w(F)(ae^i - be^j) = \mathbf{x}^- - \mathbf{x}^+$, so $w(F)$ is non-negative if and only if $x_i^+ \leq x_i^-$. We will now show that this holds.

By construction, each $\mathbf{b} = (r; t; m)$ demands $i(\mathbf{b}, \mathbf{q}^-)$ at \mathbf{q}^- and at \mathbf{p} , so this good maximises $t_i(r_i - p_i)$ and $t_i(r_i - q_i^-)$ over $i \in [n]_0$ (by definition of D_b). It follows that $\pi_{\mathcal{B}}(\mathbf{p}) = \sum_{(r;t;m) \in \mathcal{B}} mt_{i(\mathbf{b}, \mathbf{q}^-)} (r_{i(\mathbf{b}, \mathbf{q}^-)} - p_{i(\mathbf{b}, \mathbf{q}^-)})$ and $\pi_{\mathcal{B}}(\mathbf{q}^-) = \sum_{(r;t;m) \in \mathcal{B}} mt_{i(\mathbf{b}, \mathbf{q}^-)} (r_{i(\mathbf{b}, \mathbf{q}^-)} - q_{i(\mathbf{b}, \mathbf{q}^-)})$

(by definition of $\pi_{\mathcal{B}}$ in Equation (4.1) of Section 4.2). Hence,

$$\pi_{\mathcal{B}}(\mathbf{p}) - \pi_{\mathcal{B}}(\mathbf{q}^-) = \sum_{(r;t;m) \in \mathcal{B}} mt_{i(b,q^-)}(q_{i(b,q^-)}^- - p_{i(b,q^-)}) = -\varepsilon \sum_{(r;t;m) \in \widehat{\mathcal{B}}^-} mt_i = -\varepsilon x_i^-.$$

The second equality uses the fact that $q_i^- = p_i - \varepsilon$ and $q_j^- = p_j$ for all $j \neq i$, so $q_{i(b,q^-)}^- = p_{i(b,q^-)}$ for all $b \notin \widehat{\mathcal{B}}^-$. Analogously, we see that $\pi_{\mathcal{B}}(\mathbf{p}) - \pi_{\mathcal{B}}(\mathbf{q}^+) = \varepsilon x_i^+$. As $\pi_{\mathcal{B}}$ is convex, it satisfies midpoint convexity $2\pi_{\mathcal{B}}(\mathbf{p}) \leq \pi_{\mathcal{B}}(\mathbf{q}^-) + \pi_{\mathcal{B}}(\mathbf{q}^+)$. Rearranging and substituting, this implies $\varepsilon x_i^+ - \varepsilon x_i^- = 2\pi_{\mathcal{B}}(\mathbf{p}) - \pi_{\mathcal{B}}(\mathbf{q}^-) - \pi_{\mathcal{B}}(\mathbf{q}^+) \leq 0$, and so $x_i^+ \leq x_i^-$ as required.

Finally, we show that (i) \implies (vii) \implies (viii) in Proposition E.4 below. Q.E.D.

Suppose A and B are two finite subsets of \mathcal{P} . The *Minkowski sum* $A + B$ of A and B consists of all points $\mathbf{x} + \mathbf{y}$ with $\mathbf{x} \in A$ and $\mathbf{y} \in B$. We use \sum to denote the Minkowski sum of multiple sets. The *Minkowski difference* $A - B$ consists of all points $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{x} + B \subseteq A$. (Note that $A - B$ is not equal to the Minkowski sum of A and $-B$.)

If \mathbf{p} are prices at which every bid demands a unique bundle, then $D_{\mathcal{B}}(\mathbf{p})$ is a single bundle consisting of the (signed) sum of the bundles demanded by all the bids. Proposition E.4 tells us that, in general, $D_{\mathcal{B}}(\mathbf{p})$ is the Minkowski difference between the demands of the positive and negative bids of \mathcal{B} at \mathbf{p} . To prove Proposition E.4, we make use of the following technical lemma about Minkowski differences. We also recall some facts from polyhedral theory.

FACT E.1—See, e.g., Baldwin et al. (2024) Lemma 3 and Schneider (2013) Lemma 3.1.11: Suppose $A, B \subseteq \mathbb{R}^n$ are nonempty finite sets satisfying $\mathbb{Z}^n \cap (\text{conv } A) = A$ and $\mathbb{Z}^n \cap (\text{conv } B) = B$. Then $\mathbb{Z}^n \cap (\text{conv } A - \text{conv } B) = A - B$. Moreover, if $C, D \subseteq \mathbb{R}^n$ are non-empty, compact, convex sets, then $(C + D) - D = C$.

FACT E.2—See e.g. Grünbaum 1967, Theorem 2.4.9 and Exercise 15.1.1: For any convex polytope P and vertex \mathbf{v} of P , there exists $\mathbf{d} \in \mathbb{R}^n$ such that \mathbf{v} is the unique minimiser of the linear function $\mathbf{d} \cdot \mathbf{x}$ over P . Moreover, \mathbf{v} also minimises $\mathbf{d}' \cdot \mathbf{x}$ over P for any \mathbf{d}' obtained by infinitesimally perturbing entries of \mathbf{d} . Conversely, for any generic $\mathbf{d} \in \mathbb{R}^n$, the minimiser \mathbf{v} of $\mathbf{d} \cdot \mathbf{x}$ over P is unique and a vertex. If P is the Minkowski sum of two (convex) polytopes A and B , then \mathbf{v} is the unique sum of two vertices of A and B , and these two vertices uniquely minimise $\mathbf{d} \cdot \mathbf{x}$ over A and B , respectively.

Fact E.2 allows us to prove:

LEMMA E.3: If \mathcal{B}^0 and \mathcal{B}^1 are valid bid collections, then $\mathcal{B}^0 \cup \mathcal{B}^1$ is valid and, for all $\mathbf{p} \in \mathcal{P}$,

$$\text{conv } D_{\mathcal{B}^0 \cup \mathcal{B}^1}(\mathbf{p}) = \text{conv } D_{\mathcal{B}^0}(\mathbf{p}) + \text{conv } D_{\mathcal{B}^1}(\mathbf{p})$$

PROOF: We show that the extremal points of each side of this equation are contained in the other. Write $\mathcal{B} := \mathcal{B}^0 \cup \mathcal{B}^1$. We know \mathcal{B} is valid by Proposition 4.3 (viii) \Leftrightarrow (i), and by noting that the weight of any facet of $\mathcal{L}_{\mathcal{B}}$ is the sum of the weights of facets of $\mathcal{L}_{\mathcal{B}^0}$ and $\mathcal{L}_{\mathcal{B}^1}$ that contain it.

By Fact E.2, every extremal point of $\text{conv } D_{\mathcal{B}^0}(\mathbf{p}) + \text{conv } D_{\mathcal{B}^1}(\mathbf{p})$ uniquely minimises $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over $\text{conv } D_{\mathcal{B}^0}(\mathbf{p}) + \text{conv } D_{\mathcal{B}^1}(\mathbf{p})$ for some $\mathbf{q} \in Q(\mathbf{p})$, and we can choose \mathbf{q} so that the minimum of $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over $\text{conv } D_{\mathcal{B}}(\mathbf{p})$ is also unique. Conversely, every extremal point of $\text{conv } D_{\mathcal{B}}(\mathbf{p})$ uniquely minimises $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over $\text{conv } D_{\mathcal{B}}(\mathbf{p})$ for some $\mathbf{q} \in Q(\mathbf{p})$, and we can choose \mathbf{q} so that the minimum of $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over $\text{conv } D_{\mathcal{B}^0}(\mathbf{p}) + \text{conv } D_{\mathcal{B}^1}(\mathbf{p})$ is also

unique. Moreover, if \hat{x} uniquely minimises $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over $\text{conv } D_{\mathcal{B}^0}(\mathbf{p}) + \text{conv } D_{\mathcal{B}^1}(\mathbf{p})$ for some $\mathbf{q} \in Q(\mathbf{p})$ then $\hat{x} = \mathbf{x}^0 + \mathbf{x}^1$ where \mathbf{x}^0 uniquely minimises $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over $\text{conv } D_{\mathcal{B}^0}(\mathbf{p})$ and \mathbf{x}^1 uniquely minimises $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over $\text{conv } D_{\mathcal{B}^1}(\mathbf{p})$.

So we can consider all extremal points of $\text{conv } D_{\mathcal{B}^0}(\mathbf{p}) + \text{conv } D_{\mathcal{B}^1}(\mathbf{p})$ and of $\text{conv } D_{\mathcal{B}}(\mathbf{p})$ by considering unique minimisers \mathbf{x}^0 , \mathbf{x}^1 and \mathbf{x} of $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ for appropriately chosen \mathbf{q} over, respectively, domains $\text{conv } D_{\mathcal{B}^0}(\mathbf{p})$, $\text{conv } D_{\mathcal{B}^1}(\mathbf{p})$ and $\text{conv } D_{\mathcal{B}}(\mathbf{p})$. We will show that $\mathbf{x}^0 + \mathbf{x}^1 = \mathbf{x}$. As extremal points of $\text{conv } D_{\mathcal{B}'}(\mathbf{p})$ are in $D_{\mathcal{B}'}(\mathbf{p})$ for any bid collection \mathcal{B}' , it follows that $\mathbf{x}^0 + \mathbf{x}^1 \in D_{\mathcal{B}}(\mathbf{p}) \subseteq \text{conv } D_{\mathcal{B}}(\mathbf{p})$ and that $\mathbf{x} \in D_{\mathcal{B}^0}(\mathbf{p}) + D_{\mathcal{B}^1}(\mathbf{p}) \subseteq \text{conv } D_{\mathcal{B}^0}(\mathbf{p}) + \text{conv } D_{\mathcal{B}^1}(\mathbf{p})$, as required.

So fix \mathbf{q} such that we have unique minimisers \mathbf{x} , \mathbf{x}^0 and \mathbf{x}^1 of $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ over, respectively, domains $\text{conv } D_{\mathcal{B}}(\mathbf{p})$, $\text{conv } D_{\mathcal{B}^0}(\mathbf{p})$ and $\text{conv } D_{\mathcal{B}^1}(\mathbf{p})$. Moreover, it is without loss of generality to choose $\mathbf{q} \in Q(\mathbf{p})$ arbitrarily close to \mathbf{p} (as it is only the vector direction of $\mathbf{q} - \mathbf{p}$ which is important). As these are vertices of integer polytopes, we know that $\mathbf{x}, \mathbf{x}^0, \mathbf{x}^1 \in \mathcal{X}$. We now argue that \mathbf{x} is the unique bundle demanded by \mathcal{B} at \mathbf{q} . Firstly, $\mathbf{x} \in D_{\mathcal{B}}(\mathbf{p})$ implies $v(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x} \geq v(\mathbf{x}') - \mathbf{p} \cdot \mathbf{x}'$ for all $\mathbf{x}' \in \mathcal{X}$, where v is the valuation with the same demand as \mathcal{B} whose existence follows from validity of \mathcal{B} . Secondly, $(\mathbf{q} - \mathbf{p}) \cdot \mathbf{x} < (\mathbf{q} - \mathbf{p}) \cdot \mathbf{x}'$ for any $\mathbf{x}' \in D_{\mathcal{B}}(\mathbf{p})$ thus implies $v(\mathbf{x}) - \mathbf{q} \cdot \mathbf{x} > v(\mathbf{x}') - \mathbf{q} \cdot \mathbf{x}'$ for all $\mathbf{x}' \in D_{\mathcal{B}}(\mathbf{p})$. As \mathbf{q} is close to \mathbf{p} , we know that $D_{\mathcal{B}}(\mathbf{q}) \subseteq D_{\mathcal{B}}(\mathbf{p})$ and so can conclude that $D_{\mathcal{B}}(\mathbf{q}) = \{\mathbf{x}\}$. That $D_{\mathcal{B}^0}(\mathbf{q}) = \{\mathbf{x}^0\}$ and $D_{\mathcal{B}^1}(\mathbf{q}) = \{\mathbf{x}^1\}$ follow in exactly the same way. By our definition of demand for bid collections (Equation (2.1)) it now follows that $\mathbf{x} = \mathbf{x}^0 + \mathbf{x}^1$. *Q.E.D.*

Let \mathcal{B}^+ be the positive bids of \mathcal{B} , and $|\mathcal{B}^-|$ denote the negative bids of \mathcal{B} , with the absolute value taken for their multiplicities.

PROPOSITION E.4: *For any valid bid collection \mathcal{B} ,*

$$D_{\mathcal{B}}(\mathbf{p}) = D_{\mathcal{B}^+}(\mathbf{p}) - D_{|\mathcal{B}^-|}(\mathbf{p}). \quad (\text{E.3})$$

Conversely, if Equation (E.3) holds for all $\mathbf{p} \in \mathcal{P}$, then the weight $w_{\mathcal{B}}(F)$ of every facet F of $\mathcal{L}_{\mathcal{B}}$ is positive.

PROOF OF PROPOSITION E.4: We now have the tools to derive:

$$\begin{aligned} D_{\mathcal{B}^+}(\mathbf{p}) - D_{|\mathcal{B}^-|}(\mathbf{p}) &= \mathcal{X} \cap [\text{conv } D_{\mathcal{B}^+}(\mathbf{p}) - \text{conv } D_{|\mathcal{B}^-|}(\mathbf{p})] \\ &= \mathcal{X} \cap [(\text{conv } D_{\mathcal{B}}(\mathbf{p}) + \text{conv } D_{|\mathcal{B}^-|}(\mathbf{p})) - \text{conv } D_{|\mathcal{B}^-|}(\mathbf{p})] \\ &= \mathcal{X} \cap \text{conv } D_{\mathcal{B}}(\mathbf{p}) \\ &= D_{\mathcal{B}}(\mathbf{p}), \end{aligned}$$

Here, the first and third equality follow from Fact E.1. For the second inequality, we observe that removing redundancies from $\mathcal{B} \cup |\mathcal{B}^-|$ yields \mathcal{B}^+ and apply Lemma E.3 to see that $\text{conv } D_{\mathcal{B}^+}(\mathbf{p}) = \text{conv } D_{\mathcal{B}}(\mathbf{p}) + \text{conv } D_{|\mathcal{B}^-|}(\mathbf{p})$. The final equality follows by definition of $D_{\mathcal{B}}(\mathbf{p})$ (see Equation (2.1)).

We prove the contrapositive of the second statement, so suppose that $\mathcal{L}_{\mathcal{B}}$ has a facet F such that $w_{\mathcal{B}}(F) < 0$. Then $\mathcal{L}_{|\mathcal{B}^-|}$ has a facet F^- containing F , while $\mathcal{L}_{\mathcal{B}^+}$ may have a facet F^+ containing F , or may not, in which case we write $w_{\mathcal{B}^+}(F^+) = 0$. In either case, it follows immediately from the definition of $w_{\mathcal{B}}$ that $w_{\mathcal{B}}(F) = w_{\mathcal{B}^+}(F^+) - w_{|\mathcal{B}^-|}(F^-)$.

Choose a price $\tilde{\mathbf{p}}$ in the relative interior of F , and such that $\tilde{\mathbf{p}} \notin \mathcal{L}_{\mathcal{B}^+}$ if F^+ does not exist. Then $\tilde{\mathbf{p}}$ is in the relative interior of F^- ; and of F^+ , if it exists, as these both contain F . But then

$|D_{|\mathcal{B}^-|}(\tilde{\mathbf{p}})| = w_{|\mathcal{B}^-|}(F^-) + 1$, by consideration of Proposition C.6 and Equation (2.1); similarly $|D_{\mathcal{B}^+}(\tilde{\mathbf{p}})| = w_{\mathcal{B}^+}(F^+) + 1$. Now $w_{\mathcal{B}^+}(F^+) - w_{|\mathcal{B}^-|}(F^-) = w_{\mathcal{B}}(F) < 0$ implies $|D_{\mathcal{B}^+}(\tilde{\mathbf{p}})| < |D_{|\mathcal{B}^-|}(\tilde{\mathbf{p}})|$. But then $\mathbf{x} + D_{|\mathcal{B}^-|}(\tilde{\mathbf{p}}) \subseteq D_{\mathcal{B}^+}(\tilde{\mathbf{p}})$ is impossible, and so $D_{\mathcal{B}^+}(\tilde{\mathbf{p}}) - D_{|\mathcal{B}^-|}(\tilde{\mathbf{p}}) = \emptyset$, that is, Equation (E.3) fails. So the second statement must hold. Q.E.D.

It also follows from Proposition 4.3 that the UDRs of valid bid collections form open convex polyhedra that constitute all the prices at which some bundle is demanded, a property inherited from UDRs of the LIPs of valuations. (For invalid bid collections, the set of prices at which some bundle is demanded need not be convex.)

APPENDIX F: SUPPLEMENTAL MATERIAL FOR SECTION 5

Recall from Section 5.2 that the auctioneer expresses their preferences using bids \mathcal{B}^0 ; to implement reserve prices they must “buy back” the entire supply s at some (sufficiently low) price. Each bidder $j \in [m]$ submits a bid collection \mathcal{B}^j . The setting of Section 5.3 makes the following standard assumptions:

ASSUMPTION F.1: The auctioneer expresses their preferences using bids \mathcal{B}^0 and demands the entire supply s at some price. The bidders $[m]$ submit bid collections $\mathcal{B}^1, \dots, \mathcal{B}^m$. All bids demand $\mathbf{0}$ at sufficiently high prices.

These mild assumptions guarantee existence of equilibrium when goods are divisible or agent preferences are strong-substitutes (Proposition F.2). Moreover, when all bids are positive, the LP in Section 5.3 finds an equilibrium with divisible goods (Corollary F.4); and a network flow formulation finds an equilibrium with indivisible goods when all bid collections are additionally strong-substitutes (Corollary F.5).

PROPOSITION F.2: *Under Assumption F.1, the PMA with divisible goods, or with indivisible goods and strong-substitutes bid collections $\mathcal{B}^0, \dots, \mathcal{B}^m$, admits a competitive equilibrium that clears supply. Moreover, equilibrium allocations maximise welfare among all feasible allocations.*

We next present an alternative way of formulating the valuation $v_{\mathcal{B}}$ (defined in Section 4.1) of a valid bid collection, when all bids are positive. This uses the convex extensions \hat{v}_b of the valuations v_b that are associated with every $b \in \mathcal{B}$ by Lemma 2.1. In the strong-substitutes case, we can work directly with the v_b .

For convenience, we write X_b and $X_{\mathcal{B}}$ for the domains of v_b and $v_{\mathcal{B}}$ respectively, so that $X_{\mathcal{B}} := \bigcup_{p \in \mathcal{P}} D_{\mathcal{B}}(p)$. The domains of their convex extensions \hat{v}_b and $\hat{v}_{\mathcal{B}}$ are respectively $\text{conv } X_b$ and $\text{conv } X_{\mathcal{B}}$.

LEMMA F.3: *For any collection \mathcal{B} of positive bids, $v_{\mathcal{B}}(\mathbf{x}) = \hat{v}_{\mathcal{B}}(\mathbf{x})$ for all $\mathbf{x} \in X_{\mathcal{B}}$, and*

$$\hat{v}_{\mathcal{B}}(\mathbf{x}) = \max \left\{ \sum_{b \in \mathcal{B}} \hat{v}_b(\mathbf{x}^b) \mid \mathbf{x}^b \in \text{conv } X_b \text{ and } \sum_{b \in \mathcal{B}} \mathbf{x}^b = \mathbf{x} \right\}, \forall \mathbf{x} \in \text{conv } X_{\mathcal{B}}. \quad (\text{F.1})$$

If, additionally, all bids in \mathcal{B} are strong-substitutes, then

$$v_{\mathcal{B}}(\mathbf{x}) = \max \left\{ \sum_{b \in \mathcal{B}} v_b(\mathbf{x}^b) \mid \mathbf{x}^b \in X_b \text{ and } \sum_{b \in \mathcal{B}} \mathbf{x}^b = \mathbf{x} \right\}, \forall \mathbf{x} \in X_{\mathcal{B}}. \quad (\text{F.2})$$

To understand Lemma F.3, consider the individual bids as each coming from a different individual “bid-agent”. Each is allocated a bundle \mathbf{x}^b , with \mathbf{x} available in total. The definition of demand for a bid collection in Equation (2.1) implicitly allows fractional allocations to individual bids, and so to these bid-agents. The right-hand side of Equation (F.1) is the maximum social welfare which can be obtained in this way, when \mathbf{x} is available. Moreover, a minor generalisation of Proposition F.2 tells us that competitive equilibrium always exists when goods are divisible, so this is the social welfare under competitive equilibrium. So, by Equation (F.1), $v_{\mathcal{B}}$ values a bundle \mathbf{x} at the welfare it achieves under competitive equilibrium with divisible goods.

When all bids are strong substitutes, then competitive equilibrium also exists with indivisible goods. So, in Equation (F.2), we do not need to assume that bid-agents might be allocated fractional quantities. But without the additional assumption of strong substitutes, competitive equilibrium can fail with indivisibilities, and so an integer allocation may not achieve the maximum welfare that is achievable with divisibilities. Hence Equation (F.2) does not necessarily hold when bids in \mathcal{B} are not strong substitutes.

Lemma F.3 underpins the methods to find equilibrium prices, described in Section 5.3. When bids are positive, we can use linear programming to find competitive equilibrium, in the cases of divisible goods and strong substitutes (Corollaries F.4 and F.5 respectively). When there are negative bids, we can find these prices by minimising the difference of such LPs, as Proposition F.6 shows.

COROLLARY F.4: *If all participants submit positive bid collections for divisible goods satisfying Assumption F.1, then the LP from Section 5.3 finds a competitive equilibrium allocation, and the shadow prices on the LP’s supply constraints are corresponding equilibrium prices.*

COROLLARY F.5: *If all participants submit positive strong-substitutes bid collections (for indivisible goods) satisfying Assumption F.1, then the network simplex algorithm efficiently finds a competitive equilibrium, and the shadow prices on the LP’s supply constraints are corresponding equilibrium prices.*

Recall that \mathcal{B}^+ is the subset of positive bids in \mathcal{B} , and $|\mathcal{B}^-|$ is the subset of negative bids with absolute values taken for their multiplicities.

