

A Graph Theoretic Approach to Markets for Indivisible Goods

Andrew Caplin and John Leahy*
New York University and N.B.E.R.

October 2010

Abstract

Many important markets, such as the housing market, involve goods that are both indivisible and of budgetary significance. We introduce new graph theoretic techniques ideally suited to analyzing such markets. In this paper and its companion (Caplin and Leahy [2010]), we use these techniques to fully characterize the comparative static properties of these markets and to identify algorithms for computing equilibria.

1. Introduction

While many important goods are indivisible, technical barriers continue to limit our understanding of markets for trading these goods. The best-studied cases are so-called allocation markets, in which each agent can consume at most one unit of one of the available indivisible goods. Shapley and Shubik (1972) provided a complete characterization of equilibria in such markets when utility is transferable. Yet understanding of the corresponding markets with non-transferable utility (NTU), a necessary feature when the goods in question are of budgetary significance, has advanced more slowly. Kaneko (1982) was first to establish conditions for existence of equilibria, while Demange and Gale (1985) showed under much the same conditions that the set of equilibrium prices is a lattice with maximal and minimal elements.¹

*We thank Vincent Crawford, Mamoru Kaneko, John Leahy Sr., Victor Norman, Ennio Stacchetti, and Ivan Werning for helpful comments. Leahy thanks the NSF for financial support.

¹They also established that the minimum price equilibrium cannot be manipulated by buyers, as well as some basic comparative static properties of the minimum price equilibrium (e.g. minimum equilibrium prices rise when more buyers are introduced).

The main barrier holding back understanding of markets for indivisible goods is technical. As Scarf [1994] has stressed, indivisibilities render the calculus of limited value in characterizing allocation markets. The reason that the transferable utility case is susceptible to analysis is precisely because it is equivalent to a problem without indivisibilities. It is tractable only because the tools of linear programming can be employed in modeling how the allocation switches with model parameters (Koopmans and Beckmann [1957]). The NTU case forces one to face many of the same complexities that render integer programming so notoriously complex. There may, for example, be “butterfly effects,” in which the smallest of parameter changes causes a global reallocation of goods. Against this technical backdrop, the limits to our understanding of markets in which significant goods are indivisible is readily understood.

In this paper and its companion (Caplin and Leahy [2010]) we introduce new mathematical structures for analyzing equilibria in NTU allocation markets. Our “GA-structures” combine an allocation of goods with a graph theoretic structure that represents indifference relations. In addition to having rich mathematical properties, GA-structures connect with a long-standing economic tradition, in particular the “rent gradient” models of Ricardo (Ricardo (1817), Alonso (1964), and Roback (1982), Kaneko, Ito, and Osawa [2006]). We show these GA-structures to have five properties that make them ideally suited to analyzing equilibria in NTU allocation markets.

1. We establish equivalence between minimum price equilibria and a class of optimization problems on GA-structures. Experience with the fundamental theorem of welfare economics shows how useful such a link between equilibrium theory and optimization theory can be. Optimization problems are simpler and better understood. They also do not require one to explicitly consider demand, supply, or the balance between them.
2. We use the link between optimization theory and equilibrium theory to invoke the standard theorem of the maximum, and thereby to fully characterize local comparative statics. We provide a chain rule for such local comparative statics. This chain rule establishes that small discrete shocks have local effects in the model. In contrast, with divisible goods, infinitesimal shocks have global effects: even the smallest change to the supply or demand for one good tends to affect the price of every other good in the economy.
3. We show that GA-structures can be used to identify the entire set of competitive equilibria, not only the minimum price equilibrium.
4. We use GA-structures to study how goods are reallocated as model parameters change.

In a generic case, we show that there are five and only five distinct forms of change in the equilibrium allocation in response to local parameter changes.

5. We show that GA-structures can be used algorithmically to identify minimum price equilibria.²

Properties 1-3 above are established in this paper, the last two in the companion paper. While we limit our attention in both papers to allocation markets, the GA-structures and associated techniques that we introduce may be relevant in other settings. We are currently working on reallocation markets and on general equilibrium market dynamics using many of the same mathematical tools. There may also be applications to auction markets, to matching markets, and to models of network formation.

The remainder of the paper is structured as follows. Section 2 discusses some related literature. Section 3 presents the basic model. Section 4 presents an example that illustrates the main objects of our analysis. Section 5 introduces GA-structures. Section 6 characterizes the minimum equilibrium price as the solution to an optimization problem on these structures. Section 7 characterizes the minimum price equilibrium allocation in a similar manner. Section 8 uses these characterizations to study the local dependence of minimum price equilibria on the economic environment. Section 9 defines a dual to the allocation problem, and uses it to characterize the complete set of equilibria. Section 10 concludes.

2. Related Literature

The standard approaches to indivisibilities either assume linear utility or make assumptions that smooth away the discreteness.

An example of the first approach is the model of Shapley and Shubik (1972). They showed that with linear utility the competitive equilibrium allocation in a market for heterogeneous, indivisible goods is equivalent to the problem of a social planner allocating goods so as to maximize the sum of utilities. This social planner's problem takes the form of the linear programming problem studied by Koopmans and Beckmann (1957).

The assumption of linear utility and the resulting absence of wealth effects may not be appropriate in many applications, especially if the good in question is an expensive one such as

²The transferable utility case is well covered in this regard: the Hungarian algorithm of Kuhn [1955] and Munkres [1957] can be used to compute the equilibrium allocation, while the ascending auction mechanism of Demange, Gale and Sotomayor (1986) solves for the minimum price equilibrium in a discretized version of the model.

a house. In the linear case, the social planner allocates goods based on some fixed notion of how much each agent desires each good. If a poor agent values a sea-view more than a rich agent, the planner will allocate a mansion by the sea to the poor agent. We do not, however, see many poor agents living in sea-side mansions. What is missing is the effect of diminishing marginal utility of wealth that leads the rich to be willing to pay more than the poor for the nicest homes. To include these effects it is necessary to consider utility functions that are non-linear in wealth.

An example of smoothing is Rosen's (1974) hedonic pricing model. It also prices heterogeneous goods given heterogeneous buyers and sellers. While goods themselves are indivisible, Rosen makes the assumption that there is a continuous density over characteristic bundles and that within this space one can adjust each characteristic while fixing the others. This assumption smoothes the type space allowing the use of the tools of calculus. In many applications, however, the type space may not be dense enough to allow such adjustments. In housing markets, for example, location is one of the most important characteristics. It is not generally possible to adjust location while keeping all other characteristics fixed, nor is it generally possible to alter characteristics of homes while maintaining a fixed location without incurring substantial costs. There is ample evidence in the urban economics literature that hedonic prices vary with location.³

There are some theoretical results in the case with indivisibilities and wealth effects. Kaneko (1982) established conditions for the existence of an equilibrium.⁴ Demange and Gale (1985) showed that the set of equilibrium prices is a lattice with maximal and minimal elements. They also established that the minimum price equilibrium cannot be manipulated by buyers, as well as some basic comparative static properties of the minimum price equilibrium. We extend on these results by illustrating the structure of minimum equilibrium prices.

Allocation problems arise naturally in a number of areas in economics. In the housing literature, the minimum equilibrium price vector is similar to the rent gradient found in Ricardo (1817), Alonso (1964), and Roback (1982). Models in this tradition tend to limit the heterogeneity in buyers or houses in order to keep the model tractable. At the same time, however, this simplicity allows them to go further than we will in modelling the supply side of the market.

In the auction and mechanism design literature, our equilibrium is similar to a second price auction or a Vickrey-Groves-Clark mechanism. These models almost always assume transferable utility. An exception is the paper by Demange and Gale (1985) cited above.

³See, for example, Meese and Wallace (1994).

⁴Quinzii (1984), Gale (1984), and Kaneko and Yamamoto (1986) also provide existence proofs. Crawford and Knoer (1981) sketch a proof of existence for a version of their model with non-transferable utility.

3. The Model

We work with a variant of the model in Demange and Gale (1985). Demange and Gale simplify the exposition and the analysis of allocation markets by removing all reference to budget constraints. This removes the need to discuss what transactions are feasible for each agent at each set of prices and ensures that the choice correspondences are continuous.⁵

There is a set of buyers $x_a \in X$, $1 \leq a \leq m$, and a set of indivisible goods $y_i \in Y$, $1 \leq i \leq n$. The goods are initially held by the sellers. Buyers may purchase the indivisible goods from sellers by making a transfer in terms of a homogeneous, perfectly divisible, numeraire good, which may be thought of as money. Sellers choose only whether or not to sell. They do not purchase the indivisible goods from other sellers. We assume that $n \geq m$ so that it is possible to match each buyer with a good.⁶

We assume that buyers can derive utility from at most one element of Y . The payoff for buyer x_a depends on the good that buyer purchases and the size of the transfer that the buyer makes to the seller. This payoff is summarized by the utility function $U_a : Y \times \mathbb{R}^n \rightarrow \mathbb{R}$, where $U_a(y_i, p_i)$ is the utility to x_a from the purchase of y_i at the price p_i .

