NBER WORKING PAPER SERIES

A GRAPH THEORETIC APPROACH TO MARKETS FOR INDIVISIBLE GOODS

Andrew Caplin
John V. Leahy

Working Paper 16284
http://www.nber.org/papers/w16284

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2010

We thank Mamoru Kaneko, John Leahy Sr., Jeffrey Mensch, Victor Norman, Ennio Stacchetti, and Ivan Werning for helpful comments. Leahy thanks the NSF for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2010 by Andrew Caplin and John V. Leahy. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Graph Theoretic Approach to Markets for Indivisible Goods
Andrew Caplin and John V. Leahy
NBER Working Paper No. 16284
August 2010
JEL No. C63,D40,E1,R31

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.

Andrew Caplin
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
andrew.caplin@nyu.edu

John V. Leahy
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
john.leahy@nyu.edu

A Graph Theoretic Approach to Markets for Indivisible Goods

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

August 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 Mamoru Kaneko, John Leahy Sr., Jeffrey Mensch, Victor Norman, Ennio Stacchetti, and Ivan Werning for helpful comments. Leahy thanks the NSF for financial support.

[1] 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 (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

[2]The 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. One exception is the paper by Demange and Gale (1985) cited above.

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

[4] Quinzii (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 $U_a : Y \times \mathbb{R}^n \to \mathbb{R}$, 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:

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

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

[5] With 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 to be made to ensure that the constraints are not binding in equilibrium, but these do not add insight.

[6] This 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.

[7] Since 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 $\sigma : Y \to X \cup \varnothing$ such that $\mu(\sigma(y_i)) = y_i$ when $\sigma(y_i) \neq \varnothing$.

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 $\bar{p} \in \Pi$:

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

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. $\lim_{p_i \to \infty} U_a(y_i, p_i) = -\infty$ and $\lim_{p_i \to -\infty} U_a(y_i, p_i) = \infty$.

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.

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

[9] This 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 $\mu^1_a = y_1$ and $\mu^1_b = y_2$, and $\mu^2_a = y_2$ and $\mu^2_b = y_1$. 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

[10] A 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 $y_i \in Y \backslash R$, there exists a unique root good $y_r \in R$ and a corresponding unique directed path $\{(y_r, y_1), (y_1, y_2), \ldots (y_{i-1}, y_i)\} \subset E$ 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, $q : \mathcal{G} \to \mathbb{R}^n$. 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 ith 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,

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

   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.

[11] This 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:

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

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:

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

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 a 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 $\Phi : \Lambda \to \mathcal{G}$ denote the mapping from parameters to the set of GA-structures that generate minimum price competitive equilibrium:

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

and let $\underline{p}$: $\Lambda \to \mathbb{R}^n$ 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 $\underline{p}(\lambda)$ 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.

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

[13] See Frantz and Parks [2002].

[14] It 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 $d\underline{p}_k/d\lambda = 0$ for all $y_k \in Y$.
2. If y_k is not a successor of y_0, then $d\underline{p}_k/d\lambda = 0$ for all $y_k \in Y$.
3. If y_k is a successor of y_0 and $\{(y_0, y_1), (y_1, y_2), \ldots (y_n, y_k)\} \subset E(F)$, then

$$\frac{d\underline{p}_k}{d\lambda} = \frac{d\underline{p}_k}{d\underline{p}_{k-1}} \frac{d\underline{p}_{k-1}}{d\underline{p}_{k-2}} \cdots \frac{d\underline{p}_1}{d\underline{p}_0} \tag{8.1}$$

where

$$\frac{d\underline{p}_i}{d\underline{p}_{i-1}} = \frac{dU_{\sigma_{i-1}}(y_{i-1}, p_{i-1})}{d\underline{p}_{i-1}} \bigg/ \frac{dU_{\sigma_{i-1}}(y_i, p_i)}{d\underline{p}_i}$$

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 $\underline{p}_k$ 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 $d\underline{p}_k/d\lambda = 0$.
2. If y_k is a successor of y_0 and $\{(y_0, y_1), (y_1, y_2), \ldots (y_n, y_i)\} \subset E(F)$ is a path in F from y_0 to y_k then

$$\frac{d\underline{p}_k}{d\lambda} = \frac{d\underline{p}_k}{d\underline{p}_{k-1}} \frac{d\underline{p}_{k-1}}{d\underline{p}_{k-2}} \cdots \frac{d\underline{p}_1}{d\lambda}$$

where

$$\frac{d\underline{p}_i}{d\underline{p}_{i-1}} = \frac{dU_{\sigma_{i-1}}(y_{i-1}, p_{i-1})}{d\underline{p}_{i-1}} \bigg/ \frac{dU_{\sigma_{i-1}}(y_i, p_i)}{d\underline{p}_i}$$

and

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

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 υ_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

[15] We thank Victor Norman for this observation.

[16] We 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 sellers 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 $\upsilon \in \mathbb{R}_+^{m+n}$ denote buyer utility. Let υ_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 υ_a^R.

Sellers choose buyers to maximize the price that they receive. The payoff of seller y_i, $p_i(x_a, \upsilon_a)$, depends on the buyer that they sell to x_a and υ_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, $p_i(x_a, \upsilon_a)$ is defined as the price that would have to charged for good y_i to provide buyer x_a with utility υ_a,

$$U_a\left[y_i, p_i(x_a, \upsilon_a)\right] = \upsilon_a.$$

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

An allocation of goods is a one-to-one mapping $\sigma : \bar{Y} \to \bar{X}$ such that 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 $\hat{r} \in [\underline{p}, \bar{p}]$.

10. Concluding Remarks

In this paper and it 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.

11. References

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 Indivisibilitie," *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 Thoery* 1, 111-130.

[22] Quinzii, Martine, (1984), "Core and Competitive Equilibria with Indivisibilities," *International Journal of Game Theory* 13, 41–60.

12. 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,

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

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 $Y \backslash A_s \subset U(p^*)$. Lemma 4 in Demange and Gale (1985) states that there exists $x_a \in X$ such that $\mu_a^* \in A_s$ and $D_a(p^*) \cap Y \backslash A_s \neq \varnothing$. Choose $y_i \in D_a(p^*) \cap Y \backslash A$. 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 $Y \backslash R^*$, 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 $\{y_r, \ldots y_1\}$ 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.□

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 vertix 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 $y_j \in Y \backslash R(F)$, there exists x_d and y_1 such that $\mu_d^* \in A_s$, $(\mu_d^*, y_1) \in E(F^*)$, and $y_1 \in Y \backslash A_s$. 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 $\mu_d \in Y \backslash A_s$. Now for each $y_k \in Y \backslash A_s$, there exists x_b such that $\mu_b^* = y_k$. Hence if

$\mu_d^* \notin Y \backslash A_s$ and $\mu_d \in Y \backslash A_s$, there exists x_e such that $\mu_e^* \in Y \backslash A_s$, but $\mu_e \notin Y \backslash A_s$. 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

$$U_a(y_j, p_j^*) \leq U_a(\mu_a^*, p_{\mu_a^*}^*) \tag{12.1}$$

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. (12.1) 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

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

This completes the proof of the lemma.□

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

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

Lemma 3 implies

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

It follows that

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

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 $\max_{F\in\mathcal{F}_\mu} \sum_{i\in\{1\ldots n\}} q_i(\mu, F)$. 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 $\{y_r, \ldots y_i\}$ 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 $\sum_{i\in\{1\ldots n\}} q_i(\hat{\mu}, \hat{F}) = \max_{F\in\mathcal{F}_\mu} \sum_{i\in\{1\ldots n\}} q_i(\mu, F)$. 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.□

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

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.□

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 equlibrium. 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 $\mathcal{F}_\mu$ are trivially continuous correspondences in λ, implying that the $H(\mu, \lambda)$ are continuous in λ. Again M is trivially continuous in λ.□

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 $(\mu', \hat{p})$ 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, $\hat{\mu} : X^B \to Y^B$ is a bijection. Since $x_a \notin X^B \Longrightarrow \bar{\mu}_a \notin Y^B$, we know that $\bar{\mu} : X \backslash X^B \to Y^A$. 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 $H(\hat{p}) \backslash Y^B$, μ' is onto $H(\hat{p})$.

It remains to show $(\mu', \hat{p})$ 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 $x_a \in X \backslash X^B$ 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 $X \backslash X^B$ into Y^A. This contradiction completes the proof.$\square$

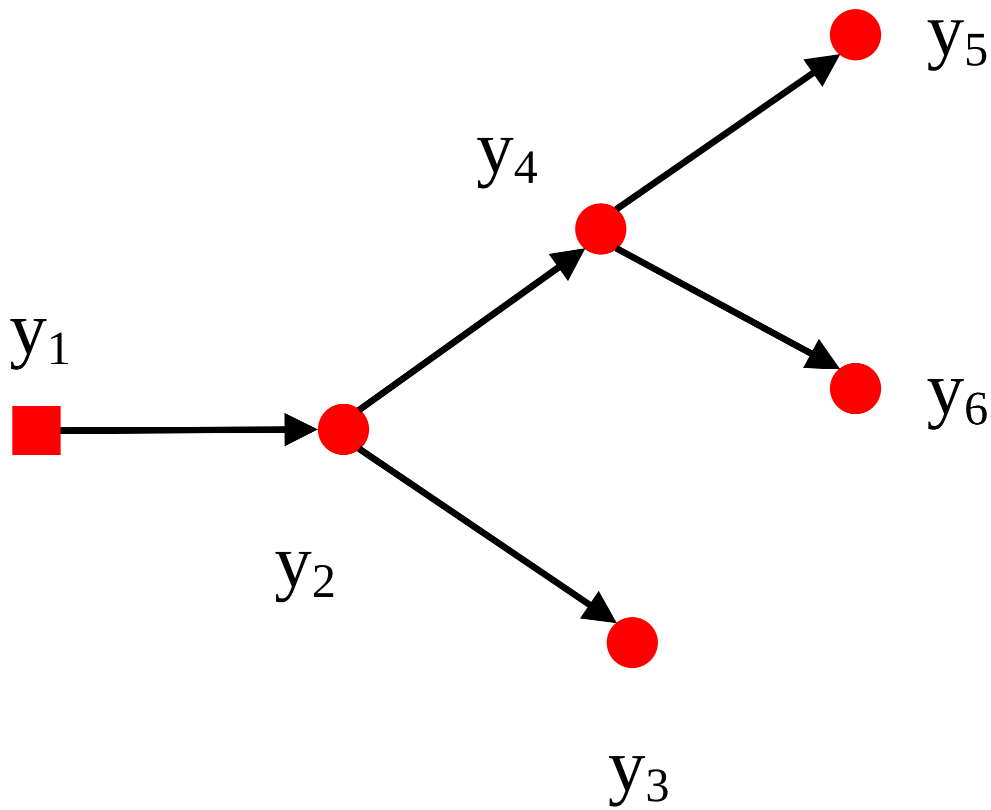

Figure 1: A directed rooted tree with edges directed away from the root good (y_1)

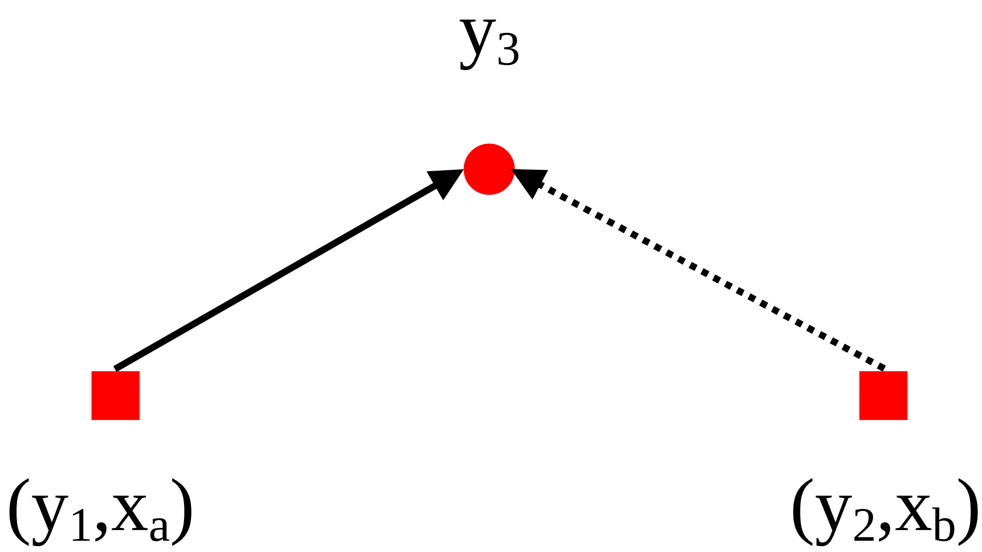

Figure 2: An example of a GA structure that generates the competitive equilibrium price but involves a non-equilibrium allocation. The solid arrow represents the graph *F*. The dashed arrow represents the indifference of x_b.

NBER WORKING PAPER SERIES

COMPARATIVE STATICS IN MARKETS FOR INDIVISIBLE GOODS

Andrew Caplin
John V. Leahy

Working Paper 16285
http://www.nber.org/papers/w16285

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2010

We thank Mamoru Kaneko, John Leahy Sr., Jeffrey Mensch, Victor Norman, Michael Reiter, and Ennio Stacchetti, for helpful comments. Leahy thanks the NSF for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2010 by Andrew Caplin and John V. Leahy. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Comparative Statics in Markets for Indivisible Goods
Andrew Caplin and John V. Leahy
NBER Working Paper No. 16285
August 2010
JEL No. C63,D40,E1,R31

ABSTRACT

We complete the study of comparative statics initiated in Caplin and Leahy [2010], which introduced a new mathematical apparatus for understanding NTU allocation markets, as such covering the housing market and other markets for large indivisible goods. We introduce homotopy methods to characterize how equilibrium changes in response to arbitrary parameter changes. Generically, we show that there can be five and only five qualitatively distinct forms of market transition: Graft; Prune and Plant; Prune and Graft; Cycle and Reverse; and Shift and Replant. Our path-following methods identify new algorithms for computing market equilibria.

Andrew Caplin
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
andrew.caplin@nyu.edu

John V. Leahy
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
john.leahy@nyu.edu

Comparative Statics in Markets for Indivisible Goods

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

August 2010

Abstract

This paper builds upon Caplin and Leahy [2010], which introduced a new mathematical apparatus for understanding NTU allocation markets, as such covering the housing market and other markets for large indivisible goods. In the current paper we complete the study of comparative statics initiated therein. We introduce homotopy methods to characterize how equilibrium changes in response to arbitrary parameter changes. Generically, we show that there can be five and only five qualitatively distinct forms of market transition: Graft; Prune and Plant; Prune and Graft; Cycle and Reverse; and Shift and Replant. Our path-following methods identify new algorithms for computing market equilibria.

