Comparative Statics with Indivisible Goods

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

April 2010 (Preliminary)

Abstract

Indivisible all-or-nothing purchase decisions and wealth effects are crucial features of many real world markets, such as housing markets and many auction markets. Together, indivisibilities and wealth effects render classical techniques of marginal analysis unsuitable for comparative static analysis. We introduce a new mathematical apparatus ideally suited to analyzing the impact of changes in the economic environment on market for large individisible goods. In addition to characterizing distinctive forms of market adjustment, our apparatus produces an algorithm for identifying market equilibria.

1. Introduction

Indivisible all-or-nothing purchase decisions and wealth effects characterize many real world markets, such as housing markets and many auction markets. Yet for reasons of mathematical tractability, economists typically abstract away from either wealth effects or indivisibilities in studying these markets. As a result, the comparative static analysis of markets in which significant goods are indivisible remains in its infancy.

The chief barrier to progress is technical. The standard comparative static tool-kit is based on marginal logic. To approximate the impact of policy changes, one weights up derivatives of the utility and profit functions for those who purchase or sell each good in positive quantity. Indivisibilities lead to discrete adjustment.

In this paper we introduce new mathematical structures ideally suited to comparative static analysis of markets with individisibilities. Our "GA-structures" combine an allocation of goods

*We thank Eric Budish, Mamoru Kaneko, Jeffrey Mensch, Parag Pathak, Alvin Roth, Ennio Stacchetti, and Ivan Werning for helpful comments. Leahy thanks the NSF for financial support.

with a graphical 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]).

Demange and Gale (1985) provide an important building block for our work by showing that the set of equilibrium prices for allocation markets with non-transferable utility (NTU) is a lattice with maximal and minimal elements. We show that the minimum price equilibria can be identified by solving an optimization problem on the class of GA-structures. The statement of this theorem makes no explicit reference either to the supply or the demand side of the market. This link between optimization theory and equilibrium theory enables us provide a generic characterization of comparative static effects of changes in the economic environment.

The most novel comparative static effects occur when a change in the environment causes the equilibrium allocation to change. We find possible changes to the corresponding GA-structure to be lawful. Generically, there are five and only five distinct forms of switch in the pattern of allocation and indifference, each associated with a distinctive alteration to the identified GA-structure. In addition to characterizing the various forms of comparative static, we identify a path through GA-structures that computes minimum price equilibria from a starting point with a trivial equilibrium. The algorithmic use of GA-structures may help overcome computational barriers to application of NTU allocation model.[1]

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 uses this characterization to study the dependence of minimum price equilibria on the economic environment. Section 8 illustrates the five generic forms of comparative static transition, and identifies an algorithm that uses a simple path through GA-structures to identify minimum prices equilibria. Section 9 defines a dual to the allocation problem, and uses it to characterize the complete set of equilibria. Section 10 concludes with directions of current research interest. In addition to applications in the housing market, there are potential applications to the study of auctions with budget constraints, and auctions in which there is diminishing marginal utility.

[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.

2. Related Literature

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

With linear utility, Shapley and Shubik (1972) showed that 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 Beckman (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 a house. In the linear case, the social planner allocates goods based only on some fixed notion of how much each agent desires each good. If a poor agent enjoys a sea-view more than a rich agent, the planner will prefer to 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.

There are some theoretical results in the non-transferable utility case. Kaneko (1982) established conditions for the existence of an equilibrium.[2] 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, and established some basic comparative static properties of the minimum price equilibrium. There is no parallel, however, to the computational simplicity of the linear case.

Rosen's (1974) hedonic pricing model offers a simple way forward. It also prices heterogeneous goods given heterogeneous buyers. Rosen makes the simplifying assumption that there is a continuous density over characteristic bundles and that 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, in particular, one can rarely 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.

[2] 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.

Allocation problems arise naturally in a number of areas in economics. In the housing literature, our 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.[3] In the auction and mechanism design literature, our equilibrium is similar to a second price auction or a Vickery-Groves-Clark mechanism. These models almost always assume transferable utility. One exception is the paper by Demange and Gale (1985) in which they show that in the minimum price equilibrium buyers would truthfully report their preferences.

3. The Model

We work with a variant of the model in Demange and Gale (1985). Demange and Gale simply the exposition and the analysis by removing all reference to budget constraints and treating prices as transfers that may be positive or negative.[4] This removes the need to discuss what allocations and what price vectors are feasible. In Section 8, we recast the problem in a more familiar light.

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 sellers. Buyers 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. We assume that $n \geq m$ so that it is possible to match each buyer with a good.[5]

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, the size of the transfer that the buyer makes to the seller, and a taste parameter that depends on the good that is purchased. This payoff is summarized by the utility function $U_a : Y \times \mathbb{R} \times \Xi \to \bar{\mathbb{R}}$, where $U_a(y_i, p_i, \xi_{ai})$ is the utility to x_a from the purchase of y_i at the price p_i when the taste parameter is $\xi_{ai} \in \Xi \subseteq \mathbb{R}$. Note that ξ_{ai} may depend on both the buyer and the good. Let ξ_a denote the vector of ξ_{ai}.

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.

[3] Alonso and Roback also consider the supply side of the market, which introduces further complications.

[4] With budget constraints, consumers' choice correspondence may cease to 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 lead to tedious discussions of feasibility that have no baring on the analysis.

[5] The possibility that a buyer may choose not to make a purchase can be captured by associating a subset of goods with exit.

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 $r_i = p_i$.[6]

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

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

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

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 $p \in R^n$, let $U(p) \equiv \{y_i \in Y|p_i > r_i\}$ denote the set of goods with prices above seller reservation levels.

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

1. $\hat{\mu}_a \in D_a(\hat{p}, \xi_a)$ for all $x_a \in X$.
2. $\hat{p}_i \geq r_i$.
3. If $y_i \in U(\hat{p})$, then there exists $x_a \in X$ such that $\hat{\mu}_a = y_i$.

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

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

[6] 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.

We make assumptions on preferences that guarantee the existence of an equilibrium.

Assumption A For each buyer $x_a \in X$,

1. $U_a(y_i, p_i, \xi_{ai})$ is continuously differentiable in p_i and ξ_{ai}, strictly decreasing in p_i and strictly increasing in ξ_{ai} for each $y_i \in Y$ and $x_a \in X$.
2. $\lim_{p_i \to \infty} U_a(y_i, p_i, \xi_{ai}) = -\infty$ and $\lim_{p_i \to -\infty} U_a(y_i, p_i, \xi_{ai}) = \infty$ for all $x_a \in X$ and $y_i \in Y$.

The first assumption is a straight forward regularity assumption that will allow us to use the implicit function theorem.[7] The second assumption 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. Demange and Gale (1985) prove that under these conditions the set of equilibrium prices is a closed lattice, so that a minimum and a maximum price equilibrium exist.

Let $\{\xi_{ai}\}_{(a,i)} \in \mathbb{R}^{mn}$ denote the vector of preference parameters, $r \in \mathbb{R}^n$ the vector of reservation prices, and $\rho = (\{\xi_{ai}\}_{(a,i)}, r) \in \Lambda$ where Λ is a bounded, open subset of $\mathbb{R}^{mn+n}$.[8] In the next three sections, we develop properties of minimum price competitive equilibria for a given ρ. To simplify the exposition, we suppress reference to the taste parameter ξ_{ai}. In the following sections, we use these properties to study comparative statics in ρ.

4. An Example

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

Consider a market composed of two goods y_1 and y_2 and two buyers x_a and x_b. Suppose that $U_a(y_1, p_1) = 2 - p_1$ and $U_a(y_2, p_2) = 1 - p_2$ so that when the prices of the two goods are equal x_a prefers good y_1, and suppose that $U_b(y_1, p_1) = 1 - p_1$ 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 $r_1 = r_2 = 0$.

The minimum price competitive equilibrium in this example is trivial: The price of each good is set equal to its reservation value and each buyer is allocated to the good he prefers.

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

[7] Strict monotonicity simplifies the later analysis but is a stronger condition than needed for existence.

[8] Taking Λ as open avoids the question of how to do comparative statics at the boundary of the parameter space.

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 all of the prices in the set to find 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. If we knew which buyer's 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.

In the current example, there are two possible allocations: μ^1 and μ^2 where $\mu^1_a = y_1$ and $\mu^1_b = y_2$, and $\mu^2_a = y_2$ and $\mu^2_b = y_3$. 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.

We will show that the desired price vector can always be found by first for each allocation maximizing prices, good by good, across all potential chains, and then minimizing this outcome across all potential allocations. In the current example, Table 1 reports the price vector that results from each chain and each allocation. For example, if we use allocation μ^1 and Chain 1, then the price of y_1 is set to $r_1 = 0$, and we use the indifference of buyer x_a to price y_2. x_a is indifferent between the two goods if the price of y_2 is -1.

TABLE 1

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

Fixing the allocation and maximizing these price vectors across chains leads to the price vector in the last column. It is important that this maximization be done price by price, since in the

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

case of μ^2 there is no single price vector that is unambiguously the largest; the rows in Table 1 are not fully ordered and do not form a lattice. Given the results of this maximization, the last row reports the minimum of the price vectors in last column. This result is equal to the minimum competitive equilibrium price vector.

Intuitively, two forces are at work. First, minimum competitive 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 lowers 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 then 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 priced at reservation. In graph theory, this property of there being a unique path from any vertex to 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 on vertex set Y with the following properties:

1. F is a forest of trees.
2. Let $E(F)$ denote the set of edges of F. There exists a subset of vertices $R(F) \subset Y$, the root goods, such that for each component tree of F there corresponds a unique element of $R(F)$, and each edge in $E(F)$ is directed away from the corresponding element of $R(F)$.

Figure 1 illustrates a directed, rooted tree. The nodes are shown as circles, except for the root node which is shown as a square. Each node corresponds to an indivisible good y_i. The

[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.

edges are shown as arrows connecting one node 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**]

Letting $(y_i, y_k) \in E(F)$ denote the edge directed from good $y_i \in Y$ to good $y_k \in Y$, we 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 for each non-root good $y_i \in Y \backslash R(F)$, there exists a unique root good $y_r \in R(F)$ and a corresponding unique directed path $\{(y_r, y_1), (y_1, y_2), \ldots (y_n, y_i)\} \subset E(F)$ 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 $(y_i, y_k) \in E(F)$, 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 $F \in \mathcal{F}$ and an allocation $\mu \in M$ such that, if $(y_i, y_k) \in E(F)$, then there exists $x_a \in X$ such that $\mu_a = y_i$.

