OPERATIONS RESEARCH

Iterative Auction Design for Graphical Valuations Part I: Tree Valuations

SCHOLARONE™
Manuscripts

Iterative Auction Design for Graphical Valuations
Part I: Tree Valuations[*]

Ozan Candogan, Asuman Ozdaglar and Pablo Parrilo[†]

November 25, 2013

Abstract

In this work, we design simple and efficient iterative auctions for selling multiple items in settings where bidders' valuations belong to a special class (tree valuations) that can exhibit both value complementarity and substitutability. Our first contribution is to provide a compact linear programming formulation of the efficient allocation problem and use it to establish the existence of a Walrasian equilibrium for tree valuations that satisfy an additional technical condition. This result reveals a new class of valuations for which a Walrasian equilibrium exists in the presence of value complementarities. We then provide an iterative algorithm that can be used for the solution of this linear programming formulation. Complementing the algorithm with an appropriate payment rule, we obtain an iterative auction which implements the efficient outcome (at an ex-post perfect equilibrium). This auction relies on a simple pricing rule, compact demand reports, and uses a novel (interleaved) price update structure to assign final payments to bidders that guarantee truthful bidding.

Keywords: Auction design, efficient auctions, iterative auctions, graphical valuations, Walrasian equilibrium, primal-dual algorithms.

1 Introduction

Iterative auctions are a class of mechanisms that are commonly employed in practice. In these auctions, the auctioneer sets prices for the items she is selling, bidders report which items they are interested in at the given prices, and in response to these reports the auctioneer updates the prices. This process terminates when the auctioneer determines a final allocation of items to bidders.

[†]Ozan Candogan is with Duke University, Fuqua School of Business. Asuman Ozdaglar and Pablo Parrilo are with the Laboratory for Information and Decision Systems, MIT. E-mails: ozan.candogan@duke.edu, asuman@mit.edu, parrilo@mit.edu.

1

Examples of iterative auctions include the auctions used for selling electricity, natural resources, bus routes, spectrum, art, antiquities, wine, jewelry, and some auctions used for procurement.

The well-known English and Dutch auctions can be viewed as examples of single-item iterative auctions. When bidders have independent private values, these auctions allocate the item efficiently, i.e., the bidder with the highest value receives the item. On the other hand, in more general multi-item settings (such as spectrum or procurement auctions) the iterative auctions that are present in the literature do not always have efficiency guarantees. More precisely, they either implement the efficient outcome under restrictive assumptions (such as the gross substitutes assumption, Gul and Stacchetti (2000); Ausubel (2006)), or they require complex pricing structures that involve a different price for each bundle of items (Ausubel and Milgrom, 2002; Bikhchandani and Ostroy, 2002; Vohra, 2011). The auctions in the first category do not allow for value complementarity between different items, which is observed in various auction environments. Those in the second category may not be practical, since they require the number of different prices that are reported to the bidders at each stage of the auction to be exponential in the number of items.

Motivated by these considerations, in this paper we develop novel and simple iterative auctions for multi-item settings. We obtain our results by focusing on a special class of bidder valuations, called graphical valuations, for which the value of a bundle of items is given by the sum of individual values of items and values for pairs of items that capture the complementarity or substituability between these items. This valuation class can naturally be represented in terms of a weighted undirected graph, where the nodes correspond to items and edges link pairs of items that exhibit complementarity/substitutability. In this paper, we further restrict attention to graphical valuations, where the underlying graph is a tree (i.e., has no cycles) and satisfies an additional technical condition (sign-consistency) which involves all bidders to view a given pair of items either as substitutes or as complements. While the tree valuations are a special case, they are fundamental for understanding the iterative auction design problem, and the results we obtain for this case extend to more general graphical valuations. For sign-consistent tree valuations, we design efficient iterative auctions that rely on the *anonymous item pricing rule* and terminate when a natural market clearance condition holds (i.e., a *Walrasian equilibrium* is reached).[1]

Our auction design approach has three main steps. First, we establish that a Walrasian equilibrium exists for sign-consistent tree valuations, and provide a *compact linear programming (LP) formulation of the efficient allocation problem* whose optimal primal and dual solutions respectively correspond to the Walrasian equilibrium allocations and prices. Second, we provide a primal-dual algorithm that solves this LP by iteratively updating the dual variables according to bidders' valuations (Figure 1a), and converges to a Walrasian equilibrium. We then interpret a subset of the dual variables as *anonymous item prices*, and show that it is possible to identify the dual update direction and the step size in this algorithm without the explicit knowledge of bidders' valuations, when bidders *compactly report their demand* in response to price updates. This allows us to slightly

---

[1]In our companion paper Candogan et al. (2013), we study general graphical valuations and show that the efficient outcome can be implemented using auctions that rely on slightly more complex pricing rules than item pricing. In particular, these pricing rules involve both item prices and discounts/markups for pairs of items.

alter the algorithm in Figure 1a, and obtain an alternative algorithm that relies only on the demand information (Figure 1b). The third step involves obtaining a novel *interleaved tree auction* (Figure 1c) by modifying the second algorithm and complementing it with an appropriate payment scheme which guarantees that bidders truthfully reveal their demand (and the efficient outcome is implemented) at an (ex-post perfect) equilibrium. In particular, the modified algorithm identifies the Walrasian equilibria for markets that involve all bidders, as well as all bidders but one, by interleaving the demand queries associated with different markets. These Walrasian equilibria are then used to determine and assign the VCG payments (i.e. the externality bidders create on each other). These payments guarantee that at equilibrium bidders truthfully report their demand, and the efficient outcome is implemented.

(a) A primal-dual algorithm for obtaining Walrasian equilibria.

(b) A demand-based algorithm for obtaining Walrasian equilibria.
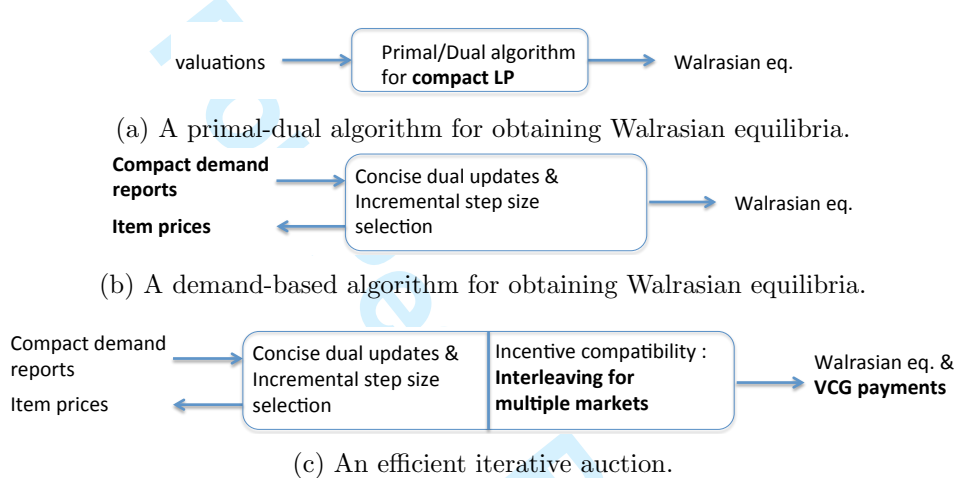

(c) An efficient iterative auction.

Figure 1: In order to obtain our auction, we first provide a primal-dual algorithm that solves the LP formulation of the efficient allocation problem and converges to a Walrasian equilibrium (a). This algorithm assumes the knowledge of bidders' valuations. We then provide a slight modification of this algorithm that relies only on the demand information (b). Finally, by modifying this algorithm to identify Walrasian equilibria for multiple markets (that consist of all bidders and all bidders but one), we obtain an iterative auction (c).

In the rest of this section, we explain our auction design approach and main contributions in more detail, and discuss the related literature.

Existence of Walrasian equilibrium: Anonymous item pricing rule associates a price with each item, and this price is the same for all bidders. Walrasian equilibrium provides a natural termination condition for auctions that rely on this pricing rule, since at this outcome all bidders demand disjoint sets of items, and hence no bidder needs to compete with the remaining bidders to acquire the set of items that she demands. In Section 3, we show that for sign-consistent tree valuations a Walrasian equilibrium exists. This result identifies a new class of valuations, for which a Walrasian equilibrium exists even in the presence of complementarities. Additionally, the result is tight, in the sense that if the sign-consistency or the tree assumption is relaxed (and more general graphical valuations are allowed), then a Walrasian equilibrium need not exist. In

Section 3, we also provide a compact LP formulation of the efficient allocation problem, which (i) involves polynomially many (in number of items and bidders) variables and constraints, and (ii) identifies the Walrasian equilibria at optimal primal-dual solutions. The existence of such an LP formulation implies that for sign-consistent tree valuations a Walrasian equilibrium can be obtained in a computationally efficient manner (through the solution of the LP).

A demand-based iterative algorithm: In Section 4, we provide a primal-dual algorithm that terminates with the efficient allocation and Walrasian equilibrium prices in finite time. This algorithm starts with a dual feasible solution (of the LP formulation of the efficient allocation problem), and at each iteration updates it in an improvement direction obtained using information on the active constraints (in the dual LP). The step size associated with these updates is equal to the largest step in the update direction that preserves feasibility. The algorithm (Figure 1a) uses the knowledge of bidders' valuation both for obtaining an update direction and choosing the step size, and terminates with the primal-dual optimal solutions of the LP formulation (and the Walrasian equilibria). A subset of these dual variables can be interpreted as the *anonymous* prices that are associated with the items. We can write down the dual update direction concisely in terms of these variables, when information on bidders' demand (at these prices) is available. In particular, given the prices, the active constraints in the underlying LP can be determined from bidders' demand information, which in turn provides an update direction. Additionally, we show that updating the prices in this direction incrementally until a bidder starts demanding a new bundle, guarantees that the largest update in the given direction that preserves feasibility is made, thus determining the stepsize of our algorithm. This structure allows for running the algorithm by only relying on bidders' demand reports (Figure 1b), and suggests a natural iterative auction that updates the prices according to these reports, and converges to the Walrasian equilibria (and the efficient outcome) when bidders are truthful. Moreover, by exploiting the structure of sign-consistent tree valuations we show that the demand reports take the compact form that involves specifying the items that belong to every (or no) demanded bundle, and pairs of items where at least (at most) one of the items is demanded.

An incentive compatible auction: We establish in Section 5 that when final payments of bidders are equal to the sum of the prices of items a bidder receives at the end of the algorithm in Figure 1b, bidders may have incentive to misreport their demand. However, we also show that this algorithm can be modified in a way that guarantees that bidders truthfully report their demand. The high-level idea behind our approach is to use the aforementioned algorithm to identify Walrasian equilibria for the markets that consist of (i) all bidders, and (ii) all bidders but one. An important and novel feature of our approach is the interleaved price update structure that avoids running a separate auction for these markets, and instead allows for finding these Walrasian equilibria jointly by updating the prices greedily giving priority to markets that are "closer" to a Walrasian equilibrium. We establish that the price/surplus difference between these equilibria can be used to identify the VCG payments of bidders, which align bidders' utilities with the goal of

4

efficiency. For such final payments, we prove that it is an ex-post perfect equilibrium for bidders to truthfully reveal their demand. That is, no bidder has incentive to misreport her demand after any (bidding) history, for any valuations of her opponents. Thus, with the aforementioned modification of our algorithm (Figure 1c), we obtain an iterative auction that implements the efficient outcome for sign-consistent tree valuations.

Our auction relies on a simple (anonymous item) pricing rule, compact demand reports, and uses the interleaved structure to guarantee truthful bidding. The results of this paper indicate that when valuations of bidders can be represented using a simple graphical model, the auctioneer can systematically exploit this structure to implement the efficient outcome by iterative auctions that rely on simple pricing rules. This suggests that it is possible to develop practical iterative auctions with provable efficiency guarantees by better understanding and exploiting the structure of bidders' valuations.

Related literature: In standard multi-item auction settings, the VCG mechanism can be used to implement the efficient outcome in dominant strategy equilibria. Despite this desirable strategic feature, VCG mechanisms are rarely used in practice. On the other hand, iterative auctions which share similar equilibrium properties to VCG mechanisms are prevalent (see for instance Ausubel and Milgrom (2006); Rothkopf et al. (1990); Engelbrecht-Wiggans and Kahn (1991)), and have found applications in spectrum auctions, electricity auctions, online markets (such as eBay) (McAfee et al., 2010; Ausubel and Cramton, 2004; Ausubel, 2004), as well as procurement settings (Hohner et al., 2003; Cramton et al., 2006). This has stimulated significant interest in recent literature, and led to development of a number of novel multi-item iterative auctions. Examples include the package bidding auction (Ausubel and Milgrom, 2002), iBundle auction (Parkes, 1999), clinching auction and its variants (Ausubel, 2004, 2006), auctions that rely on universally competitive equilibria (UCE) (Mishra and Parkes, 2007, 2009), and the best response auction (Nisan et al., 2011).

The standard approach (Vohra, 2011) for developing iterative auctions involves three main steps. Efficient iterative auctions implicitly solve an optimization problem, and find the welfare maximizing allocation. The first step involves providing a linear programming formulation of the efficient allocation problem. Then, iterative algorithms for solutions of this optimization problem are developed. Finally, complementing these iterative algorithms with appropriate payment schemes that guarantee truthful demand revelation at equilibrium, iterative auctions are obtained. A number of works employ such iterative solutions of the efficient allocation problem for the design of iterative auctions (Bikhchandani et al., 2002; De Vries and Vohra, 2003; Parkes, 2006; De Vries et al., 2007; Vohra, 2011; Mishra and Parkes, 2007). On the other hand, the iterative algorithms (and the corresponding iterative auctions) present in the literature often rely on using exponentially many prices at each step (a price for every bundle of items), as the underlying valuations are general and the corresponding LP formulations do not admit a simple structure. In this paper, we follow a similar approach, but by exploiting the special structure of graphical valuations, we obtain provably efficient auctions that rely on simple pricing rules.

Many of the existing iterative auctions that implement the efficient outcome by relying on a simple pricing rule are applicable only in settings where a Walrasian equilibium exists. Bikhchandani and Mamer (1997) provide a linear programming characterization of the existence of a Walrasian equilibrium (which we discuss in more detail in Section 2). A notable class of valuations where a Walrasian equilibrium exists and simple iterative auctions guarantee efficiency, is the class of gross substitutes (GS) (Kelso and Crawford, 1982; Gul and Stacchetti, 1999). The GS property, on the other hand, does not allow for any value complementarity between different items. A generalization of this class which allows for a very specific value complementarity structure is the class of gross substitutes and complements (GSC), see Sun and Yang (2006, 2009). The GSC structure suggests that items can be grouped into two sets so that all items in a given set are gross substitutes, and items that belong to different sets are complements. It is possible to establish the existence of a Walrasian equilibrium and provide simple iterative auctions (Sun and Yang, 2006, 2009) for such valuations. However, these results are limited to the particular complementarity structure imposed by the GSC valuations, and do not overlap with our contributions (see Appendix A). Baldwin and Klemperer (2012) extends the existence result for GSC and identifies more general demand structures for which a Walrasian equilibrium exists. However, the results of this paper are not directly comparable to ours since while our characterization is in terms of the conditions on the valuation functions, the aforementioned result is in terms of bidders' demand.[2]

The idea of using simple pricing rules for implementing the efficient outcome is also explored in Mishra and Parkes (2009); Bikhchandani et al. (2011); Lahaie (2009, 2011). The first paper focuses on settings, where each bidder has unit demand (i.e., never demands more than a single item) or items are homogeneous (i.e., valuations are only a function of *number of items* acquired by the bidders). The second work provides an auction that employs only a single price, and guarantees efficiency in settings with an additional combinatorial structure (a matroid structure) is present. We note that these auctions are not directly applicable in our setting, since the valuations we consider do not exhibit the structures imposed in these papers. The last two papers use kernel methods from machine learning for obtaining iterative algorithms that solve for the efficient allocation. These algorithms rely on pricing rules that (potentially) become progressively more complex over time in order to identify (nonlinear) market clearing prices, but do not lead to an incentive compatible auction that implements the efficient outcome. In our setting, it is not necessary to rely on pricing rules that are more complex than anonymous item pricing, and our objective is to provide an auction that implements the efficient outcome at an ex-post perfect equilibrium.

The idea of identifying Walrasian equilibria for markets with all bidders, and all bidders but one, and using these for the computation of VCG payments (and auction design) was also employed in Ausubel (2006). This paper focuses on gross substitutes, and identifies such Walrasian equilibria by running a separate auction for each market. In our setting, despite the fact that the gross substitutes assumption does not hold, we are able to compute the VCG payments and implement the efficient

---

[2]Consequently, it is possible to obtain examples where bidders' valuations are sign-consistent tree valuations (hence a Walrasian equilibrium exists), but these valuations are associated with a demand structure for which a Walrasian equilibrium does not always exist.

outcome. Additionally, we obtain our result without explicitly running separate auctions, but by using a novel interleaved price update structure (as explained in Section 5).

The problem of finding the efficient outcome in a general combinatorial setting is hard both from a computational complexity and a communication complexity point of view (Lehmann et al., 2006; Nisan and Segal, 2006; Cramton et al., 2006; Blumrosen and Nisan, 2010). This motivates considering classes of value functions with additional structure (Blumrosen and Nisan, 2010; Cramton et al., 2006). Recently, Conitzer et al. (2005); Sandholm and Boutilier (2006); Zhou et al. (2009); Abraham et al. (2012) considered graphical valuation structures that are similar to those that we consider in this paper (referred to as k-wise valuations in the first two works), characterized the computational complexity of the efficient auction design problem, and provided approximately efficient direct mechanisms. Additionally, the first two works studied the problem of eliciting preferences for this class of valuations, and established that bidders' preferences can be elicited using polynomially many "value queries". In our paper, we adopt a similar valuation model, but focus on developing iterative auctions that implement the efficient outcome at an ex-post perfect equilibrium, by relying on natural demand queries and a simple anonymous item pricing rule.

2   Model and Preliminaries

In this section, we introduce our graphical valuation model (Section 2.1). We also discuss structural properties (e.g. pricing rules, termination conditions) of iterative auctions, and introduce the Walrasian equilibrium concept that plays a key role in the design of such auctions (Section 2.2).

2.1   Graphical Valuations

We consider settings where an auctioneer sells N (heterogeneous) items to M bidders. We denote the set of items by $\mathcal{N}$ and the set of bidders by $\mathcal{M}$. For each bidder $m \in \mathcal{M}$, the value function $v^m : 2^{\mathcal{N}} \to \mathbb{R}^+$, captures the value $v^m(S)$ this bidder has for any set $S \subseteq \mathcal{N}$ of items. In the *auction setup* we consider, we assume that bidders' value functions are *private*, i.e., each bidder knows her own value function, but not the value functions of other bidders. We make two standard assumptions about the value functions:

Assumption 2.1. *Bidders have value zero for not receiving any items, i.e., $v^m(\emptyset) = 0$. Moreover, the value functions are* monotone, *i.e., $v^m(S_1) \leq v^m(S_2)$ if $S_1 \subseteq S_2$.*

In the multi-item setting considered here, the value functions may not be *additive*, i.e., the value a bidder has for a set S need not be equal to the sum of the values of items that are contained in this set ($v^m(S) \neq \sum_{i \in S} v^m(\{i\})$). If for any bidder m items i and j satisfy $v^m(\{i,j\}) \geq v^m(\{i\}) + v^m(\{j\})$, we say that these items are *(pairwise) complements*. On the other hand, if i and j are such that $v^m(\{i,j\}) \leq v^m(\{i\}) + v^m(\{j\})$ for all m, we refer to them as *(pairwise) substitutes*.[3]

---

[3]In this work we are mainly interested in pairwise complemenatary/substitutability, and unless noted otherwise, we refer to these simply as complementarity/substitutability.

We next introduce the graphical valuation model that we focus on in this paper.

Definition 2.1 (Graphical Valuations). *Let $G = (\mathcal{N}, E)$ be a graph such that the set of nodes corresponds to the set of items $\mathcal{N}$ and there are edges between nodes (items) that may exhibit value complementarity or substitutability. We refer to G as a* value graph *for set of items $\mathcal{N}$.*

We say that the value function $v : 2^{\mathcal{N}} \to \mathbb{R}^+$ is graphical *(with respect to G) if it satisfies $v(S) = \sum_{i \in S} w_i + \sum_{(i,j) \in E | i,j \in S} w_{ij}$ for all $S \subset \mathcal{N}$, where $\{w_i\}_{i \in \mathcal{N}}$ represent the nonnegative weights associated with nodes, and $\{w_{ij}\}_{(i,j) \in E}$ represent the weights associated with the edges.*

This definition implies that a value function is graphical if there exist node weights and edge weights associated with the underlying value graph, such that the value of any bundle S is given by the sum of the weights of nodes and edges contained in an induced subgraph of G with set of nodes S.[4] See Figure 2 for an example.
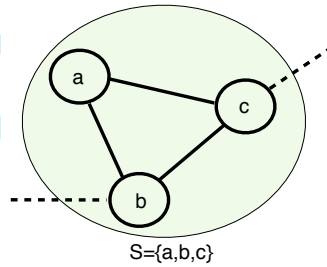

$$S = \{a,b,c\}$$

Figure 2: For a graphical valuation v, the value of bundle $S = \{a, b, c\}$ can be given as $v(S) = w_a + w_b + w_c + w_{ab} + w_{ac} + w_{bc}$.

A value function associates a value with each bundle of items, and hence can be thought of as a vector of length 2^N. On the other hand, the definition of graphical valuations suggests that these value functions can be uniquely defined by specifying N node weights and at most N^2 edge weights. This implies that the set of graphical valuations has smaller dimension than the set of general value functions, and hence is not fully general. Despite not being fully general, graphical valuations can naturally represent pairwise complementarity and substitutability. For instance, assume that i and j are two items such that for graphical valuation v, we have $w_{ij} \geq 0$. Then it can be seen that $v(\{i\}) + v(\{j\}) = w_i + w_j \leq w_i + w_j + w_{ij} = v(\{i, j\})$, and hence i, j are pairwise complements. Conversely, if $w_{ij} \leq 0$, then $v(\{i\}) + v(\{j\}) = w_i + w_j \geq w_i + w_j + w_{ij} = v(\{i, j\})$ and i, j are pairwise substitutes.

In this work, we assume that all bidders have graphical valuations with respect to a common graph $G = (\mathcal{N}, E)$. Additionally, in order to simplify the notation we allow for variables w_{ij}^m for $(i, j) \notin E$, and follow the convention $w_{ij}^m = 0$ for such variables.

---

[4]We provide a discussion of the differences between graphical valuations and other related special valuation classes such as gross substitutes, gross substitutes and complements (Kelso and Crawford, 1982; Gul and Stacchetti, 1999; Sun and Yang, 2006); sub/superadditive valuations, and sub/supermodular valuations (Blumrosen and Nisan, 2010) in Appendix A.

Assumption 2.2. *There exists a value graph $G = (\mathcal{N}, E)$ such that the value function of each bidder is a graphical valuation with respect to G, i.e., for each bidder $m \in \mathcal{M}$, there exist weights $\{w_i^m\}$ and $\{w_{ij}^m\}$ such that $v^m(S) = \sum_{i \in S} w_i^m + \sum_{(i,j) \in E | i, j \in S} w_{ij}^m = \sum_{i \in S} w_i^m + \sum_{i,j \in S} w_{ij}^m$ for all $S \subset \mathcal{N}$.*

Graphical valuations satisfying Assumption 2.2 capture the value complementarity/substitutability in many practical auction settings including spectrum auctions, truck route auctions, and real estate auctions. In these settings, the items that are auctioned correspond to different geographical regions, and there are value complementarities and substitutabilities between neighboring regions. For instance, in spectrum auctions, complementarities between adjacent geographical regions are present due to roaming and interference (Cramton et al., 1997; Moreton and Spiller, 1998). Similarly, different bands in the same geographical region can be viewed as substitutes as bidders may only have limited demand for spectrum in each geographical region. Such complementarities and substitutabilities exhibit a similar structure for all bidders, and can naturally be captured by graphical valuations by associating a node with each spectrum band - geographical region pair, and an edge with pairs of spectrum bands in adjacent (or the same) geographical regions (see Figure 3).[5]
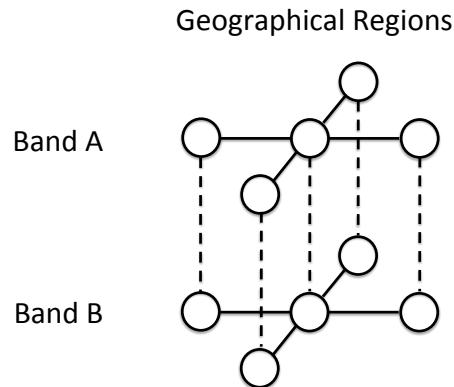

Figure 3: Consider a spectrum auction where two bands (A & B), over one central and four peripheral geographical regions are sold. Agents view the bands in neighboring geographical regions as complements, while they view different bands in the same geographical region as substitutes. This can be captured using the graphical model in the above figure, and assigning positive weights to the solid lines, and negative weights to the dashed ones.

2.2 Efficient Iterative Auctions and Walrasian Equilibrium

In this paper our objective is to design efficient auctions for graphical valuations. Given bundles of items $S^m \subset \mathcal{N}$ for all $m \in \mathcal{M}$, we say that $\{S^m\}_{m \in \mathcal{M}}$ is a *feasible allocation* if (i) each bidder

---

[5]Graphical valuations can be generalized to capture multilateral complementarities and substitutabilities among multiple items. In particular, we can associate a weight with each clique of the value graph with cardinality at most $r \in \{1, \ldots, N\}$, and represent the value a bidder has for a bundle S by the sum of the weights of all such cliques contained in this bundle. It can be seen that graphical valuations correspond to the case of $r = 2$. A discussion of this generalization and its applications in the context of iterative auction design can be found in Candogan (2013).

$m \in \mathcal{M}$ receives a bundle of items $S^m \subset \mathcal{N}$, (ii) each item is assigned to at most one bidder, i.e., $S^m \cap S^l = \emptyset$ for $m, l \in \mathcal{M}$ with $m \neq l$. An *efficient allocation* is a feasible allocation $\{S^m\}_{m \in \mathcal{M}}$ that maximizes the *welfare* or *total value*, i.e., $\sum_m v^m(S^m) = \max_{\{\{Z^m\}_m | Z^m \cap Z^l = \emptyset\}} \sum_m v^m(Z^m)$. An auction that terminates with an efficient allocation for any value functions is an *efficient auction*.

We focus on designing efficient auctions that have an iterative structure. In these auctions, the auctioneer sells items to bidders through a dynamic process, whereby she posts prices, and collects responses from bidders for demanded bundles of items. She uses this information to update the prices until a final allocation of bundles to bidders is determined. We refer to such auctions as *iterative auctions*.

Since final allocations are in terms of bundles, a natural structure for iterative auctions involves a different price for every bundle. However, due to the presence of exponentially many bundles (in the number of items) this pricing rule is informationally intensive. This motivates focusing on simpler pricing rules for iterative auction design. In the literature such a pricing rule, *anonymous item pricing*, has been employed for efficient iterative auction design in settings where valuations do not exhibit complementarities (Ausubel, 2004, 2006; Gul and Stacchetti, 1999, 2000). This pricing rule suggests offering a price p_i for item $i \in \mathcal{N}$ to all bidders, and it compactly captures the price of every bundle as a summation of prices of items contained in it. In this work, we design iterative auctions that rely on anonymous item pricing and guarantee efficiency for subclasses of graphical valuations which exhibit complementarity (in particular tree valuations, see Section 3).

At given anonymous item prices, we refer to the quantity $v^m(S) - \sum_{i \in S} p_i$ as the *surplus* bidder m associates with bundle S. We say that a bundle S^* is *demanded* by bidder m, if maximum surplus is achieved for this bundle, i.e., $v^m(S^*) - \sum_{i \in S^*} p_i = \max_S v^m(S) - \sum_{i \in S} p_i$. We denote the set of bundles bidder m demands by D^m, i.e., $D^m = \arg\max_S v^m(S) - \sum_{i \in S} p_i$.

Auctions that rely on anonymous item prices can be terminated when bidders demand disjoint bundles of items. Observe that this is a natural termination point for the auction, since bidders do not compete with each other for the items that they demand. This outcome coincides with the classical Walrasian equilibrium concept from microeconomic theory (Mas-Collel et al., 1995):

Definition 2.2 (Walrasian equilibrium). *Consider prices $\{p_i\}_{i \in \mathcal{N}}$, and allocation $\{S^m\}_m$ where $S^m \subset \mathcal{N}$ for every m. The tuple $(\{p_i\}_i, \{S^m\}_m)$ is a* Walrasian equilibrium *if*

(i) $p_i \geq 0$, for $i \in \mathcal{N}$,

(ii) $\{S^m\}_{m \in \mathcal{M}}$ is a feasible allocation, i.e., $S^k \cap S^m = \emptyset$, for $k \neq m$,

(iii) S^m is demanded by bidder m for all $m \in \mathcal{M}$, i.e., $v^m(S^m) - \sum_{i \in S^m} p_i \geq v^m(S) - \sum_{i \in S} p_i$ for any $S \subset \mathcal{N}$, and $m \in \mathcal{M}$,

(iv) $p_i = 0$ if $i \notin \cup_m S^m$.

Observe that conditions (ii) and (iii) suggest that bidders demand disjoint bundles of items. Conditions (i) and (iv) guarantee that prices are nonnegative, and the price of an item that is not

demanded is equal to zero. It can be shown (by aggregating inequality (ii) over all bidders) that the allocation $\{S^m\}$ associated with a Walrasian equilibrium is efficient. This suggests that auctions that terminate when a Walrasian equilibrium is reached can implement the efficient outcome. In this work we focus on such iterative auctions.

A Walrasian equilibrium need not exist for all classes of valuations. Hence, a prerequisite for the design of auctions that terminate at a Walrasian equilibrium is establishing its existence. A necessary and sufficient condition for the existence of a Walrasian equilibrium is given in Bikhchandani and Mamer (1997). They consider the following linear programming formulation of the efficient allocation problem (LP1) and its dual (DLP1):

$$
(LP1) \quad
\begin{aligned}
\max \ & \sum_m \sum_S x^m(S) v^m(S) \\
s.t. \ & \sum_S x^m(S) \leq 1 \quad \forall m \\
& \sum_m \sum_{S | i \in S} x^m(S) \leq 1 \quad \forall i \\
& x^m(S) \geq 0.
\end{aligned}
\qquad
(DLP1) \quad
\begin{aligned}
\min \ & \sum_i p_i + \sum_m \pi^m \\
s.t. \ & \pi^m \geq v^m(S) - \sum_{i \in S} p_i \quad \forall \ S, m \\
& p_i, \pi^m \geq 0.
\end{aligned}
$$

At integral feasible solutions of LP1, $x^m(S)$ takes value 1 if bundle S is assigned to bidder m. The first two constraints respectively ensure that each bidder m receives at most one bundle, and each item i can be present in at most one bidder's bundle. The objective function is the total value generated by an allocation of items according to $\{x^m(S)\}_m$. In DLP1, we have a variable p_i for each item i, which can be interpreted as the price of the relevant item. At optimal solutions of DLP1 it can be seen that π^m is equal to the maximum surplus of bidder m at prices $\{p_i\}$, i.e., $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$.

Bikhchandani and Mamer (1997) show that a Walrasian equilibrium exists if and only if LP1 has an optimal solution that is integral. This result follows by establishing that at integral feasible solutions of LP1 the complementary slackness conditions in LP1/DLP1 are identical to the Walrasian equilibrium conditions in Definition 2.2. Moreover, when LP1 has optimal solutions that are integral, the allocation $\{S^m\}_m$ with bundles S^m satisfying $x^m(S^m) = 1$, together with optimal prices $\{p_i\}_i$ constitute a Walrasian equilibrium. We employ these results in the next section to characterize the existence of a Walrasian equilibrium for graphical valuations. Subsequently, we use this characterization to develop efficient iterative auctions that terminate at this equilibrium.

3 Tree Valuations and Walrasian Equilibrium

In the rest of this paper, we focus on the special setting where the value functions satisfy the following assumptions:[6]

---

[6]We study the efficient allocation problem for general graphical valuations in our companion paper (Candogan et al., 2013).

Assumption 3.1 (Tree valuation). *Let $G = (\mathcal{N}, E)$ be the value graph associated with the graphical valuations of bidders. We assume that G is a tree graph, i.e., it contains no cycles.*

Assumption 3.2 (Sign-consistency). *For some $(i, j) \in E$ and $m \in \mathcal{M}$, if $w_{ij}^m > 0$, then $w_{ij}^k \geq 0$ for all $k \in \mathcal{M}$, and similarly if $w_{ij}^m < 0$, then $w_{ij}^k \leq 0$ for all $k \in \mathcal{M}$.*

The tree graph assumption imposes a natural hierarchy in the complementarity/substitutability structure: each node (item) can exhibit pairwise complementarity/substitutability only with its parent/children in the underlying tree.[7] The sign-consistency requirement, on the other hand, suggests that two items i and j are either substitutes or complements for all bidders. Note that this assumption still allows for the presence of both complementary and substitutable items in the set of items, but it prevents having two items as substitutes for some bidders and complements for the remaining ones. As an example of valuations that satisfy these assumptions, consider the spectrum auction in Figure 3, and assume that the auctioneer sells only a single spectrum band (band A) in multiple regions. In this case, the valuations can be represented using a tree graph (where nodes correspond to different geographical regions, and there is a natural hierarchy between the central node and the peripheral ones). Moreover, if spectra in the neighboring regions are complementary for a bidder (due to reduced interference), we expect them to be complementary for the remaining bidders as well. Thus, the sign-consistency assumption also holds.

In this section, we establish that a Walrasian equilibrium exists for sign-consistent tree valuations (Section 3.1). We also provide a compact LP formulation (that involves polynomially many variables and constraints in the number of items and bidders) of the efficient allocation problem, and show that Walrasian equilibria can be obtained through the solution of this LP and its dual. In Section 3.2, we show that the sign-consistency and tree graph assumptions are indispensable for the existence of a Walrasian equilibrium. For ease of exposition, we delegate the proofs to Appendix B. The compact LP formulation that we provide in this section plays a key role in subsequent sections for the design of efficient iterative auctions that terminate at a Walrasian equilibrium.

3.1   A Compact LP Formulation for Walrasian equilibria

In this section, we establish the existence of a Walrasian equilibrium for sign-consistent tree valuations. In order to obtain the existence result, we first provide a compact linear programming formulation (LP2) of the efficient allocation problem that exploits the graphical structure of bidders' valuations. We then establish that for sign-consistent tree valuations, this formulation has integral optimal solutions. This result allows for showing that LP1 (introduced in Section 2.2) also has integral optimal solutions, and implies the existence of a Walrasian equilibrium for sign-consistent tree valuations (Theorem 3.1). We conclude the section by showing that the compact LP formulation and its dual can be used to find the Walrasian equilibrium allocation and prices (Theorem 3.2).

---

[7]We note that the results presented in this paper still hold when the underlying graph is a forest, i.e., a collection of disjoint trees.

Before we provide the compact LP formulation of the efficient allocation problem, we state a related integer programming formulation (recall that $w_{ij}^m = 0$ for $(i,j) \notin E$):

$$\max \quad \sum_m \sum_i x_i^m w_i^m + \sum_m \sum_{i,j} y_{ij}^m w_{ij}^m$$

$$s.t. \quad \sum_m x_i^m \le 1, \qquad\qquad \forall i$$

(IP)

$$y_{ij}^m \le x_i^m, x_j^m, \qquad\qquad \forall m, i, j$$

$$x_i^m + x_j^m - 1 \le y_{ij}^m, \qquad\qquad \forall m, i, j$$

$$x_i^m \in \{0,1\}, y_{ij}^m \in \{0,1\}, \qquad \forall m, i, j.$$

In this formulation, the variable x_i^m takes the value 1 if item i is allocated to bidder m. The first constraint guarantees that each item is assigned to at most one bidder. The second and third constraints jointly imply that $y_{ij}^m = x_i^m x_j^m$, i.e., $y_{ij}^m = 1$ if and only if both i and j are assigned to bidder m. It can be immediately seen that every feasible solution $\{x_i^m, y_{ij}^m\}$ of this integer program corresponds to a feasible allocation $\{S^m\}$, where $S^m = \{i | x_i^m = 1\}$ (and vice versa). Moreover, since $y_{ij}^m = 1$ if and only if $x_i^m = x_j^m = 1$, for graphical valuations the associated objective value is the welfare generated by this allocation, i.e., $\sum_m v^m(S^m)$. Since every feasible solution corresponds to a feasible allocation (and vice versa) and the efficient allocations are defined as the feasible allocations that maximize welfare, it immediately follows that the optimal solutions of this integer program correspond to efficient allocations.

The LP relaxation of this integer program (obtained by replacing the constraints $x_i^m \in \{0,1\}, y_{ij}^m \in \{0,1\}$ with $x_i^m \in [0,1], y_{ij}^m \in [0,1]$), leads to a linear programming formulation of the efficient allocation problem (hereafter referred to as LP2):

$$\max \quad \sum_m \sum_i x_i^m w_i^m + \sum_m \sum_{i,j} y_{ij}^m w_{ij}^m$$

$$s.t. \quad \sum_m x_i^m \le 1, \qquad\qquad \forall i$$

$(LP2)$

$$y_{ij}^m \le x_i^m, x_j^m \qquad\qquad \forall m, i, j$$

$$x_i^m + x_j^m - 1 \le y_{ij}^m, \qquad\qquad \forall m, i, j$$

$$x_i^m \le 1, \qquad\qquad \forall m, i$$

$$0 \le x_i^m, y_{ij}^m \qquad\qquad \forall m, i, j.$$

In this formulation the constraint $y_{ij}^m \le 1$ is implied by the second and fourth constraints and hence is omitted.[8] Since integral solutions of this LP are also solutions of IP, it follows that these solutions correspond to feasible allocations, and their objective value is the associated welfare. Observe that

---

[8]In fact, the fourth constraint is also redundant, as it is implied by the first constraint and nonnegativity of x_i^m. This constraint is not omitted in order to ensure consistency with a related LP formulation $(LP - D^m)$ that we discuss in Section 4, where this constraint is necessary.

LP2 uses the graphical structure of the valuations (and employs decision variables x_i^m, y_{ij}^m that can be associated with the nodes and edges of the underlying graph) to obtain a compact formulation of the efficient allocation problem, i.e., it involves polynomially many variables and constraints (in the number of bidders and items).

In general this LP relaxation is not exact, i.e., the relaxation may have non-integer solutions that lead to a higher objective value than the optimal objective value of IP (in fact this is the case even for tree valuations that are not sign-consistent, see Section 3.2 for examples). Interestingly, we show in the following theorem that for sign-consistent tree valuations LP2 has optimal solutions that are integral. Moreover, this result has interesting consequences. In particular, it allows for establishing the existence of integral optimal solutions to LP1, and hence the existence of a Walrasian equilibrium.

Theorem 3.1.   *(i) For sign-consistent tree valuations LP2 has integral optimal solutions.*

(ii) If LP2 has integral optimal solutions, then so does LP1.

(iii) For sign-consistent tree valuations LP1 has integral optimal solutions, and a Walrasian equilibrium exists.

The first part of this result is established by exploiting the tree structure to obtain an alternative linear programming formulation of the efficient allocation problem, whose optimal objective value is equal to the welfare associated with the efficient allocation. This formulation relies on the fact that for trees the efficient allocation can be constructed recursively (by first identifying the efficient allocations in subtrees rooted at the children of a given node, and then using the allocations associated with subtrees to construct a solution to the efficient allocation problem for the tree that consists of all nodes), and has a similar structure to linear programs that can be used for solving dynamic programs (see Bertsekas (1996)). Under the sign-consistency assumption, we show that the optimal solutions of LP2 can be mapped to feasible solutions of this LP with the same objective value. These results imply that the objective value of LP2 is bounded from above by the welfare associated with the efficient allocation. Theorem 3.1(i) follows since IP, and hence LP2, have feasible integral solutions that correspond to the efficient allocation and have objective value equal to the associated welfare.

Note that while LP1 is applicable for all valuations, LP2 makes use of the graphical structure and is applicable only for graphical valuations. The second part of Theorem 3.1 suggests that these formulations are closely related, and the existence of integral solutions in the latter formulation implies the existence of such solutions in the former. The main idea behind the proof is to establish that for any feasible solution of LP1 it is possible to construct a feasible solution of LP2 with the same objective value, and conversely for any feasible integer solution of LP2 it is possible to construct a feasible integer solution of LP1, again with the same objective value (see Figure 4). These two facts immediately imply that when LP2 has an optimal solution that is integral, this solution leads to a (weakly) larger objective value than all feasible solutions of LP1. Moreover,

there exists a feasible integer solution of LP1 with the same objective value. Thus, this solution is optimal in LP1.
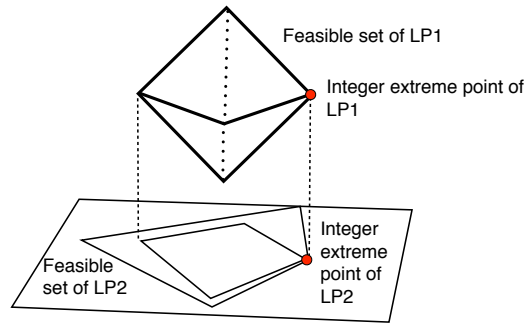

Figure 4: A feasible solution $\{x_i^m, y_{ij}^m\}$ to LP2 can be constructed from a feasible solution $\{x^m(S)\}$ of LP1 (by setting $x_i^m = \sum_{S|i\in S} x^m(S), y_{ij}^m = \sum_{S|ij\in S} x^m(S)$). These solutions have the same objective values in the corresponding optimization problems. Additionally, any feasible integer point $\{\hat{x}_i^m, \hat{y}_{ij}^m\}$ of LP2 corresponds to a feasible integer point $\{\hat{x}^m(S)\}$ of LP1 (where $\hat{x}^m(S) = 1$ only for $S = \{i|\hat{x}_i^m = 1\}$). Thus, if LP2 admits an optimal solution that is integral, then so does LP1.

Finally, the last part of Theorem 3.1 immediately follows from parts (i) and (ii), since a Walrasian equilibrium exists if and only if LP1 has integral optimal solutions (see Section 2). In the literature, the existence of a Walrasian equilibrium is established mainly in settings with no value complementarities (e.g., Gul and Stacchetti (1999)). Interestingly, our result implies that despite the fact that tree valuations exhibit both value complementarity and substitutability, when all bidders have sign-consistent tree valuations a Walrasian equilibrium always exists.[9]

We next focus on obtaining the Walrasian equilibrium allocation and prices using LP2. We start by presenting the dual of LP2:

$$
\begin{aligned}
\min \quad & \sum_{m,i} \pi_i^m + \sum_{m,i,j} p_{ij}^m + \sum_i p_i \\
(DLP2) \quad s.t. \quad & \pi_i^m \geq w_i^m - p_i + \sum_{j|j\neq i} q_{ij}^{m,i} - \sum_{j|j\neq i} p_{ij}^m \quad \forall m,i \\
& q_{ij}^{m,i} + q_{ij}^{m,j} - p_{ij}^m \geq w_{ij}^m \quad \forall m,i,j \\
& q_{ij}^{m,i}, p_{ij}^m, \pi_i^m, p_i \geq 0 \quad \forall m,i,j.
\end{aligned}
$$

In this formulation, the variables $p_i, q_{ij}^{m,i}, p_{ij}^m, \pi_i^m$ respectively correspond to the constraints $\sum_m x_i^m \leq 1, y_{ij}^m \leq x_i^m, x_i^m + x_j^m - 1 \leq y_{ij}^m$, and $x_i^m \leq 1$ of LP2. The variable p_i can be interpreted as the price of item i. Note that since both LP1 and LP2 have integral optimal solutions for sign-consistent tree

---

[9]A Walrasian equilibrium also exists for the class of gross substitutes, and its generalization gross substitutes and complements (see Gul and Stacchetti (1999) and Sun and Yang (2006)). We show in Appendix A that the class of graphical valuations that satisfy Assumptions 3.1 and 3.2 is not contained in these classes. Moreover, graphical valuations that also belong to these classes constitute a small subset of tree valuations, where the underlying graph consists of connected components of size at most two. Thus, our result here establishes the existence of a Walrasian equilibrium for a distinct and broad class of valuations.

valuations (by Theorem 3.1), the optimal objective values of DLP1 and DLP2 are both equal to the welfare associated with the efficient allocation. Since the objective of DLP2 can be expressed as $\sum_m (\sum_i \pi_i^m + \sum_{i,j} p_{ij}^m) + \sum_i p_i$, comparing the objective value of DLP2 with that of DLP1 (given by $\sum_m \pi^m + \sum_i p_i$), it follows that the quantity $\sum_i \pi_i^m + \sum_{i,j} p_{ij}^m$ can be interpreted as the surplus of bidder m (i.e., corresponds to π^m in DLP1). Intuitively, DLP2 associates variables with the nodes and edges of the underlying graph to capture bidders' surplus.

We next establish that the prices $\{p_i\}$ at an optimal solution of DLP2 and the allocation suggested by an integral optimal solution of LP2 constitute a Walrasian equilibrium. Since these LP formulations have polynomially many variables and constraints in the number of items and bidders, our result also implies that by solving these LP formulations a Walrasian equilibrium can be obtained in a computationally efficient way.

Theorem 3.2. *Assume that bidders have sign-consistent tree valuations.*

  (i) *Let $\{p_i\}$ be the prices that appear at an optimal solution of DLP2, and allocation $\{S^m\}$ be such that $S^m = \{i | x_i^m = 1\}$ for an integral optimal solution $\{x_i^m, y_{ij}^m\}$ of LP2. The prices $\{p_i\}$ and allocation $\{S^m\}$ constitute a Walrasian equilibrium.*

  (ii) *If the node and edge weights are integral, then a Walrasian equilibrium can be obtained in time polynomial in the number of bidders and items.*

The proof of the first part of this theorem follows by showing that complementary slackness between the integral optimal solutions of LP2 and optimal solutions of DLP2 imply the Walrasian equilibirum conditions given in Definition 2.2 for the aforementioned allocation and prices. The second part relies on the observation that LP2/DLP2 are LP formulations with polynomially many constraints and variables. Since such LP formulations can be efficiently solved (e.g., using the ellipsoid algorithm), the result follows from the first part of the theorem.

Remark: As explained in Section 2, the allocation-price pair that constitutes a Walrasian equilib-rium can be obtained through the optimal solutions of LP1 and DLP1 as well. However, these prob-lems have exponentially many variables/constraints (in the number of items), and unlike LP2/DLP2 in general they cannot be solved in a computationally efficient way.

3.2 Nonexistence of a Walrasian Equilibrium

In this section, we show that the characterization provided in the previous section is tight, in the sense that if Assumption 3.1 or Assumption 3.2 is relaxed, then a Walrasian equilibrium need not exist for graphical valuations. In particular, we provide examples in which only the tree (Example 3.1) or the sign-consistency assumption (Example 3.2) holds, and establish that in these examples a Walrasian equilibrium does not exist.

Example 3.1. *Consider a setting with bidders m, k and items i, j. Assume that the valuations of bidders are represented with the tree valuation in Figure 5, which is not sign-consistent.*
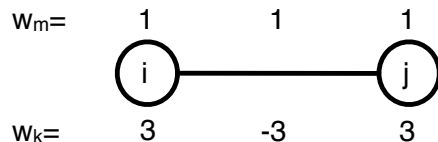
16

Figure 5: The weights for bidder m (k) is given above (below) the nodes/edges.

The optimal integer solutions of LP1 result in a total welfare of 4 (this can be obtained by choosing $x^m(\{i\}) = x^k(\{j\}) = 1$ and $x^l(S) = 0$, for remaining S and $l \in \{m, k\}$). On the other hand, consider the following solution of LP1: $x^m(\{i, j\}) = x^m(\emptyset) = 1/2$, $x^k(\{i\}) = x^k(\{j\}) = 1/2$, and $x^m(\{i\}) = x^m(\{j\}) = x^k(\{i, j\}) = x^k(\emptyset) = 0$. Feasibility of this solution in LP1 can be immediately checked, and the corresponding objective value is given by:

$$x^m(\{i, j\})v^m(\{i, j\}) + x^m(\emptyset)v^m(\emptyset) + x^k(\{i\})v^k(\{i\}) + x^k(\{j\})v^k(\{j\}) = \frac{3 + 0 + 3 + 3}{2} = 4.5.$$

Note that the objective value associated with this solution is larger than the objective of the optimal integer solution of LP1. Hence, LP1 does not have integral optimal solutions. Since a Walrasian equilibrium exists if and only if LP1 has integral optimal solutions, it follows that a Walrasian equilibrium does not exist for this example.

Example 3.2. We next focus on the 3-cycle example in Figure 6, which satisfies the sign-consistency assumption but not the tree assumption.[10] For this example, it can be seen that the optimal integer solution of LP1 results in an objective value of 1. On the other hand, $x^1(\{AB\}) = x^2(\{BC\}) = x^3(\{CA\}) = 1/2$ and $x^1(\emptyset) = x^2(\emptyset) = x^3(\emptyset) = 1/2$, is a feasible solution of LP1 with objective value $\frac{3}{2}$. Hence, LP1 does not have an optimal solution that is integral, and a Walrasian equilibrium does not exist.
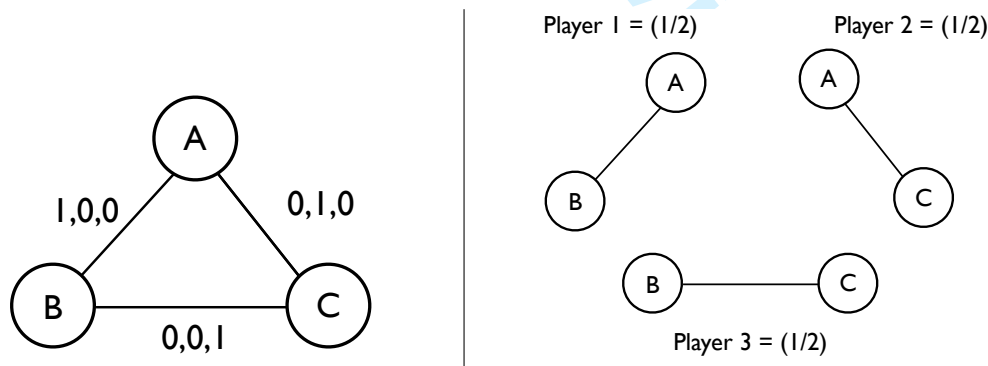
Figure 6: A value graph for three items is given on the left. We assume that there are three bidders, and node weights are equal to zero for all bidders. The labels associated with the edges designate the edge weights for different bidders. A feasible (fractional) solution to LP1 with larger objective value than any feasible integer solution of LP1 is given on right.

[10]This example generalizes to any graph containing a k-cycle (and any $k > 2$). See Candogan (2013) for details.

17

These examples suggest that both the tree and the sign-consistency assumptions are indispensable for the existence of a Walrasian equilibrium. Hence, it is not possible to design iterative auctions that terminate at this outcome for more general graphical valuations. This implies that different termination conditions or pricing rules are needed for such valuations. We discuss iterative auction design using more general pricing rules in our companion paper Candogan et al. (2013).

4 An Iterative Algorithm for Obtaining a Walrasian Equilibrium

In this section, we obtain an iterative algorithm for the solution of LP2/DLP2, which for sign-consistent tree valuations terminates with a Walrasian equilibrium in finite time. This algorithm is based on primal-dual algorithms, and relies on setting prices for different items and updating them according to bidders' demand at the given prices. In Section 5, we use this algorithm for iterative auction design, and establish (after complementing the algorithm with an appropriate payment rule) that it is an equilibrium for bidders to truthfully report their demand in these auctions.

We start this section by providing an overview of primal-dual algorithms (Section 4.1). In Section 4.2, we show how they can be applied to the solution of the efficient allocation problem LP2, and provide intuition for the associated dual variable updates. In particular, we establish that the algorithm increases the prices (i.e., variables $\{p_i\}$ in DLP2) of overdemanded items, and decreases those of underdemanded ones at each update. This algorithm uses the parameters of LP2/DLP2 (such as node and edge weights) in dual variable updates. In Section 4.3 we show that it is possible to slightly modify the algorithm and run it relying solely on bidders' demand reports (and not on the node/edge weights). Moreover, the resulting algorithm converges to a Walrasian equilibrium in finite time. We exploit this structure in the next section for the design of iterative auctions. The proofs of the results of this section are delegated to Appendix C.

4.1 Primal-dual Algorithm

In this section, we provide an overview of primal-dual algorithms, and introduce a slight generalization of these algorithms which we subsequently use for iterative auction design. The results of this section follow immediately from the existing literature on primal-dual algorithms (for instance Papadimitriou and Steiglitz (1998); Bertsimas and Tsitsiklis (1997); Bazaraa et al. (2011); Vohra (2011)), and are presented without proofs.

We start by focusing on a generic primal linear program (left) and its dual (right):

$$
\begin{array}{ll}
\max_{y} & d^T y \\
s.t. & Ay \leq a \\
& By \leq b \\
& y \geq 0,
\end{array}
\qquad
\begin{array}{ll}
\min_{\lambda,\mu} & a^T \lambda + b^T \mu \\
s.t. & A^T \lambda + B^T \mu \geq d \\
& \lambda, \mu \geq 0.
\end{array}
\tag{1}
$$

In these formulations, $a \in \mathbb{R}^{m_1}, b \in \mathbb{R}^{m_2}, d \in \mathbb{R}^n$ and $A \in \mathbb{R}^{m_1 \times n}, B \in \mathbb{R}^{m_2 \times n}$ are given, and $y \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}^{m_1}, \mu \in \mathbb{R}^{m_2}$ are respectively the decision variables for the primal and the dual. We denote the element in the ith row and jth column of matrix A by A_{ij}, and the ith element of a given vector y by y_i (similarly for other vectors and matrices). The feasible primal-dual pair $(y^*, [\lambda^*, \mu^*])$ is optimal if and only if it satisfies the complementary slackness (hereafter referred to as CS) conditions (Bertsimas and Tsitsiklis, 1997):

$$\lambda_i^* > 0 \rightarrow \sum_j A_{ij} y_j^* = a_i , \quad \mu_i^* > 0 \rightarrow \sum_j B_{ij} y_j^* = b_i,$$
$$y_j^* > 0 \rightarrow \sum_i A_{ij} \lambda_i^* + \sum_i B_{ij} \mu_i^* = d_j. \tag{2}$$

The high-level idea behind primal-dual algorithms is to start with a dual feasible solution $[\lambda^*, \mu^*]$, and check if there exists a primal feasible solution that satisfies the CS condition with the given dual solution. If this is the case, optimal solutions to both problems are found and the algorithm terminates. Otherwise $[\lambda^*, \mu^*]$ is not optimal, and the primal-dual algorithm iteratively updates the dual solution to a dual feasible solution with improved objective value.

The CS conditions can be checked using the *restricted primal/dual problems* (hereafter RP/RD):

$$\text{(RP)} \qquad\qquad\qquad\qquad\qquad \text{(RD)}$$

$$\min_{\gamma, h, y} \quad \sum_i \gamma_i + \sum_i h_i$$

$$s.t. \quad \sum_j A_{ij} y_j + \gamma_i - h_i = a_i \qquad \forall i | \lambda_i^* > 0 \qquad \max_{\bar{\lambda}, \bar{\mu}} \quad - a^T \bar{\lambda} - b^T \bar{\mu}$$

$$\sum_j A_{ij} y_j + \gamma_i - h_i \leq a_i \qquad \forall i | \lambda_i^* = 0 \qquad s.t. \quad \sum_i A_{ij} \bar{\lambda}_i + \sum_i B_{ij} \bar{\mu}_i \geq 0$$

$$\sum_j B_{ij} y_j = b_i \qquad \forall i | \mu_i^* > 0 \qquad\qquad\qquad \forall j | \sum_i A_{ij} \lambda_i^* + \sum_i B_{ij} \mu_i^* = d_j$$

$$\sum_j B_{ij} y_j \leq b_i \qquad \forall i | \mu_i^* = 0 \qquad -1 \leq \bar{\lambda}_j \leq 1 \qquad \forall j$$

$$y_j = 0 \qquad \forall j | \sum_i A_{ij} \lambda_i^* + \sum_i B_{ij} \mu_i^* > d_j \qquad \bar{\lambda}_i \geq 0 \qquad \forall i | \lambda_i^* = 0$$

$$y, \gamma, h \geq 0. \qquad\qquad\qquad\qquad\qquad \bar{\mu}_i \geq 0 \qquad \forall i | \mu_i^* = 0.$$

RD is the dual of the optimization problem obtained after omitting the variables in RP that are set equal to zero (i.e. variables that appear in the fifth constraint of RP). For $\gamma = h = 0$, the first and third constraints of RP correspond to the CS conditions in (2), while the remaining constraints guarantee that the component y of the solution of RP is feasible in the original primal problem. The variables γ, h capture the deviation from the CS conditions (associated with the first constraint of the original LP), and the objective of RP is to minimize the aggregate deviation.[11] In the standard version of primal-dual algorithms (e.g., Papadimitriou and Steiglitz (1998)), such variables are

---

[11]RP can equivalently be formulated by replacing the objective with min $\sum_i |z_i|$, where $z_i = \gamma_i - h_i$.

introduced for all constraints of the original LP (i.e., $m_2 = 0$), and this guarantees that RP is always feasible.[12] On the other hand, we allow for introducing such variables only for a subset of the constraints, since this structure allows for obtaining simpler formulations of RP/RD (with fewer variables/constraints), and plays an important role in our iterative auction design approach (as discussed in the next section).

It can be seen that for a given dual feasible solution $[\lambda^*, \mu^*]$ the optimal objective of RP is zero if and only if the optimal solution (y^*, γ^*, h^*) of RP is such that $\gamma^* = h^* = 0$. Moreover, in this case feasibility in RP implies that $(y^*, [\lambda^*, \mu^*])$ satisfy the CS conditions in (2), and hence are optimal in the original primal/dual problems (1). On the other hand, if the optimal objective of RP is strictly positive, then the optimal solution $[\bar{\lambda}, \bar{\mu}]$ of RD is an *improvement direction*, i.e., $a^T \lambda^* + b^T \mu^* > a^T(\lambda^* + \epsilon\bar{\lambda}) + b^T(\mu^* + \epsilon\bar{\mu})$ and $[(\lambda^* + \epsilon\bar{\lambda}), (\mu^* + \epsilon\bar{\mu})]$ is feasible in the original dual problem for sufficiently small $\epsilon > 0$. This follows since by strong duality we have $a^T\bar{\lambda} + b^T\bar{\mu} < 0$, and the (first, third, and fourth) constraints of RD (which are satisfied by $[\bar{\lambda}, \bar{\mu}]$) guarantee feasibility of $[(\lambda^* + \epsilon\bar{\lambda}), (\mu^* + \epsilon\bar{\mu})]$ for sufficiently small ϵ. Thus, we conclude that by solving RP/RD the CS conditions can be tested, and either an optimal solution to the original primal/dual problems, or an improvement direction can be obtained.

Given the improvement direction, the dual solution $[\lambda^*, \mu^*]$ is updated to $[\lambda^* + \theta^*\bar{\lambda}, \mu^* + \theta^*\bar{\mu}]$ using the *primal-dual stepsize* of $\theta^* \triangleq \max\{\theta \geq 0 | [\lambda^* + \theta^*\bar{\lambda}, \mu^* + \theta^*\bar{\mu}]$ is feasible in the dual$\}$. In other words, the dual is updated by taking the largest step in the improvement direction that preserves feasibility. The algorithm terminates when the dual feasible solution $[\lambda^*, \mu^*]$ is such that the corresponding formulation of RP has an optimal solution (y^*, γ^*, h^*) with objective value zero, and returns $(y^*, [\lambda^*, \mu^*])$ at termination. The improvement direction and stepsize defined here guarantee that termination occurs in finite time and $(y^*, [\lambda^*, \mu^*])$ are optimal respectively in the original primal-dual problems in (1):

Proposition 4.1. *If RP is feasible after each dual update, then the primal-dual algorithm terminates in finite time with optimal solutions of the optimization problems in* (1).

The proof follows by examination of the finite-convergence proofs of primal-dual algorithms given in Papadimitriou and Steiglitz (1998); Bazaraa et al. (2011), and is omitted.[13] The key step in the proof is to show that after dual updates the optimal solution of RP corresponds to a different extreme point of the polytope obtained by replacing the equality constraints in RP with inequality constraints. Hence, the extreme point corresponding to an optimal solution of the original primal problem is found in finitely many updates.

Applications of primal-dual algorithms to auction design can be found in the literature (Bikhchandani et al., 2002; Parkes, 2006; Mishra and Parkes, 2007; De Vries et al., 2007; Vohra, 2011). How-

---

[12]In the standard version of primal-dual algorithms it is also assumed that the original primal problem (1) is in standard form (i.e., consists of equality and nonnegativity constraints), and $a, b \geq 0$. Consequently, restricted primal can be formulated using only γ_i variables (i.e., setting $h_i = 0$). Here we allow for inequality constraints in the original primal, as LP2 involves such constraints, and provide a more general RP.

[13]In these works, the γ_i, h_i variables are associated with all constraints of the original LP, and hence RP is always feasible. Their analysis immediately extends to the case where such variables are associated with only some constraints, provided that RP remains feasible after each dual update.

ever, the existing auction formats either disallow complementarities or rely on complex bundle pricing rules. In subsequent sections, we use primal-dual algorithms with LP2/DLP2, and develop simple efficient iterative auction formats that are applicable when complementarities are present.

4.2 A Primal-Dual Algorithm for the Efficient Allocation Problem

In this section, we discuss how the primal-dual algorithm can be used for the solution of the efficient allocation problem LP2, and characterize the corresponding dual update direction and the stepsize. In particular, we show that the algorithm increases (decreases) the prices of overdemanded (underdemanded) items, and chooses the smallest stepsize which causes a bidder to demand a new bundle. Throughout this section, we use the shorthand notations $\pi = \{\pi_i^m\}_{m \in \mathcal{M}, i \in \mathcal{N}}$, $p = \{p_i\}_{i \in \mathcal{N}}$, $p_E = \{p_{ij}^m\}_{m \in \mathcal{M}; i, j \in \mathcal{N}}$, $q_E = \{q_{ij}^{m,i}\}_{m \in \mathcal{M}; i, j \in \mathcal{N}}$ to compactly represent the solutions of DLP2.

Formulating RP/RD and Obtaining an Improvement Direction: We start by providing the CS conditions (in LP2/DLP2) associated with a given feasible solution (p, π, p_E, q_E) of DLP2:

$$\sum_m x_i^m = 1 \qquad \forall i \mid p_i > 0, \quad \text{and} \quad C^m \quad \forall m, \tag{3}$$

where for all $m \in \mathcal{M}$, C^m represents the following set of constraints:

$$C^m \begin{cases} y_{ij}^m = x_i^m & \forall i, j \mid q_{ij}^{m,i} > 0 \\ x_i^m + x_j^m - 1 = y_{ij}^m & \forall i, j \mid p_{ij}^m > 0 \\ x_i^m = 1 & \forall i \mid \pi_i^m > 0 \\ y_{ij}^m = 0 & \forall i, j \mid q_{ij}^{m,i} + q_{ij}^{m,j} - p_{ij}^m > w_{ij}^m \\ x_i^m = 0 & \forall i \mid \pi_i^m > w_i^m - p_i + \sum_j q_{ij}^{m,i} - \sum_j p_{ij}^m. \end{cases}$$

Similarly we use the shorthand notation F^m to represent the feasibility constraints other than $\sum_i x_i^m \leq 1$ (in LP2) that are not implied by the CS conditions:

$$F^m \begin{cases} y_{ij}^m \leq x_i^m & \forall i, j \mid q_{ij}^{m,i} = 0 \\ x_i^m + x_j^m - 1 \leq y_{ij}^m & \forall i, j \mid p_{ij}^m = 0 \\ x_i^m \leq 1 & \forall i \mid \pi_i^m = 0 \\ y_{ij}^m \geq 0 & \forall i, j \mid q_{ij}^{m,i} + q_{ij}^{m,j} - p_{ij}^m = w_{ij}^m \\ x_i^m \geq 0 & \forall i \mid \pi_i^m = w_i^m - p_i + \sum_j q_{ij}^{m,i} - \sum_j p_{ij}^m. \end{cases}$$

With slight abuse of notation, we denote by $F^m(k)$ the set of (i, j) (only i in the case of the third constraint) for which the kth constraint is present in F^m, e.g., $F^m(1) = \{(i, j) | q_{ij}^{m,i} = 0\}$, and $F^m(2) = \{(i, j) | p_{ij}^m = 0\}$.

Using this notation we next formulate RP/RD that can be used to check the CS conditions (3).

21

$$\underline{\text{(RP)}} \hspace{6cm} \underline{\text{(RD)}}$$

$$\min_{\{x_i^m, y_{ij}^m, \gamma_i, h_i\}} \quad \sum_i \gamma_i + \sum_i h_i$$

$$\max_{\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E} \quad -\sum_i \bar{p}_i - \sum_{m,i,j} \bar{p}_{ij}^m - \sum_{m,i} \bar{\pi}_i^m$$

$$s.t. \quad \sum_m x_i^m + \gamma_i - h_i = 1 \quad \forall i | p_i > 0$$

$$s.t. \quad \bar{\pi}_i^m \geq -\bar{p}_i + \sum_j \bar{q}_{ij}^{m,i} - \sum_j \bar{p}_{ij}^m \quad \forall m, i, j | (i,j) \in F^m \, (5)$$

$$\sum_m x_i^m + \gamma_i - h_i \leq 1 \quad \forall i | p_i = 0$$

$$\bar{q}_{ij}^{m,i} + \bar{q}_{ij}^{m,j} - \bar{p}_{ij}^m \geq 0 \quad \forall m, i, j | (i,j) \in F^m \, (4)$$

$$C^m, F^m \quad \forall m \in \mathcal{M}$$

$$\bar{\pi}_i^m \geq 0 \quad \forall m, i | i \in F^m \, (3)$$

$$\gamma_i, h_i \geq 0 \quad \forall i.$$

$$\bar{p}_{ij}^m \geq 0 \quad \forall m, i, j | (i,j) \in F^m \, (2)$$

$$\bar{q}_{ij}^{m,i} \geq 0 \quad \forall m, i, j | (i,j) \in F^m \, (1)$$

$$\bar{p}_i \geq 0 \quad \forall i | p_i = 0$$

$$-1 \leq \bar{p}_i \leq 1 \quad \forall i.$$

In this formulation, following the approach in Section 4.1, we associate γ_i, h_i variables only with the first constraint of LP2 (hence they appear only in the first two constraints in RP), while we impose the remaining CS and feasiblity constraints C^m, F^m. We next introduce a special property (Property 4.1) of the solutions of DLP2, and establish (Lemma 4.1) that when this property is satisfied the RP/RD given above are feasible, and they can be used to check the CS conditions and identify an improvement direction.

Property 4.1 (Acceptable dual solutions)**.** *The dual feasible solution (p, π, p_E, q_E) of DLP2 is acceptable if there exists a primal solution to LP2 that satisfies the constraints C^m, F^m for all m.*

Lemma 4.1. *Assume that an acceptable dual solution (p, π, p_E, q_E) of DLP2 is given. Then, the following are true:*

(i) RP has feasible solutions.

(ii) If the optimal objective value of RP is zero, then (p, π, p_E, q_E) is optimal in DLP2, and the restriction of the optimal solutions of RP to $\{x_i^m, y_{ij}^m\}_{m,i,j}$ gives optimal solutions of LP2.

(iii) If the optimal objective value of RP is not zero, then no feasible solution of LP2 satisfies the CS conditions, and the optimal solutions of RD are dual improvement directions in DLP2.

The proof of the first part immediately follows since Property 4.1 implies that solutions to RP that satisfy C^m, F^m always exist, and choosing γ_i, h_i appropriately the remaining constraints of RP can be satisfied. The second and third part follow by noting that the optimal objective value of RP is zero only for $\gamma = h = 0$, and such optimal solutions of RP satisfy the CS conditions (3), i.e., both the conditions in C^m, and the CS condition $\sum_m x_i^m = 1$ for $p_i > 0$. This observation implies that RP has objective value zero if and only if a feasible solution of LP2 satisfies CS conditions with (p, π, p_E, q_E), and $\gamma = h = 0$ together with this solution of LP2 constitute an optimal solution to RP. When the objective value of RP is strictly positive, by strong duality the optimal solution of RD suggests a direction that improves the dual objective in DLP2, moreover the constraints

of RD suggest that this is a feasible update direction. The details are omitted since they follow immediately from Property 4.1 and the CS conditions in LP2/DLP2.

Stepsize Selection and the Primal-Dual Algorithm: We next focus on the stepsize associated with a dual update direction obtained from an optimal solution of RD, and establish that this stepsize leads to updated dual solutions that are acceptable in DLP2. Additionally, we show that the updated solution and stepsize can be characterized in terms of bidders' demand, and complete the description of our primal-dual algorithm.

We begin by providing a linear programming formulation that can be used to identify the demanded bundles, and subsequently exploit it for characterizing the stepsize and properties of dual updates. Recall that a bundle S^* is demanded by bidder m if $v^m(S^*) - \sum_{i \in S^*} p_i = \max_S v^m(S) - \sum_{i \in S} p_i$. Given prices $\{p_i\}$, the following primal/dual problems (that have a similar structure to LP2/DLP2) can be used for identifying such bundles:

$$\underline{(LP - D^m)}$$

$$\max_{\{x_i^m, y_{ij}^m\}} \quad \sum_i x_i^m(w_i^m - p_i) + \sum_{i,j} y_{ij}^m w_{ij}^m$$

$$s.t. \quad y_{ij}^m \leq x_i^m, x_j^m \quad \forall i,j$$
$$x_i^m + x_j^m - 1 \leq y_{ij}^m, \quad \forall i,j$$
$$x_i^m \leq 1 \quad \forall i$$
$$0 \leq x_i^m, y_{ij}^m \quad \forall i,j.$$

$$\underline{(DLP - D^m)}$$

$$\min_{\{\pi_i^m, p_{ij}^m, q_{ij}^{m,i}\}_{i,j}} \quad \sum_i \pi_i^m + \sum_{ij} p_{ij}^m$$

$$s.t. \quad \pi_i^m \geq (w_i^m - p_i) + \sum_{j|j \neq i} q_{ij}^{m,i} - \sum_{j|j \neq i} p_{ij}^m \quad \forall i$$
$$q_{ij}^{m,i} + q_{ij}^{m,j} - p_{ij}^m \geq w_{ij}^m \quad \forall i,j$$
$$q_{ij}^{m,i}, p_{ij}^m, \pi_i^m \geq 0 \quad \forall i,j.$$

In these optimization problems bidder m and prices $\{p_i\}$ are fixed. The decision variables in $LP - D^m$ are $\{x_i^m, y_{ij}^m\}_{i,j}$, and those in $DLP - D^m$ are $\{\pi_i^m, p_{ij}^m, q_{ij}^{m,i}\}_{i,j}$ (note that the decision variables are present only for bidder m). Given any bundle S, a feasible solution to $LP - D^m$ can be constructed by setting $x_i^m = y_{ij}^m = 1$ for $i,j \in S$ (and setting remaining variables equal to zero). Moreover, it can be checked that the corresponding objective value is the surplus of the aforementioned bundle, i.e., $v^m(S) - \sum_{i \in S} p_i$.

$LP - D^m$ can be viewed as a special case of LP2, with only a single bidder m, and node/edge weights $\{w_i^m - p_i, w_{ij}^m\}_{i,j}$. This observation implies that for sign-consistent tree valuations, this optimization problem has optimal solutions that are integral (see Theorem 3.1). Consider an integral optimal solution of $LP - D^m$, and let $S^m = \{i | x_i^m = 1\}$. Primal feasibility implies that $y_{ij}^m = 1$ if and only if $i,j \in S^m$. Observe that the objective value associated with this solution is equal to the surplus of bidder m for bundle S^m, i.e., $\sum_{i \in S^m}(w_i^m - p_i) + \sum_{i,j \in S^m} w_{ij}^m = v^m(S^m) - \sum_{i \in S^m} p_i$. Since $LP - D^m$ has a feasible solution associated with any bundle and the corresponding objective value is the surplus of this bundle, it follows that the bundle S^m associated with an optimal solution of $LP - D^m$ maximizes bidder m's surplus. That is, the integral optimal solutions of $LP - D^m$ are associated with the set of items bidder m demands.

Optimal solutions of $LP - D^m$ are closely related to those of RP. Let a dual feasible solution (p, π, p_E, q_E) of DLP2, and the corresponding optimal solution $\{\gamma_i, h_i, x_i^m, y_{ij}^m\}_{m,i,j}$ of RP be given.

Consider the restriction of these solutions respectively to $\{\pi_i^m, p_{ij}^m, q_{ij}^{m,i}\}_{i,j}$ and $\{x_i^m, y_{ij}^m\}_{i,j}$. Since RP is constructed by imposing C^m for all m, it follows that the optimal solution of RP satisfies these constraints. It can be checked from $LP - D^m/DLP - D^m$ that the CS conditions the optimal solutions of these problem satisfy are identical to the aforementioned constraints, i.e., they are captured by C^m. Thus, it follows that the solutions $\{x_i^m, y_{ij}^m\}_{i,j}$ and $\{\pi_i^m, p_{ij}^m, q_{ij}^{m,i}\}_{i,j}$ are optimal respectively in $LP - D^m/DLP - D^m$. Since optimal solutions of $LP - D^m$ identify the bundles bidder m demands, this suggests that the optimal solution of RP (which has $\{x_i^m, y_{ij}^m\}_{i,j}$ as one of its components) aggregates the demand information from all bidders. Moreover, the objective value of RP is equal to the violation of the first constraint when $\{x_i^m, y_{ij}^m\}_{i,j}$ captures the demand of each bidder m.

We show in Lemma 4.2 that the relation between RP and $LP - D^m$ generalizes to the duals of these optimization problems. In particular, optimal solutions of RD (or dual update directions) can be obtained in terms of those of $DLP - D^m$. Moreover, this relation can be further exploited to characterize the corresponding stepsize for dual updates in terms of bidders' demand.[14]

Lemma 4.2. *Let (p, π, p_E, q_E) denote an acceptable solution of DLP2, and D^m denote the set of bundles bidder $m \in \mathcal{M}$ demands at price vector p. Assume that an optimal solution of RD associated with this dual feasible solution is $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$, and the corresponding objective value is nonzero. Define $\theta^* = \min\{\theta_1, \theta_2\}$, where*

$\theta_1 = \min\{\theta \geq 0 |\ S \notin D^m$ *enters the demand set of some bidder m at prices $p + \theta\bar{p}\}$, and*

$\theta_2 = \min\{\theta \geq 0 | p_i + \theta\bar{p}_i = 0$ *for some i such that $\bar{p}_i < 0\}$.*

Then the following are true:

(i) θ^* is nonzero and bounded, i.e., $0 < \theta^* < \infty$.

(ii) For every bidder m consider an optimal solution $(\tilde{\pi}^m, \tilde{p}_E^m, \tilde{q}_E^m)$ of $DLP - D^m$ at the price vector $p + \theta^*\bar{p}$. Let $\tilde{\pi} = \{\tilde{\pi}^m\}_m$, $\tilde{p}_E = \{\tilde{p}_E^m\}_m$, and $\tilde{q}_E = \{\tilde{q}_E^m\}_m$. The tuple $(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$, where $(\hat{\pi}, \hat{p}_E, \hat{q}_E) = ((\tilde{\pi}, \tilde{p}_E, \tilde{q}_E) - (\pi, p_E, q_E))/\theta^*$, is an optimal solution of RD.

(iii) θ^* is the primal-dual stepsize associated with this solution.

(iv) The updated solution $(p, \pi, p_E, q_E) + \theta^*(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is acceptable in DLP2.

This lemma implies that a dual update can be made by first obtaining an optimal solution $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ to RD. Then, the primal-dual stepsize θ^* can be chosen as the smallest update step, where a bidder starts demanding a new bundle (or the price of an item becomes zero), when prices are updated in the $\bar{p}$ direction. Observe that this choice of the stepsize relies only on $\bar{p}$ and not on the remaining components of the dual optimal solution $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$. Moreover, the lemma suggests

---

[14] The relation between $LP - D^m/DLP - D^m$ and RP/RD is a consequence of associating the γ_i, h_i variables with only the first constraint of LP2 in the formulation of RP, and leads to the stepsize/dual update direction provided in Lemma 4.2.

that there exists an optimal solution $(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ of RD (potentially different from $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$) such that θ^* together with this optimal solution constitute a valid dual improvement direction and a primal-dual stepsize (at dual feasible solution (p, π, p_E, q_E)). Part (iv) of the lemma implies that the dual updates defined by this lemma are acceptable in DLP2. Thus, Lemma 4.1 holds and RP/RD can be used to check the CS conditions and identify an improvement direction after each dual update. This observation leads to a primal-dual algorithm (see Algorithm 1) that uses the update direction/stepsize given in Lemma 4.2 for dual updates.

Algorithm 1 A centralized algorithm for the solution of LP2/DLP2

S0 (Initialize): Start with an acceptable solution (p, π, p_E, q_E).

S1 (Find Improvement Direction): Formulate RP/RD. Solving RD identify an update direction $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$.

   If the objective value of this problem is equal to zero, then go to Step S3. Otherwise go to Step S2.

S2 (Update Dual): Compute θ^* given in Lemma 4.2. For all $m \in \mathcal{M}$, solve $DLP - D^m$ at price vector $p + \theta^* \bar{p}$. Denote the optimal solution of $DLP - D^m$ by $(\tilde{\pi}^m, \tilde{p}_E^m, \tilde{q}_E^m)$, and let $(\hat{\pi}^m, \hat{p}_E^m, \hat{q}_E^m) = ((\tilde{\pi}^m, \tilde{p}_E^m, \tilde{q}_E^m) - (\pi^m, p_E^m, q_E^m))/\theta^*$.

   Update the dual solution to $(p, \pi, p_E, q_E) + \theta^*(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$. Go to Step S1.

S3 (Terminate): Terminate returning the current dual solution (p, π, p_E, q_E), and the component $\{x_i^m, y_{ij}^m\}_{m,i,j}$ of an integral optimal solution of RP.

Step S3 is reached and Algorithm 1 terminates, when in Step S1 RP has objective value zero. Thus, Lemma 4.1 implies that at termination $\{x_i^m, y_{ij}^m\}_{m,i,j}$ and (p, π, p_E, q_E) are respectively optimal in LP2/DLP2. To see that RP has an integral optimal solution at Step S3, note that LP2 has an integral optimal solution for sign-consistent tree valuations (Theorem 3.1) and RP has a corresponding optimal solution (obtained by choosing $\{x_i^m, y_{ij}^m\}_{m,i,j}$ as in the optimal solution of LP2, and $\gamma = h = 0$).[15] Moreover, the allocation suggested by this solution (given by $S^m = \{i | x_i^m = 1\}$ for all m) and the price vector p (which is also obtained at Step S3, and is a component of an optimal solution of DLP2) constitute a Walrasian equilibrium (Theorem 3.2). Since primal-dual algorithms converge in finite time and Algorithm 1 employs the primal-dual update direction and stepsize (as established in Lemma 4.2) for the solution of LP2/DLP2, it immediately follows from Proposition 4.1 that Algorithm 1 terminates in finite time with optimal solutions of LP2 and DLP2 and identifies Walrasian equilibrium allocation/prices. Observe that this algorithm relies on the knowledge of bidders valuations (to formulate the sets of constraints C^m, F^m and optimization problems RP/RD), and corresponds to the algorithm outlined in Figure 1a.

---

[15]Feasibility of this solution in RP is immediate since optimal solutions of LP2 and DLP2 satisfy the CS conditions, and hence the aforementioned solution satisfies all constraints of RP with $\gamma = h = 0$.

Intuition on price updates: We conclude this section by providing intuition on the price up-dates employed in our primal-dual algorithm. If at an optimal solution of RP, $h_i > 0$, then the first two constraints of RP imply that $\sum_m x_i^m > 1$. Since, optimal solutions of RP correspond to optimal solutions of $LP - D^m$ for all m (as explained before Lemma 4.2) and the solution $\{x_i^m, y_{ij}^m\}_{i,j}$ of the latter characterizes demanded bundles, it follows that if $h_i > 0$ then item i belongs to demand sets of multiple bidders. In this case we say that item i is *overdemanded*. Similarly, if at an optimal solution $\gamma_i > 0$, then we say that i is *underdemanded*. Note that RP may have multiple optimal solutions, but it is never the case that for some item i, $h_i > 0$ in some optimal solutions, and $\gamma_i > 0$ in the others. This is because, when $h_i > 0$ at an optimal solution, the complementary slackness conditions suggest that in the optimal solutions of RD the constraint $\bar{p}_i \le 1$ is active. Conversely if $\gamma_i > 0$ at an optimal solution, then the constraint $\bar{p}_i \ge -1$ is active. Since at a given dual optimal solution at most one of these inequalities can be active, an item can be either overdemanded or underdemanded but not both. These observations immediately imply the following lemma:

Lemma 4.3. *Let (p, π, p_E, q_E) be a feasible solution of DLP2 that is not optimal, and $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ be a corresponding optimal solution of RD. The following are true:*

(i) *An item cannot be overdemanded with respect to one optimal solution of RP and underde-manded with respect to another one.*

(ii) *If i is overdemanded, then $\bar{p}_i = 1$.*

(iii) *If i is underdemanded, then $\bar{p}_i = -1$.*

This lemma suggests that Algorithm 1 updates the prices in an intuitive way: At each stage the prices of overdemanded items are increased, and those of underdemanded items are decreased. Moreover, it can be seen from the algorithm that the price update at a given direction continues until a bidder starts demanding a new bundle, or the price of an item decreases to zero.