1. Introduction

Among the most significant of all markets are those, such as the housing market, that allocate large indivisible goods to households. These goods are best modeled as allocation markets with non-transferable utility (NTU). Caplin and Leahy [2010] (henceforth CL) introduced a new mathematical apparatus for understanding these markets. In the current paper we complete the study of comparative statics initiated therein. We use homotopy methods to characterize how the minimum price competitive equilibrium changes in response to arbitrary parameter changes. Given an initial equilibrium, we follow a path through parameter space, building up the discrete change in equilibrium from infinitesimal ones. We show that along the generic path equilibria evolve in a very controlled manner. There are five and only five qualitatively distinct forms of

*We thank Mamoru Kaneko, John Leahy Sr., Jeffrey Mensch, Victor Norman, Michael Reiter, and Ennio Stacchetti, for helpful comments. Leahy thanks the NSF for financial support.

market transition. The homotopy path decomposes all comparative statics into these five basic operations.

In addition to identifying market transitions, we show that our path-following methods are ideally suited to algorithmic use. They can be used in principle to compute minimum price equilibria from a starting point with a trivial equilibrium. The resulting algorithms may help to overcome computational barriers to application of the NTU allocation model.[1] It is also not hard to envision adopting our methods to studying the dynamic properties of markets for indivisible goods by recasting the homotopy paths as sample paths which characterize the evolution of the market over time.

Demange and Gale [1985] show that there exists a minimum price equilibrium in the class of allocation markets that we consider, those in which each agent is either a buyer or a seller. CL showed that these equilibria correspond to the solutions to a certain optimization problem. This optimization is over a set of mathematical structures, called GA-structures, that combine an allocation of goods to buyers along with a particular class of directed graphs that summarize key indifference relations among goods. CL show for small changes in parameter values the allocation and the indifference relations that characterize the minimum price competitive equilibrium are unlikely to change. Given that "local" comparative statics involve a fixed GA-structure, they are able to show that changes in model parameters that cause the price of a given good to change impact only those goods that "follow" it in the corresponding directed graph.

In this paper we use GA-structures to characterize comparative statics in the large. We study how the allocation of goods and the pattern of local interdependence changes when there is a discrete shift in the character of the minimum price equilibrium. While there are examples in which the most minor change in parameters cause the entire structure of the equilibrium to change in arbitrary ways, we show such cases to be the exception rather than the rule. We introduce a natural definition of regularity for comparative static paths and show generically small parameter changes can induce five and only five distinct forms of market transitions. We use GA-structures to illustrate the precise nature of each such transition. For reasons that will be clear, the market transitions are labeled as: Grafting; Pruning; Prune and Graft; Cyclic Reversal; and Shift and Re-plant.

Markets for indivisible goods are characterized by multiple equilibria. Due to the presence of non-convexities, small changes in prices often do not alter the equilibrium allocation. We focus

[1] The 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.

on the minimum equilibrium price for a number of reasons. First, we learn about the entire range of possible equilibria. The minimum and maximal equilibrium prices define the boundaries of all equilibria. The two prices have the same structure. The maximal equilibrium price is simply the dual of the minimal, in which the roles of buyers and sellers are reversed. Moreover, CL show that all equilibrium prices may be thought of as minimal equilibrium prices for some set of parameter values. Finally, Demange and Gale [1985] show that the minimal equilibrium price vector is not manipulable by buyers, making it a natural benchmark for multi-unit auctions.

Section 2 presents the general model and summarizes relevant results from CL. The additional results in this paper rest in part on our ability to count the number of distinct GA-structures. The key cardinality results are in section 3. They are derived by connecting the number of GA-structures with the structure of market demand at minimum equilibrium prices. We show that there is typically only one GA-structure, and that the most important points of market transition involve two and only two such structures. Section 4 introduces the domain in which we study comparative statics, which involves paths through a rich space of model parameters. Section 5 identifies the sense in which comparative static transitions are almost always "regular", in that the replacement for a given GA-structure is a unique second element that appears at a point of transition. Section 6 identifies the five generic forms of market adjustment. Section 7 shows how to use path-following methods to algorithmically identify the minimum equilibrium price. Section 8 concludes. All proofs are in the Appendix.

2. Background

2.1. Model

There is a set of buyers $x_a \in X$, $1 \leq a \leq m$, and a set of indivisible goods $y_i \in Y$, $1 \leq i \leq n$ with $n > m$. 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.

Buyers derive utility from at most one element of Y. The payoff for buyer x_a is summarized by the utility function $U_a : Y \times \mathbb{R} \to \mathbb{R}$, where $U_a(y_i, p_i)$ is the utility to x_a from the purchase of y_i with a transfer of p_i of the numeraire good.

The supply side is trivial. Sellers choose only whether or not to sell. They do not purchase the indivisible goods from other sellers. Each seller wishes only to obtain the highest possible price above a reservation level. Let $r \in \mathbb{R}^n_+$ denote the vector of seller reservation prices. Overall, we can limit attention to a set Π in the search for competitive equilibria,

$$\Pi = \left\{ p \in \mathbb{R}^n_+ | p \geq r \text{ and } p_{n-a} = r_{n-a} = 0 \text{ all } 0 \leq a \leq m-1 \right\}.$$

Given any price vector $p \in \Pi$, the (non-empty) demand correspondence $D_a(p)$ specifies members of Y that maximize utility the utility of buyer x_a,

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

An *allocation* $\mu : X \to Y$ is a one-to-one mapping from buyers to goods. The set of all allocations is M

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.

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 for 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 make continuity and range assumptions that guarantee that the set of equilibria is a closed lattice (see Demange and Gale [1985]). The first assumption is a straight forward regularity assumption. 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 price for the second that makes the buyer indifferent between the two goods.

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. $\lim_{p_i \to \infty} U_a(y_i, p_i) = -\infty$ and $\lim_{p_i \to -\infty} U_a(y_i, p_i) = \infty$.

With Assumption A we know not only that there is a competitive equilibrium allocation, but also that there is a minimum equilibrium price (henceforth MEP) $\underline{p} \in \Pi$, which has the property that it is an equilibrium price and that, for any equilibrium price,

$$p_i \geq \underline{p}_i \text{ all } y_i \in Y.$$

Note we have presented the model without explicit reference to buyers' endowment of the numeraire good. We have implicitly assumed that buyers hold enough of the numeraire good to buy the goods that they are assigned at the competitive equilibrium prices. An alternative approach would have been to have expressed utility in terms of the endowment left over after purchase and to have included the additional assumption that each buyer preferred exit to spending all of his endowment on a purchase. This condition would hold for any utility function that satisfied the Inada conditions. This is the approach taken by Kaneko (1982) and is equivalent to the current one.

2.2. GA-Structures

CL introduce GA-structures, which combine an allocation of goods with a particular class of graph, to study the structure of equilibria in the allocation model. The graphs in question are directed graphs on vertex set Y that: are forests of trees (i.e. contain no cycles); that are rooted, in the sense that each component tree has a unique element that is specified as its root, and in which each edge is directed away from the root. We let $\mathcal{F}$ denote the class of directed rooted forests on vertex set Y. Note the insistence that all goods that are at the "tail" of a directed edge are allocated.

Definition *A graph-allocation structure (GA-structure)* comprises:

1. A directed graph $F = (Y, R, E)$ with vertex set Y, root set $R \subseteq Y$, and edge set $E \in Y^2$ in which F is a forest of trees, each component of F contains a unique element of R and each edge $(y_i, y_k) \in E$ is directed away from the corresponding element of R.
2. An allocation $\mu \in M$ such that, if $(y_i, y_k) \in E$, then there exists $x_a \in X$ such that $\mu(x_a) = y_i$.

We let $\mathcal{G}$ denote the set of all GA-structures.

CL construct a natural mapping from $\mathcal{G}$ into prices by first pricing root goods at reservation, and then iteratively pricing successors by indifference of the individual allocated to their unique predecessor. Assumption A guarantees that this construction is well-defined.

Definition The function $q : \mathcal{G} \to \mathbb{R}^n$, *the price generated by* graph-allocation structure $(\mu, F) \in \mathcal{G}$ is defined by setting $q_i(\mu, F) = r_i$ for all $y_i \in R(F) \equiv A_0$, with iterative pricing of successor goods $y_k \in S$ of A_s for $s \geq 1$ based on the indifference condition,

$$U_a(y_i, q_i(\mu, F)) = U_a(y_k, q_k(\mu, F)). \tag{2.1}$$

where $\mu_a = y_i$, and $(y_i, y_k) \in E(F)$, so that $y_i \in A_s$ is the direct predecessor of $y_k \in S$.

2.3. The Min-Max Theorem

CL establish that with Assumption A, a solution to an optimization problem on these generated prices identifies the minimum equilibrium price. It also establishes that any allocation that forms part of a minimum price equilibrium (MPE) must also be identifiable from this optimization problem.

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

$$\sum_{i \in \{1 \ldots n\}} q_i(\mu^*, F^*) = \min_{\mu \in M} \max_{F \in \{F \in \mathcal{F} | (\mu, F) \in \mathcal{G}\}} \sum_{i \in \{1 \ldots n\}} q_i(\mu, F). \tag{2.2}$$

Moreover, let (μ^*, F^*) be an argument that solves (2.2), if $q_i(\mu^*, F^*) > r_i$ for all $y_i \notin R(F)$ then $(\mu^*, q_i(\mu^*, F^*))$ is a MPE.

Comparative statics are particularly simple when there is one and only one element $(\mu^*, F^*) \in \mathcal{G}$ that solves (2.2) for some range of parameters. In such cases, a continuous change in the model's parameters do not change in the allocation μ or the graph F. Changes in sellers reservation have a direct effect on prices of root goods and changes in a buyer utility directly affects the prices of goods that are direct successors of the good that that buyer is allocated to. These changes propagate through the graph structure affecting the price of all successors in F.

3. Cardinality and the Demand Graph

The focus of this paper is on cases in which a change in parameters is large enough to force a change in the equilibrium GA-structure. It turns out that this generally happens in a very controlled manner. There is a sense in which it is rare to come across more than two solutions to (2.2). For most parameter values there is only one solution. As one moves through the parameter space, the generic path follows this single solution until at some point there a second solution appears along with the first.[2] The path then follows one of these two solutions for a while. The current section formalizes these arguments using counting arguments.

3.1. The GAME set

We begin by stating exactly what we mean to count. It is the set of $(\mu, F) \in \mathcal{G}$ which generate the minimum equilibrium price and also have the property that μ allocates buyers not only to the tails but to the heads of all edges in $E(F)$.[3]

Definition: The *graph allocation minimum price equilibrium (GAME) set* is,

$$\Phi = \{(\mu, F) \in \mathcal{G} | (\mu, q(\mu, F)) \text{ is a MPE and for all } (y_i, y_k) \in E(F) \Longrightarrow \mu_a = y_k \text{ some } x_a \in X\}.$$

The condition that μ allocates a buyer to the head of each edge in $E(F)$ rules out some uninteresting cases. Suppose, for example, that two goods y_1 and y_2 are both priced at reservation, and that there is only one buyer who is indifferent between them at these prices. In this case, there are four elements of $\mathcal{G}$ that generate the minimum price competitive equilibrium and hence "solve" the min-max problem. There are two potential allocations, and for each allocation we may either price both goods at reservation or use the buyer's indifference to price the unallocated good. The condition that μ allocates a buyer to the head of each edge in $E(F)$ rules out the use of indifference to price the unallocated good. In this case, $\Phi = 2$.

Note that Φ depends on the parameters of the model through the minimum price equilibrium. We suppress this dependence, as we are considering a fixed set of parameters at this point.

3.2. Demand Graphs and Condition M

The cardinality of the GAME set Φ is basic in the work that follows. We derive conditions under which Φ is unique and under which it has two elements. We develop these cardinality results using properties of demand at minimum equilibrium prices. We introduce the MPE

[2] Throughout this paper we use the term generic in the topological sense. A property is generic if it holds on a dense open set.

[3] The definition of Φ in this paper differs from that in CL. In that paper we did not need the restriction that the heads of all edges are allocated.

demand graph and a class of its subgraphs that turn out to be of particular importance in this characterization.

Definition: The MPE *demand graph* $D(\underline{p}) = (X, Y, \mathcal{E}^*)$ is a bipartite graph with partition X and Y, and an edge set defined by,

$$\mathcal{E}^* = \{(x_a, y_j) \in X \times Y | y_j \in D_a(\underline{p})\}.$$

Definition: The edge set $\mathcal{E} \subset \mathcal{E}^*$ is said to satisfy *Condition M*, written $\mathcal{E} \in \mathcal{M}(\underline{p})$, if there exists a partition $\{Y_l^{\mathcal{E}} | 0 \leq l \leq L\}$ of Y such that,

M1. $Y_0^{\mathcal{E}}$ is comprised of the elements of Y that isolated in the demand graph.

M2. For $l \geq 1$, each element of $Y_l^{\mathcal{E}}$ is connected to the other elements of $Y_l^{\mathcal{E}}$, but not to elements of $Y \backslash Y_l^{\mathcal{E}}$.

M3. For $l \geq 1$, there is one and only one element, $\sigma_l^{\mathcal{E}} \in Y$, of degree (valence) 1 in each partition set $Y_l^{\mathcal{E}}$. All other elements of $Y_l^{\mathcal{E}}$ have degree 2.

M4. For $l \geq 1$, $|X_l^{\mathcal{E}}| = |Y_l^{\mathcal{E}}|$ where $X_l^{\mathcal{E}} = \left\{x_a \in X | (x_a, y_i) \in \mathcal{E} \text{ some } y_i \in Y_l^{\mathcal{E}}\right\}$.

M5. $\underline{p}_i > r_i$ implies that the degree of y_i is equal to two.

Both the demand graph D and the set of subgraphs $\mathcal{M}$ depend on the price vector $\underline{p}$ which determines the set of goods in each buyer's demand set. To save on notation, we suppress this dependence in much of what follows as we are considering a fixed set of parameters at this point.

3.3. A Counting Lemma

The reason for introducing these particular graphs is that they will correspond in a precise manner with elements of Φ. $\mathcal{M}$ and Φ are essentially equivalent. Given a MEP $\underline{p}$, we introduce a bijection $\eta : \mathcal{M} \to \Phi$ between Φ and $\mathcal{M}$ that provides structure to counting arguments. The proofs of all propositions are in the appendix.