Let $p \in \mathbb{R}^n$ denote the vector of goods prices. Each seller wishes to obtain the highest possible price above a reservation level. Let $r \in \mathbb{R}^n$ denote the vector of seller reservation prices. The supply side is trivial: each seller prefers to hold on to their good for any $p_i < r_i$ and to sell for any $p_i > r_i$. The seller is indifferent when $r_i = p_i$.⁷ Choosing $r \geq 0$ will ensure that all prices are positive if so desired.

Given any price vector $p \in \mathbb{R}^n$, the demand correspondence $D_a(p)$ specifies members of Y that maximize utility the utility of x_a :

$$D_a(p) = \{y_i \in Y \mid U_a(y_i, p_i) \geq U_a(y_k, p_k) \text{ for all } y_k \in Y\}.$$

An *allocation* is a one-to-one mapping $\mu : X \rightarrow Y$ from buyers to goods. It simplifies later

⁵With budget constraints, consumers' choice correspondences may not be continuous, and therefore the demand correspondence may fail to be upper-hemicontinuous. Assumptions (such as the Inada conditions) may be made to ensure that the constraints are not binding in equilibrium, but these do not add insight.

⁶This is without loss of generality. The possibility that a buyer may choose not to make a purchase can be captured by associating a subset of goods with exit.

⁷Since we will be interested in minimum price competitive equilibria, the exact form of a seller's utility does not matter so long as it is increasing in the transfer and there is a point r_i at which seller i is indifferent between selling and holding.

notation to let μ_a denote the good assigned to buyer x_a by the allocation μ ,

$$\mu_a \equiv \mu(x_a).$$

The set of all allocations is M . It will sometimes prove useful to have available the inverse mapping $\sigma : Y \rightarrow X \cup \emptyset$ such that $\mu(\sigma(y_i)) = y_i$ when $\sigma(y_i) \neq \emptyset$.

A competitive equilibrium is a price vector and an allocation such that all buyers choose optimally and all goods with prices above their reservation level are allocated. Given $p \in \mathbb{R}^n$, let $U(p) \equiv \{y_i \in Y | p_i > r_i\}$ denote the set of goods with prices strictly above seller reservation levels.

Definition A *competitive equilibrium* is a pair (p^*, μ^*) with $p^* \in \mathbb{R}^n$ and $\mu^* \in M$ such that:

1. $\mu_a^* \in D_a(p^*)$ for all $x_a \in X$.
2. $p_i^* \geq r_i$ for all $y_i \in Y$.
3. If $y_i \in U(p^*)$, then there exists $x_a \in X$ such that $\mu_a^* = y_i$.

The first condition is buyer optimality: the allocation must maximize the utility of each buyer. The second condition is seller optimality: no seller will part with a good for less than the reservation price. The third states that all goods with prices above reservation must be allocated. This ensures that supply is equal to demand.

We are interested in Π , the set of equilibrium prices, and, should they exist, the minimum and maximum equilibrium prices, respectively $\underline{p} \in \Pi$ and $\bar{p} \in \Pi$:

$$\begin{aligned} \Pi &= \{p \in \mathbb{R}^n | \exists \mu \in M \text{ s.t. } (p, \mu) \text{ an equilibrium}\}; \\ \underline{p} &\in \Pi \text{ is such that } p \in \Pi \implies p_i \geq \underline{p}_i \text{ all } i; \\ \bar{p} &\in \Pi \text{ is such that } p \in \Pi \implies p_i \leq \bar{p}_i \text{ all } i; \end{aligned}$$

We make assumptions on preferences that guarantee that utility is well behaved and an equilibrium exists.

Assumption A For each buyer $x_a \in X$ and good $y_i \in Y$,

1. $U_a(y_i, p_i)$ is continuously differentiable in p_i and strictly decreasing in p_i .
2. $\lim_{p_i \rightarrow \infty} U_a(y_i, p_i) = -\infty$ and $\lim_{p_i \rightarrow -\infty} U_a(y_i, p_i) = \infty$.

The first assumption is a regularity assumption that will allow us to use the implicit function theorem. Strict monotonicity simplifies the later analysis but is a stronger condition than needed for existence. The second assumption in combination with the first ensures that given any buyer, any two goods, and a price for one of the goods, there is a unique price for the second that makes the buyer indifferent between the two goods. Demange and Gale (1985) prove that under these conditions the set of equilibrium prices is a lattice and that there exists a minimum equilibrium price.⁸

4. A Motivating Example

A simple example will introduce some of the main objects of our analysis and some of the logic behind our characterization of minimum price equilibria.

Consider a market composed of two goods y_1 and y_2 and two buyers x_a and x_b . Suppose that the preferences of x_a are described by the utility functions $U_a(y_1, p_1) = 2 - p_1$ and $U_a(y_2, p_2) = 1 - p_2$. These have the property that when the prices of the two goods are equal x_a prefers good y_1 . Similarly, suppose that the preferences of x_b are $U_b(y_1, p_1) = 1 - p_1$ and $U_b(y_2, p_2) = 3 - p_2$ so that when the prices are equal x_b prefers good y_2 . Finally, suppose that sellers' reservation prices are $r_1 = r_2 = 0$.

The minimum price competitive equilibrium in this example is trivial: the price of each good is set equal to its reservation value and buyers purchase the goods they prefer.

We now discuss how to use “chains of indifference” to characterize the minimum price competitive equilibrium in this example. The idea behind a chain of indifference is that in any minimum price competitive equilibrium, any set of goods whose prices are strictly above reservation must contain a good that is demanded by some buyer allocated to a good outside of the set.⁹ Otherwise, we could reduce the prices of all the goods in this set and obtain a competitive equilibrium with lower prices. An implication is that each good is connected by indifference to a good whose price is the reservation price. Any good that is priced above its reservation value must be demanded by a buyer allocated to another good. If we knew which buyers were indifferent to which goods in equilibrium, we could build up the equilibrium price vector, starting with the goods priced at their reservation values and using the appropriate “chains of indifference” to price all other goods. The complication is knowing which buyers to assign to which goods and which goods should be connected through indifference.

⁸Demange and Gale assume that buyers may exit the market and therefore have a maximum willingness to pay. This ensures that Π has a maximal element as well.

⁹This is Lemma 4 in Demange and Gale (1985).

In the current example, there are two possible allocations: buyer x_a is matched either to y_1 or y_2 and buyer x_b is matched with the other good. Denote these allocations by μ^1 and μ^2 where $\mu_a^1 = y_1$ and $\mu_b^1 = y_2$, and $\mu_a^2 = y_2$ and $\mu_b^2 = y_1$. There are three potential chains of indifference, if we characterize chains by the goods that are to be connected through indifference. The first sets the price of y_1 to its reservation value r_1 and allows the price of y_2 to be set so that the buyer allocated to y_1 is indifferent between the two goods. The second reverses these roles: y_2 is set at its reservation value r_2 and the price of y_1 is set so that the buyer allocated to y_2 is indifferent between the two goods. The third possibility is that the prices of both goods are set at their reservation values.

An allocation and a chain together determine a price vector. For example, allocation μ^1 and chain 1 specify that y_1 is priced at reservation and the indifference of buyer x_a should be used to price y_2 . The price of y_1 is therefore set to $r_1 = 0$, and the price of y_2 is -1 since x_a is indifferent between the two goods when the price vector is $(0, -1)$. Table 1 reports the price vector that results from each chain and each allocation.

TABLE 1

Allocation\Chain	1	2	3	max sum
μ^1	(0,-1)	(-2,0)	(0,0)	(0,0)
μ^2	(0,2)	(1,0)	(0,0)	(1,2)
p^*				(0,0)

Our main result is that the minimum equilibrium price vector can be found by first maximizing the sum of prices across all potential chains for a given allocation, and then minimizing this across all potential allocations. In the current example, fixing the allocation and choosing the price vector that maximizes the sum of prices leads to the price vector in the last column. Minimizing this result with respect to the allocation leads to the price vector in last row. This price vector is the minimum equilibrium price vector.

Intuitively, two forces are at work. First, minimum equilibrium prices are determined by the willingness to pay of the next most interested buyer. Picking the wrong chain results in using the willingness to pay of a less interested buyer. This tends to lower the resulting price vector. This is why we take the maximum across chains. Second, allocating a consumer to the wrong good increases that consumer's willingness to pay for other goods. This tends to raise the resulting price vector. This is why we minimize across allocations.

