Modeling House Values: Extremal Equilibria in Allocation Markets with Non-Transferable Utility

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

June 2009

Abstract

Wealth effects impact the valuation and allocation of large indivisible goods such as houses. Yet no general method exists for quantifying these effects. We introduce an algorithm for identifying equilibria in allocation markets with non-transferable utility, thereby filling this gap. In cases with multiple equilibria, the algorithm can be used to solve for maximum and minimum equilbrium prices, and thereafter to uncover the full set of equilibria.

1. Introduction

Houses and other large durable goods are indivisible, heterogeneous, and expensive. Each of these characteristics poses a challenge for model builders. Indivisibility imposes integer constraints. Heterogeneity leads to a non-trivial problem of matching houses and buyers. Expense makes it difficult to justify a local linear approximation to buyer utility.

The difficulty in modeling the impact of wealth has particularly constrained economic modeling of house values, as transferable utility greatly simplifies the analysis of allocation markets. If utility is linear in money, 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

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

problem takes the form of the linear programming problem studied by Koopmans and Beckman (1957). This programming problem is, in turn, easily solved by the Hungarian algorithm of Kuhn (1955) and Munkres (1957) or by the ascending auction mechanism of Demange, Gale and Sotomayor (1986).

The assumption of transferable utility, however, may not lead to reasonable outcomes in housing markets. 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 would 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.[1] 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. The inability to solve for the set of equilibria has held back practical application of these models.

In this paper, we present an algorithm for constructing minimum price competitive equilibria in allocation markets with non-transferable utility. The 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 not completely alter the balance between supply and demand, while at the same time keeping track of the resulting changes in the allocation.

At its base, our algorithm rests on recognition of the graph theoretic structure of minimum price equilibria. Minimum price equilibria are characterized by chains of indifference: if any price is above a seller's reservation value, then there must be a potential buyer, other than the buyer allocated the good, who would demand the good at any lower price. Any lower price will

[1] 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.

then lead to excess demand. In addition, these chains of indifference must begin somewhere: they must be rooted in a good that is priced at reservation or supported by a buyer that has exited the market. Otherwise we could reduce prices along the entire chain.

We show how to build these chains of indifference by progressively adding buyers to the market and tracking the resulting change in the allocation. Our algorithm begins by artificially raising the utility that each buyer receives from purchasing nothing until no purchase is a competitive equilibrium. We then reduce this reservation utility one buyer at a time to its true level. Along the way we keep track of minimum equilibrium prices and the resulting chains of indifference, as well as the equilibrium allocation. We call the algorithm the GAME-algorithm since it uses Graphs and Allocations to track Minimum Equilibrium prices.

There are several nice features of the GAME-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)).

Second, the algorithm is likely to be relatively fast in many practical applications. We show that for several useful classes of buyer utility, such as when utility is exponential in wealth or log in wealth, the algorithm does not revisit previous allocations as reservation utilities are lowered. More generally, there is a sense in which the algorithm is minimal: it searches only through the set of potential solutions for some set of reservation utilities 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.

Third, it can be used to characterize the entire set of equilibria. The minimum price equilibrium provides one bound. The other bound, the maximal price equilibrium, is the solution to a natural dual to the allocation problem. We show that each equilibrium in between these two extremal equilibria is a minimum price equilibrium with sellers' reservation prices set between the values that buyers pay in these two extremes.

Fourth, the algorithm provides insights into the comparative static structure of minimum price equilibria. For each good with prices above seller's reservation, the algorithm identifies a potential buyer who is not allocated to the good, but would demand the good at any lower price. It also identifies a set of "root" goods, some of which are at reservation price, others of which represent the outside options of buyers on the verge of entry into the market. The equilibrium set of root goods and the indifference relations characterize local comparative statics. Small changes in model parameters work their way through the chain in a well-defined order, while larger changes may change the equilibrium allocation. The equilibrium structure of these chains says much about how shocks to one sector of the market propagate to others, and about which sectors of the market are interconnected and which are distinct.

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.[2] 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) in which they show that in the minimum price equilibrium buyers would truthfully report their preferences.

The remainder of the paper is structured as follows. Section 2 presents the model of a market for a collection of heterogeneous, indivisible goods. Section 3 presents two graph theoretic characterizations of equilibrium prices. The first provides necessary and sufficient conditions for the existence of a competitive equilibrium allocation. It is an extension of Hall's theorem on the existence of an allocation. The second theorem strengthens the first. It provides necessary and sufficient conditions for the existence of a minimum price competitive equilibrium allocation. Section 4 introduces the graph theoretic basis for our algorithm. We show how to calculate prices given an allocation and a chain of indifference. The section ends with a theorem that provides necessary and sufficient conditions for an allocation and a chain of indifference to support a minimum price equilibrium. Section 5 presents the algorithm and proves that it identifies the minimum price equilibrium. Section 6 presents two classes of utility functions for which the algorithm has strong convergence properties. Section 7 shows how to adapt the algorithm to solve for the maximum price equilibrium in the dual problem, and thereafter to compute all equilibria. Section 8 concludes with directions for future work.

2. The Model

There is a set of buyers $x_a \in X$, $1 \leq a \leq m$, and a set of indivisible goods $y_i \in Y$, $1 \leq i \leq n$. There is also a homogeneous, perfectly divisible, numeraire good, which may be thought of as money.

The indivisible goods are initially held by sellers. Each seller wishes to obtain the highest possible price above a reservation level. Let $p \in \mathbb{R}^n_+$ and $r \in \mathbb{R}^n_+$ denote the vectors of indivisible goods prices and seller reservation prices respectively, with p_i the price of y_i and r_i the reserva-

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

tion price of its seller.[3] The supply side is trivial: each seller prefers to hold for any $p_i < r_i$ and to sell for any $p_i > r_i$. The seller is indifferent when $r_i = p_i$.

Each buyer is endowed with a stock of the numeraire good. Let $w \in \mathbb{R}^m_{++}$ denote the vector of endowments, with $w_a > 0$ denoting the endowment of buyer x_a. Each buyer's utility depends on the indivisible good that they purchase, as well as their holdings of the numeraire good. We assume that buyers can derive utility from at most one element of Y. They also have the option of purchasing nothing and exiting the market at no cost. We represent this choice by a set of null goods $Y_\varnothing = \{y_{n+1}, \dots, y_{n+m}\}$. Good y_{n+a} represents the outside option of buyer x_a. Let $\bar{Y} = Y \cup Y_\varnothing$ represent the expanded goods space and $\bar{Y}_a = Y \cup \{y_{n+a}\}$ represent the goods available to x_a. We normalize the prices and the reservation prices of the goods in $Y_\varnothing$ to zero. Overall, we can limit attention to a set Π in the search for competitive equilibria,

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

The payoff for buyer x_a is summarized by the utility function $U_a : \bar{Y}_a \times \mathbb{R}_+ \to \bar{\mathbb{R}}$. It will be useful to define a vector of reservation utilities, $\bar{u}^R \in \mathbb{R}^m$, corresponding to the utility that each buyer receives from the outside option:

$$\bar{u}^R_a \equiv U_a(y_{n+a}, w_a).$$

Given any price vector $p \in \Pi$, the (non-empty) demand correspondence $D_a(p)$ specifies members of $\bar{Y}_a$ that are affordable and maximize utility:

$$D_a(p) = \left\{y_i \in \bar{Y}_a | p_i \leq w_a \text{ and } U_a(y_i, w_a - p_i) \geq U_a(y_k, w_a - p_k) \text{ for all } y_k \in \bar{Y}_a \text{ with } p_k \leq w_a\right\}.$$

An *allocation* is a one-to-one mapping $\mu : X \to \bar{Y}$ from buyers to goods such that each buyer chooses from their feasible set, $\mu(x_a) \in \bar{Y}_a$. It simplifies later notation to let μ_a denote the good assigned to buyer x_a according to the allocation μ,

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

The set of all allocations is M.

A competitive equilibrium comprises a price vector and an allocation such that all buyers choose optimally and all goods with prices above their reservation level are allocated. Given

[3] $\mathbb{R}^n_+$ denotes the non-negative real numbers and $\mathbb{R}^n_{++}$ denotes the positive real numbers.

$p \in \Pi$, let $H(p) \equiv \left\{y_i \in \bar{Y} | p_i > r_i\right\}$ denote the set of goods with prices above seller reservation levels.

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

1. $\hat{\mu}_a \in D_a(\hat{p})$ for all $x_a \in X$.
2. If $y_i \in H(\hat{p})$, 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 subject to affordability. The second states that all goods with prices above reservation must be allocated. This ensures that supply is equal to demand.

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

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

We make assumptions on preferences that guarantee the existence of an equilibrium (Kaneko, 1982). Note that these assumptions are not used in full force in the next section.

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

1. $U_a(y_i, c)$ is a continuous and strictly increasing function of c for each $y_i \in \bar{Y}_a$.
2. $U_a(y_{n+a}, w_a) \geq U(y_i, 0)$.

The first assumption is a straight forward regularity assumption.[4] The second assumption states that agents would prefer to exit the market rather than spend all of their endowment in purchasing one of the indivisible goods.[5] 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.

In addition to these assumptions, we will need to make several genericity assumptions. We postpone the discussion of these assumptions until we develop the structure of the algorithm more completely.

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

[5] This assumption eliminates discontinuities in utility that arise from liquidity constraints. Note that this condition is satisfied when utility functions satisfy Inada conditions.

3. The Characterization Theorems

Our algorithms are based on properties of the demand graph associated with equilibrium and minimum equilibrium prices. Given a price vector, this graph links buyers with the goods in their demand set.

Definition Given a price vector $p \in \Pi$, the *demand graph* G_p is the bipartite graph with vertex set $V = X \cup \bar{Y}$, bipartition $X, \bar{Y}$, and edge set E_p where

$$E_p = \{(x_a, y_i) \in X \times \bar{Y} | \ y_i \in D_a(p)\}.$$

Given $A \subset X$, let $\Gamma_p(A) \subset Y$ denote the set of neighbors of A in G_p,

$$\Gamma_p(A) = \{y_i \in \bar{Y} | (x_a, y_i) \in E_p \text{ for some } x_a \in A\}.$$

Similarly, given $B \subset \bar{Y}$, let $\Gamma_p(B)$ denote the set of neighbors of B in G_p,

$$\Gamma_p(B) = \{x_a \in X | (x_a, y_i) \in E_p \text{ for some } y_i \in B\}.$$

A useful tool for proving the existence of allocations is Hall's Theorem.[6] This theorem states that an allocation μ that maximizes buyers' utility at price p exists if and only if $|\Gamma_p(A)| \geq |A|$ for all $A \subset X$, where $|S|$ denotes the cardinality of the set S. Therefore, if we can find a price vector such that the demand graph satisfies this condition, we can immediately conclude that there is an allocation $\mu \in M$ that is optimal from the buyer viewpoint.

Competitive equilibrium requires both that the allocation μ is optimal from the buyers' viewpoint given $p \in \Pi$ and that it is onto the set $H(p)$ of goods with prices above their reservation value. Theorem 1 provides an extension of Hall's theorem characterizing existence of a stable match that is onto a particular subset of $\bar{Y}$. Note that this result does not require Assumption A. All proofs are contained in the Appendix.

Theorem 1 A price vector $p \in \Pi$ is a competitive equilibrium price vector if and only if,

1. $|\Gamma_p(A)| \geq |A|$ all $A \subset X$;
2. $|\Gamma_p(B)| \geq |B|$ all $B \subset H(p)$.

[6] Demange, Gale and Sotomayor (1986) base their auction mechanism on Hall's theorem.

This theorem is entirely "natural" and in the spirit of Hall's theorem itself. In order for there to exist an equilibrium allocation, it is clearly necessary that Hall's theorem apply not only in terms of buyer preferences, but also in terms of seller preferences. What this means is that there must be enough buyers interested in any subset of $H(p)$ for it to be possible to create an allocation match that is onto this set. Hence the conditions of Hall's theorem must be satisfied not only for buyers, but also for above-reservation-price goods. What the theorem establishes is that the combination of Hall applied to buyers and to above-reservation-price goods is not only necessary for a given price vector to be an equilibrium, but also sufficient.

To characterize the minimum equilibrium price, we strengthen the condition on the seller side, insisting that any set of goods $B \subset H(p)$ is connected with a strictly larger set of buyers. Note that we invoke strict monotonicity of the utility function in this proof.[7] Note that a converse theorem characterizes the maximum equilibrium price, which has the property that all subsets of buyers $A \subset X$ who receive strictly more than reservation utility have neighborhood sets $\Gamma_p(A)$ of strictly higher cardinality.

Theroem 2 If $U_a(y_i, c)$ is strictly increasing in c for all $x_a \in X$, then $p = \underline{p}$ if and only if:

1. $|\Gamma_p(A)| \geq |A|$ all $A \subset X$;
2. $|\Gamma_p(B)| > |B|$ all $B \subset H(p)$.

4. GA-Structures and Minimum Equilibrium Prices

The algorithm is based on identifying a demand graph that satisfies the conditions of Theorem 2. Rather than searching for a (p, μ) pair such that Theorem 2 holds, we search through a related class of structures. This class has the advantage of being finite. These structures are motivated by the idea that, according to Theorem 2, the minimum equilibrium price is characterized by the fact that any subset of goods in $H(p)$ must be in demand by a buyer not allocated to that subset. We use this idea to construct "chains of indifference" built upon combining a specific class of graphs on the vertex set $\bar{Y}$ and an allocation μ. The graphs are all forests of directed, rooted trees in which all edges point away from the root, and in which all null goods are in the root set.[8]

[7] Demange and Gale (1985, lemma 4) established that $|\Gamma_p(B)| > |B|$ for all $B \subset H(p)$ is necessary for p to be a minimal equilibrium price.

[8] 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.

Definition The class $\mathcal{F}$ comprises all directed graphs F on vertex set $\bar{Y}$ with the following properties:

1. F is a forest of trees.
2. Each tree in F has a unique root good and each edge in $E(F)$ is directed away from the corresponding root.
3. The set of root goods $R(F) \subset \bar{Y}$ contains all null goods,

$$Y_{\varnothing} \subset R(F).$$

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 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 $y_i \in \bar{Y}$ 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 $y_i \in \bar{Y} \backslash R(F)$, 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 F and an allocation μ to create price vectors. To do this, we first limit attention to cases in which μ allocates buyers to all goods that are end-points of edges in $E(F)$.