Lemma 1: Let $\underline{p}$ be a minimum equilibrium price vector. There exists a bijection $\eta : \mathcal{M} \to \Phi$ so that

$$|\Phi| = |\mathcal{M}|$$

Given $\mathcal{E} \in \mathcal{M}$, the graph $(X, Y, \mathcal{E})$ is a bipartite graph with partition $\{X, Y\}$. Since it is a selection from the demand graph at $\underline{p}$, each buyer is matched only to goods that are optimal at these prices. With conditions M1-M4, $(X, Y, \mathcal{E})$ is acyclic. Given one good in each component with degree 1, a cycle would either lead to fewer buyers than sellers or require a good with degree greater than two. Each partition element $Y_l^{\mathcal{E}}$ for $l \geq 1$ therefore corresponds to a distinct tree. In each such tree, the root good has degree 1 and all others have degree 2. F is constructed recursively, beginning with the valence one goods and then at each step adding as direct successors all goods that are two edges distant in $\mathcal{E}$. Finally, $Y_0^{\mathcal{E}}$ is added to the root set. The equality between the number of goods and buyers in each such set generates a natural mapping μ. The valence one goods are allocated to their neighbors in $\mathcal{E}$. Removing these buyers creates a new set of valence one goods, who are then allocated to their neighbors, and so on. Note that the equality of buyers and sellers in M4 guarantees that all goods in $Y_l^{\mathcal{E}}$ are allocated. This implies that all non-root goods are allocated and that the heads and tails of all edges in $E(F)$ are allocated. Finally, it is easy to show that the (μ, F) generated in this way generates $\underline{p}$, so that $(\mu, F) \in \Phi$.

This bijection is illustrated in Figure 1 for a case in which there are five buyers and six goods. Figure 1(a) illustrates $\mathcal{E} \in \mathcal{M}$, with the partition $\{Y_l^{\mathcal{E}} | 0 \leq l \leq L\}$. The sixth good is unassigned, $Y_0^{\mathcal{E}} = y_6$. The other two partition sets are $Y_1^{\mathcal{E}} = \{y_1, y_2, y_3\}$ and $Y_2^{\mathcal{E}} = \{y_4, y_5\}$. To draw the corresponding allocation $\eta(\mathcal{E})$ we iterate from the goods of degree 1, so that $\mu_i = y_i$ for $1 \leq i \leq 5$. The forest is found by placing down the three root goods, y_1, y_4, and y_6 (which is unoccupied) and then drawing all edges corresponding to additional demands of their matched buyers, x_1 and x_4. Buyer x_4 demands y_5 in addition to y_4, while buyer x_1 demands y_2 and y_3 in addition to y_1. Adding these as directed edges from the respective root good completes the forest, as illustrated in Figure 1(b). The inverse mapping is also clear: if one starts from the forest of directed rooted trees, all goods are demanded by the individual allocated to them, and any individual allocated to a good at the origin of a directed edge demands all goods to which such edges are directed.

3.4. Uniqueness

A characterization of demand graphs that give rise to uniqueness of the GAME set immediately from Lemma 1. In the statement, $d^*(y_i)$ is the degree of vertex y_i associated with demand graph D.

Corollary 1: $\mathcal{E}^* \in \mathcal{M}$ and $d^*(y_i) = 2 \Longrightarrow \underline{p}_i > r_i$ then $|\Phi| = 1$.

Since $\mathcal{E}^* \in \mathcal{M}$ all goods have valence less than two. Any other $\mathcal{E} \in \mathcal{M}$ is a subgraph of $\mathcal{E}^*$, and therefore involves reducing the valence of some good. Reducing $d^*(y_i)$ from two to one violates M5 since $d^*(y_i) = 2 \Longrightarrow \underline{p}_i > r_i$. Reducing $d^*(y_i)$ from one to zero reduces the number of goods that are connected to buyers and violates M4.

We identify $\mathcal{U} \subset \mathcal{M}$ as the set of graphs satisfying the conditions of the uniqueness theorem. This set is of particular value in enumerating elements of Φ when there is more than one.

3.5. Cardinality Two

Corollary 1 implies that the existence of multiple equilibrium GA structures is associated either with one or more "extra" indifference relationships in the demand graph that can be removed and have an edge structure that satisfies Condition M, or with some set of goods with valence greater than one which have prices equal to their reservation values.

Our next result concerns circumstances in which $|\Phi| = 2$. This is the case whenever the demand graph associated with the minimum price equilibrium has one and only one edge more than some subgraph $\tilde{\mathcal{E}} \subset \mathcal{E}^*$ such that $\tilde{\mathcal{E}} \in \mathcal{U}$.

While the theorem is proven in the appendix in standard analytic manner, certain of the arguments are best understood in the figures that follow.

Theorem 1: $|\Phi| = 2$ if and only if there exists $(\bar{x}, \bar{y}) \in \mathcal{E}^*$ such that $\mathcal{E}^*/(\bar{x}, \bar{y}) \equiv \tilde{\mathcal{E}} \in \mathcal{U}$.

The proof considers five cases. Each case is distinguished by the position of the good $\bar{y}$ and the buyer $\bar{x}$. We illustrate how beginning with $\tilde{\mathcal{E}} \in \mathcal{U}$, the addition of an extra edge makes possible the construction of a second element $\mathcal{E}' \in \mathcal{M}$.

Figure 2 illustrates a case in which $\bar{y} \in Y_0^{\tilde{\mathcal{E}}}$. In this example, there are two buyers and three goods. Figure 2(a) illustrates the edge set $\tilde{\mathcal{E}} \in \mathcal{U}$:

$$\tilde{\mathcal{E}} = \{(x_1, y_1), (x_1, y_2), (x_2, y_2)\},$$

The additional edge in $\mathcal{E}^*$ is (x_2, y_3) as illustrated in Figure 2(b). We obtain a second element of $\mathcal{E}' \in \mathcal{M}$ by removing the edge (x_1, y_1), as illustrated in figure 2(c).

Figure 3 illustrates a case in which $\bar{y} \in Y_l^{\tilde{\mathcal{E}}}$ and $\bar{x} \in X_l^{\tilde{\mathcal{E}}}$ for some $l \geq 1$. In this example, there are three buyers and three goods. We begin with the edge set $\tilde{\mathcal{E}} \in \mathcal{U}$ comprising,

$$\tilde{\mathcal{E}} = \{(x_1, y_1), (x_1, y_2), (x_2, y_2), (x_1, y_3), (x_3, y_3)\},$$

This is illustrated in Figure 3(a). The additional edge in $\mathcal{E}^*$ is (x_3, y_2) as illustrated in Figure 3(b). $\mathcal{E}' \in \mathcal{M}$ is obtained by removing edge (x_1, y_2), as illustrated in Figure 3(c).

Figure 4 illustrates a case with $\bar{\mu} \in Y_l^{\tilde{\mathcal{E}}}$ and $\bar{y} \in Y_m^{\tilde{\mathcal{E}}}$ with $m \notin \{0, l\}$, and with $\bar{y}$ of valence 2 in $\tilde{\mathcal{E}}$. Again there are three buyers and three goods. The edge set $\tilde{\mathcal{E}} \in \mathcal{U}$ comprises,

$$\tilde{\mathcal{E}} = \{(x_1, y_1), (x_2, y_2), (x_2, y_3), (x_3, y_3)\},$$

as illustrated in Figure 4(a). The additional edge in $\mathcal{E}^*$ is (x_1, y_3) as illustrated in Figure 4(b). $\mathcal{E}' \in \mathcal{M}$ is obtained by removing (x_2, y_3), as illustrated in figure 4(c).

In the final case, $\bar{\mu} \in Y_l^{\tilde{\mathcal{E}}}$, $\bar{y} \in Y_m^{\tilde{\mathcal{E}}}$ with $m \notin \{0, l\}$, and $\bar{y}$ has valence 1. This case is similar to that in Figure 4, except that $(x_2, y_3) \notin \tilde{\mathcal{E}}$.

It is important to note that there is something special about the addition of a single edge to $\tilde{\mathcal{E}} \in \mathcal{U}$ that leads to one and only one additional element of Φ. The addition of two edges does not have a simple structure. It is easy to construct examples in which the addition of two edges to $\tilde{\mathcal{E}} \in \mathcal{U}$ can lead to Φ having three, four, or six elements depending on the position of the additional goods and buyers in the graph $(X, Y, \tilde{\mathcal{E}})$.

4. Regular Comparative Static Paths

We now consider a set of models indexed by a parameter vector $\lambda \in \Lambda$ and consider comparative statics with respect to λ. CL show that if the $q(\mu, F; \lambda)$ are smooth in a sense to be defined below, the effect of an infinitesimal change in λ on the minimum equilibrium price can almost always be analyzed using a fixed GA structure. In this case, the change in λ only affects the prices of goods directly influenced by the change in λ and their successors in the relevant graph F.

Discrete changes in λ are more complicated because they lead to shifts in the GA structure. In this section, we use homotopy methods to characterize discrete comparative statics.[4] Given two points λ_0 and λ_1 in Λ, we start with a known equilibrium at λ_0 and consider paths through Λ running from λ_0 to λ_1. We let $\Phi(\lambda)$ denote the dependence of the GAME set on λ. We show that if the $q(\mu, F; \lambda)$ are smooth and if Λ is large enough, then $\Phi(\lambda)$ has the following two properties along the generic path: (1) $\Phi(\lambda)$ has at most two elements at every point along the path; and (2) $\Phi(\lambda)$ has two elements at a only finite number of points so that it is single valued almost everywhere. We call these paths regular comparative static paths.

[4] See Judd (1999, p. 179) for an introduction to homotopy methods.

Comparative statics are conceptually simple along regular comparative static paths. One follows the implications of a change in parameters for a single GA structure until at some point a second solution to the min-max problem emerges. According to Theorem 1, this would happen if some the demand set of some buyer expanded or if the price of some good fell to its reservation value. At that point, one would choose the structure that solved the min-max problem along the next segment of the path. The following sections formalize these arguments.

4.1. Model Types

We consider a set of models indexed by a parameter vector $\lambda \in \mathbb{R}^{mn+n}$. The first mn components of λ are shifters ξ_{ai} of each buyer x_a's utility for each good y_i. We let $U_a(y_i, p_i, \xi_{ai})$ be the utility to x_a from the purchase of y_i at the price p_i when the taste parameter is ξ_{ai}. The remaining n components are the reservation prices of the sellers r_i. We restrict λ to lie in a set $\Lambda \subseteq \mathbb{R}^{mn+n}$ in which all the $U_a(y_i, p_i, \xi_{ai})$ are strictly monotonic in ξ_{ai}. This avoids "flat spots" in which a change in the parameters has no effect on the prices. In addition, we assume that Λ is non-empty, open, bounded, and convex.[5]

We make the following smoothness assumption on the utility functions.[6]

Assumption B: The utility functions $U_a(y_i, p_i, \xi_{ai})$ are analytic in p_i and ξ_{ai} for all $x_a \in X$ and $y_i \in Y$. Moreover the $U_a(y_i, p_i, \xi_{ai})$ are strictly monotonic in ξ_{ai} for all $\lambda \in \Lambda$.

Many of the objects that we have been considering become functions of λ. $q(\mu, F, \lambda)$ is the price-induced by (μ, F) in model $\lambda \in \Lambda$. $\underline{p}(\lambda)$ is the minimum equilibrium price in model $\lambda \in \Lambda$. $\Phi(\lambda)$ is the GAME set. $\mathcal{M}(\lambda)$ as the edge sets that satisfy condition M.

The next lemma follows from the properties of analytic functions and the Theorem of the Maximum applied to the max-min problem (2.2).

Lemma 2: With Assumptions A and B, each $q_i(\mu, F, \lambda)$ is analytic at λ for all $\lambda \in \Lambda$; $\Phi(\lambda)$ is non-empty, compact-valued, and upper-hemicontinuous; and $\underline{p}_i(\lambda)$ is continuous for all $y_i \in Y$.

[5] Taking Λ as open avoids the question of how to do comparative statics at the boundary of the parameter space. Boundedness will imply that closed subsets are compact. Convexity implies connectedness, which is a natural assumption when considering continuous paths. Convexity will also prove useful in constructing perturbations of paths.

[6] A function $f(x)$ is analytic at a point x_0, if its Taylor series expansion converges on a neighborhood of x_0. Almost all commonly used utilty functions are analytic almost everywhere in their domain.

Moreover, it follows from Lojasiewicz's Structure Theorem for Real Varieties (Krantz and Parks [2002], p. 168) that the GAME set is generically single valued. Let $\hat{\Lambda} \equiv \{\lambda \in \Lambda | |\Phi(\rho)| > 1\}$.

Theorem 2: With Assumptions A and B, $\hat{\Lambda}$ is open and dense in Λ.

4.2. Paths

A typical comparative static exercise involves some well defined change in model parameters from λ_0 to λ_1. We will represent these changes with a path through Λ starting from λ_0 and ending at λ_1. Technically, a path is a continuous mapping of the unit interval into parameter space, $\pi : [0, 1] \to \Lambda$, with $\pi(0) = \lambda_0$ and $\pi(1) = \lambda_1$. Any path induces a correspondence $\phi(z) : [0, 1] \to \mathcal{G}$ which maps each z into the GAME set $\phi(z)$.

$$\phi(z) = \Phi(\pi(z)).$$

Note that ϕ inherits the properties of Φ. It is upper-hemicontinuous, non-empty and compact valued.

4.3. Regular Paths

Even though Φ is almost everywhere single valued, it will not be possible in general to find a path between two sets of parameters that avoids points at which Φ takes on multiple values. We consider now what happens at such points.

A simple example shows that in general anything can happen. Suppose that at $\bar{\lambda}$ all buyers have identical preferences and that the path π passes through $\bar{\lambda}$, then, in principle, one can transition from any market situation to any other situation with the most minor of changes in model parameters.

This example in which the equilibrium allocation can change in an arbitrary manner relies on a seeming coincidence, with many individuals suddenly become indifferent to various goods the same point along the path. The results of the last section suggest that cases in which there is only one new indifference at a time, or in which one and only one good falls to reservation price at a given point in the parameter space, will be more ordered.

The simplest case is a path at which the number $|\Phi|$ was no higher than 2 all the way along the path. In fact the following cases are the easiest of all.

Definition 4.1. *A path $\pi(z)$ is regular if the induced mapping $\phi(z)$ has the following properties:*

1. $\max_{z\in[0,1]} |\phi(z)| \leq 2$.

2. $Z = \{z \in [0,\bar{z}] || \phi(z)| > 1\}$ is finite.

Along a regular path, comparative statics is conceptually easy. One travels along a path, working out the implications of a change in parameters for a given GA structure until at some point a new element of Φ arises. At this point, one chooses to continue with the structure that solves the max-min problem beyond that point. There are a finite number of such switches.

The cardinality results of the last section at least give some hope that there may be many regular paths, given that particular transitions can be shown to ensure cardinality of no more than 2. In the next section, we show in fact that regularity is a generic property of comparative static paths.

5. The Generic Path is Regular

5.1. Analytic Shapes

Paths can take many forms and the set of possible paths is quite large. In order to keep the analysis manageable, we consider a notion of genericity based on fixing the "shape" of a path and varying the initial condition. We show that given the shape of the path, the initial conditions associated with regular paths form a dense open set. In other words, the generic path is regular.

The shape is a continuous mapping $S : [0,1] \to \mathbb{R}^{mn+n}$ such that $S(0) = 0$. The initial condition is a point in parameter space $\lambda_0 \in \Lambda$. The pair (λ_0, S) define a path $\pi(z; \lambda_0, S) = \lambda_0 + S(z)$, which begins at λ_0 and ends at $\lambda_1 = \lambda_0 + S(1)$. A path is admissible if $\lambda_0 + S(z) \in \Lambda$ for all $z \in [0,1]$. Let Λ_S denote the set of λ_0 for which $\lambda_0 + S(z)$ is admissible.

For the remainder of the paper we fix the shape S and assume that Λ_S is a non-empty, open, bounded, and convex subset of Λ.

Assumption C: Each component of $S(z)$ is analytic on (0,1).

5.2. Regularity is Generic

Let $\Lambda_R = \{\lambda_0 \in \Lambda_S | \pi(z; \lambda_0, S) \text{ is regular}\}$ denote the set of λ_0 for which the path $\pi(z; \lambda_0, S)$ is regular. Theorem 3 states the sense in which almost all paths are regular.

Theorem 3: With A-C, Λ_R is open and dense in Λ_S.

The proof of Theorem 3 involves a number of steps. The upper hemicontinuity of Φ can be used to show that the set of non-regular paths is closed. To show that the Λ_R is dense in Λ_S we first show that the set of points $\Lambda_F = \{\lambda_0 | Z = \{z \in [0, \bar{z}] || \phi(z)| > 1\}$ is finite$\}$ is dense in Λ_S. To see this, note that each point in Z is associated with an intersection of at least two $q(\mu, F, \pi(z))$. We show that Assumptions B and C imply that the $q_i(\mu, F, \pi(z))$ are analytic in z for all $z \in (0,1)$, $(\mu, F) \in \mathcal{G}$, and $y_i \in Y$. Analytic functions whose intersections have accumulation points must be equal everywhere. We show that whenever $q_i(\mu, F, \pi(z))$ are identical for two (μ, F) we can find a perturbation of λ_0 such that they are different. This establishes Λ_F is dense in Λ_S. To establish that the set Λ_R is dense in Λ_F we show that given $\lambda_0 \in \Lambda_F$, we can perturb λ_0 in such a way that we reduce the number of edges in $\mathcal{E}^*$ at points in which $|\phi(z)| > 2$.

6. The Five Market Transitions

According to Theorem 3, on a regular path there are only a finite number of points at which the GA structure changes and at each of these points of transition $|\phi(z)| = 2$. According to Theorem 1, points at which $|\phi(z)| = 2$ are associated with an "extra" indifference in the demand graph. The extra indifference could arise either with the expansion of the demand set of a single buyer to a single new good or with a single price falling to r_i, thereby making superfluous the demand of the buyer who heretofore had supported the good.

The implication is that along the generic path there are only a few ways in which the structure of the market can change. In fact, there are five types of transition. Four are associated with an expansion of a buyer's demand set. They differ in the position of the good demanded, whether it is (1) unallocated, (2) allocated and a root good in another component, (3) allocated and a predecessor in the same component, or (4) none of the above. The last case is associated with a contraction in the graph satisfying the conditions of the Corollary 1. This happens when the price of a good falls to its reservation level. Below we discuss each case in terms of the impact on the GA-structure, since it is this that illustrates most clearly how, if at all, the allocation changes, and how the structure of market interdependence changes at critical transition points.

We present the cases in the order of complexity. In each case, we describe what happens at a point of market transition z. There is some GA structure (μ, F) which has characterized the market prior to z. At z either some buyer $\bar{x}$ becomes indifferent to some good $\bar{y}$ or the price of some good $\bar{y}$ falls to $\bar{r}$. We describe the relationship between (μ, F) and (μ', F'), the GA structure that characterizes the market after z. It is useful to let $\bar{\mu}$ denote the good assigned $\bar{x}$

by μ.

6.1. Graft

Grafting occurs when $\bar{x}$'s demand correspondence expands to include a good $\bar{y}$ that is the root of another tree. The tree with $\bar{y}$ as its root is incorporated in $\bar{x}$ to create a single larger tree.

An example of grafting is illustrated in Figure 5. Figure 5(a) shows the initial GA structure (μ, F). There are five goods. The only goods that are labeled are $\bar{y}$, the root of the second tree, and $\bar{\mu}$, the good in the first tree whose buyer becomes indifferent to $\bar{y}$. This indifference is illustrated by the light directed edge from $\bar{\mu}$ to $\bar{y}$. It is clear that $\bar{y}$ can be priced in two ways: it can be set at reservation or it can be priced using $\bar{x}$'s indifference. This latter GA structure is illustrated in Figure 5(b). It involves solidifying the link from $\bar{\mu}$ to $\bar{y}$ and the uprooting of $\bar{y}$, as illustrated by the circular as opposed to square node.

In general, grafting is characterized by the following rules. First, there is no change in the equilibrium allocation, $\mu' = \mu$. Second, F' differs from F in just two ways: the added edge from $\bar{\mu}$ to $\bar{y}$, $E(F') = E(F) \cup (\bar{\mu}, \bar{y})$, and the removal of $\bar{y}$ from the root set, $R(F') = R(F) \backslash \bar{y}$.

6.2. Prune and Plant

Pruning is the opposite of grafting. It involves the division of one tree into two separate trees. It occurs when some good $\bar{y}$ that used to be above reservation price falls to this price, and at that point becomes a root good. This occurrence is illustrated in Figure 6(a). The new GA structure involves severing of the proximate link to the new root good, as illustrated in Figure 6(b). Again, there is no change in the allocation.

6.3. Prune and Graft

As its name indicates, prune and graft combines both pruning and grafting. This case involves separating a branch from a tree and attaching it somewhere else. This is the "none of the above" case described above, in which $\bar{y}$ is allocated, but not a root good or a predecessor to $\bar{\mu}$. As in the cases of pruning and grafting there is no change in the allocation. As with pruning, the link between $\bar{y}$ and its direct predecessor in F is cut. As with grafting, $\bar{y}$ is grafted onto the new directed edge from $\bar{\mu}$ to $\bar{y}$.

Two different cases with this common structure are illustrated in figures 7(a) and 7(b), with the difference being that in the first case $\bar{\mu}$ and $\bar{y}$ are in different trees, which in the second case they are in the same tree, with $\bar{\mu}$ being a predecessor of $\bar{y}$. It is clear in each case, the addition of the edge $(\bar{\mu}, \bar{y})$ creates two and only two ways to price good $\bar{y}$, one corresponds to the initial graph F, the other to F'.

Prune and Graft is characterized by the following rules: $\mu' = \mu$; $R(F') = R(F)$; and $E(F') = \{E(F) \backslash (y', \bar{y})\} \cup (\bar{\mu}, \bar{y})$ where y' is the direct predecessor of $\bar{y}$ in F.

6.4. Cyclic Reversal

The two final cases both involve changes both in the allocation of goods, as well as the nature of the interdependence. Cyclic Reversal occurs when both $\bar{\mu}$ and $\bar{y}$ are elements of the same tree, and $\bar{y}$ is a predecessor of $\bar{\mu}$.

A simple example with three goods and three buyers is illustrated in Figure 8(a). Initially, each buyer is matched with the correspondingly numbered good: $\mu_i = y_i$, and the GA structure involves two directed edges: $E(F) = \{\ (y_1, y_2), (y_2, y_3)\}$. In this example, $\bar{\mu} = y_3$ and $\bar{y} = y_1$, as illustrated by the corresponding light edge. Note that the addition of the directed edge, (y_3, y_1) creates a cycle in which each buyer is indifferent between his assigned good and its direct successor. A cyclic permutation of the allocation therefore keeps each buyer in his demand set. It is clear that we can generate the same price vector as follows. First, allocate $\bar{x}$ to y_3 and use his indifference to price y_3. Next take the buyer allocated to y_2, shift him to y_3 and use his indifference to price y_2. Finally, take the buyer allocated to y_1, shift him to y_2. This GA structure is illustrated in figure 8(b).

Formally, cyclic reversal involves the following steps. Add the edge $(\bar{\mu}, \bar{y})$ to $E(F)$. This creates a directed cycle $C = \{y_1, y_2, \ldots, y_k, y_{k+1}\}$ with $y_1 = y_{k+1} = \bar{y}$, $y_k = \bar{\mu}$, and $(y_i, y_{i+1}) \in E(F) \cup (\bar{\mu}, \bar{y})$ for $i \in \{1, \ldots k-1\}$. The new allocation μ involves rotating all buyers allocated to a good in C to its direct successor

$$\mu'_a = \begin{cases} \mu_a & \text{if } \mu_a \notin C \\ y_{i+1} & \text{if } \mu_a = y_i \in C \end{cases}$$

The new graph F' is derived from the old graph F in a series of steps by first reversing all of the edges in the cycle C by replacing (y_i, y_{i+1}) for all $y_i, y_{i+1} \in C$ with (y_{i+1}, y_i), then eliminating (y_2, y_1) and finally, replacing all (y_i, y_j) with $y_i \in C$ and $y_j \notin C$ with (y_{i+1}, y_j) so that all goods not in P are priced by the same buyers in (μ', F') as in (μ, F).

6.5. Shift and Replant

Shift and Replant occurs when $\bar{y}$ is unallocated. It is similar to cyclic reversal except that instead of relocating buyers around a cycle, we relocate them along a path.

Figure 9(a) provides an example. There are four goods. The first three are allocated to like-numbered buyers. The forth good $\bar{y}$ is unallocated and hence a root good as indicated by the rectangle to which no one is assigned. The initial GA structure (μ, F) involves directed edges (y_1, y_2) and (y_2, y_3). It is buyer $x_3 = \bar{x}$ who then becomes indifferent with the unoccupied good $\bar{y}$. Note that we can support the same price vector, by first shifting the buyers along the connected path between the two root goods, with $\bar{x}$ to $\bar{y}$, x_2 to y_3 and x_1 to y_2, and then using $\bar{x}$'s indifference to price y_3, x_2's to price y_2. y_1 is unallocated and hence priced at reservation. This situation is illustrated in Figure 9(b).

Formally, shift and replant is characterized by the following operations. We first add the edge $(\bar{\mu}, \bar{y})$ and let $P = \{y_1, y_2, \ldots, \bar{\mu}, \bar{y}\}$ denote the path in $E(F) \cup (\bar{\mu}, \bar{y})$ beginning at $y_1 \in R(F)$ and ending with $\bar{y}$. The change in the allocation μ involves shifting all buyers allocated to a good in P to its direct successor:

$$\mu'_a = \begin{cases} \mu_a & \text{if } \mu_a \notin P \\ y_{i+1} & \text{if } \mu_a = y_i \in P \end{cases}$$

The new graph F' is derived from the old graph F. First, reverse all of the edges in the path P by replacing (y_i, y_{i+1}) for all $y_i, y_{i+1} \in P$ with (y_{i+1}, y_i). Then eliminate (y_2, y_1). Finally, replace all (y_i, y_j) with $y_i \in P$ and $y_j \notin P$ with (y_{i+1}, y_j), so that all goods not in P are priced by the same buyers in (μ', F') as in (μ, F).

6.6. The Fundamental Theorem of Calculus for Allocation Markets

The above results allow us to express discrete changes in prices as the integral of infinitesimal market changes. Consider a change for λ_0 to λ_1 along a regular path $\pi(z)$. Let $\hat{Z} = \{z \in [0,1] | \phi(z - \varepsilon) \neq \phi(z + \varepsilon)\}$ denote the set of potential switch points, and label the elements of $\hat{Z} = \{z_1, \ldots z_S\}$. Let (μ_s, F_s) denote the unique element of $\phi(z)$ over the interval (z_s, z_{s+1}), where it is understood that $z_0 = 0$ and $z_{S+1} = 1$. Prior results imply that the change in the price of good y_i may be obtained by integrating the change in price over each subinterval and summing over subintervals:

$$p_i(\lambda_1) - p_i(\lambda_0) = \sum_{s=0}^{S} \int_{z_s}^{z_{s+1}} \left(\left. \frac{dp_i}{d\lambda(z)} \right|_{(\mu_s, F_s)} \frac{d\lambda(z)}{dz} \right) dz$$

Here the derivative $dp_i/d\lambda(z)$ is taken with respect to the current GA structure (μ_s, F_s).

7. An Algorithm for Computing Competitive Equilibria

The maximization in problem (2.2) is over a very large set and may be difficult to perform in practice. A generalization of Cayley's theorem states that for each allocation μ there are

$$\sum_{k=1}^{m} \left(\begin{array}{c} m \\ k \end{array} \right) km^{m-1-k}$$

different forests of rooted trees on the m allocated goods.[7] On top of this there are $n!/(n-m!)$ ways to allocate buyers to goods. If $n = m = 10$, we get more than 8.5×10^{15} different GA-structures. If $n = m = 1000$, we get more than 9×10^{3609}.

We have shown how to use homotopy methods to move between any two competitive equilibria. This insight may be used to compute equilibria and avoid searching through such a large set. We only need a simple equilibrium to start from. One convenient initial equilibrium is the null equilibrium in which the preference parameters are such that each buyer prefers a different good at reservation prices.

The algorithm works in the following way. Given any set of parameters, we initialize this algorithm by raising the parameter $\xi_{a,a}$ for each buyer x_a high enough that each buyer x_a prefers y_a to all other goods when prices are set at reservation levels. At this level of reservation utility, the minimum price equilibrium has all prices at reservation level and all buyers allocated to their preferred goods. There is a unique GA-structure corresponding to this equilibrium. It involves the equilibrium allocation and a null graph in which all goods are root goods and there are no edges. This equilibrium is our starting point. We can then lower the $\xi_{a,a}$ to their original levels, tracking GA-structures that correspond to the minimum price equilibrium. Note that lowering the $\xi_{a,a}$ has the effect of raising prices monotonically to their final level.

This algorithm is related to the ascending auction mechanism of Demange, Gale and Sotomayor (1986). They consider minimum price equilibria in a model with transferable utility

[7] See Aigner and Ziegler [2003, 3rd edition, p. 178].

and discrete prices. Their algorithm involves increasing the prices of all goods in minimal overdemanded sets by one unit until supply and demand are brought into balance. The key complication that non-transferable utility introduces is that the same price change affects the demands of different buyers differently. The challenge is to find a way of raising prices that does not completely alter the balance between supply and demand, while at the same time keeping track of the resulting changes in the allocation. The GA structures provide such a mechanism.