PROPOSITION F.6: *For any valid bid collection \mathcal{B} ,*

$$v_{\mathcal{B}}(\mathbf{x}) = \min \{ v_{\mathcal{B}^+}(\mathbf{x} + \mathbf{z}) - v_{|\mathcal{B}^-|}(\mathbf{z}) \mid \mathbf{z} \in X_{|\mathcal{B}^-|} \}.$$

We now turn to the proofs of these results. Recall from Section 2.4 that we write $\widehat{D}_{\mathcal{B}}(\mathbf{p}) = \text{conv } D_{\mathcal{B}}(\mathbf{p})$. We first observe:

LEMMA F.7: *If $\mathcal{B} := \bigcup_{j=0}^m \mathcal{B}^j$ is the union of valid bid collections $\mathcal{B}^0, \dots, \mathcal{B}^m$, then \mathcal{B} is valid and $\widehat{D}_{\mathcal{B}}(\mathbf{p}) = \sum_{j=0}^m \widehat{D}_{\mathcal{B}^j}(\mathbf{p})$ for all $\mathbf{p} \in \mathcal{P}$, and $\text{conv } X_{\mathcal{B}} = \sum_{j=0}^m \text{conv } X_{\mathcal{B}^j}$.*

PROOF: We prove for $m = 1$; the general case follows immediately. That $\widehat{D}_{\mathcal{B}}(\mathbf{p}) = \widehat{D}_{\mathcal{B}^0}(\mathbf{p}) + \widehat{D}_{\mathcal{B}^1}(\mathbf{p})$ is given by Lemma E.3. Now, if $\mathbf{x} \in \text{conv } X_{\mathcal{B}}$ then for some $\mathbf{p} \in \mathcal{P}$, we know $\mathbf{x} \in \widehat{D}_{\mathcal{B}}(\mathbf{p}) = \widehat{D}_{\mathcal{B}^0}(\mathbf{p}) + \widehat{D}_{\mathcal{B}^1}(\mathbf{p}) \subseteq \text{conv } X_{\mathcal{B}^0} + \text{conv } X_{\mathcal{B}^1}$, and so $\text{conv } X_{\mathcal{B}} \subseteq \text{conv } X_{\mathcal{B}^0} + \text{conv } X_{\mathcal{B}^1}$. Conversely, consider a vertex \mathbf{x} of $\text{conv } X_{\mathcal{B}^0} + \text{conv } X_{\mathcal{B}^1}$. Then \mathbf{x} is the unique sum $\mathbf{x} = \mathbf{x}^0 + \mathbf{x}^1$ of two vertices of $\text{conv } X_{\mathcal{B}^0}$ and $\text{conv } X_{\mathcal{B}^1}$ respectively,

which both uniquely minimise some $\mathbf{p} \cdot \mathbf{x}'$ over $\text{conv } X_{\mathcal{B}^0}$ and $\text{conv } X_{\mathcal{B}^1}$ (Fact E.2). This unique minimisation still holds if we multiply \mathbf{p} by an arbitrarily large scalar, so we obtain \mathbf{p} such that $\mathbf{x}^0 \in D_{\mathcal{B}^0}(\mathbf{p})$ and $\mathbf{x}^1 \in D_{\mathcal{B}^1}(\mathbf{p})$ and thus $\mathbf{x} \in \widehat{D}_{\mathcal{B}}(\mathbf{p})$. But if all vertices of $\text{conv } X_{\mathcal{B}^0} + \text{conv } X_{\mathcal{B}^1}$ are in $\widehat{D}_{\mathcal{B}}(\mathbf{p})$ then we have the required subset inclusion, and so $\text{conv } X_{\mathcal{B}} = \text{conv } X_{\mathcal{B}^0} + \text{conv } X_{\mathcal{B}^1}$, as required. *Q.E.D.*

Both Proposition F.2 and Lemma F.3 will follow easily from the following result:

LEMMA F.8: *Given valid bid collections $\mathcal{B}^0, \dots, \mathcal{B}^m$, then for every $\mathbf{x} \in \sum_{j=0}^m \text{conv } X_{\mathcal{B}^j}$, there exists a price $\mathbf{p} \in \mathcal{P}$ and allocation $\{\mathbf{x}^j\}_{j \in [m]_0}$, such that $\mathbf{x} = \sum_{j=0}^m \mathbf{x}^j$ and $\mathbf{x}^j \in \widehat{D}_{\mathcal{B}^j}(\mathbf{p})$. Moreover, then*

$$\sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) = \widehat{v}(\mathbf{x}) := \max \left\{ \sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{y}^j) \mid \mathbf{y}^j \in \text{conv } X_{\mathcal{B}^j} \text{ and } \sum_{j=0}^m \mathbf{y}^j = \mathbf{x} \right\}. \quad (\text{F.3})$$

and $\mathbf{x} \in D_{\widehat{v}}(\mathbf{p})$. Conversely, if $\mathbf{x} = \sum_{j=0}^m \mathbf{x}^j$ where $\mathbf{x}^j \in \text{conv } X_{\mathcal{B}^j}$ satisfies $\sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) = \widehat{v}(\mathbf{x})$, then there exists $\mathbf{p} \in \mathcal{P}$ such that $\mathbf{x} \in D_{\widehat{v}}(\mathbf{p})$ and $\mathbf{x}^j \in \widehat{D}_{\mathcal{B}^j}(\mathbf{p})$ for all $j \in [m]_0$.

If, additionally, $\mathcal{B}^0, \dots, \mathcal{B}^m$ are strong substitutes bid collections and $\mathbf{x} \in \sum_{j=0}^m X_{\mathcal{B}^j}$, then in both parts above we may take $\mathbf{x}^j \in D_{\mathcal{B}^j}(\mathbf{p})$ for all $j \in [m]_0$, and

$$\widehat{v}(\mathbf{x}) = v(\mathbf{x}) := \max \left\{ \sum_{j=0}^m v_{\mathcal{B}^j}(\mathbf{y}^j) \mid \mathbf{y}^j \in X_{\mathcal{B}^j} \text{ and } \sum_{j=0}^m \mathbf{y}^j = \mathbf{x} \right\}. \quad (\text{F.4})$$

PROOF: Write $\mathcal{B} := \bigcup_{j=0}^m \mathcal{B}^j$. By Lemma F.7, \mathcal{B} is valid and $\widehat{D}_{\mathcal{B}}(\mathbf{p}) = \sum_{j=0}^m \widehat{D}_{\mathcal{B}^j}(\mathbf{p})$. This is the aggregate demand for divisible goods of the bid collections $\mathcal{B}^0, \dots, \mathcal{B}^m$. Its demand set is convex for every $\mathbf{p} \in \mathcal{P}$, and therefore by the supporting hyperplane theorem, the associated valuation $v_{\mathcal{B}}$ is concave, and thus demands every bundle in its domain (Mas-Colell et al., 1995, pp. 135-138, especially Proposition 5.C.1(v)). Its domain is $\text{conv } X_{\mathcal{B}} = \sum_{j=0}^m \text{conv } X_{\mathcal{B}^j}$ by Lemma F.7.

Fix $\mathbf{x} \in \text{conv } X_{\mathcal{B}}$. As seen above, there exists $\mathbf{p} \in \mathcal{P}$ such that $\mathbf{x} \in \widehat{D}_{\mathcal{B}}(\mathbf{p}) = \sum_{j=0}^m \widehat{D}_{\mathcal{B}^j}(\mathbf{p})$. Thus, for all $j \in [m]_0$, there exists $\mathbf{x}^j \in \widehat{D}_{\mathcal{B}^j}(\mathbf{p})$, and $\sum_{j=0}^m \mathbf{x}^j = \mathbf{x}$. Moreover, that means that for all $j \in [m]_0$, and $\mathbf{y}^j \in \text{conv } X_{\mathcal{B}^j}$, we know

$$\widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) - \mathbf{p} \cdot \mathbf{x}^j \geq \widehat{v}_{\mathcal{B}^j}(\mathbf{y}^j) - \mathbf{p} \cdot \mathbf{y}^j, \quad (\text{F.5})$$

so assume that $\sum_{j=0}^m \mathbf{y}^j = \mathbf{x}$ and sum these inequalities over $j \in [m]_0$. Subtract $\mathbf{p} \cdot \mathbf{x}$ from both sides to obtain $\sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) \geq \sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{y}^j)$, whenever $\sum_{j=0}^m \mathbf{y}^j = \mathbf{x}$. So Equation (F.3) holds.

Moreover, for any allocation $\{\mathbf{y}^j\}_{j \in [m]_0}$, which need not sum to \mathbf{x} , we have Equation (F.3) holding. So, summing, we can infer that $\widehat{v}(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x} \geq \widehat{v}(\mathbf{y}) - \mathbf{p} \cdot \mathbf{y}$ for all $\mathbf{y} \in \text{conv } X_{\mathcal{B}}$. So $\mathbf{x} \in D_{\widehat{v}}(\mathbf{p})$. Note that \mathbf{x} was arbitrary in $\text{conv } X_{\mathcal{B}}$; we can conclude that \widehat{v} is concave.

Now, suppose $\mathbf{x} = \sum_{j=0}^m \mathbf{x}^j$ where $\mathbf{x}^j \in \text{conv } X_{\mathcal{B}^j}$ satisfies $\sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) = \widehat{v}(\mathbf{x})$. Then \mathbf{x} is in the domain of \widehat{v} , so there exists $\mathbf{p} \in \mathcal{P}$ such that $\mathbf{x} \in D_{\widehat{v}}(\mathbf{p})$. Suppose, for a contradiction, that $\mathbf{x}^k \notin \widehat{D}_{\mathcal{B}^k}(\mathbf{p})$ for some $k \in [m]_0$, and let $\{\mathbf{y}^j\}_{j \in [m]_0}$ satisfy $\mathbf{y}^j = \mathbf{x}^j$ for $j \neq k$ and $\mathbf{y}^k \in \widehat{D}_{\mathcal{B}^k}(\mathbf{p})$. Then $\widehat{v}_{\mathcal{B}^j}(\mathbf{y}^j) - \mathbf{p} \cdot \mathbf{y}^j \geq \widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) - \mathbf{p} \cdot \mathbf{x}^j$ for every $j \in [m]_0$, and the inequality is strict for k . So, if $\mathbf{y} = \sum_{j=0}^m \mathbf{y}^j$, then $\sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{y}^j) - \mathbf{p} \cdot \mathbf{y} > \sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) - \mathbf{p} \cdot \mathbf{x}$. But $\widehat{v}(\mathbf{y}) \geq$

$\sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{y}^j)$ and $\widehat{v}(\mathbf{x}) = \sum_{j=0}^m \widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j)$ by assumption. So $\widehat{v}(\mathbf{y}) - \mathbf{p} \cdot \mathbf{y} > \widehat{v}(\mathbf{x}) - \mathbf{p} \cdot \mathbf{x}$, contradicting $\mathbf{x} \in D_{\widehat{v}}(\mathbf{p})$.

Now suppose $\mathcal{B}^0, \dots, \mathcal{B}^m$ are strong substitutes and $\mathbf{x} \in \sum_{j=0}^m X_{\mathcal{B}^j}$. In this case $D_{\mathcal{B}}(\mathbf{p}) = \sum_{j=0}^m D_{\mathcal{B}^j}(\mathbf{p})$ (Danilov et al., 2001), and so the \mathbf{x}^j found above satisfy $\mathbf{x}^j \in D_{\mathcal{B}^j}(\mathbf{p})$ for all $j \in [m]_0$, showing that $v(\mathbf{x}) = \widehat{v}(\mathbf{x})$. Conversely, if $\mathbf{x} = \sum_{j=0}^m \mathbf{x}^j$ where $\mathbf{x}^j \in X_{\mathcal{B}^j}$ satisfies $\sum_{j=0}^m v_{\mathcal{B}^j}(\mathbf{x}^j) = v(\mathbf{x})$, then we already know $v(\mathbf{x}) = \widehat{v}(\mathbf{x})$ and so, from the previous part, we have existence of \mathbf{p} such that $\mathbf{x}^j \in \widehat{D}_{\mathcal{B}^j}(\mathbf{p})$ for all $j \in [m]_0$. But we assumed $\mathbf{x}^j \in X_{\mathcal{B}^j} \subseteq \mathcal{X}$ and so $\mathbf{x}^j \in \mathcal{X} \cap \widehat{D}_{\mathcal{B}^j}(\mathbf{p}) = D_{\mathcal{B}^j}(\mathbf{p})$, as required. Q.E.D.

Now Proposition F.2 and Lemma F.3 are immediate corollaries.

PROOF OF PROPOSITION F.2: By Assumption F.1 we know $\mathbf{s} \in \sum_{j=0}^m \text{conv} X_{\mathcal{B}^j}$. So Lemma F.8 tells us that a competitive equilibrium exists and is welfare maximising, and that any welfare-maximising allocation forms a competitive equilibrium for some price. The case for strong substitutes follows in exactly the same way. Q.E.D.

PROOF OF LEMMA F.3: Regard \mathcal{B} as a list of valid bid collections, each containing only one (positive) bid. By Lemma F.7 we know $\text{conv} X_{\mathcal{B}} = \sum_{b \in \mathcal{B}} \text{conv} X_b$ and, for all $\mathbf{p} \in \mathcal{P}$, we have $\widehat{D}_{\mathcal{B}}(\mathbf{p}) = \sum_{b \in \mathcal{B}} \widehat{D}_b(\mathbf{p})$. So by Lemma F.8, if $\mathbf{x} \in \text{conv} X_{\mathcal{B}}$ then there exists $\mathbf{p} \in \mathcal{P}$ such that $\mathbf{x} \in \sum_{b \in \mathcal{B}} \widehat{D}_b(\mathbf{p}) = \widehat{D}_{\mathcal{B}}(\mathbf{p}) = D_{\widehat{v}_{\mathcal{B}}}(\mathbf{p})$, and this holds if and only if $\mathbf{x} \in D_{\widehat{v}}(\mathbf{p})$, where \widehat{v} is the function defined on the right-hand side of Equation (F.1). By Corollary 4.1, it follows that $\widehat{v}_{\mathcal{B}} = \widehat{v} + K$ for some constant K , but if we consider \mathbf{x} demanded at prices \mathbf{p} at which all bids demand a unique bundle, we can verify that $K = 0$ and $\widehat{v}_{\mathcal{B}} = \widehat{v}$. The case of strong substitutes follows because $\widehat{v}(\mathbf{x}) = v(\mathbf{x})$ for all $\mathbf{x} \in X_{\mathcal{B}}$, in Lemma F.8. Q.E.D.

We now prove that the PMA with divisible goods can be solved using the LP from Section 5.3 if the bid collections are positive.

PROOF OF COROLLARY F.4: Let $\mathcal{B} = \mathcal{B}^0 \cup \dots \cup \mathcal{B}^m$. Let $\mathbf{x}^0, \mathbf{x}^1, \dots, \mathbf{x}^m$ be an equilibrium allocation for the divisible PMA with $\sum_{j \in [m]_0} \mathbf{x}^j = \mathbf{s}$, which exists by Proposition F.2. By definition, \mathbf{x}^j lies in $X_{\mathcal{B}^j}$ for each participant $j \in [m]_0$. Lemma F.3 tells us that there exists a bundle $\mathbf{y}^b \in \text{conv} X_b$ for each $b \in \mathcal{B}^j$ so that $\mathbf{x}^j = \sum_{b \in \mathcal{B}^j} \mathbf{y}^b$ and $\widehat{v}_{\mathcal{B}^j}(\mathbf{x}^j) = \sum_{b \in \mathcal{B}^j} \widehat{v}_b(\mathbf{y}^b) = \sum_{b \in \mathcal{B}^j} \mathbf{r}^b \cdot \mathbf{y}^b$. Recall from Lemma 2.1 that $X_b = \mathcal{X} \cap \text{conv}\{\mathbf{m}t_i \mathbf{e}^i \mid i \in I\}$, where I is the set of goods in which b is interested and $m > 0$. So $\mathbf{y}^b \in \text{conv} X_b$ implies that the bid demand constraint of the LP and $\mathbf{y}^b \geq 0$ are satisfied, for each $b \in \mathcal{B}$. Moreover, $\mathbf{x}^j = \sum_{b \in \mathcal{B}^j} \mathbf{y}^b$ is the bidder allocation constraint of the LP, and $\mathbf{x}^0 + \dots + \mathbf{x}^m = \mathbf{s}$ immediately implies the supply constraints. So $(\mathbf{x}^j)_{j \in [m]_0}$ and $(\mathbf{y}^b)_{b \in \mathcal{B}^j}$ form a feasible solution to the LP. Moreover, the LP's objective function expresses the allocation's welfare $\sum_{j \in [m]_0} \widehat{v}(\mathbf{x}^j) = \sum_{j \in [m]_0} \sum_{b \in \mathcal{B}^j} \mathbf{r}^b \cdot \mathbf{y}^b$.

Now suppose that this equilibrium allocation $(\mathbf{x}^j)_{j \in [m]_0}$ and $(\mathbf{y}^b)_{b \in \mathcal{B}^j}$ is not an optimal solution for the LP, and that $(\tilde{\mathbf{x}}^j)_{j \in [m]_0}$ and $(\tilde{\mathbf{y}}^b)_{b \in \mathcal{B}^j}$ is a feasible solution that achieves a greater objective function value. The bid demand constraints of the LP imply that $\tilde{\mathbf{y}}^b \in \text{conv} X_b$ for each $b \in \mathcal{B}$, and the bidder allocation constraints imply that $\tilde{\mathbf{x}}^j = \sum_{b \in \mathcal{B}^j} \tilde{\mathbf{y}}^b$ for all participants $j \in [m]_0$. So Lemma F.3 tells us that $\widehat{v}_{\mathcal{B}^j}(\tilde{\mathbf{x}}^j) \geq \sum_{b \in \mathcal{B}^j} \widehat{v}_b(\tilde{\mathbf{y}}^b) = \sum_{b \in \mathcal{B}^j} \mathbf{r}^b \cdot \tilde{\mathbf{y}}^b$ for each participant $j \in [m]_0$, and so the welfare $\sum_{j \in [m]_0} \widehat{v}_{\mathcal{B}^j}(\tilde{\mathbf{x}}^j)$ achieved by allocation $(\tilde{\mathbf{x}}^j)_{j \in [m]_0}$ is lower-bounded by $\sum_{j \in [m]_0} \sum_{b \in \mathcal{B}^j} \mathbf{r}^b \cdot \tilde{\mathbf{y}}^b$. But this latter term is strictly larger than the welfare

of the equilibrium allocation $(\mathbf{x}^j)_{j \in [m]_0}$ and $(\mathbf{y}^b)_{b \in \mathcal{B}}$. This is a contradiction because equilibrium allocations maximise welfare by Proposition F.2. So the LP is bounded, and a feasible solution of the LP is optimal if and only if it is an equilibrium allocation.

To see that the shadow prices \mathbf{p} on the LP's supply constraints are supporting equilibrium prices, form the dual LP and note that the complementary slackness conditions imply $\mathbf{y}^b \in D_b(\mathbf{p})$. Q.E.D.

With positive strong-substitutes bids (for indivisible goods), we can find equilibrium by interpreting the LP from Section 5.3 as a minimum-cost network flow problem. These can be solved efficiently using standard methods such as the network simplex algorithm. See Ahuja et al. (1993) for an introduction to minimum-cost network flows and the network simplex algorithm.

PROOF OF COROLLARY F.5: As bids are strong-substitutes, the bid demand constraints of the LP simplify to $\sum_{i \in I_b} y_i^b = m^b$. We first rewrite the LP without the redundant bidder allocation constraints (and thus without variables \mathbf{x}^j). This LP can then be understood as a flow network consisting of a source node, a target node t , nodes \mathcal{B} and nodes $[n]_0$. The source node is connected to each bid $\mathbf{b} \in \mathcal{B}$ by an arc with capacity m^b and cost 0. Each bid $\mathbf{b} \in \mathcal{B}$ is connected to each good $i \in [n]_0$ (including the null good) in which it is interested by an arc with infinite capacity and cost $-r_i^b$, and each good $i \in [n]_0$ is connected to the target node by an arc with capacity s_i ($s_0 = \infty$) and cost 0. The flow constraints of the problem ensure that any feasible flow in the network is an allocation to bidders, and vice versa. Moreover, a competitive equilibrium allocation (of indivisible goods) exists by Proposition F.2, so the network problem has feasible flows and integer minimal-cost flows correspond one-to-one to equilibrium allocations. Moreover, as the arc capacities are integer, the network simplex algorithm finds an integer minimal-cost flow. Finally, the shadow prices to the supply constraints in the original LP are supporting equilibrium prices by complementary slackness conditions. Q.E.D.

Our method for proving Proposition F.6 mirrors that of Baldwin et al. (2024, Theorem 1), which is the strong substitutes special case. We first give a useful lemma summarising the relationships between valuations, convex conjugates and subdifferentials. Recall that, given a function $f : \text{dom } f \rightarrow \mathbb{R}$ where $\text{dom } f \subseteq \mathbb{R}^n$, we define its *convex conjugate* $f^* : \text{dom } f^* \rightarrow \mathbb{R}$ by $f^*(\mathbf{p}) = \sup_{\mathbf{x} \in \text{dom } f} (\mathbf{p} \cdot \mathbf{x} - f(\mathbf{x}))$, where $\text{dom } f^*$ is the set of points on which $f^*(\mathbf{p})$ is finite-valued. And the *subdifferential* ∂f of f is the correspondence

$$\partial f(\mathbf{x}) = \{\mathbf{p} \in \mathbb{R}^n \mid \mathbf{p} \cdot \mathbf{x} - f(\mathbf{x}) \geq \mathbf{p} \cdot \mathbf{y} - f(\mathbf{y}) \quad \forall \mathbf{y} \in \mathbb{R}^n\}.$$

The *domain* $\text{dom } \partial f$ of the subdifferential is those $\mathbf{x} \in \text{dom } f$ for which $\partial f(\mathbf{x}) \neq \emptyset$.

We also extend our definition of concave extension: for any valuation $v : X \rightarrow \mathbb{R}$ with finite domain $X \subseteq \mathbb{Z}^n$, the *concave extension* $\widehat{v} : \text{conv } X \rightarrow \mathbb{R}$ is the minimal concave function such that $\widehat{v}(\mathbf{x}) \geq v(\mathbf{x})$ for all $\mathbf{x} \in X$; equality holds for all \mathbf{x} v is concave). Now:

LEMMA F.9—(cf. Baldwin et al. (2024, Lemma 1)): *For a valuation $v : X \rightarrow \mathbb{R}$, with $X \subseteq \mathbb{Z}^n$, write $f_v = -\widehat{v}$, where \widehat{v} is the concave extension of v . Then, for all $\mathbf{x} \in \text{conv } X$ and $\mathbf{p} \in \mathbb{R}^n$, we have:*

- (i) $\partial f_v(\mathbf{x}) = \{\mathbf{p} \in \mathbb{R}^n \mid \mathbf{x} \in \text{conv } D_v(-\mathbf{p})\}$.
- (ii) $f_v^*(\mathbf{p}) = \pi_{\widehat{v}}(-\mathbf{p}) = \pi_v(-\mathbf{p})$ and $\partial f_v^*(\mathbf{p}) = \text{conv } D_v(-\mathbf{p})$;
- (iii) $\text{dom } \partial f_v = \text{dom } f_v = \text{conv } X$ and $\text{dom } \partial f_v^* = \text{dom } f_v^* = \mathbb{R}^n$;
- (iv) $-f_v^{**}(\mathbf{x}) = \inf_{\mathbf{p} \in \mathbb{R}^n} (\pi_v(\mathbf{p}) + \mathbf{p} \cdot \mathbf{x}) = \widehat{v}(\mathbf{x})$

PROOF: It is clear by definition of f_v and ∂f_v that $\partial f_v(\mathbf{x}) = \{\mathbf{p} \in \mathbb{R}^n \mid \mathbf{x} \in D_{\widehat{v}}(-\mathbf{p})\}$, so Part (i) follows because $D_{\widehat{v}}(-\mathbf{p}) = \text{conv } D_v(-\mathbf{p})$, cf., e.g., [Baldwin and Klemperer \(2019b, Lemma 2.17\)](#). Thus $\partial f_v(\mathbf{x}) \neq \emptyset$ if and only if \mathbf{x} is demanded for some prices; as \widehat{v} is concave this holds for every $\mathbf{x} \in \text{dom } \widehat{v} = \text{conv } X = \text{dom } f$, which proves the first part of Part (iii). Next, it is clear that $f_v^*(\mathbf{p}) = \pi_{\widehat{v}}(-\mathbf{p})$. But for all $\mathbf{q} \in \mathbb{R}^n$ there exists $\mathbf{x} \in D_v(\mathbf{q}) \subsetneq \text{conv } D_v(\mathbf{q}) = D_{\widehat{v}}(\mathbf{q})$ and so $\pi_{\widehat{v}}(\mathbf{q}) = \widehat{v}(\mathbf{x}) - \mathbf{q} \cdot \mathbf{x} = v(\mathbf{x}) - \mathbf{q} \cdot \mathbf{x} = \pi_v(\mathbf{q})$, so the first part of (ii) holds. Moreover, by [Rockafellar \(1970, Theorem 23.5\)](#), $\mathbf{x} \in \partial f_v^*(\mathbf{p})$ if and only if $\mathbf{p} \in \partial f_v(\mathbf{x})$, which combined with Part (i) is sufficient to prove Part (ii). Moreover, then $\partial f_v^*(\mathbf{p}) \neq \emptyset$ for all $\mathbf{p} \in \mathbb{R}^n$, and similarly $f_v^*(\mathbf{p})$ is well defined for all $\mathbf{p} \in \mathbb{R}^n$, so this demonstrates the second part of Part (iii).