The next two sections formalize these arguments. In the next section, we associate chains of indifference with a particular set of directed graphs on Y . We then show how to combine these graphs with allocations to generate prices such as those that appear in the cells of Table 1. The min-max theorem is presented in the succeeding section.

5. GA-Structures

What characterizes the construction of prices from a chain of indifference is that each good is either priced at reservation or it is connected by some unique path to a good that is itself priced at reservation. In graph theory, the property of there being a unique path from any vertex to the member of a set of source points is characteristic of a forest of rooted trees.¹⁰ The graphs that we are interested in are all forests of directed, rooted trees in which all edges point away from the root.

Definition The class \mathcal{F} comprises all directed graphs $F = (Y, R, E)$ with vertex set Y , root set $R \subseteq Y$, and edge set E that have the following properties:

1. F is a forest of trees.
2. E is a set of ordered pairs of vertices where for $e \in E$, $e = (y_1, y_2)$ is directed from y_1 to y_2 .
3. Each component (maximal connected subset) of F contains a unique element of R , and each edge in E is directed away from the corresponding element of R .

Figure 1 illustrates a directed, rooted tree. The vertices are shown as circles, except for the root vertex which is shown as a square. Each vertex corresponds to an indivisible good y_i . The edges are shown as arrows connecting one vertex to another. The edges are all directed away from the root node, y_1 . The absence of cycles characterizes the graph as a tree. A forest is a collection of such graphs.

[Figure 1]

We will write $E(F)$ and $R(F)$ when it is necessary to indicate to which graph the edge set and the root set belong. Let $e = (y_i, y_k) \in E$ denote the edge directed from good y_i to good y_k . We say that y_i is the tail of e , and y_k is the head of e . We also say that y_i is the direct predecessor of y_k and y_k is the direct successor of y_i . A standard and valuable observation is that

¹⁰A tree is a graph with no cycles. A forest is a graph whose components are trees. A rooted tree is a tree with one vertex denoted as the root.

for each non-root good $y_i \in Y \setminus R$, there exists a unique root good $y_r \in R$ and a corresponding unique directed path $\{(y_r, y_1), (y_1, y_2), \dots, (y_{i-1}, y_i)\} \subset E$ connecting the root set to y_i . We say that $y_k \neq y_i$ is a predecessor of y_i if y_k lies on this path between y_r and y_i . If y_k is a predecessor of y_i , we say that y_i is a successor of y_k .

We now show how to use a graph $F \in \mathcal{F}$ and an allocation μ to create a price vector. To do this, we limit attention to cases in which if $(y_i, y_k) \in E$, then μ allocates a buyer to y_i , the tail of the edge (y_i, y_k) .

Definition A graph-allocation structure (GA-structure) comprises a graph $F = (Y, R, E) \in \mathcal{F}$ and an allocation $\mu \in M$ such that, if $(y_i, y_k) \in E$, then there exists $x_a \in X$ such that $\mu_a = y_i$.

We let $\mathcal{G} \subset M \times \mathcal{F}$ denote the class of all such GA-structures.

We construct a mapping from GA-structures to prices, $q : \mathcal{G} \rightarrow \mathbb{R}^n$. The price mapping is derived by induction on the set of goods that we have priced. The idea is to first set the root goods at their reservation prices, and then to use the allocation μ and the graph F to construct chains of indifference. We price each non-root good using the indifference of the buyer allocated to its direct predecessor.¹¹ We let $q_i(\mu, F)$ denote the i th element of the vector $q(\mu, F)$.

Construction of $q(\mu, F)$ We construct $q(\mu, F) \in \mathbb{R}^n$ iteratively:

1. Define $A_0 \equiv R(F)$ and set $q_i(\mu, F) = r_i$ for all $y_i \in A_0$.
2. Given $s \geq 0$ and $q_i(\mu, F)$ for all $y_i \in A_s \subset Y$, let S comprise the set of direct successors of A_s ,

$$S = \{y_k \in Y \setminus A_s \mid \exists y_i \in A_s \text{ with } (y_i, y_k) \in E(F)\}.$$

For each $y_k \in S$, consider its direct predecessor $y_i \in A_s$ with $(y_i, y_k) \in E(F)$. Consider x_a such that $\mu_a = y_i$. Then $q_k(\mu, F)$ is defined implicitly by the indifference condition:

$$U_a(y_i, q_i(\mu, F)) = U_a(y_k, q_k(\mu, F)). \tag{5.1}$$

3. Set $A_{s+1} = A_s \cup S$. If $A_{s+1} = Y$, stop. Otherwise repeat the induction step.

¹¹This is similar to the rent gradient in Ricardo (1871) or the differential rent vector of Kaneko, Ito and Osawa (2006). Kaneko, Ito and Osawa make assumptions that guarantee that F has only one component that is not null, and that goods in this component have at most one direct successor. See also Miyake (2003) for a similar construction.

It is easy to see that this construction is well defined with Assumption A. Since every good is connected to the root set S will be non-empty so long as $A_s \neq Y$. Since $(\mu, F) \in \mathcal{G}$, there always exists $x_a \in X$ with $\mu_a = y_i$ in step 2. It then follows from Assumption A that there exists a unique $q_k(\mu, F) \in \mathbb{R}$ that satisfies (5.1). Finally, given the finite number of goods, this process will end after a finite number of steps with $A_s = Y$. Since F is a forest, there is a unique path to any good from the root set, so each element of S in step 2 has a unique direct predecessor. It follows that this construction defines a unique price vector $q(\mu, F) \in \mathbb{R}^n$.

6. The Min-Max Theorem

We are now in a position to present our main characterization theorem which relates GA-structures to minimum price competitive equilibria. The proofs of all of the Theorems and Lemmas are contained in the Appendices.

Theorem 1: $q(\mu^*, F^*)$ is a minimum equilibrium price if and only if:

$$\sum_{i \in \{1..n\}} q_i(\mu^*, F^*) = \min_{\mu \in M} \max_{F \in \mathcal{F}_\mu} \sum_{i \in \{1..n\}} q_i(\mu, F) \quad (6.1)$$

where $\mathcal{F}_\mu = \{F \in \mathcal{F} | (\mu, F) \in \mathcal{G}\}$.

We establish this result through a series of lemmas. To prove that the minimum equilibrium price vector solves (6.1), we first show that for any minimum price competitive equilibrium (μ^*, p^*) there exists a GA-structure $(\mu^*, F^*) \in \mathcal{G}$ with $q(\mu^*, F^*) = p^*$. Next we show that altering the graph only lowers the implied price, $q_i(\mu^*, F^*) \geq q_i(\mu^*, F)$ for all $(\mu^*, F) \in \mathcal{G}$. If this were not the case, (μ^*, p^*) could not be a competitive equilibrium, since there would be some buyer willing to bid more than p^* for a good that they are not allocated under μ^* . Finally, we show that if μ is not associated with a competitive equilibrium then there exists some F such that $q_i(\mu, F) \geq q_i(\mu^*, F^*)$. Again the intuition is that allocating a buyer a good that is not in their demand set increases their willingness to pay for other goods. The converse follows from the fact that we know from Demange and Gale (1985) that there exists a unique minimum equilibrium price. It is shown that this implies that any solution to (6.1) is a competitive equilibrium. Along the way, we prove an alternate version of Theorem 1.

Corollary 1 $q(\mu^*, F^*)$ is a minimum equilibrium price if and only if:

$$q_i(\mu^*, F^*) = \min_{\mu \in M} Q_i(\mu) \text{ for all } i \quad (6.2)$$

where $Q_i(\mu) = \max_{F \in \mathcal{F}_\mu} q_i(\mu, F)$.

The difference between the two formulations is that in Theorem 1 we choose an allocation to minimize the sum of the components of $q(\mu, F)$, whereas in Corollary 1 we minimize each component individually. Corollary 1 also allows for the maximization over graphs to take place component by component. The advantage of Theorem 1 is its simplicity. The advantage of Corollary 1 is that it shows that the equilibrium allocation not only minimizes the sum of prices but each price individually. This will prove useful when discussing comparative statics below.

Most GA-structures generate prices and allocations that are inconsistent with optimization by buyers or sellers. Some generate prices that lie below sellers' reservation prices; others allocate goods to buyers who would prefer to purchase other goods. Buyer and seller optimality are enforced through the maximization and minimization. On the sellers' side, maximizing over F guarantees that all prices are above sellers' reservation, since we can always choose F such that a given good is part of the root set. Minimizing over μ guarantees that all goods in $U(p)$ are potentially allocated. For example, suppose that for a given allocation maximizing over F leads to a situation in which an unallocated good is priced above reservation. This can only happen if that good is in the demand set of a buyer allocated to another good. Often reallocating that buyer to the unallocated good solves the problem. Note that this also tends to reduce the price vector by removing the indifference that was driving up the price of the unallocated good in the first place. On the buyers' side, given the equilibrium allocation, maximizing over F guarantees that no buyer prefers any good to the good that they are allocated. Minimizing over μ avoids raising prices through misallocations. Of course, the above intuitive arguments are incomplete, and the proof itself is as a result somewhat intricate.

