The Dynamics of Pure Market Exchange

Herbert Gintis*

December 16, 2011

> The problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess.
>
> F. A. Hayek (1945), p. 519.

Abstract

This paper investigates the out-of-equilibrium price and quantity adjustment process for a decentralized market economy. The economy is modeled as a Markov process. Each agent has a set of *private prices* that are updated through experience, less successful agents copying the strategies of more successful agents, as well as varying private prices in response to personal trading experience.

With an initial random assignment of private price vectors to agents, this Walrasian economy with Markov dynamics quickly moves to an ergodic subprocess with a set of *quasi-public prices*, in which the standard error of private prices across agents becomes very small. In the long run, over a wide range of parameters, the stationary distribution of this Markov process approximates a Walrasian equilibrium of the system. We call this stationary distribution a *quasi-Walrasian equilibrium*.

When agents are permitted to trade in goods they neither produce nor consume, a money good appears in the stationary distribution. The resulting dynamical system is, moreover, highly resilient in the face of exogenous shocks.

These findings suggest that the Markov process is an appropriate analytical tool for modeling the dynamics of a Walrasian market economy, and a fully decentralized Walrasian economy in which information is derived purely from individual trading experience can be stable under a wide range of conditions.

*Santa Fe Institute and Central European University

1 Introduction

Adam Smith (2000[1759]) envisioned a decentralized economy that sustains an efficient allocation of resources through the "invisible hand" of market competition. Smith's vision was formalized by Léon Walras (1954 [1874]), and a proof of existence of equilibrium for a simplified version of the Walrasian economy was provided by Wald (1951 [1936]). Soon after, Debreu (1952), Arrow and Debreu (1954), Gale (1955), Nikaido (1956), McKenzie (1959), Negishi (1960), and others contributed to a rather complete proof of the existence of equilibrium in Walrasian economies. Such economies are particularly attractive because they capture the basic structural characteristics of market economies, and because a Walrasian equilibrium is Pareto-efficient (Arrow 1951, Debreu (1951, 1954), Hurwicz 1960).

The question of stability of the Walrasian economy was a central research focus in the years immediately following the existence proofs (Arrow and Hurwicz, 1958, 1959, 1960; Arrow, Block and Hurwicz, 1959; Nikaido 1959; McKenzie, 1960; Nikaido and Uzawa 1960; Uzawa 1960a,b). The models investigated assumed that out of equilibrium, there is a system of *public prices* shared by all agents, the time rate of change of prices being a function of excess demand. The public price system was implemented by a single agent (the "auctioneer") acting outside the economy to update prices in the current period on the basis of the current pattern of excess demand, using a process of "tâtonnement," as was first suggested by Walras (1954 [1874]) himself.

The quest for a general stability theorem was derailed by Herbert Scarf"s simple examples of unstable Walrasian equilibria (Scarf 1960). There were attempts soon after to continue the analysis of tâtonnement by adding trading out of equilibrium (Uzawa 1959, 1961, 1962; Negishi 1961; Hahn 1962; and Hahn and Negishi 1962), but with only limited success.

General equilibrium theorists in the early 1970's harbored some expectation that plausible restrictions on utility functions might entail stability, because gross substitutability was known to imply global stability (Arrow et al. 1959) and gross substitutability was known to hold for Cobb-Douglas and many other utility functions (Fisher 1999). However, gross substitutability is not a property of constant elasticity of substitution (CES) and more general utility functions. Moreover, Sonnenschein (1973), Mantel (1974, 1976), and Debreu (1974) showed that any continuous function, homogeneous degree zero in prices, and satisfying Walras' Law, is the excess demand function for some Walrasian economy.

However, Hahn and Negishi (1962) showed that if out-of-equilibrium trade is permitted and the so-called *Hahn condition* obtains, then the Walrasian equilibrium is stable under tâtonnement. The Hahn condition says that markets are sufficiently informationally complete that if there is aggregate excess demand then no individual experiences excess supply, and if there is aggregate excess supply, then no individual experiences excess demand. Fisher (1983) significantly broadened this model and proved stability assuming a expectational condition ("no favorable surprise") that should logically hold in any Walrasian equilibrium, plus a weakened Hahn condition, according to which lower-price sellers realize their plans before higher-price sellers do.

Nevertheless, surveying the state of the art some quarter century after Arrow and Debreu's seminal existence theorems, Fisher (1983) concluded that little progress had been made towards a cogent model of Walrasian market dynamics. More recent studies have shown that the tâtonnement dynamic is stable only under extremely stringent conditions (Kirman 1992). Indeed, chaos in price movements is the generic case for the tâtonnement adjustment processes (Saari 1985, Bala and Majumdar 1992). Saari (1995) and others have shown that the information needed by a price adjustment mechanism that can ensure stability include virtually complete knowledge of all cross-elasticities of demand in addition to excess demand quantities.