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

We construct a mapping from GA-structures to prices, $q : \mathcal{G} \rightarrow \mathbb{R}^n$. The price mapping is derived by induction on the set of goods that we have priced. The idea is first to 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 $q(\mu, F)$.

Definition Given Assumption A1 and $(\mu, F) \in \mathcal{G}$ we define $q(\mu, F) \in \mathbb{R}^n$ 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,

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

[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 successor. See also Miyake (2003).

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.

It is easy to see that this construction s well defined. First, since $(\mu, F) \in \mathcal{G}$, there always exists $x_a \in X$ with $\mu_a = y_i$ in step 2. Second, it follows from Assumption A that there exists a unique $q_k(\mu, F) \in \mathbb{R}$ that satisfies (5.1). Third, given the finite number of goods, this process will end after a finite number of steps with $A_s = Y$. Finally, it defines a unique price vector $q(\mu, F) \in \mathbb{R}^n$. 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.

6. A Min-Max Theorem

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

Theorem (μ^*, p^*) is a MPCE if and only if:

$$p_i^* = \min_{\mu \in M} \max_{F \in \mathcal{F}_\mu} q_i(\mu, F) \quad \text{all } i \in \{1, \dots n\}; \tag{6.1}$$

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

$$\mu^* \in \arg\min H(\mu) \tag{6.2}$$

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

We establish this result through a series of lemmas. We know from Demange and Gale (1985) that there exists a minimum price competitive equilibrium in this model. 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, $q(\mu^*, F^*) \geq q(\mu^*, F)$ 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 $q(\mu, F) \geq q(\mu^*, F^*)$. 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. Note that the converse to (6.2) may not hold. There are GA structures that generate the minimum equilibrium price, but are not competitive equilibrium allocations. A partial converse would state that if $\mu \in \arg\min H(\mu)$ and for all $y_i \in U(p^*)$ there exist x_a such that $\mu_a = y_i$ then (μ, p^*) is a MPCE.

Most GA structures generate prices and allocations that are inconsistent with optimization by buyers or sellers. Some generate prices that lie below sellers' reservation; 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 to be the null graph. Minimizing over μ guarantees that all goods in $U(p)$ are potentially allocated; for example, if such a good were unallocated and were demanded by only one buyer, then we could lower prices by reallocating that buyer to that good. 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.[12]

Many comparative static results from the literature follow immediately from Theorem 1. Demange and Gale (1985) show that minimal equilibrium prices are 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 restricting one buyer be allocated to a null tree.

7. The Minimum Equilibrium Price Correspondence and Paths through Parameter Space

7.1. The Minimum Equilibrium Price Correspondence

To characterize the dependence of minimum price equilibria on the parameters, we introduce $\Phi : \Lambda \to \mathcal{G}$, a mapping from parameters to the set of GA-structures that generate minimum

[12] Note that $\mu \in \arg\min H(\mu)$ does not imply that μ is part of a MPCE. A partial converse would state that if $\mu \in \arg\min H(\mu)$ and for all $y_i \in U(p^*)$ there exist x_a such that $\mu_a = y_i$ then (μ, p^*) is a MPCE.

price competitive equilibrium:

$$\Phi(\rho) = \{(\mu, F)|\mu \in \arg\min H(\mu) \text{ and } H(\mu) = q(\mu, F, \rho)\}$$

Theorem 1 establishes that $\Phi(\rho)$ is non-empty. Moreover, the price vector $q(\mu, F, \rho)$ induced by a fixed GA-structure is continuous in ρ given Assumption A which states that $U_a(y_i, q_i, \xi_{ai})$ is continuous in ρ.

Lemma 4 Given Assumption A, $(\mu, F) \in \mathcal{G}$ and $\rho \in \Lambda$, $q(\mu, F, \rho)$ is continuous in ρ.

Given the continuity of the $q(\mu, F, \rho)$, it follows directly from the Theorem of the Maximum applied to (6.1) that $\Phi(\rho)$ is upper-hemicontinuous.

Lemma 5: Given Assumption A, $\Phi(\rho)$ is non-empty, compact-valued, and upper-hemicontinuous at $\rho \in \Lambda$.

Let $\underline{p}$: $\Lambda \rightarrow \mathbb{R}^n$ denote the equilibrium price vector associated with a minimum price equilibrium.

$$\underline{p}(\rho) = \{p \in \mathbb{R}^n | p = q(\mu, F, \rho) \text{ and } (\mu, F) \in \Phi(\rho)\}$$

The Theorem of the Maximum also establishes that $\underline{p}(\rho)$ is continuous.

Lemma 6: Given Assumption A, $\underline{p}(\rho)$ is continuous in ρ.

7.2. MPCE Paths

We will perform comparative statics by analyzing how the equilibrium changes as we follow a path $\rho(z)$ from one set of parameters to another. While it might be impossible to find a path between two sets of parameters that avoids points at which Φ takes on multiple values. We now show that under fairly general conditions almost all paths avoid parameters with more than two equilibria.

In order to make explicit statements like "on almost every path", we need a measure on paths. To this end, fix the "shape" of a path and vary the initial condition. The shape is a continuous mapping $S : [0, 1] \rightarrow \mathbb{R}^{mn+n}$ such that $S(0) = 0$. The initial condition is a point in parameter space $\rho_0 \in \Lambda$. The pair (ρ_0, S) define a path $\rho(z; \rho_0, S) = \rho_0 + S(z)$, which begins at ρ_0 and ends at $\rho_0 + S(1)$. A path is admissible if $\rho_0 + S(z) \in \Lambda$ for all $z \in [0, 1]$. Let Λ_S denote the set of ρ_0 for which $\rho_0 + S(z)$ is admissible. For the remainder of the section we fix the shape

S and assume that Λ_S has positive Lebesgue measure. Our measure over paths is the Lebesgue measure on Λ_S.

A path (ρ_0, S) induces a correspondence $\phi(z; \rho_0, S) : [0,1] \to \mathcal{G}$ which maps each z into the set of GA-structures that generate the minimum price competitive equilibrium for the model with parameters $\rho(z)$

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

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

In order to make statements about ϕ, we need to restrict the set of paths as well as the prices generated along these paths in order to avoid anomalies such as space filling curves or closed sets with no interior and positive measure. We therefore limit ourselves to analytic functions.

Assumption B: The following functions are analytic:

1. Each component of S is analytic.
2. 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 $y_i \in Y$.

Assumption B implies that there cannot be too many points at which $\Phi(\rho)$ is multi-valued. Let $\hat{\Lambda} \equiv \{\rho \in \Lambda \text{ s.t. } |\Phi(\rho)| > 1\}$ be the set of parameters for which multiple GA structures are associated with the generate the minimum price competitive equilibrium. Here $|z|$ denotes the number of elements in the set z. With Assumption B, $\hat{\Lambda}$ has Lebesgue measure zero. Intuitively, every point at which $|\Phi(\rho)| > 1$ is the intersection of two or more $q(\mu, F, \rho)$. Assumption B implies that the $q(\mu, F, \rho)$ are analytic. Analytic functions cannot intersect too often without being identical on there entire range. Our ability to alter the payoffs of individual buyers and individual sellers implies, however, that we can perturb ρ in a way that affects prices generated by one GA structure and not another. In this way we can show that the $q(\mu, F, \rho)$ are not equal everywhere and therefore cannot intersect too often.

Lemma 7 $\hat{\Lambda}$ has zero Lebesgue measure.

Assumption B also implies Theorem 2 which is very useful in characterizing comparative statics along paths.

Theorem 2: Given $\rho_0 \in \Lambda_S$, either $\phi(z; \rho_0, S)$ has the following properties or ρ_0 lies on the boundary of an open set O and $\phi(z; \hat{\rho}_0, S)$ has the following properties for all $\hat{\rho}_0 \in O$:

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

2. $\phi(z)$ has at most 2 elements for all $z \in [0,1]$.

3. The graph of $\phi(z)$ contains no isolated points.

4. Given $(\mu, F) = \phi(z)$ and $y_i \notin R(F)$, $q_i(\mu, F, \rho(z)) > r_i(z)$ for all $z \notin Z$.

We establish Theorem 2 through a sequence of lemmas that appear in the Appendix. We first show that Assumption B implies that the $q_i(\mu, F, \rho(z))$ are analytic. Each point in Z is associated with an intersection of at least two $q(\mu, F, \rho(z))$. Analytic functions whose intersections have accumulation points must be equal everywhere, and the ρ_0 such that two $q(\mu, F, \rho(z))$ are identical are measure zero. The generic lack of accumulation points establishes the first part of the Theorem.

The other parts follow from the observation that for ρ in the interior of $\Lambda(S)$ we can perturb ρ to reduce the size of $|\phi(z)|$ whenever $|\phi(z)| > 2$. It follows that the set of undesirable points is on the boundary of an open set. These perturbations result from the following two lemmas.

Lemma 8: Suppose $(\mu, F) \in \Phi(\rho)$ and $q_i(\mu, F, \rho) > r_i$ for $y_i \notin R(F)$. There exists $(\mu', F') \in \Phi(\rho)$ with $(\mu', F') \neq (\mu, F)$ if and only if there exists x_b and $y_0 \in D_b(q(\mu, F, \rho)$ such that $y_0 \neq \mu_b$ and $(\mu_b, y_0) \notin E(F)$.

Lemma 9: Under conditions of the preceding lemma. $|\Phi(\rho)| = 2$ if there exists only one such x_b and y_0.

Given $|\phi(z)| > 2$, these lemmas imply that the set $V_b = \{(x_b, y_0) | y_0 \in D_b(q(\mu, F, \rho(z)), y_0 \neq \mu_b$, and $(\mu_b, y_0) \notin E(F)\}$ has at least two elements. We can perturb ρ_0 to eliminate all but one, thereby reducing $|\phi(z)|$ to 2. The complication in the proof is keeping track of what happens at other z. If $q_i(\mu, F, \rho) = r_i$ for $y_i \notin R(F)$, we can perturb r_i, to reduce $|\phi(z)|$.

Note that Lemmas 8 and 9 imply that the existence of multiple equilibrium GA structures is associated with "extra" indifference relationships. Together with Theorem 2 they imply that the normal state of affairs is one equilibrium GA structure and no extra indifference relationships and occasional points and at which there is one extra indifference relationship.