Many comparative static results from the literature follow from Theorem 1. Demange and Gale (1985) show that minimal equilibrium prices are weakly increasing in seller reservation, that increasing the number of sellers does not raise prices, and that increasing the number buyers does not lower prices. In our framework, an increase in reservation prices can only raise $q(\mu, F)$; an increase in the number of sellers is equivalent to an expansion in the set of potential matches; and reducing the number of buyers is equivalent to a restriction on the set of graphs, namely the restriction that one buyer be allocated to a null tree.

7. Competitive Equilibrium Allocations

Theorem 1 concerns the price vector. Allocations are more complicated. The arguments used to prove Theorem 1 establish that if μ^* is a minimum price competitive equilibrium allocation

then there is exists a GA-structure involving μ^* which solves (6.1). The converse, however, is not true. There exist GA-structures that solve (6.1) that do not involve competitive equilibrium allocations.

The following example illustrates such a situation. The example involves a good that is unallocated and priced above reservation. Normally such a GA-structure would not solve the min-max problem. Reallocating the buyer assigned to the unallocated good's direct predecessor to the unallocated good would lower prices, since the unallocated good would lose the indifference supporting its high price. In the example, however, there are multiple buyers interested in the unallocated good. When one buyer is reallocated, the others' indifference continues to support the good's high price.

Example: There are three goods and two buyers. The minimum equilibrium price has goods y_1 and y_2 priced at reservation and good y_3 is priced above reservations. At these prices x_a is indifferent between y_1 and y_3 and x_b is indifferent between y_2 and y_3 . There are two competitive equilibrium allocations: either x_a is allocated to y_1 and x_b is allocated to y_3 or x_a is allocated to y_3 and x_b is allocated to y_2 . The key point is that in any competitive equilibrium y_3 must be allocated since it is priced above reservation. It is not necessary, however, that y_3 be allocated for a GA-structure to price it. The GA-structure with x_a allocated to y_1 and x_b allocated to y_2 , together with a graph that includes the edge (y_1, y_3) generates the minimum equilibrium price vector and hence solves (6.1). This GA-structure is illustrated in Figure 2. If x_b were not indifferent to y_3 , then reallocating x_a to y_3 would lower prices.

[Figure 2]

The property of competitive equilibrium that fails in the example is that the allocation is not onto the set $U(p)$. It turns out that if an allocation solves the min-max problem and is onto $U(p)$, then the allocation is a competitive allocation.

Theorem 2: If (p^*, μ^*) is a minimum price competitive equilibrium then,

$$\mu^* \in \arg \min H(\mu), \tag{7.1}$$

where $H(\mu) = \max_{F \in \mathcal{F}_\mu} \sum_{i \in \{1 \dots n\}} q_i(\mu, F)$. Moreover, if $\mu \in \arg \min H(\mu)$ and for all μ is onto $U(p^*) = \{y_i \in Y | p_i^* > r_i\}$, then (p^*, μ) is a minimum price competitive equilibrium.

The theorem follows from the observation that given any μ , if there exists a buyer x_a who strictly prefers some good y_i to μ_a at the minimum equilibrium price vector p^* , then there exists

a graph F such that $p_i^* \leq q_i(\mu, F)$ with strict inequality for y_i . Hence any allocation that solves the min-max problem satisfies buyer optimality. If it is also onto $U(p^*)$ then it satisfies the other conditions for a competitive equilibrium as well.

Although the min-max problem does not pin down the equilibrium allocation, it is easy to construct an equilibrium allocation given any solution to the min-max problem. Since the solution to the min-max problem gives the equilibrium price vector, the values of all goods are known. The problem becomes one of finding for fixed payoffs an allocation that both maximizes buyers' utility and is onto the set $U(p^*)$. To construct such an allocation, one begins with any allocation that solves the min-max problem. One then identifies an unallocated good whose price is above reservation. Since the allocation solves the min-max problem, there exists a chain of indifference extending from the unallocated good to the root set. Next, shift each buyer allocated to a good in that chain to its immediate successor in the chain. This operation leads to another allocation that solves the min-max problem and reduces the set of unallocated goods by one. Caplin and Leahy (2010) discuss such reallocations in greater detail.

8. Local Comparative Statics

In this section we consider a collection of models indexed by the parameter $\lambda \in \Lambda$, where Λ is an open set in \mathbb{R}^T for some constant T . Depending on the application, λ may parametrize a shift in the reservation prices of sellers, $r(\lambda)$, and/or it may reflect some aspect of buyers' utility, $U_a(y_i, p_i; \lambda)$. Hence given any GA-structure (μ, F) , the price vector will be $q(\mu, F; \lambda)$. Let $\Phi : \Lambda \rightarrow \mathcal{G}$ denote the mapping from parameters to the set of GA-structures that generate minimum price competitive equilibrium:

$$\Phi(\lambda) = \{(\mu, F) \mid \mu \in \arg \min H(\mu) \text{ and } F \in \arg \max_{F \in F_\mu} q(\mu, F; \lambda)\}$$

and let $p : \Lambda \rightarrow \mathbb{R}^n$ denote the minimum equilibrium price vector given λ .

If we assume that $q(\mu, F; \lambda)$ is continuous in λ for all $(\mu, F) \in \mathcal{G}$, it follows directly from the Theorem of the Maximum applied to (6.2) that $\Phi(\lambda)$ is upper-hemicontinuous and the minimum price competitive equilibrium price vector is continuous.

Theorem 3 If $q(\mu, F; \lambda)$ is continuous in λ for all $(\mu, F) \in \mathcal{G}$, then all of the components of $p(\lambda)$ are continuous and $\Phi(\lambda)$ is non-empty, compact-valued, and upper-hemicontinuous at $\lambda \in \Lambda$.

The picture that emerges is one of a finite collection of surfaces $q(\mu, F; \lambda)$ in \mathbb{R}^n , one surface for each (μ, F) . Given λ , the minimum competitive equilibrium price vector is associated with one of these surfaces. As we alter λ , we move along one surface until it intersects with another and the min-max problem may tell us to switch and follow the other surface. The characterization of these switches is the subject of Caplin and Leahy (2010). The point that we want to make here is that so long as these intersections are not too frequent, local comparative statics will almost everywhere involve a fixed GA-structure.

A natural smoothness assumption that leads to sparse switches is to assume that the $q(\mu, F; \lambda)$ are analytic functions of λ .¹² This assumption would be satisfied in almost any practical application of the model, as it only requires that the utility functions be analytic functions of p and λ and that the reservation prices be analytic functions of λ .¹³

Assumption B $q(\mu, F; \lambda)$ is an analytic function of λ for all $(\mu, F) \in \mathcal{G}$

If the $q(\mu, F; \lambda)$ are analytic, then Lojasiewicz's Structure Theorem for Real Varieties (Krantz and Parks [2002], p. 168) implies that the set of switch points is at most dimension $T - 1$. It follows that at almost every point in the parameter space λ_0 there will be a neighborhood $N(\lambda_0)$ such that $(\mu, F) \in \Phi(\lambda_0)$ implies $(\mu, F) \in \Phi(\lambda)$ for $\lambda \in N(\lambda_0)$.¹⁴

We use this insight to discuss the local effects of various parameter changes. First, suppose that there is an increase in the reservation price of a good y_0 . In this case $\lambda \equiv r_0$. Suppose that $(\mu, F) \in \Phi(\lambda_0)$. It is immediate from the construction of $q(\mu, F)$ that for almost all λ , a change in r_0 impacts $q(\mu, F)$ only if y_0 is part of the root set, and even then the effect is limited to the component of F containing y_0 . If y_0 is not an element of the root set then r_0 is almost surely inframarginal in the sense that it is strictly below the minimum equilibrium price. If y_0 is part of the root set, then the effects on the successors of y_0 work through the graph F . An increase in r_0 has a direct effect on the price of y_0 which then affects the willingness of the buyer allocated to y_0 to pay for other goods. This alters the price of the direct successors of y_0 , and by induction their successors. Prices of goods in other components are not affected by the change in r_0 since they are not connected in any way to y_0 .

Proposition 1 summarizes the effect of a change in the reservation price of y_0 .

¹²A function $f(x)$ is analytic at a point x_0 if its Taylor series expansion converges on a neighborhood of x_0 .

¹³See Frantz and Parks [2002].

¹⁴It is possible that multiple allocations support the competitive equilibrium for $\lambda \in N(\lambda_0)$. In this case the equilibrium prices will be unique but the equilibrium allocation will be indeterminate. Caplin and Leahy (2010) discuss conditions under which $\Phi(\lambda_0)$ is generically unique.

Proposition 1 Suppose $\lambda \equiv r_0$ for some $y_0 \in Y$. Suppose further that (μ, F) solves (6.1) at $\lambda \in \Lambda$ and that Assumption B holds. For almost all $\lambda \in \Lambda$ the following are true:

1. If $y_0 \notin R(F)$, then $dp_k/d\lambda = 0$ for all $y_k \in Y$.
2. If y_k is not a successor of y_0 , then $dp_k/d\lambda = 0$ for all $y_k \in Y$.
3. If y_k is a successor of y_0 and $\{(y_0, y_1), (y_1, y_2), \dots, (y_n, y_k)\} \subset E(F)$, then

$$\frac{dp_k}{d\lambda} = \frac{dp_k}{dp_{k-1}} \frac{dp_{k-1}}{dp_{k-2}} \dots \frac{dp_1}{dp_0} \quad (8.1)$$

where

$$\frac{dp_i}{dp_{i-1}} = \frac{dU_{\sigma_{i-1}}(y_{i-1}, p_{i-1})}{dp_{i-1}} \bigg/ \frac{dU_{\sigma_{i-1}}(y_i, p_i)}{dp_i}$$

Equation (8.1) is a type of “chain rule” for local comparative statics in allocation markets. This chain rule works along the chain of indifference connecting p_k to the root set. The movement in each price in the chain affects the price of its direct successor.

A shock to the utility of a buyer is slightly more complicated. There is no equivalent to the buyer being inframarginal, since we are assuming all buyers are allocated. The effect of a shock to the utility of buyer x_0 is limited to the prices of all goods that are successors of μ_0 . Note that the price of μ_0 does not change since x_0 's utility is used to price the direct successors of μ_0 given the price of μ_0 .

Proposition 2 Suppose λ shifts the utility of buyer x_0 who is assigned to y by μ . Suppose further that (μ, F) solves (6.1) at $\lambda \in \Lambda$ and that Assumption B holds. For almost all $\lambda \in \Lambda$ the following are true:

1. If y_k is not a successor of y , then $dp_k/d\lambda = 0$.
2. If y_k is a successor of y_0 and $\{(y_0, y_1), (y_1, y_2), \dots, (y_n, y_i)\} \subset E(F)$ is a path in F from y_0 to y_k then

$$\frac{dp_k}{d\lambda} = \frac{dp_k}{dp_{k-1}} \frac{dp_{k-1}}{dp_{k-2}} \dots \frac{dp_1}{d\lambda}$$

where

$$\frac{dp_i}{dp_{i-1}} = \frac{dU_{\sigma_{i-1}}(y_{i-1}, p_{i-1})}{dp_{i-1}} \bigg/ \frac{dU_{\sigma_{i-1}}(y_i, p_i)}{dp_i}$$

and

$$\frac{d\underline{p}_1}{d\lambda} = \left(\frac{dU_0(y_0, p_0)}{d\lambda} - \frac{dU_0(y_1, p_1)}{d\lambda} \right) \bigg/ \frac{dU_0(y_1, p_1)}{d\underline{p}_1}$$

The picture that emerges from these two propositions contrasts with the case of divisible goods. With divisible goods every good is connected to every other good through indifference. Even the smallest change to the supply or demand for one good tends to affect the price of every other good in the economy. With indivisibility, small changes in the supply or demand for a good, only affect the prices of that good and its successors. With indivisible goods small discrete shocks have local effects, whereas with divisible goods, infinitesimal shocks have global effects.¹⁵

Global comparative statics with indivisible goods are more complicated. Larger shocks induce changes in the equilibrium GA-structures. Caplin and Leahy (2010) show how to build up the effects of large shocks from the effects of small shocks. They show that transitions among GA-structures are orderly. Along the generic path in parameter space only certain $q(\mu, F; \lambda)$ surfaces intersect. Therefore only certain changes in the graph and allocation ever need to be considered.

9. The Dual Problem and the Equilibrium Set

9.1. Exit and the Primal Problem

In order to simplify the analysis, we have assumed to this point that all buyers make a purchase from Y . In many applications buyers may have outside options. Moreover to model the dual problem, we will need a notion of maximal willingness to pay.¹⁶

We introduce exit by expanding the choice sets of buyers and sellers. Let y_a denote the outside option of buyer x_a . We assume that buyer x_a chooses from the set $Y \cup y_a$. No other buyer demands y_a . Let v_a denote the value of the outside option to buyer x_a . We normalize the reservation prices of y_a to zero. Let $\bar{Y} = Y \cup \{y_a\}_{a=1}^m$.

Similarly let x_i denote the outside option of seller y_i and $\bar{X} = X \cup \{x_i\}_{i=1}^n$. This preserves the one-to-one nature of the allocation. With these amendments there are $m + n$ buyers and sellers. We normalize the utility of the phantom buyers to zero. Demange and Gale prove that

¹⁵We thank Victor Norman for this observation.

¹⁶We introduce exit without introducing budget constraints. Budget constraints introduce discontinuities in utility at the point resources are exhausted. Minimum price competitive equilibria may not exist. To eliminate these problems it is often assumed that agent exit before resources are exhausted. See Kaneko (1982). This assumption would be justified by any model in which the Inada conditions held.

there exists a minimum price competitive equilibrium in this model, which we now refer to as the “primal model.”

The only impact that introducing exit has is to restrict the set of admissible GA-structures. The set of admissible allocations is reduced in two ways. An allocation must not assign a buyer to another buyer’s outside option, nor can an allocation assign one seller’s outside option to another seller. The set of admissible graphs is also reduced. The outside options of buyers must be root goods, since no other buyer can demand them. Appropriately amended versions of Theorem 1, Corollary 1 and Theorem 2 follow immediately.

9.2. The Dual

In the dual one switches the positions of buyers and sellers and reinterprets equilibrium as taking place in the space of buyer utilities rather than in the space of seller prices. In the dual buyers’ utility plays the role that prices play in the primal. The utility of buyers’ outside options plays the role of seller’s reservation prices. Let $v \in \mathbb{R}_+^{m+n}$ denote buyer utility. Let v_a^R denote the utility x_a receives from exit. Sellers can only make a sale to x_a if they offer utility greater than v_a^R .

Sellers choose buyers to maximize the price that they receive. The payoff of seller y_i , $p_i(x_a, v_a)$, depends on the buyer that they sell to x_a and v_a , the utility received by x_a . In the case of a null buyer

$$p_i(x_i, 0) \equiv r_i.$$

In the case of a non-null buyer, $p_i(x_a, v_a)$ is defined as the price that would have to be charged for good y_i to provide buyer x_a with utility v_a ,

$$U_a [y_i, p_i(x_a, v_a)] = v_a.$$

Given $v_a \geq v_a^R$, this solution exists and is unique due to strict monotonicity and continuity of the utility function. The supply correspondence $S_i(v)$ includes those buyers who generate maximum values for this “indirect profit function” $p_i(x_a, v_a)$.

An allocation of goods is a one-to-one mapping $\sigma : \bar{Y} \rightarrow \bar{X}$ such that each good is assigned a feasible buyer,

$$\sigma_i = \sigma(y_i) \in X \cup x_i.$$

σ is the inverse of μ defined on the extended set of goods and buyers. We let M^{-1} denote the set of such allocations.

Definition 9.1. A competitive equilibrium in the dual model is a pair $(\hat{v}, \hat{\sigma})$ such that:

1. $\hat{\sigma}_i \in S_i(\hat{v})$ for all $y_i \in Y$.
2. $v_a \geq \hat{v}_a^R$
3. If $\hat{v}_a > \hat{v}_a^R$, then there exists $y_i \in Y$ such that $\hat{\sigma}_i = x_a$.

9.3. Maximum Price Equilibria

We construct maximum price competitive equilibrium using the dual of the GA-structure. This is an allocation $\sigma \in M^{-1}$ and a graph $T \in \mathcal{T}$ on the set of buyers \bar{X} . The graphs T are of the form (\bar{X}, R, E) and satisfy all of the conditions of the class \mathcal{F} with the additional restriction that all of the null buyers satisfy $x_i \in R$. The allocations allocate sellers to buyers, respecting the restriction that no seller can be allocated to the outside option of another seller. As in the case of minimal price competitive equilibria we need to limit attention to cases in which if $(x_a, x_b) \in E$, then σ allocates a seller to x_a . We denote the class of admissible (σ, T) pairs \mathcal{H} , and the class of admissible graphs given an allocation σ , $\mathcal{T}(\sigma)$.

