Gross Substitutes and Endowed Assignment Valuations*

Michael Ostrovsky† Renato Paes Leme‡

August 26, 2014

Abstract

We show that the class of preferences satisfying the Gross Substitutes condition of Kelso and Crawford (1982) is strictly larger than the class of Endowed Assignment Valuations of Hatfield and Milgrom (2005), thus resolving the open question posed by the latter paper. In particular, our result implies that not every substitutable valuation function can be "decomposed" into a combination of unit-demand valuations.

1 Introduction

The notion of Gross Substitutes (GS) for preferences over bundles of indivisible goods (Kelso and Crawford, 1982) plays a critical role in a wide variety of theoretical and practical settings. When agents' preferences satisfy the GS condition, stable matchings are guaranteed to exist in two-sided matching markets (Kelso and Crawford, 1982; Roth, 1984; Hatfield and Milgrom, 2005); competitive equilibria in exchange economies are guaranteed to exist (Bikhchandani and Mamer, 1997; Gul and Stacchetti, 1999); the efficient, incentive-compatible Vickrey-Clarke-Groves mechanism has many additional attractive properties in combinatorial auction environments (Ausubel and Milgrom, 2006); and the resulting settings have many other useful characteristics, such as tractable and well-behaved comparative statics. In contrast, when some agents' preferences do not satisfy the GS condition, these results typically do not hold, substantially complicating both the theoretical analysis of such settings and the practical design of markets for such allocation problems.[1]

*We are grateful to Rakesh Vohra for sharing with us the unpublished results of his work with Rudolf Müller and Sven de Vries, and to the editor and two anonymous referees for helpful comments and suggestions. The first author is grateful to Google for its hospitality during the 2013–14 academic year, when this work was completed.

†Graduate School of Business, Stanford University; ostrovsky@stanford.edu.

‡Google Research; renatoppl@google.com.

[1]Two classes of settings in which many positive results hold despite the violations of the GS condition are exchange economies with two complementary classes of substitutable goods (Sun and Yang, 2006, 2009) and

For this reason, the question of understanding what classes of preferences satisfy the GS condition has direct practical and theoretical importance, and has attracted considerable attention in the literature, which proposed various alternative characterizations of the condition as well as classes of preferences satisfying it.[2] However, a question that has remained open is whether the class of preferences satisfying the GS condition is limited to those that can be "decomposed" into combinations of single-unit demands. Hatfield and Milgrom (2005) formalized this question by introducing a rich class of preferences, based on single-unit demands, that they called "Endowed Assignment Valuations" (EAV), showed that all valuation functions in that class satisfy the GS condition, and posed the question of whether this class exhausts the set of GS preferences.[3] Müller, Vohra, and de Vries (personal communication; subsequently Müller et al.) constructed an example of a GS valuation function that they conjectured does not belong to EAV, and showed that the results of matroid theory imply a weaker version of this result (but not the actual result itself; see below for details). We show that the valuation function constructed by Müller et al. is in fact outside of EAV, thus proving that the scope of GS is strictly larger than that of EAV. The main step in the proof is to identify a property, "strong exchangeability," that every EAV function must satisfy. We then show that the valuation function in the example does not satisfy this property. While the inspiration for the proof also comes from matroid theory (specifically, the theory of strongly exchangeable matroids due to Brualdi (1969)), our exposition is completely self-contained and involves only elementary mathematical arguments.

2 Setup

There is a finite set S of objects in the economy. An agent's *valuation function* $v : 2^S \to \mathbb{R}$ assigns a value to every subset of S. Without loss of generality, $v(\varnothing) = 0$. Also, valuation functions are monotone: for any sets X and Y such that $X \subseteq Y$, $v(X) \leq v(Y)$.

For a vector of prices $p \in \mathbb{R}^{|S|}$ and a bundle of objects $X \subseteq S$, denote by $p(X) = \sum_{i \in X} p_i$ the price of bundle X. The demand of the agent with valuation function v given prices p is

economies with agents who can both buy and sell goods, with selling goods being complementary to buying other goods (Ostrovsky, 2008; Hatfield et al., 2013). In both of those settings, however, the preferences can be "transformed" into those satisfying the GS condition (see, e.g., Section 3 of Sun and Yang, 2006, and Section III.A of Hatfield et al., 2013), and these transformations are precisely what makes it possible to achieve the positive results.

[2]See, e.g., Kelso and Crawford (1982); Gul and Stacchetti (1999); Reijnierse et al. (2002); Fujishige and Yang (2003); Bing et al. (2004); Hajek (2008); Hatfield et al. (2012).