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 $x_a, x_b \in X$ such that $\mu_a = y_i$ and $\mu_b = y_k$. 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, $q : \mathcal{G} \to \bar{\mathbb{R}}^n_+$, 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. In formalizing this inductive process, one must allow for infeasibility to be uncovered at some point in the computational process, either because the induced price of some good is below the seller's reservation or because it is above the buyer's ability to pay. To this end, we introduce a "null" price $p_\varnothing$ that has the interpretation that no price exists with the desired properties.

Definition Given Assumption A1 and $(F, \mu) \in \mathcal{G}$ we define the generated price $q(F, \mu) \in \Pi \cup p_\varnothing$ iteratively as follows:

1. Define $A_0 \equiv R(F)$ and set $q_i = r_i$ on A_0. Note that $q_i = r_i \equiv 0$ for $y_i \in Y_\varnothing$. If there exists x_a such that $\mu_a = y_i \in A_0$ and $w_a \leq r_i$, set $q(F, \mu) = p_\varnothing$ and stop the iteration.
2. Given $s \geq 0$, assume $q_i \geq r_i$ has been specified for all $y_i \in A_s \subset \bar{Y}$. Let S comprise direct successors of A_s,

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

Since F is a forest, any $y_k \in S$ has a unique corresponding $y_i \in A_s$ with $(y_i, y_k) \in E(F)$. Since $(F, \mu) \in \mathcal{G}$, there exists $x_a, x_b \in X$ with $\mu_a = y_i$ and $\mu_b = y_k$. Consider the following indifference condition:

$$U_a(y_i, w_a - q_i) = U_a(y_k, w_a - q_k). \tag{4.1}$$

It follows from Assumption A1 that if there exists a q_k that satisfies (4.1), it will be unique. If for each $y_k \in S$, there exists a $q_k \in [r_k, w_b]$ satisfying (4.1), define $A_{s+1} = A_s \cup S$ and proceed to the next step. Otherwise set $q(F, \mu) = p_\varnothing$ and stop the induction.

Since F is a forest, there is a unique path to any good from the root set. Hence this process will end after a finite number of steps with $A_s = \bar{Y}$ and with prices $q(F, \mu) \in \Pi \cup p_\varnothing$.

The following theorem is the critical result driving our algorithmic search for the minimum price equilibrium. It relates GA-structures to minimum price equilibria.

Theorem 3 Given Assumption A, (p, μ) is a minimum price equilibrium if and only if there exists $(F, \mu) \in \mathcal{G}$ with $q(F, \mu) = p$ and $\mu_a \in D_a(q(F, \mu))$ for each $x_a \in X$.

Given a model satisfying Assumption A, we know as a result of the necessity part of the theorem that there exists $F \in \mathcal{F}$ such that $(F, \mu) \in \mathcal{G}$ with $\mu_a \in D_a(q(F, \mu))$ and $q(F, \mu) = p$. Whereas there are a continuum of potential (p, μ) pairs, there are a finite number of allocations μ in M and a finite number of possible graphs F in $\mathcal{G}$.

One possible way forward is to use a brute force algorithm to identify the minimum price equilibrium, and to sample potential combinations (F, μ) until one finds a GA-structure that satisfies the conditions of Theorem 3. The problem with this strategy is that the set of GA-structures is very large. A simple thought experiment leads to a lower bound. Suppose that there are at least as many goods as buyers. Then with n goods and m buyers there are $n!/(n-m!)$ ways to allocate each buyer a good. A generalization of Cayley's theorem states that for each allocation there are

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

different rooted trees on the m allocated goods.[9] 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}. This calculation ignored the null goods, so the actual number of GA-structures is larger.

5. The GAME-Correspondence and GAME-Algorithm

Rather than undertake a massive unguided search through all GA-structures, our algorithm follows a highly structured path that is indexed by the level of reservation utility. The key insight underlying construction of this path is that the equilibrium in the full model, which may be highly complex, can be approached by solving a sequence of simpler models. Having found GA-structures that generate the minimum price equilibrium for the very simplest cases, one can follow a path through these structures that successively reduces the distance to the true model, all the while tracking the associated minimum price equilibrium, ultimately converging on the solution to the true model.

We now provide an intuitive description of the algorithm. In this description, we ignore complications that arise when there are multiple potential moves. We postpone discussion of these complications until after description of the algorithm. In the next section we present a genericity assumption that handles these cases, and is motivated to ensure that the intuitive description given below applies.

5.1. An Intuitive Description of the Algorithm

We initialize the algorithm by raising reservation utilities high enough so that all buyers exit the market. At this level of reservation utility, the minimum price equilibrium has all prices at

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

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. We then lower the reservation utility of each buyer, one at a time, and keep track of the GA-structure that corresponds to the minimum price equilibrium corresponding to that level of reservation utility.[10]

Given a GA-structure (F, μ), lowering the vector of reservation utilities v^R has the effect of raising the corresponding minimum equilibrium price vector. This is clear from the construction of q. Lowering the reservation utility of buyer x_s makes y_{n+s} less desirable. This increases x_s's willingness to pay for any direct successors of y_{n+s}, which increases the willingness of buyers allocated to those goods to pay for their direct successors, and so on. If we lower the reservation utility of only buyer x_s then only prices of successors of y_{n+s} rise. Moreover, given Assumption A1, this increase in prices will be continuous so long as the minimum equilibrium price is in the interior of Π.

We lower the reservation utility of x_s until either it reaches its true level $\bar{u}_s^R$ or some buyer becomes indifferent to some good that he did not previously demand. If it reaches $\bar{u}^R$ then we move on to lower the reservation utility of x_{s+1}. If some buyer becomes indifferent to some good that he did not previously demand, we exchange (F, μ) for a new GA-structure (F', μ'). The exact form that (F', μ') takes depends on the nature of the new indifference. Let x^* denote the buyer who becomes indifferent to a new good; μ^* the good assigned to that buyer at 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_{n+s} and its successors, as well as 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_{n+s} and y^* are both root goods. y_{n+s} is a root good because it is an element of $Y_\varnothing$. 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_{n+s}. In general, μ^* may also be y_{n+s}. μ^* must be y_{n+s} or one of its successors, since only buyers of these goods see the value of their goods fall (y_{n+s} because we are lowering reservation utility and the others because their prices are rising), and so only buyers of these goods alter their demands. Since y^* is initially unallocated, it is also a null tree in F. The dashed arrow

[10] The order in which we lower buyer's utility does not matter for convergence, since there is a unique minimum price equilibrium. All orderings converge to the same equilibrium. Ordering, however, may affect the speed of convergence.

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_{n+s} 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 (F', μ'), 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. The algorithm calls for a switch from the current (F, μ) to (F', μ'). Note that the prices induced by (F', μ') no longer depend on the reservation utility of x_s, since y_{n+s} is unallocated. As we lower the reservation utility of x_s further, x_s will strictly prefer y_2 to y_{n+s}, so (F', μ') captures the structure of demand. We continue to lower the reservation utility of x_s until it reaches its true level $\bar{u}_s^R$.

[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 (F', μ') by maintaining $\mu' = \mu$ and by replacing the edge (y_j, y^*) in F oriented towards y^* with the edge (μ^*, y^*). As we reduce the reservation utility of x_s further, the prices of successors of y_{n+s} will rise. In Figure 3(a), we see that x_j will no longer demand y^*, so that the (F', μ') accurately captures demands. The situation in Figure 3(b) is more complicated. As we have lowered the reservation utility to the level that generated the indifference of x^*, x^*'s willingness to pay for y^* has risen faster than x_c's. For the (F', μ') to capture demand as we lower reservation utility further, we need x^*'s willingness to pay for y^* to continue rising faster than x_c's. This amounts to ruling out tangencies in willingness to pay. We will argue below that in a wide class of models such tangencies are extremely rare. With such tangencies ruled out, we continue lowering the reservation utility of x_s until it either reaches $\bar{u}_s^R$ or a new indifference arises given the new GA-structure.

[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). Since y^* is a predecessor of μ^*, it must be a successor of y_{n+s}.[11] 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

[11] Note that y^* cannot be y_{n+s} as only x_s demands y_{n+s}.

circuit. Figure 4(b) depicts the GA-structure (F', μ') associated with this alternative allocation. At the given vector of reservation utilities, both GA-structures generate the same price vector. The algorithm calls for a switch from the current (F, μ) to (F', μ'). At this point we run into a situation similar to that of Figure 3(b), as we have lowered the reservation utility to the level that generated the indifference of x^*, $q(F, \mu; \upsilon^R)$ has risen faster than $q(F', \mu'; \upsilon^R)$. For the new GA-structure to capture the minimum price equilibrium we need $q(F, \mu; \upsilon^R) > q(F', \mu'; \upsilon^R)$ as we lower reservation utility further. Again this amounts to ruling out tangencies, and we will argue below that in a wide class of models such tangencies are extremely rare.

[**Figure 4**]

There are several issues that may arise in implementing this algorithm. First, there may come a point 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; multiple indifference would lead to confusion over which action to take. Second, the algorithm may take an infinite number of steps to converge. Assumption G below rules out all such cases. In the next section, we show that this assumption holds generically in the cases that we consider.

We call this algorithm the GAME-algorithm, given that it operates by identifying combinations of GA-structures that correspond to minimum equilibrium prices. Of course, we are not studying this algorithm for its own sake. The idea is that we are following a path that allows us to uncover the actual minimum equilibrium price.

To establish that we are successful in this regard, we first introduce the correspondence that we would ideally like to uncover, and impose seemingly natural genericity conditions. This GAME-correspondence associates with each relevant utility vector the set of GA-structures that generate the minimal equilibrium price vector. We then detail the idealized path of the algorithm. The description is theoretical since we assume that we may reduce reservation utilities continuously. The critical result that pulls this all together is presented in the last sub-section, which establishes that, under the genericity conditions we impose, the GAME-algorithm is always a selection from the GAME-correspondence. Under our assumptions this implies that it identifies a minimum price equilibrium of the original model once we have reduced reservation utility to the initially given value of $\bar{u}^R$. In the next section, we discuss practical issues of implementation, including the issue of translating a continuous path to a computer and a number of special cases in which the algorithm simplifies greatly.

5.2. The GAME-Correspondence

Recall that $\bar{u}^R \in R^m$ denotes the initial vector of reservation utilities. We set maximum reservation utility $\bar{U}^R$ at some level such that each buyer x_a strictly prefers the outside option to any good priced at its reservation price,

$$\bar{U}_a^R > \max\left[\bar{u}_a^R, \max_{y_i \in Y} U_a(y_i, \max\{w_a - r_i, 0\})\right].$$

Define $\Delta_a = \left(\bar{U}_a^R - \bar{u}_a^R\right)$ as the gap between the x_a's initial level of reservation utility $\bar{U}_a^R$ and the true level $\bar{u}_a^R$ and $\bar{z} = \sum_{a=1}^m \Delta_a$ as the sum of the gaps. We associate with each $z \in [0, \bar{z}]$ a vector of reservation utilities $\upsilon^R \in [\bar{u}^R, \bar{U}^R]$. Specifically, if $z \in \left[\sum_{i=s+1}^m \Delta_a, \sum_{a=s}^m \Delta_a\right]$, then:

$$\upsilon_a^R(z) = \begin{cases} \bar{U}_a^R \text{ for } a \geq s+1; \\ \bar{u}_a^R + \left(z - \sum_{a=s+1}^m \Delta_a\right) \text{ for } a = s; \\ \bar{u}_a^R \text{ for } a \leq s-1. \end{cases} \tag{5.1}$$

As we reduce z from $\bar{z}$ to zero, we lower the reservation utility of buyers, one by one, from $\bar{U}_a^R$ to $\bar{u}_a^R$. In this way we trace out models with lower and lower reservation utility. The GAME-correspondence associates with each $z \in [0, \bar{z}]$ the set of all GA-structures that correspond to the minimum price equilibrium for a model with reservation utility $\upsilon^R(z)$.

Definition The GAME-correspondence, $\Phi : [0, \bar{z}] \longrightarrow 2^{\mathcal{G}}$ gives for each $z \in [0, \bar{z}]$ the set of all GA-structures (F, μ) such that $(q(z, F, \mu), \mu)$ is a minimum price equilibrium given reservation utility $\upsilon^R(z)$.

Given Assumption A and Theorem 3, we know that the GAME-correspondence is nonempty. The upper hemicontinuity of demand will guarantee that it is upper hemicontinuous. Our algorithm, however, relies on the following stronger conditions.

Assumption G: $\Phi(z)$ has the following properties:

1. $\Phi(z)$ has at most 2 elements for all $z \in [0, \bar{z}]$.
2. $\{z \in [0, \bar{z}] || \Phi(z)| > 1\}$ is finite.
3. $\{z \in [0, \bar{z}] || \Phi(z)| > 1\}$ does not intersect $\{z \in [0, \bar{z}] | z = \sum_{a=s}^m \Delta_a \text{ some } 1 \leq s \leq m\}$.
4. The graph of $\Phi(z)$ contains no isolated points.

Assumption G addresses some of the potential problems discussed in the previous section. Assumption G1 rules out multiple potential moves. There can be only two potential GA-structures at any given time. Assumption G2 guarantees that there are only a finite number of points that we may want to make changes to the GA-structures. Assumption G3 ensures that there is no confusion at the end of stages regarding which GA-structure to use going forward. Assumption G4 rules out tangencies, that is situations in which an agent demands a good at one price but not at prices immediately above and below. In Section 6, we prove that Assumption G holds generically in several special cases. The structure of the proofs suggest that the result generalizes.

Assumption G implies that Φ has a simple structure. As z falls, Φ is equal to some GA-structure until at some point a second GA-structure appears. Given that there are no isolated points, Φ then switches to the new GA-structure beyond that point. Our algorithm is designed to uncover just this path.

5.3. The GAME-Algorithm

In this section we construct a mapping $\Theta : [0, \bar{z}] \longrightarrow \mathcal{G}$ which maps z into a pair $(F(z), \mu(z))$. We invoke Assumption A. We do not invoke Assumption G at this point. We use Assumption G in the next section to prove that $\Theta(z) \in \Phi(z)$ for all $z \in [0, \bar{z}]$.