We construct vectors of buyer utilities from $(\sigma, T) \in \mathcal{H}$ in the same way that we constructed prices from (μ, F) . We set the utilities of all buyers $x_a \in R$ equal to their reservation value v_a^R . We then proceed by induction using the indifference of sellers to assign utilities to the direct successors of buyers whose utility we already know. If $(x_a, x_b) \in E$ and $\sigma_i = x_a$, then

$$p_i(x_a, v_a) = p_i(x_b, v_b)$$

gives v_b as a function of v_a .

The maximal price competitive equilibrium is characterized by maximizing the sum of buyer utility over admissible graphs T and then minimizing over allocations σ .

$$\min_{\sigma \in M^{-1}} \max_{T \in \mathcal{T}(\sigma)} \sum_{a=1}^m v_a$$

We can back out the equilibrium price of any good y_i from any solution to this maximization problem. Let $x_a = \sigma_i$, then p_i solves,

$$U_a(y_i, p_i) = v_a.$$

If $\sigma_i = x_i$, then $p_i = r_i$.

9.4. The Equilibrium Set

Having solved the original model to identify the minimum equilibrium price, and the dual to identify the maximum equilibrium price, one can characterize the set of competitive equilibria. Every competitive equilibrium is associated with a set of reservation prices that lie between the identified minimum and maximum equilibrium prices. The formal statement is in Theorem 4.

Theorem 4 A price vector p is a competitive equilibrium price vector if and only if it is the minimum price competitive equilibrium price vector for a model with reservation prices $\hat{r} \in [p, \bar{p}]$.

10. Concluding Remarks

In this paper and its companion, we present a new mathematical apparatus for understanding allocation markets with NTU. We are currently extending the work to a dynamic context and solving for the reallocation of objects over time. Buyers may become sellers or agents may act simultaneously as buyers and sellers.

The housing market is particularly promising in terms of applications. With regard to theory, many questions concerning housing markets require the introduction of trading frictions. In housing markets only a small fraction of homes are traded in any given period of time. What do minimum price equilibria look like in this case? What influence do non-traded homes have on current transactions? With regard to empirical implementation, to what extent do prices reflect local income and to what extent local amenities? How do shocks to one location such as the location of a factory or school propagate through space and time? To what extent does the revealed pattern of movements over the housing life cycle connect housing prices and housing returns in geographically disconnected areas? Other applications, e.g. in auction markets, are also of interest.

References

- [1] Alonso, William (1964), *Location and Land Use*, Cambridge: Harvard University Press.
- [2] Caplin, Andrew, and John Leahy (2010), “Comparative Statics in Markets for Indivisible Goods,” NYU working paper.
- [3] Crawford, Vincent, and Elsie Knoer (1981), “Job Matching with Heterogeneous Firms and Workers,” *Econometrica* 49, 437-50.
- [4] Demange, Gabrielle, and David Gale (1985), “The Strategy Structure of Two-Sided Matching Markets,” *Econometrica* 53, 873-888.
- [5] Demange, Gabrielle, and David Gale, and Marilda Sotomayor (1986), “Multi-Item Auctions,” *The Journal of Political Economy* 94, 863-873.
- [6] Gale, David, (1984), “Equilibrium in a Discrete Economy with Money.” *International Journal of Game Theory* 13, 61–64.
- [7] Kaneko, Mamoru, (1982), “The Central Assignment Game and the Assignment Markets,” *Journal of Mathematical Economics* 10, 205–232.
- [8] Kaneko, Mamoru, Tamon Ito, and Yu-ichi Osawa (2006), “Duality in Comparative Statics in Rental housing markets with Indivisibilities,” *Journal of Urban Economics* 59, 142-170.
- [9] Kaneko, Mamoru, and Yoshitsugu Yamamoto (1986), “The Existence and Computation of Competitive Equilibria in Markets with an Indivisible Commodity”, *Journal of Economic Theory* 38, 118-136
- [10] Koopmans, Tjalling and Martin Beckmann (1957), “Assignment Problems and the Location of Economic Activities”, *Econometrica* 25, 53-76.
- [11] Krantz, Steven, and Harold Parks (2002), “A Primer of Real Analytic Functions,” Boston: Birhauser.
- [12] Kuhn, Harold (1955), “The Hungarian Method for the Assignment Problem,” *Naval Research Logistics Quarterly* 2, 83-97.
- [13] Meese, Richard, and Nancy Wallace (1994), “Testing the Present Value Relation for Housing Prices: Should I Leave my House in San Francisco,” *Journal of Urban Economics*, Volume 35, 245-266.

- [14] Miyake, Mitsunobu (1994), “Comparative Statics of Assignment Markets with General Utilities,” *Journal of Mathematical Economics* 23, 519-531.
- [15] Miyake, Mitsunobu (2003), “Precise Computation of a Competitive Equilibrium of the Discrete Land Market Model,” *Regional Science and Urban Economics* 33, 721-743.
- [16] Munkres, James (1957), “Algorithms for the Assignment and Transportation Problems,” *Journal of the Society of Industrial and Applied Mathematics* 5, 32-38.
- [17] Ricardo, David, (1817), *Principles of Political Economy and Taxation*, London, John Murray.
- [18] Roback, Jennifer (1982), “Wages, Rents, and the Quality of Life,” *Journal of Political Economy* 90, 1257-1278.
- [19] Rosen, Sherwin (1974), “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition,” *Journal of Political Economy* 92, 34-55.
- [20] Scarf, Herbert E. (1994), “The Allocation of Resources in the Presence of Indivisibilities,” *Journal of Economic Perspectives*, Vol. 8, 111-128.
- [21] Shapley, Lloyd, and Martin Shubik (1972), “The Assignment Game I: The Core.” *International Journal of Game Theory* 1, 111-130.
- [22] Quinzii, Martine, (1984), “Core and Competitive Equilibria with Indivisibilities,” *International Journal of Game Theory* 13, 41–60.

11. Appendix

The following lemmas are used in the proof of Theorem 1.

Lemma 1: Given a minimum price competitive equilibrium (μ^*, p^*) , there exists $F^* \in \mathcal{F}$ such that $(\mu^*, F^*) \in \mathcal{G}$ and $p^* = q(\mu^*, F^*)$.

Proof: Consider a minimum price competitive equilibrium (μ^*, p^*) . We construct $F \in \mathcal{F}$ such that $(\mu^*, F) \in \mathcal{G}$ and $p^* = q(\mu^*, F)$.

The first stage in the construction of graph F^* on Y is to identify the root set as all goods that are at reservation prices,

$$R^* = \{y_k \in Y \mid p_k^* = r_k\}.$$

The graph is completed by induction. Let $A_1 = R^*$ and let F_1 denote the null graph on the vertex set A_1 . At stage $s > 1$ of the construction, suppose we have identified $A_s \subset Y$ and a graph $F_s = (A_s, R^*, E_s)$ on the vertex set A_s such that F_s is a forest of rooted trees with root set R^* and all of the edges in E_s are directed away from the roots. By construction, $R^* \subseteq A_s$ and $Y \setminus A_s \subset U(p^*)$. Lemma 4 in Demange and Gale (1985) states that there exists $x_a \in X$ such that $\mu_a^* \in A_s$ and $D_a(p^*) \cap Y \setminus A_s \neq \emptyset$. Choose $y_i \in D_a(p^*) \cap Y \setminus A_s$. Let $A_{s+1} = A_s \cup \{y_i\}$, $E_{s+1} = E_s \cup (\mu_a^*, y_i)$, and $F_{s+1} = (A_{s+1}, R^*, E_{s+1})$. By construction, F_{s+1} is a forest of rooted trees with root set R^* and all edges directed away from the roots. Given that there are a finite number of elements in $Y \setminus R^*$, this construction converges in a finite number of steps to a graph $F^* = (Y, R^*, E^*)$.

To see that $(\mu^*, F^*) \in \mathcal{G}$, note that the construction implies that $(y_i, y_k) \in E^*$ only if there exists $x_a \in X$ such that $\mu_a^* = y_i$. To see that $q(\mu^*, F^*) = p^*$, note first that, by construction, all goods in R^* are at reservation prices. Furthermore note that for any edge $(y_i, y_k) \in E^*$, the buyer $x_a \in X$ with $\mu_a^* = y_i$ is indifferent between y_i and y_k at p^* . In light of Assumption A, the fact that all implied indifferences hold at p^* is sufficient to complete the demonstration that $q(\mu^*, F) = p^*$. \square

Lemma 2: For any minimum price competitive equilibrium (μ^*, p^*) ,

$$p_i^* = \max_{F \in \mathcal{F}_\mu} q_i(\mu^*, F)$$

Proof: The proof is by contradiction. Lemma 1 states that if (μ^*, p^*) is a minimum price competitive equilibrium then there exists $(\mu^*, F^*) \in \mathcal{G}$ such that $q(\mu^*, F^*) = p^*$. Suppose that

$q(\mu^*, F^*) = p^*$, and consider any $F \neq F^*$ such that $q_1(\mu^*, F) > p_1^*$ for some good $y_1 \in Y$. $p_1 > p_1^*$ implies $y_1 \notin R(F)$. Let $y_r \in R(F)$ denote the root good in the component of F containing y_1 . y_r is a predecessor of y_1 in F . Since $y_r \in R(F)$, $q_r(\mu^*, F) = r_r \leq q_r(\mu^*, F^*)$.