There are several attractive features of this algorithm. First, it terminates by precisely identifying the minimum equilibrium price. This is not the case with approximation methods that are often employed in computing economic equilibria (e.g. Scarf (1973)).[8] Second, the algorithm is likely to be relatively fast in many practical applications. There is a sense in which the algorithm is minimal: it searches only through the set of potential solutions for some set of utility parameters and by-passes the mass of entirely unsuitable price vectors. This mirrors the situation with the simplex method, in which one searches only through the set of extreme points of the feasible set, all of which are optimal for some vector of resources.

7.1. From Continuous to Discrete

Practical implementation of the algorithm requires discretizing z. Figure 10 illustrates one possibility. We begin with a known solution to to the model (μ_1, F_1) at $z_1 = 0$. At each step in the algorithm our current candidate solution (μ_1, F_1) which is known to generate the minimum equilibrium price at some point z_1, $(\mu_1, F_1) \in \Phi(z_1)$, and test whether it generates the minimum equilibrium price at some further point z_2. Initially, we take $z_2 = 1$.

If $(\mu_1, F_1) \in \Phi(z_2)$, this test succeeds and we reset $z_1 = z_2$ and $z_2 = 1$. If $z_1 = 1$, we are finished. Otherwise, we test (μ_1, F_1) at $z_2 = 1$.

We will say that there is a "violation of competitive equilibrium" at z_2 if, given $q(\mu_1, F_1, \pi(z_2))$, a buyer prefers a good to the one that he is allocated or if there is good whose price falls below reservation. The test succeeds if there is no violations. If there is a single violation of competitive equilibrium when applying (μ_1, F_1) at z_2, meaning a single buyer who prefers a single good or a single good whose price falls below reservation, then we construct (μ_2, F_2), which is the alternative GA structure suggested by treating the violation as an indifference and making the appropriate market transition as in the last section, and test (μ_2, F_2) at z_2. If $(\mu_2, F_2) \in \Phi(z_2)$, we update $(\mu_1, F_1) = (\mu_2, F_2)$, as well as $z_1 = z_2$, and $z_2 = 1$ and proceed as before. If there are multiple violations of competitive equilibrium at $q(\mu_1, F_1; \lambda(z_2))$ or if $(\mu_2, F_2) \notin \Phi(z_2)$, then

[8] Miyake (2003) also provides an algorithm with this property.

we stepped over multiple market transitions. We then try taking a smaller step: we update $z_2 = (z_1 + z_2)/2$ and try (μ_1, F_1) again.

Given Theorem 3, which states that there are a finite number of points of transition on a generic path, this algorithm converges. Eventually, the step size is reduced to the point that (z_1, z_2) contains a single market transition and $(\mu_2, F_2) \in \Phi(z_2)$

8. Conclusion

In this paper and it companion, we have introduced a new mathematical apparatus for understanding allocation markets with nontransferable utility. We are currently extending the work to a dynamic context and solving for the reallocation of objects over time. Our methods may also apply in other areas.

References

[1] Caplin, Andrew, and John Leahy (2010), "A Graph Theoretic Approach to Allocation Markets," New York University working paper.

[2] Demange, Gabrielle, and David Gale (1985), "The Strategy Structure of Two-Sided Matching Markets," *Econometrica* 53, 873-888.

[3] Demange, Gabrielle, and David Gale, and Marilda Sotomayor (1986), "Multi-Item Auctions," *The Journal of Political Economy* 94, 863-873.

[4] Judd, Kenneth, (1999) *Numerical Methods in Economics,* Cambridge: MIT Press.

[5] Kaneko, Mamoru, (1982), "The Central Assignment Game and the Assignment Markets, *Journal of Mathematical Economics* 10, 205–232.

[6] Krantz, Steven, and Harold Parks (2002), "A Primer of Real Analytic Functions," Boston: Birhauser.

[7] Kuhn, Harold (1955), "The Hungarian Method for the Assignment Problem," *Naval Research Logistics Quarterly* 2, 83-97.

[8] Miyake, Mitsunobu (2003), "Precise Computation of a Competitive Equilibrium of the Discrete Land Market Model," *Regional Science and Urban Economics* 33, 721-743.

[9] Munkres, James (1957), "Algorithms for the Assignment and Transportation Problems," *Journal of the Society of Industrial and Applied Mathematics* 5, 32-38.

[10] Scarf, Herbert (1973), *The Computation of Economic Equilibria,* New Have: Yale University Press.

9. Appendix: Proofs

Lemma 1: There exists a bijection $\eta : \mathcal{M} \to \Phi$ so that

$$|\Phi| = |\mathcal{M}|$$

Proof: We identify a mapping $\eta : \mathcal{M} \to \Phi$ and show it to be a bijection. Given $\mathcal{E} \in \mathcal{M}$, we identify $\eta(\mathcal{E}) = (\mu(\mathcal{E}), F(\mathcal{E}))$ starting with the allocation. We first match all goods of degree one.

If there are goods remaining in any set $Y_l^{\mathcal{E}}$ after the removal of these goods and their matched buyers, there is as least one of degree one, since at least one edge to the unmatched goods is removed when the prior match is created, due to the connectedness of the components. This process ends by uniquely specifying $\mu(\mathcal{E}) \in M$ in no more than $\max|Y_l^{\mathcal{E}}|$ steps. We construct the graph $F(\mathcal{E}) = (Y, R, E)$ by first setting R equal to the union of $Y_0^{\mathcal{E}}$ and the valence 1 good in each set $Y_l^{\mathcal{E}}$ for $l \geq 1$. We then define $E = \{(y_i, y_k)|(x_a, y_i), (x_a, y_k) \in \mathcal{E}$ and $\mu_a = y_i\}$. Since each non-root good is allocated and has degree two, it has a unique predecessor in E. Since each is connected to the root good by construction, each component is a tree. To establish that $\eta : \mathcal{M} \to \Phi$, note by construction $q(\mu(\mathcal{E}), F(\mathcal{E})) = \underline{p}$, that $(\mu(\mathcal{E}), F(\mathcal{E}))$ identifies a minimum price equilibrium, and that all goods to which a directed edge points are allocated.

It remains to establish that $\eta : \mathcal{M} \to \Phi$ is onto Φ. Consider any $(\tilde{\mu}, \tilde{F}) \in \Phi$. We identify the element $\gamma(\tilde{\mu}, \tilde{F}) = \eta^{-1}(\tilde{\mu}, \tilde{F}) \in \mathcal{M}$ such that $\eta(\tilde{\mathcal{E}}) = (\tilde{\mu}, \tilde{F})$ in stages. First, we generate a bipartite graph on $X \cup Y$ with bipartition (X, Y) by joining each buyer with the good to which they are matched, placing $(x_a, \tilde{\mu}_a) \in \tilde{\mathcal{E}}$, and then placing $(x_a, y_b) \in \tilde{\mathcal{E}}$ if $(\tilde{\mu}_a, y_b) \in F$. To see that $\gamma(\tilde{\mu}, \tilde{F}) \in \mathcal{M}$, note first that $q(\tilde{\mu}, \tilde{F}) = \underline{p}$, hence this graph is a subset of $\mathcal{E}^*$. To see that it satisfies the remaining aspects of condition M, look at each tree in the forest and note that there every node of the tree has one and only one buyer due to the insistence that downstream as well as upstream goods are allocated. Each such tree in the forest thereby is mapped to a connected set $(X_l^{\mathcal{E}}, Y_l^{\mathcal{E}})$ for $1 \leq l \leq L$ that is disconnected from all others, and in which the equal cardinality condition M2 is satisfied. The isolated set $Y_0^{\mathcal{E}}$ comprises all goods that are not in demand by anyone. By construction, there is one and only one good of degree 1 in each partition set $Y_l^{\mathcal{E}}$ for $l \geq 1$, and all others have degree 2. Moreover, by construction all goods of valence 1 are root goods for which $q_j^* = r_j$. By construction, goods of degree 0 also satisfy $q_j^* = r_j$. Hence $\underline{p}_i > r_i$ implies that the degree of y_i is 2. This confirms that $\gamma(\tilde{\mu}, \tilde{F}) \in \mathcal{M}$. Finally, it is direct from the construction that $\eta(\tilde{\mathcal{E}}) = (\tilde{\mu}, \tilde{F})$, establishing that indeed $\gamma : \Phi \to \mathcal{M}$ is the inverse function of $\eta : \mathcal{M} \to \Phi$.□

Corollary 1: $\mathcal{E}^* \in \mathcal{M}$ and $d^*(y_i) = 2 \Longrightarrow \underline{p}_i > r_i$ then $|\Phi| = 1$.

Proof: Suppose that $\mathcal{E}^* \in \mathcal{M}$ and all y_i with $d^*(y_i) = 2$ have $\underline{p}_i > r_i$. We show that $\mathcal{M}$ has one element. Lemma 1 then implies $|\Phi| = 1$. That there is no second element follows from the observation that since $\mathcal{E}^* \in \mathcal{M}$ all goods have degree less than or equal to two. Removing any edge from the demand graph D reduces the degree of some good. If the degree falls from one to zero, then the number of goods with degree greater than zero is less than the number of buyers in contradiction of condition M4. If a good has degree two in D, then by assumption its price is above reservation. Reducing the degree of such a good contradicts condition M5.□

Theorem 1: $|\Phi| = 2$ if and only if there exists $(\bar{x}, \bar{y}) \in \mathcal{E}^*$ such that $\mathcal{E}^*/(\bar{x}, \bar{y}) \in \mathcal{U}$.

Proof: (If) Suppose that there exists $(\bar{x}, \bar{y}) \in \mathcal{E}^*$ such that $\mathcal{E}^*/(\bar{x}, \bar{y}) \in \mathcal{U}$. To establish $|\Phi| = 2$, let $\mathcal{E}_1 = \mathcal{E}^*/(\bar{x}, \bar{y})$. We consider a number of sub-cases and in each such case identify $\mathcal{E}_2 \neq \mathcal{E}_1$ with $\mathcal{E}_2 \in \mathcal{M}$. We then show that there is no third member $\mathcal{E}_3 \in \mathcal{M}$. The argument depends on features of $\bar{\mu} \in Y$, the good matched to $\bar{x}$ in the unique match associated with $\mathcal{E}_1 \in \mathcal{U}$.

1. If $\bar{y} \in Y_0^{\mathcal{E}_1}$, let $\bar{\sigma} \in Y$ be the root good in the set $Y_l^{\mathcal{E}_1}$ such that $\bar{x} \in X_l^{\mathcal{E}_1}$, and let $\mu^{-1}(\bar{\sigma}) \in X_l^{\mathcal{E}_1}$ be the buyer matched to that good. Define $\mathcal{E}_2 = \mathcal{E}^*/\{\mu^{-1}(\bar{\sigma}), \bar{\sigma}\}$. Note that $\mathcal{E}_2 \in \mathcal{U}$ by construction, since the only change is the replacement of $\bar{\sigma} \in Y_l^{\mathcal{E}_1}$ with the good $\bar{y} \in Y_l^{\mathcal{E}_2}$, also of valence 1. Now suppose that there exists some third element $\mathcal{E}_3 \in \mathcal{M}$, with $\mathcal{E}_3 \notin \{\mathcal{E}_1, \mathcal{E}_2\}$. Note that this set must contain both $\{\mu^{-1}(\bar{\sigma}), \bar{\sigma}\}$ and $(\bar{x}, \bar{y})$ given that uniqueness holds with either element excluded. Hence some other edge from those joining $X_l^{\mathcal{E}_1}$ and $Y_l^{\mathcal{E}_1}$ must be removed, thereby reducing the degree of some element other than $\bar{\sigma}$ or $\bar{y}$ either to 1 or to zero. This contradicts the fact that this good must have $\underline{p}_i > r_i$, which implies that its degree in $\mathcal{E}_2$ must be 2 by M4.

2. Suppose $(\bar{\mu}, \bar{y}) \subset Y_l^{\mathcal{E}_1}$ for some $l \geq 1$. Note in this case that $\mathcal{E}^*$ itself has a cycle, since it $X_l^{\mathcal{E}_1} \cup Y_l^{\mathcal{E}_1}$ has as many edges as vertices in $\mathcal{E}^*$: this cycle is unique, since all cycles must contain $\{\bar{x}, \bar{y}\}$, so that if there were two of them, then there would have been one prior to the addition of $\{\bar{x}, \bar{y}\}$ (geometrically clear). Now define $\{x^C, \bar{y}\}$ as the unique edge other than $(\bar{x}, \bar{y})$ heading to $\bar{y}$ in this unique cycle in the graph $(X \cup Y, \mathcal{E}^*)$. Define $\mathcal{E}_2 = \mathcal{E}^*/\{x^C, \bar{y}\}$. By construction, $\mathcal{E}_2 \neq \mathcal{E}_1$ satisfies $\mathcal{E}_2 \in \mathcal{U}$, since all goods have the same valence in $\mathcal{E}_2$ and $\mathcal{E}_1$, and the associated partitions are identical. Now suppose that there exists some third element $\mathcal{E}_3 \in \mathcal{M}$, with $\mathcal{E}_3 \notin \{\mathcal{E}_1, \mathcal{E}_2\}$. This set must contain both $\{x^C, \bar{y}\}$ and $(\bar{x}, \bar{y})$ given that uniqueness holds with either element excluded. Hence some other edge in the unique cycle must be removed. In so doing, one reduces the degree of some good other than $\bar{y} \in Y_l^{\mathcal{E}_1}$ below 2. This again contradicts the fact that this good must have $p_i > r_i$, which implies that its degree in $\mathcal{E}_2$ must be 2 by M4.

3. Suppose $\bar{\mu} \in Y_l^{\mathcal{E}_1}$ and $\bar{y} \in Y_m^{\mathcal{E}_1}$ with $m \notin \{0, l\}$. There are two sub-cases depending on the valence of $\bar{y}$.

1. If $\bar{y}$ is of valence 2 in $\mathcal{E}_1$ then define $\mathcal{E}_2 = \mathcal{E}^*/\{x^N, \bar{y}\}$, where x^N is the element in $Y_l^{\mathcal{E}}$ that is not matched to $\bar{y}$ in unique match $\bar{\mu}(\mathcal{E}_1)$. Note that $\mathcal{E}_2 \in \mathcal{U}$, since one now has two sets with root goods of valence 1 at reservation and all others of valence 2, in line with the uniqueness condition. Again any third element $\mathcal{E}_3 \in \mathcal{M}$, with $\mathcal{E}_3 \notin \{\mathcal{E}_1, \mathcal{E}_2\}$ must include both $\{x^N, \bar{y}\}$ and $(\bar{x}, \bar{y})$, but exclude some other element of the edge set corresponding to $\left(X_l^{\mathcal{E}} \cup X_m^{\mathcal{E}}, Y_l^{\mathcal{E}} \cup Y_m^{\mathcal{E}}\right)$. However this happens, it calls for there to be at least two elements of valence 1 since the resulting set is disconnected, but there is by assumption only one good at reservation price in the set, a contradiction.
2. If $\bar{y}$ is of valence 1 in $\mathcal{E}$, then $\mathcal{E}^* \in \mathcal{M}$ since the joint set $Y_l^{\mathcal{E}} \cup Y_m^{\mathcal{E}}$ has all but one good of valence 1, which good is at reservation price. Now suppose that there exists some third element $\mathcal{E}_3 \in \mathcal{M}$, with $\mathcal{E}_3 \notin \{\mathcal{E}_1, \mathcal{E}^*\}$. Since $\mathcal{E}^*/(\bar{x}, \bar{y})$ satisfies the uniqueness condition, this set must contain $(\bar{x}, \bar{y})$ and exclude something from the edge set corresponding to $\left(X_l^{\mathcal{E}} \cup X_m^{\mathcal{E}}, Y_l^{\mathcal{E}} \cup Y_m^{\mathcal{E}}\right)$. As before, it will create some good with $\underline{p}_i > r_i$ yet with degree strictly, a contradiction.

4. An adaptation of case 3(b) works in case 2, with $\mathcal{E}^* \in \mathcal{M}$ and existence of one and only one $y_j \in Y$ such that $d^*(y_j) = 2$ and $\underline{p}_j = r_j$. Since $\mathcal{E}^* \in \mathcal{M}$, there is a unique allocation $\mu(\mathcal{E}^*)$ consistent with $\mathcal{E}^*$. Identify the unique element $x_j^N \in X$ that is to linked y_j but not matched to it according to this allocation . Note that $\mathcal{E}^*/\{x_j^N, y_j\} = \mathcal{E}_1 \in \mathcal{U}$, since y_j has valence 1 and is a root good with the removal of its second edge, with all of the matches preserved. Now suppose that there exists some third element $\mathcal{E}_3 \in \mathcal{M}$, with $\mathcal{E}_3 \notin \{\mathcal{E}_1, \mathcal{E}^*\}$. This set must contain $\{x_j^N, y_j\}$ given that uniqueness holds with it excluded, and it cannot be the whole set, since this is $\mathcal{E}^*$. Hence some other edge must be removed from the partition element $X_l^{\mathcal{E}^*} \cup Y_l^{\mathcal{E}^*}$ that contains y_j. If this is of degree 1 , then there is one too few good: if it is any of the degree 2 goods, then one reduces the degree of some good with $\underline{p}_i > r_i$ below 2, a contradiction.

(Only If) If $|\Phi| = 2$ then certainly $\mathcal{E}^*$ has some subset that satisfies the uniqueness condition of corollary 1. It can have at most one more edge, since if it had two, then we know from above proof how to create at least three members of $\mathcal{M}$. If it has no more edges than $\mathcal{E}^*$, then it must contain at least one more good at reservation price than in that minimum, and we know that it cannot contain more than that, since then the proof above shows can find at least one for each additional good at reservation. This completes the proof.□

Lemma 2: With Assumptions A and B, each $q_i(\mu, F, \lambda)$ is analytic in $\lambda \in \Lambda$; $\Phi(\lambda)$ is non-empty, compact-valued, and upper-hemicontinuous; and each $\underline{p}_i(\lambda)$ is continuous.

Proof: We prove the $q_i(\mu, F, \lambda(z))$ are analytic by induction. Consider first $y_i \in R(F) \equiv A_0$. $g_i = r_i$ which is trivially analytic in λ. Now suppose that for all $y_i \in A_n$, g_i are analytic functions of λ. Consider S, the set of direct successors to A_n. $S = \{y_i \notin A_n | (y_j, y_i) \in E(F)$ for some $y_j \in A_n\}$. Let $\mu_a = y_j$. Now q_j is defined implicitly by the indifference of x_a:

$$U_a(y_i, q_i, \xi_{ai}) = U_a(y_j, q_j, \xi_{aj})$$

where according to Assumption B, $U_a(y_j, q_j, \xi_{ai})$ and $U_a(y_i, q_i, \xi_{aj})$ are analytic functions of their second and third arguments and strictly monotonic in the third. Assumption A ensures that U_a is strictly monotonic in the second. It follows from the Real Analytic Implicit Function Theorem (Krantz and Parks, 2002, p. 35) that q_i is an analytic function of λ. This completes the induction step.

That $\Phi(\lambda)$ is non-empty, compact-valued, and upper-hemicontinuous; and each $\underline{p}_i(\lambda)$ is continuous follows from Theorem 3 in CL.□

The following Lemmas are used in the proof of Theorem 3. The proof of Theorem 3 is divided between three propositions which together establish that Λ_R is open and dense in Λ_S.

Lemma 3: For each y_i, $q_i(\mu, F, \lambda(z))$ is real analytic on $z \in (0, 1)$.

Proof: According to Lemma 2, Assumptions A and B imply that $q_i(\mu, F, \lambda)$ is real analytic on $\lambda \in \Lambda$. Assumption C states that λ is real analytic on $z \in (0, 1)$. That the $q_i(\mu, F, \lambda(z))$ are analytic follows from the observation that compositions of analytic functions are analytic (Kranz and Parks, p.19).□

Lemma 4: Given any three GA structures, (μ_1, F_1), (μ_2, F_2), and (μ_3, F_3), let $\tilde{\Lambda}$ denote the set λ_0 such that $|\phi(z; \lambda_0, S) \cap \{(\mu_1, F_1), (\mu_2, F_2), (\mu_3, F_3)\}| \leq 2$ for all z. $\tilde{\Lambda}$ is open.

Proof: Consider $\bar{\lambda}_0 \in \tilde{\Lambda}$ and let $\Lambda^B = \Lambda\{\lambda \in \Lambda | \{(\mu_1, F_1), (\mu_2, F_2), (\mu_3, F_3)\} \in \Phi(\lambda)\}$. The upper hemicontinuity of Φ implies then implies Λ^B is closed. Since Λ^B is bounded, it is compact.

Given $\lambda, \lambda' \in \Lambda$, let $d(\lambda, \lambda')$ denote the Euclidian distance between λ and λ' and let $d(z) = \min_{\lambda \in \Lambda^B} d(\pi(z; \bar{\lambda}_0, S), \lambda)$. Λ^B is compact so the minimum problem is well defined. $d(\lambda, \lambda')$ is continuous in both arguments, so the Theorem of the Maximum states that $d(z)$ is continuous.

Since $z \in [0,1]$, $d(z)$ has a minimal value. Let $\delta = \min d(z)$. Since the path π and the set Λ^B are disjoint $\delta > 0$. It follows that a ball around $\bar{\lambda}_0$ with radius $\delta/2$ is contained in $\tilde{\Lambda}$. This establishes that $\tilde{\Lambda}$ is open.□

Lemma 5: Given $\lambda \in \Lambda$ such that $|\Phi(\lambda)| > 2$, there exists $\{(\mu^1, F^1), (\mu^2, F^2)(\mu^3, F^3)\} \subseteq \Phi(\lambda)$ and $l \in \mathbb{R}^{mn+n}$ with $d(l, 0) = 1$ such that given any $\chi > 0$ and $\delta = \chi l$ such that $\lambda + \delta \in \Lambda$, $\{(\mu^1, F^1), (\mu^2, F^2)\} \subseteq \Phi(\lambda + \delta)$ and $(\mu^3, F^3) \notin \Phi(\lambda + \delta)$.

Proof: Fix $\lambda \in \Lambda$ and such $|\Phi(\lambda)| > 2$. Let $(\mu^1, F^1) \in \Phi(\lambda)$. Since $|\Phi(\lambda)| > 2$, the arguments of Theorem 1 can be extended to show that there must be either to goods $y_1$and y_2 such that $y_j \notin R(F^1)$ and $q_j(\mu^1, F^1, \lambda) = r_j(\lambda)$ for $j = \{1, 2\}$; or two sets of indifference (x_a, y_1) and (x_a, y_2) such that $\{(x_a, y_1), (x_b, y_2)\} \in \mathcal{E}^*$, y_1 is not equal to μ^1_a or any of its direct successors in F^1, y_2 is not equal to μ^1_b or any of its direct successors in F^1; or one of each.

(1) Construct (μ^2, F^2) as equal to (μ^1, F^1) except that $y_1 \in R(F^2)$. Similarly, (μ^3, F^3) has $y_1, y_2 \in R(F^2)$. Now let $(y_k, y_1) \in E(F^1)$ and $y_k = \mu^1(x_a)$. Let δ be a vector of zeros with a value $\hat{\delta} < 0$ in the place associated with ξ_{a1}. $q(\mu^2, F^2, \lambda + \delta) = q(\mu^2, F^2, \lambda + \delta) = q(\mu^2, F^2, \lambda)$. $q_1(\mu^1, F^1, \lambda + \delta) < r_1(\mu^1, F^1, \lambda + \delta)$. It follows that $(\mu^1, F^1) \notin \Phi(\lambda + \delta)$. Need $\{(\mu^1, F^1), (\mu^3, F^3)\} \in \Phi(\lambda + \delta)$.

(2) Construct (μ^2, F^2) as the other structure given $(x_a, y_1) \in \mathcal{E}^*$. Similarly (μ^3, F^3) as the other structure given $(x_b, y_2) \in \mathcal{E}^*$. Note that y_2 is not equal to μ^2_a or any of its direct successors in F^2. Let δ be a vector of zeros with a value $\hat{\delta} < 0$ in the place associated with ξ_{b2}. $(\mu^3, F^3) \notin \Phi(\lambda + \delta)$. $q(\mu^1, F^1, \lambda + \delta) = q(\mu^2, F^2, \lambda + \delta) = q(\mu^1, F^1, \lambda)$.

(3) Construct (μ^2, F^2) as equal to (μ^1, F^1) except that $y_1 \in R(F^2)$, and construct (μ^3, F^3) as the other structure given $(x_a, y_1) \in \mathcal{E}^*$. Let δ be a vector of zeros with a value $\hat{\delta} < 0$ in the place associated with ξ_{a1}. $(\mu^3, F^3) \notin \Phi(\lambda + \delta)$. $q(\mu^1, F^1, \lambda + \delta) = q(\mu^2, F^2, \lambda + \delta) = q(\mu^1, F^1, \lambda)$.

The proof is completed by the observation that the only property of δ used in the proof is its direction.□

Proposition (Theorem 3, Part 1): Λ_R is open.

Proof: The proof is in two parts. We first prove that the set of initial conditions for which $|\phi(z)| \leq 2$ is open. We then prove that the set of initial conditions for which Z is finite is open. The intersection of two open sets is open.

(1) *The set of* λ_0 *for which* $|\phi(z; \lambda_0, S)| \leq 2$ *is open.*

Given S, let Λ_1 be the set of initial conditions λ_0 such that $|\phi(z;\lambda_0,S)| \leq 2$ for all $z \in [0,1]$. Let $\tilde{\Lambda}[(\mu_1,F_1),(\mu_2,F_2),(\mu_3,F_3)]$ denote the set λ_0 such that there exists z such that $|\phi(z;\lambda_0,S)| \cap \{(\mu_1,F_1),(\mu_2,F_2),(\mu_3,F_3)\} \leq 2$. $\Lambda_1 = \cap\tilde{\Lambda}[(\mu_1,F_1),(\mu_2,F_2),(\mu_3,F_3)]$ where the intersection is over all triplets of GA structures. Each $\tilde{\Lambda}[(\mu_1,F_1),(\mu_2,F_2),(\mu_3,F_3)]$ is open by Lemma 4. Since finite intersections of open sets are open, Λ_1 is open.

(2) *The set of λ_0 for which Z is finite is open.*

Given λ_0, let $Z = \{z \in [0,\bar{z}]||\phi(z;\lambda_0,S)| > 1\}$. Given S, let Λ_2 be the set of λ_0 such that Z is finite. Consider $\bar{\lambda}_0 \in \Lambda_2$. The proof is by contradiction. Given $\varepsilon > 0$, let $B_\varepsilon(\lambda) = \{\lambda'|d(\lambda,\lambda') < \varepsilon\}$. Suppose that for all $\varepsilon > 0$, there exists $\lambda' \in B_\varepsilon(\bar{\lambda}_0)$ such that $Z(\lambda')$ is infinite. Consider a sequence $\{\varepsilon_n\}$ converging to zero and consider a sequence λ_n such that $\lambda_n \in B_\varepsilon(\bar{\lambda}_0)$ and $Z = \{z \in [0,\bar{z}]||\phi(z;\lambda_n,S)| > 1\}$ is infinite. Given that there are a finite number of GA structures, there exists a subsequence $\{\lambda_m\} \subseteq \{\lambda_n\}$ and two GA structures, (μ_1,F_1) and (μ_2,F_2), such that $q(\mu_1,F_1,\pi(z;\lambda_m,S)) = q(\mu_2,F_2,\pi(z;\lambda_m,S))$ at infinitely many points $z \in (0,1)$. By Lemma 3, $q_i(\mu_1,F_1,\pi(z;\lambda_m,S))$ and $q_i(\mu_2,F_2,\pi(z;\lambda_m,S))$ are real analytic on $z \in (0,1)$. Given $a < 0$ and $b > 1$, let $\hat{S}(z) : (a,b) \to \mathbb{R}^{mn+n}$ be the (component by component) analytic continuation of S. Given the continuity of S, we can choose a, b, and ε_m small enough that $\lambda_m + \hat{S}(z)$ are contained in Λ for all $z \in (a,b)$. Let $\hat{q}_i(\mu,F,z,m) = q_i(\mu,F,\lambda_m + \hat{S}(z))$. Standard arguments establish that and $\hat{q}_i(\mu,F,z,m)$ is analytic and $\hat{q}_i(\mu,F,z,m) = q_i(\mu,F,\pi(z;\lambda_m,S))$ on $[0,1]$. It follows that $\hat{q}_i(\mu_1,F_1,z,m) = \hat{q}_i(\mu_2,F_2,z,m)$ at an infinite number of points in $[0,1]$. Let U denote the set of intersections. It has an accumulation point in $[0,1] \subset (a,b)$. Corollary 1.2.7 in Krantz and Parks implies that $\hat{q}_i(\mu_1,F_1,z,m) = \hat{q}_i(\mu_2,F_2,z,m)$ for all $z = (a,b)$. It follows that $q(\mu_1,F_1,\pi(z;\lambda_m,S)) = q(\mu_2,F_2,\pi(z;\lambda_m,S))$ for $z \in [0,1]$. The continuity of the $q(\mu,F)$ then implies that $q(\mu_1,F_1,\pi(z;\bar{\lambda}_0,S)) = q(\mu_2,F_2,\pi(z;\bar{\lambda}_0,S))$. But $\bar{\lambda}_0 \in \Lambda_2$. This contradiction establishes that Λ_2 is open.