We divide the construction of Θ into a series of stages and steps. Stages are indexed by $s \in \{1, \ldots, m\}$, one for each buyer. Stage s corresponds to $z \in \left(\sum_{a=s+1}^{m} \Delta_a, \sum_{a=s}^{m} \Delta_a\right]$ and involves lowering the reservation utility of buyer x_s, v_s^R, from $\bar{U}_s^R$ to $\bar{u}_s^R$. Steps correspond to adjustments that are made to (F, μ) during the stage. The algorithm involves a number T_s of steps in stage s: the steps are ordered with higher steps corresponding to lower levels of z. While we allow in principle for an infinite number of steps, we show in the next section that the algorithm is finite when Assumption G is invoked. The defining feature of a step within a stage is that the mapping Θ is constant, enabling us to index the corresponding GA-structure by stage and step as $(F(s,t), \mu(s,t))$.

We construct Θ inductively. We initalize with,

$$\Theta(\bar{z}) = (F(1,1), \mu(1,1)),$$

where $\mu(1,1)$ matches all buyers with their reservation goods and $F(1,1)$ is the null forest with no edges. Note that $\Theta(\bar{z}) \in \mathcal{G}$. Define $z(1,1) = \bar{z}$.

The induction step begins with a point $z(s,t) \in \left(\sum_{a=s+1}^{m} \Delta_a, \sum_{a=s}^{m} \Delta_a\right]$ and $\Theta(z) \in \mathcal{G}$

on $z \geq z(s,t)$. Where possible we now identify $z(s,t+1) \in [0, z(s,t))$ and extend $\Theta(z)$ to $[z(s,t+1), z(s,t))$ such that $\Theta(z) \in \mathcal{G}$.

Define $\hat{Y}(s,t) \subset \bar{Y}$ as the set y_{n+s} and its successors in $F(s,t)$ and $\hat{X}(s,t) \subset X$ as the set of buyers allocated to goods in $\hat{Y}(s,t)$,

$$\hat{X}(s,t) \equiv \{x_a \in X | \mu_a(s,t) \in \hat{Y}(s,t)\}.$$

For $x_a \in \hat{X}(s,t)$, let $J_a(s,t) = \{y_i \in \bar{Y}_a | y_i \neq \mu_a(s,t), (\mu_a(s,t), y_i) \notin E(F(s,t))\}$ denote the set of goods in $\bar{Y}_a$ excluding μ_a and its direct successors. Denote the value of $\mu_a(s,t)$ to x_a by $U_a(z,s,t) \equiv U_a(y_k, w_a - q_k(z,s,t)))$ where $y_k = \mu_a(s,t)$ and q_k is the price of y_k according to $q(z,s,t) \equiv q(z, F(s,t), \mu(s,t))$.

Let $\mathcal{Z}(s,t)$ denote the set of all $z \in (\sum_{a=s+1}^{m} \Delta_a, z(s,t))$ such that a buyer in $\hat{X}(s,t)$ is indifferent between $\mu_a(s,t)$ and a good in $J_a(s,t)$:

$$\mathcal{Z}(s,t) \equiv \left\{ z \in \left(\sum_{a=s+1}^{m} \Delta_a, z(s,t) \right) \middle| U_a(z,s,t) = U_a(y_i, w_a - q_i(z,s,t)) \right. \\ \left. \text{some } x_a \in \hat{X}(s,t), y_i \in J_a(s,t), q(z,s,t) \in \Pi \right\}.$$

If $\mathcal{Z}(s,t) = \varnothing$, we set $z(s,t+1) = \sum_{a=s+1}^{m} \Delta_a$. In this case, if $z(s,t+1) = 0$, we set

$$\Theta(z) = (F(s,t), \mu(s,t))$$

for all $z \in [0, z(s,t))$, and the algorithm ends. Otherwise, we set

$$\Theta(z) = (F(s,t), \mu(s,t))$$

for all $z \in [z(s,t+1), z(s,t))$, the stage ends, and $z(s+1,1) = z(s,t+1)$.[12]

If $\mathcal{Z}(s,t) \neq \varnothing$, we set $z(s,t+1) = \sup \mathcal{Z}(s,t)$. If $z(s,t+1) = z(s,t)$, we set

$$\Theta(z) = (F(s,t), \mu(s,t))$$

for all $z \in [0, z(s,t))$, and the algorithm ends.[13]

[12] When Assumption G3 is later invoked, there will be no ambiguity concerning how to choose (F, μ) at $z(s,t+1)$.

[13] When Assumption G2 is later invoked, this case will not arise.

If $z(s, t+1) < z(s, t)$, then $z(s, t+1) \in \mathcal{Z}(s, t)$ follows from Assumption A1 which implies that buyers' demands are upper hemi-contintinuous. In this case, we set $\Theta(z) = \{F(s, t), \mu(s, t)\}$ for all $z \in (z(s, t+1), z(s, t))$. We alter $\Theta(z)$ at $z = z(s, t+1)$. We select x^* from the set of $x_b \in \hat{X}(s, t)$ such that $U_b(\mu_b(s, t), q(z, s, t)) = U_b(y_i, q(z, s, t))$ for some y_i that is not a direct successor of $\mu_b(s, t)$.[14] Let $\mu^* = \mu(x^*, s, t)$, and y^* the good that x^* values equally with μ^* and its direct successors. The changes that we make in $(F(s, t), \mu(s, t))$ correspond to the three cases described intuitively in Section 5.1.

Case 1: y^* is unallocated under $\mu(s, t)$ as in Figure 2(a) above. Let $Y^P = (y_1^P, y_2^P, \ldots y_k^P)$ denote the path from $y_{n+s} = y_1^P$ to $\mu^* = y_k^P$ in $F(s, t)$ (this path may be trivial if $\mu^* = y_{n+s}$). 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:

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

We alter the edges $E(F(s, t))$ accordingly. We add the edge (y^*, μ^*). 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 $F(s, t+1)$.

Case 2: y^* is allocated under $\mu(s, t)$, but there exists no path directed from y^* to μ^*. This situation was depicted in Figure 3. In this case we make no changes to μ: $\mu(s, t+1) = \mu(s, t)$. We alter F by replacing the edge (y_j, y^*) oriented towards y^*, with the edge (μ^*, y^*). We label the new graph $F(s, t+1)$.

Case 3: y^* is allocated under $\mu(s, t)$ and there exists a path directed from y^* to μ^*, as in Figure 4. We add the edge (μ^*, y^*) to $F(s, t)$. This creates a directed circuit $Y^C = (y_1^C, \ldots y_k^C, y_{k+1}^C)$ where $y_1^C = y_{k+1}^C = y^*$ and $y_k^C = \mu^*$. We shift each buyer who is matched an element of Y^C in the direction of the circuit: if $\mu_a(s, t) = y_z^C \in Y^C$, then $\mu_a(s, t+1) = y_{z+1}^C$. Otherwise $\mu_a(s, t+1) = \mu_a(s, t)$. We delete the edge (y_1^C, y_2^C). We reverse the orientation of all other $(y_z^C, y_{z+1}^C) \in Y^C$. 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 $\left(y_{z+1}^C, y_j\right)$. We label the new graph $F(s, t+1)$.

To complete the induction step we establish in the appendix that $(F(s, t+1), \mu(s, t+1)) \in \mathcal{G}$.

Lemma 1 With Assumption A, $(F(s, t+1), \mu(s, t+1)) \in \mathcal{G}$.

This completes the description of the algorithm.

[14] When Assumption G1 is later invoked, this set will have a single element.

5.4. Equivalence

The critical result is that assumptions A and G ensure the GAME–algorithm is contained within the GAME-correspondence, so that its end-point $\Theta(0)$ identifies a minimum price equilibrium of the original model.

Theorem 4 Given Assumptions A and G,

$$\Theta(z) \in \Phi(z),$$

and the GAME-algorithm converges in a finite number of steps to $\Theta(0) \in \Phi(0)$ which is a minimum price equilibrium of the model.

6. Implementation

We first present two classes of utility function that are especially easy to work with and that can be shown to satisfy Assumption G almost surely over a standard parameter space. We then discuss the issue of implementing the algorithm taking account of the fact that computer search is discrete, hence unable to directly uncover the entire GAME-correspondence, depending as it does on a continuous parameter.

6.1. The No Switches Property

The algorithm converges very quickly for utility functions that satisfy the following "no switches" property.

No Switches (NS) Property: A set of utility functions $\{U_a\}_{a=1}^m$ such that $U_a : \bar{Y}_a \times \mathbb{R}_+ \to \bar{\mathbb{R}}$ satisfies property NS if for any $w \in \mathbb{R}^m_{++}$, any pair of individuals x_a and x_b, any pair of goods y_i and y_k, and any price vectors p and p' such that x_a is indifferent between y_i and y_k both at p and at p', x_b either strictly prefers y_i to y_k at both prices, is indifferent between y_i and y_k at both prices, or strictly prefers y_k to y_i at both prices.

Property NS states that if we raise the prices of two goods to keep one buyer indifferent, then no other buyer alters his demand for these goods. For any class of utility functions satisfying condition NS, the path of prices implied by $\Theta(z)$ never causes any group of buyers to want to reallocate goods amongst themselves. It implies that each step t of stage s ends either with $v_s^R = u_s^R$ or when a buyer assigned to a good in the component containing y_{n+s} demands some good that is in another component. The number of steps is bounded above since the component containing y_{n+s} expands with each step. The number of steps is kept small because there is no reallocation except when a buyer becomes indifferent to an unallocated good. At this point the stage effectively ends as y_{n+s} is no longer allocated.

6.1.1. Transferable Utility

The case with linear or transferable utility, in which the distribution of wealth is irrelevant to demand, satisfies NS. In fact, our algorithm in this case can be seen as an extension of Demange, Gale, and Sotomayor (1986) to continuous offers.

Formally, the utility function for buyer $x_a \in X$ takes the following form,

$$U_a(y_i, c) = h_{ia} + c,$$

where h_{ia} represents the value of good y_i to buyer x_a. For x_a to be indifferent between y_i and y_k, the price differential must equal the difference in h:

$$h_{ia} - h_{ka} = p_i - p_k.$$

That property NS is satisfied follows immediately from the fact that the h's are fixed: the price differential that makes x_a indifferent is fixed, and at a fixed price differential the demands of other buyers are fixed.

It is easy to show that this specification satisfies Assumption G generically. First, Assumption G2 is satisfied since during stage s each step involves an expansion of the component containing y_{n+s}, the number of steps is finite. Second, if h is drawn from a continuous probability distribution on R^{mn+m} (there are m agents and $n+1$ goods for each agent to consider), this utility function satisfies Assumption G3 almost surely as it requires a particular alignment of h's for there to exist an indifference at the end of a stage not encoded in F. Third, given any z and any GA structure (F, μ), it requires a particular alignment of the h's for there to exist more than one indifference not encoded in F; according to Lemma 3 one additional indifference is associated with at most two elements of Φ (Assumption G1). Finally, since each step within a stage ends with some buyer allocated to a good within the component containing y_{n+s} indifferent to a good outside that component. This implies that there are no isolated elements of Φ (Assumption G4) as the old GA structure can no longer support the minimal price equilibrium as z falls further.[15]

6.1.2. Exponential utility

We show that property NS is satisfied by utility functions that are additively separable between the indivisible goods and the numeraire good, exponential in the numeraire good, and have the same coefficient of absolute risk aversion. Formally, the assumption is that,

$$U_a(y_i, c) = h_{ia} - e^{-\alpha c}, \tag{6.1}$$

where $\alpha > 0$ is the coefficient of absolute risk aversion.

Given the wealth vector w, we can transform utility by multiplying U_a by $e^{\alpha w_a}$:

$$\hat{U}_a \equiv U_a e^{\alpha w_a} = h_{ia} e^{\alpha w_a} - e^{\alpha p_i}.$$

Defining $\hat{p}_i = e^{\alpha p_i}$ the resulting utility function is,

$$\hat{U}_a = h_{ia} e^{\alpha w_a} - \hat{p}_i.$$

Since utility is additively separable in prices, this function inherits properties of the linear case. In particular it satisfies NS and Assumption G almost surely when h is drawn from a continuous probability distribution on R^{mn+m}.

With this form of additive-exponential utility, changes in wealth alter the relative desirabilities of goods. Given wealth, however, changes in prices work very much as they do in the linear case.

[15] If the outside good is unallocated, buyer x_s shifts off of y_{n+s}. As reservation utility falls further y_{n+s} becomes even more undesirable. The new GA structure therefore holds over an interval of z below the point at which the switch is made. If the outside good is allocated, it and its sucessors are added to the component containing y_{n+s}; its price rises, and it is no longer desired by the buyer allocated to its former direct predecessor.

6.2. Log Utility

A useful class of utility functions are those that are additively separable between the indivisible goods and the numeraire good and logarithmic in the numeraire good:[16]

$$U_a(y_i, c) = h_{ia} + \ln c, \tag{6.2}$$

This class of utility functions has the property that given any subset of goods $\tilde{Y} = \{y_1, y_2 \ldots y_K\} \subseteq Y$ and any mapping $\alpha : \tilde{Y} \to X$ of goods to agents (not necessarily one-to-one) either there exists at most one price at which there is a cycle of indifference in which $\alpha(y_i)$ is indifferent between goods y_i and y_{i+1} (where $K + 1 = 1$) or there is an indifference cycle at all prices. This property implies that the algorithm will not cycle through the same graph structures as z falls.

To see this let $\tilde{p}_1$ denote the price of y_1 and define $\tilde{p}_{k+1}$ recursively according to the indifference of $\alpha(k)$:

$$h_{k,\alpha(k)} + \ln\left(w_{\alpha(k)} - \tilde{p}_k\right) = h_{k+1,\alpha(k)} + \ln\left(w_{\alpha(k)} - \tilde{p}_{k+1}\right). \tag{6.3}$$

Suppose that all the $\tilde{p}_k$ are positive and that $\tilde{p}_k < \min\{w_{\alpha(k)}, w_{\alpha(k-1)}\}$ so that affordability conditions are met.

The key observation is that $\partial\tilde{p}_{k+1}/\partial\tilde{p}_k$ is constant. Rearranging (6.3),

$$\tilde{p}_{k+1} = (1 - \chi)w_{\alpha(k)} + \chi\tilde{p}_k, \tag{6.4}$$

where $\chi = \exp[h_{k,\alpha(k)} - h_{k+1,\alpha(k)}]$. Hence $\partial\tilde{p}_{k+1}/\partial\tilde{p}_k = \chi$. The chain rule implies that $\partial\tilde{p}_{k+1}/\partial\hat{p}_1$ is constant. That log utility has the desired property follows immediately: either there exits a price at which indifference holds or there is not, and if so either $\partial\tilde{p}_{K+1}/\partial\hat{p}_1 = 1$ or it does not.

It takes a special choice of parameters for $\partial\tilde{p}_{K+1}/\partial\hat{p}_1 = 1$. This will not generically be the case. We show below that this log utility function also satisfies Assumption G with probability 1 if the h_{ia} and w_a are drawn from a continuous probability distribution.

Theorem 5: Consider the utility function.

$$\ln h_{ai} + \ln(w_a - p_i).$$

Fix the value of the null goods. Select $(w, h) \in \mathcal{A} \subset R^{m+nm}$ where $\mathcal{A}$ is a bounded set

[16] This utility function is likely to be of practical use in housing markets. Davis and Ortalo-Magne (2008) have shown that the share of housing expenditure is constant over time which is consistent with a Cobb-Douglas utility function.

with positive Borel measure. Then the GAME-correspondence $\Phi(z)$ satisfies Assumption G almost surely.

6.3. From Continuous to Discrete

Implementation of the algorithm requires discretizing z. Suppose that we have arrived in the course of the algorithm with $z(s,t-1)$ and associated GA-structure $(F(s,t-1),\mu(s,t-1))$. We proceed as follows. Recall $\hat{Y}(s,t-1)$ ($\hat{Y}$ in what follows) is the component of $F(s,t-1)$ containing y_{n+s}. For each x_a such that $\mu_a(s,t-1)\in\hat{Y}$, we calculate the best feasible option outside of $\hat{Y}$. Let $YF_a=\{y_i|y_i\in\bar{Y}_a\backslash\hat{Y}$ and $w_a>q_i\}$ denote the set of feasible, outside options where q_i is the price of y_i under $q(z,s,t-1)$, and let ψ_a denote the best choice from this set:

$$\psi_a=\{y_i\in YF_a|U_a(y_i,w_a-q_i)\geq U_a(y_j,w_a-q_j)\text{ all }y_j\in YF_a\}.$$

Note that since $y_i\notin\hat{Y}$, q_i is not affected by z during stage $(s,t-1)$.

Now for each ψ_a we can calculate the level of z at which x_a becomes indifferent between $\mu_a(s,t-1)$ and ψ_a. This involves calculating the price of $\mu_a(s,t-1)$ that makes x_a indifferent and then working backward to y_{n+s}. Label the resulting level z_a.[17] We then consider $z_b=\max\{z_a\}$. We then check to see if $(F(s,t-1),\mu(s,t-1),z_b)$ satisfies Theorem 3. If it does then $z(s,t)=z_b$ and the step ends with x_b demanding ψ_b.

If $(F(s,t-1),\mu(s,t-1),z_b)$ fails Theorem 3, then we know that the step should have ended sooner with some rearrangement of $F(s,t-1)$ or $\mu(s,t-1)$ on T. We look for $\hat{z}\in(z_b,z(s,t-1))$ in which there is only one change in demand. The natural approach is to bisect the interval $(z_b,z(s,t-1))$, calculate $q(z,s,t-1)$ at the resulting $\hat{z}$, and check to see how many x_a allocated to goods in T prefer goods other than $\mu_a(s,t-1)$. If there is only one such x_a, then we make the adjustments to $(F(s,t-1),\mu(s,t-1))$ associated with x_a demanding the good in question, and test to see if the resulting $(F(s,t),\mu(s,t),q(z,F(s,t),\mu(s,t)))$ satisfies Theorem 3. If there is more than one such x_a, we continue bisecting the interval $(\hat{z},z(s,t-1))$. If there is no such x_a, we continue bisecting the interval $(z_b,\hat{z})$. We proceed until we obtain $(F(s,t),\mu(s,t),q(z,F(s,t),\mu(s,t)))$ that satisfies Theorem 3.

[17] ψ_s may be empty if x_s cannot afford any of the outside options. In this case we set z_s to $\sum_{a=s+1}^{m}\Delta_a$. Assumption A2 ensures that ψ_a is nonempty for all other $x_a\in\hat{Y}$.

7. The Equilibrium Set

In section 7.1 we outline a variation on the GAME-algorithm that solves for the maximum price equilibrium. The algorithm solves a dual problem in which sellers' reservation prices are adjusted down from an artificially high level to their true level. In section 7.2 we show how to use GAME-algorithms to uncover all equilibria of the model.

7.1. Maximum Price Equilibrium

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: we let x_{m+i} denote the null buyer for seller $y_i \in Y$, defining $\bar{X}_i = X \cup \{x_{m+i}\}$ as the feasible set of choices for y_i, and $\bar{X} = \cup \bar{X}_i$, the overall extended set of buyers, by direct analogy with the original (primal) case.

Buyers' utility υ in the dual plays the role of prices in the primal. Buyers are willing to purchase a good if the specified utility is at least at reservation level, and are unwilling to not purchase if the inequality is strict. Define $\bar{\upsilon} \in \mathbb{R}_+^m$ by $\bar{\upsilon}_a = \max_{y_i \in \bar{Y}_a} U_a(y_i, w_a)$. $\bar{\upsilon}_a$ is the maximal utility that can be promised to buyer x_a from the consumption of all goods.[18] Let $\tilde{\Pi} = \left\{ \upsilon \in \mathbb{R}_+^m | \bar{u}^R \leq \upsilon \leq \bar{\upsilon} \right\}$.

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" $V_i(x_a, \upsilon_a)$.

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

[18] This upper bound will never be exceeded in equilibrium since it is (weakly) above the value buyers receive in the minimal price equilibrium.

assigned a feasible buyer,

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

Definition 7.1. *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.

Assumption A is 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.

7.2. 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 then use the original algorithm with reservation prices that lie between the identified minimum and maximum equilibrium prices to uncover the complete set of 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 $\hat{r} \in [\underline{p}, \bar{p}]$.

8. Concluding Remarks

We have presented an algorithm for constructing minimum price competitive equilibria in allocation markets with non-transferable utility. The solution suggests 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 requires 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?

9. References

References

[1] Aigner, Martin, and Günter Ziegler (2003), *Proofs from THE BOOK*, New York: Springer-Verlag, third 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] Davis, Morris, and Francois Ortalo-Magne, (2008) "Household Expenditures, Wages, Rents," University of Wisconsin working paper.

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

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

[7] Gale, David, (1984), "Equilibrium in a Discrete Economy with Money." *International Journal of Game Theory* 13, 61–64.

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

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

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

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

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

[15] Scarf, Herbert, (1973), *The Computation of Economic Equilibria*, New Haven, Yale University Press.

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

10. Appendix

Proofs of Lemmas and Theorems appear in the order in which they appear in the text. Additional lemmas used only in the appendix follow.

Proof of Theorem 1: (only if) Consider any $p > r$ for which there exists $\mu \in M$ such that (p, μ) is an equilibrium. By Hall's theorem, we know that existence of such a match ensures that (1) is satisfied. To establish (2), consider $B \subset H(p)$, and note that, from property 2 of the equilibrium,

$$y_i \in H(p) \Longrightarrow \mu_a(x) = y_i \text{ some } x_a \in x.$$

Hence,

$$|\Gamma_p(B)| \geq |\{x_a \in x | \mu_a(x) \in B\}| = |B|.$$

(if) Consider price $p > r$ such that conditions (1) and (2) hold. The proof that we can find a match $\mu \in M$ such that (p, μ) is an equilibrium is inductive based on the cardinality $H(p)$. With $|H(p)| = 1$, condition (1) allows Hall's theorem to be used to create a match $\mu \in M$ from X into $\bar{Y}$. Once this initial match is set up, we see whether or not $\exists x_a \in X$ such that $H(p) = \mu_a(x)$. If so, we are done. If not, we are done once we have found any buyer who is linked in the graph to $H(p)$ and switched this individual to the unique good in $H(p)$: existence of such a link is guaranteed by condition (2) according to which at least one individual demands this good at the given prices.

For the inductive step, we begin by assuming that the theorem holds for all sets $H(p)$ of cardinality $K \geq 1$, and prove that it extends to all cases with $|H(p)| = K + 1$.

The nature of the inductive step depends on which of the following exhaustive conditions hold:

CASE 1: $|\Gamma_p(\tilde{B})| = |\tilde{B}|$ some $\tilde{B} \subset H(p)$;

CASE 2: $|\Gamma_p(\tilde{A})| = |\tilde{A}|$ some $\tilde{A} \subset X$ with $\Gamma_p(\tilde{A}) \cap H(p) \neq \phi$.

CASE 3: $|\Gamma_p(A)| > |A|$ all $A \subset X$ with $\Gamma_p(A) \cap H(p) \neq \phi$ and $|\Gamma_p(B)| > |B|$ all $B \subset H(p)$.

CASE 1: Consider a set $\tilde{B} \subset H(p)$ such that $|\Gamma_p(\tilde{B})| = |\tilde{B}|$. The key observation is that subgraph G'_p of G_p defined by removing $\tilde{B}$ from $\bar{Y}$ and $\Gamma^p(\tilde{B})$ from X satisfies conditions (1) and (2). Define,

$$\begin{aligned} X' &= X \backslash \Gamma_p(\tilde{B}); \\ H'(p) &= H(p) \backslash \tilde{B}; \\ \bar{Y}' &= \bar{Y} \backslash \tilde{B}. \end{aligned}$$

Note that $|H'(p)| \leq K$, so that it will be our goal to establish that we can create an appropriate match in this subgraph based on the induction hypothesis. To this end, what we must establish is,

$$\begin{aligned} |\Gamma'_p(A')| &\geq |A'| \text{ all } A' \subset X'; \\ |\Gamma'_p(B')| &\geq |B'| \text{ all } B' \subset H'(p); \end{aligned}$$

where Γ'_p is the neighborhood function in the subgraph G'_p of G_p.

By construction of A', no member has any neighbors removed by removal of $\Gamma_p(\tilde{B})$, hence by condition (1),

$$|\Gamma'_p(A')| = |\Gamma_p(A')| \ \geq |A'| \text{ all } A' \subset X'.$$

As for $B' \subset H'(p)$, note that,

$$\Gamma'_p(B') \supset \Gamma_p(\tilde{B} \cup B') \backslash \Gamma_p(\tilde{B}).$$

Hence,

$$\begin{aligned} |\Gamma'_p(B')| &\geq |\Gamma_p(\tilde{B} \cup B')| - |\Gamma_p(\tilde{B})| \\ &= |\Gamma_p(\tilde{B} \cup B')| - |\tilde{B}|. \end{aligned}$$

Since $\tilde{B} \cup B' \subset H'(p)$ and we know from (2) that $|\Gamma_p(\tilde{B} \cup B')| \geq |\tilde{B} \cup B'|$. Finally, since $\tilde{B} \cap B' = \phi$ and application of (2) reveals that $|\Gamma_p(\tilde{B} \cup B')| \geq |\tilde{B} \cup B'|$, we derive the desired inequality,

$$|\Gamma'_p(B')| \geq |\tilde{B} \cup B'| - |\tilde{B}| = |B'|.$$

In light of the induction hypothesis, we know that we can create a match function μ' from X' into $\bar{Y}'$ which has the required properties that $\mu' \in M'(p)$ (the match function in the submarket) and such that $y_i \in H'(p) \Longrightarrow \exists x_a \in X$ such that $y_i = \mu'(x_a)$. With respect to the set $\tilde{B}$, note that by condition (2) we can apply Hall's theorem to produce a match between elements of $\tilde{B}$ and those of $\Gamma_p(\tilde{B})$, and that since these two sets have the same number of elements, the map is invertible to a map $\tilde{\mu}$ that takes $\Gamma_p(\tilde{B})$ onto $\tilde{B}$, so that $y_i \in \tilde{B} \Longrightarrow \exists x_a \in \Gamma_p(\tilde{B})$ such that $y_i = \tilde{\mu}_a(x)$. To complete the construction of the required matching function $\mu \in M$, we simply patch together the separate matching functions μ' from X' into $\bar{Y}'$ and $\tilde{\mu}$ from $\Gamma_p(\tilde{B})$ into $\tilde{B}$, noting that $\mu_a(x) \in D_a(p)$ by construction, as is the property that $y_i \in H(p) \Longrightarrow \exists x_a \in X$ such that $y_i = \mu_a(x)$. To complete the match function in the desired manner, we simply match x_k

with y_k, and are done.□

Proof of Theorem 2: (only if) This is a direct implication of Lemma 4 in Demange and Gale (1985).

(if) Fix price $p \in \Pi$ such that conditions (A1) and (S2) hold. Now consider any price vector $q \in \Pi$ such that $q_j < p_j$ for some j, and let $L(q,p)$ be the set of goods with prices strictly lower in q than in p, noting that all such must be contained in $H(p)$ since other prices are at their global minimum levels. Let $K \geq 1$ denote the cardinality of $L(q,p)$,

$$K = |L(q,p)| = |\{y_j \in H(p) | q_j < p_j\}|.$$

Define $A(q) \subset X$ as the set of buyers all of whose most preferred points at prices $q \in \Pi$ lie within set $L(q,p)$,

$$A(q) = \{x_i \in X | D_i(q) \subset L(q,p)\}.$$

To prove that $q \in \Pi$ is not an equilibrium price vector, we show that $A(q)$ has cardinality of at least $(K+1)$, contradicting necessary condition (A1), since in this case $Z_q(A(q)) = 0$ and $|\Gamma_q(A(q))| = K < |A(q)|$. To establish that $|A(q)| \geq K+1$, note first that with (S2) and with $L(q,p) \subset H(p)$,

$$|\Gamma_p(L(q,p))| = |\{x_i \in X | D_i(p) \cap L(q,p) \neq \phi\}| \geq K+1.$$

To complete the proof, note that, in light of strict monotonicity of the utility function in consumption,

$$\Gamma_p(L(q,p)) \subset A(q).$$

To see this, note that since $D_i(p) \cap L(q,p) \neq \phi$, there exists $y_{\tilde{j}} \in L(q,p)$ such that,

$$U_i(y_{\tilde{j}}, w_i - p_{\tilde{j}}) \geq U_i(y_k, w_i - p_k) \text{ all } y_k \in B_i(p).$$

Note that if we now compare the utility of this element $y_{\tilde{j}}$, which remains affordable at the new price vector q, to that of any alternative element that was in $B_i(p)$ but is not in $L(q,p)$, then due to strict monotonicity in the utility function,

$$U_i(y_{\tilde{j}}, w_i - p_{\tilde{j}}) > U_i(y_k, w_i - p_k) \text{ all } y_k \in B_i(p) \backslash L(q,p).$$

The final check we need to run is to confirm that there is no member of $B_i(q) \backslash L(q,p)$ that was not affordable at previous prices $B_i(p)\backslash L(q,p)$, or that,

$$B_i(q)\backslash L(q,p) \subset B_i(p)\backslash L(q,p).$$

But this is immediate, since goods $y_k \in B_i(q)\backslash L(q,p)$ are defined by the twin properties,

$$w_i > q_k \text{ and } q_k \geq p_k,$$

so that, a fortiori, $w_i > p_k$, implying $y_k \in B_i(p)\backslash L(q,p)$, and completing the proof that $\Gamma_p(L(q,p)) \subset A(q)$, and with it the result.□

Proof of Theorem 3: (only if) Consider a pair $(\hat{p}, \hat{\mu})$ that form a competitive equilibrium, and in which the equilibrium price $\hat{p}$ is minimal among all prices in any equilibrium. We set $\mu = \hat{\mu}$ and use the demand graph to construct $\hat{F} \in \mathcal{F}$ such that $(\hat{F}, \hat{\mu}) \in \mathcal{G}$ and for which $q(\hat{F}, \hat{\mu}) = \hat{p}$: the fact that $\hat{\mu}_a \in D_a(q(\hat{F}, \hat{\mu}))$ will follow immediately given that $(\hat{p}, \hat{\mu})$ is an equilibrium.

The first stage in the construction of graph $\hat{F}$ on $\bar{Y}$ is to identify the root set as all goods that are at reservation prices,

$$R(\hat{F}) = \left\{ y_k \in \bar{Y} | \hat{p}_k = r_k \right\}.$$

Note that $Y_\varnothing \subset R(\hat{F})$ as required for $\hat{F} \in \mathcal{F}$. The graph is completed by induction. Let $A_1 = R(\hat{F})$ and let F_1 denote the null graph on the vertex set A_1. At stage $s \geq 1$ of the construction, suppose we have identified $A_s \subset \bar{Y}$ and a graph F_s on the vertex set A_s such that F_s is a forest of rooted trees with root set $R(\hat{F})$ and with all of the edges of F_s are directed away from the roots. By construction, $R(\hat{F}) \subset A_s$ and $\bar{Y}\backslash A_s \subset H(\hat{p})$. Given the latter, theorem 2 implies that there exists x_a such that $\hat{\mu}_a \in A_s$ and $D_a(\hat{p}) \cap \bar{Y}\backslash A_s \neq \varnothing$. Choose $y_i \in D_a(\hat{p}) \cap \bar{Y}\backslash A$, define $A_{s+1} = \{y_i\} \cup A_s$, add the edge $(\hat{\mu}_a, y_i)$ 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 $R(\hat{F})$ with all edges directed away from the roots. Given that there are a finite number of elements in $\bar{Y}\backslash R(\hat{F})$, this construction converges in a finite number of steps to $A_S = \bar{Y}$ to graph F_S, and we define $\hat{F} = F_S$ as the final forest of rooted trees on $\bar{Y}$.

To establish that $(\hat{F}, \hat{\mu}) \in \mathcal{G}$, note $(y_i, y_k) \in E(\hat{F})$ implies both that there exists $x_a \in X$ such that $\hat{\mu}_a = y_i$ and that $\hat{p}_k > r_k$. Since $\hat{p} \in \Pi^E$, this implies by definition that there exists $x_b \in X$ such that $\hat{\mu}_b = y_k$ as required. To see that $q(\hat{F}, \hat{\mu}) = \hat{p}$, note first that, by construction, all goods in $R(\hat{F})$ are at reservation prices. Furthermore note that for any edge $(y_i, y_k) \in E(\hat{F})$, the buyer $x_a \in X$ with $\hat{\mu}_a = y_i$ is indifferent between y_i and y_k at $\hat{p}$. In light of assumption A, the fact that all implied indifferences hold at $\hat{p}$ is sufficient to complete the demonstration that $q(\hat{F}, \hat{\mu}) = \hat{p}$ that which is precisely the condition used in generating the function $q(F, \hat{\mu})$.

(if) Suppose there exists $(F, \mu) \in \mathcal{G}$ with $\mu_a \in D_a(q(F, \mu))$. Construct the bipartite graph $\hat{G}$ with vertex set $V = X \cup \bar{Y}$, bipartition $X, \bar{Y}$, and edge set $E(\hat{G}) \subset X \times \bar{Y}$ with $(x_a, y_i) \in E(\hat{G})$ if either $y_i = \mu_a$ or $(\mu_a, y_i) \in E(F)$. It is immediate that $\hat{G}$ is a subgraph of the demand graph $G_{q(F,\mu)}$. We show that the conditions of Theorem 2 apply to $\hat{G}$, hence that they apply to $G_{q(F,\mu)}$, whereupon it follows from Theorem 2 that $(\mu, q(F, \mu))$ is a minimum price competitive equilibrium

Given $A \subset X$, it is immediate that $\Gamma_{\hat{G}}(A)$, the neighbors of A in $\hat{G}$, satisfies $|\Gamma_{\hat{G}}(A)| \geq |A|$, since μ is one-to-one. It remains only to show that, given $B \subset H(q(F, \mu))$, set $\Gamma_{\hat{G}}(B)$, the neighbors of B in $\hat{G}$, satisfies $|\Gamma_{\hat{G}}(B)| > |B|$. The weak version of this inequality, $|\Gamma_{\hat{G}}(B)| \geq |B|$, follows from the fact that if $y_i \in H(q(F, \mu))$, then $y_i \notin R(F)$ so that there exists $y_k \in \bar{Y}$ with $(y_k, y_i) \in E(F)$. By the definition of a GA-structure, all such y_i are matched under μ so that the corresponding edges (x_a, y_i) with $y_i = \mu_a$ are in $E(\hat{G})$ by construction. To establish that the inequality is strict, $|\Gamma_{\hat{G}}(B)| > |B|$, note that there exists at least one $y_i \in B$ that is not a successor in F to any of the other goods in B (any good in B with the shortest path to the root set satisfies this condition). Given that $y_i \notin R(F)$ and that $(F, \mu) \in \mathcal{G}$, there exist $y_k \in \bar{Y}$ such that $(y_k, y_i) \in E(F)$ with $\mu_a = y_k$ for some $x_a \in X$. By construction of $\hat{G}$, this implies that $(\mu_a, y_i) \in E(\hat{G})$, adding one to the cardinality of $\Gamma_{\hat{G}}(B)$ and completing the proof.□

Proof of Lemma 1: The proof is by induction. Since $F(1,1)$ is the null tree and $\mu(1,1)$ allocates buyers to their null goods. $(F(1,1), \mu(1,1)) \in \mathcal{G}$.

Given that $(F(s,t), \mu(s,t)) \in \mathcal{G}$, if the next step ends the stage $(F(s+1,1), \mu(s+1,1)) = (F(s,t), \mu(s,t)) \in \mathcal{G}$. We therefore assume that the next step ends with an added indifference before the end of the stage, and show that $(F(s,t+1), \mu(s,t+1)) \in \mathcal{G}$.

There are three cases to consider. For each case we show that $\mu(s,t+1) \in M$, that $F(s,t+1) \in \mathcal{F}$, and then $(F(s,t+1), \mu(s,t+1) \in \mathcal{G}$.

The first case is when y^* is unallocated under $\mu(s,t)$.

In this case, the only change in allocation is to shift x_b to y^* and to shift each other buyer assigned by $\mu(s,t)$ to a good on the path Y^P to its direct successor. Given $\mu(s,t)$ is an injection, it is clear that $\mu(s,t+1)$ is an injection. Moreover, given $y^* \in Y$, the set of buyers allocated to their null goods under $\mu(s,t+1)$ is a subset of the set of buyers allocated to their null goods under $\mu(s,t)$. It follows that $\mu(s,t+1) \in M$.

We now show that $F(s,t+1) \in \mathcal{F}$ given $F(s,t) \in \mathcal{F}$. We run through the conditions in the definition of $\mathcal{F}$ one by one. First, the only additional edge is (y^*, μ^*). Since y^* is unallocated, it is a null tree in $F(s,t)$. Hence the new edge does not create a cycle. Conditional on $F(s,t)$ being a forest, $F(s,t+1)$ is a forest. Second, both y_{n+s} and y^* are elements of $R(F(s,t))$. All successors of y_{n+s} in $F(s,t)$ become successors of y^* in $F(s,t+1)$. y_{n+s} is unallocated and hence a null tree in $F(s,t+1)$. Hence each component of $F(s,t+1)$ has a unique root good and the edges of $F(s,t+1)$ are directed away from that root good. Finally, the third condition holds, since the root set has not changed.

For $(F(s,t+1), \mu(s,t+1)) \in \mathcal{G}$ we need that any good that is the endpoint of an edge is allocated. This follows from $(F(s,t), \mu(s,t)) \in \mathcal{G}$, that y_{n+s} is no longer either allocated or an endpoint, and that y^* is now both allocated and an endpoint.

The second case is when y^* is allocated under $\mu(s,t)$, but that there exists no path directed from y^* to μ^*. In this case, no changes are made to μ: $\mu(s,t) \in M$ implies $\mu(s,t+1) \in M$.

We now show that $F(s,t+1) \in \mathcal{F}$ given $F(s,t) \in \mathcal{F}$. First, suppose that $F(s,t+1)$ contains a cycle. Since $F(s,t)$ is a forest, it must be the case that the addition of (μ^*, y^*) created the cycle and that there are two paths in $F(s,t+1)$ (not necessarily directed) between μ^* and y^*. Since there was only one edge directed toward y^* in $F(s,t)$ and this edge is eliminated from $F(s,t+1)$, y^* has no direct predecessors in $F(s,t+1)$ other than μ^*. Since y^* is not a predecessor of μ^*, there is no cycle involving y^* and μ^*. Hence $F(s,t+1)$ is a forest. Second, as y^* and its successors become successors of μ^*, all edges are directed away from the root set. Finally, the root set is unchanged. Hence $F(s,t+1) \in \mathcal{F}$.

Any good that is the endpoint of an edge in $F(s,t+1)$ is allocated. This follows from $(F(s,t), \mu(s,t)) \in \mathcal{G}$ and that the only new edge connects two allocated goods. Hence $(F(s,t+1), \mu(s,t+1)) \in \mathcal{G}$.

The third case is when y^* is allocated under $\mu(s,t)$ and y^* is a predecessor or μ^*.

Given $\mu(s,t) \in M$, that $\mu(s,t+1) \in M$ follows from the fact that the only change in allocation is to rotate buyers around the directed circuit created by the addition of the edge (μ^*, y^*).

We now show that $F(s,t+1) \in \mathcal{F}$ given $F(s,t) \in \mathcal{F}$. First, since $F(s,t)$ is a forest, it contains no cycles. The addition of the edge (y^*, μ^*) creates one cycle. This cycle is destroyed with the removal of the edge (y^*, y_2), where y_2 is the successor of y^* on the path in F between y^* and μ^*. Hence $F(s,t+1)$ is a forest. The adjustments ensure that all edges are directed away from the root set. Finally, no change was made to the root set. $F(s,t+1) \in \mathcal{F}$.

It follows from $(F(s,t), \mu(s,t)) \in \mathcal{G}$ and that the only new edge connects two allocated goods that $(F(s,t+1), \mu(s,t+1)) \in \mathcal{G}$.$\square$

Proof of Theorem 4: The proof is by induction. Given $\Theta(z) \in \Phi(z)$ over the interval $(z(s,t), \bar{z}]$, we show that $\Theta(z) \in \Phi(z)$ at $z(s,t)$ and in the case that $z(s,t) > 0$, we show that $\Theta(z) \in \Phi(z)$ over an interval $(z(s',t'), \bar{z}]$ where $z(s',t') < z(s,t)$.

The initial step is to show that $\Theta(z) \in \Phi(z)$ for $z \in (z(1,2), \bar{z}]$. $\Theta(\bar{z}) = \Phi(\bar{z})$ by construction. By assumption $F(1,1)$ is the null graph, so y_{n+1} has no successors and $\hat{Y}(1,1) = y_{n+1}$. $\hat{X}(1,1)$, the set of buyers allocated to $\hat{Y}(1,1)$, is simply x_1. We next characterize $z(1,2)$, the end of the step. Since x_1 is the only buyer in $\hat{X}$ and given the continuity of preferences, the next step occurs either at the end of the stage or when x_1 is indifferent between y_{n+1} and the good that he prefers most at the price vector r. Using equation (5.1),

$$z(1,2) = \sum_{a=2}^{m} \Delta_a + \max(0, \max_{y_i \in Y} U_1(y_i, w_1 - r_i) - \bar{u}_1^R\}.$$

$\bar{U}_1^R > \max_{y_i \in \bar{Y}_i} U_1(y_i, w_1 - r_i)$ by construction; it follows that $z(1,2) < \bar{z}$. We now show that $\Theta(z) \in \Phi(z)$ on $z \in (z(1,2), \bar{z})$. Since $F(1,1)$ is the null graph $q(z,1,1) = r$ for $z \in (z(1,2), \bar{z})$. As all agents strictly prefer their null goods at this price over this range of z, it follows from Theorem 3 that $(\mu(1,1), q(z,1,1))$ is a competitive equilibrium. We have shown that $\Theta(z) \in \Phi(z)$ for $z \in (z(1,2), \bar{z}]$.

We turn to the induction step. Suppose that $\Theta(z) \in \Phi(z)$ over the interval $(z(s,t), \bar{z}]$. We consider three cases.

The first case is $z(s,t) = z(m,t) = 0$. In this case, the algorithm ends with $\Theta(0) \equiv \{\mu(m,t), F(m,t), q(0,m,t)\}$. Assumption A implies that demand is upper hemicontinuous; it follows that $\mu_a(s,t) \in D_a(q(0,m,t)$. Theorem 3 then implies $\Theta(0) \in \Phi(0)$.

The second case is $z(s,t) = \sum_{a=s+1}^{m} \Delta_a$ and $s < m$. At this point the stage ends, $z(s,t) = z(s+1,1)$. We repeat the arguments of the initial induction step with $\hat{Y}(s+1,1) = y_{n+s+1}$, $\hat{X}(s+1,1) = x_{s+1}$, and $z(s+1,2) = \sum_{a=s+2}^{m} \Delta_a + \max(0, \max_{y_i \in Y} U_{s+1}(y_i, w_{s+1} - r_i) - \bar{u}_{s+1}^R\} < z(s,t+1)$. $\Theta(z) = \Phi(z)$ for $z \in (z(s+1,2), z(s+1,1)]$.

The third case is $z(s,t) > \sum_{a=s+1}^{m} \Delta_a$. In this case we make adjustments to (F,μ). We first show that $\Theta(z(s,t)) \in \Phi(z(s,t))$. Then we show that $z(s,t+1) < z(s,t)$ and $\Theta(z(s,t)) \in \Phi(z)$ for $z \in (z(s,t+1), z(s,t))$.

In each of the three cases that we make adjustments to (F,μ), $\mu_a(s,t) \in D_a(q(z(s,t), F(s,t-1), \mu(s,t-1)))$ as the alterations in μ shift buyers within their demand sets. Theorem 3 then implies $\Theta(z(s,t)) \in \Phi(z(s,t))$.

Assumption G2 implies that there exists $\varepsilon_1 > 0$ such that $\Theta(z(s,t-1))$ is the only element of $\Phi(z)$ over the interval $(z(s,t), z(s,t) + \varepsilon_1)$. Otherwise there would be a continuum of points at which $\Phi(z)$ is multivalued. The upper hemicontinuity of demand and Theorem 3 imply that $\Theta(z(s,t-1)) \in \Phi(z(s,t))$. Therefore both $\{\Theta(z(s,t-1)), \Theta(z(s,t)\} \in \Phi(z(s,t))$. Assumption G1 implies $\Phi(z(s,t))$ takes on no other values. Assumption G4 states $\Theta(z(s,t)$ is not isolated. This along with Assumption G2 implies that there exists $\varepsilon_2 > 0$ such that $\Theta(z(s,t)) = \Phi(z)$ for $z \in (z(s,t) - \varepsilon_2, z(s,t))$. $z(s,t+1) \leq z(s,t) - \varepsilon_2$. Otherwise $\mathcal{Z}(s,t+1) \cap (z(s,t) - \varepsilon_2, z(s,t)) = \varnothing$ and Lemma 3 implies that $\Phi(z(s,t))$ is not single valued.

Finally, we need to show that $\Theta(z(s,t)) = \Phi(z)$ for $z \in (z(s,t+1), z(s,t))$. Suppose not. Then there exists $\hat{z} \in (z(s,t+1), z(s,t))$ such that $\Theta(z(s,t)) \neq \Phi(z)$. Let $\hat{Z}$ denote the set of all such $\hat{z}$ and let $\tilde{z}$ denote the supremum of $\hat{Z}$. Suppose first that $\tilde{z} \in \hat{Z}$. The upper hemicontinuity of demand, Theorem 3 and the fact that $\Theta(z(s,t)) = \Phi(z)$ for all $z \in (\tilde{z}, z(s,t))$ implies that $\Theta(z(s,t)) \in \Phi(\tilde{z})$, but $\tilde{z} \in \hat{Z}$ implies $\Theta(z(s,t)) \notin \Phi(\tilde{z})$, a contradiction. Suppose instead that $\tilde{z} \notin \hat{Z}$, then $\Theta(z(s,t)) \in \Phi(\tilde{z})$. Consider a sequence $\{\hat{z}_k\} \in \hat{Z}$ such that $\hat{z}_k \to \tilde{z}$. Given that there are a finite number of GA structures, consider $\{\hat{z}'_k\}$, a subsequence of $\{z_k\}$, such that $\hat{z}'_k \to \tilde{z}$ and $(\hat{F}, \hat{\mu}) \in \Phi(\hat{z}_k)$. Given the upper hemicontinuity of demand, and the continuity of prices in z, $(\hat{F}, \hat{\mu}) \in \Phi(\tilde{z})$. Hence $\{(\hat{F}, \hat{\mu}), \Theta(z(s,t))\} \subset \Phi(\tilde{z})$. The uniqueness of the minimum competitive equilibrium price vector implies that both $(\hat{F}, \hat{\mu})$ and $\Theta(z(s,t))$ generate the same price vector at $\tilde{z}$. Lemma 2 then implies that $\tilde{z} \in \mathcal{Z}(s,t)$, but

$$\begin{aligned}
\tilde{z} &= \sup\{z \in (z(s,t+1), z(s,t)) | \Theta(z(s,t)) \neq \Phi(z)\} \\
&> z(s,t+1) \\
&= \sup \mathcal{Z}(s,t)
\end{aligned}$$

This contraction establishes $\Theta(z(s,t)) = \Phi(z)$ for $z \in (z(s,t+1), z(s,t))$.

By Assumption G3, there are only a finite number of steps, and the algorithm reaches $z = 0$, with $\Theta(0) \in \Phi(0)$.□

Proof of Theorem 5: We begin with a few housekeeping matters. First, we normalize the measure of $\mathcal{A}$ to one. Second, it is difficult to work with z, since changes in w alter the mapping between z and (w,h).[19] We therefore work with the utility associated with the outside option, which we denote $h'_{a,n+a}$ to distinguish it from the true value $h_{a,n+a}$. In order to facilitate comparison across (w,h), we choose the maximal value of $h'_{a,n+a}$ such that buyer x_a chooses the outside option for all $(w,h) \in \mathcal{A}$. This is feasible because $\mathcal{A}$ is bounded, and without loss of generality since raising the value of $\bar{U}_a^R$ has no effect on the properties of $\Phi(z)$. Let $\bar{h}_{a,n+a}$, denote

[19] If we alter (w,h), then we alter utilty at the end of the stage and hence the range of z. The mapping between z and (w,h) is therefore affected by (w,h).

the maximal value of $h_{a,n+a}$. We let $\bar{\xi} = \sum_{a=1}^{m}(\bar{h}_{a,n+a} - h_{a,n+a})$. We associate with each $\xi \in [0, \bar{\xi}]$ a vector $\{h'_{a,n+a}\}_{a=1}^{m}$: specifically, if $z \in \left[\sum_{i=s+1}^{m}(\bar{h}_{a,n+a} - h_{a,n+a}), \sum_{a=s}^{m}(\bar{h}_{a,n+a} - h_{a,n+a})\right]$, then,

$$h'_{a,n+a} = \begin{cases} \bar{h}_{a,n+a} \text{ for } a \geq s+1; \\ h_{a,n+a} + \left(z - \sum_{a=s+1}^{m}(\bar{h}_{a,n+a} - h_{a,n+a})\right) \text{ for } a = s; \\ h_{a,n+a} \text{ for } a \leq s-1. \end{cases}$$

For each $(w, h) \in \mathcal{A}$, each $z \in [0, \bar{z}]$ can be mapped into an element of $[0, \bar{\xi}]$. Let $\Phi(\xi)$ denote the mapping from ξ to the set of minimal price GA-structures. Any properties of $\Phi(\xi)$ that hold over $[0, \bar{\xi}]$ will be inherited by $\Phi(z)$ on the appropriate range.

Assumption G has four parts. We begin with G2: there exists a measure one subset of $\mathcal{A}$ such that $\Phi(\xi)$ is single-valued except for a finite set $\{\xi_k\} \in [0, \bar{\xi}]$.

Let $\Xi(w, h)$ denote the set of ξ for which $\Phi(\xi)$ is multi-valued. We say that any set of GA structures that give rise to the same price vector for all values of (w, h) and all values of ξ form an equivalence class and we let $\mathcal{G}^{ec}$ denote the set of equivalence classes.[20] Note that given two GA structures $(F, \mu), (F', \mu') \in \mathcal{G}$, if $F \neq F'$ then it is easy to see that there exist (w, h) such that the two structures generate different prices: find a directed edge in F' that is not in F; suppose that this edge is from y_1 and to y_2; alter the value of y_1 to the buyer allocated by μ' to y_2; this alters the price of y_2 generated by (F', μ') but not that generated by (F, μ). Therefore every element of $\mathcal{G}^{ec}$ involves a single forest F and a collection μ^{ec} of allocations such that any (F, μ) with $\mu \in \mu^{ec}$ generates the same price for all (w, h). We therefore let (F, μ^{ec}) denote a generic element of $\mathcal{G}^{ec}$.

Lemma 4 states that if $\xi \in \Xi(w, h)$ then there exist at least two elements of $\mathcal{G}^{ec}$ generating the minimal price equilibrium. It is therefore sufficient to show that there exist finite ξ at which any two elements of $\mathcal{G}^{ec}$ generate the same price vector. According to Lemma 5 this is the case on a subset $\mathcal{A}'$ of $\mathcal{A}$ with measure one. This establishes G2.

Lemma 5 also establishes that the ξ that end steps are not elements of $\Xi(w, h)$ for $(w, h) \in \mathcal{A}'$. This establishes G3.

We now establish G1: there exists a measure one subset of $\mathcal{A}$ such that $\Phi(\xi)$ has at most 2 elements for all $\xi \in [0, \bar{\xi}]$.

It is sufficient to prove that $\Phi(\xi)$ cannot have three elements. Suppose that there exists a subset $\mathcal{A}'' \subseteq \mathcal{A}'$ with positive measure such that for all $(w, h) \in \mathcal{A}''$, there exists a ξ such that $\Phi(\xi)$ has three elements. Since this set has positive measure, its interior is non-empty. Suppose

[20] Many GA structures give rise to the same price vector. This is because buyers allocated to goods with no successors play no role in pricing. These buyers may be redistributed without affecting the price vector.

$(w_0, h_0) \in$int$\mathcal{A}''$. We construct a perturbation in (w_0, h_0) that lies outside of $\mathcal{A}''$ contradicting the assumption that $(w_0, h_0) \in$int$\mathcal{A}''$.

Let $\{\xi_k^0\}$ denote the set of points in $[0, \bar{\xi}]$ at which $\Phi(\xi)$ has more than two elements. $\{\xi_k^0\}$ has a finite number of elements. Consider first the largest element in $\{\xi_k^0\}$. Call it $\xi_{\max}^0$.

If all three elements of $\Phi(\xi_{\max}^0)$ do not give the same allocation, then one must allocate a good to a buyer that the others do not. Suppose that this structure is GA', the buyer is x_b and the good is y_j. Lower h_{bj} by $\varepsilon_1 > 0$. The other two GA structures remain minimal price competitive equilibria. GA' is no longer associated with a competitive equilibrium since x_b now prefers the good or goods that he is allocated by the other two structures.

If all three elements of $\Phi(\xi_{\max}^0)$ give the same allocation, the graphs must differ. Suppose $(F, \mu), (F', \mu) \in \Phi(\xi_{\max}^0)$. Suppose that the only difference between F and F' is that $(y_j, y_k) \in F$ whereas y_k is a root good in F'. Suppose $\mu(x_b) = y_j$. Lower h_{bk} by $\varepsilon_1 > 0$. This has no effect on (F', μ), so (F', μ) is still a minimum price competitive equilibrium. (F, μ), however, is not, since it generates a price for y_k that is below r_k.

Finally, suppose all three elements of $\Phi(\xi_{\max}^0)$ give the same allocation and have the same root set. By Lemma 6 one element has a directed edge that does not belong to the other two. Suppose that this structure is GA', that the buyer allocated to the source of this edge is x_b, and that this edge is directed towards y_j. Lower h_{bj} by $\varepsilon_1 > 0$. This raises the price of y_j generated by GA'. It does not affect the price of y_j according to the other GA-structures, hence they continue to give the same price. It follows that GA' is not an element of $\Phi(\xi_{\max}^0)$ at this value of (w, h).

Consider now (w_1, h_1) which is equal to (w_0, h_0) except that h_{bj} has been replaced by $h_{bj} - \varepsilon$. Let $\{\xi_k^1\}$ denote the set of points at which $\Phi(\xi)$ has more than two elements. Given the continuity of preferences ε_1 can be chosen small enough that $\{\xi_k^1\}$ has one fewer element than $\{\xi_k^0\}$: every element of $\{\xi_k^1\}$ is in a neighborhood of an element of $\{\xi_k^0\}$ and no element is in the neighborhood of $\xi_{\max}^0$.

We then repeat the arguments above until $\{\xi_k^s\}$ is empty. Since the ε_1 can be arbitrarily small, it follows that $(w_0, h_0) \notin$int$\mathcal{A}''$. This contradiction establishes G1.

We now establish G4: there exists a measure one subset of $\mathcal{A}$ such that the graph of $\Phi(\xi)$ contains no isolated points.

Let $\mathcal{A}''' \subseteq \mathcal{A}'$ denote the measure one set of (w, h) for which $\Phi(\xi)$ takes on one or two values. Fix $(w, h) \in \mathcal{A}'''$. Suppose $\Phi(\xi_0) = \{(F, \mu), (F', \mu')\}$. Assumption A implies $\Phi(\xi)$ is upperhemicontiuous in ξ at ξ_0. This upperhemicontinity and the finite number elements of $\Xi(w, h)$ imply that there exists $\varepsilon > 0$, such that $\Phi(\xi)$ is equal to (F, μ) or (F', μ') on $(\xi_0, \xi_0 + \varepsilon)$.

Without loss of generality, suppose $\Phi(\xi) = (F, \mu)$. (F, μ) is not isolated.

We now show that there exists ε' such (F, μ) is not a competitive equilibrium for $z \in (\xi_0 - \varepsilon', \xi_0)$. If $F = F'$, then $\mu \neq \mu'$, which implies that there are indifferences not captured by F. By Lemma 3, there are more than two elements of $\Phi(\xi_0)$, which contradicts the assumption $(w, h) \in \mathcal{A}'''$. Hence $F \neq F'$. It follows from Lemma 3 that F' contains only one edge not contained in F, otherwise $(w, h) \notin \mathcal{A}'''$. Let x_b be the buyer allocated to the source of this edge. Suppose that the edge is directed towards y_j. Consider pricing y_j in two ways: using (F, μ) and using (F, μ) to price $\mu(x_a)$ and x_a to price y_j. The first way must lead to a higher price for $\xi \in (\xi_0, \xi_0 + \varepsilon)$, since $(F', \mu) \notin \Phi(\xi)$ on this interval. Both lead to the same price at ξ_0. Since derivatives of these prices with respect to ξ are constant by equations (6.3) and (6.4), the second way leads to a higher price for $z \in (\xi_0 - \varepsilon', \xi_0)$. Hence x_a prefers y_0 to $\mu(x_a)$ for $z \in (\xi_0 - \varepsilon', \xi_0)$, and $(F, \mu) \notin \Phi(\xi)$ on this interval.

The upperhemicontinuity of Φ at ξ_0 implies that we can choose ε' such that $(F', \mu') \in \Phi(\xi)$ for $(\xi_0 - \varepsilon', \xi_0)$. Hence (F', μ') is not isolated. This completes G4 and the proof of the theorem.□

Proof of Theorem 6: (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, $\hat{\mu} : X^B \rightarrow Y^B$ is a bijection. Since $x_a \notin X^B \Longrightarrow \bar{\mu}_a \notin Y^B$, we know that $\bar{\mu} : X \backslash X^B \rightarrow Y^A$.

Since $\bar{\mu}$ is an allocation it is 1-1 on this domain and only assigns null goods appropriately, ensuring that μ' itself one-to-one and is an allocation. Moreover, since $\bar{\mu}$ is onto $H(\hat{p})\backslash Y^B$, 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.□

Lemma 2 Suppose $(F, \mu) \in \Theta(z(s,t))$ and $(F, \mu), (F', \mu') \in \Phi(z')$ where $z' < z(s,t)$. Then, given Assumptions A and G3, there exists at z' a buyer x_b who is indifferent between to goods y_1 and y_2 such that $(y_1, y_2) \notin F$.

Proof: First, suppose that $\mu \neq \mu'$, and consider $x_a \in X$ such that $\mu_a \neq \mu'_a$. If there exists no x_b such that $\mu_b = \mu'_a$, then we are done: $(\mu_a, \mu'_a) \notin F$ since (μ, F) is a GA-structure and the endpoints of each edge in a GA structure are allocated. If there exists $\mu_b = \mu'_a$, consider μ'_b. If there exists no x_c such that $\mu_c = \mu'_b$ we are done. If $\mu'_b = \mu_a$, we are done, as we cannot have $(\mu_a, \mu'_a) \in F$ and $(\mu'_a, \mu_a) \in F$ since F is a forest. If there exists $\mu_c = \mu'_b$, we continue as above until we eventually find a cycle or an unallocated good.

Now suppose $\mu = \mu'$. Suppose that the lemma is false. Then F' is a subforest of F, and there exists a $\hat{y} \in R(F')$ such that $\hat{y} \notin R(F)$. Given Assumptions A and G3, however, every $y_i \notin R(F)$ has $p_i > r_i$ for all $z < z(s,t)$. This contradiction establishes the lemma.□

Lemma 3 Consider $(F, \mu) \in \Phi(\xi)$. Suppose that for $y_i \notin R(F)$, $q_i(\xi, F, \mu) > r_i$. Suppose that there exists a buyer x_b who is indifferent between $\mu(x_b)$ and y_0 at $q(\xi, F, \mu)$, and that the edge $(\mu(x_b), y_0) \notin F$. Suppose that this is the only such buyer. Then $\Phi(\xi)$ has exactly two elements.

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

Case 1: y_0 is unallocated by μ.

Let $\hat{Y}$ denote the component 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, $q_i(\xi, F, \mu) > r$ 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, and that given any $(F', \mu') \in \Phi(\xi)$, μ' must be onto $\hat{Y}\backslash y_R$. Since there are no indifferences not encoded in $F \cup (\mu(x_b), y_0)$, only buyers in $\hat{X}$ demand goods in $\hat{Y} \cup y_0$. It follows that no $(F', \mu') \in \Phi(\xi)$ can be onto both y_R and y_0. To price all goods with price above reservation one needs as many indifferences as goods. It follows 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: y_0 is allocated but not a predecessor to $\mu(x_b)$.

In this case, there is only one way to price all goods that are not y_0 or successors of y_0. There are 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-structures.

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

The addition of the edge $(\mu(x_b), y_0)$ creates a circuit. 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.

Cases 1 through 3 are exhaustive. This completes the lemma.□

Lemma 4 Suppose that $\Phi(\xi)$ takes on multiple values at $\xi \in [0, \bar{\xi}]$, then there exist at least two elements of $\mathcal{G}^{ec}$ that generate the minimal price competitive equilibrium.

Proof: We begin with some observations that we will use to identify when two GA structures belong to different elements of $\mathcal{G}^{ec}$. Recall given two GA structures (F, μ) and (F', μ') if $F \neq F'$ then (F, μ) and (F', μ') belong to different elements of $\mathcal{G}^{ec}$. Moreover, if $F = F'$ and $\mu(x_a) \neq \mu'(x_a)$ where $\mu'(x_a)$ has a successor in F', then we can the value to x_a of the successor of $\mu'(x_a)$. This alters the price generated by (F', μ') but not that generated by (F, μ). Again (F, μ) and (F', μ') belong to different elements of $\mathcal{G}^{ec}$.

Now let $(F, \mu) \neq (F', \mu') \in \Phi(\xi)$. If $F \neq F'$, the result is immediate. If $F = F'$, then $\mu \neq \mu'$, which implies that there exists a buyer x_a who is allocated to y_j under μ and $y_k \neq y_j$ under F'. If either y_j or y_k have successors, then the result is immediate. If neither y_j nor y_k have successors, x_a's indifference between y_j and y_k is not reflected in either F or F'. Consider F'' which replaces the edge ending in y_j by (y_j, y_k). $(F'', \mu) \in \Phi(\xi)$ and $F'' \neq F$. The result is immediate using (F, μ) and (F'', μ).□

Lemma 5 Suppose that buyer's utility takes the form of equation (6.2). Consider (F, μ^{ec}),$(\hat{F}, \hat{\mu}^{ec}) \in \mathcal{G}^{ec}$. There exists a subset $\mathcal{A}' \subseteq \mathcal{A}$ with measure one such that the prices generated by (F, μ^{ec}) and $(\hat{F}, \hat{\mu}^{ec})$ are equal for at most two values of ξ for all $(w, h) \in \mathcal{A}'$. Moreover, neither of these values lie in the set $\left\{0, \left\{\sum_{a=1}^{k} (\bar{h}_{a,n+a} - h_{a,n+a})\right\}_{k=1}^{m}\right\}$.

Proof: Fix (w, h). Consider (F, μ^{ec}),$(\hat{F}, \hat{\mu}^{ec}) \in \mathcal{G}^{ec}$.

We first define a few useful objects. Let $q(\xi, F, \mu^{ec})$ denote the price generate by (F, μ^{ec}) at z and $q'(\xi, \hat{F}, \hat{\mu}^{ec})$ the price generated by $(\hat{F}, \hat{\mu}^{ec})$. Let Z and Z' denote the subsets of $[0, \bar{\xi}]$ on which q and q' are real valued respectively. For each $y_s \in Y$ define $\Lambda(y_s, F, \mu^{ec}) = \{y_0, x_{a_0}, y_1, x_{a_1}, y_2, x_{a_3} \dots y_s\}$ such that (1) $y_0 \in R(F)$; (2) for $v = \{0, \dots, s-1\}$, y_v is a successor of y_{v-1}, and (3) $y_v = \mu(x_{a_v})$.

Suppose that there exists $\xi \in Z \cap Z'$ such that $q = q'$. Given that $(F, \mu^{ec}) \neq (\hat{F}, \hat{\mu}^{ec})$, there must exist some good $y_s \in Y$ such that $\Lambda(y_s, F, \mu^{ec}) \neq \Lambda(y_s, \hat{F}, \hat{\mu}^{ec})$. A necessary condition for $q = q'$ is that the price of y_s is the same in each case. We show that this occurs at a finite number of z.

We consider three cases.

Case 1: For both (F, μ^{ec}) and $(\hat{F}, \hat{\mu}^{ec})$ the root good in the component containing y_s is an element of Y.

Given $\Lambda(\hat{y}, F, \mu^{ec}) = \{y_0, x_{a_0}, y_1, x_{a_1}, y_2, x_{a_3} \dots y_s\}$, the price of y_s is:

$$\left(1 - \frac{h_{a_{s-1},s-1}}{h_{a_{s-1},s}}\right) w_{a_{s-1}} + \frac{h_{a_{s-1},s-1}}{h_{a_{s-1},s}} \left(1 - \frac{h_{a_{s-2},s-2}}{h_{a_{s-2},s-1}}\right) w_{a_2} + \dots \\ + \left(\frac{h_{a_{s-1},s-1}}{h_{a_{s-1},s}} \frac{h_{a_{s-2},2}}{h_{a_{s-2},1}} \dots \frac{h_{a_1,1}}{h_{a_1,0}}\right) \left(1 - \frac{h_{a_0,0}}{h_{a_0,1}}\right) w_{a_0} \tag{10.1}$$

This depends only on (w, h) and is independent of z. It is clear that (F, μ^{ec}) and $(\hat{F}, \hat{\mu}^{ec})$ generate the same prices on a measure zero subset of (w, h). Call this subset $A_1(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{ec}))$.

Case 2: For $\hat{F}'$ the root good in the component containing y_s is an element of Y, but that for F it is a null good.

In this case only q' is fixed. The price of y_s according to (F, μ^{ec}) takes a form similar to (10.1) except that $h_{a_0,0} = h_{a,n+a}$, the value of the outside option. It is clear that prices of goods in the component containing y_s are strictly increasing in z during the stage in which the value of y_0 is falling. The strict monotonicity of q in z, implies that if $q = q'$ in the interior of a stage, then there is only one z for which $q = q'$. Suppose that $q = q'$ at the end of a stage. This is just like the case above. It happens on only a measure zero subset of A. Denote this measure zero subset $A_2(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{ec}))$.

Case 3: For both (F, μ^{ec}) and $(\hat{F}, \hat{\mu}^{ec})$ the root good in the component containing y_s is an element of $Y_\varnothing$. If these roots are different, then the analysis is similar to the above. Since the value of the root goods are lowered one at a time, only one of $p_s(F, \mu^{ec})$ or $p_s(\hat{F}', \hat{\mu}')$ will rise and then later the other will rise. It follows that there are at most two z for which the prices of

y_s are equal.

If the roots are the same, then we compare the derivatives. Let p_s denote the price of y_s:

$$\frac{dp_s}{d\xi} = \begin{array}{cc} \left(\frac{w_0}{h_{a_0,1}}\right)\left(\frac{h_{a_1,1}}{h_{a_1,2}}\right)\cdots\left(\frac{h_{a_{s-1},s-1}}{h_{a_{s-1},s}}\right) & \text{if } \frac{dh_{a_0,0}}{d\xi}=1 \\ 0 & \text{otherwise} \end{array}$$

It is clear that these will be equal only for a measure zero set of (w, h). If the roots are different, then there can be at most two points at which the prices of y_s are equal, unless they are equal at the end of a stage which happens on a measure zero set as above. Collect these measure zero sets and label them $A_3(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{ec}))$.

The set on which (F, μ^{ec}) and $(\hat{F}, \hat{\mu}^{ec})$ give the same prices at more than a finite z is a subset of

$$\begin{aligned} A(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{ec})) &= A_1(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{ec})) \cup \ldots \\ & A_2(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{ec})) \cup A_3(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{ec})). \end{aligned}$$

Summing over goods and elements of $\mathcal{G}^{ec}$ gives a measure zero set $A_4 = \cup_{\{(F,\mu^{ec}),(\hat{F},\hat{\mu}^{ec})\}} \left\{\cup_s A(y_s, (F, \mu^{ec}), (\hat{F}, \hat{\mu}^{e\right.$
We conclude that for all but a measure zero set of (w, h) any two q,q' coincide at zero, one or two ξ. Let $\mathcal{A}' = \mathcal{A} \backslash A_4$.

The other statements were proved in the course of this proof. □

Lemma 6 Suppose that $\Phi(\xi)$ has three elements that give the same allocation and have the same set of root goods, then one element has a directed edge that does not belong to the other two.

Proof: Suppose that $(F, \mu), (F', \mu), (F'', \mu) \in \Phi(z)$. Since $F \neq F'$, there exists an edge in one that is not in the other. Without loss of generality suppose that there exits and edge $(y_j, y_k) \in F$ such that $(y_j, y_k) \notin F'$.

Since y_k is not a root good, there must be an edge directed towards y_k. Hence there exists y_l such that $(y_l, y_k) \in F'$. Note $(y_l, y_k) \notin F$ since each good has at most one incoming edge. The observation that (y_j, y_k) and (y_l, y_k) cannot both be in F'' completes the this case.□

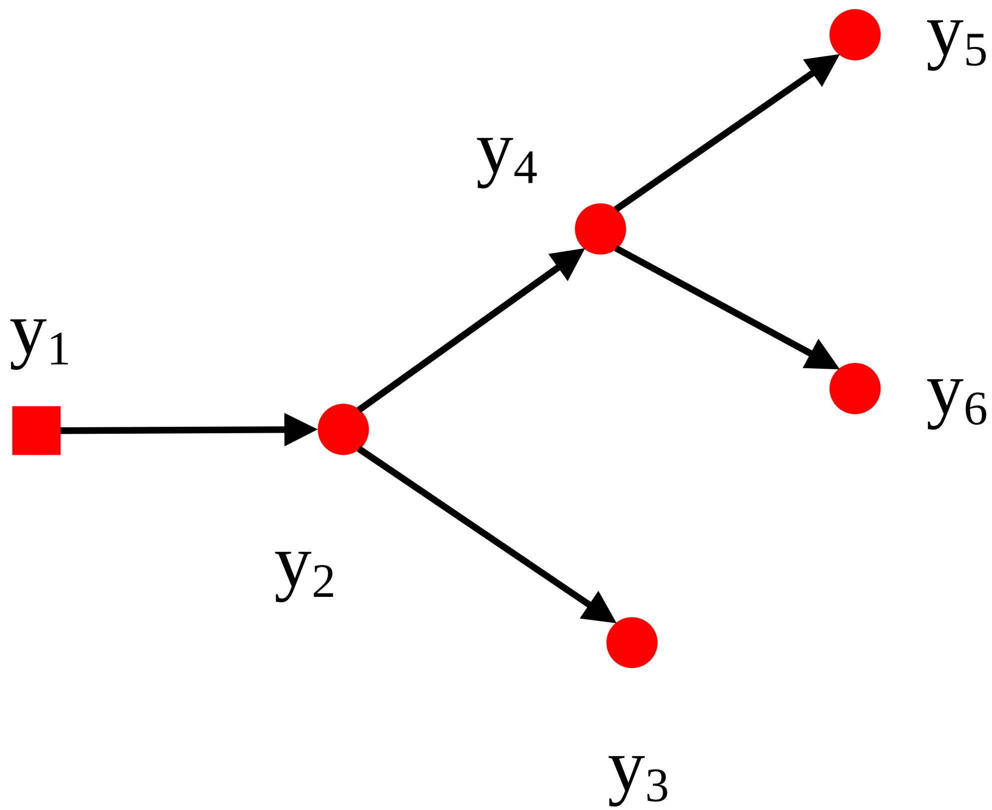

Figure 1: A directed rooted tree with edges directed away from the root good (y_1)

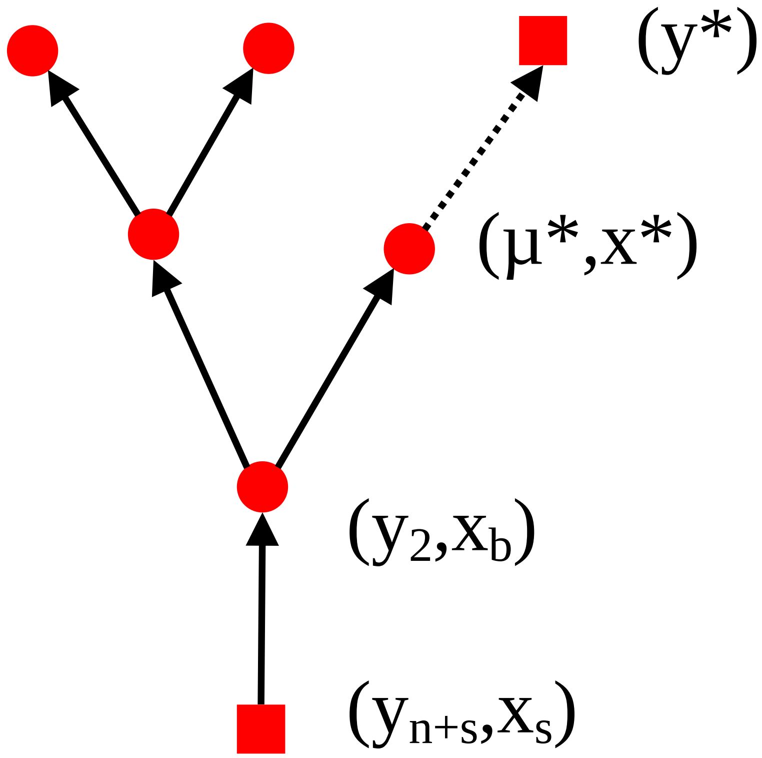

Figure 2(a)

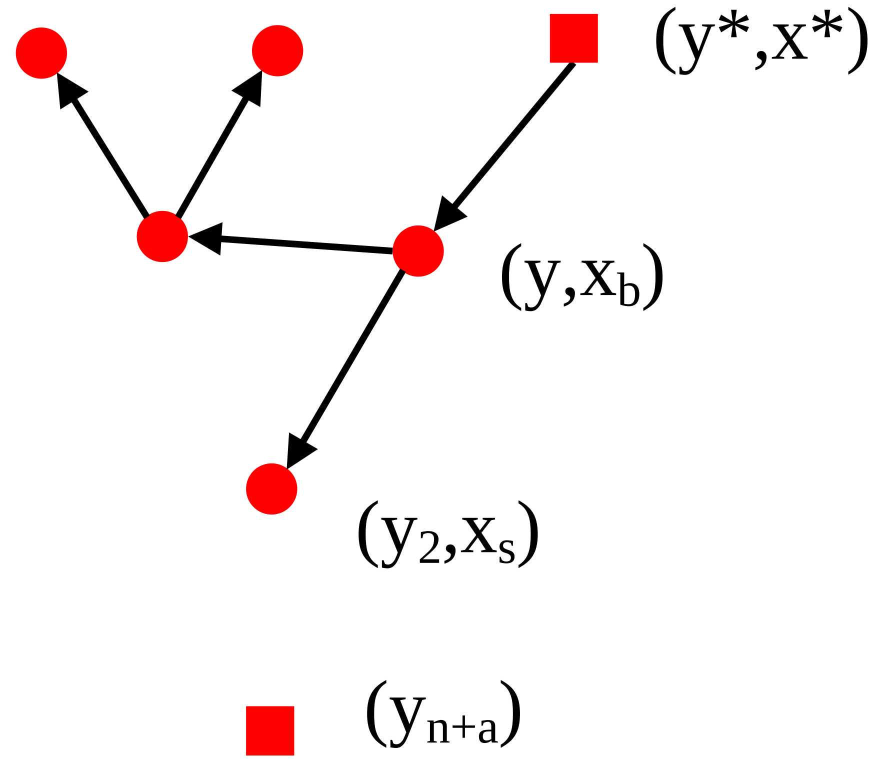

Figure 2(b)

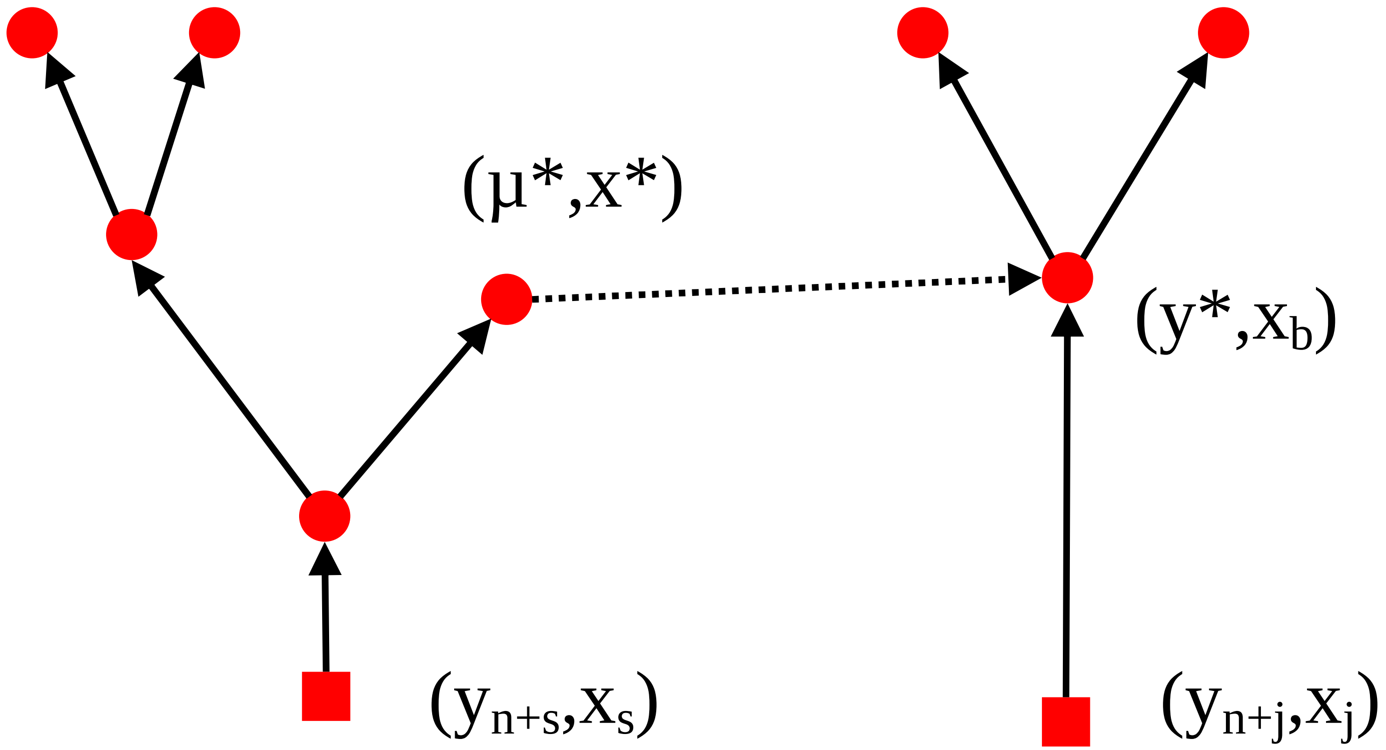

Figure 3(a)

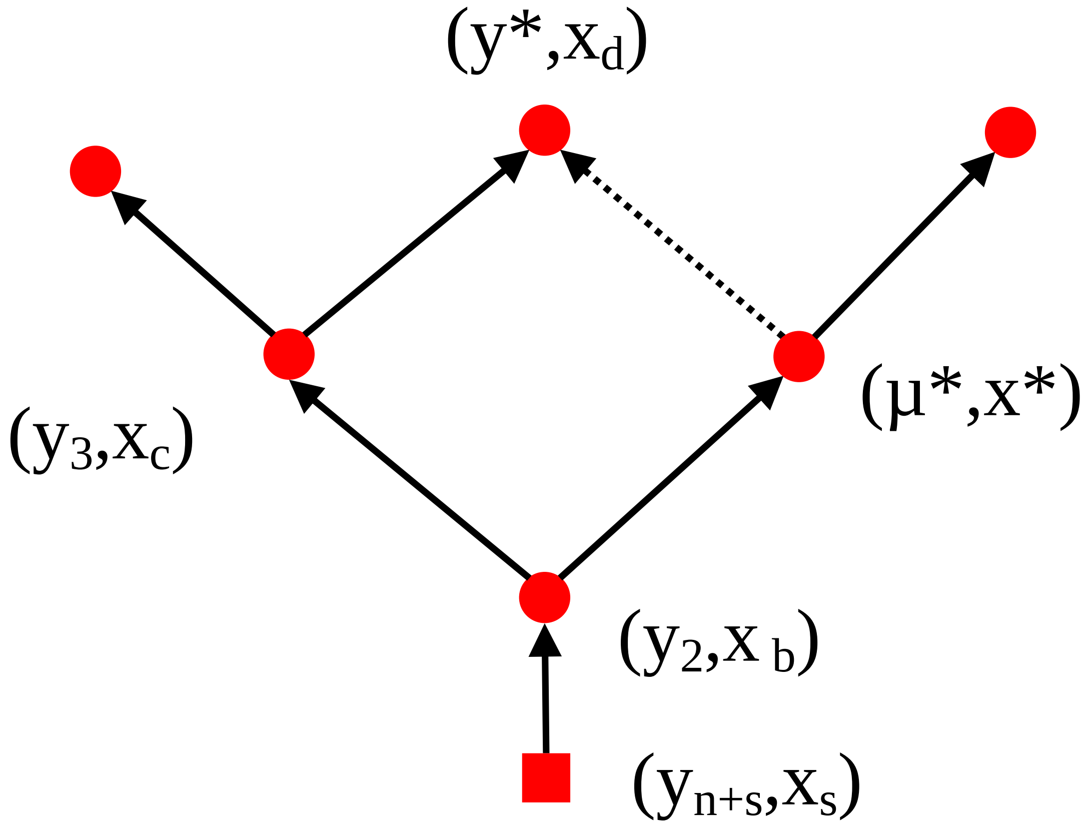

Figure 3(b)

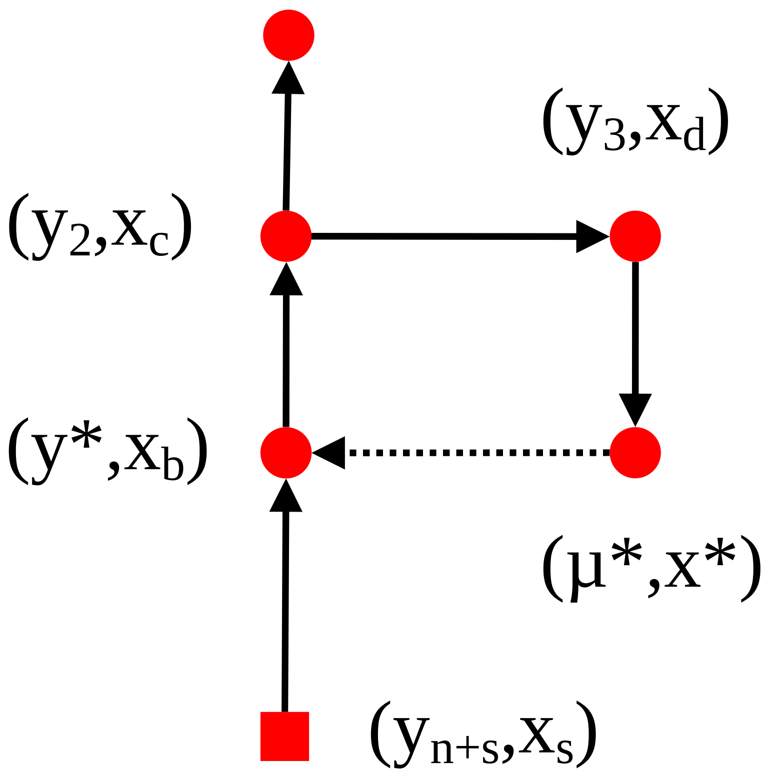

Figure 4(a)

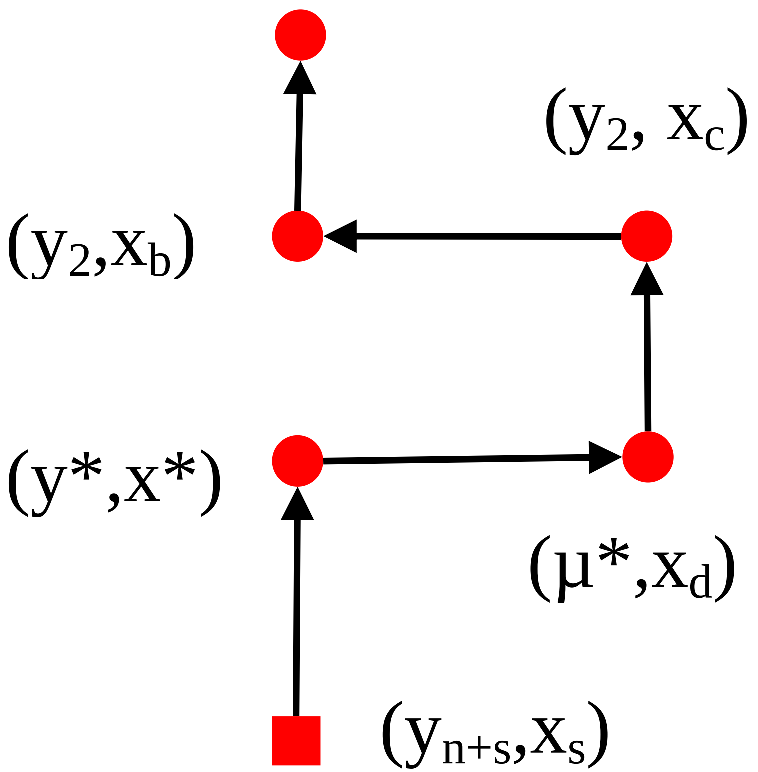

Figure 4(b)