[3]Hatfield and Milgrom (2005, p. 927) say: "To the best of our knowledge, all of the substitutes valuations that have been used or proposed for practical applications are included among the endowed assignment valuations. Indeed, the question of whether all substitutes valuations are endowed assignment valuations is an open one."

the collection of bundles of objects that maximize the agent's payoff, net of prices:

$$D(p) = \underset{X \subseteq S}{\operatorname{argmax}} \{v(X) - p(X)\}.$$

Note that for some price vectors p, the agent may be indifferent between two or more bundles of objects, and in that case $D(p)$ contains multiple bundles. For any price vector p, $D(p)$ contains at least one bundle (possibly the empty one, if prices are too high).

2.1 Gross Substitutes

Valuation function v satisfies the *Gross Substitutes* condition if raising the prices of some of the objects does not decrease the demand for other objects. Formally,

Definition 1 *Valuation function v satisfies the* Gross Substitutes *condition if for any pair of price vectors p and p' such that $p' \geq p$, for any bundle $X \in D(p)$, there exists bundle $X' \in D(p')$ such that for all objects $j \in X$ such that $p_j = p'_j$, we have $j \in X'$.*[4]

2.2 Endowed Assignment Valuations

The notion of *Assignment Valuations* (AV), due to Shapley (1962), can be described as follows. Consider a firm that has a set of positions, J. There is a set of objects, S. An object i assigned to position j generates output α_{ij}. The valuation of the firm for a set of objects X is equal to the highest amount of output it can produce by matching some of the objects in X to some positions in J (only one object can be matched to a position, and vice versa, and some objects and positions can remain unmatched).[5] Formally,

Definition 2 *Valuation function v over set S of objects is an* Assignment Valuation *if there exists a set of positions, J, and a matrix α of dimension $|S| \times |J|$ such that for any set $X \subseteq S$,*

$$v(X) = \max_{z} \sum_{i \in X,\, j \in J} \alpha_{ij} z_{ij},$$

where z varies over all possible assignments of elements in X to elements in J.[6]

The notion of *Endowed Assignment Valuations* (EAV), due to Hatfield and Milgrom (2005), extends AV as follows. Consider a firm that has a set of positions, J, and *has already*

[4] Expression "$p' \geq p$" in the definition means that for all $i \in S$, $p'_i \geq p_i$.

[5] Shapley (1962) talks about assigning "men" to "machines." For consistency with the rest of our paper, we talk about assigning "objects" to "positions" instead.

[6] That is, $z_{ij} \in \{0, 1\}$ for all i and j; $\sum_{j \in J} z_{ij} \in \{0, 1\}$ for all $i \in X$; and $\sum_{i \in X} z_{ij} \in \{0, 1\}$ for all $j \in J$.

purchased a set of objects T, i.e., set T is that firm's *endowment*. There is also a set of objects, S, which the firm can purchase in addition to T. The valuation of the firm over set $S \cup T$ is an assignment valuation, as defined above. Then the *incremental* valuation of the firm over bundles in S is an endowed assignment valuation. Formally:

Definition 3 *Valuation function v over set S of objects is an* Endowed Assignment Valuation *if there exists another set of objects, T, and an Assignment Valuation function w over set $S \cup T$ such that for all $X \subseteq S$, $v(X) = w(X \cup T) - w(T)$.*

Hatfield and Milgrom (2005) show that every endowed assignment valuation v satisfies the GS condition, and also show that EAV is precisely the family of valuations that is obtained by starting out with unit-demand valuations (i.e., values α_{ij} of "matches" between individual objects and individual positions) and then constructing richer valuations out of these simple ones by repeatedly "merging" valuations together (e.g., "merging" two one-position firms to obtain a two-position one) and using the "endowment" operation (as in the transition from AV to EAV above).[7] Thus, if classes of GS and EAV preferences were equal, that would imply that every GS valuation can be "decomposed" into a combination of simple unit-demand valuations. Our main result shows that this is not the case: some GS valuations cannot be decomposed in this fashion.

3 Main Result

Theorem 1 *The class of Gross Substitutes valuations is strictly larger than the class of Endowed Assignment Valuations.*

The rest of this section contains the proof of Theorem 1. The proof proceeds as follows. First, we identify a property, "strong exchangeability," that every EAV function satisfies. Second, we show that the valuation function constructed by Müller et al. does not satisfy this property, but does satisfy the GS condition.

Step 1: Strong Exchangeability

Definition 4 *Valuation function v is* strongly exchangeable *if for every price vector $p \in \mathbb{R}^{|S|}$ such that $|D(p)| \geq 2$, for every pair of inclusion-minimal[8] bundles X and Y in $D(p)$ such that $X \neq Y$, there exists a bijective function σ from the set $X \setminus Y$ to the set $Y \setminus X$ such that for every $i \in X \setminus Y$, both bundles $X \cup \{\sigma(i)\} \setminus \{i\}$ and $Y \cup \{i\} \setminus \{\sigma(i)\}$ are in $D(p)$.*

[7] Theorems 13 and 14 in Hatfield and Milgrom (2005).