8. Comparative Statics and Computation

According to Theorem 2, as z rises ϕ is equal to some GA-structure (μ, F) until at some point $\hat{z}$ a second GA-structure (μ', F') appears. Given that there are no isolated points, ϕ then switches to the new GA-structure beyond $\hat{z}$. Lemmas 8 and 9 imply that these points of transition are associated either with a single price falling to r_i or with the expansion of the demand set of a single buyer to a single new good.

The model makes a distinction between local and global comparative statics. Local comparative statics involve shifts in $q(\mu, F, \rho)$ for a given GA structure. It is very unlikely that a small change in the parameters will trigger a new indifference or destroy an existing one. Global comparative statics may be decomposed into a finite number of adjustments in the GA structure. These adjustments may be classified into one of four different types depending on the nature of the indifference that is being added or removed.

The first case is one in which the price of a non-root good $y_i \notin R(F)$falls to its reservation level r_i. In this case, to arrive at the next GA structure (μ', F'), we simply eliminate the edge pointing towards y_i from F. y_i and its successors in F separate off from their component of F and form a new tree in F'.

The other three cases involve the expansion of the demand set of some buyer. Let x^* denote the buyer who becomes indifferent to a new good; μ^* the good assigned to that buyer up to this point; and y^* the good not previously demanded. There are three mutually exclusive and exhaustive cases to consider: y^* is unallocated; y^* is allocated, but there exists no path directed from y^* to μ^*; and y^* is allocated and there exists a path directed from y^* to μ^*.

The first case is illustrated in Figure 2(a). The figure depicts y_r and its successors where y_r is the root good in the component of F containing y^*. We have labeled some of the goods of interest, and also indicated the buyers allocated to those goods. We have depicted root goods as squares. y_r and y^* are both root goods. y^* is a root good because it is unallocated. The solid arrows indicate the directed edges of F. In figure 2(a) μ^* is depicted as a successor of y_r. In general, μ^* may also be y_r. Since y^* is initially unallocated, it is also a null tree in F. The dashed arrow indicates the indifference of x^* between μ^* and y^*.

At the given price level, there are two allocations that satisfy buyer optimality. Buyers assigned to a good along the unique path connecting y_r to μ^* are happy either with the good that they receive under μ or with that good's direct successor, where we think of y^* as the successor of μ^*. Figure 2(b) illustrates the relevant portion of (μ', F'), the GA-structure associated with this alternative allocation. Note that directed edges adjust to reflect the reallocation of the buyers: edges are now directed away from y^*. At the given vector of reservation utilities, both GA-structures generate the same price vector. Theorem 2 calls for a switch from the current (μ, F) to (μ', F').

Formally, let $Y^P = (y_1^P, y_2^P, \dots y_k^P)$ denote the path from $y_r = y_1^P$ to $\mu^* = y_k^P$ in F (this path may be trivial if $\mu^* = y_r$). We shift x^* to y^*, and we shift all other buyers matched to a good in Y^P to that good's successor in Y^P:

$$\mu_a(s,t) = \begin{cases} y^* & \text{if } x_a = x^*; \\ y^P_{z+1} \in Y^P & \text{if } \mu_a(s,t) = y^P_z \text{ and } z = \{1, \dots k-1\}; \\ \mu_a(s,t) & \text{otherwise.} \end{cases}$$

We alter the edges $E(F)$ accordingly. We add the edge (y^*, μ^*). We delete the edge (y^P_1, y^P_2). We reverse the orientation of all other edges $(y^P_z, y^P_{z+1}) \in Y^P$. We replace all edges $\left(y^P_z, y_j\right)$ with $y^P_z \in Y^P$ and $y_j \notin Y^P$ with $\left(y^P_{z+1}, y_j\right)$ where $y^P_{k+1} = y^*$. We label the new graph $F(s, t+1)$.

[Figure 2]

The second case occurs when y^* is allocated under μ, but there exists no path directed from y^* to μ^* in F. Figures 3(a) and 3(b) illustrate two variants of this case. In Figure 3(a), μ^* and y^* are in different components of F. In Figure 3(b), μ^* and y^* are in the same component. In both instances we arrive at the new GA-structure (μ', F') by maintaining $\mu' = \mu$ and by replacing the edge (y_j, y^*) in F oriented towards y^* with the edge (μ^*, y^*). Theorem 2 calls for a switch from the current (μ, F) to (μ', F').

Formally, we make no changes to μ: $\mu' = \mu$. We alter F by replacing the edge (y_j, y^*) oriented towards y^*, with the edge (μ^*, y^*). We label the new graph F'.

[Figure 3]

The third case is when y^* is allocated and there exists a path directed from y^* to μ^*. This case is depicted in Figure 4(a). In this case, y^* could be y_r. The addition of x^*'s indifference between μ^* and y^* creates a directed circuit. At the current price vector, there are two allocations that satisfy buyer optimality: μ and an allocation in which each buyer assigned to a good in this circuit is instead assigned to that good's direct successor in this circuit. Figure 4(b) depicts the GA-structure (μ', F') associated with this alternative allocation. At the given vector of reservation utilities, both GA-structures generate the same price vector. Theorem 2 calls for a switch from the current (μ, F) to (μ', F').

Formally, we add the edge (μ^*, y^*) to F. This creates a directed circuit $Y^C = (y^C_1, \dots y^C_k, y^C_{k+1})$ where $y^C_1 = y^C_{k+1} = y^*$ and $y^C_k = \mu^*$. We shift each buyer who is matched an element of Y^C in the direction of the circuit: if $\mu_a = y^C_z \in Y^C$, then $\mu'_a = y^C_{z+1}$. Otherwise $\mu'_a = \mu_a$. We delete the edge (y^C_1, y^C_2). We reverse the orientation of all other $(y^C_z, y^C_{z+1}) \in Y^C$. We replace all edges $\left(y^C_z, y_j\right)$ with $y^C_z \in Y^C$ and $y_j \notin Y^C$ with $\left(y^C_{z+1}, y_j\right)$. We label the new graph F'.

[Figure 4]

Theorem 2 rules out several complications as unlikely. First, it rules out tangencies in which x^* becomes indifferent to y^* at z but then ceases to demand y^*. Second, it rules out points at which the demand correspondences of multiple buyers simultaneously expand or that the demand correspondence of a single buyer expands to include multiple new goods.

8.1. An Algorithm for Computing Competitive Equilibria

The maximization in problem (6.1) is over a very large and may be difficult to perform in practice. A generalization of Cayley's theorem states that for each allocation μ there are

$$\sum_{k=1}^{m} \binom{m}{k} km^{m-1-k}$$

different forests of rooted trees on the m allocated goods.[13] 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 previously shown how 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 ξ_{aa} for each buyer x_a high enough so that each buyer x_a prefers y_a to all other goods.[14] At this level of reservation utility, the minimum price equilibrium has all prices at reservation level and all buyers allocated to their null goods. There is a unique GA-structure corresponding to this equilibrium. It involves the equilibrium allocation and a null graph. This equilibrium is our starting point. We can then lower the ξ_{aa} to their original levels, tracking GA-structures that correspond to the minimum price equilibrium. Note that lowering the ξ_{aa} 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 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

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

[14] The range of the parameter space must be large enough that this is possible.

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 nice 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)).[15] 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.

9. The Dual Problem and the Equilibrium Set

9.1. Budget Constraints

We have constructed our GA structures off of the reservation prices of sellers. The dual is built on buyers' maximal willingness to pay. To this point we have assumed that buyers are willing and able to pay any amount to sellers. To introduce the dual problem, we first reinterpret the model as a model of buyers who hold money and chose whether to enter the market and purchase a good. Wealth will then constrain the set of prices buyers are willing to accept.

If wealth is to constrain choice, we need to specify what happens when a buyer cannot choose anything. We associate a subset of Y with these null or exit choices. Let y_a for $a \in \{1, \ldots m\}$ denote the exit option of buyer x_a. We set the reservation price of each y_a equal to zero, so that they are always affordable. We assume that the utility parameters are such that only x_a demands y_a at these prices.

Next we introduce money into the utility function by making a restriction on the form of U_a. We assume that $U_a(y_i, p_i, \xi_{ai}) = U_a(y_i, m_a - p_i, \xi_{ai})$ where we interpret m_a as x_a's initial holdings of a numeraire good. Assumptions A and B apply to this function. In order to ensure the existence of an equilibrium in which buyers do not spend more than m, we assume that each buyer x_a prefers his outside good y_a to any other good y_i when $p_i = m_a$.

Assumption C: For each buyer $x_a \in X$, $U_a(y_a, w_a, \xi_{aa}) \geq U_a(y_i, 0, \xi_{ai})$.

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

Assumption C replaces Assumption A2. Kaneko (1982) proves that there exists an equilibrium under Assumptions A and C.[16]

This alternative market setting introduces some technical complications. $q(\mu, F, \rho)$ is not defined when a buyer is allocated to a good whose price exceeds the buyer's holdings of the numeraire good. The maximum in (6.1) must be taken over the set $\tilde{\mathcal{F}}(\mu, \rho) \subseteq \mathcal{F}_\mu$ where $\tilde{\mathcal{F}}(\mu, \rho)$ is the set of F for which $q(\mu, F, \rho)$ is well defined. This amendment works because, all buyers can afford their purchases at the minimum price competitive equilibrium; this complication does not arise.

With these amendments, Theorems 1 and 2 apply in this environment.

9.2. The Dual Problem

We can define the dual to our alternative market setting. In the dual one switches the positions of buyers and sellers and reinterprets the search for equilibrium as taking place in the space of buyer utilities rather than in the space of seller prices. To characterize supply in the dual we introduce a set of n null buyers, one for each seller: we let x_{m+i} denote the null buyer for seller $y_i \in Y$. Let $\bar{X} = X \cup \{x_{m+i}\}_{i=1}^n$.

Buyers' utility υ in the dual plays the role of prices in the primal. The best a buyer can do is to purchase their favorite good at a zero price. Define $\bar{\upsilon} \in \mathbb{R}^m$ by $\bar{\upsilon}_a = \max_{y_i \in Y} U_a(y_i, w_a, \xi_{ai})$. The worst that a buyer can do is stick with his null good. Define $\upsilon^R \in \mathbb{R}^m$ by $\upsilon_a^R = U_a(y_i, w_a, \xi_{aa})$. Let $\tilde{\Pi} = \left\{\upsilon \in \mathbb{R}_+^m | \bar{u}^R \leq \upsilon \leq \bar{\upsilon}\right\}$ denote the range of buyer utilities.