Consider the path $\{y_r, \dots, y_1\}$ in F . Let y_2 denote the first predecessor of y_1 along this path such that $q_2(\mu^*, F) \leq p_2^*$. As $q_r(\mu^*, F) \leq q_r(\mu^*, F^*)$ and $q_1(\mu^*, F) > q_1(\mu^*, F^*)$, y_2 is well defined. Let y_3 denote the direct successor of y_2 along this path. Consider x_2 such that $\mu_2^* = y_2$. $(y_2, y_3) \in F$, implies that $U(y_2, q_2(\mu^*, F)) = U(y_3, q_3(\mu^*, F))$. Since x_2 is indifferent between y_2 and y_3 at $q(\mu^*, F)$, the fact that $q_2(\mu^*, F) \leq p_2^*$ and $q_3(\mu^*, F) > p_3^*$ implies that x_2 strictly prefers y_3 to y_2 at the price vector p^* . But (μ^*, p^*) is a minimum price competitive equilibrium. This contradiction establishes the lemma. \square

Lemma 3: Given a minimum price competitive equilibrium (μ^*, p^*) and any allocation μ , then there exists $F \in \mathcal{F}$ such that $(\mu, F) \in \mathcal{G}$ and $p_i^* \leq q_i(\mu, F)$ for all i . If there exists a buyer x_a who strictly prefers some good y_i to μ_a at p^* , then the inequality is strict for good y_i .

Proof: Let (μ^*, p^*) be a minimum price competitive equilibrium and μ any allocation. We show how to construct an F such that $q(\mu, F) \geq p^*$. The lemma follows immediately.

By Lemma 1, there exists, F^* such that $q(\mu^*, F^*) = p^*$. Note that we may pick F^* such that $R(F^*)$ contains all y_i such that $p_i^* = r_i$ by simply removing all of the edges (y_k, y_i) such that $p_i^* = r_i$. Given this choice of F^* , all goods that are not root goods must be allocated by μ^* , that is for all $y_i \notin R(F^*)$ there exists x_b such that $\mu_b^* = y_i$. Otherwise (μ^*, p^*) would not be a competitive equilibrium.

To construct F , we first construct the directed graph K as follows.

(1) For each edge $(y_i, y_k) \in E(F^*)$ find $x_b \in X$ such that $\mu_b^* = y_i$. Then, if $\mu_b \neq y_k$, include (μ_b, y_k) in $E(K)$. Intuitively, every good is being priced by the same person in K as in F^* .

(2) If $\mu_c^* \neq \mu_c$, include (μ_c, μ_c^*) in $E(K)$.

K may not be a tree. It is possible that the second step creates a vertex with a indegree of two. To construct F we will make a selection from K .

We construct F by induction. Let $R(F) = R(F^*)$. Delete all edges $(y_i, y_k) \in K$ in which $y_k \in R(F)$. Note also that $q_i(\mu, F) = r_i = p_i^*$ for all $y_i \in R(F)$.

Now suppose that $p_i^* \leq q(\mu, F)$ all $y_i \in A_s \subset Y$. Since $p_j^* > r_j$ for all $y_j \in Y \setminus R(F)$, there exists x_d and y_1 such that $\mu_d^* \in A_s$, $(\mu_d^*, y_1) \in E(F^*)$, and $y_1 \in Y \setminus A_s$. We consider two cases. First, if $\mu_d \in A_s$, then $(\mu_d, y_1) \in K$ by rule (1) above, and we add (μ_d, y_1) to F . The second case is $\mu_d \in Y \setminus A_s$. Now for each $y_k \in Y \setminus A_s$, there exists x_b such that $\mu_b^* = y_k$. Hence if

$\mu_d^* \notin Y \setminus A_s$ and $\mu_d \in Y \setminus A_s$, there exists x_e such that $\mu_e^* \in Y \setminus A_s$, but $\mu_e \notin Y \setminus A_s$. By rule (2), $(\mu_e, \mu_e^*) \in E(K)$ and we add it to $E(F)$.

Let (y_j, y_k) denote the edge that we have added to F at this stage and suppose that $\mu_a = y_j$. q_k is determined by

$$U_a(y_k, q_k) = U_a(y_j, q_j).$$

But $y_j \in A_s$ implies $q_j \geq p_j^*$

$$U_a(y_j, q_j) \leq U_a(y_j, p_j^*)$$

and the definition of competitive equilibrium implies

$$U_a(y_j, p_j^*) \leq U_a(\mu_a^*, p_{\mu_a^*}^*) \tag{11.1}$$

Finally by construction

$$U_a(\mu_j^*, p_{\mu_a^*}^*) = U_a(y_k, p_k^*)$$

Note the last step is redundant in the case of rule (2) as $\mu_a^* = y_k$. It follows from the monotonicity of U_a that $q_k \geq p_k^*$. This completes the induction step.

If μ is not a competitive equilibrium allocation then some buyer strictly prefers μ^* to μ . (11.1) becomes a strict inequality and $q_k \geq p_k^*$ for all y_k with strict inequality for at least one y_k . It follows that in this case

$$\sum_{i \in \{1 \dots n\}} q_i(\mu, F) > \min_{\mu \in M} \max_{F \in \mathcal{F}_\mu} \sum_{i \in \{1 \dots n\}} q_i(\mu, F).$$

This completes the proof of the lemma. \square

Proof of Theorem 1: We first show that the minimum competitive equilibrium price is a solution to (6.1).

Suppose that (μ^*, p^*) is a minimum price competitive equilibrium and that $q(\hat{\mu}, \hat{F}) = p^*$. By Lemma 1, there exists F^* such that

$$q(\mu^*, F^*) = p^* = q(\hat{\mu}, \hat{F}).$$

Lemma 2 implies

$$q_i(\mu^*, F^*) = \max_{F \in \mathcal{F}_\mu} q_i(\mu^*, F)$$

It follows immediately that

$$\sum_{i \in \{1 \dots n\}} q_i(\mu^*, F^*) = \sum_{i \in \{1 \dots n\}} \max_{F \in \mathcal{F}_\mu} q_i(\mu^*, F)$$

Lemma 3 implies

$$\sum_{i \in \{1 \dots n\}} q_i(\mu^*, F^*) \leq \max_{F \in \mathcal{F}_\mu} \sum_{i \in \{1 \dots n\}} q_i(\mu, F)$$

It follows that

$$\sum_{i \in \{1 \dots n\}} q_i(\hat{\mu}, \hat{F}) = \sum_{i \in \{1 \dots n\}} q_i(\mu^*, F^*) = \min_{\mu \in M} \max_{F \in \mathcal{F}_\mu} \sum_{i \in \{1 \dots n\}} q_i(\mu, F).$$

To establish the converse, we need only show that the solution to (6.1) is unique, for then the preceding arguments establish equivalence. Suppose that there is a solution $(\hat{\mu}, \hat{F})$ to min-max that is not a competitive equilibrium. This could happen in one of three ways: either some price is below reservation; some good whose price is above reservation is unallocated; or some buyer is allocated to a good that is not in his or her demand correspondence. We discuss each case in turn.

In the first case it is clear that $(\hat{\mu}, \hat{F})$ does not solve $\max_{F \in \mathcal{F}_\mu} \sum_{i \in \{1 \dots n\}} q_i(\mu, F)$. We can raise the price of any good to its reservation value by adding it to the root set.

Second, suppose that there exists y_1 such that $q_1(\hat{\mu}, \hat{F}) > r_1$ and there exists no $x_a \in X$ such that $\hat{\mu}_a = y_1$. Since $q_1(\hat{\mu}, \hat{F}) > r_1$, y_1 has a direct predecessor in \hat{F} , call it y_2 , and there exists x_b such that $\hat{\mu}_b = y_2$ and x_a is indifferent between y_1 and y_2 .

Consider μ' such that $\mu'_a = \mu_a$ for all $x_a \neq x_b$, $\mu'_b = y_1$, and y_2 is unallocated.

Suppose that there exists y_i and F' such that $q_i(\mu', F') > q_i(\hat{\mu}, \hat{F})$. Let y_r denote the root good associated with the component of F' containing y_i . Consider the path $\{y_r, \dots, y_i\}$ in F' . Note that $q_r(\mu', F') = r_r \leq q_r(\hat{\mu}, \hat{F})$. Let y_3 denote the good closest to y_r on this path such that $q_3(\mu', F') > q_3(\hat{\mu}, \hat{F})$ and let y'_4 denote the immediate predecessor of y_3 in F' and \hat{y}_4 the immediate predecessor in \hat{F} . Let $\hat{\sigma}$ and σ' denote the inverses of $\hat{\mu}$ and μ' . If $\hat{\sigma}(y'_4) = \sigma'(y'_4)$, then $y'_4 \neq \hat{y}_4$, otherwise the edge (y'_4, y_3) is in $E(\hat{F})$ and $q_4(\mu', F') \leq q_4(\hat{\mu}, \hat{F})$ implies $q_3(\mu', F') \leq q_3(\hat{\mu}, \hat{F})$. But in this case if $y'_4 \neq \hat{y}_4$ we can replace the edge $(\hat{y}_4, y_3) \in E(\hat{F})$ with the edge (y'_4, y_3) thereby raising the price of y_3 and its successors in \hat{F} without reducing the price of any other good. This contradicts the assumption that $\sum_{i \in \{1 \dots n\}} q_i(\hat{\mu}, \hat{F}) = \max_{F \in \mathcal{F}_\mu} \sum_{i \in \{1 \dots n\}} q_i(\mu, F)$. It follows that $\hat{\sigma}(y'_4) \neq \sigma'(y'_4)$. This implies that $y_4 = y_1$ since this is the only good that is allocated to a different buyer under μ' . In this case we can replace $(\hat{y}_4, y_3) \in E(\hat{F})$ with the edge (y_2, y_3) .

Since there $q_1(\mu', F') \leq q_1(\hat{\mu}, \hat{F})$ and x_a is indifferent between y_1 and y_2 at $q_1(\hat{\mu}, \hat{F})$, this change raises the price of y_3 and its successors without altering any other price. This again leads to a contradiction. As all cases are exhausted, it follows that $q_i(\mu', F') \leq q_i(\hat{\mu}, \hat{F})$.

Since $\hat{\mu}$ minimizes the sum of prices it follows that $q_i(\mu', F') = q_i(\hat{\mu}, \hat{F})$. Since $q_1(\mu', F') > r_1$, y_1 has a direct predecessor in F' , call it y_5 , and there exists x_b such that $\hat{\mu}_b = y_5$ and x_b is indifferent between y_5 and y_2 . Note that $y_5 \neq y_2$ as y_2 is unallocated, and $x_b \neq x_a$ since x_a is allocated to y_1 . We may therefore take F' to be equal to F except that the edge (y_2, y_1) is replaced with the edge (y_5, y_1) . The new graph has one less that is unallocated and has price above reservation. Repeating the above arguments until this number is zero establishes the lemma.

The third case is covered by Lemma 4, which states any allocation that does not satisfy buyer optimality leads to a strictly larger value for $H(\mu)$.

It follows that any solution to the min-max problem is a competitive equilibrium. \square

Proof of Corollary 1: Lemmas 1 through 3 establish that the minimum equilibrium price satisfies (6.2). There can be no other solution to (6.2). \square

Proof of Theorem 2: The first statement follows directly from the arguments used to establish Theorem 1. If $q(\mu, F)$ solves (6.1), then $q(\mu, F)$ is a competitive equilibrium price vector. Hence $q(\mu, F) \geq r$. By assumption if $q_i(\mu, F) > r_i$, there exists x_a such that $\mu_a = y_i$. Finally, it follows from Lemma 3, that if there exists a buyer x_a and a good y_i such that x_a prefers y_i to μ_a at $q(\mu, F)$, then $q(\mu, F)$ does not solve (6.1). This establishes the second statement. \square

Proof of Theorem 3: Consider $\lambda_0 \in \Lambda$. Assumption A guarantees the existence of a minimum price competitive equilibrium. Lemma 1 guarantees that there exists a GA-structure that supports a competitive equilibrium. Hence $\Phi(\lambda_0)$ is non-empty.

\mathcal{G} is a discrete set. Hence $\Phi(\lambda_0)$ is compact.

The upper-hemicontinuity of $\Phi(\lambda)$ and the continuity of $\underline{p}(\lambda)$ follow from applying the theorem of the maximum first to

$$H_i(\mu, \lambda) = \max_{F \in \mathcal{F}_\mu} q_i(\mu, F, \lambda)$$

and then to

$$\underline{p}_i(\lambda) = \min_{\mu \in M} H_i(\mu, \lambda)$$

for each y_i . Note the $q(\mu, F, \lambda)$ are continuous in λ by assumption, and the \mathcal{F}_μ are trivially continuous correspondences in λ , implying that the $H(\mu, \lambda)$ are continuous in λ . Again M is trivially continuous in λ . \square

Proof of Theorem 4: (only if) Given a competitive equilibrium price vector p , we know that $p \in [\underline{p}, \bar{p}]$. If we take $\hat{r} = p$, then p is a minimal price competitive equilibrium price vector.

(if) Let $(\hat{\mu}, \hat{p})$ denote a minimal price competitive equilibrium for a model with reservation prices $\hat{r} \in [\underline{p}, \bar{p}]$, and let $(\bar{\mu}, \bar{p})$ denote a maximal price competitive equilibrium for the original model. We show that there exists (μ', \hat{p}) that is a equilibrium for the original model.

First note that $(\bar{\mu}, \bar{p})$ is a competitive equilibrium for the model with reservation price vector \hat{r} , since raising the vector of reservation prices from r to a point in $(r, \bar{p}]$ only weakens the second condition in the definition of a competitive equilibrium. It follows that $\hat{p} \leq \bar{p}$, since \hat{p} is the minimal equilibrium price vector based on reservation utilities \hat{r} .

Let $Y^A = \{y_i \in \bar{Y} | \hat{p}_i = \bar{p}_i\}$ denote the set of goods for which \hat{p} and \bar{p} agree and $Y^B = \{y_i \in Y | \bar{p}_i > \hat{p}_i\}$ denote the set on which they disagree (note either set may be empty). Define $X^B = \{x_a \in X | \bar{\mu}_a \in Y^B\}$. Let μ' be defined as follows

$$\mu'_a = \begin{cases} \hat{\mu}_a & \text{if } x_a \in X^B \\ \bar{\mu}_a & \text{otherwise} \end{cases}$$

We first show that μ' is an allocation and is onto $H(\hat{p})$. Since $\bar{p}_i > r_i$ for all $y_i \in Y^B$, it follows that all $y_i \in Y^B$ are allocated at \bar{p} and that $|X^B| = |Y^B|$. Since $\hat{p}_i < \bar{p}_i$ if and only if $y_i \in Y^B$, it follows that $D_a(\hat{p}) \subset Y^B$ for all $x_a \in |X^B|$. Since $(\hat{\mu}, \hat{p})$ is a competitive equilibrium, $\hat{\mu} : X^B \rightarrow Y^B$ is a bijection. Since $x_a \notin X^B \implies \bar{\mu}_a \notin Y^B$, we know that $\bar{\mu} : X \setminus X^B \rightarrow Y^A$. Since $\bar{\mu}$ is an allocation it is 1-1 on this domain and only assigns null goods appropriately, ensuring that μ' itself one-to-one and is an allocation. Moreover, since $\bar{\mu}$ is onto $H(\hat{p}) \setminus Y^B$, μ' is onto $H(\hat{p})$.

It remains to show (μ', \hat{p}) satisfies buyer optimality. Since $(\hat{\mu}, \hat{p})$ is a competitive equilibrium, this is clear for $x_a \in X^B$. Suppose that there exists $x_a \in X \setminus X^B$ such that $\mu'_a \notin D_a(\hat{p})$. Now since $(\bar{\mu}, \bar{p})$ is a competitive equilibrium, $\mu'_a = \bar{\mu}_a \in D_a(\bar{p})$. Since $\bar{p} = \hat{p}$ on Y^A , it follows that $D_a(\hat{p}) \in Y^B$. The fact that $\hat{\mu}$ maps X^B onto Y^B implies that $\hat{\mu}$ maps $X \setminus X^B$ into Y^A . This contradiction completes the proof. \square




Figure 1: A directed rooted tree with edges directed away from the root good (y_1)




Figure 2: An example of a GA structure that generates the competitive equilibrium price but involves a non-equilibrium allocation. The solid arrow represents the graph F . The dashed arrow represents the indifference of x_b .