[8] Bundle X is inclusion-minimal in $D(p)$ if $X \in D(p)$ and for every strict subset X' of set X, $X' \notin D(p)$.

Lemma 1 *Every endowed assignment valuation function v is strongly exchangeable.*

Proof.

Consider an EAV function v and a vector of prices p, and suppose there exist bundles X and Y such that $X \neq Y$, $v(X) - p(X) = v(Y) - p(Y) = \max_{Z \subseteq S} v(Z) - p(Z)$, and for every $X' \subsetneq X$ and $Y' \subsetneq Y$, $v(X') - p(X')$ and $v(Y') - p(Y')$ are both strictly smaller than $v(X) - p(X) = v(Y) - p(Y)$.

Consider the set of "endowed objects" T, the set of "positions" J, and the matrix of match values $\alpha \in \mathbb{R}^{(|S|+|T|)\times|J|}$, as in Definitions 2 and 3 above. We thus have $w(T) = \max_z \sum_{i \in T,\, j \in J} \alpha_{ij} z_{ij}$, where z varies over all possible assignments of elements in T to elements in J, and for any bundle $\tilde{S} \subseteq S$, we have $v(\tilde{S}) = \max_{\tilde{z}} \sum_{i \in \tilde{S} \cup T,\, j \in J} \alpha_{ij} \tilde{z}_{ij} - w(T)$, where $\tilde{z}$ varies over all possible assignments of elements in $\tilde{S} \cup T$ to elements in J.

Take a profit-maximizing assignment between the objects in $X \cup T$ and positions in J under prices p corresponding to the subset X being chosen from S, and call this assignment z_X. Formally,

$$z_X \in \operatorname*{argmax}_z \sum_{i \in X \cup T,\, j \in J} \alpha_{ij} z_{ij},$$

where z varies over all possible assignments of elements in $X \cup T$ to elements in J. Similarly, take a profit-maximizing assignment z_Y corresponding to the subset Y being chosen:

$$z_Y \in \operatorname*{argmax}_z \sum_{i \in Y \cup T,\, j \in J} \alpha_{ij} z_{ij},$$

where z varies over all possible assignments of elements in $Y \cup T$ to elements in J.

Note that we allow prices to be negative, and for any object i such that $p(i) < 0$, we have $i \in X$ even if object i is not matched to any positions under z_X. If, however, $p(i) \geq 0$ and $i \in X$, then i has to be matched to a position under z_X: If $p(i) \geq 0$, $i \in X$, and i is not assigned to any position to under z_X, then $v(X \setminus \{i\}) - p(X \setminus \{i\}) \geq v(X) - p(X)$, contradicting either the optimality or the inclusion-minimality of bundle X. Analogously, we have $i \in Y$ whenever $p(i) < 0$, and when $p(i) \geq 0$, then $i \in Y$ only if i is assigned to some position under z_Y.

Now, construct the following colored bipartite graph with objects from $S \cup T$ as vertices on one side and positions from J as vertices on the other side. For every assignment between an object i in $S \cup T$ and a position j in J under z_X, draw a red edge connecting i and j. For every assignment between an object i' in $S \cup T$ and a position j' in J under z_Y, draw a blue edge connecting i' and j'. (Some object-position pairs—those that are matched to each other under both z_X and z_Y—are connected by two different edges.)

In the graph, each node has degree zero, one, or two. Thus, the graph can be decomposed into disjoint paths and cycles that have no vertices in common. Take any object $i \in X \setminus Y$. Note that it has degree one in the graph.[9] Therefore, it is an end of a path (this path can be as short as just one red edge, or it can consist of a chain of edges of alternating colors, starting with a red one). The other end of the path can be a position in set J (if the number of edges in the path is odd) or an object in set T, an object in set $Y \setminus X$, or an object in set $X \cap Y$ (if the number of edges in the path is even).

Moreover, note that the sum of "match values" $(\alpha_{ij} - p_i)$ over the red edges in this path has to be equal to the sum of match values over the blue edges in this path.[10,11] The reason for this equality is that if it did not hold, we could "swap" one set of edges for another and obtain an assignment of objects to positions with a profit higher than that of z_X and z_Y.[12]

Next, note that the other end of the path has to be an object in $Y \setminus X$. In all other cases (the other end is a position in J, or an object in T, or an object in $X \cap Y$), *all* objects involved in this path would be in $X \cup T$, and by "swapping" the red edges in this path for the blue edges in it,[13] we would obtain an assignment with the same total profit as that of set X and assignment z_X, but with a set of non-endowed objects $X \setminus \{i\}$ instead of X—which

[9] Object i cannot have degree two, because it does not belong to bundle Y. Also, $i \notin Y$ implies that $p(i) \geq 0$. As we observed above, any object $i \in X$ such that $p(i) \geq 0$ has be to matched to some position under assignment z_X. Thus, it has degree one in the graph.

[10] The same statement is true for the cycles in the graph—the sum of match values over the red edges in any cycle is equal to the sum of match values over the blue edges in that cycle. Our proof, however, only relies on the observation of the equality of sums of match values for paths.

[11] In the case when the path has no blue edges, the sum of match values over the set of blue edges is equal to zero. In the case when the other end of the path, i^*, is an object in set $X \cap Y$, the sum of match values over the red edges includes $(-p(i^*))$ (note that as we observed above, if $i^* \in X$ and i^* is not assigned to any position under z_X, it has to be the case that $p(i^*) < 0$).

[12] Formally: Suppose the sum of "match values" over the set of red edges in the path is strictly larger than that over the set of blue ones. Note that object i is the only object involved in the path that belongs to $X \setminus Y$, and let B denote the set of objects in the path that belong to $Y \setminus X$ (set B may consist of one or zero elements: if the other end of the path is an object in $Y \setminus X$, then $|B| = 1$; if the other end of the path is a position in J or an object in $X \cap Y$ or T, then $|B| = 0$). Let $Y' = Y \cup \{i\} \setminus B$, and let z' be an assignment of objects in Y' to positions in J that coincides with assignment z_Y (i.e., corresponds to the blue edges) on all elements of Y' that are not involved in the path, and coincides with z_X (i.e., corresponds to the red edges) on all elements of Y' that *are* involved in the path. Then

$$v(Y') - p(Y') \geq \sum_{k \in Y', j \in J} \alpha_{kj} z'_{kj} - w(T) - p(Y') > \sum_{k \in Y, j \in J} \alpha_{kj} (z_Y)_{kj} - w(T) - p(Y) = v(Y) - p(Y),$$

where the first inequality holds because z' is a feasible assignment of objects to positions for set Y' (though not necessarily the optimal one), and the second inequality holds because the difference between $\sum_{k \in Y', j \in J} \alpha_{kj} z'_{kj} - p(Y')$ and $\sum_{k \in Y, j \in J} \alpha_{kj} (z_Y)_{kj} - p(Y)$ is precisely the difference between the sum of match values over the set of red edges in the path and the sum of match values over the set of blue edges. Thus, $v(Y') - p(Y') > v(Y) - p(Y) = v(X) - p(X)$, contradicting the assumption that bundles X and Y were optimal given the vector of prices p. The case in which the sum of match values is larger over the set of blue edges in the path than over the set of red ones is completely analogous.

[13] "Swapping" can again be formalized in a way analogous to footnote 12.

would violate the assumption that X was inclusion-minimal.

Thus, we can establish a one-to-one mapping σ between the objects in $X \setminus Y$ and $Y \setminus X$ by following the paths connecting them in the red–blue graph. The fact that for every $i \in X \setminus Y$, both bundles $X \cup \{\sigma(i)\} \setminus \{i\}$ and $Y \cup \{i\} \setminus \{\sigma(i)\}$ are in $D(p)$ follows from the observation that the sum of match values along the red edges is equal to the sum of match values along the blue edges, for every path. ■

Step 2: Not-strongly-exchangeable GS Valuation Function

Consider the following valuation function.[14] There are six objects in set S, graphically represented as the edges of the complete graph with four vertices (see Figure 1).

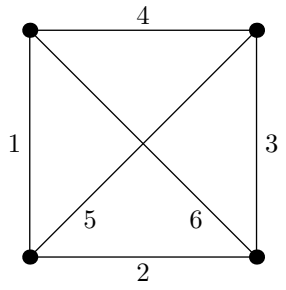

Figure 1: Graphical representation of the six objects in set S

Define valuation r on this set S as follows. For a set of objects (i.e., edges) $X \subset S$, $r(X)$ is equal to the size of the largest subset of X that does not contain any cycles. In other words, for any set X of size at most 2, $r(X) = |X|$; for any set X of size at least 4, $r(X) = 3$, and for any set X of size 3, $r(X) = 2$ if the three objects in X form a cycle, and $r(X) = 3$ if the three objects in X do not form a cycle. (One interpretation of this valuation function is that it represents some network, and additional edges are only valuable when they allow new connections that are not already available without them.)

Lemma 2 *Valuation function r is not strongly exchangeable.*

Proof. Let $p = 0$, so the profit from any bundle is equal to its valuation. Note that $r(\{1, 2, 3\}) = r(\{4, 5, 6\}) = 3$, and so both bundles are inclusion-minimal maximizers. However, there is no "strongly exchangeable" one-to-one mapping between the two. Indeed, edge 1 can only be exchanged for edge 4 (it cannot be exchanged for edge 5, because

[14] This valuation function was conjectured by Müller et al. to be a potential example of a GS valuation that is not an EAV. Müller et al. showed that the results of matroid theory imply that this valuation function satisfies the GS condition, and that it cannot be represented as an Endowed Assignment Valuation if all the elements in matrix α of match values are constrained to be 0 or 1. Those results, however, do not imply the impossibility of such a representation with general EAV functions, which is what our main result shows.

$r(\{5,2,3\}) = 2 < 3$; and it cannot be exchanged for edge 6, because $r(\{4,5,1\}) = 2 < 3$). Likewise, edge 3 can only be exchanged for edge 4. Since there is no one-to-one mapping under which both edge 1 and edge 3 are mapped to edge 4, valuation function r is not strongly exchangeable. ■

Lemma 3 *Valuation function r satisfies the Gross Substitutes condition.*

Proof. This result follows from the fact that function r is a matroid rank function.[15] However, for completeness, we provide a self-contained proof of Lemma 3, which does not rely on any results or definitions of matroid theory (although of course the ideas of the proof are closely related to that theory). The proof also illustrates that function r is not an "edge case": there is a rich class of GS valuations that are not EAV. The self-contained proof is in the Appendix. ■

Thus, by Lemmas 1 and 2, valuation r does not belong to the class of Endowed Assignment Valuations. By Lemma 3, valuation r satisfies the Gross Substitutes condition. Combined with the fact that EAV $\subseteq$ GS (Hatfield and Milgrom, 2005), these two observations conclude the proof of Theorem 1.

4 Matroid-Based Valuations

Our result shows that not all GS valuations can be "built from" unit-demand valuations. It thus leads to a natural question: What *are* the fundamental building blocks for GS valuations? We now describe one class of valuations, which by our main result is strictly larger than the set of EAV, and conjecture that this class of valuations is in fact equal to those that satisfy the GS condition. We call this class *Matroid-Based Valuations*. To define this class of valuation functions, we need to recall the definitions of a matroid and a weighted matroid.

Definition 5 *A matroid is a pair $(S, \mathcal{I})$ where S is a finite set of objects (called the ground set of the matroid) and $\mathcal{I}$ is a collection of subsets of set S (called the independent sets of the matroid) such that*

- *$\varnothing \in \mathcal{I}$ (i.e., the empty set of objects is independent);*
- *if $X \in \mathcal{I}$ and $X' \subseteq X$, then $X' \in \mathcal{I}$ (i.e., any subset of an independent set is independent);*

[15] Specifically: r is the rank function of matroid $M(K_4)$ (Oxley, 1992); every matroid rank function is $M^\natural$-concave (Murota, 1996; Murota and Shioura, 1999); and every $M^\natural$-concave function satisfies the Gross Substitutes condition (Fujishige and Yang, 2003).

- *if $X, Y \in \mathcal{I}$ and $|X| < |Y|$, then there exists object $y \in Y \setminus X$ such that $X \cup \{y\} \in \mathcal{I}$.*

Definition 6 *A weighted matroid is a tuple $(S, \mathcal{I}, w)$, where $(S, \mathcal{I})$ is a matroid and w is a weight function $w : S \to \mathbb{R}_+$ that assigns a non-negative value to every object in the ground set E.*

A valuation function v induced by weighted matroid $(S, \mathcal{I}, w)$ is constructed as follows. Take any set of objects $X \subset S$. Then the valuation of this set is given by

$$v(X) = \max_{Y \subseteq X,\, Y \in \mathcal{I}} \sum_{y \in Y} w(y).$$

By the results of Murota (1996), Murota and Shioura (1999) and Fujishige and Yang (2003), every valuation function v induced by a weighted matroid satisfies the GS condition. The opposite, however, is not true. To see that, note that every valuation function v induced by a weighted matroid satisfies the following property. Take any set $X \subseteq S$ such that for every $X' \subsetneq X$, $v(X') < v(X)$. Then it is the case that $v(X) = \sum_{x \in X} v(x)$. Consider now set $S = \{a, b\}$ and valuation function u given by $u(\varnothing) = 0$, $u(\{a\}) = u(\{b\}) = 1$, $u(\{a, b\}) = 1.5$. Valuation function u satisfies the GS condition, but it is not the case that $u(\{a, b\}) = u(\{a\}) + u(\{b\})$.

Consider now the two substitutability-preserving operations used to generate the class of EAV preferences from unit-demand ones: "merging" and "endowment".[16] Formally, the merging operation takes two valuation functions v_1 and v_2 over set S of objects, and defines the "merged" valuation function v^* as follows: for any set $X \subseteq S$,

$$v^*(X) = \max_{Y \subseteq X} \left(v_1(Y) + v_2(X \setminus Y) \right).$$

The "endowment" operation takes a valuation function v over some set $S \cup T$ of objects (with $S \cap T = \varnothing$) and defines valuation v' over subsets X of set S as

$$v'(X) = v(X \cup T) - v(T).$$

Both of these operations preserve substitutability (i.e., if functions v_1, v_2, and v above satisfy the GS condition, then functions v^* and v' satisfy it as well). Thus, if one starts with a class of substitutable valuation functions and repeatedly applies these two operations, the resulting class of valuation functions will also be substitutable. If one starts with unit-demand valuations, then as Hatfield and Milgrom (2005) show, the resulting class of valuation

[16] The "merging" operation is also known in the literature as "convolution" (Murota, 1996) and "OR" operation (Lehmann et al., 2006). The "endowment" operation is also known as the "marginal valuation".

functions is EAV. As the main result of our paper shows, EAV is strictly smaller than the class of GS valuations; in particular, there are valuation functions induced by weighted matroids (like the valuation function r induced by the matroid in Step 2 of the proof of main result, with all weights set to 1) that are not in EAV.

But what if one instead starts with the class of valuation functions induced by weighted matroids? Formally, let MBV (for *Matroid-Based Valuations*) be the smallest class of valuation functions that contains all valuation functions induced by weighted matroids and is closed under the operations of "merging" and "endowment." We know that MBV is strictly larger than EAV, and we know that MBV is a subset of valuation functions that satisfy the GS condition. It is an open question whether MBV = GS.

Conjecture 1 *The class of Matroid-Based Valuations is equal to the class of valuations that satisfy the Gross Substitutes condition.*

Appendix: Proof of Lemma 3

We will prove the following generalization of Lemma 3. Take any graph G. Let S be the set of edges of graph G. Consider the following valuation r over the subsets X of S:

$$r(X) = \text{the number of edges in the largest subset of } X \text{ that does not contain cycles.}$$

Then valuation function r satisfies the Gross Substitutes condition.

We will prove the following statement about valuation r. Take any vector of prices $p \in \mathbb{R}^{|S|}$ such that for any sets $X_1, X_2 \subseteq S$, $r(X_1) - p(X_1) \neq r(X_2) - p(X_2)$ (and in particular, $D(p)$ is single-valued, and so slightly abusing notation, we will denote by $D(p)$ the unique payoff-maximizing bundle). Increase the price of one object, i: take vector of prices $p' \in \mathbb{R}^{|S|}$ such that $p'_i > p_i$ and $p'_j = p_j$ for all $j \neq i$, and for any sets $X_1, X_2 \subseteq S$, $r(X_1) - p'(X_1) \neq r(X_2) - p'(X_2)$ (and in particular, $D(p')$ is again single-valued). Take any object $j \neq i$. Then if $j \in D(p)$, then $j \in D(p')$.[17]

The proof of the above statement will rely on the following graph-theoretic observation. Take any set of edges $X \subsetneq S$ and three distinct edges a, b, and c that are not in X. Suppose (i) the set of edges $X \cup \{a, b\}$ contains a cycle that contains edges a and b, (ii) the set of edges $X \cup \{b\}$ does not contain a cycle that contains edge b, (iii) the set of edges $X \cup \{a, c\}$ contains a cycle that contains edges a and c, and (iv) the set of edges $X \cup \{c\}$ does not contain a cycle that contains edge c. Then the set of edges $X \cup \{b, c\}$ contains a cycle that contains edges b and c.[18]

Observe now that for both vectors of prices p and p', the demands under those vectors can be constructed using the following "greedy" procedure. First, order the objects from

[17] This statement appears to be weaker than the definition of Gross Substitutes, because (a) it only considers vectors of prices under which indifferences between bundles do not arise and (b) it involves raising the price of only one object, rather than several ones. However, this definition is in fact equivalent to Definition 1. To address issue (a), one can perturb the prices by a very small amount in such a way that indifferences disappear and bundle X in Definition 1 becomes the unique demanded set. The corresponding unique set X' will then have the desired property—and will survive as a demanded set in the limit as the size of the perturbation is taken to zero. To address issue (b), one can simply raise prices one by one, and note that every time one of the prices increases, the previously demanded objects for which the price did not increase remain demanded in at least one optimal bundle after the increase.

[18] To see this, let e_1 and e_2 denote the two endpoints of edge a. Let X_1 denote the set of edges connecting edge b to endpoint e_1 in a cycle in $X \cup \{a, b\}$ that contains edges a and b; X_2 denote the set of edges connecting b to endpoint e_2 in that same cycle; X_3 denote the set of edges connecting edge c to endpoint e_1 in a cycle in $X \cup \{a, c\}$ that contains edges a and c; and finally let X_4 denote the set of edges connecting c to endpoint e_2 in that same cycle. Note that $(X_1 \cup X_3) \cap (X_2 \cup X_4) = \varnothing$ (because otherwise $X \cup \{b\}$ would contain a cycle that contains edge b or $X \cup \{c\}$ would contain a cycle that contains edge c). This, in turn, implies that one can find sets of edges $Y_1 \subseteq (X_1 \cup X_3)$ and $Y_2 \subseteq (X_2 \cup X_4)$ such that the set of edges $\{b\} \cup Y_1 \cup \{c\} \cup Y_2$ is a cycle.

the cheapest to the most expensive.[19] Next, go down the list of objects in that ordering. For each object, if adding it to the list of those already in the demanded set increases the total payoff (i.e., the incremental value of that object is higher than its price), then do add it. Otherwise, do not.[20,21,22]

Next, if $i \notin D(p)$, or $i \in D(p')$, then it is immediate that $D(p') = D(p)$, and thus $j \in D(p')$. Suppose $i \in D(p)$, $i \notin D(p')$. Suppose also that $j \in D(p)$ and $j \notin D(p')$—we will show that this will lead to a contradiction.

Without loss of generality, suppose j is the cheapest object (other than i) that is chosen in $D(p)$ but not in $D(p')$. Given that the demands can be constructed using the "greedy" procedure above, two statements must be true. First, $p_i < p_j$ (otherwise, increasing the price of objects i would not have had any effect on whether object j is demanded). Second, there is exactly one object, k, with the price between p_i and p_j that is not chosen under price p but is chosen under price p'.[23]

[19]Our "no indifferences" condition on vectors p and p' implies that all objects have different prices, so for each of these two price vectors, this ordering is unique.

[20]The "no indifferences" condition implies that the incremental value of an object is never equal to its price.

[21]To see that this procedure indeed produces the bundle with the highest payoff for the agent, suppose $X \neq D(p)$ is the bundle generated by the procedure (given price vector p). Note first that bundle X cannot be a strict subset of $D(p)$, because any object $i \in (D(p) \setminus X)$ has a positive incremental contribution when added to $D(p) \setminus \{i\}$, and thus also has a positive incremental contribution when added to any subset of $D(p) \setminus \{i\}$, and would therefore also have had a positive incremental contribution when it was encountered on the path of the greedy procedure.

Next, take the cheapest object $j \in X \setminus D(p)$. Let

$$K = \{k \in D(p) \setminus X \ : \ D(p) \cup \{j\} \text{ contains a cycle that contains } j \text{ and } k\}.$$

Note that j and every $k \in K$ have prices strictly between zero and one (otherwise, they would belong to either both X and $D(p)$ or neither). Note also that set K is not empty, and that there are no cycles containing any objects from K in $D(p)$. Also, all objects in K are more expensive than j (because j was the cheapest object in $X \setminus D(p)$, and so any object j' cheaper than j was considered prior to j by the greedy procedure—and if the incremental value of j' was not found to be positive by the greedy procedure, it could not be positive in the bundle $D(p)$).

Consider the set of cycles containing j in $D(p) \cup \{j\}$. From this set, pick a cycle, C, with the smallest number of objects from K. Take any $k \in K \cap C$. Consider the set $D(p) \cup \{j\} \setminus \{k\}$. The total cost of the objects in this bundle is cheaper than that in $D(p)$. And the total valuation of the bundle is the same: there is no cycle in $D(p) \cup \{j\}$ containing j and objects from $(K \cap C) \setminus \{k\}$ but not $K \setminus C$ (because of how C was chosen), and there is also no cycle in $D(p) \cup \{j\}$ containing j and objects from $(K \cap C) \setminus \{k\}$ and some object k' from $K \setminus C$ (because in that case, by the graph-theoretic observation above, $D(p)$ would have contained a cycle that contained k and k'). Thus, the net payoff from bundle $D(p) \cup \{j\} \setminus \{k\}$ is higher than that from bundle $D(p)$, contradicting the definition of $D(p)$.

[22]We could end the proof here: the fact that for any price vector, the optimal demand can be constructed using the "greedy" procedure implies that the valuation function satisfies the GS condition; in fact the two statements are equivalent (see e.g., Paes Leme, 2014). However, since the purpose of this Appendix is to provide a self-contained proof of Lemma 3, we include the additional arguments that conclude the proof.

[23]Clearly, there has to be at least one such object—otherwise, simply removing object i from the bundle could not have led to a decrease in the incremental value of object j, and so it would continue to be chosen

Consider now edges i, j, k, and the set of edges $T \subset S \setminus \{i, j, k\}$ that are cheaper than j and that are chosen under the vector of prices p (and thus also under the vector of prices p': for the objects that are cheaper than j, $D(p)$ and $D(p')$ only differ by i and k). Note that the prices of objects j and k are positive (otherwise, they would always be demanded, under both p and p'). Also, there is a cycle in the set $T \cup \{i, k\}$ that contains objects i and k, and no cycle containing k in the set $T \cup \{k\}$ (otherwise, the presence of i could not have affected the incremental value of k), and there is also a cycle in the set $T \cup \{k, j\}$ that contains objects k and j, and no cycle containing j in the set $T \cup \{j\}$ (otherwise, the presence of k could not have affected the incremental value of j). But these observations imply that there is a cycle in the set $T \cup \{i, j\}$ that contains objects i and j, which contradicts the assumption that $j \in D(p)$.

by the greedy procedure under p'. To see that there cannot be two (or more) such objects, assume the contrary, and take the two cheapest such objects, k_1 and k_2 (with k_1 being the cheaper of the two). Let $X = \{x \in (D(p) \setminus \{i\}) \, : \, p_x < p_{k_2}\}$, i.e., the set of objects chosen by the greedy procedure prior to object k_2 under the vector of prices p. Note that k_1 and k_2 must have prices strictly between zero and one. Note also that k_1 must belong to a cycle that also contains object i and objects in X, and cannot belong to any cycle that only contains k_1 and objects in X (but not i). Likewise, k_2 must belong to a cycle that also contains object i and objects in X, and cannot belong to any cycle that only contains k_2 and objects in X. The graph-theoretic observation then implies that k_1 and k_2 belong to a cycle that also contains objects in X, contradicting the assumption that both were chosen by the greedy procedure under the vector of prices p'.

References

Ausubel, Lawrence M. and Paul Milgrom (2006) "The lovely but lonely Vickrey auction," in *Combinatorial Auctions*: MIT Press.

Bikhchandani, Sushil and John W. Mamer (1997) "Competitive equilibrium in an exchange economy with indivisibilities," *Journal of Economic Theory*, Vol. 74, pp. 385–413.

Bing, Meir, Daniel J. Lehmann, and Paul Milgrom (2004) "Presentation and structure of substitutes valuations," in *ACM Conference on Electronic Commerce*, pp. 238–239.

Brualdi, Richard A. (1969) "Comments on bases in dependence structures," *Bulletin of the Australian Mathematical Society*, Vol. 1, pp. 161–167.

Fujishige, Satoru and Zaifu Yang (2003) "A note on Kelso and Crawford's gross substitutes condition," *Mathematics of Operations Research*, Vol. 28, pp. 463–469.

Gul, Faruk and Ennio Stacchetti (1999) "Walrasian equilibrium with gross substitutes," *Journal of Economic Theory*, Vol. 87, pp. 95–124.

Hajek, Bruce (2008) "Substitute valuations: Generation and structure," *Performance Evaluation*, Vol. 65, pp. 789–803.

Hatfield, John W., Nicole Immorlica, and Scott D. Kominers (2012) "Testing substitutability," *Games and Economic Behavior*, Vol. 75, pp. 639–645.

Hatfield, John W., Scott D. Kominers, Alexandru Nichifor, Michael Ostrovsky, and Alexander Westkamp (2013) "Stability and competitive equilibrium in trading networks," *Journal of Political Economy*, Vol. 121, pp. 966–1005.

Hatfield, John W. and Paul R. Milgrom (2005) "Matching with contracts," *American Economic Review*, Vol. 95, pp. 913–935.

Kelso, Alexander S., Jr. and Vincent P. Crawford (1982) "Job matching, coalition formation, and gross substitutes," *Econometrica*, Vol. 50, pp. 1483–1504.

Lehmann, Benny, Daniel J. Lehmann, and Noam Nisan (2006) "Combinatorial auctions with decreasing marginal utilities," *Games and Economic Behavior*, Vol. 55, pp. 270–296.

Murota, Kazuo (1996) "Convexity and Steinitz's exchange property," *Advances in Mathematics*, Vol. 124, pp. 272 – 311.

Murota, Kazuo and Akiyoshi Shioura (1999) "M-convex function on generalized polymatroid," *Mathematics of Operations Research*, Vol. 24, pp. 95 – 105.

Ostrovsky, Michael (2008) "Stability in supply chain networks," *American Economic Review*, Vol. 98, pp. 897–923.

Oxley, James G. (1992) *Matroid Theory*: Oxford University Press.

Paes Leme, Renato (2014) "Gross substitutability: An algorithmic survey," Working paper, Google Research.

Reijnierse, Hans, Anita van Gellekom, and Jos A. M. Potters (2002) "Verifying gross substitutability," *Economic Theory*, Vol. 20, pp. 767–776.

Roth, Alvin E. (1984) "Stability and polarization of interests in job matching," *Econometrica*, Vol. 52, pp. 47–57.

Shapley, Lloyd S. (1962) "Complements and substitutes in the optimal assignment problem," *Naval Research Logistics Quarterly*, Vol. 9, pp. 45–48.

Sun, Ning and Zaifu Yang (2006) "Equilibria and indivisibilities: gross substitutes and complements," *Econometrica*, Vol. 74, pp. 1385–1402.

——— (2009) "A double-track adjustment process for discrete markets with substitutes and complements," *Econometrica*, Vol. 77, pp. 933–952.