$\Lambda_R = \Lambda_1 \cap \Lambda_2$ each of which is open. This completes the proof.□

Proposition (Theorem 3, Part 2) Let $Z^1(\lambda_0) = \{z \in [0,\bar{z}]||\phi(z)| > 1\}$ and let $\Lambda_F = \{\lambda_0|Z(\lambda_0) \text{ is finite}\}$. With A-C, Λ_F is dense in Λ_S.

Proof: Given λ_0, let $Z = \{z \in [0,\bar{z}]||\phi(z;\lambda_0,S)| > 1\}$. Let $\Lambda^I \in \Lambda_S$ denote the set of $\lambda_0 \in \Lambda_S$ such that Z is infinite. We show that given any $\lambda_0 \in \Lambda^I$, we can perturb λ_0 and obtain a path for which Z is finite.

We first prove that for each y_i, $\underline{p}_i(z)$ is real analytic on $(0,1)$ except at a finite number of points. The proof is by contradiction. Let $A = \{z \in (0,1)|\underline{p}_i(\lambda(z)) \text{ is not analytic at } z\}$. Suppose

$|A| = \infty$. Since $|A| = \infty$, there exists an accumulation point $z_1 \in [0,1]$.[9] Consider an arbitrary $z_1 \in A$. There exists (μ_1, F_1) such that $\underline{p}_i(\lambda(z_1)) = q_i(\mu_1, F_1, \lambda(z_1))$. Moreover, given any neighborhood Ω of z_1, if $\underline{p}_i(\lambda(z)) = q_i(\mu_1, F_1, \lambda(z))$ for all $z \in \Omega$ then $\underline{p}_i$ is analytic at z. It follows that there exists there a $z_2 \in \Omega$ such that $\underline{p}_i(\lambda(z_2)) \neq q_i(\mu_1, F_1, \lambda(z_2))$. Given the continuity of the $q_i(\mu, F, \lambda(z))$, this implies that there are two GA structures, (μ_1, F_1) and (μ_2, F_2), such that $q_i(\mu_1, F_1, \lambda(z)) = q_i(\mu_2, F_2, \lambda(z))$ at some point in Ω. As z_1 is arbitrary and A is finite, there exist an infinite number of these intersections. Given that there are a finite number of GA structures, there are two GA structures, (μ_3, F_3) and (μ_4, F_4) such that $q_i(\mu_3, F_3, \lambda(z)) = q_i(\mu_4, F_4, \lambda(z))$ at infinitely many points and $q_i(\mu_3, F_3, \lambda(z)) \neq q_i(\mu_4, F_4, \lambda(z))$ at infinitely many points. Given $q_i(\mu_3, F_3, \lambda(z))$ and $q_i(\mu_4, F_4, \lambda(z))$ are both analytic, $q_i(\mu_3, F_3, \lambda(z)) = q_i(\mu_4, F_4, \lambda(z))$ at infinitely many points implies $q_i(\mu_3, F_3, \lambda(z)) = q_i(\mu_4, F_4, \lambda(z))$ at all $z \in (0,1)$. This contradiction establishes that A is finite.

We return to the proof of the proposition. Suppose $\lambda_0 \in \Lambda^I$. Since $\underline{p}$ is piecewise analytic, we can divide $[0,1]$ into a finite number of sub-intervals $\{[0, z_1), (z_1, z_2), \ldots (z_T, 1]\}$ such that $\underline{p}$ is analytic on each subinterval.[10] Note that on each interval (z_i, z_{i+1}), $\phi(z)$ is constant except at a finite number of points, for if there is any GA structure (μ_1, F_1) such that $\underline{p}_i(\lambda(z)) = q_i(\mu_1, F_1, \lambda(z))$ at an infinite number of points in (z_1, z_2), then either $\phi(z) = \{(\mu_1, F_1)\}$ or $(\mu_2, F_2) \in \phi(z)$ and $q_i(\mu_1, F_1, \lambda(z)) = q_i(\mu_2, F_2, \lambda(z))$ so that $(\mu_1, F_1) \in \phi(z)$ as well since both are analytic on $[0,1]$.

Let c_i denote the number of elements of $\phi(z)$ on interval (z_i, z_{i+1}) except possibly at a finite number of points. Given an interval in which $c_i > 1$, we show that we can perturb $\bar{\lambda}_0$ and reduce c_i by one without raising any other c_j. Given that there are finitely many intervals with values of $c_i > 1$, a finite number of such perturbations will reduce Z to a finite set.

Suppose $c_i > 1$. We consider two cases. First, suppose that $q_i(\mu_1, F_1, \lambda(z)) > r_i(\lambda(z))$ for $y_i \notin R(F_1)$ and $z \in (z_i, z_{i+1})$. According to Corollary 1, $|\Phi(z)| > 1$ if and only if there exists x_b and $y_0 \in D_b(q(\mu_1, F_1, \lambda)$ such that $y_0 \neq \mu_b$ and $(\mu_b, y_0) \notin E(F_1)$. Consider a small perturbation of λ_0 in which we reduce the element associated with ξ_{b0} by ε_{b0} for all such y_0. Call this perturbation λ_0' and let $\phi'(z)$ denote the new path originating from λ_0'. We choose the ε_{b0} small enough that $\lambda_0' \in \Lambda_S$. It follows from Corollary 1, that $(\mu_2, F_2) \notin \Phi(\lambda_1(z))$ for $z \in (z_i, z_{i+1})$. Note that since $q_i(\mu_1, F_1, \lambda(z)) = q_i(\mu_2, F_2, \lambda(z))$ on $[0,1]$, $(\mu_2, F_2) \notin \Phi(\lambda_1(z))$ for all z. The perturbation does not raise $\phi(z)$ at any other point.

[9]If $z_1 = \{0,1\}$ we will need to extend all functions as in the previous proposition to $(a,b) \supset [0,1]$ so that the accumulation point lies in an open set.

[10]Note that if needed we can imbed the closed interval in an open interval and extend all functions to the open interval.

Second, suppose that $q_i(\mu_1, F_1, \lambda(z)) = r_i(\lambda(z))$ for some $y_i \notin R(F_1)$. In this case we consider a perturbation in which we increase the component of λ_0 associated with r_i. Call this perturbation λ_0' and let $\phi'(z)$ denote the new path originating from λ_0'. It is clear that $(\mu_1, F_1) \notin \Phi(\lambda(z))$ for $z \in (z_i, z_{i+1})$. Similarly the perturbation does not raise ϕ at any other point.

This completes the proof of the proposition.□

Proposition (Theorem 3, Part 3) With A-C, Λ_R is dense in Λ_F.

Proof: Let $\pi(z;\lambda_0)$ denote the path beginning at λ_0, let $\phi(z;\lambda_0)$ denote the GAME correspondence along this path, and let $Z^2(\lambda_0) = \{z \in [0,1] || \phi(z;\lambda_0)| > 2\}$ denote the set of points along the path at which $\phi(z;\lambda_0)$ takes more than two values.

To show that Λ_R is dense in Λ_F we show that given an arbitrary $\bar{\lambda}_0 \in \Lambda_F$ such that $Z^2(\bar{\lambda}_0)$ is not empty, we can construct a perturbation of $\bar{\lambda}_0$, $\bar{\lambda}_0' \in \Lambda_F$, such that $Z^2(\bar{\lambda}_0')$ is empty.

The construction is inductive. Let Ψ denote the set possible three pairs of

$$\Psi = \{\{(\mu_1, F_1), (\mu_2, F_2), (\mu_3, F_3)\} | \{(\mu_1, F_1), (\mu_2, F_2), (\mu_3, F_3)\} \in \mathcal{G}\}$$

Since $|\mathcal{G}|$ is finite, $|\Psi| = N < \infty$. At each step we choose an initial condition $\bar{\lambda}_0^n \in \Lambda_F$ and a radius $\delta^n > 0$, such that (1) for some ψ_n not previously considered the set $\{z \in [0,1] \mid \psi_n \subset \phi(z;\lambda_0)\}$ is empty for all $\lambda_0 \in B_{\delta^n}(\bar{\lambda}_0^n)$; and (2) $B_{\delta^n}(\bar{\lambda}_0^n) \subseteq B_{\delta^{n-1}}(\bar{\lambda}_0^{n-1})$, where $B_\delta(\lambda_0) = \{\lambda | d(\lambda, \lambda_0) < \delta\}$. The second condition ensures that the sets $\{z \in [0,1] \mid \psi_i \subset \phi(z;\lambda_0)\}$ is empty for all $\lambda_0 \in B_{\delta^n}(\bar{\lambda}_0^n)$ and $i < n$. The iteration stops when $Z^2(\bar{\lambda}_0^n)$ is empty. This must happen in fewer than N steps.

Initially, we choose $\bar{\lambda}_0^0 = \bar{\lambda}_0$ and $\delta^0 > 0$ such that $B_{\delta^0}(\bar{\lambda}_0^0) \subset \Lambda_S$. This choice is possible given that Λ_S is open. It will be useful below to choose δ^0 such that the closure of $B_{\delta^0}(\bar{\lambda}_0^0)$, $\bar{B}_{\delta^0}(\bar{\lambda}_0^0)$, is contained in Λ_S. Let $X = \{\lambda | \lambda = \pi(z, \lambda_0) \text{ for some } \lambda_0 \in \bar{B}_{\delta^0}(\bar{\lambda}_0^0)\}$. $X \subset \Lambda$ is compact since Λ is bounded. It follows that the function $f : X \to \mathbb{R}$ such that $f(\mu, F, \lambda) = \sum_i g_i(\mu, F, \lambda)$, being a continuous function from a compact metric space X to a metric space $\mathbb{R}$, is uniformly continuous on X.

The induction step begins with $\delta^{n-1} > 0$, $\bar{\lambda}_0^{n-1} \in \Lambda_F$, such that $B_{\delta^n}(\bar{\lambda}_0^{n-1}) \subseteq B_{\delta^{n-1}}(\bar{\lambda}_0^{n-2}) \subseteq B_{\delta^0}(\bar{\lambda}_0^0)$ and $Z^2(\bar{\lambda}_0^{n-1})$ is not empty.

Given $Z^2(\bar{\lambda}_0^{n-1})$ is not empty, there exists z_1 such that $|\phi(z_1;\bar{\lambda}_0^{n-1})| > 2$. Let $\hat{\lambda} = \pi(z_1;\bar{\lambda}_0^n)$. It follows from Lemma 5 and the convexity of Λ that there exists $\{(\mu^1, F^1), (\mu^2, F^2)(\mu^3, F^3)\} \equiv \psi^n \subseteq \Phi(\hat{\lambda})$ and $l \in \mathbb{R}^{mn+n}$ such that $d(l, 0) = 1$ such that for $\theta \in (0, \delta^{n-1})$, $\{(\mu^1, F^1), (\mu^2, F^2)\} \subseteq \Phi(\hat{\lambda} + \theta l)$ and $(\mu^3, F^3) \notin \Phi(\hat{\lambda} + \theta l)$. .

Any $\theta < \delta^{n-1}$ ensures that $\hat{\lambda} + \theta l \in B_{\delta^{n-1}}(\bar{\lambda}_0^{n-1})$. We must be careful, however, that we do not shift λ_0 so far that we create new points at which $\psi^n \subseteq \Phi(\pi(z, \lambda_0)$. To this end, let $\hat{Z} = \{z | (\mu^3, F^3) \notin \phi(z; \bar{\lambda}_0^{n-1})$ and there exist $\{(\mu^4, F^4), (\mu^5, F^5)\} \subseteq \phi(z; \bar{\lambda}_0^{n-1})\}$ and let $\hat{\varepsilon} = \min_{z \in \hat{Z}} \left| f(\mu^3, F^3, \pi(z; \bar{\lambda}_0^{n-1})) - f(\mu^1, F^1, \pi(z; \bar{\lambda}_0^{n-1})) \right|$. $\bar{\lambda}_0^{n-1} \in \Lambda_F$ implies $Z^1(\bar{\lambda}_0^{n-1})$ is finite. $\hat{Z} \subseteq Z^1(\bar{\lambda}_0^{n-1})$ implies $\hat{Z}$ is finite. $\hat{Z}$ finite implies that $\hat{\varepsilon} > 0$. Since $f(\mu^3, F^3, \lambda)$ is uniformly continuous on X, there exists δ_1 such that $d(\lambda, \lambda') < \delta_1$ implies $d(f(\mu^3, F^3, \lambda), f(\mu^3, F^3, \lambda')) < \hat{\varepsilon}/2$. Since all paths $\pi(z, \lambda_0)$ have the same shape, $d(\lambda, \lambda') < \delta_1$ implies $d(f(\mu^3, F^3, \pi(z, \bar{\lambda}_0^{n-1})), f(\mu^3, F^3, \pi(z, \lambda_0)) < \hat{\varepsilon}/2$ for all $z \in [0, 1]$. If we choose $\theta > 0$ such that $\theta < \min\{\delta^{n-1}, \delta_1\}$, then $\{z \in [0, 1] \mid \psi^n \subset \phi(z; \bar{\lambda}_0^{n-1} + \theta l)\}$ is empty. We fix $\theta \in (0, \min\{\delta^{n-1}, \delta_1\})$.