Finally by definition $-f_v^{**}(\mathbf{x}) = -\sup_{\mathbf{p}' \in \text{dom } f_v^*} (\mathbf{p}' \cdot \mathbf{x} - f_v^*(\mathbf{p}')) = \inf_{\mathbf{p}' \in \mathbb{R}^n} (-\mathbf{p}' \cdot \mathbf{x} + \pi_v^*(-\mathbf{p}'))$, where we apply parts (ii) and (iii). But, letting $\mathbf{p} = -\mathbf{p}'$, this is equal to $\inf_{\mathbf{p} \in \mathbb{R}^n} (\pi_v(\mathbf{p}) + \mathbf{p} \cdot \mathbf{x})$. Moreover, $f_v^{**} = f_v$ because f is closed and convex ([Rockafellar, 1970, Theorem 12.2](#)) and so $f_v^{**}(\mathbf{x}) = \widehat{v}(\mathbf{x})$. *Q.E.D.*

We make use of Toland-Singer duality, in the following form:

FACT F.10—([Tao and An \(1997, Theorem 1\)](#)): Let $f : \text{dom } f \rightarrow \mathbb{R}$ and $g : \text{dom } g \rightarrow \mathbb{R}$ be convex continuous functions with closed and convex domains $\text{dom } f \subseteq \text{dom } g \subseteq \mathbb{R}^n$ and such that $\text{dom } g^* \subseteq \text{dom } f^* \subseteq \mathbb{R}^n$. If one of the differences $f(\mathbf{x}) - g(\mathbf{x})$ and $g^*(\mathbf{y}) - f^*(\mathbf{y})$ has a minimum in $\text{dom } f$, respectively $\text{dom } g^*$, the other difference also has one, and

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

Moreover, if $\tilde{\mathbf{x}}$ minimises $f(\mathbf{x}) - g(\mathbf{x})$, then any $\tilde{\mathbf{y}} \in \partial g(\tilde{\mathbf{x}})$ minimises $g^*(\mathbf{y}) - f^*(\mathbf{y})$. Conversely, if $\tilde{\mathbf{y}}$ minimises $g^*(\mathbf{y}) - f^*(\mathbf{y})$, then any $\tilde{\mathbf{x}} \in \partial f^*(\tilde{\mathbf{y}})$ minimises $f(\mathbf{x}) - g(\mathbf{x})$.

We need to show:

LEMMA F.11—(cf. [Baldwin et al. \(2024, Lemma 4\)](#)): *Let \mathcal{B} be a valid collection of bids. Then $X_{\mathcal{B}} + X_{|\mathcal{B}^-|} \subseteq X_{\mathcal{B}^+}$.*

PROOF: Observe that removing redundancies from $\mathcal{B} \cup |\mathcal{B}^-|$ gives \mathcal{B}^+ , so $\text{conv } X_{\mathcal{B}} + \text{conv } X_{|\mathcal{B}^-|} = \text{conv } X_{\mathcal{B}^+}$ by Lemma F.7. But $X_{\mathcal{B}} + X_{|\mathcal{B}^-|} \subseteq \mathcal{X} \cap (\text{conv } X_{\mathcal{B}} + \text{conv } X_{|\mathcal{B}^-|})$. Since $v_{\mathcal{B}^+}$ is concave, we know that $\mathcal{X} \cap \text{conv } X_{\mathcal{B}^+} = X_{\mathcal{B}^+}$. So $X_{\mathcal{B}} + X_{|\mathcal{B}^-|} \subseteq X_{\mathcal{B}^+}$. *Q.E.D.*

PROOF OF PROPOSITION F.6: To prove this result, we drop the 0th coordinates and assume that bundles $\mathbf{x} \in \mathbb{Z}^n$ and prices $\mathbf{p} \in \mathbb{R}^n$. This makes no difference to the result, as $p_0 = x_0 = 0$ in every case, but it makes it more straightforward to apply the tools of Lemma F.9 and Fact F.10. Since \mathcal{B}^+ and $|\mathcal{B}^-|$ consist exclusively of positive bids, they are valid, as is \mathcal{X} by assumption, and so all parts of Proposition 4.3 apply to all three bid collections.

Fix $\mathbf{x} \in X_{\mathcal{B}}$. Let $f : \text{conv } X_{|\mathcal{B}^-|} \rightarrow \mathbb{R}$ be the convex extension of $\mathbf{z} \mapsto -v_{|\mathcal{B}^-|}(\mathbf{z})$ and let $g : \text{conv } X_{|\mathcal{B}^-|} \rightarrow \mathbb{R}$ be the convex extension of $\mathbf{z} \mapsto -v_{\mathcal{B}^+}(\mathbf{x} + \mathbf{z})$. By Lemma F.11 we know g is well defined on this domain, and observe that $\text{dom } f = \text{dom } g$. Moreover, $\text{dom } g^* = \text{dom } f^* = \mathbb{R}^n$ by Lemma F.9. By Lemma F.9 Part (ii) and Proposition 4.3 Part (v), we know that $f^*(\mathbf{p}) = \pi_{|\mathcal{B}^-|}(-\mathbf{p})$, and infer that $g^*(\mathbf{p}) = \pi_{\mathcal{B}^+}(-\mathbf{p}) - \mathbf{p} \cdot \mathbf{x}$. So we apply Fact F.10 to $f - g$: if the minimum exists on either side, then

$$\min_{\mathbf{z} \in \text{conv } X_{|\mathcal{B}^-|}} \widehat{v}_{\mathcal{B}^+}(\mathbf{x} + \mathbf{z}) - \widehat{v}_{|\mathcal{B}^-|}(\mathbf{z}) = \min_{\mathbf{p} \in \mathbb{R}^n} \pi_{\mathcal{B}^+}(-\mathbf{p}) - \pi_{|\mathcal{B}^-|}(-\mathbf{p}) - \mathbf{p} \cdot \mathbf{x} \quad (\text{F.6})$$

But we can re-write the right-hand side of Equation (F.6) as

$$\min_{\mathbf{p} \in \mathcal{P}} \pi_{\mathcal{B}^+}(\mathbf{p}) - \pi_{|\mathcal{B}^-|}(\mathbf{p}) + \mathbf{p} \cdot \mathbf{x} = \min_{\mathbf{p} \in \mathcal{P}} \pi_{\mathcal{B}}(\mathbf{p}) + \mathbf{p} \cdot \mathbf{x} = \min_{\mathbf{p} \in \mathcal{P}} \pi_{v_{\mathcal{B}}}(\mathbf{p}) + \mathbf{p} \cdot \mathbf{x} = \widehat{v}_{\mathcal{B}}(\mathbf{x})$$

where we apply the definition of $\pi_{\mathcal{B}}$ from Section 4.2; Part (v) from Proposition 4.3; and Lemma F.9 Part (iv). So the left-hand side of Equation (F.6) also has a solution, which is equal to $\widehat{v}_{\mathcal{B}}(\mathbf{x})$, which we can also write as $v_{\mathcal{B}}(\mathbf{x})$, since $\mathbf{x} \in \mathcal{X}$.

It remains to show that we can simply minimise the left-hand side of Equation (F.6) over $X_{|\mathcal{B}^-|}$. But, by Fact F.10, if \mathbf{p} minimises the right-hand side of Equation (F.6) then any $\mathbf{z} \in \partial f^*(\mathbf{p}) = \text{conv } D_{|\mathcal{B}^-|}(\mathbf{p})$ minimises the left-hand side of Equation (F.6). And as $D_{|\mathcal{B}^-|}(\mathbf{p}) \neq \emptyset$, we know $\text{conv } D_{|\mathcal{B}^-|}(\mathbf{p})$ always contains an integer bundle, so we may assume that $\mathbf{z} \in \mathcal{X}$, and so $\mathbf{z} \in \mathcal{X} \cap \text{conv } X_{|\mathcal{B}^-|} = X_{|\mathcal{B}^-|}$ (see above the statement of Lemma F.3). So Equation (F.6) does indeed reduce to $\min_{\mathbf{z} \in X_{|\mathcal{B}^-|}} v_{\mathcal{B}^+}(\mathbf{x} + \mathbf{z}) - v_{|\mathcal{B}^-|}(\mathbf{z})$, as required. Q.E.D.

F.1. Material for Section 5.6

We write $\|\mathbf{x} - \mathbf{y}\|_1 = \sum_{i \in [n]} (x_i - y_i)$ to mean the L_1 -distance (or *Taxicab distance*) between two vectors \mathbf{x} and \mathbf{y} .

LEMMA F.12: *Suppose $D_{(\mathbf{r}; \beta)}$ is the demand correspondence of a single arctic bid $(\mathbf{r}; \beta)$. For any $\varepsilon > 0$, we can create a collection \mathcal{B} of standard substitutes PMA bids (for divisible goods) so that the following holds at any prices $\mathbf{p} \in \mathcal{P}$ with $\mathbf{p} \geq \varepsilon$: For every $\mathbf{x} \in D_{(\mathbf{r}; \beta)}(\mathbf{p})$, there exists $\mathbf{y} \in \widehat{D}_{\mathcal{B}}(\mathbf{p})$ with $\|\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon$. Likewise, for every $\mathbf{y} \in \widehat{D}_{\mathcal{B}}(\mathbf{p})$, there exists $\mathbf{x} \in D_{(\mathbf{r}; \beta)}(\mathbf{p})$ with $\|\mathbf{x} - \mathbf{y}\|_1 \leq \varepsilon$.*

PROOF: Fix $\varepsilon > 0$ and prices $\mathbf{p} \geq \varepsilon$. Define $\alpha = \max\{\lambda > 0 \mid \lambda r_i \leq \varepsilon, \forall i \in [n]\}$ so that $\mathbf{p} \geq \varepsilon$ implies $p_i \geq \alpha r_i$ for all $i \in [n]$. Now construct a bid collection $\mathcal{B} = \{\mathbf{b}^0, \mathbf{b}^1, \dots, \mathbf{b}^K\}$ of $K + 1 \gg 0$ standard PMA bids $\mathbf{b}^k := (\mathbf{r}^k; \mathbf{t}; m^k)$ with roots evenly spaced along the line segment from $\alpha \mathbf{r}$ to \mathbf{r} . Define the bids' tradeoff vector \mathbf{t} as $t_i := \frac{1}{r_i}$ for all $i \in [n]$. The bids' roots are $\mathbf{r}^k := \alpha \mathbf{r} + \frac{K-k}{K}(1 - \alpha)\mathbf{r}$, and the multiplicities are $m^k := \frac{\beta}{r_i^k t_i} - \sum_{j=0}^{k-1} m^j$. (The choice of i for the definition of multiplicity is irrelevant, as $r_i^k t_i = \alpha + \frac{K-k}{K}(1 - \alpha)$ for any i .)

Suppose first that the arctic bid and \mathcal{B} both demand a unique bundle at \mathbf{p} . If the arctic bid demands $\mathbf{0}$, then $\mathbf{p} > \mathbf{r}$ and all standard bids in \mathcal{B} also demand $\mathbf{0}$, so we are done. So suppose this is not the case. The arctic bid demands bundle $\frac{\beta}{p_i} \mathbf{e}^i$ for some unique good $i \in [n]$. Let l be the largest $k \in [K - 1]_0$ with $r_i^k > p_i$. Such a k exists because $r_i^K = \alpha r_i < p_i < r_i = r_i^0$. The standard bids $\mathbf{b}^0, \dots, \mathbf{b}^l$ demand bundle $m^k t_i \mathbf{e}^i$, while the bids $\mathbf{b}^{l+1}, \dots, \mathbf{b}^K$ demand $\mathbf{0}$. So the aggregate bundle demanded by \mathcal{B} is $\sum_{k=0}^l m^k t_i \mathbf{e}^i = \frac{\beta}{r_i^l} \mathbf{e}^i$. By our choice of l , we have $\frac{\beta}{r_i^l} < \frac{\beta}{p_i} < \frac{\beta}{r_i^{l+1}}$, so the L_1 -distance between the bundles demanded by the arctic bid and \mathcal{B} is thus $\Delta := \frac{\beta}{p_i} - \frac{\beta}{r_i^l} < \frac{\beta}{r_i^{l+1}} - \frac{\beta}{r_i^l}$. As the derivative of function $f(k) = \frac{\beta}{r_i^{k+1}} - \frac{\beta}{r_i^k}$ is non-negative, $f(k)$ on domain $[K - 1]_0$ is maximised at $k = K - 1$, and so $f(K - 1) = \frac{1}{\alpha r_i} - \frac{1}{\alpha r_i + \frac{1}{K}(1 - \alpha)r_i}$ is an upper bound on $f(l)$ and thus on Δ . It is clear that $f(K - 1) < \frac{\varepsilon}{n+1}$ for sufficiently large K , so $\Delta \leq f(0) < \frac{\varepsilon}{n+1}$.

Now suppose the arctic bid is indifferent between two or more bundles at \mathbf{p} , and let $\mathbf{x} \in D_{\mathbf{r}; \beta}(\mathbf{p})$. Then \mathbf{x} is the convex combination of $m \leq n + 1$ bundles $\mathbf{x}^j \in D_{\mathbf{r}; \beta}(\mathbf{p}^j)$ for generic prices $\mathbf{p}^1, \dots, \mathbf{p}^m$ infinitesimally close to \mathbf{p} , and $\mathbf{x} = \sum_{j \in [m]} \lambda_j \mathbf{x}^j$ where $\lambda_j = [0, 1]$ and $\sum_{i \in [m]} \lambda_j = 1$. By the previous paragraph, we know that there exists $\mathbf{y}^j \in D_{\mathcal{B}}(\mathbf{p}^j)$ with

$\|\mathbf{x}^j - \mathbf{y}^j\|_1 \leq \frac{\varepsilon}{n+1}$. Moreover, $\sum_{j \in [m]} \lambda^j \mathbf{y}^j =: \mathbf{y} \in D_{\mathcal{B}}(\mathbf{p})$ by the definition of $D_{\mathcal{B}}$. It is straightforward that

$$\|\mathbf{x} - \mathbf{y}\|_1 = \left\| \sum_{j \in [m]} (\mathbf{x}^j - \mathbf{y}^j) \right\|_1 \leq \sum_{j \in [m]} \|\mathbf{x}^j - \mathbf{y}^j\|_1 \leq \varepsilon.$$

The argument for the other direction is analogous.

Q.E.D.

APPENDIX G: FULL PROOFS FOR THEOREM 3.1 AND COROLLARIES 3.2, 3.3 AND 3.7

Here we prove the remaining results from Section 3.

THEOREM 3.1: *For any substitutes concave valuation v , there exists a unique bid collection \mathcal{B} such that $D_{\mathcal{B}} = D_v$.*

COROLLARY 3.2: *For any strong-substitutes valuation v , there exists a unique bid collection \mathcal{B} such that $D_{\mathcal{B}} = D_v$, and all bids in \mathcal{B} are strong-substitutes bids.*

COROLLARY 3.3: *For any substitutes concave valuation v , there exists a regular bid collection \mathcal{B} for any $\underline{\mathbf{p}} \in \mathcal{P}$, such that $D_{\mathcal{B}}(\mathbf{p}) = D_v(\mathbf{p})$ for all $\mathbf{p} \geq \underline{\mathbf{p}}$, if there exists $\bar{\mathbf{p}}$ such that $D_v(\mathbf{p}) = \{\mathbf{0}\}$ for all $\mathbf{p} \geq \bar{\mathbf{p}}$.*

COROLLARY 3.7: *For any regular valuation v , there exists a unique bid collection \mathcal{B} such that $D_{\mathcal{B}} = D_v$, and all bids in \mathcal{B} are regular bids.*

G.1. Preliminaries

Recall that the roots of standard substitutes PMA bids can take real values together with $-\infty$, so we define $\mathbb{R} := \mathbb{R} \cup \{-\infty\}$ and the set of all possible roots $\mathcal{R} := \{\mathbf{r} \in \mathbb{R}^{n+1} \mid r_0 \in \{0, -\infty\}\}$. Recall that $I := \{i \in [n]_0 \mid r_i > -\infty\}$ is the set of goods in which the bid is interested. Recall that $\mathcal{T} \subseteq \mathbb{Z}^{[n]_0} \geq 0$ is the set of tradeoff vectors, for which $t_0 = 1$ and \mathbf{t}_{-0} is a primitive integer vector. Moreover, a bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m) \in \mathcal{R} \times \mathcal{T} \times \mathbb{Z}$ satisfies $r_i = -\infty \implies t_i = 0$ for $i \in [n]$. This appendix uses a different normalisation to the body text. Having found \mathcal{B} as described in this appendix, they can then be re-normalised as described in Section 2.2. More generally, Lemma 2.2 tells us that a bid can be normalised by defining any consistent procedure for changing its root from \mathbf{r} to an \mathbf{r}' in the set $\{\mathbf{r}' \in \mathcal{P} \mid t_i(r_i - r'_i) = t_j(r_j - r'_j), \forall i, j \in I\}$.

Let $N_\varepsilon(\mathbf{p}) = \{\mathbf{q} \in \mathcal{P} \mid \|\mathbf{q} - \mathbf{p}\| < \varepsilon\}$ denote the ε -neighbourhood in price space around price vector \mathbf{p} . (The norm $\|\cdot\|$ we choose is arbitrary, but we assume the L_1 -norm for concreteness.)

We define the characteristic vector $\chi^J \in \{-1, 1\}^{[n]_0}$ of set $J \subseteq [n]_0$ as $\chi_k^J := 1$ if $k \in J$ and $\chi_k^J := -1$ otherwise. We also use the shorthand $\chi_{0k}^J := \chi_0^J \chi_k^J$ (so $\chi_{0k}^J = -1$ if exactly one of 0 and k is in J , and it is 1 otherwise).

A *substitutes hyperplane* is a hyperplane in price space \mathcal{P} that is normal to $a\mathbf{e}^i - b\mathbf{e}^j$ for co-prime $a, b \in \mathbb{N}$ and distinct $i, j \in [n]_0$. We call (i, j) the *orientation* of the hyperplane with distinct $i, j \in [n]_0$, and the co-prime integers $a, b \in \mathbb{N}$ define its *slope* $\frac{a}{b}$. If i or j are 0, this slope is redundant and denoted 1. Throughout, we make extensive use of the notation

$$H(\mathbf{r}; i, j; \frac{a}{b}) := \{\mathbf{q} \in \mathcal{P} \mid (a\mathbf{e}^i - b\mathbf{e}^j) \cdot (\mathbf{q} - \mathbf{r}) = 0\}$$

with $\mathbf{r} \in \mathcal{R}$ and $r_i, r_j < -\infty$ to denote the substitutes hyperplane normal to $ae^i - be^j$ containing point $r_i e^i + r_j e^j$. The half-spaces on either side of $H(\mathbf{r}; i, j; \frac{a}{b})$ are denoted

$$\begin{aligned} H^-(\mathbf{r}; i, j; \frac{a}{b}) &:= \{\mathbf{q} \in \mathcal{P} \mid (ae^i - be^j) \cdot (\mathbf{q} - \mathbf{r}) \leq 0\}, \\ H^+(\mathbf{r}; i, j; \frac{a}{b}) &:= \{\mathbf{q} \in \mathcal{P} \mid (ae^i - be^j) \cdot (\mathbf{q} - \mathbf{r}) \geq 0\}. \end{aligned}$$

Note that $H(\mathbf{r}; i, j; \frac{a}{b}) = H(\mathbf{r}; j, i; \frac{b}{a})$ and $H^+(\mathbf{r}; i, j; \frac{a}{b}) = H^-(\mathbf{r}; j, i; \frac{b}{a})$.

It is useful to describe explicitly the LIP \mathcal{L}_b of an individual bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ interested in goods I . Recall that \mathcal{L}_b is the set of prices at which \mathbf{b} is indifferent between two or more bundles. By definition of D_b , the prices $R^i := \{\mathbf{p} \in \mathbb{R}^n \mid t_i(p_i - r_i) \geq t_k(p_k - r_k), \forall k \in I\}$ at which \mathbf{b} demands good i form a convex region (and the interior of this region is the UDR within which good i is uniquely demanded). So a bid is indifferent between goods i and j at all prices in facet $F_b^{ij} := R^i \cap R^j$ of \mathcal{L}_b , and \mathcal{L}_b is the union of the facets F_b^{ij} taken over all $i, j \in [n]_0$. We can also define the facets F_b^{ij} using our hyperplane notation:

$$F_b^{ij} = H(\mathbf{r}; i, j; \frac{t_i}{t_j}) \cap \bigcap_{k \in I \setminus \{i, j\}} H^+(\mathbf{r}; k, j; \frac{t_k}{t_j}). \quad (\text{G.1})$$

The vector $t_i e^i - t_j e^j$ is normal to F_b^{ij} . But, although \mathbf{t} is a primitive integer vector, t_i and t_j need not be coprime when $n > 2$. So the primitive integer vector normal to F_b^{ij} is $\frac{1}{\gcd(t_i, t_j)}(t_i e^i - t_j e^j)$. Meanwhile the change in demand from crossing this F_b^{ij} is $mt_i e^i - mt_j e^j$. We conclude that $w_b(F_b^{ij}) = m \gcd(t_i, t_j)$.


G.1.1. Extending LIPs with hyperplanes

Suppose (\mathcal{L}, w) is the weighted LIP of a substitutes valuation v or a collection \mathcal{B} of standard substitutes PMA bids. We can take the union of the (substitutes) hyperplanes which are the affine spans of each facet in \mathcal{L} , and call the resulting point set the *hyperplanes of indifference prices (HIP)* \mathcal{H} . As \mathcal{H} is the union of a collection of substitutes hyperplanes, it is immediate that it is the union of an $(n - 1)$ -dimensional polyhedral complex. So, analogously to LIPs, \mathcal{H} inherits from this complex the *faces, facets*, and *vertices*. More explicitly, the *faces* of a HIP \mathcal{H} are the subsets of hyperplanes $H \subseteq \mathcal{H}$ given by taking the intersection with either (or both) of the positive or negative half-spaces defined by H for all other hyperplanes $H' \subseteq \mathcal{H}$; its *facets* are the $(n - 1)$ -dimensional faces; and its *vertices* are the intersection of n linearly independent hyperplanes of \mathcal{H} .

Note that any facet of \mathcal{H} is either a subset of a facet of \mathcal{L} , or does not meet any facet of \mathcal{L} $(n - 1)$ -dimensionally. Conversely, any facet of \mathcal{L} is either a facet of \mathcal{H} or is subdivided into several facets of the corresponding HIP \mathcal{H} by other hyperplanes of the \mathcal{H} . This allows us to associate with (\mathcal{L}, w) a *weighted HIP* (\mathcal{H}, w) with the following weight function w . We write $w(F) = w(F')$ for any $(n - 1)$ -dimensional subset F of a facet F' of \mathcal{L} , and $w(F) = 0$ for any $(n - 1)$ -dimensional linear subset of \mathcal{P} that has at most $(n - 2)$ -dimensional intersection with \mathcal{L} . As argued above, in particular this gives a weight on every facet F of \mathcal{H} (and also on some other sets, which will be useful to us later). A weighted LIP for two goods and its corresponding weighted HIP are illustrated in Figure G.1.

We write (\mathcal{H}_v, w_v) for the weighted HIP of valuation v and (\mathcal{H}_B, w_B) for the weighted HIP of bid collection \mathcal{B} .

OBSERVATION G.1: Suppose (\mathcal{H}, w) is the weighted HIP of a valuation or bid collection. For every $(n - 2)$ -dimensional face G of \mathcal{H} , the weights $w(F^k)$ of the facets F^1, \dots, F^l that



(a) The LIP of a valuation. Facets are labelled with their weights.

(b) The corresponding weighted HIP is drawn in grey and dashed. Facets with non-zero weight are labelled with their weights.

FIGURE G.1.—Illustrations of a weighted LIP and weighted HIP of an ordinary substitutes valuation in price space with two goods. Two facets of \mathcal{L} with non-zero weight have each been divided into two facets of \mathcal{H} , so two non-zero facet labels have been added.

contain G , and primitive integer normal vectors $\mathbf{n}^1, \dots, \mathbf{n}^l$ for these facets defined by a fixed rotational direction about G , satisfy $\sum_{k=1}^l w(F^k) \mathbf{n}^k = 0$.

Note that our weighted HIPs are *parsimonious* in that every hyperplane contains at least one facet with non-zero weight (cf. Section 3.1). However, it is sometimes convenient to extend them by inserting additional “dummy” hyperplanes (which are still in substitutes directions); all facets within a dummy hyperplane are weighted zero. Such dummy hyperplanes do not interfere with the balancing property, but a HIP with dummy hyperplanes is not parsimonious, so we sometimes clarify that this assumption is needed.

The HIP of a valuation or bid collection divides price space \mathcal{P} into n -dimensional regions. In direct analogy to LIPs, we also call these regions UDRs and note that the UDRs of a HIP \mathcal{H} are subsets of UDRs of the corresponding LIP \mathcal{L} . The dividing facet F between two UDRs of \mathcal{H} is either contained in a facet of \mathcal{L} , or the two UDRs are subsets of the same UDR of \mathcal{L} , in which case F has weight 0 in \mathcal{H} . So in either case, the change in demand between the two UDRs is captured by the facet normal and facet weight of F . This follows directly from the analogous results for the weighted LIPs of valuations and bid collections in Section 3.1.

PROPOSITION G.2: *At all prices in a given UDR of (\mathcal{H}, w) , the same unique bundle is demanded. The change in demand as we change prices to cross a facet F of \mathcal{H} is $w(F)\mathbf{n}$, where \mathbf{n} is the primitive integer vector that is normal to F and points in the opposite direction to the price change.*

It is useful to write out the HIPs \mathcal{H}_b and \mathcal{H}_B of an individual bid \mathbf{b} and bid collection \mathcal{B} explicitly using our hyperplane notation. Using our explicit expression for F_b^{ij} in Equation (G.1), it is straightforward that

$$\mathcal{H}_b = \bigcup_{i,j \in I} H(\mathbf{r}; i, j; \frac{t_i}{t_j}). \quad (\text{G.2})$$

Note that the facet F_b^{ij} of prices at which \mathbf{b} is indifferent between goods i and j is not necessarily a facet of \mathcal{H}_b , as it is subdivided by additional hyperplanes of \mathcal{H}_b when $n \geq 4$; and conversely, not all facets of \mathcal{H}_b are contained in the sets F_b^{ij} . The weighted HIP (\mathcal{H}_B, w_B) of a bid collection can be constructed by initially letting $\mathcal{H}_B = \bigcup_{b \in \mathcal{B}} \mathcal{H}_b$ and $w_B = \sum_{b \in \mathcal{B}} w_b$. We then remove from \mathcal{H}_B all hyperplanes containing only zero-weighted facets, to ensure parsimony. This construction mirrors that of \mathcal{L}_B in Appendix C.

We conclude with two technical observations about HIPs that we will use below.

OBSERVATION G.3: Let \mathcal{H} be a HIP. For any point $\mathbf{p} \in \mathcal{P}$, there exists a sufficiently small ε such that every hyperplane $H \subseteq \mathcal{H}$ either contains \mathbf{p} or does not intersect $N_\varepsilon(\mathbf{p})$, the ε -neighbourhood of \mathbf{p} .

LEMMA G.4: Let G be an $(n-2)$ -dimensional face of HIP \mathcal{H} .

- (i) If G is contained in two hyperplanes $H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ and $H(\mathbf{p}; k, l; \frac{s_k}{s_l})$ of \mathcal{H} with distinct $i, j, k, l \in [n]_0$, then G is not contained in any further hyperplane H of \mathcal{H} .
- (ii) If G is contained in three hyperplanes $H(\mathbf{p}; i, j; \frac{s_i}{s_j})$, $H(\mathbf{p}; i, k; \frac{s_i}{s_k})$ and $H(\mathbf{p}; k, j; \frac{s_k}{s_j})$ of \mathcal{H} with distinct $i, j, k \in [n]$, then G is not contained in any further hyperplane H of \mathcal{H} .
- (iii) If G is contained in two hyperplanes $H(\mathbf{p}; i, 0; 1)$ and $H(\mathbf{p}; j, 0; 1)$, then G is also contained in all hyperplanes of \mathcal{H} of the form $H(\mathbf{p}; i, j; \frac{a}{b})$, and not contained in any further hyperplane H of \mathcal{H} .

PROOF: Without loss of generality, let $\mathbf{p} \in \text{relint}(G)$. Write an additional hyperplane of \mathcal{H} passing through \mathbf{p} as $H = H(\mathbf{p}; x, y; \frac{t_x}{t_y})$ with $x > y$.

Consider statement (i). First, if $x, y \notin \{i, j, k, l\}$, then G contains point $\mathbf{p} + \delta e^x$ but H does not, so G is not contained in H . Second, suppose $x = i$. Allow y to be any good in $[n]_0 \setminus \{i\}$, but observe that if $y = j$ then the case $\frac{t_x}{t_y} = \frac{t_i}{t_j} = \frac{s_i}{s_j}$ means that H is identical to one of the existing facets containing G , so exclude that case. Then $\mathbf{p} + \delta(s_j e^i + s_i e^j)$ lies in G but not in H , for sufficiently small $\delta > 0$ (note that i and j cannot both be 0).

Statement (ii) is proved analogously with the same two possibilities for H using respectively points $\mathbf{p} + \delta e^x$ and $\mathbf{p} + \delta(\frac{1}{s_i} e^i + \frac{1}{s_j} e^j + \frac{1}{s_k} e^k)$.

Finally, consider statement (iii). $H(\mathbf{p}; i, 0; 1) \cap H(\mathbf{p}; j, 0; 1) \subseteq H(\mathbf{p}; i, j; \frac{a}{b})$ is immediate for any $\frac{a}{b} \in \mathbb{Q}_S$, so any hyperplanes $H(\mathbf{p}; i, j; \frac{a}{b})$ in \mathcal{H} contain G . The argument that $H(\mathbf{p}; x, y; \frac{a}{b})$ does not contain G if $\{x, y\} \neq \{i, j\}$ is analogous to the proof of Part (i). Q.E.D.

G.2. Constructing the bid collection

Let $v : X \rightarrow \mathbb{R}$ be a concave substitutes valuation for n goods with finite domain $X \subseteq \mathcal{X}$. We now construct the corresponding bid collection \mathcal{B} . Appendix G.3 then proves $D_B = D_v$.

The bounding box. Define a *bounding box* (or *cube*) $[\mathcal{P}] := \{\mathbf{p} \in \mathcal{P} \mid \underline{C}_i \leq p_i \leq \overline{C}_i, \forall i \in [n]\}$, choosing boundary values $(\underline{C}_i, \overline{C}_i)_{i \in [n]}$ so that the interior of the box intersects all faces of \mathcal{H}_v . We call the points $\{\mathbf{x} \in [\mathcal{P}] \mid x_i = \underline{C}_i\}$ and $\{\mathbf{x} \in [\mathcal{P}] \mid x_i = \overline{C}_i\}$ the *lower* and *upper i -boundaries* of the box. Let $[\mathcal{L}_v], [\mathcal{H}_v]$ be the union of respectively $\mathcal{L}_v, \mathcal{H}_v$ with the facets of the box extended to hyperplanes. For any facet F of \mathcal{L}_v or \mathcal{H}_v , $[F]$ is the intersection of F with $[\mathcal{P}]$. As in Section 3.3 and Appendix A, we use the vertices of $[\mathcal{L}_v]$ to define the roots of our bids. Vertices of \mathcal{L}_v in the interior of the bounding box lead to regular bids (interested in all goods). If a vertex lies on a lower i -boundary, the corresponding bid is not interested in good i and we set its i th root entry to $r_i = -\infty$. If a vertex lies on any upper boundary, the bid is not

interested in the null good 0, and we set $r_i = -\infty$. All other root entries are set to correspond to the vertex's entries. We therefore record the lower i -boundaries of the bounding box on which point $\mathbf{p} \in [\mathcal{P}]$ lies, as well as whether it lies on any upper boundaries, by defining:

$$I(\mathbf{p}) := \{k \in [n]_0 \mid p_k > \underline{C}_k, \text{ or } k = 0 \wedge (\forall l \in [n] : p_l < \overline{C}_l)\}.$$

If \mathbf{p} lies in the interior of $[\mathcal{P}]$, for instance, we have $I(\mathbf{p}) = [n]_0$. Define $\tau : [\mathcal{P}] \rightarrow \mathcal{R}$ by $\tau(\mathbf{p})_i = p_i$ if $i \in I(\mathbf{p})$ (recalling that $p_0 = 0$) and $\tau(\mathbf{p})_i = -\infty$ if $i \in [n]_0 \setminus I(\mathbf{p})$. This function fixes prices of goods in $I(\mathbf{p})$, reduces the prices of goods not in $I(\mathbf{p})$ to $-\infty$, and records whether $0 \in I(\mathbf{p})$. It is clearly injective, and so invertible on its image. Note that bids in its image which are not interested in good 0 are normalised in a different way from that used in the body text (see Lemma 2.2 and following paragraphs).

Orientations and slope vectors. Fix some $S \in \mathbb{N}$ so that the (rational) slopes of all hyperplanes in \mathcal{H}_v are contained in $\mathbb{Q}_S := \{\frac{a}{b} \mid 1 \leq a, b \leq S\}$. We will refer to this value S throughout Appendices G.2 and G.3. The smallest and largest slopes in \mathbb{Q}_S are called *extremal*. We also define two functions $\eta^\chi(\frac{a}{b})$ with $\chi \in \{-1, 1\}$ on the set of slopes $\mathbb{Q}_S \setminus \{S^\times\}$, such that $\eta^{-1}(\frac{a}{b})$ is the next-smaller and $\eta^1(\frac{a}{b})$ is the next-larger slope for $\frac{a}{b} \in \mathbb{Q}_S$.⁵⁸ An *orientation at prices* $\mathbf{p} \in [\mathcal{P}]$ is an ordered pair (i, j) of distinct goods $i, j \in I(\mathbf{p})$. We now define *slope vector(s) for orientation (i, j) at \mathbf{p}* . If $0 \in \{i, j\}$, then the (unique) *slope vector for (i, j) at \mathbf{p}* consists of the vector $\mathbf{s} \in \mathcal{T}$ with $s_k = 1$ for $k \in I(\mathbf{p})$ and $s_k = 0$ for $k \in [n] \setminus I(\mathbf{p})$. If $0 \notin \{i, j\}$, then the set \mathcal{S} of *slope vectors for (i, j) at \mathbf{p}* consists of the vectors $\mathbf{s} \in \mathcal{T}$ whose positive entries are those that are indexed by $I(\mathbf{p}) \cup \{0\}$ and that satisfy $\frac{s_k}{s_l} \in \mathbb{Q}_S$ for all $k, l \in I(\mathbf{p})$.

Creating the bid collection \mathcal{B} . First we consider each vertex \mathbf{p} of $[\mathcal{L}_v]$ with $|I(\mathbf{p})| \geq 2$. Fix such a vertex \mathbf{p} and the orientation (i, j) at \mathbf{p} with the two largest goods $i > j$ in $I(\mathbf{p})$. For every slope vector \mathbf{s} for (i, j) at \mathbf{p} , add bid $(\tau(\mathbf{p}); \mathbf{s}; m)$ to \mathcal{B} . The bid's multiplicity m is the value $m^{ij}(\mathbf{p}; \mathbf{s})$ of the multiplicity function m^{ij} that will be defined in Appendix G.2.2 below. The multiplicity function computes the weighted sum of weights of certain facets containing \mathbf{p} in the hyperplane $H(\mathbf{p}; i, j; \frac{s_i}{s_j})$. These facets overlap with 'sliver regions' of this hyperplane, which are introduced in Appendix G.2.1. Corollary G.18 in Appendix G.3 shows that m^{ij} is integral, so the multiplicities of the bids we construct are well-defined. To avoid redundancy, we only add bids for which $m^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$. As we will show later (xx), this holds only if \mathcal{L}_v contains a facet with normal $s_k e^l - s_l e^k$ that meets \mathbf{p} , for all $k, l \in [n]$.

Once we have considered all vertices of $[\mathcal{L}_v]$ with $|I(\mathbf{p})| \geq 2$, we add a single-minded bid to \mathcal{B} for each good $i \in [n]$. Fix generic prices \mathbf{p} so that $D_v(\mathbf{p}) = \{\mathbf{x}\}$ and $D_B(\mathbf{p}) = \{\mathbf{y}\}$. For each good $i \in [n]$, we add to \mathcal{B} the bid $(\mathbf{r}; e^i; x_i - y_i)$ interested only in good i , where the root $\mathbf{r} \in \mathcal{R}$ satisfies $r_i = 0$ and $r_k = -\infty$ for $k \in [n]_0 \setminus \{i\}$.

Our procedure to compute \mathcal{B} is formally stated in Algorithm 1. As a final step one can re-normalise these bids as in the body text (see Lemma 2.2 and the following paragraphs) but we will use the normalisation provided by Algorithm 1 throughout this appendix.

G.2.1. Sliver regions

Suppose we wish to add a bid \mathbf{b} for vertex \mathbf{p} , orientation (i, j) and slope vector \mathbf{s} , as described in Algorithm 1. In order to determine its multiplicity m , we consider the weights of

⁵⁸The next-smaller slope in \mathbb{Q}_S for $\frac{a}{b}$ is the largest fraction $\sigma \in \mathbb{Q}_S$ with $\sigma < \frac{a}{b}$, and likewise for the next-larger slope. Note that \mathbb{Q}_S does not contain a next-smaller or next-larger slope if $\frac{a}{b} = S^{-1}$ or $\frac{a}{b} = S$, respectively, and the functions η^χ are not defined for these slopes.

Algorithm 1 Constructing the bid collection \mathcal{B} .

-
- 1: Initialise $\mathcal{B} = \emptyset$.
 - 2: **for** every vertex \mathbf{p} of $[\mathcal{L}_v]$ with $|I(\mathbf{p})| \geq 2$ **do**
 - 3: Fix orientation (i, j) with the two largest goods $i > j \in I(\mathbf{p})$.
 - 4: **for** every slope vector \mathbf{s} for (i, j) at \mathbf{p} **do**
 - 5: Add bid $\mathbf{b} = (\tau(\mathbf{p}); \mathbf{s}; m_v^{ij}(\mathbf{p}; \mathbf{s}))$ to \mathcal{B} if $m_v^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$.
 - 6: Fix generic prices $\mathbf{p} \in \mathcal{P}$ such that $D_v(\mathbf{p}) = \{\mathbf{x}\}$ and $D_{\mathcal{B}}(\mathbf{p}) = \{\mathbf{y}\}$.
 - 7: **for** every good $i \in [n]$ **do**
 - 8: Add bid $\mathbf{b} = (\mathbf{r}, \mathbf{e}^i; x_i - y_i)$ to \mathcal{B} with $r_i = 0$ and $r_k = -\infty$ for all $k \in [n]_0 \setminus \{i\}$.
 - 9: **return** \mathcal{B} .
-

carefully selected facets in $H := H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ meeting at \mathbf{p} . As there are potentially many facets in H meeting at \mathbf{p} and only a subset of these are considered when computing m , we now systematically define ‘sliver regions’ of H around \mathbf{p} so that every facet of interest intersects $(n-1)$ -dimensionally with at least one of these slivers. We will see in Lemma G.11 that any two facets overlapping with the same sliver region have the same weight. This allows us to associate each sliver region with the weight of its overlapping facets. In Appendix G.2.2, we will formally define the multiplicity m as the signed sum of the weights of the sliver regions in H around \mathbf{p} .

Let \mathbf{p} be an arbitrary point in $[\mathcal{P}]$ and (i, j) be an orientation at \mathbf{p} with $i > 0$. For every slope vector \mathbf{s} for (i, j) at \mathbf{p} and every set $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$, we define a corresponding sliver region $R^{ij}(\mathbf{p}; \mathbf{s}; J)$ around \mathbf{p} contained in $H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ as follows.

DEFINITION G.5: Fix a point $\mathbf{p} \in [\mathcal{P}]$, orientation (i, j) at \mathbf{p} with $i > 0$, slope vector \mathbf{s} for (i, j) at \mathbf{p} , set $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$, and $\varepsilon > 0$ such that Observation G.3 holds. For each $k \in [n]_0 \setminus \{i, j\}$, define

$$Q^{kj}(\mathbf{p}; \mathbf{s}; J) = \begin{cases} H^+(\mathbf{p}; k, j; \frac{1}{s}) & \text{if } k \notin I(\mathbf{p}) \text{ and } k \neq 0, \\ H^{-x_0^J}(\mathbf{p}; 0, j; 1) & \text{if } k = 0, \\ H^{-x_k^J}(\mathbf{p}; k, 0; 1) & \text{if } k \in I(\mathbf{p}) \setminus \{0\} \text{ and } j = 0, \\ H^{-x_k^J}(\mathbf{p}; k, j; \frac{s_k}{s_j}) & \text{if } k \in I(\mathbf{p}) \setminus \{0\} \text{ and } \frac{s_k}{s_j} = S^{x_0^J}, \\ H^{-x_k^J}(\mathbf{p}; k, j; \frac{s_k}{s_j}) \cap H^{x_k^J}(\mathbf{p}; k, j; \eta^{x_0^J}(\frac{s_k}{s_j})) & \text{else.} \end{cases} \quad (\text{G.3})$$

and the sliver region

$$R^{ij}(\mathbf{p}; \mathbf{s}; J) = [\mathcal{P}] \cap N_\varepsilon(\mathbf{p}) \cap H(\mathbf{p}; i, j; \frac{s_i}{s_j}) \cap \bigcap_{k \in [n]_0 \setminus \{i, j\}} Q^{kj}(\mathbf{p}; \mathbf{s}; J). \quad (\text{G.4})$$

A sliver region with $0 \notin \{i, j\}$ is *extremal* if $\frac{s_k}{s_j} = S^{x_0^J}$ for every $k \in I(\mathbf{p}) \setminus \{i, j\}$.

To understand the intuition of this construction, suppose we include a bid $\mathbf{b} = (\mathbf{r}; \mathbf{s}; m)$ (where $\mathbf{r} = \tau(\mathbf{p})$ for some $\mathbf{p} \in [\mathcal{P}]$) in an existing set of bids \mathcal{B} , and consider the effect this will have on the weights of facets in $\mathcal{H}_{\mathcal{B}}$ lying in $H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ where $i, j \in I = I(\mathbf{p})$. As we showed in and just below Equation (G.1), it will provide an additional weight of $m \gcd(s_i, s_j)$ on any facet with $(n-1)$ -dimensional intersection with (and so contained in) $F_{\mathbf{b}}^{ij} = H(\mathbf{p}; i, j; \frac{s_i}{s_j}) \cap$




FIGURE G.2.—An illustration of sliver regions contained in a plane in a setting with $n = 3$ goods when $j = 0$ and when $0 \notin \{i, j\}$. (a) Example with $(i, j) = (3, 0)$. The local neighbourhood of \mathbf{p} in the plane with orientation $(3, 0)$ is partitioned into 4 sliver regions $R^{30}(\mathbf{p}; \mathbf{s}; J)$ by the planes normal to \mathbf{e}^1 and \mathbf{e}^2 (the planes $H(\mathbf{p}; k, 0; 1)$ for $k = 1, 2$). They are labelled by the corresponding set J . (b) Example with $(i, j) = (3, 2)$ and $S = 2$. The local neighbourhood of \mathbf{p} in a plane with orientation $(3, 2)$ and slope 1 is partitioned into 8 sliver regions R^1, \dots, R^8 by the ‘horizontal’ plane normal to \mathbf{e}^2 and the three planes with orientation $(1, 2)$ and slopes $\mathbb{Q}_S = \{\frac{1}{2}, 1, 2\}$.

$\bigcap_{k \in I \setminus \{i, j\}} H^+(\mathbf{p}; k, j; \frac{s_k}{s_j})$. By construction, the sliver region $R^{ij}(\mathbf{p}; \mathbf{s}; \emptyset)$ intersects $(n - 1)$ -dimensionally with F_b^{ij} . However, other bids already in \mathcal{B} will also introduce weights on facets corresponding in the same way to their own root and trade-offs, and these facets may overlap with those of \mathbf{b} . The multiplicities of these bids will be tuned so that the weights of the other facets meeting \mathbf{p} are correct. So, to find the correct multiplicity for \mathbf{b} , we need to take a signed sum of weights of facets containing \mathbf{p} —considering whether they are on the other side of the hyperplanes bounding F_b^{ij} . So for each $k \in I(\mathbf{p})$, inclusion of $k \in J$ denotes whether we have swapped to the other side of the hyperplane bounding F_b^{ij} with orientation (k, j) .

However, because there may be multiple hyperplanes with orientation (k, j) , we must not go “too far” on either side of these hyperplanes. Consider $\mathbf{q} \in H(\mathbf{p}; i, j; \frac{s_i}{s_j})$. Whether \mathbf{q} is in the positive or negative half-space defined by $H(\mathbf{p}; k, j; \frac{s_k}{s_j})$ depends on whether $s_k(q_k - p_k) \geq s_j(q_j - p_j)$ or $s_k(q_k - p_k) \leq s_j(q_j - p_j)$. And, for generic \mathbf{q} , this depends on whether $\frac{(q_k - p_k)}{(q_j - p_j)}$ is greater than or less than $\frac{s_k}{s_j}$; but note that the way in which the latter inequalities correspond with the former, depends on the sign of $q_j - p_j$, and so on whether $\mathbf{q} \in H^+(\mathbf{p}; 0, j; 1)$. But moreover, if $\frac{(q_k - p_k)}{(q_j - p_j)}$ strays too far from equality with $\frac{s_k}{s_j}$ then \mathbf{q} will be on the other side of another hyperplane with orientation (i, j) ; this happens whenever $\frac{(q_k - p_k)}{(q_j - p_j)}$ crosses the threshold of any $\sigma \in \mathbb{Q}_S$. So we must specify that \mathbf{q} lies on the correct side of the next hyperplane along, with respect to the slopes possible, that is, those in \mathbb{Q}_S . That is what the fifth form for $Q^{kj}(\mathbf{p}; \mathbf{s}; J)$ achieves.

Figure G.2 illustrates the sliver regions for three goods when $I(\mathbf{p}) = \{0, 2, 3\}$ in both the cases $j = 0$ and $j > 0$. Figure G.2 (left) shows the case of $j = 0$: here \mathbf{s} is uniquely defined and S is not relevant. H is partitioned by one hyperplane of orientation $(k, 0)$ for every $k \in [n]_0 \setminus \{i, j\}$. It follows for every sliver region that the Q^{kj} consist of the single half-space $H^-(\mathbf{p}; k, 0; 1)$ or $H^+(\mathbf{p}; k, 0; 1)$; it is the positive half-space when $k \notin I(\mathbf{p})$. Thus the sliver regions $R^{30}(\mathbf{p}; \mathbf{s}; J)$ are the four regions of Figure G.2 (left), where they are labelled with the corresponding set J .

When $j > 0$, and $S \geq 2$, H is partitioned by multiple hyperplanes of the same orientation (k, j) (with $k \neq 0$) and different slopes in \mathbb{Q}_S . And not all sliver regions are extremal. In the example of Figure G.2 (right), where $S = 2$, $(i, j) = (3, 2)$ and $\frac{s_3}{s_2} = 1$ (so H has slope 1), the extremal sliver regions are the regions labelled R^1 , R^4 , R^5 and R^8 . Region R^5 is identified by pair $((2, 1, 1), \emptyset)$ and thus can be written as $R^{32}(\mathbf{p}; (2, 1, 1); \emptyset)$. Region R^4 is identified by pair $((1, 2, 2), \{0\})$, region R^1 by pair $((2, 1, 1), \{0, 1\})$, and region R^8 by pair $((1, 2, 2), \{1\})$. The regions labelled R^2 , R^3 , R^6 and R^7 are not extremal, and are each identified by two pairs. The sliver region R^2 , for instance, can be written as $R^{32}(\mathbf{p}; (2, 1, 1); \{0\})$ and $R^{32}(\mathbf{p}; (1, 1, 1); \{0, 1\})$; and sliver region R^7 can be written as $R^{32}(\mathbf{p}; (1, 1, 1); \{1\})$ and $R^{32}(\mathbf{p}; (1, 2, 2); \emptyset)$.

We now develop important properties of sliver regions analytically. The construction of the sets $Q^{kj}(\mathbf{p}; \mathbf{s}, J)$ and the sliver regions allows us to make the following three observations.

OBSERVATION G.6: For any $k \in I(\mathbf{p})$, we have $Q^{kj}(\mathbf{p}; \mathbf{s}, J) \cap H(\mathbf{p}; k, j; \frac{s_k}{s_j}) = H(\mathbf{p}; k, j; \frac{s_k}{s_j})$.

OBSERVATION G.7: For any prices \mathbf{p} and slopes \mathbf{s} , the subset relationship $R^{ij}(\mathbf{p}; \mathbf{s}, J) \subseteq H(\mathbf{p}; i, j; \frac{s_i}{s_j}) \cap \bigcap_{k \in I(\mathbf{p}) \setminus \{i, j\}} H^+(\mathbf{p}; i, k; \frac{s_k}{s_j})$ holds if and only if $J = \emptyset$.

OBSERVATION G.8: Suppose (i, j) is an orientation with $0 \notin \{i, j\}$. For any $k \in I(\mathbf{p}) \setminus \{0, i, j\}$ with $\frac{s_k}{s_j} \neq S^{x_{0k}^J}$, we have

$$R^{ij}(\mathbf{p}; \mathbf{s}, J) = \begin{cases} R^{ij}(\mathbf{p}; \mathbf{s}', J \setminus \{k\}) & \text{if } k \in J, \\ R^{ij}(\mathbf{p}; \mathbf{s}', J \cup \{k\}) & \text{if } k \notin J, \end{cases}$$

where \mathbf{s}' is the slope vector for (i, j) at \mathbf{p} with $\frac{s'_k}{s'_j} = \eta^{x_{0k}^J}(\frac{s_k}{s_j})$ and $\frac{s'_l}{s'_j} = \frac{s_l}{s_j}$ for all $l \in I(\mathbf{p}) \setminus \{j, k\}$.

From Observation G.8 we can conclude:

COROLLARY G.9: A sliver region is extremal if and only if it is identified by a single pair (\mathbf{s}, J) . A non-extremal sliver region is identified by multiple pairs (\mathbf{s}, J) , and the number of identifying pairs with even and odd set cardinality is the same.

LEMMA G.10: $R^{ij}(\mathbf{p}; \mathbf{s}, J)$ is $(n - 1)$ -dimensional for any $\mathbf{p} \in [\mathcal{P}]$, orientation (i, j) at \mathbf{p} with $i > 0$, slope vector \mathbf{s} for (i, j) at \mathbf{p} , and $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$. If $j \neq 0$ then $R^{ij}(\mathbf{p}; \mathbf{s}, J)$ and $R^{ji}(\mathbf{p}; \mathbf{s}, J)$ have $(n - 1)$ -dimensional intersection.

PROOF: Fix $\mathbf{p} \in [\mathcal{P}]$ and suppose first that $\varepsilon > 0$ is such that $\varepsilon < p_k - \underline{C}_k$ for all $k \in [n]$ whenever this is greater than 0, and similarly $\varepsilon < \overline{C}_k - p_k$ for all $k \in [n]$ whenever this is greater than 0. Fix infinitesimal $\delta > 0$ such that $\delta < \frac{\varepsilon}{2nS}$. We will use δ to construct a point $\mathbf{q} \in [\mathcal{P}]$ that lies in $R^{ij}(\mathbf{p}; \mathbf{s}, J)$, and also in $R^{ji}(\mathbf{p}; \mathbf{s}, J)$ when $j \neq 0$. Moreover, we see that it lies in the interior of every bounding half-space intersecting to define these sliver region(s), apart from $H^{ij}(\mathbf{p}; \mathbf{s}, J)$, and thus so does any point in an infinitesimal neighbourhood of \mathbf{q} in $H^{ij}(\mathbf{p}; \mathbf{s}, J)$. This demonstrates that the sliver region, and where relevant the intersection, has dimension $(n - 1)$.

Take the case $0 \notin \{i, j\}$. First, set $q_i = p_i + \chi_0^J \frac{\delta}{s_i}$ and $q_j = p_j + \chi_0^J \frac{\delta}{s_j}$. This ensures that $\mathbf{q} \in H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ and that \mathbf{q} lies in the interiors of $H^{-x_0^J}(\mathbf{p}; 0, i, 1)$ and $H^{-x_0^J}(\mathbf{p}; 0, j, 1)$.

Second, for any $0 \neq k \notin I(\mathbf{p})$, let $q_k = p_k + 2S\delta$, which ensures that \mathbf{q} lies in the interiors of $Q^{kj}(\mathbf{p}; \mathbf{s}; J) = H^+(\mathbf{p}; k, j; \frac{1}{S})$ and $Q^{ki}(\mathbf{p}; \mathbf{s}; J) = H^+(\mathbf{p}; k, i; \frac{1}{S})$.

Third, for any $k \in I(\mathbf{p}) \setminus \{0, i, j\}$ we let $q_k = p_k + (\chi_0^J - \delta\chi_k^J) \frac{\delta}{s_k}$. This ensures that \mathbf{q} lies in the interior of $H^{-\chi_k^J}(\mathbf{p}; k, j; \frac{s_k}{s_j})$ and $H^{-\chi_k^J}(\mathbf{p}; k, i; \frac{s_k}{s_i})$. Moreover, \mathbf{q} lies in the interior of $H^{\chi_k^J}(\mathbf{p}; k, j; \eta^{\chi_{0k}^J}(\frac{s_k}{s_j}))$ if $\frac{s_k}{s_j} \neq S^{\chi_{0k}^J}$, and in the interior of $H^{\chi_k^J}(\mathbf{p}; k, i; \eta^{\chi_{0k}^J}(\frac{s_k}{s_i}))$ if $\frac{s_k}{s_i} \neq S^{\chi_{0k}^J}$. This fully specifies \mathbf{q} , and shows that it lies in the interior of $Q^{kj}(\mathbf{p}; \mathbf{s}; J)$ and $Q^{ki}(\mathbf{p}; \mathbf{s}; J)$ for all $k \in [n]_0 \setminus \{i, j\}$.

Finally, note that $\mathbf{q} \in N_\varepsilon(\mathbf{p})$ by definition of δ . To demonstrate that $\mathbf{q} \in [\mathcal{P}]$, first note that by definition of ε , for all $k \in [n]$ we have $\underline{C}_k < p_k \implies \underline{C}_k < q_k$ and $p_k < \overline{C}_k \implies q_i < \overline{C}_k$. If $p_k = \underline{C}_k$ for some $k \in [n]$ then $k \notin I(\mathbf{p})$ (and so in particular $k \neq i, j$) and so $q_k > p_k = \underline{C}_k$. And if $p_k = \overline{C}_k$ for any $k \in [n]$ then $k \in I(\mathbf{p})$, $0 \notin I(\mathbf{p})$ and so $q_k < p_k = \overline{C}_k$. So indeed $\mathbf{q} \in [\mathcal{P}]$, which completes this case.

Now turn to the case $j = 0$. Now the slope vector \mathbf{s} is unique with $s_k = 1$ for all $k \in I(\mathbf{p}) \cup \{0\}$. Let $q_i = p_i$ so that $\mathbf{q} \in H(\mathbf{p}; i, 0; 1)$, and let $q_k = p_k - \delta\chi_k^J$ for $k \in [n] \setminus \{i\}$. If $k \notin I(\mathbf{p})$ then $k \notin J$ so \mathbf{q} lies in the interior of $H^+(\mathbf{p}; k, 0; 1) = Q^{k0}(\mathbf{p}; \mathbf{s}; J)$. And if $k \in I(\mathbf{p}) \setminus \{i, 0\}$ then $\mathbf{q} \in H^{-\chi_k^J}(\mathbf{p}; k, 0; 1) = Q^{k0}(\mathbf{p}; \mathbf{s}; J)$. This fully specifies \mathbf{q} , and shows that it lies in the interior of $Q^{kj}(\mathbf{p}; \mathbf{s}; J)$ and $Q^{ki}(\mathbf{p}; \mathbf{s}; J)$ for all $k \in [n]_0 \setminus \{i, j\}$.

Again, $\mathbf{q} \in N_\varepsilon(\mathbf{p})$ by definition of δ . Being in the case $j = 0$ implies that $0 \in I(\mathbf{p})$ and so $p_k < \underline{C}_k$ for all $k \in [n]$. So, $\mathbf{q} \in N_\varepsilon(\mathbf{p})$ and the definition of ε is sufficient for $q_k < \underline{C}_k$ for all $k \in [n]$. We similarly have $\underline{C}_k < p_k \implies \underline{C}_k < q_k$. If $p_k = \underline{C}_k$ for some $k \in [n]$ then $k \notin I(\mathbf{p})$ (and so in particular $k \neq i$) and so $q_k > p_k = \underline{C}_k$. So indeed $\mathbf{q} \in [\mathcal{P}]$, which completes this case.

It was more convenient to prove this result using assumptions on the value of ε defining the sliver regions; observe that the result still holds if we relax these assumptions, as we will simply get larger sets containing the same $(n-1)$ -dimensional neighbourhoods of \mathbf{q} . *Q.E.D.*

G.2.2. Computing multiplicities

We now introduce the multiplicity function, which maps any combination of a point, orientation and slope vector to an integer. This definition is more general than required for the algorithm constructing the bid collection \mathcal{B} , and will be used in full generality when we prove Theorem 3.1. Throughout this section, let (\mathcal{H}, w) be a balanced weighted HIP whose hyperplanes have slopes in \mathbb{Q}_S .

Fix some point $\mathbf{p} \in \mathcal{P}$, orientation (i, j) at \mathbf{p} with $i > 0$, and slope vector \mathbf{s} for (i, j) at \mathbf{p} . We assume that $H := H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ is a hyperplane in \mathcal{H} by temporarily adding H as a dummy hyperplane with zero-weighted facets if necessary. Each sliver region $R^{ij}(\mathbf{p}; \mathbf{s}; J)$ around \mathbf{p} in H has $(n-1)$ -dimensional intersection with one or more facets of \mathcal{H} contained in H . Lemma G.11 tells us that all such facets have the same weight, and we denote this weight by $w^{ij}(\mathbf{p}; \mathbf{s}; J)$. We also introduce the shorthand $w^{ij}(\mathbf{p}; J) := w^{ij}(\mathbf{p}; \mathbf{s}; J)$ when $j = 0$, as the slope vector at \mathbf{p} for $(i, 0)$ is unique.

LEMMA G.11: *Fix $\mathbf{p} \in [\mathcal{P}]$, orientation (i, j) at \mathbf{p} with $i > 0$, and slope vector \mathbf{s} for (i, j) at \mathbf{p} . All facets that have $(n-1)$ -dimensional intersection with sliver region $R := R^{ij}(\mathbf{p}; \mathbf{s}; J)$ have the same weight $w^{ij}(\mathbf{p}; \mathbf{s}; J)$.*

PROOF: Let F and F' be two adjacent facets of \mathcal{H} contained in hyperplane $H := H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ that intersect full-dimensionally with R . Suppose we can show that F and F' are not separated by any hyperplanes of orientation (k, j) with $k \in [n]_0 \setminus \{j\}$. Then F and

F' are also not separated by any hyperplanes of orientation (k, i) with $k \in [n]_0 \setminus \{i\}$, as $H^+(\mathbf{p}; k, j; \sigma) \cap H = H^+(\mathbf{p}; k, i; \sigma \frac{s_j}{s_i}) \cap H$ and $H^-(\mathbf{p}; k, j; \sigma) \cap H = H^-(\mathbf{p}; k, i; \sigma \frac{s_j}{s_i}) \cap H$. Hence, the face G of \mathcal{H} separating F and F' is contained in a hyperplane H' of orientation (k, l) with $\{k, l\} \cap \{i, j\} = \emptyset$. By Lemma G.4, G is contained in no other hyperplanes of \mathcal{H} . Since \mathcal{H} is balanced around G (cf. Observation G.1) and the normal vectors to H and H' are linearly independent, it follows that $w(F) = w(F')$.

It remains to show that hyperplane $H' := H(\mathbf{p}; k, j; \sigma)$ does not separate any two points in R if $k \in [n]_0$ and σ is any slope in \mathbb{Q}_S . Fix $\mathbf{q}, \mathbf{q}' \in \text{relint } R$.

Suppose first that $j = 0$ or $k = 0$. Then, as there is only one hyperplane with orientation $(k, 0)$ or respectively $(0, j)$ passing through \mathbf{p} , the assumption that $\mathbf{q}, \mathbf{q}' \in \text{relint } R$ implies that they lie on the same side of H' . So we assume from here that $j, k > 0$.

Next suppose that $k = i$ and $\sigma \neq \frac{s_i}{s_j}$. As $\mathbf{q}, \mathbf{q}' \in H$, we have $(q_i - p_i) = \frac{s_j}{s_i}(q_j - p_j)$ and $(q'_i - p_i) = \frac{s_j}{s_i}(q'_j - p_j)$. Suppose \mathbf{q} and \mathbf{q}' are separated by H' and without loss of generality that $\sigma(q_i - p_i) < (q_j - p_j)$ and $\sigma(q'_i - p_i) > (q'_j - p_j)$. Hence, $\frac{s_j}{s_i}\sigma(q_j - p_j) < (q_j - p_j)$ and $\frac{s_j}{s_i}\sigma(q'_j - p_j) > (q'_j - p_j)$. These equations together imply that $q_j - p_j$ and $q'_j - p_j$ do not have the same sign, that is, \mathbf{q} and \mathbf{q}' lie in different half-spaces of $H(\mathbf{p}; 0, j; 1)$. But $R \subseteq Q^{0j}(\mathbf{p}; \mathbf{s}; J)$ which is either $H^+(\mathbf{p}; 0, j; 1)$ or $H^-(\mathbf{p}; 0, j; 1)$, so this contradicts our assumption that \mathbf{q} and \mathbf{q}' both lie in R .

Now we consider $k \notin I(\mathbf{p})$. Recall that $k \notin I(\mathbf{p})$ means that $p_k = C_k$, and so $R \subseteq [\mathcal{P}]$ means that $q_k, q'_k > p_k$. Now $\mathbf{q} \in Q^{kj}(\mathbf{p}; \mathbf{s}; J) = H^+(\mathbf{p}; k, j; \frac{1}{S})$ implies $q_k - p_k \geq S(q_j - p_j)$; since $q_k - p_k > 0$ we can re-write this as $\frac{q_j - p_j}{q_k - p_k} \leq \frac{1}{S}$. It follows by definition of S that $\frac{q_j - p_j}{q_k - p_k} \leq \sigma$ for any $\sigma \in \mathbb{Q}_S$, and so that $\mathbf{q} \in H^+(\mathbf{p}; k, j; \sigma)$. As the same holds for \mathbf{q}' , the points \mathbf{q}, \mathbf{q}' cannot be separated by H' .

Finally, consider $k \in I(\mathbf{p}) \setminus \{i, j, 0\}$ and suppose that $0, k \in J$. If $\sigma \leq \frac{s_k}{s_j}$, then $\mathbf{q}, \mathbf{q}' \in H^-(\mathbf{p}; k, j; \frac{s_k}{s_j}) \cap H^-(\mathbf{p}; 0, j; 1) \subseteq H^-(\mathbf{p}; k, j; \sigma)$. If $\sigma > \frac{s_k}{s_j}$, then $\frac{s_k}{s_j} \neq S = S^{X_{0k}^J}$ and so also $\mathbf{q}, \mathbf{q}' \in H^+(\mathbf{p}; k, j; \eta^{+1}(\frac{s_k}{s_j})) \cap H^-(\mathbf{p}; 0, j; 1) \subseteq H^+(\mathbf{p}; k, j; \sigma)$. So \mathbf{q}, \mathbf{q}' are not separated by H' . The other possibilities for inclusion of $0, k \in J$ may be seen similarly. Q.E.D.

COROLLARY G.12: $R^{ij}(\mathbf{p}; \mathbf{s}; J)$ and $R^{ji}(\mathbf{p}; \mathbf{s}; J)$ have the same weight $w^{ij}(\mathbf{p}; \mathbf{s}; J) = w^{ji}(\mathbf{p}; \mathbf{s}; J)$ if $0 \notin \{i, j\}$.

PROOF: Lemma G.10 tells us there is some facet F that has $(n-1)$ -dimensional intersection with both sliver regions, so the claim follows by Lemma G.11. Q.E.D.

The *multiplicity function* introduced in Definition G.13 computes a weighted sum of the sliver region weights. Corollary G.18 in Appendix G.3 shows that this function is integral, ensuring that the multiplicities of the bids in \mathcal{B} are well-defined.

DEFINITION G.13: For any point $\mathbf{p} \in [\mathcal{P}]$, orientation (i, j) at \mathbf{p} with $i > 0$, and slope vector \mathbf{s} for (i, j) at \mathbf{p} , the *multiplicity* (w.r.t. any HIP (\mathcal{H}, w)) is given by

$$m^{ij}(\mathbf{p}; \mathbf{s}) := \frac{1}{\gcd(s_i, s_j)} \sum_{J \subseteq I(\mathbf{p}) \setminus \{i, j\}} (-1)^{|J|} w^{ij}(\mathbf{p}; \mathbf{s}; J). \quad (\text{G.5})$$

We also refer to the multiplicity functions with respect to the weighted HIPs (\mathcal{H}_v, w_v) and $(\mathcal{H}_{\mathcal{B}'}, w_{\mathcal{B}'})$ for any bid collection \mathcal{B}' as m_v^{ij} and $m_{\mathcal{B}'}^{ij}$.

Note that $m^{ij}(\mathbf{p}; \mathbf{s}) = 0$ if the LIP \mathcal{L} defining \mathcal{H} does not contain a facet in $H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ at \mathbf{p} , as in that case all facets of \mathcal{H} at \mathbf{p} have weight zero. The weight function of $(\mathcal{H}_{\mathcal{B}'}, w_{\mathcal{B}'})$ satisfies $w_{\mathcal{B}'} = \sum_{\mathbf{b} \in \mathcal{B}'} w_{\mathbf{b}}$, (see Section 3.1 and Appendix C) so we have $m_{\mathcal{B}'}^{ij} = \sum_{\mathbf{b} \in \mathcal{B}'} m_{\mathbf{b}}^{ij}$. Recall that $m^{0j}(\mathbf{p}; \cdot)$ has not been defined; with this exception, the order of i and j does not matter in m^{ij} , as Corollary G.12 implies that:

OBSERVATION G.14: For any $\mathbf{p} \in [\mathcal{P}]$ and any distinct $i, j \in I(\mathbf{p}) \setminus \{0\}$, we have $m^{ij}(\mathbf{p}; \cdot) = m^{ji}(\mathbf{p}; \cdot)$.

G.3. Showing correctness

In Appendix G.3.1, we prove $(\mathcal{H}_v, w_v) = (\mathcal{H}_{\mathcal{B}}, w_{\mathcal{B}})$ for the bid collection \mathcal{B} constructed in Appendix G.2. By Proposition G.2, this implies that the difference in demand as we move from one price to another is the same for D_v and $D_{\mathcal{B}}$. Appendix G.3.2 then concludes the proof of Theorem 3.1 by showing that $D_v = D_{\mathcal{B}}$, and arguing that \mathcal{B} is the only bid collection (up to normalisation) satisfying this equality. In order to prove these results, we first establish two key properties of the multiplicity functions m_v^{ij} and $m_{\mathcal{B}}^{ij}$. Lemma G.15 characterises $m_{\mathcal{B}}^{ij}$, and Lemma G.16 shows that the two multiplicity functions coincide.

LEMMA G.15: Fix prices $\mathbf{p} \in [\mathcal{P}]$ and bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m) \in \mathcal{B}$.

- i) For any orientation $(i, 0)$ at \mathbf{p} , we have $m_{\mathbf{b}}^{i0}(\mathbf{p}) = m$ if $\mathbf{r} = \tau(\mathbf{p})$ and $m_{\mathbf{b}}^{i0}(\mathbf{p}) = 0$ otherwise.
- ii) For any orientation (i, j) with $0 \notin \{i, j\}$ at \mathbf{p} and slope vector \mathbf{s} for (i, j) at \mathbf{p} , we have $m_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{s}) = m$ if $\mathbf{r} = \tau(\mathbf{p})$ and $\mathbf{s} = \mathbf{t}$, and $m_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{s}) = 0$ otherwise.

LEMMA G.16: For any $\mathbf{p} \in [\mathcal{P}]$ with $|I(\mathbf{p})| \geq 2$, orientation (i, j) at \mathbf{p} with $i > 0$, and slope vector \mathbf{s} for (i, j) at \mathbf{p} , we have $m_v^{ij}(\mathbf{p}; \mathbf{s}) = m_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s})$.

In order to prove Lemmas G.15 and G.16, we first develop properties of the multiplicity function in technical Lemmas G.17, G.19 and G.20. Throughout this section, let (\mathcal{H}, w) be a balanced weighted HIP whose hyperplanes have slopes in $\mathbb{Q}_{\mathcal{S}}$. Note that Observation G.8 implies that the multiplicity function $m^{ij}(\mathbf{p}; \mathbf{s})$ defined in Definition G.13 can be written, for any \mathbf{p} at which $I(\mathbf{p})$ contains at least three distinct goods and any $k \in I(\mathbf{p}) \setminus \{i, j\}$, as

$$m^{ij}(\mathbf{p}; \mathbf{s}) = \frac{1}{\gcd(s_i, s_j)} \sum_{J \subseteq I(\mathbf{p}) \setminus \{i, j, k\}} (-1)^{|J|} [w^{ij}(\mathbf{p}; \mathbf{s}; J) - w^{ij}(\mathbf{p}; \mathbf{s}; J \cup \{k\})]. \quad (\text{G.6})$$

LEMMA G.17: Fix a point $\mathbf{p} \in \mathcal{P}$ and two orientations (i, j) and (k, l) with $0 \notin \{i, j, k, l\}$. For any slope vector \mathbf{s} for (i, j) at \mathbf{p} , we have $m^{ij}(\mathbf{p}; \mathbf{s}) = m^{kl}(\mathbf{p}; \mathbf{s})$.

PROOF: The claim is immediate if $I(\mathbf{p}) = \{i, j\}$, so suppose $I(\mathbf{p})$ contains three distinct non-zero goods i, j and k . We prove $m^{ij}(\mathbf{p}; \mathbf{s}) = m^{kj}(\mathbf{p}; \mathbf{s})$, which implies the lemma by Observation G.14. Define $H^{ij} := H(\mathbf{p}; i, j; \frac{s_i}{s_j})$, $H^{kj} := H(\mathbf{p}; k, j; \frac{s_k}{s_j})$ and $H^{ik} := H(\mathbf{p}; i, k; \frac{s_i}{s_k})$, and assume that these hyperplanes are in \mathcal{H} by adding dummy hyperplanes with zero-weighted facets if necessary. Using Equation (G.6), we prove $m^{ij}(\mathbf{p}; \mathbf{s}) = m^{kj}(\mathbf{p}; \mathbf{s})$ by showing, for every set $J \subseteq I \setminus \{i, j, k\}$, that

$$\frac{w^{ij}(\mathbf{p}; \mathbf{s}; J) - w^{ij}(\mathbf{p}; \mathbf{s}; J \cup \{k\})}{\gcd(s_i, s_j)} = \frac{w^{kj}(\mathbf{p}; \mathbf{s}; J) - w^{kj}(\mathbf{p}; \mathbf{s}; J \cup \{i\})}{\gcd(s_k, s_j)}. \quad (\text{G.7})$$

Fix $J \subseteq I \setminus \{i, j, k\}$. By Observation G.6, we have $Q^{ij}(\mathbf{p}; \mathbf{s}; J) \cap H^{ij} = H^{ij}$ and $Q^{kj}(\mathbf{p}; \mathbf{s}; J) \cap H^{kj} = H^{kj}$. Writing out sliver regions $R^{ij}(\mathbf{p}; \mathbf{s}; J)$ and $R^{kj}(\mathbf{p}; \mathbf{s}; J)$ according to Definition G.5, we see that $R^{ij}(\mathbf{p}; \mathbf{s}; J) \cap H^{kj} = R^{kj}(\mathbf{p}; \mathbf{s}; J) \cap H^{ij}$. Let $R := R^{ij}(\mathbf{p}; \mathbf{s}; J) \cap H^{kj} = R^{kj}(\mathbf{p}; \mathbf{s}; J) \cap H^{ij}$. Lemma G.10 implies that R is $(n-2)$ -dimensional and contained in $H^{ij} \cap H^{kj}$. Let $G \subseteq H^{ij} \cap H^{kj}$ be a face of \mathcal{H} that has $(n-2)$ -dimensional intersection with R . By Lemma G.4 (part 2), all facets F of \mathcal{H} that contain G lie in H^{ij} , H^{ik} or H^{kj} . Each of these hyperplanes contains exactly two such facets, one on either side of G . By Observation G.6, again, $Q^{kj}(\mathbf{p}; \mathbf{s}; J) \cap H^{kj} = Q^{kj}(\mathbf{p}; \mathbf{s}; J \cup \{k\}) \cap H^{kj}$; and $Q^{lj}(\mathbf{p}; \mathbf{s}; J) = Q^{lj}(\mathbf{p}; \mathbf{s}; J \cup \{k\})$ for all $l \in [n]_0 \setminus \{i, j, k\}$; so $R^{ij}(\mathbf{p}; \mathbf{s}; J \cup \{k\}) \cap H^{kj} = R^{ij}(\mathbf{p}; \mathbf{s}; J) \cap H^{kj} = R$. Hence, one of the facets in H^{ij} containing G has $(n-1)$ -dimensional intersection with $R^{ij}(\mathbf{p}; \mathbf{s}; J)$ and so weight $w^{ij}(\mathbf{p}; \mathbf{s}; J)$, while the other facet has $(n-1)$ -dimensional intersection with $R^{ij}(\mathbf{p}; \mathbf{s}; J \cup \{k\})$ and so weight $w^{ij}(\mathbf{p}; \mathbf{s}; J \cup \{k\})$. Analogous statements hold for H^{kj} . Moreover, let w_1 and w_2 denote the weights of the facets in H^{ik} on either side of G . Choosing a consistent orientation around G , the balancing property of (\mathcal{H}, w) implies the following vector equation whose j th component is Equation (G.7).

$$\begin{aligned} \mathbf{0} &= \frac{w^{ij}(\mathbf{p}; \mathbf{s}; J) - w^{ij}(\mathbf{p}; \mathbf{s}; J \cup \{k\})}{\gcd(s_i, s_j)} (s_i \mathbf{e}^i - s_j \mathbf{e}^j) \\ &+ \frac{w^{kj}(\mathbf{p}; \mathbf{s}; J) - w^{kj}(\mathbf{p}; \mathbf{s}; J \cup \{i\})}{\gcd(s_k, s_j)} (s_j \mathbf{e}^j - s_k \mathbf{e}^k) \\ &+ \frac{w_1 - w_2}{\gcd(s_i, s_k)} (s_k \mathbf{e}^k - s_i \mathbf{e}^i). \end{aligned}$$

Q.E.D.

COROLLARY G.18: *For any prices $\mathbf{p} \in [\mathcal{P}]$, orientation (i, j) at \mathbf{p} with $i > 0$, and slope vector \mathbf{s} for (i, j) at \mathbf{p} , we have $m^{ij}(\mathbf{p}; \mathbf{s}) \in \mathbb{Z}$.*

PROOF: If $j = 0$, then $\gcd(s_i, s_j) = 1$ (as $s_j = s_0 = 1$ by definition), so integrality of $m^{ij}(\mathbf{p}; \mathbf{s})$ follows immediately from Definition G.13 and the integrality of facet weights. So suppose $i, j > 0$. Fix \mathbf{p} and \mathbf{s} , apply Lemma G.17 to see that $m^{ij}(\mathbf{p}; \mathbf{s}) = m^{ik}(\mathbf{p}; \mathbf{s})$ for all $k \in I(\mathbf{p}) \setminus \{0, i\}$, and write $m^{ij}(\mathbf{p}; \mathbf{s}) = \frac{a}{b}$ with co-prime $a, b \in \mathbb{Z}$. For any $k \in I(\mathbf{p}) \setminus \{0, i\}$, we observe by Definition G.13 that $\gcd(s_i, s_k) m^{ik}(\mathbf{p}; \mathbf{s}) = \gcd(s_i, s_k) \frac{a}{b}$ is integral. So $b \mid \gcd(s_i, s_k)$ for all $k \in I(\mathbf{p}) \setminus \{0, i\}$. It follows that $b \mid s_k$ for all $k \in I(\mathbf{p}) \setminus \{0\}$. But as $s_l = 0$ for all $l \in [n] \setminus I(\mathbf{p})$ and $\mathbf{s} \in \mathcal{T}$, so the non-zero entries of \mathbf{s} indexed by $[n]$ are co-prime, we have $b = 1$ and so $m^{ij}(\mathbf{p}; \mathbf{s}) = a \in \mathbb{Z}$. *Q.E.D.*

LEMMA G.19: *Fix a point $\mathbf{p} \in \mathcal{P}$ with $0 \in I(\mathbf{p})$. For any distinct $i, j \in I(\mathbf{p}) \setminus \{0\}$, we have $m^{i0}(\mathbf{p}) = \sum_{\mathbf{s} \in \mathcal{S}} s_i m^{ij}(\mathbf{p}; \mathbf{s})$, where \mathcal{S} is the set of all slope vectors for (i, j) at \mathbf{p} .*

PROOF: The claim is vacuous if $|I(\mathbf{p})| = 2$, so suppose $I(\mathbf{p}) \setminus \{0\}$ contains distinct i, j . For any slope $\sigma \in \mathbb{Q}_S$ and set $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$, define the *extremal slope vector* $\mathbf{s}(\sigma, J)$ for (i, j) at \mathbf{p} as the vector $\mathbf{s} \in \mathcal{T}$ satisfying, for every $k \in I(\mathbf{p}) \setminus \{0, j\}$,

$$\frac{s_k}{s_j} = \begin{cases} \sigma & \text{if } k = i, \\ S^{\times J}_{\sigma^k} & \text{otherwise.} \end{cases}$$

Note that $\mathbf{s}(\sigma, J)$ is a primitive vector, but coordinates $s_i(\sigma, J)$ and $s_j(\sigma, J)$ need not be coprime. Clearly, any extremal sliver region is identified by an extremal slope vector (cf. Corollary G.9).

First we show

$$\sum_{\substack{\mathbf{s} \in \mathcal{S} \\ J \subseteq I(\mathbf{p}) \setminus \{i, j\}}} \frac{s_i}{\gcd(s_i, s_j)} (-1)^{|J|} w^{ij}(\mathbf{p}; \mathbf{s}; J) = \sum_{\substack{\sigma \in \mathbb{Q}_S \\ J \subseteq I(\mathbf{p}) \setminus \{i, j\}}} \frac{s_i(\sigma, J)}{\gcd(s_i(\sigma, J), s_j(\sigma, J))} (-1)^{|J|} w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J); J). \quad (\text{G.8})$$

We first rearrange the sum on the left-hand side of Equation (G.8) by grouping together all terms with pairs (\mathbf{s}, J) identifying the same sliver region. By Corollary G.9, each such sliver region is either extremal and thus associated with a single identifying set (\mathbf{s}, J) or it is identified by the same number of pairs (\mathbf{s}, J) with even and odd cardinality of set J . Moreover, if (\mathbf{s}, J) and (\mathbf{s}', J') define the same sliver region then $\frac{s'_j}{s'_i} = \frac{s_j}{s_i}$ (by Definition G.5) and so $\frac{s'_i}{\gcd(s'_i, s'_j)} = \frac{s_i}{\gcd(s_i, s_j)}$. It follows that the terms corresponding to non-extremal regions cancel out, and we are left with the weighted sum of the weights of all extremal sliver regions around \mathbf{p} in H , the right-hand side of Equation (G.8).

We now prove, for any $J \subseteq I(\mathbf{p}) \setminus \{i, j, 0\}$, that

$$\begin{aligned} & w^{i0}(\mathbf{p}; J) - w^{i0}(\mathbf{p}; J \cup \{j\}) \\ &= \sum_{\sigma \in \mathbb{Q}_S} \frac{s_i(\sigma, J)}{\gcd(s_i(\sigma, J), s_j(\sigma, J))} [w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J); J) - w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J \cup \{0\}); J \cup \{0\})]. \end{aligned} \quad (\text{G.9})$$

To show this, we proceed similarly to the proof of Lemma G.17. Fix $J \subseteq I(\mathbf{p}) \setminus \{i, j, 0\}$ and define the hyperplanes $H^{i0} := H(\mathbf{p}; i, 0; 1)$ and $H^{j0} := H(\mathbf{p}; j, 0; 1)$. We can assume that these hyperplanes exist in \mathcal{H} by adding dummy hyperplanes with zero-weighted facets, if necessary. From Definition G.5, it follows that $R := R^{i0}(\mathbf{p}; J) \cap H^{j0} = R^{j0}(\mathbf{p}; J) \cap H^{i0} \subseteq H^{i0} \cap H^{j0}$, and R is $(n-2)$ -dimensional. Let G be an $(n-2)$ -face of \mathcal{H} that has $(n-2)$ -dimensional intersection with R . By Lemma G.4, a facet of \mathcal{H} contains G only if it lies in H^{i0} , H^{j0} or in a hyperplane $H(\mathbf{p}; i, j; \sigma)$ for some $\sigma \in \mathbb{Q}_S$. Each such hyperplane contains two facets with G as a bounding face, one on either side. $R^{i0}(\mathbf{p}; J)$ and $R^{i0}(\mathbf{p}; J \cup \{j\})$ both contain R and lie on either side of G , as we can see using the same methods as in the proof of Lemma G.17. So they each have $(n-1)$ -dimensional intersection with one of these two facets in H^{i0} . Hence the facets have weights $w^{i0}(\mathbf{p}; J)$ and $w^{i0}(\mathbf{p}; J \cup \{j\})$. An analogous statement holds for H^{j0} .

We now turn to the hyperplanes with orientation (i, j) containing G . Fix slope $\sigma \in \mathbb{Q}_S$ and let $H^{ij} := H(\mathbf{p}; i, j; \sigma)$. As above, we assume that $H^{ij} \subseteq \mathcal{H}$. Recall that $H^{ij} \cap H^{i0} = H^{ij} \cap H^{j0} = H^{i0} \cap H^{j0}$. We prove that

$$R \subseteq R^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J); J) \cap H^{i0} = R^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J \cup \{0\}); J \cup \{0\}) \cap H^{i0}. \quad (\text{G.10})$$

To see the set inclusion in Equation (G.10), fix some $\mathbf{q} \in \text{relint } R$. For any $k \in I(\mathbf{p}) \setminus \{i, j, 0\}$,

$$\begin{aligned} \mathbf{q} \in Q^{k0}(\mathbf{p}; \mathbf{1}; J) \cap H^{i0} \cap H^{j0} &= H^{-x_k^J}(\mathbf{p}; k, 0; 1) \cap H^{i0} \cap H^{j0} \\ &= H^{-x_k^J}(\mathbf{p}; k, j; S^{x_{0k}^J}) \cap H^{i0} \cap H^{ij} \\ &= Q^{kj}(\mathbf{p}; \mathbf{s}(\sigma, J); J) \cap H^{i0} \cap H^{ij}. \end{aligned}$$

Similarly, $\mathbf{q} \in Q^{k0}(\mathbf{p}; \mathbf{1}; J) \cap H^{i0} \cap H^{j0} = Q^{kj}(\mathbf{p}; \mathbf{s}(\sigma, J); J) \cap H^{i0} \cap H^{ij}$ for any $k \notin I(\mathbf{p}) \cup \{0\}$. Finally, $\mathbf{q} \in H^{j0} \cap H^{i0} = Q^{j0}(\mathbf{p}; \mathbf{s}(\sigma, J); J) \cap H^{ij} \cap H^{i0}$. Observe that $\mathbf{q} \in [\mathcal{P}] \cap N_\varepsilon(\mathbf{p})$ because $\mathbf{q} \in R$. Together, this implies $\mathbf{q} \in R^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J); J) \cap H^{i0}$.

To show the equality in Equation (G.10), first note that, by definition of \mathbf{s} , for $k \in [n]_0 \setminus \{i, j\}$ and any $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$, the set $Q^{kj}(\mathbf{p}; \mathbf{s}(\sigma, J'); J')$ is a single half-space of orientation (k, j) . In the definition of R^{ij} , we intersect the Q^{kj} sets with H^{ij} , and, in Equation (G.10), additionally with H^{i0} , but $H^{ij} \cap H^{i0} = H^{j0} \cap H^{i0}$. Observe that, for any $\chi \in \{0, 1\}$, any $k \in I \setminus \{i, j, 0\}$ and any slope $\frac{a}{b}$, we have $H^\chi(\mathbf{p}; k, j; \frac{a}{b}) \cap H^{j0} = H^\chi(\mathbf{p}; k, 0; 1) \cap H^{j0}$, which is independent of the slope, while $H^\chi(\mathbf{p}; 0, j; \frac{a}{b}) \cap H^{j0} = H^{j0}$, which is independent of χ . We can conclude that, for every $k \in [n]_0 \setminus \{i, j\}$, we have $Q^{kj}(\mathbf{p}; \mathbf{s}(\sigma, J); J) \cap H^{i0} = Q^{kj}(\mathbf{p}; \mathbf{s}(\sigma, J \cup \{0\}); J \cup \{0\}) \cap H^{i0}$. The infinitesimal neighbourhood of \mathbf{p} is the same in each case. So we have demonstrated the equality in Equation (G.10).

Hence, the weights of the facets in $H^{ij} = H(\mathbf{p}; i, j; \sigma)$ on either side of G are $w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J); J)$ and $w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J \cup \{0\}); J \cup \{0\})$. Note that $\frac{1}{\gcd(s_i(\sigma, J), s_j(\sigma, J))}(s_i(\sigma, J)e^i - s_j(\sigma, J)e^j)$ is a primitive vector normal to $H(\mathbf{p}; i, j; \sigma)$. So, choosing a consistent orientation around G , the balancing property of (\mathcal{H}, w) thus implies the following vector equation whose i th component is Equation (G.9).

$$\begin{aligned} 0 &= [w^{i0}(\mathbf{p}; J) - w^{i0}(\mathbf{p}; J \cup \{j\})]e^i - [w^{j0}(\mathbf{p}; J) - w^{j0}(\mathbf{p}; J \cup \{i\})]e^j \\ &\quad - \sum_{\sigma \in \mathbb{Q}_S} \frac{w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J); J) - w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J \cup \{0\}); J \cup \{0\})}{\gcd(s_i(\sigma, J), s_j(\sigma, J))} (s_i(\sigma, J)e^i - s_j(\sigma, J)e^j). \end{aligned}$$

Finally, by combining Equation (G.6), Equation (G.9) Equation (G.8) and Equation (G.6) again in turn, we see that

$$\begin{aligned} m^{i0}(\mathbf{p}) &= \sum_{J \subseteq I(\mathbf{p}) \setminus \{i, j, 0\}} (-1)^{|J|} [w^{i0}(\mathbf{p}; J) - w^{i0}(\mathbf{p}; J \cup \{j\})] \\ &= \sum_{\substack{J \subseteq I(\mathbf{p}) \setminus \{i, j, 0\} \\ \sigma \in \mathbb{Q}_S}} \frac{s_i(\sigma, J)}{\gcd(s_i(\sigma, J), s_j(\sigma, J))} (-1)^{|J|} [w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J); J) - w^{ij}(\mathbf{p}; \mathbf{s}(\sigma, J \cup \{0\}); J \cup \{0\})] \\ &= \sum_{\substack{J \subseteq I(\mathbf{p}) \setminus \{i, j, 0\} \\ \mathbf{s} \in S}} \frac{s_i}{\gcd(s_i, s_j)} (-1)^{|J|} [w^{ij}(\mathbf{p}; \mathbf{s}; J) - w^{ij}(\mathbf{p}; \mathbf{s}; J \cup \{0\})] \\ &= \sum_{\mathbf{s} \in S} s_i m^{ij}(\mathbf{p}; \mathbf{s}), \end{aligned}$$

as required. Q.E.D.

LEMMA G.20: Fix point $\mathbf{p} \in [\mathcal{P}]$, orientation (i, j) at \mathbf{p} with $i > 0$, and slope vector \mathbf{s} for (i, j) at \mathbf{p} . If $m^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$, then $[\mathcal{L}_v]$ has a facet in $H(\mathbf{p}; k, l; \frac{s_i}{s_j})$ with a vertex at \mathbf{p} , for all $k, l \in I(\mathbf{p})$.

PROOF: Assume $|I(\mathbf{p})| \geq 2$, as otherwise the statement is vacuous. In order to show that \mathbf{p} is a vertex of $[\mathcal{L}_v]$, we show that it is the unique point of intersection of facets of \mathcal{L}_v and the bounding box. By definition of $I(\mathbf{p})$, \mathbf{p} lies in the lower boundary hyperplane $e^k \cdot \mathbf{q} = \underline{C}_k$ for every $k \in [n] \setminus I(\mathbf{p})$.

If $I(\mathbf{p}) = \{i, 0\}$ then $j = 0$ and $m^{i0}(\mathbf{p}) \neq 0$ implies that \mathcal{L}_v contains a $(i, 0)$ -facet F of non-zero weight meeting \mathbf{p} . This facet F and the $n - 1$ boundary hyperplanes containing \mathbf{p} (the hyperplanes $e^k \cdot \mathbf{q} = \underline{C}_k$ for all $k \in [n] \setminus \{i\}$) intersect at \mathbf{p} , as required.

So suppose $I(\mathbf{p}) \neq \{i, 0\}$. Then if $j = 0$, there exists $k \in I(\mathbf{p}) \setminus \{i, 0\}$, and by Lemma G.19, $m^{i0}(\mathbf{p}) \neq 0$ implies there exists $\mathbf{s} \in \mathcal{T}$ such that $m^{ik}(\mathbf{p}, \mathbf{s}) \neq 0$. So we may assume that $j > 0$. Now, by Lemma G.17, we have $m^{kl}(\mathbf{p}; \mathbf{s}) = m^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ for all distinct $k, l \in I(\mathbf{p}) \setminus \{i, j, 0\}$. It follows that $H(\mathbf{p}; k, l; \frac{s_k}{s_l})$ contains a facet of \mathcal{L}_v meeting \mathbf{p} . Set up a system of equations with $(s_k e^k - s_j e^j) \cdot (\mathbf{q} - \mathbf{p}) = 0$ for each hyperplane $H(\mathbf{p}; k, j; \frac{s_k}{s_j})$ with $k \in I(\mathbf{p}) \setminus \{j, 0\}$, and $e^k \cdot \mathbf{q} = \underline{C}_k$ for each boundary hyperplane with $k \in [n] \setminus I(\mathbf{p})$. It is straightforward that this system has $n - 1$ equations and rank $n - 1$.

We now argue that $[\mathcal{L}_v]$ has an additional facet or boundary hyperplane meeting \mathbf{p} such that adding its equation to our system increases the rank to n . This will immediately imply that \mathbf{p} is the unique intersection of all these facets and boundary hyperplanes, and so is a vertex of $[\mathcal{L}_v]$, as required. If $0 \notin I(\mathbf{p})$, then there exists some $k \in I(\mathbf{p}) \setminus \{0\}$ such that \mathbf{p} lies in the upper boundary hyperplane associated with equation $e^k \cdot \mathbf{q} = \overline{C}_k$, and we are done. So suppose $0 \in I(\mathbf{p})$. If \mathcal{L}_v contains a (k, j) -facet with slope $\sigma \neq \frac{s_k}{s_j}$ meeting \mathbf{p} , for some $k \in I(\mathbf{p}) \setminus \{j\}$, then we are done. Hence, suppose this is not the case. Lemma G.17 implies that \mathbf{s} is the only slope vector \mathbf{s}' for (i, j) at \mathbf{p} with $m^{ij}(\mathbf{p}; \mathbf{s}') \neq 0$. Then $m^{i0}(\mathbf{p}) = s_i m^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ by Lemma G.19 and so \mathcal{L}_v contains an $(i, 0)$ -facet meeting \mathbf{p} . Q.E.D.

We are now ready to prove Lemmas G.15 and G.16.

PROOF OF LEMMA G.15: We will argue parts (i) and (ii) together by showing that the statement must hold for any orientation (i, j) at \mathbf{p} . Assume that $|I(\mathbf{p})| \geq 2$, as otherwise the statement is trivial, and let $I = \{i \in [n]_0 \mid r_i > -\infty\}$ be the goods in which \mathbf{b} is interested. Recall that by construction of \mathcal{B} (Algorithm 1) we know that $\mathbf{r} = \tau(\mathbf{p}')$ for some $\mathbf{p}' \in [\mathcal{P}]$, and so $r_i > \underline{C}_i$ for all $i \in I$; and $r_i \leq \overline{C}_i$ for all $i \in [n]$; and $r_0 = -\infty$ if and only if there exists $i \in [n]$ with $r_i = \overline{C}_i$; otherwise $r_0 = 0$.

By Definition G.13, we know that $m_b^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ implies that $H(\mathbf{p}; i, j; \frac{s_i}{s_j})$ contains a non-0-weighted facet, and so is contained in \mathcal{H}_b , for any $i, j \in I(\mathbf{p})$. Moreover, the description of \mathcal{H}_b of Equation (G.2) then implies that $i, j \in I$ and $H(\mathbf{p}; i, j; \frac{s_i}{s_j}) = H(\mathbf{r}; i, j; \frac{t_i}{t_j})$, and therefore that $\frac{s_i}{s_j} = \frac{t_i}{t_j}$ and $t_i(r_i - p_i) = t_j(r_j - p_j)$.