The utility of seller y_i, $V_i(x_a, \upsilon_a)$, 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

$$V_i(x_{n+i}, \upsilon_{n+i}) \equiv r_i.$$

In the case of a non-null buyer, $V_i(x_a, \upsilon_a)$ is defined as the price that would have to charge for good y_i to provide buyer x_a with utility υ_a,

$$U_a\left[y_i, w_a - V_i(x_a, \upsilon_a)\right] = \upsilon_a.$$

Given $q \in \tilde{\Pi}$, 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" $V_i(x_a, \upsilon_a)$.

[16] Assumption C eliminates discontinuities in utility that arise from liquidity constraints. Note that this condition is satisfied when utility functions satisfy Inada conditions.

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

$$\lambda_i = \lambda(y_i) \in \bar{X}_i.$$

Definition 9.1. *A competitive equilibrium in the dual model is a pair $(\hat{\upsilon}, \hat{\lambda})$ with $\hat{\upsilon} \geq \tilde{\Pi}$ such that:*

1. *$\hat{\lambda}_i \in S_i(\hat{\upsilon})$ for all $y_i \in Y$.*

2. *If $\hat{\upsilon}_a > \bar{u}_a^R$, then there exists $y_i \in Y$ such that $\hat{\mu}_i = x_a$.*

Note that there is a natural 1-1 correspondence between equilibria in the primal and the dual that has the property of inverting the ordering. In particular the minimum equilibrium price in the primal corresponds to the maximum equilibrium utility in the dual and vice versa.

Assumptions A1 and C are sufficient to establish that the indirect profit function has key properties required for the remaining results of the paper to apply, in particular the existence results and graph theoretic characterizations of the minimum utility equilibrium. The analog of A1 is that $V_a(x_a, \upsilon_a)$ is continuous and strictly decreasing on $\upsilon_a \in \tilde{\Pi}_a$: these are immediate implications of the corresponding assumptions on the consumer utility functions. There is no need for an analog for A2, which is designed to ensure that wealth constraints do not create a discontinuity at the point of zero wealth, so that there is a point of indifference between staying in the market and leaving. This is universally valid in the dual problem, since there are no wealth constraints.

Given that the same underlying structure is valid, analogous results hold concerning the graph theoretic structure of the minimum-utility equilibrium and its identification through appropriately redefined match-edge structures based on directed graphs on $\bar{X}$, with edges indicating indifference on the supply side of the market. This means that the natural adaptation of our algorithm to the dual model identifies minimum equilibrium utilities and the corresponding equilibrium allocations, from which we can immediately infer the maximum equilibrium price. The algorithm is written as a function of the reservation level of seller price which is first taken so high that all sellers wish to leave the market, and is then lowered to its true value. Note that one can replace true reservation prices with the minimum equilibrium prices without impacting the working of the algorithm.

9.3. 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 6.

Theorem 6 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

We have presented a new mathematical apparatus for comparative statics of allocation markets with NTU. This apparatus enables us to characterize comparative statics and to provide a new class of algorithm for identifying equilibria. The results suggest many directions for future research. One broad line of such research involves placing the model in a dynamic context. This requires 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 of interest.

11. References

References

[1] Aigner, Martin, and Günter Ziegler (2003), *Proofs from THE BOOK*, New York: Springer-Verlag, second edition.

[2] Alonso, William (1964), *Location and Land Use*, Cambridge: Harvard University Press.

[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 Beckman (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] Miyake, Mitsunobu (1994), "Comparative Statics of Assignment Markets with General Utilities," *Journal of Mathematical Economics* 23, 519-531.

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

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

[16] Ricardo, David, (1817), *Principles of Political Economy and Taxation,* London, John Murray.

[17] Roback, Jennifer (1982), "Wages, Rents, and the Quality of Life," *Journal of Political Economy* 90, 1257-1278.

[18] Rosen, Sherwin (1974), "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy* 92, 34-55.

[19] Scaft, Herbert E. (1994), "The Allocation of Resources in the Presence of Indivisibilitie," *Journal of Economic Perspectives*, Vol. 8, 111-128.

[20] Shapley, Lloyd, and Martin Shubik (1972), "The Assignment Game I: The Core." *International Journal of Game Thoery* 1, 111-130.

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

12. Appendix

12.1. Theorem 1

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

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

Proof: Consider a MPCE (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,

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

The graph is completed by induction. Let $A_1 = R(F)$ 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 on the vertex set A_s such that F_s is a forest of rooted trees with root set $R(F)$ and with all of the edges of F_s are directed away from the roots. By construction, $R(F) \subset A_s$ and $Y \backslash A_s \subset U(p^*)$. Lemma 4 in Demange and Gale (1985) states that there exists x_a such that $\mu_a^* \in A_s$ and $D_a(p^*) \cap Y \backslash A_s \neq \varnothing$. Choose $y_i \in D_a(p^*) \cap Y \backslash A$, define $A_{s+1} = \{y_i\} \cup A_s$, add the edge (μ_a^*, y_i) to $E(F_s)$, and denote the resulting graph F_{s+1}. By construction, F_{s+1} is a forest of rooted trees with root set $R(F)$ with all edges directed away from the roots. Given that there are a finite number of elements in $Y \backslash R(\hat{F})$, this construction converges in a finite number of steps to $A_S = Y$ to graph F_S, and we define $F = F_S$ as the final forest of rooted trees on Y.

To establish that $(\mu, F^*) \in \mathcal{G}$, note $(y_i, y_k) \in E(F)$ implies that 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 $R(F)$ are at reservation prices. Furthermore note that for any edge $(y_i, y_k) \in E(F)$, 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 $\hat{p}$ is sufficient to complete the demonstration that $q(\mu, F^*) = p^*$ that which is precisely the condition used in generating the function $q(\mu, F^*)$.□

Lemma 2: For any MPCE (μ^*, p^*),

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

Proof: 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 $q_i(\mu^*, F) > p_i^*$ for some good $y_i \in Y$. $p_i > p_i^*$ implies $y_i \notin R(F)$. Let $y_r \in R(F)$ denote the root good that is a predecessor of y_i in F. $y_r \in R(F)$ implies $q_r(\mu^*, F) = r_r \leq q_r(\mu^*, F^*)$. This implies that we can choose y_i such that $q_j(\mu^*, F) \leq p_j^*$ for all predecessors of y_i in F, in particular y_k the direct predecessor of y_i. Consider x_b such that $\mu_b^* = y_k$. $(y_k, y_i) \in F$, implies that $U(y_k, q_k(\mu^*, F)) = U(y_i, q_i(\mu^*, F))$. As $q_k(\mu^*, F) \leq p_k^*$, $q_i(\mu^*, F) > p_i^*$, it follows that x_b strictly prefers y_i to y_k at the price vector p^*. But (μ^*, p^*) is a MPCE. This contradiction establishes the lemma.□

Lemma 3: Given a MPCE (μ^*, p^*) and an allocation μ that is not part of a MPCE, then

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

with strict inequality for one $y_i \in Y$.

Proof: Let (μ^*, p^*) be a MPCE and fix μ such that μ is not part of any MPCE. First we show how to construct an F such that $q(\mu, F) \geq p^*$. Then we show how to construct F' such that $q_i(\mu, F') > p_i^*$ for some good y_i.

By Lemma 1, there exists, F^* such that $q(\mu^*, F^*) = p^*$. It is easy to see that we can pick F such that $p_i^* = r_i$ implies $y_i \in R(F)$, which in turn implies by the definition of a competitive equilibrium that for all $y_i \notin R(F)$ there exists x_b such that $\mu_b^* = y_i$.

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 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 y_i has no successors in F, and if $\mu_b^* = y_i \neq \mu_b$, then include (μ_b, μ^*) in $E(K)$.

Note that K may not be a tree. To construct F we will make a selection from K.

We construct F by induction. Let $R(F) = \{y_i | \text{such that the indegree of } y_i \text{ is zero in } K\}$. Note that $R(F) \subseteq R(F^*)$, since in constructing K we did not reduce the indegree of any node from its value in 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 $y_i \in A_s \subset Y$. We claim that there exists x_c such that $\mu_c \in A_s$, $(\mu_c, y_k) \in E(K)$, and that $y_k \in Y \backslash A_s$. Since $p_j^* > r_j$ for all $y_j \in Y \backslash A_s$, there exists y_j and x_d such that $\mu_d^* \in A_s$, $(\mu_d^*, y_j) \in E(F^*)$.If $\mu_d \in A_s$, we add (μ_d^*, y_j) to F. Otherwise, $\mu_d \in Y \backslash A_s$. Since for all $y_i \in Y \backslash A_s$, there exists x_b such that $\mu_b^* = y_i$, there exists x_e such that $\mu_e^* \in Y \backslash A_s$, but $\mu_e \notin Y \backslash A_s$. If μ_e^* has a direct successor y_l in F^*, then $(\mu_e, y_l) \in E(K)$ satisfies the claim and we add (μ_e, y_l) to F. If not then, (μ_e, μ_e^*) satisfies the claim add (μ_e, μ_e^*) to 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

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

Finally by construction

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

It follows from the monotonicity of U_a that $q_k \geq p_k^*$. This completes the induction step.

If $q_k > p_k^*$ for some $y_k \in Y$. Then we take $F'' = F$. Suppose $q_k = p_k^*$ for all y_k. Since μ is not a competitive equilibrium allocation there exists $x_a \in X$ such that $\mu_a \notin D_a(p^*)$. Suppose $y_j \in D_a(p^*)$. $\mu_a \notin D_a(p^*)$ implies $(\mu_a, y_j) \notin E(F')$. Construct F'' as follows: Begin with F'; if there exits $(y_i, y_j) \in E(F')$, eliminate it; add the directed edge (μ_a, y_j); if there exists. It is easy to see that this implies a strictly higher price for y_j and its successors in F'. This completes the proof of the Lemma.□

Proof of Theorem 1: Suppose that (μ^*, p^*) is a minimum price competitive equilibrium. Lemma 1 implies

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

Lemma 3 implies that

$$H(\mu^*) \leq H(\mu)$$

for all $\mu \in M$. It follows that

$$\mu^* \in \arg\min H(\mu).$$

Suppose that $(\hat{\mu}, \hat{p})$ satisfy

$$\hat{p} = \min_{\mu \in M} \max_{F \in \mathcal{F}_\mu} q(\mu, F) \tag{12.1}$$

and

$$H(\hat{\mu}) \leq H(\mu)$$

Since the solution to (12.1) is unique and a minimum price competitive equilibrium exists by Assumption A, it follows that $(\hat{p}, \hat{\mu})$ is a competitive equilibrium.□

12.2. Properties

We first present some properties of q.

Lemma 4: Given $(\mu, F) \in \mathcal{G}$ and $\rho \in \Lambda$, $q(\mu, F; \rho)$ is continuous in ρ.

Proof: Consider $\bar{\rho}$ in Λ. Let $\{\rho_n\}$ be any sequence in Λ converging to $\bar{\rho}$. The proof is by induction. It is useful in what follows to write to drop the dependence of $q_i(\mu, F; \rho)$ on (μ, F). Consider $A_0 \equiv R(F)$ and set $q_i(\rho_n) = r_i(\rho_n)$ on A_0. Since $r_i(\rho_n)$ converges to $r_i(\bar{\rho})$, it follows that $q_i(\rho_n)$ converges to $q_i(\bar{\rho})$ for all goods in A_0. Suppose that $q_i(\rho_n)$ converges to $q_i(\bar{\rho})$ for all goods in A_s, $s \geq 0$. Let S comprise direct successors of A_s in F. For each $y_k \in S$, there exists $y_i \in A_s$ and x_a such that $\mu_a = y_i$, such that

$$U_a(q_i(\rho_n), y_i, \rho_n) = U_a(q_k(\rho_n), y_k, \rho_n)$$

The continuity and strict monotonicity of U then imply that $q_k(\rho_n)$ converges to $q_k(\bar{\rho})$.□

Lemma 5: Given Assumption A and the continuity of U_a in ξ_a for all x_a, $\Phi(\rho)$ is non-empty, compact valued, and upper-hemicontinuous at $\rho \in \Lambda$.

Proof: Consider $\bar{\rho} \in \Lambda$. Assumption A guarantees the existence of a minimum price competitive equilibrium. Theorem 3 guarantees the existence of a GA-structure. Hence $\Phi(\bar{\rho})$ is non-empty.

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

That $\Phi(\rho)$ is upper-hemicontinuous follows from applying the theorem of the maximum first to

$$H(\mu, \rho) = \max_{F \in \mathcal{F}_\mu} q(\mu, F, \rho)$$

and then to

$$\hat{p} = \min_{\mu \in M} H(\mu, \rho)$$

Lemma 4 shows that the $q(\mu, F, \rho)$ are continuous functions in ρ, and the $\mathcal{F}_\mu$ are trivially continuous correspondences in ρ, implying that the $H(\mu, \rho)$ are continuous in ρ. Again M is trivially continuous in ρ.□

Lemma 6 $\underline{p}(\rho)$ is continuous in ρ.

Proof: Follows directly from the theorem of the maximum. See proof to Lemma 5.

Lemma 7 Let $\hat{\Lambda}$ denote the subset of Λ for which Φ has more than one value. $\hat{\Lambda}$ has zero Lebesgue measure.

Proof: (sketch) $\Phi(\rho)$ is multi-valued only if there are (μ, F) and (μ', F') that generate the minimum equilibrium price. At any such ρ there exists a ρ nearby such that Φ is single valued. Hence $q(\mu, F, \rho)$ and $q(\mu', F', \rho)$ are not globally equal. As they are analytic, their intersections must be therefore be measure zero.$\square$

12.3. Theorem 2

We first present two useful lemmas which will be used throughout the proof of Theorem 2. The first establishes that if we have multiple equilibrium GA-structures at some parameter vector ρ, then given an equilibrium GA-structure in which prices of all non-root goods exceed sellers' reservation, there exists a buyer who demands a good other than the good she is assigned or its direct successors. The second states that if there are two GA-structures there is only one such good.

Lemma 8: Suppose $(\mu, F) \in \Phi(\rho)$ and $q_i(\mu, F, \rho) > r_i$ for $y_i \notin R(F)$. There exists $(\mu', F') \in \Phi(\rho)$ with $(\mu', F') \neq (\mu, F)$ if and only if there exists x_b and $y_0 \in D_b(q(\mu, F, \rho)$ such that $y_0 \neq \mu_b$ and $(\mu_b, y_0) \notin E(F)$.

Proof: We begin with the if statement. We consider three cases depending on the position of y_0 in (μ, F).

Case 1: y_0 is unallocated by μ.

We construct $(\mu', F') \in \Phi(\rho) \backslash (\mu, F)$. Let $\hat{Y}^P = (y_1^P, y_2^P, \ldots y_k^P)$ denote the path from $y_1^P = y_R$ to $\mu_b = y_k^P$ in F (this path may be trivial if $\mu_b = y_r$). We shift x_b to y_0, and we shift all other buyers matched to a good in Y^P to that good's direct successor in Y^P:

$$\mu_a' = \begin{cases} y_0 & \text{if } x_a = x_b; \\ y_{z+1}^P \in Y^P & \text{if } \mu_a = y_z^P \text{ and } z = \{1, \ldots k-1\}; \\ \mu_a & \text{otherwise.} \end{cases}$$

We alter the edges $E(F)$ accordingly. We add the edge (y_0, μ_b). We delete the edge (y_1^P, y_2^P). We reverse the orientation of all other edges $(y_z^P, y_{z+1}^P) \in Y^P$. We replace all edges $\left(y_z^P, y_j\right)$ with

$y_z^P \in Y^P$ and $y_j \notin Y^P$ with $\left(y_{z+1}^P, y_j\right)$ where $y_{k+1}^P = y^*$. We label the new graph F'. It is easy to show that $q(\mu, F, \rho) = q(\mu', F', \rho)$, so that $(\mu', F') \in \Phi(\rho)$.

Case 2: y_0 is allocated but not a predecessor to μ_b.

In this case, we make no changes to μ: $\mu' = \mu$. We construct F' from F by replacing the edge $(y_j, y_0) \in F$ oriented towards y_0, with the edge (μ_b, y_0). It is easy to show that $q(\mu, F, \rho) = q(\mu', F', \rho)$, so that $(\mu', F') \in \Phi(\rho)$.

Case 3: y_0 is a predecessor to $\mu(x_b)$.

To construct (μ', F'), we add the edge (μ_b, y_b) to F. This creates a directed circuit $Y^C = (y_1^C, \ldots y_k^C, y_{k+1}^C)$ where $y_1^C = y_{k+1}^C = y_0$ and $y_k^C = \mu_b$. We shift each buyer who is matched an element of Y^C in the direction of the circuit: if $\mu_a = y_z^C \in Y^C$, then $\mu'_a = y_{z+1}^C$. If $\mu_a \notin Y^C$, $\mu'_a = \mu_a$. We now delete the edge (y_1^C, y_2^C). We reverse the orientation of all other $(y_z^C, y_{z+1}^C) \in Y^C$, and we replace all edges $\left(y_z^C, y_j\right)$ with $y_z^C \in Y^C$ and $y_j \notin Y^C$ with $\left(y_{z+1}^C, y_j\right)$. The new graph is F'. It is easy to show that $q(\mu, F, \rho) = q(\mu', F', \rho)$, so that $(\mu', F') \in \Phi(\rho)$.

As these three cases are exhaustive, we have established that $\Phi(\rho)$ has at least two elements.

We now turn to the only if statement. Suppose that $(\mu, F), (\mu', F') \in \phi(\rho)$ and that $q_i(\mu, F, \rho) > r_i$ for $y_i \notin R(F)$.

If $E(F') \nsubseteq E(F)$, then there exists $(y_2, y_1) \in E(F')$ such that $(y_2, y_1) \notin E(F)$. Since $(\mu', F') \in \phi(\rho)$, there exists x_1 and x_2 such that $\mu'_1 = y_1$ and $\mu'_2 = y_2$. Suppose that $\mu_2 = \mu'_2$, then $y_1 \neq \mu_2$ and y_1 is not a successor to μ_2 in F. Suppose now that $\mu_2 \neq \mu'_2$. If $(\mu_2, y_2) \notin F$, x_2 and y_2 satisfy the theorem. If $(\mu_2, y_2) \in F$, then there must exist x_3 such that $y_3 = \mu'_3 \neq \mu_3 = y_2$. If $(\mu_3, y_3) \notin F$, x_3 and y_3 satisfy the theorem. If $(\mu_3, y_3) \in F$, we precede as before. This creates a sequence $\{y_2, y_3, \ldots\}$. As Y is finite and F is a tree, this sequence must end at some point with $(\mu_n, y_n) \notin F$. In this case, x_n and y_n satisfy the theorem.

We now show that $E(F')$ is not strictly contained in $E(F)$. Suppose now that $E(F') \subset E(F)$. Then there exists some $(y_i, y_k) \in E(F)$ such that $(y_i, y_k) \notin E(F')$. As F is a tree, each good has at most one direct predecessor. Hence y_k is a root good in F'. $q_k(\mu', F', \rho) = q_k(\mu, F, \rho) > r_k$ by assumption. It follows that $y_k \notin R(F)$, a contradiction.

Finally suppose that that $F = F'$. Then $\mu \neq \mu'$. It is easy to see that the lemma holds in this case (the only comlication is when μ and μ' differ on goods with no successors in F; in which case the buyers must be indifferent.)$\square$

Lemma 9: Under conditions of the preceding lemma. $|\Phi(\rho)| = 2$ if there exists only one such x_b and y_0.

Proof: We now show that this implies that there are only two elements of $\Phi(\rho)$. We again consider the three cases.

Case 1: Let $\hat{Y}$ denote the component of F containing $\mu(x_b)$ and let y_R denote the root good in $\hat{Y}$, and $\hat{X}$ the set of buyers allocated to $\hat{Y}$ by μ. By assumption $q_i(\mu, F, \rho') > r_i$ for $y_i \in \hat{Y} \backslash y_R$. It follows that no $y_i \in \hat{Y} \backslash y_R$ can be a root good for an F'' for which $(\mu'', F'') \in \Phi(\rho')$. Moreover, the definition of competitive equilibrium implies that for $(\mu'', F'') \in \Phi(\rho')$, μ'' must be onto $\hat{Y} \backslash y_R$.

There is no $x_c \notin \hat{X}$ and $y_1 \in \hat{Y}$, such that $y_1 \in D_c(q(\mu, F, \rho'))$. Since $(\mu_c, y_1) \notin F$, the preceding establishes that $y_1 \notin D_c(q(\mu, F, \rho))$. The statement follows from the assumption $D_a(q(\mu, F, \rho')) = D_a(q(\mu, F, \rho))$ for all $x_a \neq x_b$ and the observation that $x_c \neq x_b$. It follows that only buyers in $\hat{X}$ demand goods in $\hat{Y} \cup y_0$. As $|\hat{X}| = |\hat{Y}|$ and any μ'' with $(\mu'', F'') \in \Phi(\rho')$ must be onto $\hat{Y} \backslash y_R$, no μ'' with $(\mu'', F'') \in \Phi(\rho')$ can be onto both y_R and y_0. It is easy to show that there is only one GA-structure with y_0 as a root good and one with y_R as a root good.

Case 2: In this case, there is only one way to price all goods that are not y_0 or successors of y_0. There is exactly two ways to price y_0, and given the price of y_0 there is only one way to price its successors. Again there are exactly two GA-stuctures.

Case 3: In this case, there is only one way to price all goods that are not successors of y_0. There are two choices of which buyer to allocate to y_0: the allocation under μ and x_b. Given the allocation to y_0 there is a unique allocation to the other goods. This implies two GA-structures.

This completes the lemma.□

We now prove that the q_i are analytic functions of z.

Lemma 10: Given Assumptions A and B, $q_i(\mu, F, \rho(z))$ is an analytic function of z for all $y_i \in Y$, and $(\mu, F) \in \mathcal{G}$, and $\underline{p}_i(z)$ is analytic except at a finite number of points.

Proof: We prove $q_i(\mu, F, \rho(z))$ is analytic by induction. Consider first $y_i \in R(F) \equiv A_0$. $g_i = r_i$. Assumption B states that r_i is an analytic function of z. Now suppose that for $y_i \in A_n$, g_i are analytic functions of z. 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 implictly by the indifference of x_a:

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

where according to Assumptions A and B $\xi_{ai}(z)$, $\xi_{aj}(z)$, and $q_j(z)$ are analytic functions of z and $U_a(y_j, q, \xi_{ai})$ and $U_a(y_i, q_i, \xi_{aj})$ are analytic functions of their second and third arguments. Assumption A ensures that U_a is strictly decreasing in q_i. It follows from the Real Analytic

Implicit Function Theorem (Krantz and Parks, 2002, p. 35) that q_i is an analytic function of z. This completes the induction step and the proof.

We now prove that $\underline{p}_i(\rho(z))$ is analytic except at a finite number of points. Suppose not. Let $A = \{z \in [0,1] | \phi(z)$ is not analytic at $z\}$. $|A| = \infty$.

Consider a neighborhood Ω of $z_1 \in A$. Suppose that $\underline{p}_i(\rho(z)) = q_i(\mu, F, \rho(z_1))$. If $\underline{p}_i(\rho(z)) = q_i(\mu, F, \rho(z))$ for all $z \in \Omega$ then $\underline{p}_i$ is analytic at z. It follows that there exists there exists $z_2 \in \Omega$ such that $\underline{p}_i(\rho(z_2)) \neq q_i(\mu, F, \rho(z_2))$. By Lemma 1, there exists (μ', F') and $\underline{p}_i(\rho(z_2)) = q_i(\mu', F', \rho(z_2))$. Construct a sequence Ω_n converging to z. This leads to a sequence (μ'_n, F'_n). Since there are a finite number of GA-structures and since and Φ is upperhemicontinuous (Lemma 5), it follows that $q_i(\mu, F, \rho(z_1)) = q_i(\mu', F', \rho(z_1))$ and there exists $z \in [0,1]$ such that $q_i(\mu, F, \rho(z_1)) \neq q_i(\mu', F', \rho(z_1))$

The above argument applies at all $z \in A$. Given that $|A| = \infty$ and the number of GA-structure is finite, there exist two GA-structures (μ_1, F_1) and (μ_2, F_2) such that $q_i(\mu_1, F_1, \rho(z)) = q_i(\mu_2, F_2, \rho(z))$ at an infinite number of z and $q_i(\mu_1, F_1, \rho(z)) \neq q_i(\mu_2, F_2, \rho(z))$ for some $z \in [0,1]$. Since the q_i are analytic, that that $q_i(\mu_1, F_1, \rho(z)) = q_i(\mu_2, F_2, \rho(z))$ at an infinite number of z implies that $q_i(\mu_1, F_1, \rho(z)) = q_i(\mu_2, F_2, \rho(z))$ for all $z \in [0,1]$. This contradiction establishes the result.□

We are now in a position to prove the Theorem.

Proposition (Theroem 2, Part 1) Given a path $(\rho_0, S(z))$, the set $Z = \{z \in [0, \bar{z}] || \phi(z)| > 1\}$ is finite for almost all $\rho_0 \in \Lambda_S$.

Proof: $\bar{\Lambda}$ denote the set of ρ_0 such that Z is not finite. We show that given any $\rho_0 \in \bar{\Lambda}$, we can perturb ρ_0 and obtain a path in which Z is finite. Given the upper hemicontinuity of Φ, there exists an open set such that Z is finite and ρ_0 lies on the boundary of this set.

Suppose that, given ρ_0, $Z = \{z \in [0, \bar{z}] || \phi(z)| > 1\}$ is infinite. Given that the number of GA-structures is finite, there exist two GA-structures (μ_1, F_1) and (μ_2, F_2) such that the set $\hat{Z} \equiv \{z \in [0,1] | (\mu_1, F_1), (\mu_2, F_2) \in \phi(z)\}$ is infinite.

As $\hat{Z} \equiv \{z \in [0,1] | (\mu_1, F_1), (\mu_2, F_2) \in \phi(z)\}$ is infinite, it has an accumulation point in [0,1]. $q(\mu_1, F_1, \phi(z))$ and $q(\mu_2, F_2, \phi(z))$ are both analytic by Lemma 10. It follows that $q(\mu_1, F_1, \phi(z)) = q(\mu_2, F_2, \phi(z))$. Lemma 10 also implies that $\underline{p}$ is piecewise analytic with a finite set of non-analytic points. Repeating the arguments above, it follows that there exists an interval $[z_1, z_2]$ such that $\underline{p}(\rho) = q(\mu_1, F_1, \phi(z)) = q(\mu_2, F_2, \phi(z))$ for $z \in [z_1, z_2]$.

Suppose that $q_i(\mu_1, F_1, \rho(z)) > r_i(\rho(z))$ for $y_i \notin R(F_1)$ and $z \in [z_1, z_2]$. Lemma 8 implies that there exists x_b and $y_0 \in D_b(q(\mu, F, \rho)$ such that $y_0 \neq \mu_b$ and $(\mu_b, y_0) \notin E(F)$. Consider a perturbation of ρ_0 in we reduce the element associated with ξ_{b0} by ε_{b0} for all such y_0. Call this perturbation ρ_1 and let $\phi_1(z)$ denote the new path originating from ρ_0. It follows from Lemma 8 that $(\mu_2, F_2) \notin \Phi(\rho_1(z))$ for $z \in [z_1, z_2]$.

Suppose that $q_i(\mu_1, F_1, \rho(z)) = r_i(\rho(z))$ for some $y_i \notin R(F_1)$. In this case we consider a perturbation in which we increase the component of ρ_0 associated with r_i. This perturbation implies that $(\mu_1, F_1), (\mu_2, F_2) \notin \Phi(\rho_1(z))$ for $z \in [z_1, z_2]$. Again the instance of infinite crossings are eliminated.

As there are a finite number of pairs of GA-structures, we can eliminate all instances in which two intersect $\underline{p}$ infinitely often, with a finite number of such perturbations.□

Proposition (Theorem 2, Part 2) Along the generic path, $\phi(z)$ takes one or two values.

Proof: Given a path $(\rho_0, S(z))$, let $Z(\rho_0) = \{z \in [0, \bar{z}] || \phi(z)| > 2\}$ denote the set of points at which $\phi(z)$ takes more than two values. Theorem 2 Part 1 implies that $Z(\rho_0)$ is finite for almost al ρ_0. Consider ρ_0 such that $Z(\rho_0)$ is not empty. We construct a perturbation of ρ_0, ρ_0', such that $Z(\rho_0')$ is empty.

Consider ρ_0 such that $Z(\rho_0)$ is not empty. Let $Z(\rho_0) = \{z_1 \ldots z_k\}$ were $z_i < z_{i+1}$ for $i \in \{1, \ldots k-1\}$. Given $(\mu, F) \in \phi(z)$, call (μ, F) tangent to the price path $\underline{p}(\rho)$ if there exists a neighborhood W of z such that $q(\mu, F, z) = \underline{p}(z)$ and either $q(\mu, F, \rho) > \underline{p}(\rho)$ for all $\rho \in W \backslash z$ or $q(\mu, F, \rho) < \underline{p}(\rho)$ for all $\rho \in W \backslash z$. Let $T(z)$ denote the set of GA-structures that are tangent at z. Define the count $C(\rho_0) = \sum_i (|\phi(z_i)| - 2 + (|T(z)| - 1) \vee 0$.

Consider z_1. $|\phi(z_1)| = b_1 > 2$.

Suppose first that $q_i(\mu, F, z_1) > r_i$ for all $y_i \notin R(F)$. As $|\phi(z_1)| > 1$, Lemma 8 implies that there exists x_b and $y_0 \in D_b(q(\mu, F, \rho)$ such that $y_0 \neq \mu_b$ and $(\mu_b, y_0) \notin E(F)$. Let V_b denote the set of such (x_b, y_0). Since $|\phi(z_1)| > 2$, Lemma 9 implies that $|V_b| > 1$. Choose $(x_c, y_1) \in V_b$. For all $(x_b, y_0) \in V_b \backslash (x_c, y_1)$, reduce the element of ρ_0 associated with ξ_{b0} by ε_{b0}. Call this ρ_1. It follows that $y_0 \notin D_b(q(\mu, F, \rho_1)$ and $y_1 \in D_c(q(\mu, F, \rho_1)$. If the ε_{b0} are small enough, there will be no other x_b and $y_0 \in D_b(q(\mu, F, \rho)$ such that $y_0 \neq \mu_b$ and $(\mu_b, y_0) \notin E(F)$. It follows that $|\phi(z_1)| = 2$.