4.3   Distributed Implementation of the Primal-Dual Algorithm

Algorithm 1 relies on finding an improvement direction through a solution of RD (Step S1), and computing the primal-dual stepsize θ^* (Step S2). In the auction design setup we consider, the value function (or node/edge weights) of a bidder is private information. This observation has two important implications for Algorithm 1 (when the weight information is not readily available): (i) RP/RD cannot be formulated as the weights are required for the construction of constraints C^m, F^m, (ii) the stepsize θ^* cannot be obtained as its computation requires knowing when bidders' demand sets change (see Lemma 4.2), which in turn relies on the knowledge of weights.

In this section, we establish that despite these issues, the auctioneer can solve LP2 by appropri-ately modifying Algorithm 1. In particular, we show in Section 4.3.1 that after each dual update C^m, F^m can be constructed through bidders' demand sets, which in turn allows for formulating RP/RD and obtaining an improvement direction. In Section 4.3.2, we show that using sufficiently small update steps the auctioneer can identify when demand sets of bidders change (for a given

price update direction). Hence, the stepsize θ^* can also be obtained by using only the demand information. We modify Algorithm 1 using the update direction and stepsize described in Sections 4.3.1 and 4.3.2, and provide a new algorithm in Section 4.3.3. This algorithm does not require the knowledge of weights or valuations, and instead runs by setting prices for items and adjusting them according to bidders' demand at the given prices. Moreover, it terminates in finite time with Walrasian equilibrium allocation and prices (and corresponds to the algorithm outlined in Figure 1b). We present our results in this section under an additional assumption: Value functions of bidders are integer-valued.[16] The results of this section are subsequently used in Section 5 to obtain an efficient iterative auction.

4.3.1 Finding an Improvement Direction

In this section, we show that the dual update direction given in Step S1 of Algorithm 1 can be identified by using bidders' demand information even when information about their valuations/weights is not available. In the algorithm, the update direction is obtained by solving RP/RD. Thus, in order to obtain this update direction the first step is to formulate these optimization problems.

Formulating RP/RD: It can be seen from RP/RD that for a given dual feasible solution of DLP2, in order to formulate these problems it is necessary and sufficient to construct the constraints C^m and F^m for every $m \in \mathcal{M}$. Note that F^m involves inequality constraints that are imposed for the active constraints in the dual feasible solution of DLP2, while C^m imposes equality constraints for inactive ones. Hence, it follows that these constraints can be specified using the active constraints of the feasible solution of DLP2. Algorithm 1 and Lemma 4.2 suggest that the updated dual solution of DLP2 coincides with the optimal solution of $DLP - D^m$ (for all m) at the updated prices. It can be checked that $DLP - D^m$ and DLP2 share the same set of constraints (for a given price vector p and bidder m), and hence it follows that after each update the active constraints of the new feasible solution of DLP2 can be identified in terms of those of $DLP - D^m$. These observations imply that in order to formulate RP/RD, it suffices to have each bidder m solve $DLP - D^m$, and report to the auctioneer the constraints that are active in this solution.

Interpretation in terms of demand: Let a price vector p be given, and $\{\pi_i^m, p_{ij}^m, q_{ij}^{m,i}\}_{i,j}$ denote a corresponding optimal solution of $DLP - D^m$. Using the active constraints in this solution, the

---

[16]This assumption is commonly made in the context of iterative auction design for establishing convergence of iterative auctions in finite time. See for instance, (De Vries and Vohra, 2003; Ausubel, 2004; Mishra and Parkes, 2007; Bikhchandani et al., 2011; Vohra, 2011).

CS conditions (in $LP - D^m/DLP - D^m$) can be stated as follows:

$$
\begin{aligned}
(i) && \pi_i^m &> w_i^m - p_i + \sum_j q_{ij}^{m,i} - \sum_j p_{ij}^m &\rightarrow&\quad x_i^m = 0 \\
(ii) && q_{ij}^{m,i} + q_{ij}^{m,j} - p_{ij}^m &> w_{ij}^m &\rightarrow&\quad y_{ij}^m = 0 \\
(iii) && \pi_i^m &> 0 &\rightarrow&\quad x_i^m = 1 \\
(iv) && q_{ij}^{m,i} &> 0 &\rightarrow&\quad x_i^m = y_{ij}^m \\
(v) && p_{ij}^m &> 0 &\rightarrow&\quad x_i^m + x_j^m - 1 = y_{ij}^m
\end{aligned}
\tag{4}
$$

Since $\{\pi_i^m, p_{ij}^m, q_{ij}^{m,i}\}_{i,j}$ is an optimal solution of $DLP - D^m$, complementary slackness implies that these conditions provide an alternative characterization of the optimal solutions of $LP - D^m$: A solution of $LP - D^m$ is optimal if an only if it is feasible (which can be checked without the knowledge of weights), and satisfies the conditions (i)-(v). Moreover, since optimal integral solutions of $LP - D^m$ correspond to demanded bundles (see Section 4.2), this characterization also constitutes an alternative representation of all demanded bundles. In particular, the first constraint in (4) suggests that item i never belongs to a bundle bidder m demands, whereas the third one suggests that it always does. The second constraint suggests that at most one of items i and j can belong to the demand set (since if $x_i^m = x_j^m = 1$ primal feasibility requires $y_{ij}^m = 1$), whereas the fifth one suggests that at least one of items i and j can belong to the demand set (since if $x_i^m = x_j^m = 0$, the equality in the fifth constraint cannot hold). Finally, the fourth constraint suggests that if item i belongs to the demand set then so does item j. Since at full generality the number of bundles a bidder can demand is exponential in the number of items N (for instance consider the case where all prices and valuations are equal to zero), whereas these bundles can be reported through $O(N^2)$ constraints of the type (i)-(v), specifying the latter allows for compactly stating demand. Hence, we refer to the collection of such constrains as *a compact representation of demand*, and say that a bidder *compactly reports her demand* if she reports this collection to the auctioneer.

We conclude that the auctioneer can identify an update direction by first asking bidders to compute their demand at the given prices (or solve $LP - D^m/DLP - D^m$), and report it compactly (by specifying which of the constraints (i)-(v) hold). Then, she can use this information to formulate and solve RP/RD and obtain an update direction.[17]

Remark: Note that the compact representation of demand sets given above also corresponds to a natural logical expression for identifying such sets. Let a_i be the logical symbol that is equal to one if item i belongs to the demand set, a_i^C denote the complement of a_i, and $\wedge, \vee, \implies$ denote respectively the "and", "or", and "implies" operators. The logical expressions corresponding to each of the CS conditions in (4) can be given as follows: (i) a_i^C, (ii) $(a_i \wedge a_j)^C$, (iii) a_i, (iv) $a_i \implies a_j$,

---

[17]Alternatively, bidders can report all bundles that they demand (i.e., potentially exponentially many bundles). This allows for identifying all optimal solutions of $LP - D^m$. Strict complementary slackness guarantees that the equalities at the right hand side of (i)-(v) hold for all such solutions if and only if there exists a dual optimal solution where the corresponding inequalities are strict. Thus, it is possible to characterize all constraints that are active at a dual optimal solution of $DLP - D^m$ by using the set of all demanded bundles.

(v) $a_i \vee a_j$. A given set is demanded if and only if it satisfies all such expressions derived from the CS conditions (or the expression obtained from combining all of the aforementioned expressions with "and" operators). This suggests that the set of all demanded bundles can be characterized using the logical expressions stated above.

4.3.2 Choosing the Stepsize

For a given price update direction, Algorithm 1 chooses the stepsize such that either a new bundle is demanded by a bidder, or the price of an item decreases to zero as a result of the update with this stepsize (see Lemma 4.2). While the auctioneer can check how large of a stepsize can be taken until the price of an item decreases to zero, the same is not true for the change in demand sets if the valuations of bidders are not available. In this section, we establish that when valuations are integral, it is possible to discover the stepsize where a bidder starts demanding a new bundle (without explicitly knowing the valuations), and use this for running the primal-dual algorithm.

We start by providing a subroutine, which given a price vector p and a price update direction $\bar{p}$, returns a stepsize θ and bidders' demands at the updated price vector $p + \theta\bar{p}$. We establish subsequently in Lemma 4.4 that when $\bar{p}$ is chosen with respect to the improvement direction obtained from an optimal solution of RD, the subroutine terminates in finite time with the primal-dual stepsize θ^*, i.e., the stepsize θ at termination guarantees that at prices $p + \theta\bar{p}$, either a bidder starts demanding a new bundle, or the price of an item decreases to zero.

In this subroutine, initially θ is increased by $1/N$ (S1) provided that such an update does not lead to negative prices for some items (the definition of θ_2 and minimization in S1 imply that if such updates lead to a negative price, then a smaller stepsize is used). After this update, bidders' demand sets $\{\hat{D}^m\}$ at price vector $p + \theta\bar{p}$ are obtained. If a bidder demands none of the bundles (D^m) she demanded at the original price vector p (S1a), then (as established in the proof of Lemma 4.4) θ is greater than the primal-dual stepsize θ^* ($\theta > \theta^*$). In this case, the subroutine corrects the prices by "stepping back" by an amount of $\hat{\theta}$, using the update direction $\bar{p}$ (S2). On the other hand, if S1a does not hold, and S1b or S1c holds the subroutine terminates. In the former case, a bidder starts demanding a new bundle (in addition to some of the previously demanded bundles). In the latter case, either $\theta = \theta_2$ and the price of an item decreases to zero, or the subroutine only modifies (increases) the prices of items (in set I) that are not demanded by any bidder. Lemma 4.4 shows that when these termination conditions are met, θ is equal to the primal-dual stepsize θ^*. Finally, if S1d holds, then $\theta < \theta^*$ and θ is increased repeating Step S1, until a termination condition is met.

Lemma 4.4. *(i) Consider a price vector p, and update direction $\bar{p}$ such that $|\bar{p}_i| \leq 1$. Assume that prices are updated from p to $p + \theta\bar{p}$ by using a stepsize $\theta \leq 1/N$. Let S_1, S_2 respectively denote the sets bidder m demands before and after the price update (where $S_1 = S_2 = S$ if some bundle S is demanded both before and after the price update), and π_1, π_2 denote the*

Subroutine Stepsize computation subroutine.

S0 *(Initialize):* Let $p \geq 0$ denote the initial price vector, D^m denote the associated demand sets for bidders $m \in \mathcal{M}$, and $\bar{p}$ denote the price update direction. Set $\theta_2 = \min\{\theta \geq 0 | p_i + \theta\bar{p}_i = 0$ for some i such that $\bar{p}_i < 0\}$, and $I = \{i | p_i > 0\}$.

S1 *(Price/Demand update):* Let $\theta := \min\{\theta + 1/N, \theta_2\}$, and ask each bidder m the set of bundles $\hat{D}^m$ that she demands at price vector $p + \theta\bar{p}$.

   S1a *(Stepsize too large?):* If $\hat{D}^m \cap D^m = \emptyset$ for some m, then go to Step S2.

   S1b *(New bundle demanded?):* Else if $\hat{D}^m - D^m \neq \emptyset$ for some m, then go to Step S3.

   S1c *(Price zero / Unbounded increase ?):* Else if $\theta = \theta_2$, or $\theta_2 = \infty$ and $S \cap I = \emptyset$ for all m, $S \in \hat{D}^m$, then go to Step S3.

   S1d *(Stepsize too small?):* Otherwise, go to Step S1.

S2 *(Stepping back):* Let $S_1 = \arg\min_{S \in D^m} \sum_{i \in S} \bar{p}_i$, and $S_2 = \arg\max_{S \in \hat{D}^m} \sum_{i \in S} \bar{p}_i$. Define $\Delta = \sum_{i \in S_1} \bar{p}_i - \sum_{i \in S_2} \bar{p}_i$, and set

$$\hat{\theta} = \frac{\left(\sum_{i \in S_1}(p_i + \theta\bar{p}_i) - \sum_{i \in S_2}(p_i + \theta\bar{p}_i)\right) - \lfloor\sum_{i \in S_1}(p_i + \theta\bar{p}_i) - \sum_{i \in S_2}(p_i + \theta\bar{p}_i)\rfloor}{\Delta}.$$

Update θ to $\theta := \theta - \hat{\theta}$. If $\theta \leq 0$, then set $\theta = 0$, update the price vector to p, and go to Step S3. Otherwise go to Step S1.

S3 *(Terminate):* Terminate returning θ, and $\hat{D}^m$ for all m.

---

associated maximum surplus. Then,

$$\pi_2 - \pi_1 = \left(\sum_{i \in S_1} p_i - \sum_{i \in S_2}(p_i + \theta\bar{p}_i)\right) - \left\lfloor\sum_{i \in S_1}(p_i + \theta\bar{p}_i) - \sum_{i \in S_2}(p_i + \theta\bar{p}_i)\right\rfloor.$$

(ii) *Assume that a dual feasible solution of DLP2 with price vector p is given, and let $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ denote an associated optimal solution of RD. The subroutine (initialized with p and $\bar{p}$) terminates with the primal-dual stepsize θ^* (given in Lemma 4.2) after finitely many iterations.*

The first part of this lemma is established by using the integrality of valuations to show that for sufficiently small price updates, the change in the surplus of bidders can be tracked. In order to prove the second part of the lemma, we first establish that the termination condition in Step S2 (i.e., $\theta \leq 0$) and the condition $\theta_2 = \infty$ and $S \cap I =$ for all $S \in \hat{D}^m$ in Step S1c never hold if bidders' true demand information is available (these conditions are present to guarantee finite termination in settings where bidders can misreport their demand). Then, we exploit the result of the first part of the lemma to show that the update in Step S2 of the algorithm makes bidder m indifferent between a bundle she originally demanded, and another one she starts demanding after the price update in Step S1. Hence, after Step S2 is completed, the algorithm satisfies the

condition of S1b and terminates. If Step S2 is not reached, then θ increases until either the price of an item decreases to zero (the termination condition $\theta = \theta_2$ in Step S1c), or a bidder starts demanding a new bundle (S1b). In both cases the subroutine terminates, and Lemma 4.2 implies that the primal-dual stepsize θ^* is returned.

Observe that in the subroutine price updates are only a function of bidders' demand sets, the current price vector p, and the price update direction $\bar{p}$. Lemma 4.4 implies that using only these quantities, the auctioneer can update the prices as suggested by Lemma 4.2. That is, after finitely many price/demand updates suggested by the subroutine, the prices that emerge are identical to those that are employed by the primal-dual algorithm (as shown in Lemma 4.2).

4.3.3 Convergence to a Walrasian Equilibrium

We conclude this section by providing a new algorithm (Algorithm 2), and establishing in Theorem 4.1 that for sign-consistent tree valuations it converges to the efficient outcome and Walrasian equilibrium prices. This algorithm is almost identical to Algorithm 1, but it does not rely on dual

---

Algorithm 2 A distributed algorithm for the solution of LP2/DLP2

S0 (Initialize): Start with $p = 0$.

S1 (Find Improvement Direction): Compute a compact representation of the set of bundles D^m each bidder m demands at prices p. Using this formulate RP/RD. Solving RD identify an update direction $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$.

If the objective value of this problem is equal to zero, then go to Step S3. Otherwise go to Step S2.

S2 (Update Prices): Update prices/demanded bundles using the stepsize computation subroutine (initialized with the price vector p and update direction $\bar{p}$). Go to Step S1.

S3 (Terminate): Terminate returning the prices p, and an allocation suggested by an integral optimal solution of RP, i.e., an allocation $\{S^m\}$ such that $S^m = \{i | x_i^m = 1\}$.

---

variables other than the price vector p. Moreover, it chooses the improvement direction and stepsize as presented in Sections 4.3.1 and 4.3.2. In particular, it uses the stepsize computation subroutine for updating the prices. Consequently, it solves LP2/DLP2 by setting prices, and adjusting them according to bidders' demand at the given prices, but without using any other information about bidders' valuations. Moreover, this algorithm terminates with a Walrasian equilibrium since (as established in Section 3) the prices and allocation suggested by optimal solutions of these optimization problems correspond to Walrasian equilibria.

Theorem 4.1. *Assume bidders have sign-consistent tree valuations. Then, Algorithm 2 terminates with an efficient outcome and Walrasian equilibrium prices in finite time.*

The proof of this theorem follows by showing that the price updates in Algorithm 2 coincide with those in Algorithm 1 (i.e., an application of the primal-dual algorithm to a solution of LP2/DLP2).

This allows for establishing that Algorithm 2 terminates in finite time, and the allocation/prices obtained at termination coincide with the allocation suggested by an integral optimal solution of LP2 and the prices suggested by an optimal solution of DLP2 (since this is the case for Algorithm 1). On the other hand, Theorem 3.2 suggests that such optimal solutions of LP2/DLP2 correspond to Walrasian equilibrium allocation/prices, and hence the claim of Theorem 4.1 follows. Our result implies that when bidders truthfully reveal their demand (and have sign-consistent tree valuations), the auctioneer can use Algorithm 2 to update prices and ensure convergence to a Walrasian equilibrium even when she does not know the bidders' valuations.

5   An Iterative Auction for Tree Valuations

It was established in the previous section that when bidders truthfully reveal their demand and have sign-consistent tree valuations, the auctioneer can use Algorithm 2 to implement the efficient outcome. In this section, we consider implementing the efficient allocation with strategic bidders (who can misreport their demand) through a dynamic iterative auction.

We start this section by introducing the solution concept (ex-post perfect equilibrium) that we use for analyzing the outcome of iterative auctions (Section 5.1). We establish that if the auctioneer charges payments to bidders according to the prices that emerge at the end of Algorithm 2, then it is not an ex-post perfect equilibrium for bidders to truthfully report their demand in this algorithm. In Section 5.2, by slightly modifying the price updates of Algorithm 2, and complementing them with appropriate final payment rules, we provide a new iterative auction format ("interleaved tree auction"). In this auction the final payments of bidders correspond to their VCG payments (defined in Section 5.1), and these payments are computed by identifying the Walrasian equilibria for markets that consist of all bidders, and all bidders but one. This auction uses a novel price update structure that interleaves the demand queries for different markets to obtain such Walrasian equilibria. Moreover, we establish that for sign-consistent tree valuations it implements the efficient allocation at an ex-post perfect equilibrium where bidders truthfully report their demand. The auction relies on a simple pricing rule (anonymous item pricing), and allows bidders to report their demand in a compact way. These desirable features suggest that even in settings with value complementarities it is possible to implement the efficient outcome using a simple auction format. As in the previous section, we obtain our results under the assumption that the value functions of bidders are integer-valued, and delegate the proofs to Appendix D.

5.1   Ex-post Perfect Equilibrium

In an iterative auction, bidders participate in a multi-stage incomplete information game, where bidders do not know their opponents' valuations. We denote the history of bids revealed until step t of the auction by H_t. Consider a bidder m, whose valuation is v^m. The strategy of this bidder assigns an action to each history H_t from the set of allowable actions, $\Sigma^m(H_t)$, bidder m can use after this history. We denote this strategy by $s^m(v^m)$, and the action associated with history H_t

by $s^m(H_t, v^m) \in \Sigma^m(H_t)$.

For a realization of valuations of bidders $\{v^k\}_k$, the payoff bidder m receives at the end of the auction game is denoted by $u^m(s^m(v^m), s^{-m}(v^{-m})|v^m)$, where $s^{-m}(v^{-m})$ denotes the strategies of all bidders but m. Similarly, we denote by $u^m(s^m(v^m), s^{-m}(v^{-m})|H_t, v^m)$, the payoff bidder m, who is of type v^m, receives by using strategy $s^m(v^m)$ after history H_t, given that her opponents use strategies $s^{-m}(v^{-m})$. Using this notation, we next introduce a solution concept that is employed in the literature for the analysis of iterative auctions (e.g., see Ausubel (2004, 2006)):

Definition 5.1 (Ex-post perfect equilibrium). *A strategy profile $s = \{s^k\}$ is an ex-post perfect equilibrium, if after any history H_t, it satisfies*

$$u^m(s^m(v^m), s^{-m}(v^{-m})|H_t, v^m) \geq u^m(z^m, s^{-m}(v^{-m})|H_t, v^m), \tag{5}$$

for any valuations $\{v^k\}$ of bidders, bidder m, and strategy z^m.

This definition suggests that a strategy profile is an ex-post perfect equilibrium, if for any valuations of her opponents and after any history, given strategies of her opponents, no agent has incentive to deviate from her strategy. In other words, after any realization of the history (H_t) and valuations (v^{-m}), the given strategy profile remains a Nash equilibrium of the induced subgame, where types of agents are public knowledge.[18]

A natural payment rule for iterative auctions involves charging bidders the total prices of the items that they acquire at the end of the auction. As discussed in Section 2, in this paper we focus on iterative auctions that terminate when a Walrasian equilibrium is reached. We next illustrate that truthful demand revelation may not be an ex-post perfect equilibrium, if bidders are charged the Walrasian equilibrium prices associated with the items that they acquire in such auctions.

Example 5.1. *Consider a two bidder auction where there are two items $\{i, j\}$ and bidders have sign-consistent tree valuations with the following weights:*

	w_i	w_j	w_{ij}
Bidder m	*6*	*4*	*-4*
Bidder k	*5*	*5*	*0*

---

[18]This solution concept is closely related to the well-known solution concepts perfect Bayesian equilibrium (PBE) and dominant strategy equilibrium. The former solution concept involves a belief system that determines the posterior probability for each agent's valuation (given the history), and after any history requires agents to best respond to their opponents with respect to these beliefs. It can be seen that the ex-post perfect equilibrium concept is stronger than PBE in that it requires the actions bidders choose after any history to be a best response to others' actions with respect to the true valuations. In this sense, this equilibrium concept is similar to the dominant strategy equilibrium, which does not involve beliefs, but requires the strategy of an agent to be a best response to any strategy profile that can be chosen by her opponents. The dominant strategy equilibrium concept, on the other hand, is usually defined for single-shot games, and does not involve a perfection step. A natural analogue of ex-post perfect equilibrium for such settings is the *ex-post equilibrium* concept, which imposes the requirements of Definition 5.1 only at H_0, i.e., the starting point of the auction. Dominant strategy equilibrium is more restrictive than this solution concept in that it requires an agent's strategy to be a best response to *any* strategy of her opponents (not just the equilibrium strategy). Hence, if a strategy profile is a dominant strategy equilibrium then it is an ex-post equilibrium as well.

At the efficient outcome, item i is assigned to bidder m, and item j is assigned to bidder k. Hence, at Walrasian equilibria bidder m demands item i and bidder k demands item j. Walrasian equilibrium price of item i is at least 5 units (as otherwise this item also belongs to the demand set of bidder k). This suggests that the price of item j is at least 3 units (as otherwise this item belongs to the demand set of bidder m). Thus, if bidders bid truthfully, and the final payments are given by the Walrasian equilibrium prices, the payoff of bidder k is at most 2 units.

Assume that instead of her true valuation, bidder k (starting from time 0) bids with respect to the following value function $(w_i^k, w_j^k, w_{ij}^k) = (1, 1, 0)$. It can be seen that in this case the efficient allocation (and Walrasian equilibrium allocation) is exactly the same as before. However, at the associated Walrasian equilibrium the price of item j is at most 1 units, in order to ensure that bidder k demands this item. However, by acquiring item j at price 1, bidder k guarantees a payoff of 4 units. Thus, we conclude that by not bidding according to her true valuations (and misreporting her demand in the auction), bidder k can improve her payoffs.

This example suggests that a different payment rule is necessary for ensure that bidders truthfully report their demand in iterative auctions. We conclude this section by providing such a payment rule, which we subsequently use in Section 5.2 to obtain an iterative auction that implements the efficient outcome at an ex-post perfect equilibrium.

Definition 5.2 (VCG Mechanism)**.** *Consider a collection of value functions $\{v^m\}$. A mechanism (mapping from types/valuations to allocations and payments) is called a VCG (Vickrey - Clarke - Groves) mechanism if it*

- *chooses an efficient allocation, i.e., $\{S^m\} \in \arg\max_{\{Z^m\}|Z^m \subset \mathcal{N}, Z^k \cap Z^l = \emptyset} \sum_m v^m(Z^m)$*

- *assigns each agent m a payment $\gamma^m(\{S^k\}, \{v^k\}_{k \neq m}) = h^m(v^{-m}) - \sum_{k \neq m} v^k(S^k)$, where h^m is any real-valued function.*

If h^m is such that $h^m(v^{-m}) = \max_{\{Z^k\}|Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k)$, then we say that payments of bidders $(\gamma^m(\{S^k\}, \{v^k\}_{k \neq m}) = \max_{\{Z^k\}|Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) - \sum_{k \neq m} v^k(S^k))$ are *VCG payments with the Clarke pivot rule*. Observe that since $\{S^k\}$ is efficient we have $\sum_{k \neq m} v^k(S^k) = \max_{\{Z^k\}|Z^k \cap Z^l = Z^k \cap S^m = \emptyset} \sum_{k \neq m} v^k(Z^k)$. Intuitively, this suggests that the VCG payments capture the opportunity cost a given agent m creates on the rest of the system by acquiring bundle S^m, i.e., the difference between the maximum welfare that can be achieved by the remaining agents and the welfare those agents have when bidder m receives bundle S^m. In this paper, we only employ VCG payments with the Clarke pivot rule, and for simplicity refer to these payments as VCG payments. Additionally, for any bundle S we refer to the quantity $\max_{\{Z^k\}|Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) - \max_{\{Z^k\}|Z^k \cap Z^l = Z^k \cap S = \emptyset} \sum_{k \neq m} v^k(Z^k)$ as the VCG payment of agent m associated with bundle S.

In multi-item sealed-bid auctions (which are single-shot games), charging VCG payments to bidders guarantees that the efficient outcome can be implemented at a dominant strategy equilibrium (Nisan et al., 2007; Krishna, 2009). In the next section, we provide an iterative auction that assigns these payments to bidders at termination, and prove that it implements the efficient

outcome at an ex-post perfect equilibrium. Note that there is a major difficulty in using these payments in iterative auction design: obtaining them requires the computation of a function of valuations, which are privately known to the bidders (Definition 5.2). Our iterative auction overcomes this difficulty by using a novel price update structure that allows for determining the VCG payments.

5.2 Interleaved Tree Auction

In this section, we provide an iterative auction format, which we refer to as the *interleaved tree auction*. This auction implements the efficient outcome at an ex-post perfect equilibrium for sign-consistent tree valuations. It accomplishes this by following a novel (interleaved) price update structure which converges to a Walrasian equilibrium (see Figure 1c) and ensures that bidders' final payments correspond to the VCG payments (that guarantee that bidders have no incentive to misreport their demand). Moreover, the auction relies on a simple pricing rule, and compact demand responses.

Finding VCG payments requires computing a nontrivial function of bidders' valuations (Definition 5.2). On the other hand, bidders' valuations are private information, and hence our auction needs to "learn" the relevant information for computing these payments. In this section, we first establish that using Walrasian equilibria for sets of bidders $\mathcal{M}$, as well as $\mathcal{M} - \{m\}$ for all $m \in \mathcal{M}$ the VCG payments can be computed (see Lemma 5.1). Then, we obtain the interleaved tree auction by providing a modification of Algorithm 2 that identifies such Walrasian equilibria and uses them for the computation of VCG payments.