We will now assume that $m_b^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ for some $i, j \in I(\mathbf{p})$ and some slope vector \mathbf{s} for (i, j) at \mathbf{p} , and show that this implies that for all $k \in I(\mathbf{p}) \setminus \{j\}$ there exists \mathbf{s}' such that either $m_b^{kj}(\mathbf{p}; \mathbf{s}') \neq 0$ or $m_b^{ik}(\mathbf{p}; \mathbf{s}') \neq 0$, with \mathbf{s}' being a slope vector for (k, j) or (i, k) as appropriate. As just described, this will tell us that $I(\mathbf{p}) \subseteq I$ and

$$t_k(r_k - p_k) = t_j(r_j - p_j) \text{ for all } k \in I(\mathbf{p}). \quad (\text{G.11})$$

The claim is immediate if $|I(\mathbf{p})| = 2$, so suppose $|I(\mathbf{p})| > 2$. If $j = 0$, then by Lemma G.19, for any $k \in I(\mathbf{p}) \setminus \{i, 0\} \neq \emptyset$ there exists $\mathbf{s}' \in \mathcal{S}$ such that $m^{ik}(\mathbf{p}; \mathbf{s}') \neq 0$. So without loss of generality we may assume $j > 0$. Then, for all $k \in I(\mathbf{p}) \setminus \{0, j\}$, we know $m_b^{kj}(\mathbf{p}; \mathbf{s}) \neq 0$ (by Lemma G.17) and so $\frac{s_k}{s_j} = \frac{t_k}{t_j}$ (by the preceding paragraph). This uniquely defines $\mathbf{s} \in \mathcal{S}$ given $I(\mathbf{p})$, so Lemma G.19 also implies $m_b^{i0}(\mathbf{p}) \neq 0$ if $0 \in I(\mathbf{p})$.

We now show $I(\mathbf{p}) = I$. If $0 \in I \setminus I(\mathbf{p})$, then $r_0 = 0$ and there exists $k \in I(\mathbf{p})$ with $p_k = \overline{C}_k > r_k$, and so $t_k(r_k - p_k) < 0$. By Equation (G.11), we have $t_l(r_l - p_l) = t_k(r_k - p_k) < 0 = t_0(r_0 - p_0)$ for all $l \in I(\mathbf{p})$, and so \mathbf{b} strictly prefers the null good to goods i and j at \mathbf{p} . Now suppose $0 \neq l \in I \setminus I(\mathbf{p})$, which means that $r_l > \underline{C}_l = p_l$. Then $t_l(r_l - p_l) > 0$. However, if $0 \notin I(\mathbf{p})$, then $p_k = \overline{C}_k \geq r_k$ for some $k \in I(\mathbf{p})$; with Equation (G.11) this implies $t_k(r_k - p_k) \leq 0$ for all $k \in I(\mathbf{p})$. And if $0 \in I(\mathbf{p}) \subseteq I$, then Equation (G.11) implies that $t_k(r_k - p_k) = t_0(r_0 - p_0) = 0$

for all $k \in I(\mathbf{p})$. In either case, $t_l(r_l - p_l) > 0 \geq t_k(r_k - p_k)$ for all $k \in I(\mathbf{p})$, and so bid \mathbf{b} strictly prefers good l to either i or j at price \mathbf{p} . Each of these cases implies that $w^{ij}(\mathbf{p}; J) = 0$ for every $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$ and so that $m_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{s}) = 0$. This contradiction implies that $I(\mathbf{p}) = I$.

Now we show that $\tau(\mathbf{p}) = \mathbf{r}$. If $0 \in I = I(\mathbf{p})$ then Equation (G.11) directly implies $p_k = r_k$ for all $k \in I$. If $0 \notin I = I(\mathbf{p})$, then there exists $k \in I(\mathbf{p})$ with $p_k = \overline{C_k} \geq r_k$, and so $p_l \geq r_l$ for all $l \in I$ by Equation (G.11); and there exists $k' \in I$ with $r_{k'} = \overline{C_{k'}} \geq p_{k'}$, and so $r_l \geq p_l$ for all $l \in I$ by Equation (G.11); so again $p_l = r_l$ for all $l \in I$. And if $k \in [n]_0 \setminus I = [n]_0 \setminus I(\mathbf{p})$ then $r_k = -\infty$ and $p_k = \overline{C_k}$. So in both cases we have $\mathbf{r} = \tau(\mathbf{p})$.

We also have $\mathbf{s} = \mathbf{t}$, as $\frac{s_k}{s_j} = \frac{t_k}{t_j}$ for all $k \in I = I(\mathbf{p})$ and \mathbf{s}, \mathbf{t} are both primitive vectors.

By Equation (G.1) and Observation G.7, the sliver region $R^{ij}(\mathbf{p}; \mathbf{s}; J)$ is contained in facet $F_{\mathbf{b}}^{ij}$ of $\mathcal{H}_{\mathbf{b}}$ if and only if $J = \emptyset$. So $w_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{s}; J) = m \gcd(s_i, s_j)$ if $J = \emptyset$, and is 0 otherwise. Hence $m_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{s}) = m$. Q.E.D.

PROOF OF LEMMA G.16: Suppose first that $j > 0$ or $I(\mathbf{p}) = \{i, 0\}$ (so $j = 0$). If \mathbf{p} is not a vertex of $[\mathcal{L}_v]$, then the bid collection \mathcal{B} contains no bids with root $\mathbf{r} = \tau(\mathbf{p})$, and so $m_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{s}) = 0$ for all $\mathbf{b} \in \mathcal{B}$ by Lemma G.15. As $m_{\mathcal{B}}^{ij}(\mathbf{p}, \mathbf{s}) = \sum_{\mathbf{b} \in \mathcal{B}} m_{\mathbf{b}}^{ij}(\mathbf{p}, \mathbf{s})$ (as stated above Observation G.14), we have $m_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s}) = 0$. Lemma G.20 implies $m_v^{ij}(\mathbf{p}; \mathbf{s}) = 0$, so the claim follows.

Now suppose \mathbf{p} is a vertex of $[\mathcal{L}_v]$. Then, for every slope vector \mathbf{s} at \mathbf{p} , Algorithm 1 stipulates that either \mathcal{B} contains a single bid with root $\tau(\mathbf{p})$, tradeoffs \mathbf{s} and multiplicity $m = m_v^{kl}(\mathbf{p}; \mathbf{s})$, where (k, l) is the orientation with the two largest goods in $I(\mathbf{p})$; or \mathcal{B} contains no bid with root $\tau(\mathbf{p})$, and tradeoffs \mathbf{s} , if this associated multiplicity is zero. If $(k, l) \neq (i, j)$, we see that $k, l \geq j > 0$ because $j > 0$ by assumption, and so Lemma G.17 tells us that $m_v^{kl}(\mathbf{p}; \mathbf{s}) = m_v^{ij}(\mathbf{p}; \mathbf{s})$. Hence Lemma G.15 implies $m_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s}) = m = m_v^{ij}(\mathbf{p}; \mathbf{s})$. For all other bids $\mathbf{b}' \in \mathcal{B}$, we have $m_{\mathbf{b}'}^{ij}(\mathbf{p}; \mathbf{s}) = 0$ by Lemma G.15, so $m_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s}) = \sum_{\mathbf{b} \in \mathcal{B}} m_{\mathbf{b}}^{ij}(\mathbf{p}; \mathbf{s}) = m_v^{ij}(\mathbf{p}; \mathbf{s})$.

We now turn to the case that $j = 0$ and $|I(\mathbf{p})| \geq 3$. Fix some $k \in I(\mathbf{p}) \setminus \{i, 0\}$. Writing \mathcal{S} for the set of slopes for (i, k) at \mathbf{p} , we see that $m_{\mathcal{B}}^{i0}(\mathbf{p}) = \sum_{\mathbf{s} \in \mathcal{S}} s_i m_{\mathcal{B}}^{ik}(\mathbf{p}; \mathbf{s}) = \sum_{\mathbf{s} \in \mathcal{S}} s_i m_v^{ik}(\mathbf{p}; \mathbf{s}) = m_v^{i0}(\mathbf{p})$. The first and third equalities follow from Lemma G.19, and the second equality holds due to $m_{\mathcal{B}}^{ik}(\mathbf{p}; \mathbf{s}) = m_v^{ik}(\mathbf{p}; \mathbf{s})$ for any $\mathbf{s} \in \mathcal{S}$ as shown above. Q.E.D.

G.3.1. Proving the equivalence of the two HIPs

We now show that $(\mathcal{H}_v, w_v) = (\mathcal{H}_{\mathcal{B}}, w_{\mathcal{B}})$. Proposition G.26 proves that $\mathcal{H}_v = \mathcal{H}_{\mathcal{B}}$, and Proposition G.28 then shows that $w_v(F) = w_{\mathcal{B}}(F)$ for every facet F of $\mathcal{H}_v = \mathcal{H}_{\mathcal{B}}$. In order to prove this, we first define partial orders for points within a hyperplane, and corresponding slope vectors, and develop two technical lemmas. Suppose (\mathcal{H}, w) is a parsimonious weighted HIP whose faces all intersect the interior of the bounding box $[\mathcal{P}]$.

DEFINITION G.21—Partial Orders: Fix a hyperplane H with orientation (i, j) and slope $\sigma \in \mathbb{Q}_{\mathcal{S}}$.

- (i) For any two points $\mathbf{p}, \mathbf{p}' \in H$, let $\mathbf{p} \preceq_H \mathbf{p}'$ if and only if either $p_j > p'_j$, or $p_j = p'_j$ and $p_k \leq p'_k$ for all $k \in [n]$.
- (ii) For any price \mathbf{p} , orientation (i, j) at \mathbf{p} with $0 \notin \{i, j\}$, and slope $\sigma \in \mathbb{Q}_{\mathcal{S}}$, we define a second partial order on the set \mathcal{S}_{σ} of slope vectors \mathbf{s} for (i, j) at \mathbf{p} with $\frac{s_i}{s_j} = \sigma$. For any two such slope vectors \mathbf{s} and \mathbf{s}' , let $\mathbf{s} \preceq_{\mathcal{S}, \sigma} \mathbf{s}'$ if and only if $\frac{s_k}{s_j} \leq \frac{s'_k}{s'_j}$ for all $k \in I(\mathbf{p})$.
- (iii) We also combine \preceq_H and $\preceq_{\mathcal{S}, \sigma}$ to formulate a partial order $\preceq_{H, \mathcal{S}}$ on all pairs $(\mathbf{p}, \mathbf{s}) \in H \times \mathcal{S}_{\sigma}$: let $(\mathbf{p}, \mathbf{s}) \preceq_{H, \mathcal{S}} (\mathbf{p}', \mathbf{s}')$ if and only if either: $\mathbf{p} \preceq_H \mathbf{p}'$ and $\mathbf{p} \neq \mathbf{p}'$; or $\mathbf{p} = \mathbf{p}'$ and $\mathbf{s} \preceq_{\mathcal{S}, \sigma} \mathbf{s}'$.

For each of these partial orders, we will use \prec to denote the associated strict partial order.

For \preceq_H , no specification need be given for coordinate i because, for any two points \mathbf{p}, \mathbf{p}' contained in H with orientation (i, j) , $p_j > p'_j$ implies $p_i > p'_i$ and $p_j = p'_j$ implies $p_i = p'_i$. It is straightforward to check that $\preceq_H, \preceq_{S, \sigma}$ and $\preceq_{H, S}$ are indeed partial orders, and that (H, \preceq_H) is a lattice. We write $\inf_H X$ for the infimum of set $X \subseteq H$ with respect to \preceq_H . Observe that, for any two $\mathbf{p}, \mathbf{p}' \in H$, we have $\inf_H \{\mathbf{p}, \mathbf{p}'\} = \mathbf{p}$ if $p_j > p'_j$, and $\inf_H \{\mathbf{p}, \mathbf{p}'\}$ is the usual Euclidean infimum if $p_j = p'_j$.

LEMMA G.22: *If F is a facet of \mathcal{H} contained in a hyperplane $H \subseteq \mathcal{H}$, then $\inf_H [F] \in [F]$.*

PROOF: By definition of $[\mathcal{P}]$ we know that $[F]$ is non-empty and $(n-1)$ -dimensional. As facets are closed, and $[\mathcal{P}]$ is bounded, we know $[F]$ is compact, and so it is sufficient to check that $([F], \preceq_H)$ is a lower semi-lattice (Milgrom and Shannon, 1994, Section 2). So, for any $\mathbf{p}, \mathbf{p}' \in [F]$, we need to establish that $\mathbf{q} := \inf_H (\mathbf{p}, \mathbf{p}') \in [F]$. Write (i, j) for the orientation of H . If $p_j \neq p'_j$, then as observed above $\mathbf{q} \in \{\mathbf{p}, \mathbf{p}'\} \subseteq [F]$. So assume that $p_j = p'_j$, which implies that $q_j = p_j = p'_j$, that $q_i = p_i = p'_i$ and that $q_k = \min(p_k, p'_k)$ for $k \in [n]$. Now, if $\mathbf{q} \notin [F]$ then there exists a hyperplane $\widehat{H} = H(\widehat{\mathbf{p}}; k, l; \sigma)$ of \mathcal{H} with $\mathbf{p}, \mathbf{p}' \in \widehat{H}^+ \setminus \widehat{H}$ and $\mathbf{q} \in \widehat{H}^- \setminus \widehat{H}$, so $\sigma p_k - p_l > \sigma \widehat{p}_k - \widehat{p}_l$ and $\sigma p'_k - p'_l > \sigma \widehat{p}_k - \widehat{p}_l > \sigma q_k - q_l$. Note that if $p'_k \leq p_k$ and $p'_l \leq p_l$ then $q_k = \min(p_k, p'_k) = p'_k$ and $q_l = \min(p_l, p'_l) = p'_l$, which contradicts our assumption on the separating hyperplane. So assume without loss of generality that $p_k > p'_k$ and $p_l < p'_l$. Then $q_k = p'_k$ and $q_l = p_l$; and so $\sigma q_k - q_l > \sigma p'_k - p'_l$, again contradicting our assumption on the separating hyperplane. This contradiction implies $\mathbf{q} \in [F]$. Q.E.D.

LEMMA G.23: *Let F be a facet of \mathcal{H} contained in hyperplane $H \subseteq \mathcal{H}$ of orientation (i, j) with $i > 0$ and slope $\frac{s_i}{s_j}$. Then $\mathbf{q} := \inf_H [F]$ is a vertex of $[\mathcal{H}]$ and we have $i, j \in I(\mathbf{q})$.*

Moreover, there exist slope vector \mathbf{s} for (i, j) at \mathbf{q} and set $J \subseteq I(\mathbf{q}) \setminus \{i, j\}$ such that $R^{ij}(\mathbf{q}; \mathbf{s}; J)$ has $(n-1)$ -dimensional intersection with $[F]$, and so $w(F) = w^{ij}(\mathbf{q}; \mathbf{s}; J)$.

PROOF: By definition of $[\mathcal{P}]$, for any facet F we know that $[F]$ is non-empty and has dimension $(n-1)$. By Lemma G.22 we know $\mathbf{q} \in [F]$. Suppose $\mathbf{q} = \lambda \mathbf{p} + (1-\lambda)\mathbf{p}'$ for $\mathbf{p}, \mathbf{p}' \in [F]$ and $\lambda \in (0, 1)$. As \mathbf{q} is the infimum of $[F]$, we have $q_j \geq p_j$ and $q_j \geq p'_j$. If either of these inequalities is strict, then $q_j > \lambda p_j + (1-\lambda)p'_j$, a contradiction. So $q_j = p_j = p'_j$, which implies (by definition of \preceq_H) that $q_k \leq p_k$ and $q_k \leq p'_k$ for all $k \in [n]$. Again, if any of these inequalities is strict, we obtain the contradiction $q_k < \lambda p_k + (1-\lambda)p'_k$. Hence $\mathbf{p} = \mathbf{p}' = \mathbf{q}$, implying that \mathbf{q} is a vertex of $[F]$.

Now pick a generic point $\mathbf{p} \in \text{relint}[F] \cap N_\varepsilon(\mathbf{q})$. Since \mathbf{p} does not lie in any face of $[F]$, and $\mathbf{q} = \inf_H [F]$, it follows when $j > 0$ that $q_j > p_j$ and $q_i > p_i$. When $j = 0$ we must have $p_j = q_j = 0$ and $p_i = q_i$, and so $\mathbf{p} \in \text{relint}[F]$ implies $q_k < p_k$ for all $k \in [n] \setminus \{i\}$.

As $[F]$ has $(n-1)$ -dimensional intersection with the interior of $[\mathcal{P}]$, we know $\underline{C}_k < p_k < \overline{C}_k$ for all $k \in [n]$. So in particular $q_i \geq p_i > \underline{C}_i$; and $q_j \geq p_j > \underline{C}_j$ if $j > 0$; so $i \in I(\mathbf{q})$ (in all cases) and $j \in I(\mathbf{q})$ if $j > 0$. If $j = 0$ then $\overline{q}_k \leq p_k < \overline{C}_k$ for all $k \in [n]$ implies $j = 0 \in I(\mathbf{q})$. So $i, j \in I(\mathbf{q})$.

We now prove the second part of the lemma by finding a slope vector \mathbf{s} for (i, j) at \mathbf{q} and a set $J \subseteq I(\mathbf{q}) \setminus \{i, j\}$ so that \mathbf{p} lies in the relative interior of $Q^{kj}(\mathbf{q}; \mathbf{s}; J)$ for every $k \in [n]_0 \setminus \{i, j\}$. It follows that an infinitesimal neighbourhood N of \mathbf{p} in $[F]$ satisfies $N \subseteq Q^{kj}(\mathbf{q}; \mathbf{s}; J)$ and so, by Definition G.5, $N \subseteq R^{ij}(\mathbf{q}; \mathbf{s}; J)$.

If $j = 0$, the slope vector \mathbf{s} for (i, j) at \mathbf{q} is uniquely defined and we let $J = \emptyset$. For every $k \in [n] \setminus \{i\}$, we know $p_k > q_k$, so \mathbf{p} lies in the interior of $H^+(\mathbf{q}; k, 0; 1) = Q^{k0}(\mathbf{q}; \mathbf{s}; \emptyset)$.

Now suppose $j > 0$. Firstly, $q_j > p_j$ implies that \mathbf{p} lies in the interior of $H^+(\mathbf{q}; 0, j; 1) = Q^{0j}(\mathbf{q}; \mathbf{s}; J)$. So we stipulate $0 \notin J$ (we need not specify whether $0 \in I(\mathbf{q})$). Now observe that

since $q_j - p_j > 0$, for any $k \in [n] \setminus \{i, j\}$ we have $\mathbf{p} \in H^+(\mathbf{q}; k, j; \sigma)$ if and only if $\frac{q_k - p_k}{q_j - p_j} \leq \frac{1}{\sigma}$, with \mathbf{p} lying in the relative interior if and only if the inequality is strict. As \mathbf{p} is generic we may assume that $\frac{q_k - p_k}{q_j - p_j} \notin \mathbb{Q}_S$. Consider $k \in I(\mathbf{q}) \setminus \{i, j, 0\}$: there are three cases to consider. If $\frac{q_k - p_k}{q_j - p_j} < \frac{1}{S}$ then \mathbf{p} is in the relative interior of $H^+(\mathbf{q}; k, j; S)$, which is $Q^{kj}(\mathbf{q}; \mathbf{s}; J)$ if we stipulate that $k \notin J$ and $\frac{s_k}{s_j} = S$. If $\frac{1}{\eta^+(\sigma)} < \frac{q_k - p_k}{q_j - p_j} < \frac{1}{\sigma}$ for some $\sigma \in \mathbb{Q}_S$ then \mathbf{p} is in the relative interior of $H^+(\mathbf{q}; k, j; \sigma) \cap H^-(\mathbf{q}; k, j; \eta(\sigma))$ which is $Q^{kj}(\mathbf{q}; \mathbf{s}; J)$ if we stipulate that $k \notin J$ and $\frac{s_k}{s_j} = \sigma$. Finally, if $\frac{q_k - p_k}{q_j - p_j} > S$, then \mathbf{p} is in the relative interior of $H^-(\mathbf{q}; k, j; \frac{1}{S})$, which is $Q^{kj}(\mathbf{q}; \mathbf{s}; J)$ if we stipulate that $k \in J$ and $\frac{s_k}{s_j} = \frac{1}{S}$. So we make the appropriate specifications, for all $k \in I(\mathbf{q}) \setminus \{i, j, 0\}$, and recalling that we already knew $\frac{s_i}{s_j}$, fully specifies \mathbf{s} and J . Finally, consider $0 \neq k \notin I(\mathbf{q})$. Now $q_k = \underline{C}_k < p_k$, so $q_k - p_k < 0$ and so $\frac{q_k - p_k}{q_j - p_j} < \frac{1}{S}$, implying \mathbf{p} is in the relative interior of $H^+(\mathbf{q}; k, j; S) = Q^{kj}(\mathbf{q}; \mathbf{s}; J)$.

We have thus demonstrated that $R^{ij}(\mathbf{q}; \mathbf{s}; J)$ has $(n-1)$ -dimensional intersection with $[F]$. That $w(F) = w^{ij}(\mathbf{q}; \mathbf{s}; J)$ follows immediately. Q.E.D.

It is useful to note at this point:

LEMMA G.24: *Suppose $\mathbf{p} \in [P]$, that (i, j) is an orientation at \mathbf{p} , that \mathbf{s} is a slope vector for (i, j) at \mathbf{p} and that $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$, and write $H := H(\mathbf{p}; i, j; \frac{s_i}{s_j})$. If $J \neq \emptyset$ then there exist some $\mathbf{p}' \in [P]$, some slope vector \mathbf{s}' for (i, j) at \mathbf{p}' satisfying $\frac{s'_i}{s'_j} = \frac{s_i}{s_j}$ and some $J' \subseteq I(\mathbf{p}') \setminus \{i, j\}$ such that $(\mathbf{p}', \mathbf{s}') \prec_{H, S} (\mathbf{p}, \mathbf{s})$ and $w^{ij}(\mathbf{p}'; \mathbf{s}'; J') = w^{ij}(\mathbf{p}; \mathbf{s}; J)$.*

PROOF: We first consider cases for which we show that there exists $\mathbf{q} \in R := R^{ij}(\mathbf{p}; \mathbf{s}; J)$ with $\mathbf{q} \prec_H \mathbf{p}$. Because R is $(n-1)$ -dimensional (Lemma G.10), we have $\mathbf{q} \in [F]$ for some facet $F \subseteq H$ such that F has $(n-1)$ -dimensional intersection with R , and so such that $w(F) = w^{ij}(\mathbf{p}; \mathbf{s}; J)$. Note that $\mathbf{p}' := \inf_H [F] \preceq_H \mathbf{q} \prec_H \mathbf{p}$. By Lemma G.23 there exist slope vector \mathbf{s}' for (i, j) at \mathbf{p}' and set $J' \subseteq I(\mathbf{p}') \setminus \{i, j\}$ such that $w^{ij}(\mathbf{p}'; \mathbf{s}'; J') = w(F) = w^{ij}(\mathbf{p}; \mathbf{s}; J)$. Moreover $\mathbf{p}' \prec_H \mathbf{p}$ implies $(\mathbf{p}', \mathbf{s}') \prec_{H, S} (\mathbf{p}, \mathbf{s})$.

Observe that if $j = 0$ and $k \in J$, then $\mathbf{q} := \mathbf{p} - \delta \mathbf{e}^k \in Q^{k0}(\mathbf{p}; \mathbf{s}; J) = H^-(\mathbf{p}; k, 0; 1)$ for any $\delta > 0$. Similarly, if $j > 0$ and $k \in J \setminus \{0\}$ and $\frac{s_k}{s_j} = \frac{1}{S}$, then $\mathbf{q} := \mathbf{p} - \delta \mathbf{e}^k \in Q^{kj}(\mathbf{p}; \mathbf{s}; J) = H^-(\mathbf{p}; k, j; \frac{1}{S})$ for $\delta > 0$. In both cases, it is easy to see that $\mathbf{q} \prec_H \mathbf{p}$ and that $\mathbf{q} \in Q^{lj}(\mathbf{p}; \mathbf{s}; J)$ for all other $l \in [n]_0$, and so that $\mathbf{q} \in R$ if δ is sufficiently small. Additionally, if $j > 0$ and $0 \in J$, then $Q^{0j}(\mathbf{p}; \mathbf{s}; J) = H^-(\mathbf{p}; 0, j; 1) = H^+(\mathbf{p}; j, 0, 1)$ and so $q_j > p_j$ for \mathbf{q} in the relative interior of R (again, R is $(n-1)$ -dimensional by Lemma G.10); thus again there exists $\mathbf{q} \in R$ with $\mathbf{q} \prec_H \mathbf{p}$. So the lemma is demonstrated for all these cases.