We now show that $C(\rho_1) < C(\rho_0)$. Consider $z_i \in Z(\rho_0)$. Suppose $|\phi(z_i)| = d_1$ and $|T(z_i)| = d_2 < d_1$. $\phi(z_i) = \{(\mu_j, F_j)\}_{j=1}^{d_1}$. Let $\phi_1(\rho)$ and $\underline{p}_1$ be associated with ρ_1. Since $\Phi(\rho)$ is upper-hemicontinuous and the change in the ε_{0b} is small. There exists a neighborhood of z_i such that $\phi_1(z_i) \subseteq \phi(z_i)$. Given the continuity of the $g(\mu, F, \rho)$, $(\mu, F) \in T(z_i)$ may cross $\underline{p}_1$ at most twice and other elements of $\phi(z_i)$ at most once. Suppose that there are k intersections. Let T_1

denote the number elements of $T(z)$ that remain tangent and T_2 the number that cross twice. $|\phi(z_i)| + T_2 - 2k + (T_1 - 1) \vee 0$. This is minimized if k is minimized. There are two cases. If $T_2 = 0$, then k is minimized at 1. $|\phi(z_i)| - 2 + (T_1 - 1) \vee 0 \leq (|\phi(z_i)| - 2 + (|T(z)| - 1)$. If $T_2 > 0$, then k is minimized at 2. $|\phi(z_i)| + T_2 - 4 + (T_1 - 1) \vee 0 \leq (|\phi(z_i)| - 2 + (|T(z)| - 1)$ as $T_1 + T_2 \leq T(z)$. Note that at z_1, $k = 2$. It follows that $C(\rho_1) < C(\rho_0)$.

Suppose now that $q_i(\mu, F, z_1) = r_i$ for some $y_i \in Y \backslash R(F)$. If $V_b \neq \varnothing$, then reduce the component of ρ_0 associated with r_i by ε_i for each $y_i \in Y \backslash R(F)$. At the new ρ_0, $q_i(\mu, F, z_1) > r_i$ for all $y_i \notin R(F)$, and the preceding applies. Let ρ_1 denote the sum of these perturbations. If $V_b = \varnothing$, reduce all but one of the r_i, to arrive at ρ_1. At ρ_1, $|\phi(z_1)|= 2$.

We repeat the argument at the new ρ_1. At each step we reduce C. C falls to zero in a finite number of steps. Let ρ_n the final initial condition. $\phi(\rho)$ takes on one or two values at this point.□

Proposition (Theorem 2, part 3) Along the generic path, $\phi(\rho)$ has no isolated points.

Proof: Along the generic path $|\phi(\rho)| \leq 2$. Suppose that $\phi(z) = \{(\mu, F), (\mu', F')\}$ and that for $\rho \in W \backslash z$, where W is a neighborhood of z, $\phi(z) = (\mu, F)$.

Suppose first that $q_i(\mu, F, z) > r_i$ for all $y_i \notin R(F)$. Since, $q_i(\mu, F, z) > r_i$ and $|\phi(z)| = 2$, Lemma 8 implies that there exists x_b and $y_0 \in D_b(q(\mu, F, z)$ such that $y_0 \neq \mu_b$ and $(\mu_b, y_0) \notin E(F)$. Reduce the component of ρ_0 associated with $\xi_{b)}$ by ε. Consider the four cases considered in Lemma 8. In cases, 2 and 3, this perturbation reduces $q(\mu', F', \rho)$. Since $\mu = \mu'$ in theses cases, it follows from Theorem 1, that (μ', F') is no longer in $\phi(z)$. In cases 1 and 4, the perturbation raises $q(\mu', F', \rho)$. Suppose F'' maximizes $q(\mu', F, z)$ after the purturbation. The upperhemicontinuity of Φ implies that $(\mu', F'') \in \phi(z)$. It follows that F' maximizes $q(\mu', F, z)$. By theorem 1, (μ', F') is no longer in $\phi(z)$ after the perturbation.

Now suppose first that $q_i(\mu, F, z) = r_i$ for all some $y_i \notin R(F)$. There can be only one such y_i since $|\phi(\rho)| = 2$. Perturb ρ_0 by reducing r_i.□

Proposition (Theorem 2, part 4) If $|\phi(\rho)| = 1$, then $q_i(\mu, F, z) > r_i$ for all $y_i \notin R(F)$.

Proof: Immediate.

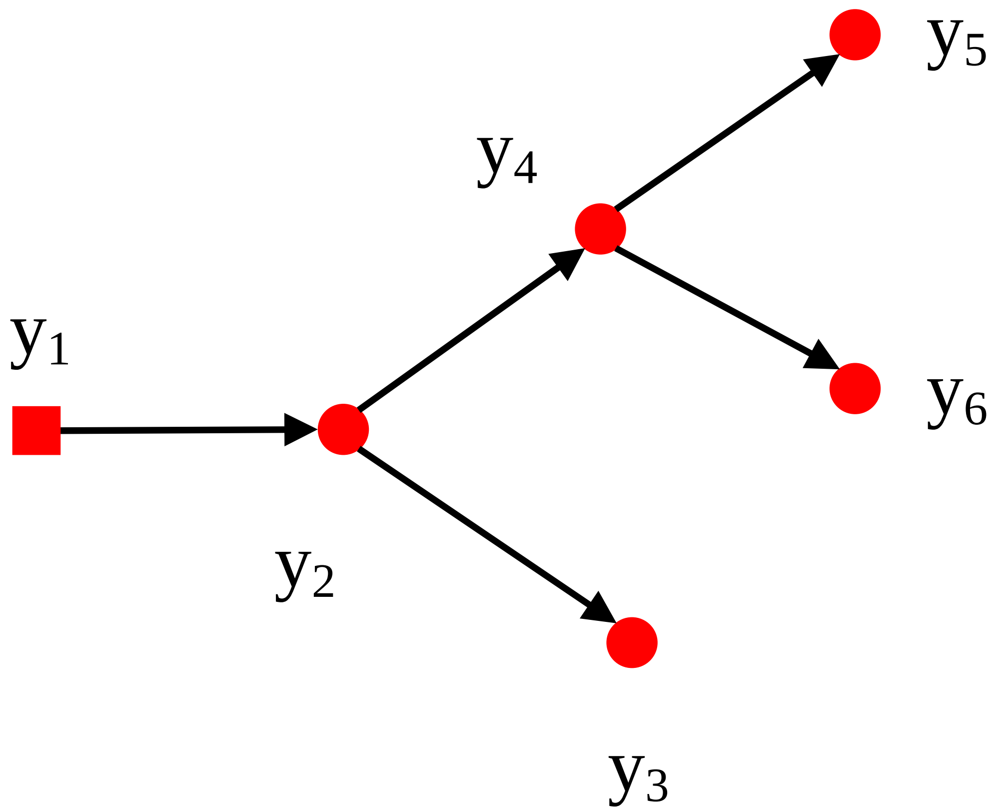

Figure 1: A directed rooted tree with edges directed away from the root good (y_1)

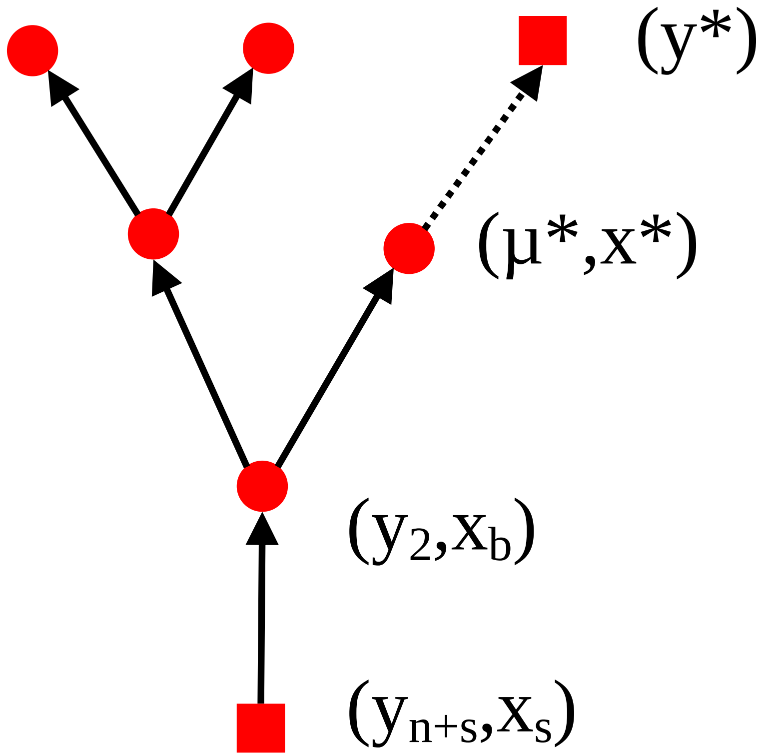

Figure 2(a)

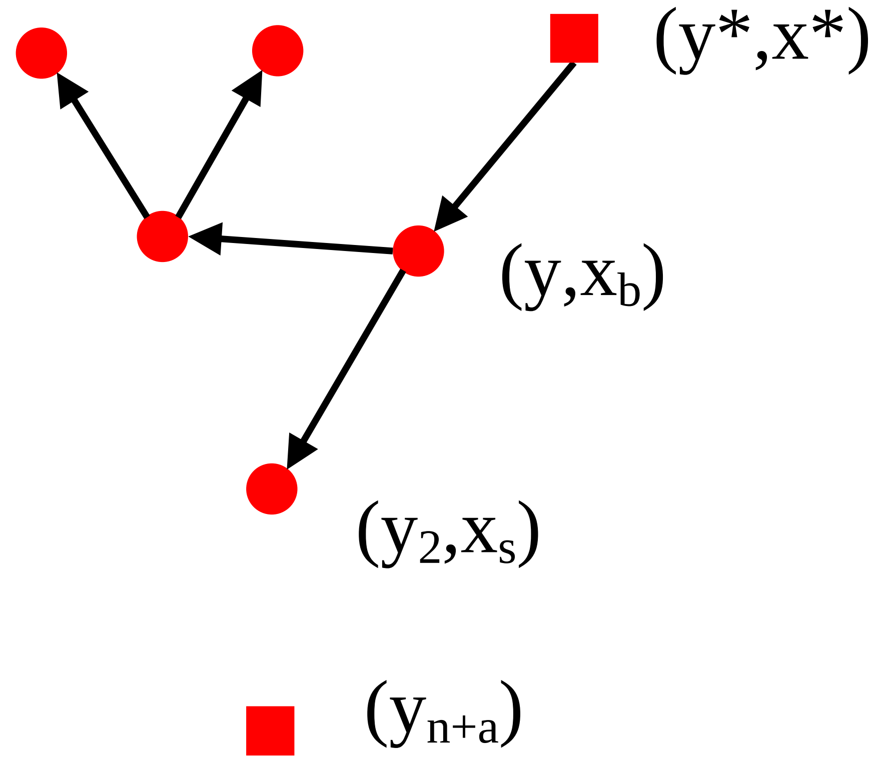

Figure 2(b)

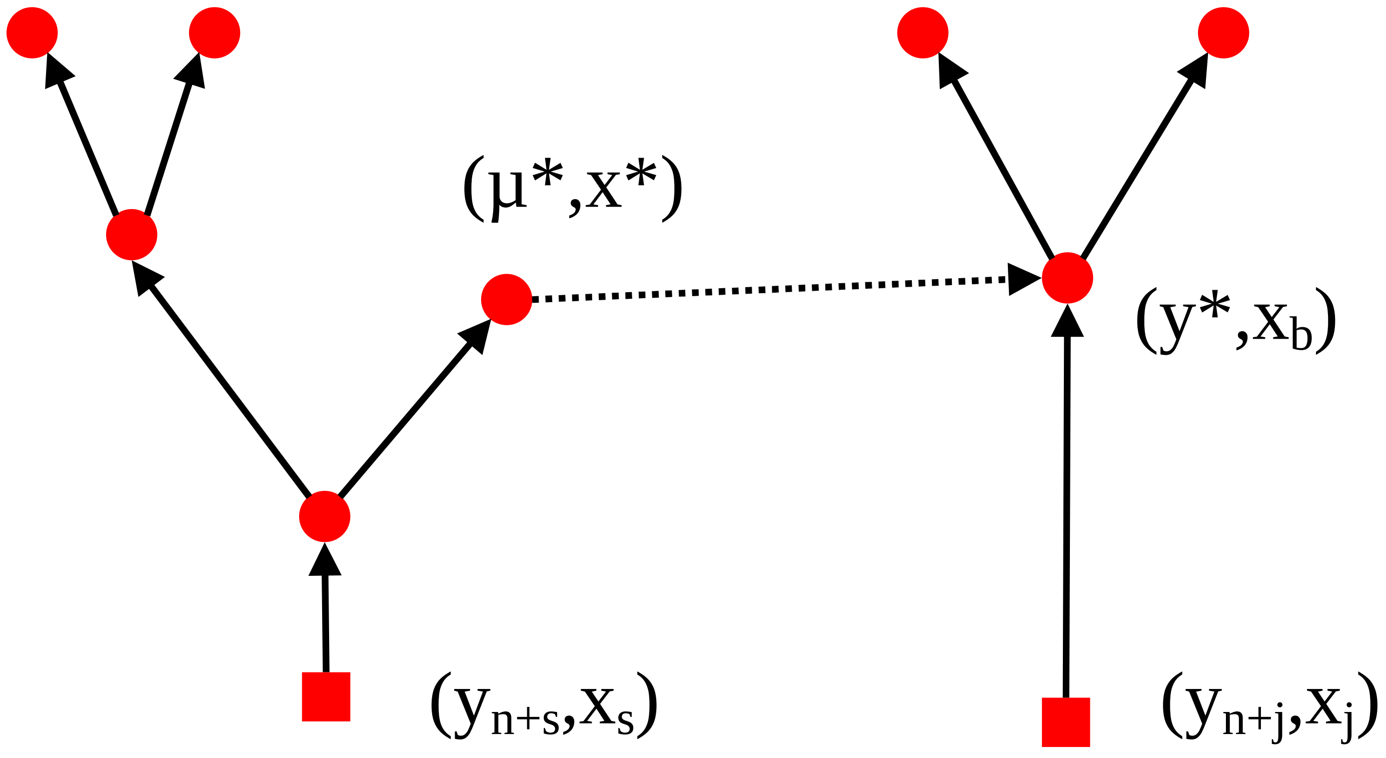

Figure 3(a)

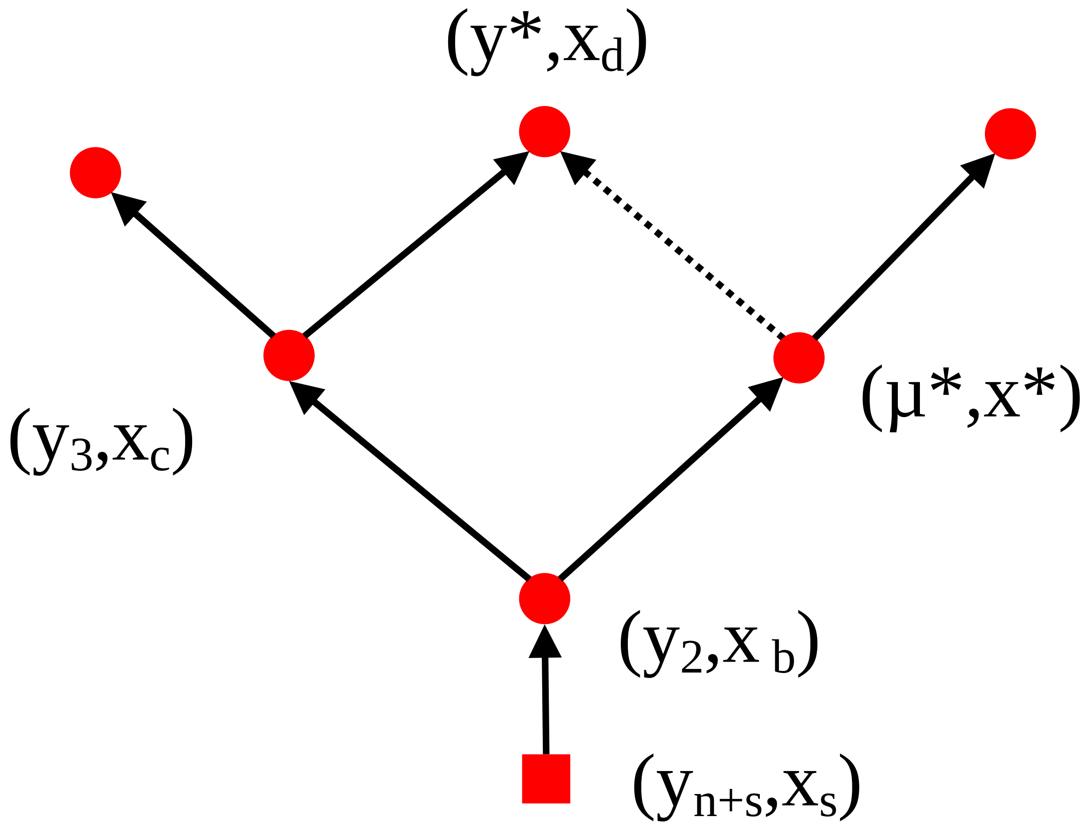

Figure 3(b)

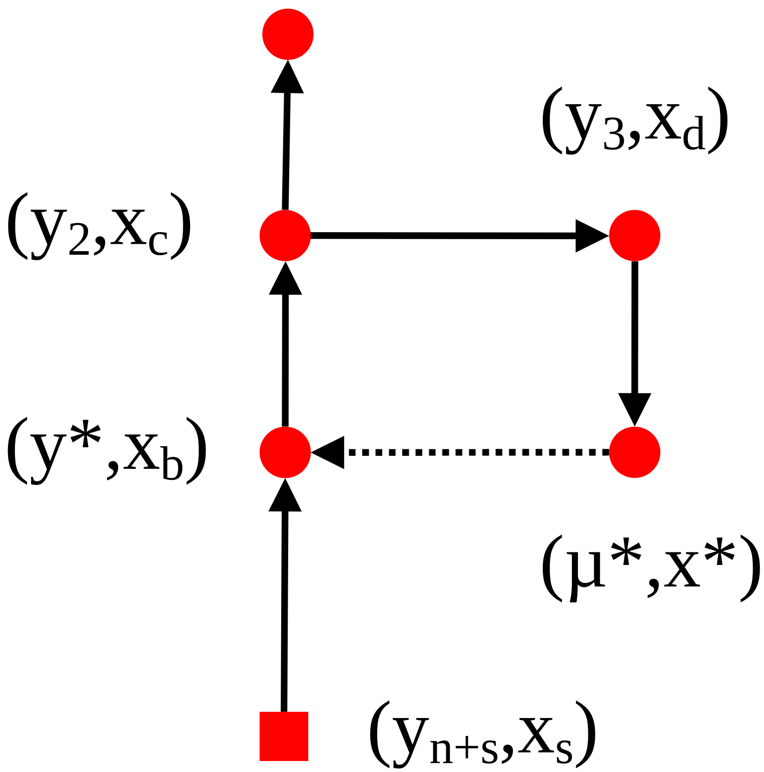

Figure 4(a)

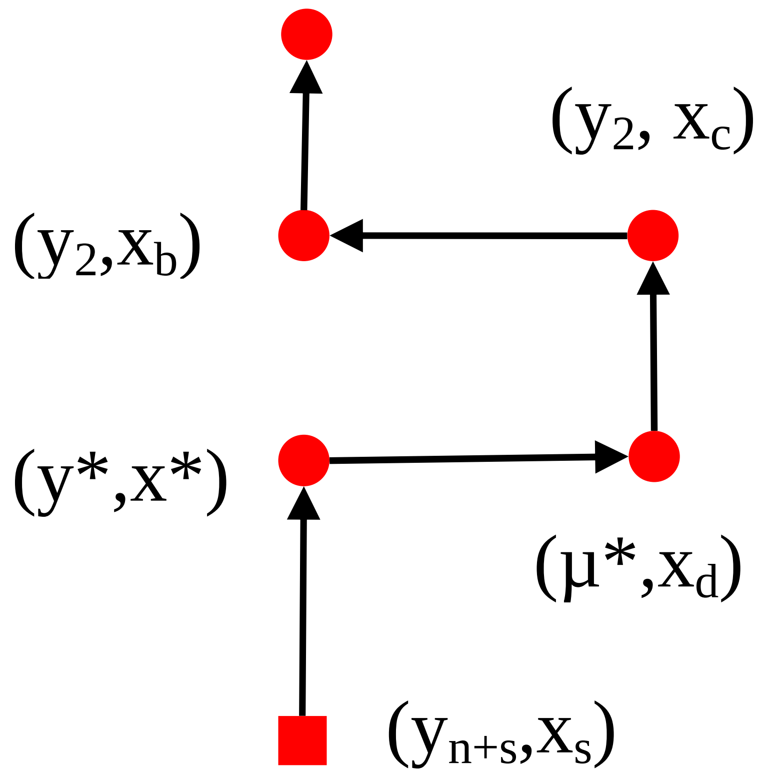

Figure 4(b)