We say that the *market clears for bidders* $\hat{\mathcal{M}} \subset \mathcal{M}$, if the given prices and bundles $\{S^m\}_{m \in \hat{\mathcal{M}}}$ demanded by bidders $m \in \hat{\mathcal{M}}$ constitute a Walrasian equilibrium (i.e., satisfy the conditions of Definition 2.2 for bidders $m \in \hat{\mathcal{M}}$). We denote the market that consist of bidders $\mathcal{M} - \{m\}$ by $\mathcal{E}_m$, and the one that consists of bidders $\mathcal{M}$ by $\mathcal{E}_\emptyset$. Our next result establishes that the VCG payments can be obtained in terms of the prices and bidders' surplus at such market clearance points.

Lemma 5.1. *Assume that markets $\mathcal{E}_m$ and $\mathcal{E}_\emptyset$ clear respectively at prices p^1 and p^2. Let $\{S_1^k\}_{k \in \mathcal{M} - \{m\}}$ and $\{S_2^k\}_{k \in \mathcal{M}}$ denote the corresponding market clearing allocations, and π_1^k and π_2^k denote the maximum surplus of bidder $k \in \mathcal{M}$ at these prices. The VCG payment of bidder m associated with bundle S_2^m is equal to $\sum_{k \neq m} \left(\pi_1^k + \sum_{i \in S_1^k} p_i^1 \right) - \sum_{k \neq m} \left(\pi_2^k + \sum_{i \in S_2^k} p_i^2 \right)$.*

Observe that by definition surplus satisfies $\pi_j^k = v^k(S_j^k) - \sum_{i \in S_j^k} p_i^j$, and hence $v^k(S_j^k) = \pi_j^k + \sum_{i \in S_j^k} p_i^j$ for $j \in \{1, 2\}$. Since Walrasian equilibrium allocations are efficient we also have that $\{S_1^k\}_{k \in \mathcal{M} - \{m\}}$ and $\{S_2^k\}_{k \in \mathcal{M}}$ respectively maximize welfare for markets $\mathcal{E}_m$ and $\mathcal{E}_\emptyset$, and the claim immediately follows from the definition of VCG payments.

We next provide the *interleaved tree auction* (Algorithm 3) that clears markets $\{\mathcal{E}_m\}_m \cup \{\mathcal{E}_\emptyset\}$ and implements the efficient outcome for sign-consistent tree valuations by keeping track of the change in bidders' surplus, and using it for charging VCG payments to bidders.

35

Algorithm 3 Interleaved Tree Auction

S0 (Initialize): Start with $p = 0$, and $\mathcal{S} = \{\mathcal{E}_m\}_{m \in \mathcal{M}}$. Set $q = 0$, and $z^m = 0$ for all $m \in \mathcal{M}$.

S1 (Find Improvement Direction): Ask each bidder m to compactly report the set of bundles D^m she demands at price vector p. Using this formulate and solve RP/RD for each market $\mathcal{E} \in \mathcal{S}$, and denote the objective value associated with market $\mathcal{E}$ by $e(\mathcal{E})$. Update $\mathcal{S} := \mathcal{S} - \{\mathcal{E} | e(\mathcal{E}) = 0\}$, and $z^m := z^m + \sum_i p_i$ if $\mathcal{E}_m$ is removed from $\mathcal{S}$.

   S1a (All markets cleared): If $\mathcal{S} = \emptyset$, then go to step S3.

   S1b (Some markets not cleared): Otherwise, let $\mathcal{E}_{m^*} \in \mathcal{S}$ be the market for which $e(\mathcal{E}) > 0$ is the smallest (break ties lexicographically), and $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ denote the associated update direction (obtained solving RD). Go to step S2.

S2 (Update Prices): Update prices/demanded bundles using the stepsize computation subroutine (initialized with market $\mathcal{E}_{m^*}$, price vector p, and update direction $\bar{p}$). At each update of the price vector from $p(t)$ to $p(t+1)$ in the subroutine (i.e., in Steps S1 and S2 of the subroutine):

   – *(Charge payments)* Charge each bidder $m \in \mathcal{M}$ a payment of $\sum_i |p_i(t+1) - p_i(t)|$, and update $q := q + \sum_i |p_i(t+1) - p_i(t)|$.
   – *(Compute rebate)* Let $\Delta \pi^k$ denote the change in surplus of bidder k as a result of the price update, given by Lemma 4.4 (i.e., equal to $\pi_2 - \pi_1$ in the lemma). For all m such that $\mathcal{E}_m \notin \mathcal{S}$ update $z^m := z^m - \sum_{k \neq m} \Delta \pi^k$.

   Update $p := p + \theta \bar{p}$, and go to Step S1.

S3 (Repeat for market $\mathcal{E}_\emptyset$): Set $\mathcal{S} = \{\mathcal{E}_\emptyset\}$, and repeat Steps S1-S2. When the condition of Step S1a holds again denote the allocation suggested by an integral optimal solution of RP by $\{S^m\}$, and set $z^m := z^m - \sum_{i \notin S^m} p_i$ for all $m \in \mathcal{M}$. Go to Step S4.

S4 (Terminate): Terminate assigning bidder $m \in \mathcal{M}$ bundle S^m, and a rebate of $q - z^m$.

Steps S1-S2 of this auction focus on clearing the markets $\{\mathcal{E}_m\}_m$. The prices and the set of markets that are not cleared at a given iteration are respectively denoted by p and $\mathcal{S}$. Given the prices, the auctioneer first collects compact demand responses (following (4)) from the bidders and solves the corresponding RP/RD for markets in $\mathcal{S}$ (S1). The set of uncleared markets $\mathcal{S}$ is updated if for some markets the objective value of these optimization problems is equal to zero (and hence a Walrasian equilibrium can be identified). The optimal solution of RD, or the dual update direction, may be different for different markets in $\mathcal{S}$. The auctioneer identifies an uncleared market $\mathcal{E}_{m^*}$ for which the objective value of RP/RD (denoted by $e(\mathcal{E})$ for market $\mathcal{E}$) is the smallest, and chooses the update direction accordingly (Step S1b). Intuitively, the objective value of RP/RD captures how close a market is to clearance (or the violation of CS conditions in the associated LP2/DLP2 as explained in Section 4), and Algorithm 3 updates the prices greedily using the direction obtained from the market that is closest to clearance.

The prices are updated in Step S2 according to the update direction given in Step S1, by using

the stepsize computation subroutine. Updates in this direction rely on a stepsize smaller than $1/N$ (as described in Step S1 of the subroutine), and continue until either the prices of some items decrease to zero, or some bidder in market $\mathcal{E}_{m^*}$ starts demanding a new bundle.[19] Steps S1-S2 terminate by identifying the market clearance points of all markets in $\{\mathcal{E}_m\}_m$ since choosing the update direction according to the market whose RP has the lowest objective value ($e(\mathcal{E})$), ensures that after each update this objective value decreases. Consequently, after each update the lowest objective value of RP for an uncleared market decreases. Since this quantity cannot decrease indefinitely (objective value of RP is lower bounded by zero), it follows that these steps eventually terminate by clearing markets $\{\mathcal{E}_m\}_m$.

After steps S1-S2 are completed, similar steps are repeated to clear the market $\mathcal{E}_\emptyset$ (step S3), starting with the price vector at which the last market in $\{\mathcal{E}_m\}_m$ clears. It is possible to merge steps S1-S2 with step S3, and clear $\mathcal{E}_\emptyset$ together with the remaining markets. However, handling this market separately ensures that the auction terminates at the Walrasian equilibrium of $\mathcal{E}_\emptyset$ (hence the efficient allocation), by assigning each bidder a bundle that she demands at the final prices. Note that bidders receive their items only after the last step is completed and the auction terminates.

Whenever prices are updated from a price vector $p(t)$ to $p(t+1)$ by (Step S1 or S2 of) the stepsize computation subroutine, the auction charges all bidders a payment $\sum_i |p_i(t+1) - p_i(t)|$ (Step S2). Charging bidders such payments ensure that they do not adopt a strategy that prevents the termination of the auction indefinitely. The quantity q captures the sum of all such payments made by a bidder throughout the auction. The parameter z^m, on the other hand, is initially updated to the sum of the prices at which market $\mathcal{E}_m$ clears (Step S1). After this market clears, the auctioneer adjusts the z^m parameter according to the change in the total surplus of all the bidders in $\mathcal{E}_m$ whenever the prices are updated (Step S2). Note that since the subroutine uses price updates that are smaller than $1/N$, the associated change in the surplus of bidders can be computed as in Lemma 4.4. Finally, when market $\mathcal{E}_\emptyset$ clears, z^m is adjusted by subtracting the total price of all items that do not belong to the bundle of bidder m at the associated allocation (Step S3). At the end of the auction, each bidder m receives a rebate that is equal to her total payment so far minus the change in the total surplus/prices for the market she is not present at, i.e., a rebate of $q - z^m$.[20] Observe that this suggests that the total payment of bidder m is given by z^m provided that the auction terminates. The updates of the z^m variable ensure that this variable captures the change in the total surplus of all bidders but m between time instants where markets $\mathcal{E}_m$ and $\mathcal{E}_\emptyset$ clear. Thus, this variable has the payment structure in Lemma 5.1, and is equal to the VCG payment of

---

[19]Here the stepsize computation subroutine allows for identifying the stepsize at which a bidder in market $\mathcal{E}_{m^*}$ starts demanding a new bundle, or the price of an item decreases to zero (if no bidder demands a new bundle). Note that this step is modular and any other procedure for identifying such a stepsize (e.g., explicitly asking bidders how large of an update in the given update direction is necessary for demanding a new bundle) could be used for the design of the auction. This subroutine, on the other hand, suggests an update structure similar to the one employed in the English auction (e.g., see Krishna (2009)), in the sense that a single price-related parameter (θ) is increased, until bidders' demand sets change.

[20]Any payment bounded away from zero could be used in Step S2 to ensure termination of the auction. We establish in the proof of Theorem 5.1 that a payment of $\sum_i |p_i(t+1) - p_i(t)|$ guarantees that the final rebate $q - z^m$ is nonnegative.

agent m when the auction terminates.

Observe that the price update structure employed by our auction closely follows the one in Algorithm 2. The key difference is that while Algorithm 2 focuses on finding the market clearing allocation (and Walrasian equilibrium) for only one market, our auction finds these allocations for all markets $\{\mathcal{E}_m\}_m \cup \{\mathcal{E}_\emptyset\}$. It accomplishes this by first focusing on markets $\{\mathcal{E}_m\}_m$, and updating the prices greedily with respect to the market that is closest to being cleared. Once these markets are cleared, it terminates by clearing $\mathcal{E}_\emptyset$. We refer to this auction as the "interleaved" tree auction, since as opposed to running a separate auction for all markets $\{\mathcal{E}_m\}_m$, we interleave the demand queries for these markets, and update the prices according to the market that is closest to being cleared (S1b). After each price update we check market clearance for all these markets, and potentially jointly clear multiple markets (S1). This structure avoids restarting the auction from the initial price vector $p = 0$ after each market clearance point, and repeatedly asking the same queries to bidders (e.g., their demand at $p = 0$), thereby potentially leading to faster termination. In the literature, there are iterative auction formats that implement the efficient outcome by using the Walrasian equilibria for markets $\{\mathcal{E}_m\}_m \cup \{\mathcal{E}_\emptyset\}$ to compute and assign the VCG payments to bidders (e.g., see Ausubel (2006)). On the other hand, these auction formats are restricted to the gross substitutes setting (i.e., do not allow for complementarities), and involve running a separate auction (in parallel or series) for each of these markets. Importantly, unlike the existing literature, thanks to the interleaved structure, we do not explicitly run multiple auctions for all the markets.

In Theorem 5.1, we establish that when bidders truthfully reveal their demand, our auction terminates in finite time by implementing the efficient allocation. Moreover, we formally show that the final rebate is nonnegative, and ensures that the total payment of each bidder is equal to her VCG payment. Exploiting this observation, we establish that bidders have no incentive to deviate from the truthful bidding strategy after any history. Hence, in our auction it is an ex-post perfect equilibrium to bid truthfully, and this equilibrium implements the efficient outcome.

Theorem 5.1. *Assume that bidders have sign-consistent tree valuations, and interleaved tree auction is used. The following are true:*

(i) *After any history H_t, if all bidders reveal their demand truthfully, then all markets $\mathcal{S} \cup \{\mathcal{E}_\emptyset\}$ clear in finite time. Moreover, the associated allocations are efficient.*

(ii) *When bidders are truthful, the total payment of every bidder m is given by z^m, and is equal to the VCG payment associated with bundle S^m assigned to her at the end of the auction. Moreover, at termination all bidders receive a nonnegative rebate, i.e., $q \geq z^m$.*

(iii) *It is an ex-post perfect equilibrium for bidders to truthfully reveal their demand in this auction. Additionally, this equilibrium implements the efficient allocation and VCG payments.*

Remarks: The special structure of sign-consistent tree valuations leads to three desirable properties of our auction.

- *Anonymous item pricing:* Our auction can implement the efficient outcome in settings where valuations exhibit complementarity (provided that they are sign-consistent tree valuations, see Example 5.2). In doing so, unlike some auction formats in the literature (e.g., Ausubel and Milgrom (2002); Parkes (2006)), it does not require offering a price for each bundle of items (or specifying exponentially many prices at each stage). Instead, it relies on a simple anonymous item pricing rule (i.e., offers a price p_i for each item i), thereby making our auction format practically appealing. Our auction is able to implement the efficient outcome by relying on anonymous item prices, since for sign-consistent tree valuations it is always possible to clear markets using this pricing rule (Theorem 3.2).

- *Compact demand queries:* At full generality there can be exponentially many bundles a bidder can demand at the given prices (requiring bidders to report all such bundles, in fact is a drawback of some of the existing iterative auctions. See e.g. Ausubel (2006); Parkes (2006)). On the other hand, as we establish in Section 4 and employ in our auction format, by exploiting the graphical structure (and using $LP-D^m/DLP-D^m$) it is possible to compactly describe the set of all bundles that are demanded by a bidder. This allows our iterative auction to converge to the efficient outcome by relying on polynomially many (as opposed to exponentially many) messages that describe the demand sets of bidders. We believe that these desirable communication requirements also make our auction format interesting from a practical point of view. This property relies on the existence of integral optimal solutions to $LP-D^m$, which corresponds to bundles demanded by bidders. Such optimal solutions always exists for sign-consistent tree valuations (since $LP-D^m$ coincides with a formulation of LP2 for a single bidder who has node weights $\{w_i^m - p_i\}$, and the latter problem has integral optimal solutions for such valuations by Theorem 3.1).

- *Interleaved structure for computing VCG payments:* Our auction clears markets $\{\mathcal{E}_m\}_m$ jointly by interleaving the demand queries associated with these markets, and employs the market clearance points in the computation of VCG payments. This result is made possible by the fact that for all these markets anonymous item pricing can be used for market clearance, which is a consequence of having sign-consistent tree valuations.

We conclude this section with an example that illustrates how our iterative auction format clears different markets and implements the efficient outcome:

Example 5.2. *Assume that there are three bidders and three items a, b, c. The underlying graph is a tree with edges (a, b) and (b, c). The weights bidders associate with the underlying nodes/edges are given as follows:*

	w_a	w_b	w_c	w_{ab}	w_{bc}
Bidder 1	*4*	*2*	*2*	*1*	*-1*
Bidder 2	*2*	*4*	*2*	*1*	*-1*
Bidder 3	*2*	*2*	*4*	*1*	*-1*

In this example all bidders view items a and b as complements, and b and c as substitutes. The auctioneer starts the auction with price vector $(p_a, p_b, p_c) = (0, 0, 0)$, and updates the prices as suggested by Algorithm 3. The resulting price updates are given in Figure 7. In the figure, the clearance round of each market $\mathcal{E}_k$ is denoted by M_k, and that of market $\mathcal{E}_\emptyset$ is denoted by M_{all}. Initially, in all
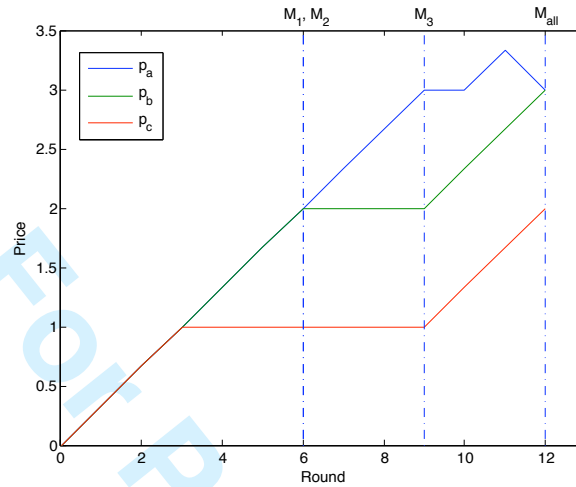

Figure 7: Price updates in the auction.

markets $\{\mathcal{E}_m\}_m$ RP/RD have the same objective value, and all items are overdemanded. Breaking the tie lexicographically, prices are updated with respect to market $\mathcal{E}_1$, and prices of all items are increased. When the price of all items reach to 1, in market $\mathcal{E}_1$, item c is no longer overdemanded (only bidder 3 demands it), and hence its price stops increasing. At round 6 markets $\mathcal{E}_1$ and $\mathcal{E}_2$ jointly clear, and between rounds 6 and 9 prices are updated using the update direction suggested by the solution of RD in $\mathcal{E}_3$. In this market, item a is overdemanded, until its price increases to 3, and at this point $\mathcal{E}_3$ clears as well. In rounds $9 - 12$, the prices are updated using the solution of RD in $\mathcal{E}_\emptyset$. At the end of the 12th round, the auction reaches a Walrasian equilibrium for $\mathcal{E}_\emptyset$ with prices $(p_a, p_b, p_c) = (3, 3, 2)$ and the efficient allocation (that involves assigning a to the first bidder, b to the second one, and c to the last one). The total payments of bidders are the corresponding VCG payments, i.e., 3, 3, and 2 units respectively for bidders $1 - 3$.

This example suggests that in the interleaved tree auction, the auctioneer can implement the efficient outcome without running a separate auction for each market in $\{\mathcal{E}_m\}_m \cup \{\mathcal{E}_\emptyset\}$. Instead, due to the interleaved structure, she can clear multiple markets in one round, and continue running the auction without restarting the auction (at price vector zero). Another property of the auction is that the prices need not always increase, and can decrease over the course of the auction. Finally, updating the prices using the update direction suggested by the market that is closest to clearance (as in Steps S1-S2 of Algorithm 3) leads to termination of the auction in finitely many steps.

6 Conclusions

In this paper, we focus on a special class of graphical valuations, where the underlying value graph is a tree, and edge weights satisfy a sign-consistency condition. We establish that for this class of valuations the efficient allocation can be identified by solving a linear optimization problem, which involves a bidder-specific variable for each item, and pair of items. This result allows for showing that for sign-consistent tree valuations even when items exhibit (pairwise) complementarity, a Walrasian equilibrium exists, and the Walrasian equilibrium allocation and prices can be obtained through the solution of the aforementioned LP and its dual. Conversely, we demonstrate that if the tree or the sign-consistency assumptions are relaxed, then a Walrasian equilibrium need not exist, and the LP formulation that we provide may not find the efficient outcome. We use iterative solutions of our LP formulation with primal-dual algorithms to obtain a new iterative auction format. We establish that for sign-consistent tree valuations this auction implements the efficient allocation at an ex-post perfect equilibrium. Importantly, our auction accomplishes this by relying on a simple anonymous item pricing rule, and allowing bidders to compactly report their demand. Additionally, it employs an interleaved price update structure which guarantees that the final payments of bidders are the VCG payments, and these payments lead to truthful demand reports in the auction. Our results suggest that in multi-item settings (with value complementarities) by exploiting the special structure of bidders' valuations, simple efficient iterative auction formats can be obtained.

Our companion paper (Candogan et al., 2013) extends the approach and results of this paper to more general graphical valuations, and provides similar iterative auction formats. As established in Section 3, for such valuations a Walrasian equilibrium need not exist. This motivates using more general pricing rules than anonymous item pricing, and employing generalizations of Walrasian equilibrium for the design of such auctions. In particular, we provide auctions that rely on simple pricing rules that are bidder-specific and involve pairwise discounts/markups for pairs of items, and the corresponding generalizations of Walrasian equilibrium. These auctions also allow bidders to compactly report their demand, clear multiple markets jointly using an interleaved structure, and guarantee efficiency (at an ex-post perfect equilibrium) for all graphical valuations. This work and the current paper open up a number of interesting future directions, including the design of simple mechanisms by exploiting more general special value structures, and the maximization of revenue and efficiency in markets with multiple bidders/sellers through the use of such mechanisms.

References

Abraham, I., Babaioff, M., Dughmi, S., Roughgarden, T., 2012. Combinatorial auctions with restricted complements. In: Proceedings of the 13th ACM Conference on Electronic Commerce. ACM, pp. 3–16.

Ausubel, L., 2004. An efficient ascending-bid auction for multiple objects. The American Economic Review 94 (5), 1452–1475.

Ausubel, L., 2006. An efficient dynamic auction for heterogeneous commodities. The American economic review 96 (3), 602–629.

Ausubel, L., Cramton, P., 2004. Auctioning many divisible goods. Journal of the European Economic Association 2 (2-3), 480–493.

Ausubel, L., Milgrom, P., 2002. Ascending auctions with package bidding. The BE Journal of Theoretical Economics 1 (1), 1–42.

Ausubel, L. M., Milgrom, P., 2006. The lovely but lonely Vickrey auction. In: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), Combinatorial Auctions. MIT Press, pp. 17–40.

Baldwin, E., Klemperer, P., 2012. Tropical geometry to analyse demand. Tech. rep., Working paper, Oxford University.

Bazaraa, M. S., Jarvis, J. J., Sherali, H. D., 2011. Linear programming and network flows. Wiley. com.

Bertsekas, D., 1996. Dynamic programming and optimal control. Vols. 1&2. Athena Scientific.

Bertsimas, D., Tsitsiklis, J., 1997. Introduction to Linear Optimization. Athena Scientific, Belmont, MA.

Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R., 2002. Linear programming and Vickrey auctions. In: Dietrich, B., Vohra, R. (Eds.), Mathematics of the Internet: E-auction and Markets. New York: Springer Verlag, pp. 75–116.

Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R., 2011. An ascending Vickrey auction for selling bases of a matroid. Operations Research 59 (2), 400–413.

Bikhchandani, S., Mamer, J., 1997. Competitive equilibrium in an exchange economy with indivisibilities. Journal of Economic Theory 74 (2), 385–413.

Bikhchandani, S., Ostroy, J., 2002. The package assignment model. Journal of Economic Theory 107 (2), 377–406.

Blumrosen, L., Nisan, N., 2010. On the computational power of demand queries. SIAM Journal on Computing 39 (4), 1372.

Candogan, O., 2013. Dynamic strategic interactions: Analysis and mechanism design. Ph.D. thesis, Massachusetts Institute of Technology.

Candogan, O., Ozdaglar, A., Parrilo, P. A., 2013. Iterative auction design for graphical valuations part ii: General graphs. working paper.

Conitzer, V., Sandholm, T., Santi, P., 2005. Combinatorial auctions with k-wise dependent valuations. In: AAAI. Vol. 5. pp. 248–254.

Cramton, P., Ausubel, L., McAfee, R., McMillan, J., 1997. Synergies in wireless telephony: Evidence from the broadband PCS auctions. Journal of Economics and Management Strategy 6 (3), 497–527.

Cramton, P., Shoham, Y., Steinberg, R., 2006. Combinatorial auctions. MIT Press.

De Vries, S., Schummer, J., Vohra, R., 2007. On ascending Vickrey auctions for heterogeneous objects. Journal of Economic Theory 132 (1), 95–118.

De Vries, S., Vohra, R., 2003. Combinatorial auctions: A survey. INFORMS Journal on Computing

15 (3), 284–309.

Engelbrecht-Wiggans, R., Kahn, C., 1991. Protecting the winner: Second-price versus oral auctions. Economics Letters 35 (3), 243–248.

Gul, F., Stacchetti, E., 1999. Walrasian equilibrium with gross substitutes. Journal of Economic Theory 87 (1), 95–124.

Gul, F., Stacchetti, E., 2000. The English auction with differentiated commodities. Journal of Economic theory 92 (1), 66–95.

Hohner, G., Rich, J., Ng, E., Reid, G., Davenport, A. J., Kalagnanam, J. R., Lee, H. S., An, C., January/February 2003. Combinatorial and quantity-discount procurement auctions benefit mars, incorporated and its suppliers. Interfaces 33 (1), 23–35.
    URL http://interfaces.journal.informs.org/content/33/1/23.abstract

Kelso, A., Crawford, V., 1982. Job matching, coalition formation, and gross substitutes. Econometrica: Journal of the Econometric Society, 1483–1504.

Krishna, V., 2009. Auction theory. Academic press.

Lahaie, S., 2009. A kernel method for market clearing. In: IJCAI. pp. 208–213.

Lahaie, S., 2011. A kernel-based iterative combinatorial auction. In: AAAI.

Lehmann, B., Lehmann, D., Nisan, N., 2006. Combinatorial auctions with decreasing marginal utilities. Games and Economic Behavior 55 (2), 270–296.

Mas-Collel, A., Whinston, M. D., Green, J., 1995. Microeconomic theory. Oxford university press Oxford.

McAfee, R. P., McMillan, J., Wilkie, S., 2010. The greatest auction in history. In: Better living through economics. Harvard University Press, pp. 168–184.

Mishra, D., Parkes, D., 2007. Ascending price Vickrey auctions for general valuations. Journal of Economic Theory 132 (1), 335–366.

Mishra, D., Parkes, D., 2009. Multi-item vickrey–dutch auctions. Games and Economic Behavior 66 (1), 326–347.

Moreton, P., Spiller, P., 1998. What's in the air: Interlicense synergies in the federal communications commission's broadband personal communication service spectrum auctions. Journal of Law and Economics 41 (2), 677–716.

Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. V., 2007. Algorithmic Game Theory. Cambridge University Press, New York, NY, USA.

Nisan, N., Schapira, M., Valiant, G., Zohar, A., 2011. Best-response auctions. In: Proceedings of the 12th ACM conference on Electronic commerce. ACM, pp. 351–360.

Nisan, N., Segal, I., 2006. The communication requirements of efficient allocations and supporting prices. Journal of Economic Theory 129 (1), 192–224.

Papadimitriou, C., Steiglitz, K., 1998. Combinatorial optimization: algorithms and complexity. Dover Publications.

Parkes, D., 1999. i-bundle: An efficient ascending price bundle auction. In: Proceedings of the 1st ACM conference on Electronic commerce. ACM, pp. 148–157.

Parkes, D., 2006. Iterative combinatorial auctions. In: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), Combinatorial Auctions. MIT Press, pp. 41–77.

Rothkopf, M., Teisberg, T., Kahn, E., 1990. Why are Vickrey auctions rare? Journal of Political Economy, 94–109.

Sandholm, T., Boutilier, C., 2006. Preference elicitation in combinatorial auctions. In: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), Combinatorial Auctions. MIT Press, pp. 41–77.

Sun, N., Yang, Z., 2006. Equilibria and indivisibilities: gross substitutes and complements. Econometrica 74 (5), 1385–1402.

Sun, N., Yang, Z., 2009. A double-track adjustment process for discrete markets with substitutes and complements. Econometrica 77 (3), 933–952.

Vohra, R. V., 2011. Mechanism Design: A linear programming approach. Cambridge University Press.

Zhou, H., Berry, R., Honig, M., Vohra, R., 2009. Complementarities in spectrum markets. In: Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on. IEEE, pp. 499–506.

A   Graphical Valuations and Other Special Valuations

In this section, we explain how graphical valuations and tree valuations are different from other special classes of valuations studied in the literature. In particular, we focus on the classes of gross substitutes and complements, gross substitutes, sub/superadditive, sub/supermodular value functions, and compare them with graphical valuations.

It was established in Section 3 that when the underlying value graph has a tree structure, and the valuations satisfy sign-consistency, a Walrasian equilibrium exists. This result allows us to identify a class of value functions which exhibit both value complementarity and substitutability, and for which a Walrasian equilibrium exists. Gross substitutes and complements (Sun and Yang, 2006, 2009), defined below, is another class of value functions that satisfies a similar property.

Definition A.1 (Gross Substitutes and Complements (GSC)). *Assume that the set of items is partitioned into two sets S_1, S_2 such that $S_1 \cap S_2 = \emptyset$, $S_1 \cup S_2 = \mathcal{N}$. Consider the value function $v : \mathcal{N} \to \mathbb{R}$. Denote by $e(k)$ the kth unit vector, and $D(p)$ the demand function associated with price vector $p \in \mathbb{R}^N$, i.e., $D(p) \triangleq \arg\max_{S \subset \mathcal{N}} v(S) - \sum_{i \in S} p_i$.*

We say that v has the gross substitutes and complements property if for $j \in \{1, 2\}$, any price vector $p \in \mathbb{R}^N$, $k \in S_j$, $\delta \geq 0$, and $D_1 \in D(p)$, there exists $D_2 \in D(p + \delta e(k))$ such that (a) $[D_1 \cap S_j] - \{k\} \subset D_2$ and (b) $D_1^c \cap S_j^c \subset D_2^c$.

Intuitively, this definition suggests that the items in sets S_1 and S_2 are substitutes among themselves (in the sense that if the price of a demanded item in one of these sets increases, the demand for the other demanded items in the same set does not decrease, $[D_1 \cap S_j] - \{k\} \subset D_2$).

Additionally items are complements across S_1 and S_2 (in the sense that if the price of a demanded item in set S_1 increases, then fewer items are demanded in set S_2, $D_1^c \cap S_j^c \subset D_2^c$).

We next illustrate that tree (and hence graphical) valuations need not satisfy the GSC property.

Example A.1. *Consider the tree valuation provided in Figure 8. Assume that this valuation satisfies the GSC property. There are three different ways of choosing sets S_1 and S_2 (due to symmetry all other cases follow from the analysis here): (i) $S_1 = \{A, B, C\}$, $S_2 = \emptyset$, (ii) $S_1 = \{A\}$, $S_2 = \{B, C\}$, (iii) $S_1 = \{A, C\}$, $S_2 = \{B\}$.*
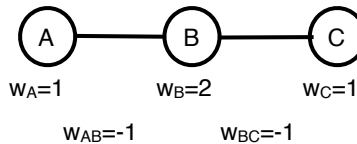

Figure 8: A tree valuation that violates the GSC property.

We will show that the GSC property fails in all of these cases, and hence the value function given in Figure 8 does not exhibit the GSC property for any choice of $\{S_j\}$. Assume that

- *Initially, the prices are $p_1(A) = 0.1$, $p_1(B) = 0.5$, $p_1(C) = 0.1$, and the corresponding demand is $D(p_1) = \{A, C\}$.*

- *Then, the price of the first item is increased, and the new prices are $p_2(A) = 1$, $p_2(B) = 0.5$, $p_2(C) = 0.1$. It follows that the demand is $D(p_2) = \{B\}$.*

This implies that the GSC property fails whenever A and C belong to the same S_j (note that by choosing $D_1 = \{A, C\}$, $D_2 = \{B\}$, $k = A$, the condition $[D_1 \cap S_j] - \{k\} \subset D_2$ fails). Thus, to check the GSC property it is sufficient to focus on case (ii). On the other hand, if $S_1 = \{A\}$, $S_2 = \{B, C\}$ then the condition $D_1^c \cap S_j^c \subset D_2^c$ fails (this can be seen by choosing $j = 1$, $D_1 = \{A, C\}$, $D_2 = \{B\}$, $k = A$). This implies that the GSC property fails in case (ii) as well. Hence, we conclude that for any choice of the $\{S_j\}$ sets, the GSC property fails for the value function in Figure 8.

This example shows that tree valuations are not contained in the class of GSC valuations. GSC generalizes the well-known gross substitutes class (Gul and Stacchetti, 1999), where Definition A.1 holds with $S_2 = \emptyset$. Thus, our results also imply that tree valuations do not necessarily satisfy the gross substitutes property.

We next investigate the additional structural assumptions under which graphical valuations exhibit the GSC property. Assume that the underlying value graph consists of connected components of size at most two. Note that in this case, valuations are additive over different connected components, and hence in order to test the GSC condition it suffices to restrict attention to subsets of the demand set that are contained in a given connected component of the graph.

Consider a pair of nodes (i, j) connected with an edge. Assume that $w_{ij}^m \leq 0$ and at a given price vector p item j belongs to a demand set D_1. We claim that if the price of item i increases, then the demand for item j cannot decrease. The claim is immediate if $i \notin D_1$, i.e., i is not demanded

at the original prices. Assume that $i \in D_1$. Observe that this implies that $w_j^m + w_{ij}^m - p_j \geq 0$, since otherwise bidder m can improve her payoff by not receiving item j at the price vector p, and hence $j \notin D_1$. On the other hand, since $w_{ij}^m \leq 0$, it follows that $w_j^m - p_j \geq w_j^m + w_{ij}^m - p_j \geq 0$. Thus, at the updated prices bidder m still maximizes her surplus by either receiving item j together with i or in isolation. Hence, item j belongs to a demand set after the price update, and condition (a) of Definition A.1 holds by assigning items (i, j) to the same set S_1 or S_2.

Conversely, assume that $w_{ij}^m \geq 0$ and at price vector p, item j does not belong to a demand set D_1. We claim that if the price of item i increases, then the demand for item j cannot increase. As before, the claim is immediate if $i \notin D_1$. If $i \in D_1$, and $j \notin D_1$, then it should be the case that $w_j^m + w_{ij}^m - p_j \leq 0$. Moreover, since $w_{ij}^m \geq 0$, this implies that $w_j^m - p_j \leq w_j^m + w_{ij}^m - p_j \leq 0$. Note that after the price update this inequality continues to hold. Thus, it should be the case that there is a demand set to which item j does not belong after the price update. Hence, condition (b) of Definition A.1 holds, by assigning items (i, j) to different sets S_1 and S_2.

These observations imply that if the underlying graph consists of components of size at most two, the GSC condition holds, by assigning items that are connected with a positive weight to different sets S_1 and S_2 (see Definition A.1), and items that are connected with a negative edge to the same set.

On the other hand, Example A.1 implies that when a connected component has at least three nodes the GSC property may not hold. Moreover, this conclusion still holds, if edge weights are not restricted to be negative as in the example, and allowed to be positive or negative; see (Candogan, 2013). Thus, unless further restriction on the weights is made, graphical valuations satisfy the GSC property only when the underlying graph consists of connected components of size two. This implies that GSC property holds for only a very restrictive subclass of graphical valuations.[21]

We conclude this section by discussing the relation of graphical valuations to subadditive/superadditive and submodular/supermodular valuations. A value function is subadditive if for any sets $A, B \subset \mathcal{N}$, it satisfies $v(A \cup B) \leq v(A) + v(B)$, and superadditive, if for disjoint A, B it satisfies $v(A \cup B) \geq v(A) + v(B)$. Similarly a value function is submodular if for any sets A, B it satisfies $v(A \cup B) + v(A \cap B) \leq v(A) + v(B)$, and supermodular if $v(A \cup B) + v(A \cap B) \geq v(A) + v(B)$. These inequalities imply that for nonnegative value functions, submodularity implies subadditivity.

It can be easily checked that if all edge weights are positive (negative), graphical valuations are superadditive and supermodular (subadditive and submodular). On the other hand, if there is an edge with negative (positive) weight, the supermodularity/superadditivity (submodularity/subadditivity) condition cannot hold (consider A, B as singletons corresponding to the end points of this edge). Since the weights of different edges in our model can be positive or negative, it follows that even for the case of trees, graphical valuations are not contained in these classes.

---

[21]A similar conclusion holds for gross substitute valuations: a graphical valuation satisfies the gross substitutes condition, if the graph consists of connected components of size at most two, and the edge weights are negative.

B   Proofs of Section 3

Proof of Theorem 3.1(i). Let W^* denote the welfare (total value) associated with an efficient allocation, and $h(i)$ denote the bidder who receives item i at this allocation. In order to establish the result, we first construct another LP formulation of the efficient allocation problem, and show that the optimal objective value of this LP is equal to W^*. We then show that optimal solutions of LP2 can be mapped to feasible solutions of this LP with the same objective value. This implies that optimal objective value of LP2 is bounded by W^*. On the other hand, LP2 has feasible integral solutions associated with the efficient allocation that have objective value W^*. Hence, it follows that such integral solutions are optimal in LP2.

We start by introducing some notation, and stating the alternative LP formulation of the efficient allocation problem. Given the underlying graphical (tree) valuation, choose a node r as the root of the tree. For any given node i in this tree, denote the set of children of node i by $C(i)$. Similarly, denote the unique parent of node i by $P(i)$. We use the convention $C(i) = \emptyset$ if i is a leaf node, and $P(i) = \emptyset$ if $i = r$. Consider the following LP formulation:

$$
\begin{aligned}
\min \quad & z^* \\
s.t. \quad & z_i^m \geq w_i^k + \sum_{j \in C(i)} z_j^k \qquad \text{for all } i \neq r, \text{ and } m, k \neq m \qquad (x_i^{m \to k}) \\
(DLP-rec) \quad & z_i^m \geq w_i^m + w_{iP(i)}^m + \sum_{j \in C(i)} z_j^m \qquad \text{for all } i \neq r, \text{ and } m \qquad (x_i^{m \to m}) \quad (6) \\
& z^* \geq w_r^m + \sum_{j \in C(r)} z_j^m \qquad \text{for all } m \qquad (x_r^{* \to m}).
\end{aligned}
$$

In DLP-rec we express the dual variable corresponding to each constraint in parenthesis. Using these variables, we next state the dual of DLP-rec (henceforth referred to as LP-rec):

$$
\begin{aligned}
\max \quad & \sum_m x_r^{\star \to m} w_r^m + \sum_{i \neq r} \sum_{m, k \neq m} x_i^{k \to m} w_i^m + \sum_{i \neq r} \sum_m x_i^{m \to m} (w_i^m + w_{iP(i)}^m) \\
s.t. \quad & \sum_k x_{P(i)}^{k \to m} = \sum_k x_i^{m \to k} \qquad \text{for all } i \neq r, \text{ m,} \qquad (z_i^m) \\
(LP-rec) \quad & \sum_m x_r^{\star \to m} = 1 \qquad (z^*) \\
& x_i^{k \to m}, x_i^{\star \to m} \geq 0.
\end{aligned} \qquad (7)
$$

The variables of DLP-rec corresponding to the constraints of LP-rec are stated in parenthesis in the above formulation.

In DLP-rec, we interpret z_i^m as the maximum welfare that can be obtained from the assignment of node i and her children, given that the parent of node i is assigned to bidder m (later we will prove that at an optimal solution this interpretation is precise). The first two constraints use the tree structure to recursively express this quantity. In particular, the first constraint suggests that this quantity is at least the sum of the value obtained from the assignment of node i to some

bidder $k \neq m$ (w_i^k) and the welfare obtained from each of the subtrees rooted at the children of i ($\sum_{j \in C(i)} z_j^k$), i.e., $z_i^m \geq w_i^k + \sum_{j \in C(i)} z_j^k$. The second quantity, similarly, implies that this quantity is weakly larger than the sum of the value obtained from the assignment of node i to bidder m ($w_i^m + w_{iP(i)}^m$), and the total welfare obtained from the subtrees ($\sum_{j \in C(i)} z_j^k$). Note that the right hand side of this constraint includes an additional term $w_{iP(i)}^m$, since it captures the case where both node i and its parent are assigned to bidder m, and hence generate an additional value of $w_{iP(i)}^m$. The final constraint suggests that the maximum welfare that can be obtained from the assignment of all nodes is at least the sum of the value obtained from the assignment of the root node to bidder m, and the welfare obtained from the assignment of the subtrees rooted at the children of this node.[22]

Note that at feasible solutions of DLP-rec where constraints 1-2 are strict, the z_i^m variable can be decreased to obtain another feasible solution, with the same objective value. This implies that DLP-rec has an optimal solution where for all $m \in \mathcal{M}$, and $i \neq r$, we have

$$z_i^m = \max \left\{ w_i^m + w_{iP(i)}^m + \sum_{j \in C(i)} z_j^m, \quad \max_{k \neq m}\{w_i^k + \sum_{j \in C(i)} z_j^k\} \right\}. \tag{8}$$

Similarly, it can be checked from DLP-rec that at all optimal solutions we have

$$z^* = \max_m \{w_r^m + \sum_{j \in C(r)} z_j^m\}. \tag{9}$$

Consider an optimal solution where (8) and (9) hold. We claim that at this solution any node $i \neq r$ satisfies the following property (or the interpretation for z_i^m stated earlier): z_i^m is equal to the maximum welfare obtained by the assignment of i and her children when the parent of this node is assigned to bidder m. It can be seen that (8) immediately implies this property for the leaf nodes, since for these nodes $z_i^m = \max\{w_i^m + w_{iP(i)}^m, \max_{k \neq m} w_i^k\}$. Assume that this property holds for any node $i \neq r$ such that the shortest path that connects i to a leaf node (that does not include r) has length at most k. It follows from (8) that z_i^m satisfies this property for any node whose children are a part of a path that has length at most k. Equivalently, the property holds for any node $i \neq r$ that belongs to a path of length at most $k + 1$. Hence, by induction the claim follows for all $\{z_i^m\}$. On the other hand, since $z^* = \max_m\{w_r^m + \sum_{j \in C(r)} z_j^m\}$, the claim implies that z^* is equal to the maximum value that can be obtained by the assignment of all nodes, i.e., W^*. This suggests that the optimal objective value of DLP-rec is weakly lower than W^*.

Similarly, consider the following solution to LP-rec: $x_r^{\star \to m} = 1$ if $h(r) = m$, and $x_r^{\star \to m} = 0$ otherwise, and $x_i^{k \to m} = 1$ if $h(i) = m$ and $h(P(i)) = k$ and $x_i^{k \to m} = 0$ otherwise. It can be immediately verified in LP-rec that this solution is feasible and the corresponding welfare is equal to W^*. Thus, the objective value of LP-rec is at least W^*. Since the optimal objective value of DLP-rec is bounded by W^*, it follows that the optimal objective values of both LP-rec, and

---

[22]This LP essentially describes a recursion that can be used for the solution of the welfare maximization problem.

DLP-rec are equal to W^*. Moreover, the solutions constructed above are optimal in the respective optimization problems.

We next establish that for sign-consistent tree valuations LP2 has optimal solutions that can be mapped to feasible solutions of LP-rec, with the same objective value. Let E^+ denote the set of edges with positive weights, and E^- denote the set of edges with negative weights (recall that by Assumption 3.2 all bidders have same-sign weights for a given edge). Observe that $LP2$ has optimal solutions where we have:

A1: $y_{ij}^m = \min\{x_i^m, x_j^m\}$ for all $m \in \mathcal{M}, ij \in E^+$,

A2: $y_{ij}^m = \max\{0, x_i^m + x_j^m - 1\}$ for all $m \in \mathcal{M}, ij \in E^-$,

A3: $\sum_m x_i^m = 1$ for all i.

Here, the first two conditions hold as given a solution where they do not hold increasing (for A1) or decreasing (for A2) y_{ij}^m a new solution with improved objective can be obtained (due to the signs of the corresponding edges). The last condition is a byproduct of Assumption 2.1.

Consider an optimal solution of LP2 satisfying A1-3. In Lemma B.1 we show that for such an optimal solution $\{x_i^m, y_{ij}^m\}$ of LP2, a feasible solution $\{x_*^{k \to m}\}$ to LP-rec with the same objective value exists. Moreover, in this solution of LP-rec, we have:

B1: $x_i^m = \sum_k x_i^{k \to m}$ for all $m, i \neq r$ (and $x_i^m = x_i^{* \to m}$ for $i = r$),

B2: $y_{ij}^m = x_i^{m \to m}$ for all m, i, j, where $j = P(i)$ in the tree rooted at r.

B3: $\sum_m \sum_k x_i^{k \to m} = 1$ for all $i \neq r$ (and $\sum_m x_i^{\star \to m} = 1$ for $i = r$).

The proof of this lemma is given at the end of this proof.

Lemma B.1. *Assume that an optimal solution $\{x_i^m, y_{ij}^m\}$ of LP2 satisfying A1-3 is given. There exists a corresponding feasible solution $\{x_i^{k \to m}, x_i^{\star \to m}\}$ of LP-rec satisfying B1-3. Moreover, these solutions have the same objective values in the respective optimization problems.*

Since optimal value of LP-rec is W^*, this Lemma implies that the objective value of LP2 is bounded by this quantity as well. On the other hand, there is an integral feasible solution of LP2 associated with the efficient allocation, which achieves this objective value. In particular, consider the solution where $x_i^m = 1$ if $h(i) = m$, $y_{ij}^m = 1$, if $h(i) = h(j) = m$, and $x_i^m = y_{ij}^m = 0$ otherwise. It can be verified from LP2 that this solution is feasible, and the associated objective value is W^*. Thus, it follows that this solution is an optimal solution of LP2 that is integral, and the claim follows. $\square$

Proof of Lemma B.1. We establish the result by constructing a feasible solution to LP-rec, and establishing that this solution satisfies B1-B2. The condition B3, on the other hand, trivially holds for any feasible solution of LP-rec. For the root node, this condition simply is the second constraint of LP-rec. For the remaining nodes summing the first constraint of LP-rec over m it can be seen

that if B3 holds for the parent of a node, it holds for the node itself as well. Thus, for all nodes, this condition follows from the feasibility of the solution we construct for LP-rec, and for the proof it suffices to focus on B1-B2.

In order to construct a solution to LP-rec, we first set $x_r^{\star \to m} = x_r^m$ for all m. Consider any node $i \neq r$, and its parent $j = P(i)$. We construct the $\{x_i^{k \to m}\}_{k,m}$ variables that satisfy conditions B1 and B2 (for all i and m) by solving a max-flow problem (see Figure 9). In this problem we maximize
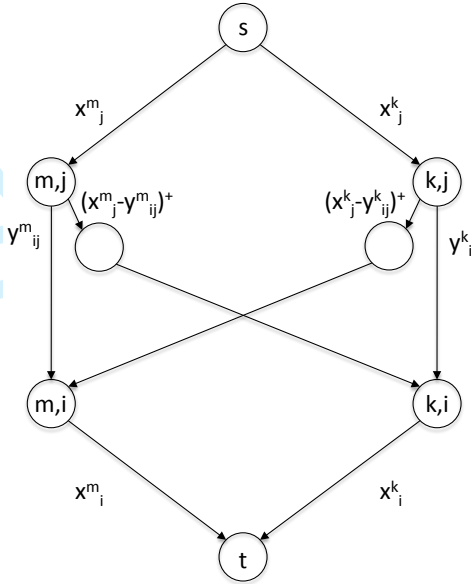

Figure 9: The max-flow problem used for constructing a solution to LP-rec.

the flow between nodes s and t. The labels associated with edges denote their capacities, and edges without labels have unlimited capacity. We associate nodes (m,j) and (m,i) with each bidder m. The (m,j) nodes have an incoming edge from s (with capacity x_j^m), and (m,i) nodes have an outgoing edge to t (with capacity x_i^m). For each (m,j) node there is an associated $(m,j)^+$ node (which is left unlabeled in the figure for simplicity). Each (m,j) has an outgoing edge connected to (m,i) node (with capacity y_{ij}^m), and the associated $(m,j)^+$ node (with capacity $(x_j^m - y_{ij}^m)^+$). The $(m,j)^+$ node, on the other hand, has outgoing edges (with unlimited capacity) connected to (k,i) nodes for $k \neq m$. Given an optimal solution of this flow problem, we will construct $\{x_i^{k \to m}\}_{k,m}$, by setting $x_i^{m \to m}$ equal to the flow between (m,j) and (m,i). Similarly, we will set $x_i^{m \to k}$ equal to the flow between $(m,j)^+$ and (k,i).

We first characterize the optimal solution of the max-flow problem. Since max-flow equals min-cut, the solution of this flow problem can be identified by focusing on the cuts (weighted by edge capacities) in the underlying graph. Since the edges between $(m,j)^+$ and (k,i) have unlimited capacity, the min-cut either involves none of the edges between the (m,j) and (m,i) nodes (for any m), or it involves both the edge $(m,j) - (m,i)$ and $(m,j) - (m,j)^+$. Since $y_{ij}^m + (x_j^m - y_{ij}^m)^+ \geq x_j^m$, the cut value is minimized by cutting edge $s - (m,j)$ instead of the edges $(m,j) - (m,i)$ and $(m,j) - (m,j)^+$. This suggests that a minimum cut in this problem is obtained by not cutting any

of the $(m,j) - (m,i)$ edges. Thus, the minimum cut value is equal to either $\sum_m x_j^m$ or $\sum_m x_i^m$. By A3 both of these quantities, hence min-cut/max-flow, is equal to 1. Note that for the total flow to be equal to $\sum_m x_j^m = \sum_m x_i^m = 1$, all edges adjacent to s and t nodes carry flow equal to their capacity. Given such an optimal solution of the flow problem, construct $\{x_i^{k \to m}\}_{k,m}$, by setting $x_i^{m \to m}$ equal to the flow between (m,j) and (m,i), and $x_i^{m \to k}$ equal to the flow between $(m,j)^+$ and (k,i). Since, the capacity of edge $(m,i) - t$ is equal to x_i^m, and the edges adjacent to node (m,i) are associated with variables $\{x_i^{k \to m}\}_k$, it follows that $\sum_k x_i^{k \to m} = x_i^m$. Consequently, B1 holds for the constructed solution.

Assume that $ij \in E^+$, hence A1 holds, i.e., $y_{ij}^m = \min\{x_i^m, x_j^m\}$. In this case, since the max-flow solution necessarily sends x_j^m amount of flow on the edge $s - (m,j)$, it follows that y_{ij}^m units of flow is sent on edge $(m,j) - (m,i)$, and $(x_j^m - y_{ij}^m)$ units on $(m,j) - (m,j)^+$. This implies that the constructed solution also satisfies B2. Thus, when $ij \in E^+$, by solving the constructed max-flow problem, we obtain $\{x_i^{k \to m}\}_{k,m}$ that satisfy B2.

Assume that $ij \in E^-$, hence A2 holds, i.e., $y_{ij}^m = \max\{0, x_i^m + x_j^m - 1\}$. If $y_{ij}^m = 0$, then on edge $(m,j) - (m,i)$ zero (or y_{ij}^m) units of flow is sent. On the other hand, if $y_{ij}^m = x_i^m + x_j^m - 1 > 0$, then $x_j^m = 1 - x_i^m + y_{ij}^m \geq y_{ij}^m$. Since, x_j^m units of flow is sent on the edge $s - (m,j)$ (at all max-flow solutions), we conclude that y_{ij}^m units of flow is sent on edge $(m,j) - (m,i)$, and $(x_j^m - y_{ij}^m)$ units of flow is sent on $(m,j) - (m,j)^+$. This implies that the constructed solution satisfies B2 when $ij \in E^-$ as well.

Consider the solution of LP-rec obtained by solving the flow problems associated with all edges of the underlying tree (and setting $x_r^{\star \to m} = x_r^m$ for all m). To complete the proof, it suffices to establish that this solution is feasible in LP-rec and establish that the objective values of the solutions we have for LP-rec and LP2 are identical.

Note that our construction trivially satisfies the second (since $\sum_m x_r^{\star \to m} = \sum_m x_r^m = 1$, where the last equality follows from A3) and third (since flows are always nonnegative) constraints of LP-rec. Thus, we focus on the first constraint. Observe that since in the optimal flow solution edges $s - (m,j)$ and $(m,i) - t$ respectively carry x_j^m and x_i^m units of flow, it follows that the solution we construct for LP-rec satisfies $\sum_k x_i^{m \to k} = x_i^m = x_{P(i)}^m$ and $\sum_k x_i^{k \to m} = x_i^m$. This implies that $\sum_k x_i^{m \to k} = x_{P(i)}^m = \sum_k x_{P(i)}^{k \to m}$. Hence, the first constraint of LP-rec is always satisfied. Therefore, we conclude that the constructed solution is also feasible in LP-rec.

We conclude the proof by showing that these solutions have the same objective value in the respective optimization problems. Note that using B1-B3 the objective value of LP-rec can be expressed as follows:

$$\sum_m x_r^{\star \to m} w_r^m + \sum_{i \neq r} \sum_{m,k \neq m} x_i^{k \to m} w_i^m + \sum_{i \neq r} \sum_m x_i^{m \to m} (w_i^m + w_{iP(i)}^m)$$

$$= \sum_m x_r^{\star \to m} w_r^m + \sum_{i \neq r} \sum_{m,k} x_i^{k \to m} w_i^m + \sum_{i \neq r} \sum_m x_i^{m \to m} w_{iP(i)}^m \qquad (10)$$

$$= \sum_m x_r^m w_r^m + \sum_{i \neq r} \sum_m x_i^m w_i^m + \sum_{i \neq r} \sum_m y_{iP(i)}^m w_{iP(i)}^m,$$

where the first equality is obtained by rearranging terms, and the second one is obtained using B1 and B2. On the other hand, the quantity in the last line of (10) is equal to the objective value of LP2, since $w_{ij}^m = 0$ unless the edge (i,j) is present in the underlying tree. Thus, the solutions we construct lead to the same objective value for the associated optimization problems, and the claim follows. $\square$

Proof of Theorem 3.1(ii). Denote the objective value at an optimal solution of LP1 by OP1, and the value at an optimal integer solution of LP1 by OPI1. Similarly, denote by OP2 and OPI2 the objective values of an optimal solution of LP2 and an optimal integer solution of LP2. If LP2 has an optimal solution that is integral, we know that $OP2 = OPI2$. Also, since the optimal objective value cannot be larger after imposing the integrality constraint, we have $OP1 \geq OPI1$. We next show that $OP1 \leq OP2$ and $OPI2 \leq OPI1$ to establish that $OP1 = OPI1$. Note that this immediately implies that LP1 has an optimal solution that is integral.

$OP1 \leq OP2$**:** Consider a feasible solution $\{x^m(S)\}$ of LP1. We will show that it is possible to construct a feasible solution of LP2 with the same objective value.

In particular, let $x_i^m = \sum_{S|i\in S} x^m(S)$, and $y_{ij}^m = \sum_{S|i,j\in S} x^m(S)$ for all $m \in \mathcal{M}$, $i \in \mathcal{N}$, and $(i,j) \in E$. Since feasible solutions of LP1 satisfy $\sum_m \sum_{S|i\in S} x^m(S) \leq 1$, we have $\sum_m x_i^m = \sum_m \sum_{S|i\in S} x^m(S) \leq 1$. Additionally, since $x^m(S) \geq 0$, we obtain $y_{ij}^m = \sum_{S|i,j\in S} x^m(S) \leq \sum_{S|i\in S} x^m(S) \leq x_i^m$. Thus, it follows that the constructed solution also satisfies $y_{ij}^m \leq x_i^m, x_j^m$. Finally, for any $ij \in E$ we have

$$x_i^m + x_j^m - y_{ij}^m = \sum_{S|i\in S} x^m(S) + \sum_{S|j\in S} x^m(S) - \sum_{S|i,j\in S} x^m(S)$$
$$= \sum_{S|i\in S} x^m(S) + \sum_{S|j\in S, i\notin S} x^m(S) \leq \sum_{S|i\in S} x^m(S) + \sum_{S|i\notin S} x^m(S) \leq 1,$$

where the last inequality follows since $\sum_S x^m(S) \leq 1$.

Summarizing, we established that the constructed x_i^m, y_{ij}^m is such that it satisfies: (i) $\sum_m x_i^m \leq 1$, (ii) $y_{ij}^m \leq x_i^m, x_j^m$, (iii) $x_i^m + x_j^m - y_{ij}^m \leq 1$. Additionally, since $x^m(S) \geq 0$, we have $x_i^m, y_{ij}^m \geq 0$. Finally, since $\sum_m x_i^m \leq 1$ and $x_i^m \geq 0$, we have $x_i^m \leq 1$, and since $y_{ij}^m \leq x_i^m$ we have $y_{ij}^m \leq 1$. These together imply that $\{x_i^m, y_{ij}^m\}$ is a feasible solution of LP2.

Observe that $\sum_{m,S} x^m(S)v^m(S)$, the objective value of LP1 corresponding to $\{x^m(S)\}$, satisfies

$$\sum_{m,S} x^m(S)v^m(S) = \sum_{m,S} x^m(S)\Big(\sum_{i\in S} w_i^m + \sum_{i,j\in S|ij\in E} w_{ij}^m\Big)$$
$$= \sum_m \Big(\sum_i w_i^m \sum_{S|i\in S} x^m(S) + \sum_{ij\in E} w_{ij}^m \sum_{S|i,j\in S} x^m(S)\Big) = \sum_m \left(\sum_i w_i^m x_i^m + \sum_{ij\in E} w_{ij}^m y_{ij}^m\right).$$

That is it is equal to the objective value of LP2 corresponding to $\{x_i^m, y_{ij}^m\}$. Hence, given a feasible solution of LP1, there exists a corresponding feasible solution of LP2, with the same objective value. Since this is true for the optimal solution of LP1 as well, we conclude $OP1 \leq OP2$.

$OPI2 \leq OPI1$: Consider a feasible integer solution $\{x_i^m, y_{ij}^m\}$ of LP2. Let $S^m = \{i | x_i^m = 1\}$. Since $\sum_m x_i^m \leq 1$, it follows that if $x_i^m = 1$ then $x_i^k = 0$ for $k \neq m$. Hence, $S^m \cap S^k = \emptyset$ for $k \neq m$.

Define $\{x^m(S)\}$ such that for all m, $x^m(S^m) = 1$, and $x^m(S) = 0$ for $S \neq S^m$. Observe that such a solution satisfies $\sum_S x^m(S) \leq 1$, and $\sum_m \sum_{S|i \in S} x^m(S) \leq 1$ (since $S^m \cap S^k = \emptyset$). Thus, it follows that $\{x^m(S)\}$ is a feasible integer solution to LP1.

Note that feasibility of $\{x_i^m, y_{ij}^m\}$ in LP2 implies that if $x_i^m, x_j^m \in \{0, 1\}$, then $y_{ij}^m \in \{0, 1\}$. More precisely for $ij \in E$, if $x_i^m = x_j^m = 1$ then $y_{ij}^m = 1$ (since $x_i^m + x_j^m - 1 \leq y_{ij}^m$). Similarly, if $x_i^m = 0$ then $y_{ij}^m = 0$ (since $y_{ij}^m \leq x_i^m$). This implies that $y_{ij}^m = 1$ if and only if $x_i^m = x_j^m = 1$.

Observe that the construction of $\{x^m(S)\}$ implies that $x_i^m = \sum_{S|i \in S} x^m(S)$. This is because, if $x_i^m = 0$, then $x^m(S) = 0$ for all S containing i, and if $x^m(i)$, there exists exactly one S (denoted by S^m) for which $i \in S^m$ and $x^m(S^m) = 1$. Similarly, our construction implies that $y_{ij}^m = \sum_{S|ij \in S} x^m(S)$. To see this, note that $x_i^m = x_j^m = 1$ if and only if $\sum_{S|i,j \in S} x^m(S) = 1$ (as before if $x_i^m = 0$, then $x^m(S) = 0$ for all $i \in S$, and if $x_i^m = x_j^m = 1$ then there exists exactly one S, denoted by S^m such that $i, j \in S^m$ and $x^m(S^m) = 1$). On the other hand, it was established before that for $ij \in E$, we have $y_{ij}^m = 1$ if and only if $x_i^m = x_j^m = 1$. These imply that $y_{ij}^m = \sum_{S|ij \in S} x^m(S)$.

Using $x_i^m = \sum_{S|i \in S} x^m(S)$ and $y_{ij}^m = \sum_{S|ij \in S} x^m(S)$, the objective value corresponding to $\{x_i^m, y_{ij}^m\}$ in LP2 (given by $\sum_m \left(\sum_i w_i^m x_i^m + \sum_{ij \in E} w_{ij}^m y_{ij}^m \right)$) and that corresponding to $\{x^m(S)\}$ in LP1 (given by $\sum_{m,S} x^m(S) v^m(S)$) can be shown to be equal:

$$\sum_m \left(\sum_i w_i^m x_i^m + \sum_{ij \in E} w_{ij}^m y_{ij}^m \right) = \sum_m \left(\sum_i w_i^m \sum_{S|i \in S} x^m(S) + \sum_{ij \in E} w_{ij}^m \sum_{S|i,j \in S} x^m(S) \right)$$

$$= \sum_m \left(\sum_S x^m(S) \sum_{i \in S} w_i^m + \sum_S x_S^m \sum_{ij \in E} w_{ij}^m S \right) = \sum_{m,S} x^m(S) v^m(S).$$

Thus, we conclude that given a feasible integer solution of LP2, there exists a corresponding feasible integer solution of LP1, with the same objective value. Since this is true for the optimal integer solution of LP2 as well, we conclude $OPI2 \leq OPI1$.

Summarizing, we have $OP1 \leq OP2$, and $OPI2 \leq OPI1$. Additionally, optimal value is weakly higher without the integrality requirement (i.e., $OPI2 \leq OP2$, $OPI1 \leq OP1$) and if LP2 has an optimal integer solution, then $OP2 = OPI2$. These imply that $OP1 \leq OP2 = OPI2 \leq OPI1 \leq OP1$, and hence $OPI1 = OP1$. That is, when LP2 has an optimal solution that is integral, then so does LP1. $\qquad \square$

Proof of Theorem 3.1(iii). Theorem 3.1(i) and (ii) immediately imply that for sign-consistent tree valuations LP1 has integral optimal solutions. On the other hand, LP1 has such solutions if and only if a Walrasian equilibrium exists (Bikhchandani and Mamer, 1997). Hence, the claim follows. $\qquad \square$

Proof of Theorem 3.2. (i) Consider the optimal solution $\{\pi_i^m, p_i, p_{ij}^m, q_{ij}^{m,i}\}$ of DLP2 and integral optimal solution $\{x_i^m, y_{ij}^m\}$ of LP2. Let $S^m = \{i | x_i^m = 1\}$ be defined as in the theorem statement.

The complementary slackness (CS) conditions in LP2/DLP2 imply that:

C1: $\pi_i^m = w_i^m - p_i + \sum_{j|j\neq i} q_{ij}^{m,i} - \sum_{j|j\neq i} p_{ij}^m$ for $i \in S^m$,

C2: $q_{ij}^{m,i} + q_{ij}^{m,j} - p_{ij}^m = w_{ij}^m$ for $i,j \in S^m$,

C3: $q_{ij}^{m,i} = 0$ if $i \in S^m$, $j \notin S^m$,

C4: $p_{ij}^m = 0$ if $i,j \notin S^m$,

C5: $\pi_i^m = 0$ if $i \notin S^m$.

C6: $p_i = 0$ if $i \notin \cup_m S^m$.

Observe that the feasibility of the integral optimal solution $\{x_i^m, y_{ij}^m\}$ in LP2 implies that $y_{ij}^m = 1$ if and only if $x_i^m = x_j^m = 1$. This observation and the fact that $S^m = \{i | x_i^m = 1\}$ can be used to see the relation between the CS conditions and C1-C6. In particular, C1 corresponds to the CS condition associated with $x_i^m = 1$ at the optimal solution of LP2, and C2 captures the CS condition associated with $y_{ij}^m = 1$. C3 is relevant when the inequality $x_i^m \geq y_{ij}^m$ is strict at the optimal solution of LP2, whereas C4 and C5 respectively capture strict inequalities $x_i^m + x_j^m - 1 \leq y_{ij}^m$ and $x_i^m \leq 1$. Finally, C6 corresponds to the case, where $x_i^m = 0$ for all i, and hence the inequality $\sum_m x_i^m \leq 1$ is strict. These inequalities imply that

$$
\begin{aligned}
\sum_{i \in S^m} \pi_i^m &= \sum_{i \in S^m} \left(w_i^m - p_i + \sum_{j|j\neq i} q_{ij}^{m,i} - \sum_{j|j\neq i} p_{ij}^m \right) \\
&= \sum_{i \in S^m} (w_i^m - p_i) + \sum_{i,j \in S^m | i \neq j} \left(q_{ij}^{m,i} + q_{ij}^{m,j} - 2p_{ij}^m \right) + \sum_{i \in S^m, j \notin S^m} \left(q_{ij}^{m,i} - p_{ij}^m \right) \qquad (11) \\
&= \sum_{i \in S^m} (w_i^m - p_i) + \sum_{i,j \in S^m | i \neq j} w_{ij}^m - \sum_{i,j \in S^m | i \neq j} p_{ij}^m - \sum_{i \in S^m, j \notin S^m} p_{ij}^m.
\end{aligned}
$$

Here, the first equality follows from C1, the second one is obtained by rearranging terms, and the third one follows from C2, and C3. Rearranging terms (and observing that $v^m(S^m) = \sum_{i \in S^m} w_i^m + \sum_{i,j \in S^m} w_{ij}^m$), this implies that

$$
\sum_{i \in S^m} \pi_i^m + \sum_{i,j \in S^m | i \neq j} p_{ij}^m + \sum_{i \in S^m, j \notin S^m} p_{ij}^m = v^m(S^m) - \sum_{i \in S^m} p_i. \qquad (12)
$$

Consider an arbitrary bundle S of items. Using dual feasibility of the given optimal solution

we can obtain:

$$
\sum_{i \in S} \pi_i^m \geq \sum_{i \in S} \left(w_i^m - p_i + \sum_{j \mid j \neq i} q_{ij}^{m,i} - \sum_{j \mid j \neq i} p_{ij}^m \right)
$$

$$
= \sum_{i \in S} (w_i^m - p_i) + \sum_{i,j \in S \mid i \neq j} \left(q_{ij}^{m,i} + q_{ij}^{m,j} - 2p_{ij}^m \right) + \sum_{i \in S, j \notin S} \left(q_{ij}^{m,i} - p_{ij}^m \right) \qquad (13)
$$

$$
\geq \sum_{i \in S} (w_i^m - p_i) + \sum_{i,j \in S \mid i \neq j} w_{ij}^m - \sum_{i,j \in S \mid i \neq j} p_{ij}^m - \sum_{i \in S, j \notin S} p_{ij}^m.
$$

Here, the equality is obtained rearranging the terms, whereas, both inequalities follow from feasibility in DLP2. Rearranging terms, it follows from (13) that

$$
\sum_{i \in S} \pi_i^m + \sum_{i,j \in S \mid i \neq j} p_{ij}^m + \sum_{i \in S, j \notin S} p_{ij}^m \geq v^m(S) - \sum_{i \in S} p_i. \qquad (14)
$$

Since $\pi_i^m = 0$ if $i \notin S^m$, $p_{ij}^m = 0$ if $i,j \notin S^m$ (C4 and C5), and the dual variables of DLP2 are nonnegative we also have

$$
\sum_{i \in S^m} \pi_i^m + \sum_{i,j \in S^m} p_{ij}^m + \sum_{i \in S^m, j \notin S^m} p_{ij}^m \geq \sum_{i \in S} \pi_i^m + \sum_{i,j \in S} p_{ij}^m + \sum_{i \in S, j \notin S} p_{ij}^m.
$$

Using this together with (12) and (14), we conclude that $v^m(S^m) - \sum_{i \in S^m} p_i \geq v^m(S) - \sum_{i \in S} p_i$.

This implies that the third condition of the Walrasian equilibrium definition (Definition 2.2) holds for the given $\{p_i\}$ and $\{S^m\}$. Condition (i) and (ii) immediately hold by construction. The last condition follows from C6. Therefore, $\{p_i\}$ and $\{S^m\}$ constitute a Walrasian equilibrium.

(ii) Note that LP2 and DLP2 have polynomially many variables and constraints in the number of bidders and items. Thus, when the problem data (i.e., weights) are integral, they can be solved in polynomial-time in the number of bidders and items using an algorithm such as ellipsoid (Bertsimas and Tsitsiklis, 1997). The result follows since as established in part (i) the optimal solutions of these problems lead to a Walrasian equilibrium. $\qquad \square$

C   Proofs of Section 4

Proof of Lemma 4.2. We first establish (i). Then, we focus on the case $\theta^* = \theta_1$ and prove (ii)-(iv). Finally, we establish that the proof of (ii)-(iv) carries over to the case where $\theta^* = \theta_2$ with minor modifications.

(i) We prove the claim by contradiction. First assume that $\theta^* = 0$. Observe that θ_1 is defined as the minimum $\theta \geq 0$ such that a bidder demands a bundle at price vector $p + \theta \bar{p}$, which she does not demand at price vector p. This immediately implies that $\theta_1 > 0$, and hence $\theta_2 = \theta^* = 0$. On the other hand, $\theta_2 = 0$ implies that an arbitrarily small update of the dual solution (π, p, p_E, q_E) in the $(\bar{\pi}, \bar{p}, \bar{p}_E, \bar{q}_E)$ direction is infeasible in DLP2. This implies that $(\bar{\pi}, \bar{p}, \bar{p}_E, \bar{q}_E)$ cannot be

an improvement direction (or an optimal solution of RD). Thus, we obtain a contradiction, and conclude that $\theta^* \neq 0$.

Next assume that $\theta^* = \infty$. Since $\theta^* = \min\{\theta_1, \theta_2\}$, it follows that this case requires having $\theta_1 = \theta_2 = \infty$. Definition of θ_2 suggests that if $\theta_2 = \infty$, then $\bar{p} \geq 0$. If $\bar{p} = 0$, it can be checked from RD that the multiplication of the optimal solution $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ with a constant larger than one leads to another feasible solution for RD. Moreover, since the objective value of this problem is nonzero, this implies that either this solution has strictly lower objective value, or the optimal objective value of RD is $-\infty$. The former case contradicts with the optimality of $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ in RD. The latter implies that RP is infeasible. However, this contradicts with the assumption that the given dual solution (π, p, p_E, q_E) satisfies Property 4.1, as this property implies that a feasible solution of RP (with appropriate choice of γ, h) always exists (Lemma 4.1). Thus, we conclude that $\bar{p} \neq 0$, and it should be the case that $\bar{p} \geq 0$, and $\bar{p}_i > 0$ for some i.

Since $\theta_1 = \infty$ it follows that for any $\theta > 0$, at prices $p + \theta\bar{p}$, every bidder m demands some bundle $S \in D^m$ that is also demanded at price vector p. This implies that as θ increases the maximum surplus of bidder m does not change. This can be seen by noting that the maximum surplus of every bidder is weakly decreasing in θ since $\bar{p} \geq 0$. If the surplus is initially equal to zero, this implies that it cannot change as θ increases. If instead it is initially positive, and the maximum surplus strictly decreases with θ, eventually the empty bundle (which has surplus equal to zero) enters the demand set. Since this case contradicts with $\theta_1 = \infty$, it follows that the maximum surplus of any bidder m does not change with θ.

Consider an optimal solution (γ, h, x, y) of RP. The feasibility of this solution in RP requires satisfying C^m (i.e., a subset of the CS conditions in LP2/DLP2 (3)) for all m. On the other hand, it can be checked that C^m also coincides with the CS conditions in a formulation of $LP - D^m/DLP - D^m$ at price vector p. Thus, we conclude that the component (π^m, p_E^m, q_E^m) of the original dual solution (of DLP2), together with the restriction of the optimal solution of RP to $\{x_i^m, y_{ij}^m\}_{i,j}$ satisfy the complementary slackness conditions in a formulation of $LP - D^m/DLP - D^m$ at price vector p. Moreover, since these solutions are derived from feasible solutions of DLP2 and RP, it can be checked that they are also feasible in $DLP - D^m/LP - D^m$. These observations imply that (π^m, p_E^m, q_E^m) is an optimal solution of $DLP - D^m$ at price vector p. Since the optimal objective of $LP - D^m$ is the maximum surplus of agent m at price vector p, strong duality implies that $\sum_i \pi_i^m + \sum_{i,j} p_{i,j}^m$ is the maximum surplus of this agent. This suggests that the objective value of $DLP2$ associated with dual solution (p, π, p_E, q_E) is the sum of the prices p and bidder surpluses at these prices.

Since $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ is an optimal solution of RD, for sufficiently small $\epsilon > 0$, $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ is a feasible solution of DLP2 with strictly lower objective value than the original dual feasible solution (p, π, p_E, q_E). Consider the price vector $p + \epsilon\bar{p}$. Let $(\tilde{\pi}^m, \tilde{p}_E^m, \tilde{q}_E^m)$ denote an optimal solution of $DLP - D^m$ at this price vector, and $(\tilde{\pi}, \tilde{p}_E, \tilde{q}_E)$ be the collection of such solutions for all m. Since $DLP - D^m$ (for all m) and $DLP2$ share the same set of constraints, it follows that $(p + \epsilon\bar{p}, \tilde{\pi}, \tilde{p}_E, \tilde{q}_E)$ is a feasible solution of DLP2. Moreover, since $(\tilde{\pi}^m, \tilde{p}_E^m, \tilde{q}_E^m)$

is optimal in $DLP - D^m$ for all m (and the objective of DLP2 is sum of prices and objectives of $DLP - D^m$ for all m), it follows that $(p + \epsilon \bar{p}, \tilde{\pi}, \tilde{p}_E, \tilde{q}_E)$ has weakly lower objective value in DLP2 than $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$. On the other hand, since the surplus of bidders at price vector $p + \theta \bar{p}$ does not change with θ (as established above), it follows that the objective value of $DLP - D^m$ is the same for price vectors p and $p + \epsilon \bar{p}$. This implies that the objective value of $(p + \epsilon \bar{p}, \tilde{\pi}, \tilde{p}_E, \tilde{q}_E)$ is equal to that of the potentially infeasible solution $(p + \epsilon \bar{p}, \pi, p_E, q_E)$. Since, $\bar{p} \geq 0$ and $p_i > 0$ for some i, this implies that $(p + \epsilon \bar{p}, \pi, p_E, q_E)$ has strictly higher objective value than (p, π, p_E, q_E). Since the objective value of the former solution is equal to that of $(p + \epsilon \bar{p}, \tilde{\pi}, \tilde{p}_E, \tilde{q}_E)$, which has objective value weakly lower than that of $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$, it follows that $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ has strictly higher objective value than (p, π, p_E, q_E). This contradicts with the fact that $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ is an improvement direction, and the claim (i) follows.

We next assume that $\theta^* = \theta_1$ and prove (ii)-(iv). Subsequently, we also extend this proof to the case of $\theta^* = \theta_2$.

(ii) Consider the tuple $(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ specified in the lemma. Since $(\tilde{\pi}^m, \tilde{p}_E^m, \tilde{q}_E^m)$ is constructed from $DLP - D^m$ (at price vector $p + \theta^* \bar{p}$) and this problem shares identical constraints with DLP2, it immediately follows that $(p + \theta^* \bar{p}, \tilde{\pi}, \tilde{p}_E, \tilde{q}_E) = (p, \pi, p_E, q_E) + \theta^*(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ satisfies all constraints of DLP2, but $p + \theta^* \bar{p} \geq 0$. It can be seen that the latter constraint also holds since by definition $\theta^* \leq \theta_2$. Thus, $(p, \pi, p_E, q_E) + \theta^*(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is feasible in DLP2. Since the set of feasible solutions is convex, this also implies that for any $\epsilon \in [0, \theta^*]$, $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is feasible in this optimization problem. Note that by definition of $\bar{p}$ we also have $-1 \leq \bar{p} \leq 1$. On the other hand, the feasible set of RD is equivalent to the set of feasible update directions (associated with the original solution of DLP2) for which $-1 \leq \bar{p} \leq 1$. These observations imply that $(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is a feasible solution of RD.

Since (p, π, p_E, q_E) is an acceptable solution, it follows that there exists a feasible solution to LP2 that satisfies the constraints C^m, F^m associated with this solution. On the other hand, as explained in the first part of the proof, C^m also coincides with the CS conditions in a formulation of $LP - D^m / DLP - D^m$ at price vector p. Thus, it follows that (π^m, p_E^m, q_E^m) is a solution of $DLP - D^m$ at price vector p. Since the optimal objective of this problem (and the associated primal problem $LP - D^m$) corresponds to the surplus bidder m has at price p, we conclude that $\sum_i \pi_i^m + \sum_{ij} p_{ij}^m$ gives bidder m's surplus at price p. Similarly, the construction of $(\tilde{\pi}, \tilde{p}_E, \tilde{q}_E)$ implies that at price vector $p + \theta^* \bar{p}$, the quantity $\sum_i \tilde{\pi}_i^m + \sum_{ij} \tilde{p}_{ij}^m$ captures the surplus of bidder m. Thus, the definition of $(\hat{\pi}, \hat{p}_E, \hat{q}_E)$ implies that $\theta^*(\sum_i \bar{\pi}_i^m + \sum_{ij} \bar{p}_{ij}^m)$ corresponds to the total change in bidder m's surplus as prices are updated from p to $p + \theta^* \bar{p}$. Moreover, since $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is feasible for all $\epsilon \in [0, \theta^*]$, and for each bidder one of the bundles demanded at price p remain demanded until $p + \theta^* \bar{p}$ (by definition of θ^*), the objective value of DLP2 corresponding to $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ captures the sum of the prices and bidders' surplus at price $p + \epsilon \hat{p}$.

Consider feasible solutions of DLP2 at price vector $p + \epsilon \bar{p}$ for sufficiently small $\epsilon > 0$. The structure of DLP2 and $DLP - D^m$ suggest that the feasible solution with the lowest objective value can be obtained by solving $DLP - D^m$ at price vector $p + \epsilon \bar{p}$ for all m, and combining the

57

corresponding optimal solutions to construct a solution to DLP2. Thus, the corresponding objective value of this solution is equal to the sum of the prices $p + \epsilon\bar{p}$, and surpluses of all bidders m at these prices. Hence, the objective value of $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ is lower bounded by this quantity as well. On the other hand, this lower bound is achieved by updating the original dual solution in the $(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ direction. Since $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ is an optimal solution of RD (hence the dual update direction that leads to the best improvement in the objective of DLP2), it follows that $(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ (which leads to a weakly better improvement) is also optimal in RD.

(iii) As established in part (ii), $(p, \pi, p_E, q_E) + \epsilon(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is feasible in DLP2 for any $\epsilon \in [0, \theta^*]$. Assume that a stepsize $\bar{\theta} > \theta^* = \theta_1$ is chosen. To prove the claim it suffices to show that $(p, \pi, p_E, q_E) + \bar{\theta}(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is not feasible in DLP2.

Observe that the bundles that are demanded can be found through integral optimal solutions of $LP - D^m$. Complementary slackness between $DLP - D^m$ and $LP - D^m$ suggests that if a new set S enters the demand set, it should be the case that in optimal solutions of $DLP - D^m$ (at price vector $p + \theta\bar{p}$), a constraint that was not active for $\theta < \theta^*$ starts to become active after $\theta = \theta^*$. On the other hand, this implies that for the given update direction and $\theta > \theta^*$, the aforementioned constraint is violated. Since $DLP - D^m$ and DLP2 share the same constraints, we conclude that $(p, \pi, p_E, q_E) + \theta(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is not feasible in DLP2 for $\theta > \theta^*$. Thus, θ^* is the largest stepsize that preserves feasibility, and the claim follows.

(iv) By construction, for all $m \in \mathcal{M}$, $(\pi^m, p_E^m, q_E^m) + \theta^*(\hat{\pi}^m, \hat{p}_E^m, \hat{q}_E^m) = (\tilde{\pi}^m, \tilde{p}_E^m, \tilde{q}_E^m)$ is an optimal solution of $DLP - D^m$ associated with the price vector $p + \theta^*\hat{p}$. Complementary slackness implies that $LP - D^m$ has a corresponding optimal solution for which the conditions C^m, F^m (formulated at $p + \theta^*\hat{p}$) hold (where C^m corresponds to the CS conditions in $LP - D^m/DLP - D^m$, and F^m corresponds to the feasibility constraints in $LP - D^m$). Since this is true for all m, it follows that the collection of such solutions of $LP - D^m$ for all $m \in \mathcal{M}$ satisfies Property 4.1. Hence, the claim follows.

Next consider the case $\theta^* = \theta_2$, and use the same construction for $(\hat{\pi}, \hat{p}_E, \hat{q}_E)$. Observe that the proof of (ii) and (iv) given above did not make use of the exact value of θ^*, hence imply the claim in this case as well. In order to complete the proof it suffices to prove (iii). Since by construction $(p, \pi, p_E, q_E) + \theta^*(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is feasible in DLP2, it suffices to show that for $\theta > \theta^*$, $(p, \pi, p_E, q_E) + \theta(\bar{p}, \hat{\pi}, \hat{p}_E, \hat{q}_E)$ is infeasible in this optimization problem. However, this immediately follows since $\theta^* = \theta_2$, and the definition of θ_2 suggests that for $\theta > \theta^*$ we have $(p + \theta\bar{p})_i < 0$ for some i. Thus, we conclude that $\theta^* = \theta_2$ is equal to the stepsize the primal-dual algorithm associates with the given improvement direction, hence (ii) also follows in this case. $\square$

Proof of Lemma 4.4. (i) First assume that $S_1 = S_2 = S$. Observe that in this case, the change in the surplus of the bidder can be simply expressed as $\pi_2 - \pi_1 = -\sum_{i \in S} \theta\bar{p}_i$. On the other hand, substituting $S_1 = S_2 = S$ it follows that the right hand side of the expression in the claim is also given by $\sum_{i \in S} p_i - (\sum_{i \in S}(p_i + \theta\bar{p}_i)) - \lfloor \sum_{i \in S}(p_i + \theta\bar{p}_i) - \sum_{i \in S}(p_i + \theta\bar{p}_i) \rfloor = -\sum_{i \in S} \theta\bar{p}_i$. Hence, the claim follows.

Next assume that $S_1 \neq S_2$. In this case, bundle S_1 is demanded before the price update, S_2 is

demanded afterwards, and these bundles are not jointly demanded. This suggests that $v^m(S_1) - \sum_{i \in S_1} p_i > v^m(S_2) - \sum_{i \in S_2} p_i$, and $v^m(S_2) - \sum_{i \in S_2}(p_i + \theta \bar{p}_i) > v^m(S_1) - \sum_{i \in S_1}(p_i + \theta \bar{p}_i)$. Note that since $|\bar{p}_i| \leq 1$ and $\theta \leq 1/N$, these inequalities imply that

$$1 > \left(v^m(S_2) - \sum_{i \in S_2}(p_i + \theta \bar{p}_i) \right) - \left(v^m(S_1) - \sum_{i \in S_1}(p_i + \theta \bar{p}_i) \right) > 0. \tag{15}$$

Observe that for any real number $a \in (0,1)$, we have $a = a - \lfloor a \rfloor$. Using this identity, and canceling out integral terms, the quantity in the middle in (15) (or the difference between the surpluses of bundles S_1, S_2 at price vector $p + \theta \bar{p}$), can equivalently be expressed as

$$\left(v^m(S_2) - \sum_{i \in S_2}(p_i + \theta \bar{p}_i) \right) - \left(v^m(S_1) - \sum_{i \in S_1}(p_i + \theta \bar{p}_i) \right)$$
$$= \sum_{i \in S_1}(p_i + \theta \bar{p}_i) - \sum_{i \in S_2}(p_i + \theta \bar{p}_i) - \lfloor \sum_{i \in S_1}(p_i + \theta \bar{p}_i) - \sum_{i \in S_2}(p_i + \theta \bar{p}_i) \rfloor. \tag{16}$$

On the other hand, $\pi_2 - \pi_1$ can be expressed as the sum of the difference between the surpluses of bundles S_1, S_2 at price vector $p + \theta \bar{p}$, and the change in the surplus of bundle S_1 between prices p and $p + \theta \bar{p}$. Since the first term is as given in (16), and the second one can be expressed as $-\theta \sum_{i \in S_1} \bar{p}_i$, it follows that $\pi_2 - \pi_1 = \sum_{i \in S_1} p_i - \sum_{i \in S_2}(p_i + \theta \bar{p}_i) - \lfloor \sum_{i \in S_1}(p_i + \theta \bar{p}_i) - \sum_{i \in S_2}(p_i + \theta \bar{p}_i) \rfloor$. Hence, the claim follows.

(ii) We start by showing that two of the termination conditions of the subroutine never hold: (a) in Step S1c the condition $\theta_2 = \infty$ and $S \cap I = \emptyset$ for all $m, S \in \hat{D}^m$, and (b) in Step S2 the condition $\theta \leq 0$. We then show that the remaining termination conditions ensure that if the subroutine terminates, it does so with θ^*. Finally, we complete the proof by establishing that finite termination occurs.

Assume by contradiction that the aforementioned condition in Step S1c holds, i.e., for some $\theta \geq 0$ the conditions in S1a and S1b do not hold, $\theta_2 = \infty$ and, we have $S \cap I = \emptyset$ for all $m, S \in \hat{D}^m$. From the definition of θ_2 it follows that $\bar{p} \geq 0$ (as otherwise $\theta_2 < \infty$). Together with $S \cap I = \emptyset$ this implies that if θ is further increased, bidders' demand sets do not change. Moreover, since bidders' surplus for any bundle is linear in the prices and the condition of S1b does not hold (i.e., no new bundle is demanded) at θ, for any ϵ such that $\theta \geq \epsilon > 0$, bidders demand set is a superset of $\hat{D}^m$. Together these facts imply that no new bundle is demanded for any $\epsilon > 0$ at price vector $p + \epsilon \bar{p}$. Since $\theta_2 = \infty$, this implies that $\theta^* = \infty$, contradicting Lemma 4.2 (i) (note that in the first part of this lemma Property 4.1 is assumed only to guarantee that RP is feasible, which trivially holds here as an optimal solution of RP is given). Thus, we conclude that the aforementioned condition in S1c cannot hold.

We next establish that the termination condition $\theta \leq 0$ of Step S2 never holds. Consider the first iteration t^* at which the condition of S1a holds for some bidder m, and S2 is reached. Let $\theta(t)$ denote the value of θ at iteration t, and $\hat{D}^m$ denote the set of bundles demanded by bidder m at

price vector $p + \theta(t^*)\bar{p}$. At Step S2, in determining $\hat{\theta}$ two sets are constructed: S_1, S_2. $S_1 \in D^m$ is a bundle that was demanded at price vector p, and whose surplus $v^m(S) - \sum_{i \in S}(p_i + \theta\bar{p}_i)$ decreases the slowest when the prices are updated in the $\bar{p}$ direction. On the other hand, S_2 is a bundle that is demanded at $\hat{D}^m$, and whose surplus increases the slowest when the same price update direction is used. Note that these definitions guarantee that for $\theta \in [0, \theta(t^*)]$, at prices $p + \theta\bar{p}$, the surplus of S_1 (S_2) is weakly larger than the surplus of any bundle in D^m ($\hat{D}^m$). Since S1a does not hold prior to t^* this implies that S_1 is demanded at price vector $p + \theta(t^* - 1)\bar{p}$

Observe that since at t^* S1a holds, for optimal solutions of RD $|\bar{p}_i| \leq 1$, and prior to t^* θ is updated with a stepsize bounded by $1/N$ (as determined by Step S1 of the subroutine), it follows that S_1 and S_2 satisfy the assumptions of part (i) for prices $p + \theta(t^* - 1)$ and $p + \theta(t^*)$ respectively. This implies that difference of the maximum surplus before and after the price updates is given by

$$\left(v^m(S_2) - \sum_{i \in S_2}(p_i + \theta(t^*)\bar{p}_i)\right) - \left(v^m(S_1) - \sum_{i \in S_1}(p_i + \theta(t^* - 1)\bar{p}_i)\right)$$
$$= \left(\sum_{i \in S_1}(p_i + \theta(t^* - 1)\bar{p}_i) - \sum_{i \in S_2}(p_i + \theta(t^*)\bar{p}_i)\right) - \left\lfloor \sum_{i \in S_1}(p_i + \theta(t^*)\bar{p}_i) - \sum_{i \in S_2}(p_i + \theta(t^*)\bar{p}_i)\right\rfloor.$$

Rearranging terms, this implies that

$$\left(v^m(S_2) - \sum_{i \in S_2}(p_i + \theta(t^*)\bar{p}_i)\right) - \left(v^m(S_1) - \sum_{i \in S_1}(p_i + \theta(t^*)\bar{p}_i)\right)$$
$$= \left(\sum_{i \in S_1}(p_i + \theta(t^*)\bar{p}_i) - \sum_{i \in S_2}(p_i + \theta(t^*)\bar{p}_i)\right) - \left\lfloor \sum_{i \in S_1}(p_i + \theta(t^*)\bar{p}_i) - \sum_{i \in S_2}(p_i + \theta(t^*)\bar{p}_i)\right\rfloor.$$

On the other hand, by changing prices in the $-\bar{p}$ direction (or stepping back) the difference between the surpluses of S_1 and S_2 decreases at a rate of $\Delta = (\sum_{i \in S_1}\bar{p}_i - \sum_{i \in S_2}\bar{p}_i)$. Thus, at the price vector $p + (\theta(t^*) - \hat{\theta})\bar{p}$ (where $\hat{\theta}$ is as specified in the subroutine), bidder m is indifferent between bundles $S_1 \in D^m$ and $S_2 \in \hat{D}^m$. It can be seen that the same argument also holds for subsequent time instants where the condition of Step S1a holds and Step S2 is reached, i.e., price updates make bidder m indifferent between S_1 and S_2. This implies that the updated value of θ at Step S2 is between the value at time instant t^* (where S_2 is demanded but S_1 is not), and $t^* - 1$ (where S_1 is demanded). Hence, it is larger than the value of θ at $t^* - 1$. Since this is true for all stages and after the first price update of the subroutine $\theta > 0$, it follows that the condition $\theta \leq 0$ of Step S2 can never hold.

These observations imply that for the subroutine to terminate the condition of S1b or the condition $\theta = \theta_2$ in S1c needs to hold. In both cases (a) S1a does not hold, and hence at price vector $p + \theta\bar{p}$ all bidders demand a bundle that they demanded at the original price vector p, (b) $p + \theta\bar{p} \geq 0$ as $\theta \leq \theta_2$. Additionally, the condition S1b suggests that at price vector $p + \theta\bar{p}$ a bidder demands a bundle that was not demanded at price vector p. Thus, Lemma 4.2 implies that when S1b holds θ is equal to the primal-dual stepsize θ^* associated with the given update direction. On the other hand, if S1c holds, then $\theta = \theta_2$, and hence the price of an item decreases to zero, but no bidder demands a new bundle (since S1b does not hold). It follows from Lemma 4.2 that in this

case $\theta = \theta^*$ as well. Thus, to complete the proof, it suffices to show that the subroutine terminates in finite time.

We prove the claim by contradiction. Assume that the subroutine does not terminate in finite time. Then, the conditions in S1b and S1c cannot hold. If S1a also does not hold, then a bundle $S \in D^m$ is demanded after every price update, and prices continue to be updated in the $\bar{p}$ direction indefinitely with a stepsize of $1/N$. On the other hand, Lemma 4.2 implies that $\theta^* < \infty$, i.e., after finitely many such steps either a new bundle enters the demand set of a bidder, or the price of an item decreases to zero. This implies that S1b or S1c eventually holds, and we obtain a contradiction. Hence, it follows that the condition of S1a eventually holds.

Let $\bar{\theta}$ denote the value of the θ parameter at the first iteration S1a holds, and $S_1 \in D^m$ and $S_2 \in \hat{D}^m$ be defined as in Step S2 of the subroutine. Recall that for $\theta \in [0, \bar{\theta}]$ at prices $p + \theta\bar{p}$, the surplus of S_1 is weakly larger than the surplus of any bundle in D^m. Thus, if a bundle in D^m is demanded, S_1 should be demanded as well. Note that at price vector $p + (\bar{\theta} - \hat{\theta})\bar{p}$ if for all m a bundle $S \in D^m$ is demanded, then in the next iteration the condition of Step S1b holds (since at least one bidder is indifferent between $S_1 \in D^m$ and S_2, hence demands a new bundle $S_2 \notin D^m$) and Step S3 is reached. In this case, we reach a contradiction. If this is not the case, in the next iteration for at least one bidder, none of the bundles in D^m is demanded. Hence, once more Step S2 is reached, and additional correction is necessary for the prices. Note that after each price correction a bundle that has higher surplus than all elements of D^m (S_2 in our original construction) is guaranteed to have weakly lower surplus than some elements of D^m. Since, the prices continue to be updated in $-\bar{p}$ direction as long as Step S2 is reached, and there are finitely many bundles, this implies that after finitely many corrections, every bidder m demands a bundle that also belongs to the original demand set D^m. Note that this implies that after finitely many corrections step S2 (and hence S1a) cannot be reached. However, as explained earlier, in this case S1b is reached, and we obtain a contradiction. Thus, we conclude that this subroutine terminates after finitely many iterations. $\qquad\square$

Proof of Theorem 4.1. Given prices p, consider an optimal solution (π^m, p_E^m, q_E^m) of $DLP - D^m$. Since DLP2 and $DLP - D^m$ share identical constraints, it follows that (p, π, p_E, q_E) constitutes a feasible solution of DLP2 (where $\pi = \{\pi^m\}, p_E = \{p_E^m\}, q_E = \{q_E^m\}$). Moreover, this dual solutions satisfies Property 4.1. This can be seen noting that the optimal solution of $DLP - D^m$ satisfies complementary slackness conditions with an optimal solution of $LP - D^m$ (which correspond to C^m in RP), and feasibility of this optimal solution in $LP - D^m$ also guarantees that F^m constraints are satisfied. Thus the collection of such optimal solutions for all m satisfies the conditions given in Property 4.1.

At each step of the algorithm, with a given price vector p we associate a dual feasible solution of DLP2, obtained by the solution of $DLP - D^m$ as described above. Observe that the compact demand reports at Step S1 reveal the active constraints of DLP2 corresponding to this solution. Thus, Lemma 4.1 implies that RP/RD can be used to check the CS condition associated with this solution, and obtain an improvement direction, if they have nonzero objective value.

Assume that the algorithm terminates, i.e., Step S3 is reached. This implies that in Step S1, RP/RD has objective value zero. Hence, Lemma 4.1 implies that the dual solution of DLP2 described above is optimal, and the optimal solution of RP gives an optimal solution of LP2. Moreover, for sign-consistent tree valuations, RP and LP2 have integral optimal solutions as well (Theorem 3.1 and Lemma 4.1). Theorem 3.2 implies that the prices at the optimal solution of DLP2, and the allocation suggested by the integral optimal solution of LP2 (or RP) constitute a Walrasian equilibrium. Thus, it follows that if the algorithm terminates, then a Walrasian equilibrium is reached. To complete the proof it suffices to show that the algorithm terminates in finitely many iterations.

Assume that for a price vector p, the associated dual solution of DLP2 is such that RP/RD has nonzero objective value, and hence Step S1 is followed by Step S2. In this case, Lemma 4.1 implies that $(\bar{p}, \bar{\pi}, \bar{p}_E, \bar{q}_E)$ of Step S1 is an improvement direction corresponding to this dual solution. Lemma 4.4, on the other hand, suggests that θ^* identified in Step S2 corresponds to the primal-dual stepsize given in Lemma 4.2. Observe that Lemma 4.2 suggests that θ^* and a dual optimal solution of RD constitute a valid primal-dual stepsize and an improvement direction. Moreover, updating the dual solution according to this stepsize and improvement direction corresponds to setting prices to $p + \theta^*\bar{p}$, and constructing (π^m, p_E^m, q_E^m) according to an optimal solution of $DLP - D^m$. That is, the dual updates suggested by Lemma 4.2 are identical to the dual feasible solution our algorithm associates with the updated price vector $p + \theta^*\bar{p}$. Since, these updates are identical to those in an application of the primal-dual algorithm to LP2/DLP2, it follows that they converge to an optimal solution of DLP2 in finite time (Proposition 4.1). In such a solution, RP/RD has objective value zero. Thus, it follows that after finitely many iterations, Step S1 of Algorithm 2 leads to Step S3, and the algorithm terminates. Hence, the claim follows. □

D   Proofs of Section 5

Proof of Theorem 5.1. (i) Let $\mathcal{S}$ denote the set of $M - 1$ bidder markets (potentially equal to $\emptyset$) that are not cleared by H_t. We first show that the stepsize computation subroutine terminates in finite time provided that bidders bid truthfully after H_t. Then, we show that the price updates obtained using the subroutine coincide with those in the primal-dual algorithm (Algorithm 1) for some market $\mathcal{E} \in \mathcal{S} \cup \{\mathcal{E}_\emptyset\}$, and lead to finite convergence.

We start by showing that θ cannot decrease/increase unboundedly in the subroutine. It immediately follows from Step S2 (i.e., the only step where θ can be decreased) of the subroutine that $\theta \geq 0$, and hence cannot decrease unboundedly. Assume by contradiction that θ increases unboundedly. In this case, the subroutine needs to reach S1d (the step which leads to an increase in θ) infinitely often. If $\theta_2 < \infty$, or if $\theta_2 = \infty$, and $S \cap I = \emptyset$ for all $S \in \hat{D}^m$, S1c implies that the subroutine eventually terminates, and S1d is not reached. Thus, due to our assumption, these cannot hold, i.e., $\theta_2 = \infty$, and $S \cap I \neq \emptyset$ for some $S \in \hat{D}^m$. Together with the definition of θ_2 and I these imply that $\bar{p} \geq 0$, and $\bar{p}_i > 0$ for some $i \in S$. Hence, when θ increases the price of bundle

S increases. On the other hand, such a bundle cannot stay demanded as θ increases unboundedly. Thus, eventually $S \notin \hat{D}^m$, the termination condition in S1c holds, and the subroutine terminates. Hence, we obtain a contradiction, and θ cannot increase unboundedly.

Thus, to prove that the subroutine terminates in finite time, it suffices to show that θ cannot increase after decreasing. Observe that for θ to decrease the condition in Step S1a of the subroutine needs to hold, i.e., $\hat{D}^m \cap D^m = \emptyset$ for some m. On the other hand, if θ starts increasing following a decrease, then S1a cannot hold, i.e., some bundle $S \in D^m$ enters the demand set $\hat{D}^m$ as a result of decreasing this parameter. Moreover, this implies that $D^m \cap \hat{D}^m \neq \emptyset$, and $\hat{D}^m - D^m \neq \emptyset$ after the last step where θ is decreased (since as shown in Lemma 4.4 decreasing θ in Step S2 makes bidder m indifferent between a bundle in D^m and one that belongs to $\hat{D}^m$ before the price decrease). However, S1b implies that the subroutine terminates in this case. Hence, we conclude that the subroutine terminates in finite time (after any history) provided that bidders truthfully reveal their demand.

We next show that in Algorithm 3 the price updates obtained from the subroutine lead to finite convergence. Steps S1 and S2 of Algorithm 3 imply that the price updates are identical to those in Algorithm 2 for a market $\mathcal{E}_{m^*}$ identified in Step S1b. On the other hand, by construction the prices that appear in Algorithm 2 (after each run of the subroutine) coincide with those in Algorithm 1 (see Theorem 4.1 and Section 4.3.3). Since the latter is a primal-dual algorithm, it follows that from time instant $t + 1$ onwards the prices that are obtained coincide with an application of the primal-dual algorithm to solution of LP2 associated with the market $\mathcal{E}$ that is closest to getting cleared. In primal-dual algorithms the optimal objective of the restricted primal (weakly) decreases at each update, and a different extreme point of the polytope (obtained by replacing the equality constraints in the restricted primal with inequality constraints) is visited (see Proposition 4.1 and Section 4.1). These facts suggest that provided that an optimal solution is not achieved, the optimal objective value of the restricted primal decreases strictly after finitely many updates. This implies that in our auction the objective value of the restricted primal associated with the market that is closest to getting cleared decreases strictly after finitely many updates (since otherwise the updates are made with respect to the same market, and decrease in finitely many iterations is guaranteed for primal-dual algorithms). Additionally, since there are finitely many extreme points of the aforementioned polytopes, there are finitely many different values the optimal objective values of restricted primals can take. Hence, it follows that all markets in $\mathcal{S}$ clear after finitely many runs of the subroutine. Since this subroutine terminates in finitely many steps (when bidders truthfully reveal their demand), it follows that all such markets clear in finite time. In addition, the same argument can be repeated for the market that consists of all bidders in $\mathcal{M}$ (following Step S3 in Algorithm 3), suggesting that all markets clear in finite time. The efficiency of the associated allocations follows from the fact that market clearance points/Walrasian equilibria (associated with subsets of bidders) are efficient (see Section 2.2).

(ii) Part (i) implies that when bidders bid truthfully the auction terminates. Observe that until termination q captures the sum of the payments of any given bidder. On the other hand, Algorithm

3 suggests that when the auction terminates, bidder m receives a rebate of $q - z^m$. This implies that z^m is the total payment of bidder m at the end of the auction.

Algorithm 3 sets $z^m = 0$ until market $\mathcal{E}_m$ clears. Let t_1 denote the time at which this market clears. Step S1 implies that at t_1, z^m is set equal to $\sum_i p_i(t_1)$, where $p_i(t)$ denotes the price of item i at time t. Moreover, after $\mathcal{E}_m$ clears, at each price update until termination, z^m is updated following Step S2, i.e., $z^m := z^m - \sum_{k \neq m} \Delta \pi^k(t)$, where $\Delta \pi^k(t)$ denotes $\Delta \pi^k$ (see Step S2) at stage t of the auction. Note that these updates cease after market $\mathcal{E}_\emptyset$ clears (step S3). Denote this time instant by t_2. After the update at t_2 (see Step S2) z^m is given by $z^m = \sum_i p_i(t_1) - \sum_{t=t_1}^{t_2-1} \sum_{k \neq m} \Delta \pi^k(t)$. Finally, after Step S3 is complete, z^m is updated to $z^m := z^m - \sum_{i \notin S^m} p_i$. This implies that at termination we have $z^m = \sum_i p_i(t_1) - \sum_{t=t_1}^{t_2-1} \sum_{k \neq m} \Delta \pi^k(t) - \sum_{i \notin S^m} p_i$. Observe that since at time t_1 and t_2, markets $\mathcal{E}_m$ and $\mathcal{E}_\emptyset$ clear respectively, it follows that $\sum_i p_i(t_1) = \sum_{k \neq m} \sum_{i \in S_1^k} p_i(t_1)$ and $\sum_{i \notin S^m} p_i(t_2) = \sum_{k \neq m} \sum_{i \in S^k} p_i(t_2)$, where $\{S_1^k\}_{k \neq m}$ and $\{S^m\}$ denote the market clearing allocations associated with t_1 and t_2. This suggests that z^m can be rewritten as

$$z^m = \sum_{k \neq m} \sum_{i \in S_1^k} p_i(t_1) - \sum_{t=t_1}^{t_2-1} \sum_{k \neq m} \Delta \pi^k(t) - \sum_{k \neq m} \sum_{i \in S^k} p_i(t_2). \tag{17}$$

When bidders truthfully reveal their demand, $\Delta \pi^k(t) = \pi^k(t+1) - \pi^k(t)$, where $\pi^k(t)$ denotes the maximum surplus of bidder k at stage t. Thus, Lemma 5.1 implies that z^m is equal to bidder m's VCG payment associated with bundle S^m. Observe that by construction q is larger than $\sum_{i \in S^m} p_i(t)$ for all m, t. This implies that at termination q is larger than the Walrasian equilibrium payments. On the other hand, it is known that the Walrasian equilibrium payments are larger than the VCG payments (Gul and Stacchetti, 1999). Thus, the claim follows, and each bidder has a nonnegative rebate.

(iii) Consider a history H_{t^*}. Assume that after this history, all agents other than m truthfully reveal their demand at each stage of the auction. We next show that bidder m maximizes her payoff by truthfully revealing her demand after H_{t^*}. There are two cases to consider (a) the market for $\mathcal{M} - \{m\}$ did not clear by time t^*, (b) it did. If after t^* bidder m bids in a way that prevents the termination of the auction, then due to the payments in Step S2, bidders receive arbitrarily low payoffs. Hence, it follows that in both cases bidder m improves her payoff by following a strategy that leads to termination of the auction (existence of such strategies follows from part (i)).

First consider case (a). Part (i) implies that if bidder m bids truthfully, all the remaining markets clear and the auction terminates. Moreover, as established in part (ii) the total payment of bidder m is given by z^m, and since bidders $k \neq m$ bid truthfully this quantity can be expressed as in (17). Lemma 5.1 implies that this payment is equal to her VCG payment, and the corresponding payoff is given by $\sum_k v^k(S_*^k) - \max_{\{Z^k\}|Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) \geq 0$, where $\{S_*^k\}$ is the efficient allocation.

Assume that bidder m does not bid truthfully and the auction terminates with some allocation $\{\hat{S}^k\}$. Since the remaining bidders bid truthfully Lemma 5.1 implies that in this case bidder m's

total payment is equal to $\max_{\{Z^k\}|Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) - \sum_{k \neq m} v^k(\hat{S}^k)$. Thus, her payoff is given by $\sum_k v^k(\hat{S}^k) - \max_{\{Z^k\}|Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) \leq \sum_k v^k(S_*^k) - \max_{\{Z^k\}|Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k)$. It follows that bidder m's payoff is weakly lower than the payoff she obtains by bidding truthfully. Hence, she cannot deviate from truthful bidding after history H_t and improve her payoff.

Consider case (b). Let t_1, t^*, t_2 respectively denote the time instant at which the market $\mathcal{E}_m$ clears, bidders $k \neq m$ start bidding truthfully, and market $\mathcal{E}_\emptyset$ clears. Also define $Q_m \triangleq \sum_{k \neq m} \sum_{t=t_1}^{t_*-1} -\Delta\pi^k(t) + \sum_i p_i(t_1)$, where $\Delta\pi^k(t)$ denotes the value of $\Delta\pi^k$ (see Step S2) at stage t of the auction. Assume that the auction terminates with a final allocation $\{\hat{S}^m\}$ determined in Step S3 of Algorithm 3. In this case, the final payment z^m of bidder m is given by (17), and leads to a payoff of:

$$
\begin{aligned}
v^m(\hat{S}^m) - z^m &= v^m(\hat{S}^m) - \left(\sum_{t=t_1}^{t_2-1} \sum_{k \neq m} -\Delta\pi^k(t) - \sum_{i \notin \hat{S}^k} p_i(t_2) + \sum_i p_i(t_1) \right) \\
&= v^m(\hat{S}^m) - \left(\sum_{t=t_1}^{t^*-1} \sum_{k \neq m} -\Delta\pi^k(t) + \sum_{t=t^*}^{t_2-1} \sum_{k \neq m} \pi^k(t) - \pi^k(t+1) - \sum_{i \notin \hat{S}^k} p_i(t_2) + \sum_i p_i(t_1) \right) \\
&= v^m(\hat{S}^m) - \left(Q_m + \sum_{k \neq m} \pi^k(t^*) - \sum_{k \neq m} \pi^k(t_2) - \sum_{i \notin \hat{S}^k} p_i(t_2) \right),
\end{aligned}
\tag{18}
$$

where the second equation follows since bidders other than m bid truthfully after t^* and hence $\Delta\pi^k(t) = \pi^k(t+1) - \pi^k(t)$ (the true change in the surplus of agent k). The definition of market clearance suggests that bidder k demands bundle $\hat{S}^k$ at time t_2. This implies that $\pi^k(t_2) + \sum_{i \in S_*^k} p_i(t_2) = v^k(\hat{S}^k)$. Thus, (18) can alternatively be expressed as

$$
v^m(\hat{S}^m) - z^m = v^m(\hat{S}^m) - \left(Q_m + \sum_{k \neq m} \pi^k(t^*) - \sum_{k \neq m} v^k(\hat{S}^k) \right) = \sum_k v^k(\hat{S}^k) - Q_m - \sum_{k \neq m} \pi^k(t^*). \tag{19}
$$

On the other hand, if bidder m truthfully bids after t^*, by (i) the final allocation $\{S_*^m\}$ is efficient. Thus, in this case (19) suggests that her payoff is given by $v^m(S_*^m) - z^m = \sum_k v^k(S_*^k) - Q_m - \sum_{k \neq m} \pi^k(t^*)$. Since $\{S_*^m\}$ is the efficient allocation, it follows that $\sum_k v^k(S_*^k) - Q_m - \sum_{k \neq m} \pi^k(t^*) \geq \sum_k v^k(\hat{S}^k) - Q_m - \sum_{k \neq m} \pi^k(t^*)$. This implies that bidder m receives a (weakly) higher payoff when she bids truthfully. Thus, we conclude that after history H_t the bidder has no incentive to deviate from truthful bidding. Since this is true for any history H_t, and bidder m, it follows that truthful bidding is an ex-post perfect equilibrium.

Finally, part (i) suggests that at this equilibrium the final allocation is achieved in finite time, and the final allocation (which is the market clearing allocation for market $\mathcal{E}_\emptyset$) is efficient. In addition, part (ii) implies that the final payments are the VCG payments, and the claim follows. $\square$