The remaining case is when $j > 0$ and $0 \notin J$ and, for every $k \in J \setminus \{0\}$ we have $\frac{s_k}{s_j} \neq \frac{1}{S}$, so $\frac{s_k}{s_j} > \frac{1}{S}$. But Observation G.8 tell us that $R^{ij}(\mathbf{p}; \mathbf{s}; J) = R^{ij}(\mathbf{p}; \mathbf{s}'; J \setminus \{k\})$, where $\frac{s'_i}{s'_j} = \eta^{-1}(\frac{s_k}{s_j})$ and $\frac{s'_l}{s'_j} = \frac{s_l}{s_j}$ for $l \in I(\mathbf{p}) \setminus \{j, k\}$, and so that $w^{ij}(\mathbf{p}; \mathbf{s}; J) = w^{ij}(\mathbf{p}; \mathbf{s}'; J \setminus \{k\})$. Note that $\mathbf{s}' \prec_{S, \frac{s_i}{s_j}} \mathbf{s}$, so $(\mathbf{p}, \mathbf{s}') \prec_{H, S} (\mathbf{p}, \mathbf{s})$. Q.E.D.

LEMMA G.25: *Let \mathcal{H} be parsimonious and let $H \subseteq \mathcal{H}$ be a hyperplane of orientation (i, j) with $i > 0$ and slope σ . There exist a point $\mathbf{p} \in [H]$ and slope vector \mathbf{s} for (i, j) at \mathbf{p} with $\frac{s_i}{s_j} = \sigma$ and $m^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$.*

PROOF: Since \mathcal{H} is parsimonious, there exist nonzero-weighted facets of $[\mathcal{H}]$ lying in H . By Lemma G.23, there exist prices $\mathbf{p} \in [H]$, slope vectors \mathbf{s} for (i, j) at \mathbf{p} with $\frac{s_i}{s_j} = \sigma$, and $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$, such that $w^{ij}(\mathbf{p}; \mathbf{s}; J) \neq 0$. Since facets are topologically closed and $[H]$ is bounded, there exists \mathbf{p} satisfying these properties which is minimal with respect to \preceq_H . Fix such \mathbf{p} ; there are finitely many slope vectors \mathbf{s} for (i, j) at \mathbf{p} , and we pick \mathbf{s} minimal with respect to $\preceq_{S, \sigma}$ such that $w^{ij}(\mathbf{p}; \mathbf{s}; J) \neq 0$ for some $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$. So (\mathbf{p}, \mathbf{s}) is minimal with respect to $\preceq_{H, S}$ such that $w^{ij}(\mathbf{p}; \mathbf{s}; J) \neq 0$ for some J , and thus by Lemma G.24 $w^{ij}(\mathbf{p}; \mathbf{s}; J) \neq 0$ if and only if $J = \emptyset$. This implies $m^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ by Definition G.13. Q.E.D.

PROPOSITION G.26: $\mathcal{H}_B = \mathcal{H}_v$.

PROOF: Fix a hyperplane H in either \mathcal{H}_B or \mathcal{H}_v , of orientation (i, j) with $i > 0$. As \mathcal{H}_v is parsimonious, Lemmas G.16 and G.25 tell us that $m_B^{ij}(\mathbf{p}; \mathbf{s}) = m_v^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ for some $\mathbf{p} \in H$ and slope vector \mathbf{s} for (i, j) at \mathbf{p} with $\frac{s_i}{s_j}$ being the slope of H . By construction of the multiplicity function, $m_B^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ implies that \mathcal{H}_B contains H and $m_v^{ij}(\mathbf{p}; \mathbf{s}) \neq 0$ implies that \mathcal{H}_v contains H . So H is in both \mathcal{H}_B and \mathcal{H}_v . As both sets are unions of hyperplanes, they are the same. Q.E.D.

We show in Proposition G.28 that $w_v = w_B$. The proof proceeds by induction on vertices and slopes according to the partial order on pairs (\mathbf{p}, \mathbf{s}) introduced in Definition G.21, and makes use of the following technical lemma.

LEMMA G.27: *For any hyperplane H of \mathcal{H} with orientation (i, j) and slope σ , the point $\mathbf{q} := \inf_H[H]$ is a vertex of $[\mathcal{H}]$, we have $I(\mathbf{q}) = \{i, j\}$, and there is a unique slope vector for (i, j) at \mathbf{p} with $\frac{s_i}{s_j} = \sigma$.*

PROOF: The hyperplane H is itself a trivial HIP \mathcal{H}' , with only one facet. So we know \mathbf{q} is a vertex of $[\mathcal{H}']$, and $i, j \in I(\mathbf{q})$, by Lemma G.23. Clearly $[\mathcal{H}'] \subseteq [\mathcal{H}]$ and so \mathbf{q} is a vertex of $[\mathcal{H}]$. For any $k \in [n] \setminus \{i, j\}$, $\mathbf{q} \in [H]$ implies $\mathbf{q}^k = \mathbf{q} + (\underline{C}_k - q_k)\mathbf{e}^k \in [H]$, but $\mathbf{q}^k \preceq_H \mathbf{q}$, so $\mathbf{q} = \mathbf{q}^k$, that is, $q_k = \underline{C}_k$ and $k \notin I(\mathbf{q})$. If $0 \notin \{i, j\}$ then $\mathbf{q} \in [H]$ implies $\mathbf{q}^\lambda = \mathbf{q} + \lambda(\mathbf{b}\mathbf{e}^i + \mathbf{a}\mathbf{e}^j) \in H$ for all $\lambda \in \mathbb{R}$, and we can choose λ so that q_j^λ is maximal such that $\mathbf{q}^\lambda \in [H]$. Then $q_i^\lambda = \overline{C}_i$ or $q_j^\lambda = \overline{C}_j$. But $\mathbf{q} \preceq_H \mathbf{q}^\lambda$ implies that $q_j \geq q_j^\lambda$ and $q_i \geq q_i^\lambda$. Thus $\mathbf{q} = \mathbf{q}^\lambda$ and so $0 \notin I(\mathbf{q})$. The fact that slope vectors are primitive imply that the slope vector for (i, j) at \mathbf{p} is as stated. Q.E.D.

PROPOSITION G.28: *We have $w_v(F) = w_B(F)$ for every facet F of \mathcal{H}_v .*

PROOF: Recall $\mathcal{H}_v = \mathcal{H}_B$ by Proposition G.26. Let $H \subseteq \mathcal{H}_v$ be a hyperplane of orientation (i, j) with $i > 0$ and slope $\sigma \in \mathbb{Q}_S$. We will use induction on the partial order $\preceq_{H, S}$ to show that $w_v^{ij}(\mathbf{p}; \mathbf{s}; J) = w_B^{ij}(\mathbf{p}; \mathbf{s}; J)$ for every vertex $\mathbf{p} \in H$ of $[\mathcal{H}_v]$ with $i, j \in I(\mathbf{p})$, every slope vector \mathbf{s} for (i, j) at \mathbf{p} with $\frac{s_i}{s_j} = \sigma$, and every $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$. The proposition follows because, by Lemma G.23, if $F \subseteq H$ is a facet of $\mathcal{H}_v = \mathcal{H}_B$ then $w_v(F) = w_v^{ij}(\mathbf{p}; \mathbf{s}; J) = w_B^{ij}(\mathbf{p}; \mathbf{s}; J) = w_B(F)$ for some such $\mathbf{p}, \mathbf{s}, J$.

Consider the base case. By Lemma G.27, the least point \mathbf{p} in $[H]$ is a vertex of $[\mathcal{H}_v]$, we have $I(\mathbf{p}) = \{i, j\}$, and the slope vector \mathbf{s} for (i, j) at \mathbf{p} with $\frac{s_i}{s_j} = \sigma$ is unique. Thus, (\mathbf{p}, \mathbf{s}) is minimal w.r.t. $\preceq_{H, S}$. By Definition G.13 and $I(\mathbf{p}) \setminus \{i, j\} = \emptyset$, as well as Lemma G.16, we have $w_v^{ij}(\mathbf{p}; \mathbf{s}; \emptyset) = m_v^{ij}(\mathbf{p}; \mathbf{s}) = m_B^{ij}(\mathbf{p}; \mathbf{s}) = w_B^{ij}(\mathbf{p}; \mathbf{s}; \emptyset)$.

Now suppose (\mathbf{p}, \mathbf{s}) is not the least pair, and assume the inductive hypothesis holds for all pairs $(\mathbf{p}', \mathbf{s}') \prec_{H, S} (\mathbf{p}, \mathbf{s})$. For any $J \subseteq I(\mathbf{p}) \setminus \{i, j\}$, with $J \neq \emptyset$, Lemma G.24 tells us that there

exist $\mathbf{p}', \mathbf{s}', J'$ such that $(\mathbf{p}', \mathbf{s}') \prec_{H,S} (\mathbf{p}, \mathbf{s})$ and $w^{ij}(\mathbf{p}'; \mathbf{s}'; J') = w^{ij}(\mathbf{p}; \mathbf{s}; J)$. By the inductive hypothesis, $w_v^{ij}(\mathbf{p}; \mathbf{s}; J) = w_v^{ij}(\mathbf{p}; \mathbf{s}'; J') = w_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s}'; J') = w_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s}; J)$. As $m_v^{ij}(\mathbf{p}; \mathbf{s}) = m_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s})$ by Lemma G.16, Definition G.13 now implies $w_v^{ij}(\mathbf{p}; \mathbf{s}; \emptyset) = w_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{s}; \emptyset)$, which completes the proof. *Q.E.D.*

G.3.2. The main theorem

We now prove that \mathcal{B} is the unique bid collection, up to normalisation, satisfying $D_v = D_{\mathcal{B}}$.

LEMMA G.29: *For any two bid collections \mathcal{B} and \mathcal{B}' , normalised in the same way but distinct, we have $D_{\mathcal{B}} \neq D_{\mathcal{B}'}$.*

PROOF: Suppose $D_{\mathcal{B}} = D_{\mathcal{B}'}$ (so $\mathcal{H}_{\mathcal{B}} = \mathcal{H}_{\mathcal{B}'}$). The demand sets are independent of the normalisation, and we may assume that the bid collections are normalised with respect to the same bounding box $[\mathcal{P}]$ so that all roots \mathbf{r} of bids are equal to $\tau(\mathbf{p})$ for some $\mathbf{p} \in [\mathcal{P}]$. Without loss of generality, there exists a bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ in \mathcal{B} that is not in \mathcal{B}' . If \mathbf{b} is unconditional (so $|I| = 1$) and interested in good i , then at any generic prices \mathbf{p} with sufficiently large p_i , demand of good i differs between \mathcal{B} and \mathcal{B}' , contradicting assumption $D_{\mathcal{B}} = D_{\mathcal{B}'}$. Now suppose $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m)$ is interested in two or more goods, and $i > j$ are the largest such goods. Let $m' \neq m$ be the multiplicity of the bid with root \mathbf{r} and tradeoffs \mathbf{t} in \mathcal{B}' , if it exists, or let $m' = 0$ if no such bid exists. Fix the unique point $\mathbf{p} \in \mathcal{P}$ for which $\mathbf{r} = \tau(\mathbf{p})$ (so $p_k = r_k$ for $k \in I(\mathbf{r})$ and $p_k = \underline{C}_k$ for $k \notin I(\mathbf{r}) \cup \{0\}$). By Lemma G.15 and $m_{\mathcal{B}}^{ij} = \sum_{\mathbf{b} \in \mathcal{B}} m_{\mathbf{b}}^{ij}$, we have $m_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{t}) = m \neq m' = m_{\mathcal{B}'}^{ij}(\mathbf{p}; \mathbf{t})$. But $D_{\mathcal{B}} = D_{\mathcal{B}'}$ implies $(\mathcal{H}_{\mathcal{B}}, w_{\mathcal{B}}) = (\mathcal{H}_{\mathcal{B}'}, w_{\mathcal{B}'})$ and hence $m_{\mathcal{B}}^{ij}(\mathbf{p}; \mathbf{t}) = m_{\mathcal{B}'}^{ij}(\mathbf{p}; \mathbf{t})$, a contradiction. *Q.E.D.*

PROOF OF THEOREM 3.1: Let \mathcal{B} be the bid collection constructed as described in Appendix G.2. Propositions G.26 and G.28 state that $(\mathcal{H}_v, w_v) = (\mathcal{H}_{\mathcal{B}}, w_{\mathcal{B}})$. By construction of the single-minded bids in \mathcal{B} , we have $D_v(\mathbf{p}) = D_{\mathcal{B}}(\mathbf{p})$ for some generic price \mathbf{p} , so $D_v(\mathbf{p}) = D_{\mathcal{B}}(\mathbf{p})$ follows for all $\mathbf{p} \in \mathcal{P}$ by Proposition G.2. Lemma G.29 shows uniqueness up to normalisation. *Q.E.D.*

PROOF OF COROLLARY 3.2: If v is a strong substitutes valuation, all hyperplanes in \mathcal{H}_v have slope 1, so we may set $S = 1$. Hence, all bids in the bid collection \mathcal{B} have one-to-one tradeoffs by construction. *Q.E.D.*

To prove Corollaries 3.3 and 3.7 we first show:

LEMMA G.30: *Let v be a concave substitutes valuation, and suppose that there exists $\bar{\mathbf{p}}$ such that $D_v(\mathbf{p}) = \{\mathbf{0}\}$ for all $\mathbf{p} \geq \bar{\mathbf{p}}$. Let bid collection \mathcal{B} satisfy $D_{\mathcal{B}} = D_v$. Then there are no unconditional bids in \mathcal{B} , and all bids in \mathcal{B} are interested in good 0.*

PROOF: First we see that Part (i) of Definition 3.6 for regular valuations is satisfied by v . Fix $i \in [n]$, let $\mathbf{p} \in \mathcal{P}$ satisfy $p_i > \bar{p}_i$ and let \mathbf{q} be generic and infinitesimally close to \mathbf{p} . Let $\mathbf{q}' \geq \bar{\mathbf{p}}$ be defined by $q'_j = \max\{q_j, \bar{p}_j\}$ for $j \in [n]$. By assumption $D_v(\mathbf{q}') = \{\mathbf{0}\}$. Moreover we can move from \mathbf{q} to \mathbf{q}' in steps by increasing the prices of goods $j \neq i$ in turn, and by genericity of \mathbf{q} we can assume that demand is unique after every such step. Because v is a substitutes valuation, demand for good i weakly increases with every such step. But demand for good i is zero at \mathbf{q}' . So demand for good i is zero at \mathbf{q} . This holds for any generic \mathbf{q} close to \mathbf{p} , so we can conclude that $x_i = 0$ for any $\mathbf{x} \in D_v(\mathbf{p})$. The choice of i was arbitrary, so Part (i) of Definition 3.6 holds.

Recall we defined a bounding box with boundaries \underline{C} and \overline{C} chosen so that the relative interior of every face of \mathcal{L}_v meets the interior of $[\mathcal{P}]$. Note that this property still holds if we weakly decrease \underline{C}_i and weakly increase \overline{C}_i for some $i \in [n]$. So allow these bounds to depend on λ , and notate the box $[\mathcal{P}]^\lambda = \{\mathbf{p} \in \mathcal{P} \mid \underline{C}_i(\lambda) \leq p_i \leq \overline{C}_i(\lambda)\}$.

Algorithm 1 uses any such box $[\mathcal{P}]^\lambda$ to generate a bid collection \mathcal{B}^λ also satisfying $D_v = D_{\mathcal{B}^\lambda}$, but these bids are normalised differently. Every bid $\mathbf{b} = (\mathbf{r}, \mathbf{t}, m) \in \mathcal{B}^\lambda$ except the unconditional bids (those for which $|I| = 1$) satisfies $\mathbf{r} = \tau(\mathbf{p})$ where \mathbf{p} is a vertex of $[\mathcal{L}_v]^\lambda$, and so in particular $\mathbf{p} \in \mathcal{L}_v \cap [\mathcal{P}]$. Here \mathbf{b} is interested in good 0 if and only if $p_i < \overline{C}_i(\lambda)$ for all $i \in [n]$ and \mathbf{b} is interested in good $i \in [n]$ if and only if $p_i > \overline{C}_i(\lambda)$.

However, the bid collection \mathcal{B} in the statement of the lemma is normalised as in Section 2.2. By Lemma G.29, \mathcal{B}^λ and \mathcal{B} are the same up to normalisation of bids (see Lemma 2.2).

Suppose, for a contradiction, there exists $\mathbf{b} = (\mathbf{r}; \mathbf{t}, m) \in \mathcal{B}$ interested in I with $|I| \geq 2$ and $0 \notin I$. Fix such a bid. For every suitable $[\mathcal{P}]^\lambda$ there exists $\mathbf{b}^\lambda \in \mathcal{B}^\lambda$ with the same demand, but potentially a different normalisation. The corresponding bid $\mathbf{b}^\lambda \in \mathcal{B}^\lambda$ satisfies $\mathbf{b}^\lambda = (\mathbf{r}^\lambda, \mathbf{t}, m)$ where $\mathbf{r}^\lambda = \tau(\mathbf{p}^\lambda)$ for some $\mathbf{p}^\lambda \in [\mathcal{P}] \cap \mathcal{L}_v$. By definition of τ and by Lemma 2.2, thus, $t_i(p_i^\lambda - r_i) = t_j(p_j^\lambda - r_j)$ for all $i, j \in I$, where $p_i^\lambda \leq \overline{C}_i$ for all $i \in [n]$ and $p_k^\lambda = \overline{C}_k(\lambda)$ for some $k \in [n]$.

In principle the k such that $p_k^\lambda = \overline{C}_k(\lambda)$ can depend on λ . Let us set $\overline{C}_i(\lambda) = \overline{C}_i + \frac{\lambda}{t_i}$ for all $i \in I$, and $\overline{C}_i(\lambda) = \overline{C}_i$ otherwise, with $\underline{C}(\lambda) = \underline{C}$. Then, for all $\lambda \geq 0$, we know $[\mathcal{P}]^\lambda$ satisfies the required properties and so there exists \mathbf{p}^λ as above. Consider the $\lambda = 0$ case: let $k \in I$ satisfy $p_k^0 = \overline{C}_k$. Now $\mathbf{p}^\lambda = \mathbf{p}^0 + \lambda \sum_{i \in I} \frac{\mathbf{e}_i}{t_i}$, as this vector has all the required properties. So $\mathbf{p}^\lambda = \mathbf{p}^0 + \lambda \sum_{i \in I} \frac{\mathbf{e}_i}{t_i} \in \mathcal{L}_v$ for all $\lambda > 0$. But, by Corollary D.2, this implies that Part (i) of Definition 3.6 does not hold. This is a contradiction to the first part of our proof. So no bid such as \mathbf{b} exists: every $\mathbf{b} \in \mathcal{B}$ which is not unconditional is interested in good 0.

Finally, we show that there are no unconditional bids in \mathcal{B} . Let prices \mathbf{p} lie above $\overline{\mathbf{p}}$ and also above every root of every regular bid in \mathcal{B} . Then $D_{\mathcal{B}}(\mathbf{p}) = D_v(\mathbf{p}) = \{\mathbf{0}\}$ by definition of \mathcal{B} and of $\overline{\mathbf{p}}$. But $D_{\mathcal{B}}(\mathbf{p})$ is equal to the sum of the demands of unconditional bids for non-null goods, because every regular bid demands $\mathbf{0}$ at prices above its root. Finally, there are no unconditional bids for good 0 because \mathcal{B} is parsimonious. So there are no unconditional bids in \mathcal{B} . *Q.E.D.*

COROLLARY 3.3: *For any substitutes concave valuation v , there exists a regular bid collection \mathcal{B} for any $\underline{\mathbf{p}} \in \mathcal{P}$, such that $D_{\mathcal{B}}(\mathbf{p}) = D_v(\mathbf{p})$ for all $\mathbf{p} \geq \underline{\mathbf{p}}$, if there exists $\overline{\mathbf{p}}$ such that $D_v(\mathbf{p}) = \{\mathbf{0}\}$ for all $\mathbf{p} \geq \overline{\mathbf{p}}$.*

PROOF: By Theorem 3.1, there exists a bid collection, \mathcal{B} , such that $D_{\mathcal{B}} = D_v$. By Lemma G.30, none of the bids in \mathcal{B} are unconditional and all are interested in good 0, that is, they have $r_0 = 0$. We replace all non-regular bids in \mathcal{B} with regular bids and show that the resulting bid collection $\tilde{\mathcal{B}}$ satisfies $D_{\mathcal{B}}(\mathbf{p}) = D_{\tilde{\mathcal{B}}}(\mathbf{p})$ for every $\mathbf{p} \geq \underline{\mathbf{p}}$. If a bid $(\mathbf{r}; \mathbf{t}; m) \in \mathcal{B}$ satisfies $r_i > -\infty$ for all true goods $i \in [n]$, it is already regular and thus included in $\tilde{\mathcal{B}}$ without modification. Otherwise, suppose the bid is interested in goods $I \subsetneq [n]_0$, so $t_i = 0$ and $r_i = -\infty$ for each good $i \in [n]_0 \setminus I$. For each such good i , we now change t_i to 1 and r_i to $\underline{p}_i - 1$ so that the bid prefers the null good to good i at any prices $\mathbf{p} \geq \underline{\mathbf{p}}$, and include this modified bid in $\tilde{\mathcal{B}}$. Moreover, the utility of receiving the goods that the bid was originally interested in is unaffected. *Q.E.D.*

PROOF OF COROLLARY 3.7: Let v be a regular substitutes valuation and \mathcal{B} the bid collection with $D_v = D_{\mathcal{B}}$, which exists by Theorem 3.1. It is clear that a regular substitutes valuation satisfies the condition of Lemma G.30, and so every bid in \mathcal{B} is interested in good 0, and is not

unconditional. By Lemma 2.2, every bid in \mathcal{B} therefore has a unique normalisation. So, by Lemma G.29, \mathcal{B} is identical to the bid collection identified by Algorithm 1.

Now suppose there exists a bid $\mathbf{b} = (\mathbf{r}; \mathbf{t}; m) \in \mathcal{B}$ which is not interested in $i \in [n]$. It is not unconditional, so by construction (Algorithm 1) its root \mathbf{r} satisfies $\mathbf{r} = \tau(\mathbf{p})$, where \mathbf{p} is a vertex of $[\mathcal{L}_v]$ such that $p_i = \underline{C}_i$, the lower face of the bounding box. In particular $\mathbf{p} \in \mathcal{L}_v$.

But if we consider instead the bounding box with lower i -bound $\underline{C}_i - \lambda$ and all other bounds the same, Algorithm 1 identifies the same bid $(\mathbf{r}; \mathbf{t}; m)$, but now $\mathbf{r} = \tau^\lambda(\mathbf{p}')$, where \mathbf{p}' is a vertex of $[\mathcal{L}_v]^\lambda$ defined by the new bounding box, satisfying $p'_i = \underline{C}_i - \lambda$. Because \mathbf{r} is unchanged and the other bounds of the box are unchanged, we have $p'_j = p_j$ for $j \neq i$. So $\mathbf{p}' = \mathbf{p} - \lambda \mathbf{e}^i \in \mathcal{L}_v$. We can perform this construction for all $\lambda > 0$. But this contradicts Definition 3.6 Part (i) by Corollary D.2. So no bid such as \mathbf{b} exists in \mathcal{B} : every bid in \mathcal{B} is interested in all goods, that is, it is regular. *Q.E.D.*