It is now more than another quarter century since Fisher's seminal contributions, but it remains the case that, despite the centrality of the general equilibrium model to economic theory, we know nothing systematic about Walrasian market dynamics in realistic settings. We show that when the market system is modeled as a Markov process rather than a system of first-order differential equations, a powerful analytical dynamic emerges.

2 A Markov Process Primer

We will model the Walrasian economy as a Markov process, starting for illustrative purposes with a very elementary and purely didactic example. Our goal is to show that even a very simple finite Markov process may have too many states to permit an analytical solution, yet the global properties of the model are easily understood.

Consider an economy with k goods, each of which could serve equally well as a money good for the economy. Suppose there are n agents in the economy, and in each period, each agent is willing to accept exactly one good as money. Suppose further that in each period, one agent switches from his own preferred money to that of another randomly encountered agent. We can describe the state of the economy as a k-vector $(w_1 \dots w_k)$, where w_i is the number of agents who accept good i as money. The total number of states in the economy is thus the number of different ways to distribute n indistinguishable balls (the n agents) into k distinguishable boxes (the k goods), which is $C(n+k-1, k-1)$, where $C(n,k) = n!/(n-k)!k!$ is the number of ways to choose k objects from a set of n objects. To verify this formula, write a particular state in the form

$$s = x \dots xAx \dots xAx \dots xAx \dots x$$

where the number of x's before the first A is the number of agents choosing good 1 as money, the number of x's between the $(i-1)^{\text{th}}A$ and the $i^{\text{th}}A$ is the number of agents choosing good i as money, and the number of x's after the final A is the number agents choosing good k as money. The total number of x's is equal to n, and the total number of A's is $k-1$, so the length of s is $n+k-1$. Every placement of the $k-1$ A's represents particular state of the system, so there are $C(n+k-1, k-1)$ states of the system.

For instance, if $n = 100$ and $k = 10$, then the number of states S in the system is $S = C(109, 9) = 4{,}263{,}421{,}511{,}271$. Suppose in each period two agents are randomly chosen and the first agent switches to using the second agent's money good as his own money good (the two money goods may in fact be the same). This gives a determinate probability p_{ij} of shifting from one state i of the system to any other state j. The matrix $P = \{p_{ij}\}$ is called a *transition probability matrix*, and the whole stochastic system is called a *Markov process*. The Markov process is finite because it has a finite number of states.

What is the long-run behavior of this Markov process? Note first that if we start in state i at time $t = 0$, the probability $p_{ij}^{(2)}$ of being in state j in period $t = 2$ is simply

$$p_{ij}^{(2)} = \sum_{k=1}^{S} p_{ik} p_{kj} = (P^2)_{ij}. \tag{1}$$

This is true because to be in state j at $t = 2$ the system must have been in some state k at $t = 1$ with probability p_{ik}, and the probability of moving from k to j is just p_{kj}. This means that the two period transition probability matrix for the Markov process is just P^2, the matrix product of P with itself. By similar reasoning, the probability of moving from state i to state j in exactly r periods, is P^r. Therefore, the time path followed by the system starting in state $s^0 = i$ at time $t = 0$ is the sequence $s^0, s^1, \dots$, where

$$\mathrm{P}[s^t = j | s^0 = i] = (P^t)_{ij} = p_{ij}^{(t)}.$$

The matrix P in our example has $S^2 \approx 1.818 \times 10^{15}$ entries. The notion of calculation P^t for even small t is quite infeasible. There are ways to reduce the calculations by many orders of magnitude (Gintis 2009, Ch. 13), but even these are completely impractical with so large a Markov process.

Nevertheless, we can easily understand the dynamics of this Markov process. We first observe as that if the Markov process is ever in the state

$$s_*^r = (0_1, \ldots, 0_{r-1}, n_r, 0_{r+1} \ldots 0_k),$$

where all n agents choose good r money, then s_*^r will be the state of the system in all future periods. We call such a state *absorbing*. There are clearly only k absorbing states for this Markov process. We next observe that from any non-absorbing state s, there is a strictly positive probability that the system moves to an absorbing state before returning to state s. For instance, suppose $w_i = 1$ in state s. Then there is a positive probability that w_i increases by 1 in each of the next $n-1$ periods, so the system is absorbed into state s_*^i without ever returning to state s. Now let $p_s > 0$ be the probability that Markov process never returns to state s. The probability that the system returns to state s at least q times is thus at most $(1 - p_s)^q$. