Assignment Preferences and Combinatorial Auctions

Robert W. Day
Operations and Information Management, School of Business,
University of Connecticut, Storrs, CT 06269-1041, Bob.Day@business.uconn.edu

S. Raghavan
The Robert H. Smith School of Business,
University of Maryland, College Park, MD 20742-1815, raghavan@umd.edu

Abstract

The bipartite matching problem and the closely related assignment problem form a cornerstone of the Networks literature. In this paper we consider the problem of a firm whose valuation function for bundles of goods is determined by an optimal assignment of goods among several agents (or tasks,) each with independent valuations for each good.

Though this model of preferences may be useful on its own for several particular markets, the underlying assignment network structure also leads us to an easy to use first approximation of preferences in more general settings. We show that because of its computational tractability and polynomial growth of user input, the resulting restricted bid language (which we call bid tables) is useful as a dynamic revelation phase in a multi-stage hybrid auction. We demonstrate that this representation is one for which linear price signals are appropriate, and show how lowest Walrasian equilibrium prices can be easily computed for preferences that can be represented as bid tables, generalizing the earlier work of Demange et al. (1986).

1 Introduction

Bipartite matching problems, and the related assignment and transportation problems form a cornerstone of the Networks literature. The economics of a bipartite set of players whose preferences are governed by matchings over ranked lists (called the marriage problem) have been studied extensively; Roth and Sotomayor (1992) provide a thorough overview of this material. Additionally, they consider the problem represented in Figure 1, in which players on one side of the market do not just rank the nodes on the opposite side of the market (to which they may be potentially matched), but instead provide a numerical weight to each potential match. For the central decision-maker the problem is now to find a bipartite matching of maximum weight, also know as the assignment problem.

As initiated in the earlier work of Demange et al. (1986), the economics of this model corresponds to a multi-item auction with unit-demand bidders, each of whose preferences can be completely described by a vector of length $N =$ the number of items being auctioned. The relevant results from this stream of literature tell us that any set of dual prices forms what is called a Walrasian equilibrium (defined formally in §5), and that the well-known Hungarian algorithm finds the unique bidder-optimal or minimal Walrasian equilibrium

Figure 1: A Bipartite Assignment Game with Unit-Demand Bidders. The bid values that a single bidder assigns to each item on the left are shown. Arcs and bid values from each item to each of the other bidders are omitted.

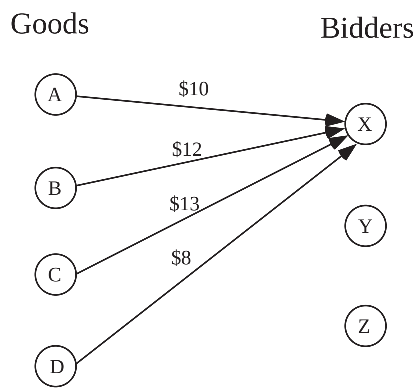

prices for the case of unit-demand bidders.

In this paper we study many of the same economic properties as Roth and Sotomayor (1992) and Demange et al. (1986), but we move from the unit-demand model of Figure 1 to the more general network depiction of Figure 2. In this new model, players are represented by several unit-demand agents on one side of the market (rather just a single unit-demand bidder), and experience the sum of the utilities each of their own agents gets for being matched to an item (node on the left side of the market.)

It is not hard to demonstrate computationally that the Hungarian Algorithm result no longer holds in this more general setting. That is, if we apply the Hungarian algorithm to the case of multiple agents the dual prices will form a Walrasian equilibrium, but they are often not minimal among all Walrasian equilibria. In this paper we demonstrate how to compute the minimal Walrasian equilibria in the new more general setting, providing a non-trivial extension of the existing line of literature on the economics of matching or assignment games.

Further, our results can be applied to a more general combinatorial auction setting, based on the beneficial properties developed from our analysis of the network structure. In general, a combinatorial auction, in which bidders bid on combinations of items, promises market efficiency but may experience computational intractability and require an exponential amount of input. To help alleviate these problems, we show how to incorporate the easy-to-use assignment structure as "bid tables" in a larger combinatorial auction framework.

The remainder of the paper proceeds as follows. We begin with a discussion of the basic properties of preference expression using Figure 2's assignment network representation in §2 and §3. In §4 and §5 we discuss the VCG and Dynamic Bid Table Auction implementations, respectively, with conclusions in §6.

Figure 2: A Bipartite Assignment Game with Unit-Demand Agents. The bid values that a single agent assigns to each item on the left are shown. Arcs and bid values from each item to each of the other agents are omitted. In contrast to the game in Figure 1, each bidder controls a *collection* of unit-demand agents, represented by the grouping of agents with dotted lines.

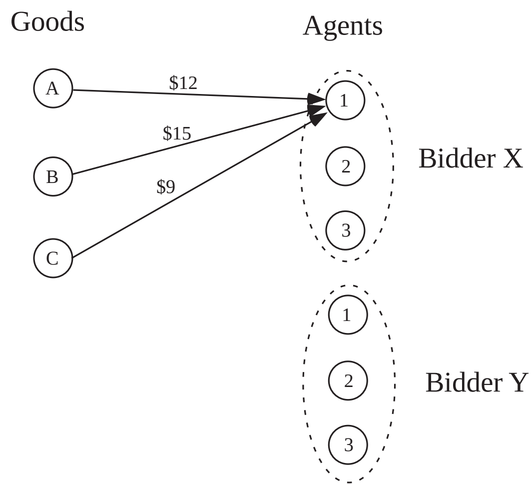

2 Linear Prices and Bid Tables

Demand revelation is desirable in many auction environments, and in multi-item auction environments bidders would often like to see *linear prices* for each item as information about the current state of the auction. Such prices allow a bidder to determine the price for a bundle by simply adding the values for the items in that bundle, making it quite simple for bidders to evaluate a number of potential packages. In order for these prices to be an accurate indicator of demand in the auction, we would like to have prices that are *accurate* and *separate* winning bundles from losing bundles. If the auction were to close immediately, we want the current prices to reflect the actual payments made by winners (accuracy) and for the current prices to be higher than the losers are willing to pay (separation). Unfortunately, linear prices which are accurate and separate are not always possible in a combinatorial auction when items are complements, as demonstrated by the following example.

Example: In a four-bidder, three-item auction let the bids on items A, B, and C be as follows:

$$b_1(\{A, B, C\}) = 6, \qquad b_2(\{A, B\}) = 5$$

$$b_3(\{A, C\}) = 5, \qquad b_4(\{B, C\}) = 5$$

Clearly, the efficient solution is to award all three items to bidder 1, but what prices can be assigned to the individual items that separate the winner from the losers? In order for the losing bidders to be satisfied, the sum of the prices of the items in a bundle should exceed any losing bid on the bundle:

$$p_A + p_B \geq 5$$

$$p_A + p_C \geq 5$$

$$p_B + p_C \geq 5$$

But this implies $p_A + p_B + p_C \geq 7.5$, a total payment that is too high for bidder 1, who will pay at most 6.

This example illustrates the well known failure of linear prices. At the conclusion of a combinatorial auction allowing for the most general expression of preferences, prices can only be expressed in terms of bundle-payments made by the winners and cannot be decomposed into meaningful individual item prices. Though this is true at the termination of a combinatorial auction, we argue that linear prices may still play a role in the early stages of a combinatorial auction, where they may be used as signals of the value of items if taken individually, before complementarity is taken into consideration. This idea is present, for example, in the clock-proxy auction proposed by Ausubel et al. (2006), who use linear prices to guide a clock auction, followed by a sealed-bid, or proxy stage.

We now begin to develop a new compact method for a bidder to write down bid information as an alternative to assigning a price to every bundle explicitly (since this expression grows exponentially in the number of items being auctioned). The approach for preference elicitation explored here make use of the concept of a *price-vector agent*. These agents represent the individual nodes on the right-hand-side of the bipartite network representation in Figure 2; referring to these nodes as agents is convenient in bridging the gap between the network representation and the underlying economics in which we are interested.

Each agent can be thought of as a fictional entity representing some portion of the preferences of a particular bidder. A price-vector agent is assigned a vector of prices (a monetary amount for each of the items in the auction) and "participates" in the auction based on these prices. An agent receiving a particular item pays at most the price vector component associated with that item. Throughout we assume that each price-vector agent is a unit-demand agent; each agent may receive at most one item. This clarifies the role of the price-vector agent in economic terms: Each agent treats all items as perfect substitutes. Expression of preferences using price-vector agents encourages each bidder to decompose his preferences into collections of substitutable items and to dedicate one or several agents to this collection depending on the level of substitutability (pure substitutability versus partial).

We call the collected set of (column) price-vectors a *bid table*. The result is any easy to read method of compactly annotating certain forms of substitutable preferences. Figure 3 shows an example of a bid table and demonstrates how a single agent can be used to name prices for pure substitutes, or how a collection of

Figure 3: A Bid Table

	agent 1	*agent* 2	*agent* 3	*agent* 4
pure substitute A_1	15	0	0	0
pure substitute A_2	16	0	0	0
pure substitute A_3	17	0	0	0
partial substitute B_1	0	23	18	12
partial substitute B_2	0	20	16	10
partial substitute B_3	0	19	15	9

agents can combine to yield a decreasing offer on partially substitutable items. To interpret a bid table, one need only keep in mind that at most one bid entry may be accepted from each row, and at most one from each column. We see therefore that in the example of Figure 3 that items A_1, A_2, and A_3 are indeed pure substitutes, because at most one can be purchased at a positive price. Similarly, if the bidder of this example receives item B_1 priced at 23, she cannot also receive (for example) item B_2 at a price of 20. If item B_2 is assigned to the bidder, the revenue maximizing auctioneer would be forced to accept a lower price from another agent of that bidder. In this case the auctioneer would charge 16 for B_2 rather than 20, verifying the partial substitutability; B_2 is worth less if taken with B_1.

In §3 we show that the set of preferences expressible in bid tables are properly contained in the set of preferences satisfying the *gross substitutes property*, connecting this theory to the economic literature on restricted preferences. We may then apply a theoretical result of Ausubel and Milgrom (2002) to elucidate a strength of the Bid Table Auction: a Vickrey-Clark-Groves (VCG) mechanism may be used with no possibility of disruption from false-name bidding or joint deviation by losing bidders, properties not satisfied for a VCG mechanism in general.

Additionally, the bid table format we present is based on the familiar assignment problem, and we show that this is a natural environment for the use of linear item prices. On one hand, if the bidders do not find the preference restrictions of the format to be restrictive, we find that this is a compact format which can be used to find either VCG outcomes (discussed in §4) or minimal Walrasian equilibrium individual item prices (discussed in §5). The latter can be implemented in either a one shot submission of bid tables or through iteratively updated bid tables in a Dynamic Bid Table Auction, guided at each stage by linear item prices.

If, on the other hand, we are in an environment requiring a more robust expression of preferences (e.g. complements are present), we argue that bid tables provide a restricted format for which linear prices make sense. In Day and Raghavan (2005), we introduce a multi-stage combinatorial auction using a Dynamic Bid Table Auction as an initial demand revelation phase. Bid tables allow us to utilize the attractive linear prices as long as it makes sense to do so, until any further bidding requires a departure from linear prices in order to treat complementary preferences. Relative to the simultaneous ascending auction (SAA), this can be accomplished without the exposure to receiving substitute goods at additive prices, and without the ability for competitors to signal among themselves.

3 Assignment Preferences and the Gross Substitutes Property

Our multi-item auction will have N items for sale, with the set of items denoted $I = \{1, 2, ...i, ...N\}$, and the set of M bidders will be referred to as $J = \{1, 2, ...j, ...M\}$. For the sake of generality we allow for several identical copies of each item to present in the auction, with the seller in the auction supplying the quantity sup_i of each item i. For any feasible bundle or package of items $x \in \mathcal{S} = \{x \in \mathbb{Z}_+^N | x_i \leq sup_i, \forall i \in I\}$, each bidder j perceives some value $v_j(x) \geq 0$ equal to her utility of receiving this bundle at zero cost. Further, in this treatment we assume non-negativity ($v_j(x \cup \{i\}) - v_j(x) \geq 0$, $\forall i, x$) and define a bidder's net utility as $u_j(x, p) = v_j(x) - \sum_{i \in S} p_i x_i$, where each component p_i of price vector p denotes a price for item i (we assume quasilinear net utility).

Next, we say a bundle of items x is *demanded by j at a price p* if x maximizes net utility for bidder j at p. By denoting the set of all bundles demanded by j at price p as $D_j(p)$, we may write this mathematically as:

$$D_j(p) = \arg\max_{x \in \mathcal{S}} u_j(x, p)$$

A bidder may demand several bundles at p if each maximizes net utility, and demands $\emptyset = [0, 0, ...0]$ when prices are too high. Additionally, we assume $v_j(\emptyset) = 0$, $\forall j$.

In a *Bid Table Auction*, each bidder j has a set of A_j agents he potentially wishes to satisfy, denoted by the set $K_j = \{1, 2, ...k, ...A_j\}$. We assume that each item in the auction may be consumed by only one agent and that each agent can consume at most one item. The interpretation of an agent varies from application to application, but in many scenarios agents may have a very natural interpretation. Motivated by the proposed use of a combinatorial auction for the allocation of airport landing slots (see Ball et al., 2006), it may be useful to think of each agent as representing a potential flight, and each item as a slot made available to an airline at an airport for landing a single plane.

Let v_{ijk} denote the utility perceived by bidder j when agent k receives item i. The preference assumption

Figure 4: Assignment Preferences in a Bid Table

	agent 1	*agent* 2	*agent* 3
Item a	0	0	2
Item b	2	5	3
Item c	4	6	4
Item d	0	3	0

of the Bid Table Auction can be specified as follows:

$$v_j(x) = \max \sum_{i \in S} \sum_{k \in K_j} v_{ijk} \cdot y_{ijk} \tag{AP}$$

$$\text{subject to} \quad \sum_{i : x_i \geq 1} y_{ijk} \leq 1, \; \forall j, k \in K_j$$

$$\sum_{j \in J} \sum_{k \in K_j} y_{ijk} \leq sup_i, \; \forall i : x_i \geq 1$$

$$\text{where } y_{ijk} = \begin{cases} 1 & \text{if item } i \text{ is assigned to bidder } j\text{'s } k\text{th agent} \\ 0 & \text{otherwise} \end{cases}$$

In short, a bidder's value for a bundle of items is a maximal value assignment of those items to her agents. Bidders modeled in this fashion may be said to have *assignment preferences,* as their preferences are governed by an assignment network. As is well known (see Nemhauser and Wolsey, 1999), the constraint matrix of this integer program is total unimodular, indicating that the problem is solved to integral optimality by its LP relaxation. This implies that a bidder could rapidly determine her value for any set of items with an LP solver or faster combinatorial algorithm designed specifically to solve assignment problems, even for a large value of N.

A bidder with assignment preferences can express her value for every bundle of items in a *bid table* containing the values of v_{ijk} with a row for each item i and column for each agent k. For example, with the bid table of Figure 4 we may determine a bidder's value for any bundle of items a, b, c, or d if assigned to any of three possible agents. Let $sup_i = 1$ for all i in this example. Considering the values in this bid table, we notice that a bidder's value for any single item is simply the maximum value for that row in the bid table; if a bidder is awarded only one item, the agent that experiences the most utility from this item will be accommodated. Thus item a by itself is worth 2 units to the bidder, item b by itself is worth 5 units, item c by itself is worth 6 units, and item d by itself is worth 3 units. The value for some collection of items is not, however, necessarily equal to the sum of his values for individual items, as there may be conflicts when the row maximums occur in the same column. For example, for bidder j with this bid table $v_j(\{b\}) = 5$,

$v_j(\{c\}) = 6$ but $v_j(\{b,c\}) = 9 \neq 5 + 6$. We find $v_j(\{a,b,c,d\}) = 11$ with an optimal assignment of a to the third agent, b to the second agent and c to the first agent. We notice that to achieve this amount we do not assign item c to the second agent, despite this being the overall highest value in the bid table, and also that we choose to assign item a and not item d, despite the fact that $v_j(\{a\}) < v_j(\{d\})$.

These observations show that assignment preference valuation functions can have some non-intuitive properties and are not contained in the class of additive valuation functions (in which the value of some set of items always equals the sum of the values for the individual items). We notice that any additive valuation function can be modelled by a bid table by taking constant rows and at least as many agents as there are items. Thus additive valuation functions are properly contained in the class of assignment preference valuation functions. One may next ask how assignment preferences relate to the larger class of valuation functions which satisfy the gross substitutes property.

Definition: The *gross substitutes property* holds if and only if the following condition holds for every bidder j: For any price vectors $p' \geq p$ with $p' \neq p$, and any $x \in D_j(p)$, there exists a set $x' \in D_j(p')$ with $x'_i \geq x_i$ for all i with $p'_i = p_i$.

Or roughly, if the prices rise on some of the items in a demanded set, then there is at least one demanded bundle at the new prices still containing all previously demanded items for which the price did not increase.

The gross substitutes property is sufficient to guarantee the convergence of several ascending price multi-item auction formats to a Walrasian equilibrium[1] (e.g. those of Ausubel and Milgrom, 2002; Gul and Stachetti, 1999a)) and has other beneficial properties which will be discussed below. Though it is a common assumption due to its attractive theoretical properties, it is uncommon for theorists to describe applications which give rise to this property, to explain under what conditions it is a safe assumption to make, or to suggest how to enforce this restriction among bidders if it is necessary for the convergence of an auction to a desirable outcome. In this paper, we start instead from the assumption of assignment preferences and show that the gross substitutes property follows as a result.

Though this fact was proven independently by Lehmann et al. (2001) for the equivalent OR-of-XOR of singletons language, the proof presented here is constructive rather than an indirect. Experience shows that us that using an LP solver to determine demanded bundles under assignment preferences may give any demanded bundle after a price increase and will often not provide a bundle that verifies the validity of the theorem. For any demanded bundle x^* before a price increase and a given demanded bundle x' after the price increase, the algorithm described in the proof constructs a new demanded bundle x^{**} after the increase, demanding (at least the same level) all previously demanded items in x^* for which the price has not risen. This is useful, for example, if the auction is demanding a price path that upholds the gross substitutes

[1] Defined formally in §4.

property throughout.

Theorem 3.1. *A valuation function determined by assignment preferences displays the gross substitutes property.*

Proof. Suppose we begin at price vector p^1 with a particular demanded bundle $x^* \in D_j(p^1)$ and consider price vector $p^2 \geq p^1$, with $p^2 \neq p^1$ and a particular demanded bundle $x' \in D_j(p^2)$. Under assignment preferences a demanded set is always supported by an assignment of items to agents which maximizes total value net of prices (sometimes more than one supporting assignment may exist). Begin by selecting select particular assignments (sets of arcs ot item agent pairs) A^* and A' supporting sets x^* and x', respectively. We will say that arc $i \to k$ is in A if assignment A has item i assigned to agent k.

Now, given an element $\bar{\imath}$ with $x^*_{\bar{\imath}} \geq 1$ with $p^1_{\bar{\imath}} = p^2_{\bar{\imath}}$ and $x'_{\bar{\imath}} < x^*_{\bar{\imath}}$, we show how to find a demanded bundle $x^{**} \in D_j(p^2)$ with $x^{**}_{\bar{\imath}} \geq x^*_{\bar{\imath}}$ and corresponding assignment A^{**}. We do this by showing how to construct two sets IN and OUT such that $A^{**} = (A' \backslash OUT) \cup IN$.

Algorithm: Constructing a set $x^{**} \in D_j(p^2)$ containing at least as much of $\bar{\imath}$
Step 0: Set $InItem = \bar{\imath}$, $OutItem = IN = OUT = \emptyset$.
Step 1: Let k be an agent assigned $InItem$ under A^*.
Step 2: Find an item (or $\emptyset$) $OutItem$ with $OutItem \to k$ in A'.
Step 3: Add arc $OutItem \to k$ (or $\emptyset$) to OUT. Add arc $InItem \to k$ to IN.
Step 4: If $OutItem \notin x^*$ (or $OutItem = \emptyset$), terminate, else set $InItem = OutItem$ and goto Step 1.
Terminate: Set $A^{**} = (A' \backslash OUT) \cup IN$.

Claim: The bundle x^{**} implied by arc set A^{**} is in $D_j(p^2)$. Suppose not. Then $x' \in D_j(p^2)$ provides $u_j(x^{**}, p^2) < u_j(x', p^2)$. Since by construction each member of OUT was in A', this yields $u_j(IN, p^2) < u_j(OUT, p^2)$, with the obvious interpretation of the utility for an arc set. That is

$$u_j(IN, p^2) = \sum_{i \to k \in IN} v_{ijk} - \sum_{i | i \to k \in IN} p^2_i < u_j(OUT, p^2) = \sum_{i \to k \in OUT} v_{ijk} - \sum_{i | i \to k \in OUT} p^2_i$$

But because every item assigned under IN is assigned under OUT, except for $\bar{\imath}$ which does not experience a price change from p^1 to p^2, we have $u_j(IN, p^1) < u_j(OUT, p^1)$. But by construction of these arc sets $(A^* \backslash IN) \cup OUT$ is a feasible assignment, and then we must have $u_j(A^*, p^1) < u_j((A^* \backslash IN) \cup OUT, p^1)$, contradicting our choice of $x^* \in D_j(p^1)$.

The validity of the claim provides a demanded set containing $\bar{\imath}$. Each execution of the algorithm finds a set in $x^{**} \in D_j(p^2)$ containing an element from x^* that was missing in x' and does not remove any elements from $x^* \cap x'$. By repeating this procedure (setting $x' = x^{**}$ each time) we arrive at a set in $D_j(p^2)$ containing all the desired elements of x^*. $\square$

To illustrate the idea of the proof consider the following example in which the items in the auction are denoted by lowercase letters. Suppose at prices p^1 the bidder in question demands the bundle $\{a, b \times 2, c, d, e, f\}$, which is found to maximize utility with the assignment of items to agents $A^* = \{a \rightarrow 1, b \rightarrow 2, c \rightarrow 3, d \rightarrow 4, e \rightarrow 5, f \rightarrow 6, b \rightarrow 7\}$. Suppose next that prices rise on items a and d. The bidder then recomputes for an optimal bundle at these prices and finds the optimal assignment $A' = \{h \rightarrow 1, r \rightarrow 2, p \rightarrow 3, q \rightarrow 4, b \rightarrow 5, d \rightarrow 6, c \rightarrow 7\}$, but this assignment clearly does not validate the gross substitutes property; items e and f did not experience a price increase but they do not appear in the newly demanded bundle, while item b did not experience a price increase but had a reduction in demand from 2 to 1.

To apply the algorithm described in the proof, first consider finding a demanded set including two copies of item b. Putting A^* above A' and designating destination agents by column, we see below that agent 2 who was assigned b in A^* is assigned r in A':

$$
\begin{array}{ccccccccc}
A^* & : & a & b & c & d & e & f & b \\
 & & & \downarrow & & & & & \\
A' & : & h & r & p & q & b & d & c \\
- & - & - & - & - & - & - & - & - \\
A^{**} & : & h & b & p & q & b & d & c
\end{array}
$$

The manipulation performed by the algorithm is in its simplest form here. The value of A^{**} at p^2 must be at least as much as the value of A' at p^2, or else it would be possible to switch r in for b in A^* and receive a higher value at p^1, contradicting the optimality of A^* as a demanded bundle. This reasoning only holds because the price of b does not change in the movement from p^1 to p^2, and because r can be switched in freely as it is not in A^*.

When we try to find a demanded bundle at p^2 containing e, it is not quite so easy; the agent assigned e under A^* is assigned b under A', and since b is already assigned under A^* a one-for-one switching argument fails. The algorithm rectifies this by tracing back a path until an item is found that was not allocated under A^*. For example, the following diagram helps us see that in moving from A^* to A' (which has now been replaced by A^{**} from the previous step) item e's agent is reassigned item b, whose agent in A^* is reassigned item c, whose agent in A^* is reassigned by item p which was unallocated in A^*. (Note that if we picked agent 2 for item b rather than agent 7, we simply add one more trivial step before finding the item p unallocated

Figure 5: Preferences that do not fit in a Bid Table

$$
\begin{array}{ll}
v(\{a\}) = 10, & v(\{a,b\}) = 20 \\
v(\{b\}) = 12, & v(\{a,c\}) = 19 \\
v(\{c\}) = 13, & v(\{b,c\}) = 17
\end{array}
$$

$$
\begin{array}{c|cc|}
a & 10 & 8? \\
b & 12 & 10? \\
c & 13 &
\end{array}
$$

in A^*.) The same optimality argument can be used to show a demanded bundle at p^2 including item e; $\{p \to 3, b \to 5, c \to 7\}$ can be replaced by $\{c \to 3, e \to 5, b \to 7\}$ or else the optimality of A^* is contradicted.

$$
\begin{array}{ccccccccc}
A^* & : & a & b & c & d & e & f & b \\
 & & & & \downarrow & & \downarrow & & \downarrow \\
A' & : & h & b & p & q & b & d & c \\
- & - & - & - & - & - & - & - & - \\
A^{**} & : & h & b & c & q & e & d & b
\end{array}
$$

Similar arguments allow us to settle on a final set A^{**}, demanding all items for which prices did not increase with at least the same levels as at price p^1. The final step is displayed in the following diagram.

$$
\begin{array}{ccccccccc}
A^* & : & a & b & c & d & e & f & b \\
 & & & & & \downarrow & & \downarrow & \\
A' & : & h & b & c & q & e & d & b \\
- & - & - & - & - & - & - & - & - \\
A^{**} & : & h & b & c & d & e & f & b
\end{array}
$$

Theorem 3.1 assures us that the gross substitutes property holds when bidders are restricted to the use of bid tables only, as in Stage I of the auction proposed in Day and Raghavan (2005), or any isolated Bid Table Auction implementation. To complete the characterization of assignment preferences with respect to the gross substitutes property, we note that the valuation function of Figure 5 maintains the gross substitutes property but cannot be expressed as an assignment preference valuation function. The gross substitutes property follows from the submodularity and positivity of the valuation function. To see that these preferences cannot be expressed as an assignment preference valuation function (or equivalently, as a bid table) notice that the value of any two items is less than the sum of the values for the individual

Figure 6: A Bid Table Auction for which VCG payments are lower than in any Walrasian equilibrium

Bidder X

A	8	0	0	0
B	8	0	0	0
C	0	8	0	0
D	0	8	0	0

Bidder Y

A	6	0	0	0
B	0	6	0	0
C	2	0	0	0
D	0	0	0	0

Bidder Z

A	0	0	0	0
B	2	0	0	0
C	0	6	0	0
D	6	0	0	0

items, thus the values for the individual items must all occur in the same column of a bid table. But now if we attempt to put in a value of 8 into column 2, row a, or a value of 10 into column 2 row b to express $v(\{a,b\}) = 20$, we either overvalue the bundle $\{a,c\}$ at 21 or the bundle $\{b,c\}$ at 23. There is therefore no way to express these gross substitute preferences as assignment preferences.

With Theorem 3.1 and earlier observations we have the following corollary, where V_{add}, V_{AP}, and V_{sub} denote the classes of valuation functions that have additive preferences, assignment preferences, and gross substitute preferences, respectively.

Corollary 3.2. $V_{add} \subset V_{AP} \subset V_{sub}$

Though all containments in Corollary 3.2 are proper, we note that assignment preferences do retain some of the interesting properties lost between V_{sub} and V_{add}. Notably, we can embed in a Bid Table Auction an example from Gul and Stachetti (1999a) that demonstrates the VCG outcome may have lower payments than the lowest Walrasian Equilibrium. This embedding is shown in Figure 6, providing an example of a Bid Table Auction for which the VCG payments are strictly less than any Walrasian Equilibrium. In one efficient allocation, bidder X gets $\{A,D\}$, bidder Y gets $\{B\}$ and bidder Z gets $\{C\}$, with VCG payments of 12, 2, and 2, respectively. The lowest Walrasian equilibrium price vector is however, $p_A = p_B = p_C = p_D = 6$, charging more to bidders Y and Z than in any VCG outcome. Having established that the VCG outcomes and the lowest Walrasian Equilibrium may differ in a Bid Table Auction, we now discuss Bid Table Auctions which achieve each of these outcomes, in §4 and §5 respectively.

4 VCG Bid Table Auction Implementations

Corollary 3.2 allows us to apply the following theorem:

Theorem 4.1. *(Ausubel and Milgrom, 2002)) For an auction with a single unit of each item and bidders with valuation functions drawn from the set V such that $V_{add} \subseteq V$, the following conditions are equivalent:*

1. $V \subseteq V_{sub}$

2. *For every profile of bidder valuations drawn from V, adding bidders can never reduce the seller's total revenues in the VCG auction.*

3. *For every profile of bidder valuations drawn from V, any shill bidding is unprofitable in the VCG auction.*

4. *For every profile of bidder valuations drawn from V, any joint deviation by losing bidders is unprofitable in the VCG auction.*

In the sealed-bid Vickrey-Clarke-Groves (VCG) auction (Clarke, 1971; Groves, 1973; Vickrey, 1961) each bidder submits her true valuation for every possible bundle of items. A winning allocation is then determined which distributes bundles to bidders so as to maximize total value. To assure that each bidder has the incentive to reveal her total value honestly, she pays not her reported value for the bundle she receives, but this value less an appropriate discount. This VCG discount assures that a bidder does not pay any more than would be necessary to receive this bundle given her opponents honest reports, and is equal to the value of the final allocation minus the maximum value attainable without this bidder. A bidder only decreases her chances of receiving her efficient allocation of items by misreporting, with no possible gain.

The assurance of honest valuation reporting is well known to be the strength of the VCG auction, but as Ausubel and Milgrom (2002) point out, there are several drawbacks. Shill bidding occurs when a bidder enters a false identity into the auction in a way that results in an inefficient allocation or alternative set of payments which is preferred by the deceitful bidder. Similarly, Ausubel and Milgrom demonstrate how two or more losing bidders in a VCG auction may increase their bids to become winning without having to pay for this increase. These problems with the VCG auction are explored in further depth elsewhere Rothkopf et al. (1990), but we note that these difficulties rely on the existence of complementary items, an impossibility when the gross substitutes condition holds. Theorem 4.1 ensures that these problems of collusive behavior and revenue reduction sometimes recognized as drawbacks to the VCG auction have no effect when the gross substitutes property holds.

Since the hypotheses for Theorem 4.1 and condition (1) of Theorem 4.1 are satisfied according to Corollary 3.2, we may conclude that the Bid Table Auction scenario is one for which the VCG auction attains its full strength. The advantages of the Bid Table Auction scenario over the general VCG auction context are that each bidder may simply submit a bid table to express her preferences (a method that is far more compact than issuing a price for every one of the $2^N - 1$ possible nonempty bundles) and the winner determination problem can be solved polynomially using LP techniques. Both of these features are important for an auction of many items. Again, we are motivated by the proposed auctions for airport landing slot rights, and notice that at LaGuardia airport, for example, over 800 slots may be available in a single day; an auction for these slots which enumerates all possible bundles would be quite impossible.

In order to run a *VCG Bid Table Auction*, one must compute both an efficient allocation and the set of VCG payments for all bidders. To determine an efficient allocation, the auctioneer, having received bid tables containing the submitted bids of b_{ijk}, needs to solve the LP:

$$z_J = \max \sum_{(i,j,k)\in I\times J\times K_j} b_{ijk}\cdot y_{ijk} \tag{P}$$

$$\text{subject to } \sum_{i\in I} y_{ijk} \le 1,\ \forall (j,k) \text{ with } j\in J \text{ and } k\in K_j$$

$$\sum_{j\in J}\sum_{k\in K_j} y_{ijk} \le sup_i,\ \forall i\in I$$

$$y_{ijk}\ge 0,\ \forall i,j,k$$

To determine VCG payments, the auctioneer can solve this problem again without bidder j to find the appropriate discount for bidder j, $z_J - z_{J\setminus j}$, where $z_{J\setminus j}$ denotes the value of the objective value of P with bidder j removed. The assignment corresponding to z_J determines a bundle S_j with bid value $b_j(S_j)$ for each bidder. Each bidder in the VCG auction receives S_j and pays $b_j(S_j) - (z_J - z_{J\setminus j})$. For the entire auction at most $M+1$ LPs must be solved. Since there are $O(MN^2)$ variables and $O(MN)$ constraints in each LP, and because it is well known that LPs can be solved in polynomial time (as a function of the number of constraints and variables), we conclude that the VCG Bid Table Auction can be solved in polynomial time.

5 Dynamic Bid Table Auctions

Though Theorem 3.1 allows the application of Theorem 4.1, implying that because of the gross substitutes property the Bid Table Auction is a good preference format for a VCG implementation, it also limits the types of possible applications of the Bid Table Auction to those in which there is no complementarity among items. It is possible however to use bid tables in the preliminary (revelation) stages of a hybrid auction concluding with a sealed-bid auction. Theorem 4.1 and the accompanying discussion by Ausubel and Milgrom (2002) make a clear case against the use of the VCG mechanism when complementarities are possible

Ignoring for now the issue of complements (i.e. maintaining the assumption that bidders have assignment preferences), it seems that the primary drawbacks to the VCG Bid Table Auction are the lack of price discovery and privacy preservation. Though privacy preservation may in principle be legally enforced, lack of price discovery may be an inherent concern and discourages the use of a sealed-bid VCG Bid Table

Auction. A bidder may not know how to fill out a bid table with honest valuations for agent/item pairs without knowing her opponents' values for various items, and would prefer a dynamic auction to learn about her competition.

Towards the goal of implementing a dynamic auction for the landing slot application, we examine the LP dual to P:

$$\min \sum_{i\in I} sup_i \cdot p_i + \sum_{j\in J}\sum_{k\in K_j} s_{jk} \tag{D}$$

$$\text{subject to} \qquad p_i + s_{jk} \geq b_{ijk},\ \forall i,j,k \tag{1}$$

$$p_i \geq 0,\ \forall i \in I$$

$$s_{jk} \geq 0,\ \forall j \in J,\ k \in K_j \tag{2}$$

Problem P has integer optimal solutions, and an optimal solution to problem D will have the same objective value. This dual formulation suggests a set of "supporting prices" p_i for each item. If we stipulate that a bidder receiving item i in an optimal solution of P pays p_i, the value of s_{jk} becomes the surplus perceived by bidder j's agent k (abbreviated (j,k)).

The complementary slackness conditions for the primal dual pair P-D are:

$$\begin{aligned} \forall i,j,k \quad & y_{ijk} > 0 \Rightarrow p_i + s_{jk} = b_{ijk} \\ & p_i + s_{jk} > b_{ijk} \Rightarrow y_{ijk} = 0 \end{aligned} \tag{3}$$

$$\begin{aligned} \forall j,k \qquad & \textstyle\sum_{i\in I} y_{ijk} < 1 \Rightarrow s_{jk} = 0 \\ & s_{jk} > 0 \Rightarrow \exists i \text{ such that } y_{ijk} = 1 \end{aligned} \tag{4}$$

$$\begin{aligned} \forall i \in I \quad & \textstyle\sum_{j\in J}\sum_{k\in K_j} y_{ijk} < 1 \Rightarrow p_i = 0 \\ & p_i > 0 \Rightarrow \textstyle\sum_{j\in J}\sum_{k\in K_j} y_{ijk} = 1 \end{aligned} \tag{5}$$

Each of these conditions (presented in equivalent pairs) carries an economic interpretation reinforcing the validity of the model. If agent (j,k) is awarded item i, condition (3) implies $p_i + s_{jk} = b_{ijk}$, validating our reference to s_{jk} as surplus. If an agent is not awarded an item in an optimal solution, by condition (4) we have $s_{jk} = 0$, and then the constraints (1) become $p_i \geq b_{ijk}$, $\forall i \in I$; if an agent is empty-handed at optimality, the optimal dual prices make any item too expensive for this agent to buy. Similarly, the second statement of (4) says that if agent (j,k) perceives any surplus, then it must be the case that (j,k) received an item. Further, if we evaluate the potential surplus that item i would bring to agent (j,k) who receives a different item at optimality, we find $s_{jk} \geq b_{ijk} - p_i$; the price of i is great enough that a change from the item awarded at optimality to i for agent (j,k) does not increase surplus. The last pair of conditions, (5), state simply that an item will have a nonzero price only if it is received by some agent.

Definition: An allocation of bundles to bidders $(x^1, x^2, ...x^j...x^M)$ and price vector p constitute a *Walrasian equilibrium* if and only if for every bidder j, $x^j \in D_j(p)$.

Theorem 5.1. *Assuming truthful demand reporting and assignment preferences, the allocation* $(x^1, x^2, ...x^j...x^M)$ *given by an optimal solution to P together with prices for items* $p = (p_1, p_2, ...p_i...p_N)$ *given by the corresponding dual solution to D constitutes a Walrasian equilibrium.*

Proof. Suppose not: there is some bidder j and some bundle $\overline{x^j}$ with j strictly preferring $\overline{x^j}$ to x^j. Equivalently,

$$b_j(\overline{x^j}) - \sum_{i \in I} \overline{x_i^j} \cdot p_i > b_j(x^j) - \sum_{i \in I} x^j \cdot p_i$$

where $b_j(x^j)$ and $b_j(\overline{x^j})$ are supported by agent sets K and $\overline{K}$, respectively. Complementary slackness conditions (3) provide $b_{ijk} - p_i = s_{jk}$ for each item i assigned to j in x^j, yielding $b_j(S_j) - \sum_{i \in I} x^j \cdot p_i = \sum_{k \in K} s_{jk}$. Similarly, constraints (1) yield $\sum_{k \in \overline{K}} s_{jk} \geq b_j(\overline{S_j}) - \sum_{i \in I} \overline{x_i^j} \cdot p_i$. Together this implies

$$\sum_{k \in \overline{K}} s_{jk} > \sum_{k \in K} s_{jk}$$

but with $s_{jk} = 0$ in the optimal solution of D for any agent $k \notin K$, and $s_{jk} \geq 0$ for all tasks, we have

$$\sum_{k \in \overline{K} \cap K} s_{jk} > \sum_{k \in K} s_{jk}$$

a contradiction. □

The existence of a Walrasian equilibrium under the gross substitutes property is established by Kelso and Crawford (1982). Among all Walrasian equilibria there exists one that is bidder-optimal, the lowest Walrasian equilibrium. The uniqueness of this bidder-optimal Walrasian equilibrium when the gross substitutes property holds follows naturally from the work of Gul and Stachetti (1999b), who demonstrate the lattice structure of Walrasian Equilibria under gross substitutes. In our case, this existence and uniqueness is guaranteed because the gross substitutes property holds by Theorem 3.1. This vector of lowest Walrasian equilibrium prices constitute what one might refer to as "good" linear prices, and their existence in the case of bid tables verifies our claim that this is a case for which linear prices make sense. It is known from the work of Demange et al. (1986) that versions of the "Hungarian algorithm" (a primal/dual method for solving assignment problems) yield this lowest Walrasian equilibrium in the special case that $A_j = 1$, $\forall j$, and we now discuss the generalization to the case of arbitrary integer values for A_j.

As in Demange et al. (1986), the Hungarian algorithm finds an efficient allocation in a Bid Table Auction and the prices used in its solution provide an optimal solution to the dual problem D, together forming a Walrasian equilibrium, from Theorem 5.1. (To apply the Hungarian Algorithm directly, it is necessary to give unique names to each identical copy of an item). In general, however, this method produces Walrasian prices that are greater than or equal to the actual lowest Walrasian equilibrium prices for a Bid Table Auction. This is because according to formulation D alone, the price a bidder pays may be determined by one of her own agents, as if her own agents must price-compete among themselves.

To avoid this self-competition problem when finding the lowest Walrasian equilibrium, we introduce the following Dual Pricing Problem DPP. This formulation specifies an LP-characterization of the minimal Walrasian equilibrium prices for a Bid Table Auction. Specifically, given a solution to the primal problem P above with objective value z, we may fix the dual objective from D at its optimal value in constraints (6) and maximize total surplus over all bidders.

$$\max \sum_{j \in J} \sum_{k \in K_j} s_{jk} \tag{DPP}$$

$$\text{subject to} \quad z = \sum_{i \in I} sup_i \cdot p_i + \sum_{j \in J} \sum_{k \in K_j} s_{jk} \tag{6}$$

$$p_i + s_{jk} = b_{ijk}, \ \forall i, j, k \text{ with } i \rightarrow j, k \tag{7}$$

$$s_{jk} = 0, \ \forall j, k \text{ with } \emptyset \rightarrow j, k \tag{8}$$

$$p_i + s_{jk} \geq b_{ijk}, \ \forall i, j, k \text{ with } i \nrightarrow j \tag{9}$$

where we expand our use of the $\rightarrow$-notation to express assignment under a specifically chosen efficient solution to problem P; for example, $i \rightarrow j, k$ expresses that item i is awarded to bidder j's kth agent in the selected efficient allocation. Similarly, $i \nrightarrow j$ signifies that item i is not awarded to bidder j.

We note that the most important distinction between DPP and D comes in constraint set (9) in which only constraints that do not involve self-competition are enforced. That is, form the set of constraints (9) from constraint set (1) of formulation D by removing all constraints $p_i + s_{jk} \geq b_{ijk}$ involving both an item i and its winner, bidder j. This approach is equivalent to re-solving the primal allocation problem P and the dual problem D with a Hungarian primal/dual method after lowering all non-winning entries in a winning row of a bid table to zero. We use the formulation DPP in the following proof as it provides an interesting interpretation as a price adjustment procedure; starting with an efficient allocation and the optimal solution of D, lower the prices on winning bidders as long as no one complains, any violated constraint from (9) interpreted as a possible complaint from another bidder. Knowing that prices may be easily and transparently adjusted to a unique minimal Walrasian equilibrium is a primary benefit of Theorem 3.1 (gross substitutes). We will use this unique price vector in a multi-round setting, using Bid Table Auctions as a mechanism of demand revelation with distinct meaningful price signals at each round of submission. We now prove that we achieve these desirable price signals formally.

Theorem 5.2. *Given an optimal solution to the primal problem P with objective z, solving DPP to maximize bidder surplus yields the lowest Walrasian equilibrium price vector p^*.*

Proof. Begin with price vector p^1 which is an optimal solution to D and therefore a Walrasian equilibrium price by Theorem 5.1. Note that this optimal solution to D gives a feasible solution to DPP with constraints (7) and (8) holding by the strong duality of P and D. Lowering any component p_i^1 and raising each corresponding component s_{jk}^1 by the same amount (where i is assigned to j, k in the solution to P) will have no effect on constraints (6), (7) or (8).

Claim: If such a shift from price to surplus does not violate constraints (9) the resulting price vector continues to support a Walrasian equilibrium. If not, some other bidder $\bar{j}$'s agent $\bar{j}, k$ not assigned item i would prefer item i to whatever item that has been assigned to $\bar{j}, k$ (if any). This implies that $b_{i\bar{j}k} - p_i > s_{jk}$ which would violate the corresponding constraint from (9).

We therefore proceed to shift price to surplus until any possible shift causes a violation of some constraint from (9), achieving a price vector p^2.

Claim: p^2 is the lowest Walrasian equilibrium price vector. Suppose not: let $p^3 \neq p^2$ be the lowest Walrasian price vector. From the lattice theory of Walrasian equilibria when the gross substitutes condition holds (as established by Gul and Stachetti, 1999b) this Walrasian Equilibrium price is unique, and for every component $p_i^3 \leq p_i^2$. Since $p^3 \neq p^2$ there must be some i for which $p_i^3 < p_i^2$. Since our price-to-surplus shifting procedure has terminated, a shift from p_i^2 to p_i^3 must violate some constraint from (9), thus we have the following inequality holding for some j, k with i not assigned to bidder j:

$$p_i^3 + s_{jk} < b_{ijk}$$

But now at price vector p^3, bidder j who is allocated bundle x^j prefers the alternative bundle with x_i^j greater by 1 over x^j, implying that p^3 does not support a Walrasian equilibrium, a contradiction. □

Now that we have shown how the lowest Walrasian price vector will be computed at each round, we propose the *Dynamic Bid Table Auction* proceeding in multiple rounds as follows. Accept bid tables from all bidders and determine a winning allocation (solution to P) using any technique (e.g. the Hungarian algorithm). Adjust the prices as suggested in Theorem 5.2 to find an optimal solution to DPP, and therefore the set of unique lowest Walrasian equilibrium prices. Announce winning prices and allow bidders to adjust their bid tables subject to the following rules:

- Any bid in a non-winning row of a bid table must be raised at least to the current price for that item (row) plus one price increment, or else it may not be altered for the remainder of the auction.
- A bidder who does not wish to increase an entry to the required amount may increase it to a price below the current price plus one increment. This is the bidder's "last-and-best" offer.

The auction then continues by computing a new set of winning bids and prices, and the process repeats until no one wishes to raise any bid table entries any further.

To show that this procedure progresses to a desirable equilibrium, assume that each bidder perceives a set of maximum bid table values, v_{ijk}. We would expect to see these entries from honest bidders in the direct revelation VCG as in §3, and now show that the dynamic game converges to the same outcome as the direct revelation game, given an assumption of straightforward bidding. Though this behavioral assumption is strong, there is evidence in practice and from the relevant literature (especially from Ausubel et al., 2006) that an appropriate bidding *activity rule* will encourage straightforward bidding behavior.

For the Dynamic Bid Table Auction scenario, we will say that a bidder bids straightforwardly if she increases a bid table b_{ijk} in a non-winning row by the minimal increment whenever the potential surplus $v_{ijk} - p_i$ for agent (j, k) is greater than the actual current surplus for agent (j, k), $\bar{s}_{jk} = v_{\bar{\imath}jk} - p_{\bar{\imath}}$, where $\bar{\imath}$ is the item currently awarded to agent (j, k). Here the modifier *actual* and 'bar' over the s emphasize that this is the actual surplus as perceived by the bidder, not the apparent surplus which may be computed using her revealed information: $s_{jk} = b_{\bar{\imath}jk} - p_{\bar{\imath}}$. In either case, this surplus will be zero if no item is awarded to agent k. We also assume that if a straightforward bidder is asked to raise a bid table value above v_{ijk}, then she will raise it to v_{ijk}, reporting her true valuation as she loses eligibility to alter the entry further.

Theorem 5.3. *A Dynamic Bid Table Auction with straightforward bidders (A) terminates at an efficient equilibrium of the direct revelation Bid Table Auction, and (B) achieves the unique lowest Walrasian equilibrium prices.*

Part A: First, we show that the allocation at termination is an efficient outcome of the direct revelation game. Suppose not. Let A be the allocation at termination of the dynamic auction with current bid table values b_{ijk}. By supposition there is an allocation $\bar{A}$ such that

$$\sum_{\bar{A}} y_{ijk} \cdot v_{ijk} > \sum_{A} y_{ijk} \cdot v_{ijk} \tag{10}$$

where the summation over an allocation signifies summation over all i, j, k with values of y_{ijk} given by that allocation. Because we have assumed that bidders bid straightforwardly and that the dynamic auction has terminated, for any bid table column j, k which is allocated item i under allocation A and $\bar{\imath}$ under $\bar{A}$, it must be the case that $v_{\bar{\imath}jk} - p_{\bar{\imath}} \leq v_{ijk} - p_i$ (if not the straightforward bidder would want to continue bidding on $\bar{\imath}$). This inequality also holds (reflexively) for any agent j, k that is awarded the same item under both allocation A and $\bar{A}$. For any j, k allocated item $\bar{\imath}$ under $\bar{A}$ which is not allocated an item under A, this condition becomes $v_{\bar{\imath}jk} - p_{\bar{\imath}} \leq 0$. Finally, by individual rationality we also have that $0 \leq v_{ijk} - p_i$ for any i allocated to j, k under A (particularly we take this inequality for any columns j, k which are for agents receiving items in both allocations, allocated items under A but not under $\bar{A}$). We then sum these three sets of inequalities, selecting (and multiplying by y_{ijk}) the appropriate one for each agent j, k who receives items in either allocation A or $\bar{A}$, or both, yielding:

$$\sum_{\bar{A}} y_{ijk} \cdot (v_{ijk} - p_i) \leq \sum_{A} y_{ijk} \cdot (v_{ijk} - p_i)$$

$$\sum_{\bar{A}} y_{ijk} \cdot v_{ijk} \leq \sum_{A} y_{ijk} \cdot v_{ijk} \tag{11}$$

with the second inequality following since each item is allocated exactly once in each allocation, allowing us to cancel the sum of all p_is from each side. But (11) contradicts (10), with the desired result following.

Part B: Next we show that the prices at termination of the dynamic schedule auction are the lowest Walrasian equilibrium prices for the direct revelation schedule auction. Using Theorem 5.2, this is equivalent to showing that the solution to the pricing problem DPP using the value $\bar{z}$ (computed using values from v_{ijk}) has equal values for all p_i to the solution of DPP using value z (computed using b_{ijk}).

Given that the optimal allocation is unchanged by increasing all b_{ijk} values to their reservation point v_{ijk} from Part A, increase all b_{ijk} values to v_{ijk} and simultaneously increase every surplus s_{jk} by the same amount wherever i is allocated to j, k. We claim that this provides a solution to DPP using value $\bar{z}$ with identical values of all p_i, as desired. Clearly, the simultaneous shift of s_{jk} values with b_{ijk} values upholds DPP constraints (6) and (7). Constraints (8) are upheld since s_{jk} only changes for columns which win items, while a violation of a constraint from (9) would imply a violation of a termination condition. Since none of the constraints of DPP are violated, our new solution must be feasible to DPP. Since any increase of the DPP objective function must be accompanied by an equivalent increase to the primal objective z, and because we have achieved all such increase as surplus, we may be assured that our new objective function is optimal. Since no p_i values have changed we have the same lowest Walrasian equilibrium price vector at the termination of the Dynamic Bid Table Auction and in the direct revelation Bid Table Auction, as desired.

Part B of this proof demonstrates the desirable privacy-preservation property of the Dynamic Bid Table Auction relative to the static revelation version: the actual surplus each bidder perceives need not be revealed in the dynamic case, only that surplus is at least one increment. By making the bid increment smaller we may in this way approach maximal privacy-preservation.

The results of this section hold under the assumption that each bidder's preferences are modeled accurately in bid tables, an assumption which we have shown to be more restrictive than the gross substitutes property. Recent trends in auction design suggest that the use of auctions which work well under the gross substitutes property (e.g. those of Gul and Stachetti, 1999a; Ausubel, 2005) may be used to reveal information necessary for bidders to make decisions in an auction allowing for more general expression of preferences, those for which gross substitutes does not hold. At the end of a Dynamic Bid Table Auction, each bidder should be comfortable that he has bid enough on individual items without being exposed to the risk of paying too much for substitute items.

6 Conclusions

Combinatorial auctions promise to increase efficiency and reduce exposure risk in an economic environment where "synergy" is significant. Applications for which combinatorial auctions achieve more desirable outcomes than traditional market mechanisms abound in both governmental allocation problems and B2B commerce, including shipping-lane and procurement auctions in the private sector, auctions for spectrum licenses by the FCC and airport landing slots by the FAA in the public sector. In each of these environments, the expression of aggregate information allows the bidders to realize synergies (e.g., economies-of-scale or owning complementary items) while the auction mechanism stimulates competition, aiding the seller of items with more competitive prices. This increased satisfaction in market outcomes comes at the cost of computational difficulty, and the design of markets which quickly and easily achieve beneficial outcomes is certain to provide interesting new avenues of research for years to come.

Here we focus primarily on the bidder's exponential bundles problem by looking at a special case for which the size of communication grows polynomially in the number of items being auctioned. Implicitly, we ask: how can a simple compact representation of preferences (price-vectors) be combined to form more elaborate statements of preference? This question follows somewhat naturally from the economic literature of Kelso and Crawford (1982), who introduce unit-demand bidders, each with preferences described by a price-vector, in a model of the job-market, in which an employee can accept at most one job. Kelso and Crawford (1982) also introduce the gross substitutes property, which has become fundamental in the study of auctions. The strength of this concept in categorizing preferences and guaranteeing the existence of a unique Walrasian linear-price equilibrium has influenced several authors, including Ausubel and Milgrom (2002), and Gul and Stachetti (1999a), all of whom provide foundational work for the research presented here.

Other results on unit-demand bidders are given by Demange et al. (1986), leading naturally into our own investigation of bid tables. Indeed, the case of unit-demand bidders is well studied, and we note that the current paper generalizes the Walrasian equilibrium results of Demange, Gale and Sotomayor (as presented, for example, by Roth and Sotomayor, 1992), from the case of unit-demand bidders to the case in which each bidder is represented by multiple unit-demand agents.

Representing bidders by multiple unit-demand agents results in the fairly natural and easy-to-read bid table format, causing us to wonder, why hasn't this been investigated before? We note that among the incentive properties of their unit-demand bidder auction model, Roth and Sotomayor (1992) show that a bidder cannot benefit from shill bidding (having someone else join the auction to distort its outcome). This may seem to imply that representing a bidder by multiple unit-demand agents would be unrewarding, but the model under which this property was proven maintains the assumption of a unit-demand valuation function for each bidder. We, on the other hand, find use for the bid table format within the more general context of multi-unit demand.

How general is the preference expression afforded by the bid table format in the multi-unit demand context? We have provided an exact characterization: assignment preferences (those which can be written in bid tables) are properly contained in the class of gross substitutes valuation functions. Applying a result of Gul and Stachetti (1999b) based on this characterization, the gross substitutes property elucidates the greatest strength of a bid table auction: unique lowest Walrasian equilibrium price signals can be computed at each round of submission. The computation of these attractive linear price signals is facilitated by our constrained optimization approach, which allows us to recognize and neutralize "self-competition" constraints. This new approach overcomes the failure of the Hungarian algorithm to provide the truly lowest Walrasian equilibrium prices in this context, as it does in the unit-demand bidder context.

While the gross substitutes property does indeed provide several strengths of the bid table environment, it also clearly exposes its weaknesses. Most notably, preferences for complementary bundles can not be expressed in bid tables alone. We therefore demonstrated the efficacy of a Dynamic Bid Table Auction which can be used as the first stage of a hybrid auction procedure, capturing substitutable preferences in bid tables and deferring to a later package auction phase for complementary expression. Further details on the design of such a hybrid auction are given in Day and Raghavan (2005), but with the current work we have established much of the foundational theory for that design, allowing us to implement a better revelation phase for a combinatorial hybrid auction. These results have included algorithms for computing better linear price signals, and for upholding the gross substitutes property in a strict sense, where earlier and more general algorithms fail to do so.

References

Ausubel, L. 2005. An efficient dynamic auction for heterogeneous commodities. *American Economic Review* To appear.

Ausubel, L., P. Cramton, and P. Milgrom 2006. The clock-proxy auction: A practical combinatorial auction design. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 5. MIT Press.

Ausubel, L. and P. Milgrom 2002. Ascending auctions with package bidding. *Frontiers of Theoretical Economics* **1**(1).

Ball, M., G. Donohue, and K. Hoffman 2006. Auctions for the safe, efficient and equitable allocation of national airspace systems resources. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 22. MIT Press, forthcoming.

Clarke, E. 1971. Multipart pricing of public goods. *Public Choice* **8**, 19–33.

Day, R. and S. Raghavan 2005. A combinatorial procurement auction featuring bundle price revelation without free riding. Working paper, University of Connecticut. Url: http://users.business.uconn.edu/bday/DayRag.pdf.

Demange, G., D. Gale, and M. Sotomayor 1986. Multi-item auctions. *J. of Political Economy* **94**(4), 863–872.

Groves, T. 1973. Incentives in teams. *Econometrica* **41**, 617–631.

Gul, F. and E. Stachetti 1999a. The english auction with differentiated commodities. *J. of Economic Theory* **92**, 66–95.

— 1999b. Walrasian equilibrium with gross substitutes. *J. of Economic Theory* **87**, 95–124.

Kelso, A. and V. Crawford 1982. Job matching, coalition formation and gross substitutes. *Econometrica* **50**, 1483–1504.

Lehmann, B., D. Lehmann, and N. Nisan 2001. Combinatorial auctions with decreasing marginal utilities. *EC 2001* .

Nemhauser, G. and L. Wolsey 1999. *Integer and Combinatorial Optimization*. John Wiley and Sons.

Roth, A. and M. Sotomayor 1992. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge Univ. Press.

Rothkopf, M., T. Teisberg, and E. Kahn 1990. Why are Vickrey auctions rare? *Journal of Political Economy* **98**, 94–109.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance* **16**, 8–37.