Lemma 4 states that the set of λ_0 such that $|\phi(z; \lambda_0, S) \cap \{(\mu_1, F_1), (\mu_2, F_2), (\mu_3, F_3)\}| \leq 2$ for all $z \in [0, 1]$ is open. Hence there exists a neighborhood of $\bar{\lambda}_0^{n-1} + \theta l$, Ω, such that $\{z \in [0, 1] \mid \psi^n \subset \phi(z; \bar{\lambda}_0^{n-1} + \theta l)\}$ is empty for all $\lambda_0 \in \Omega$ and $\Omega \subset B_{\delta^{n-1}}(\bar{\lambda}_0^{n-1})$. Given Assumptions A-C, the previous proposition states that Λ_F is dense in Λ_S and hence there exists $\lambda_0' \in \Omega \cap \Lambda_F$. We set $\bar{\lambda}_0^n = \lambda_0'$. We choose δ^n such that $B_{\delta^n}(\bar{\lambda}_0^n) \subseteq \Omega \subseteq B_{\delta^{n-1}}(\bar{\lambda}_0^{n-1})$. This complete the induction step and the proof.$\square$

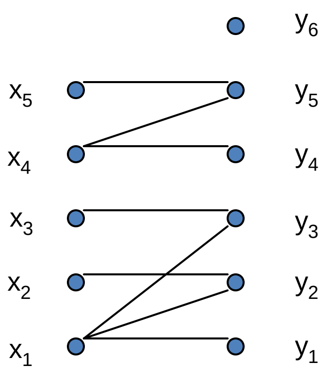

Figure 1(a)

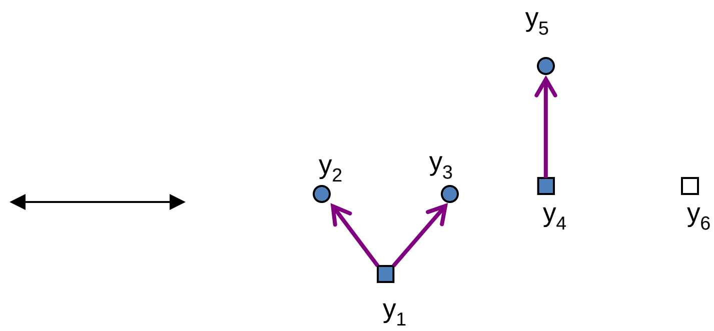

Figure 1(b)

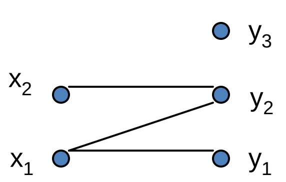

Figure 2(a)

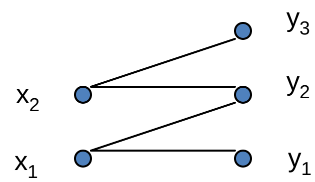

Figure 2(b)

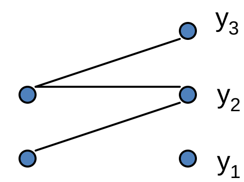

Figure 2(c)

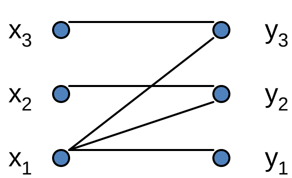

Figure 3(a)

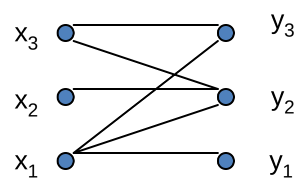

Figure 3(b)

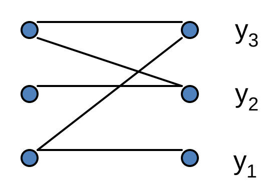

Figure 3(c)

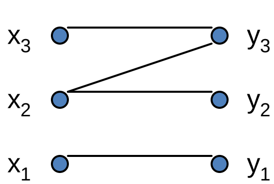

Figure 4(a)

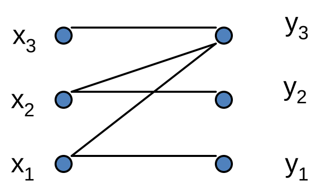

Figure 4(b)

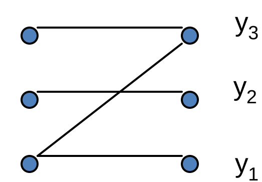

Figure 4(c)

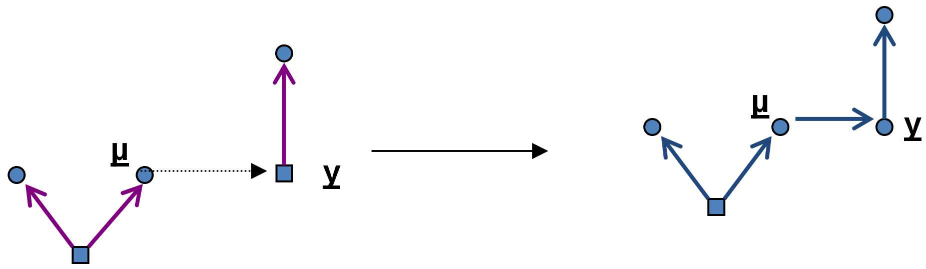

Figure 5(a)

Figure 5(b)

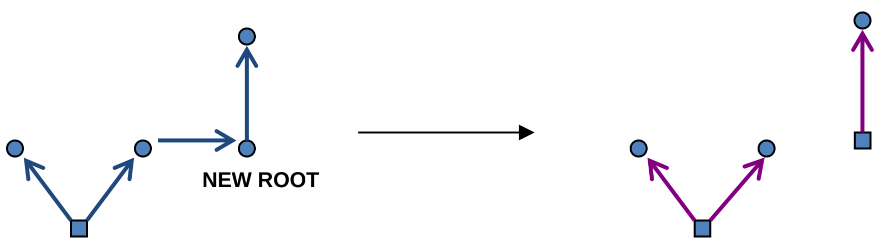

Figure 6 (a)

Figure 6(b)

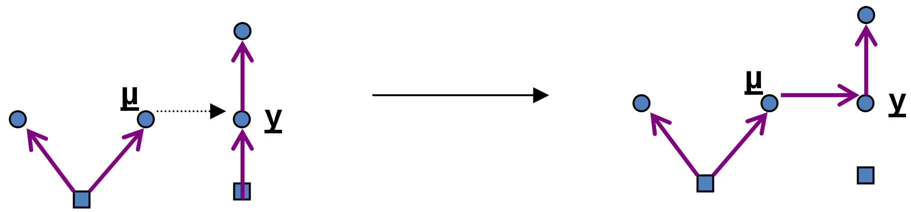

Figure 7(a)

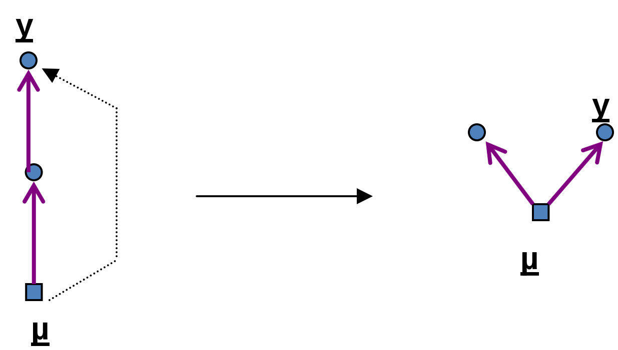

Figure 7(b)

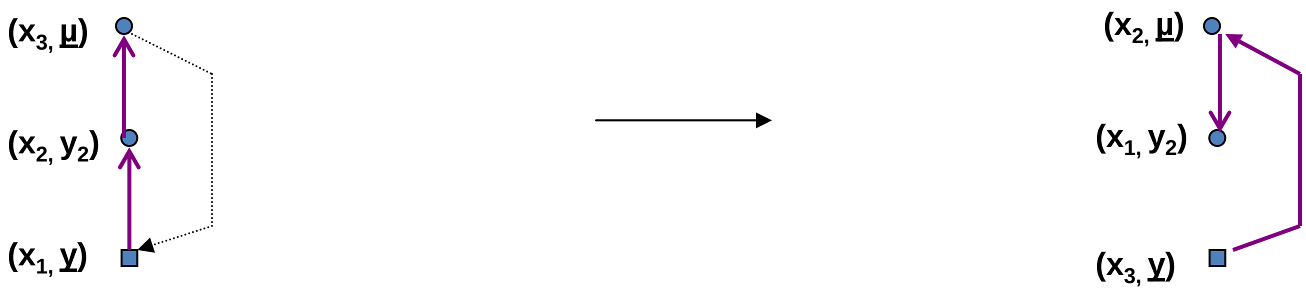

Figure 8(a)

Figure 8(b)

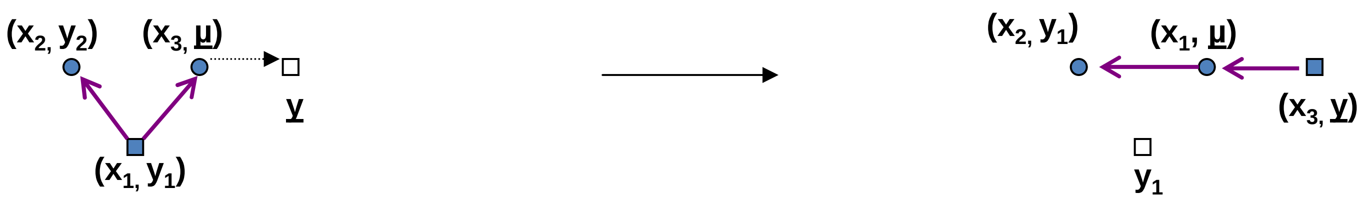

Figure 9(a)

Figure 9(b)

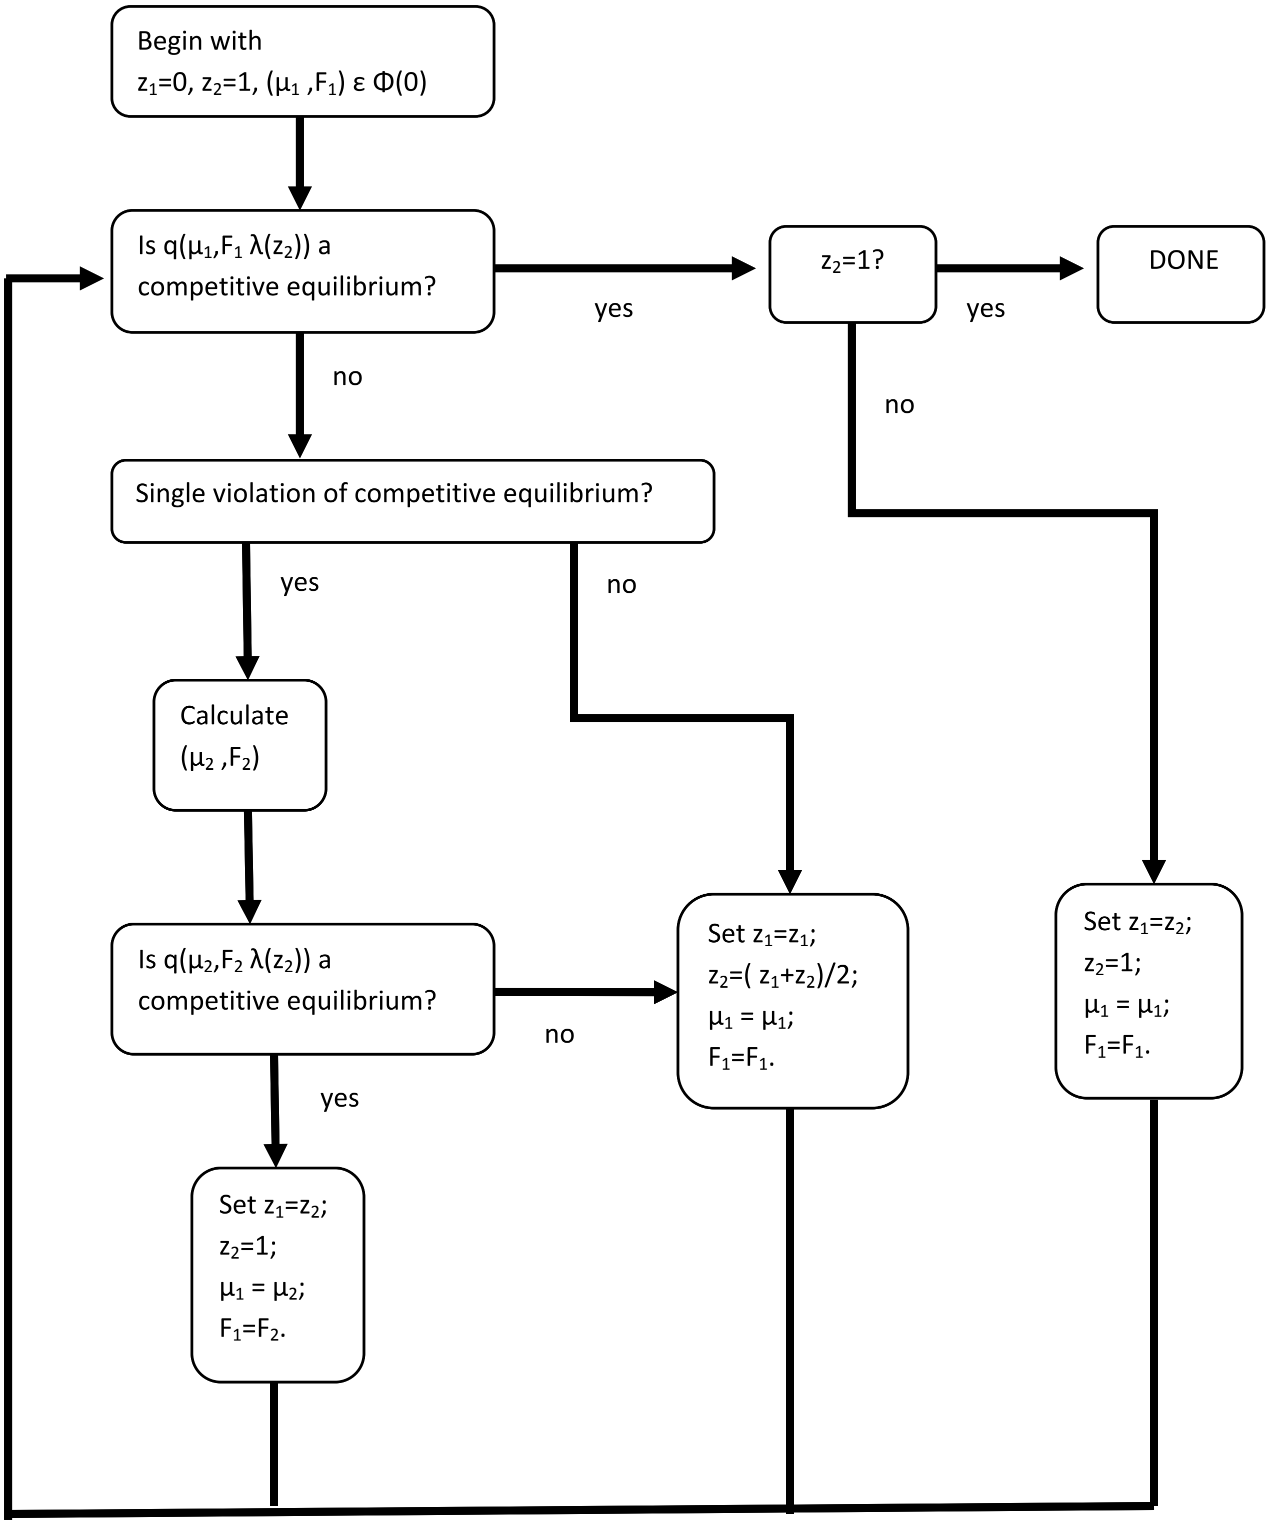

Figure 10: Flow Chart Describing Algorithm for Calculating Φ(1) from Φ(0)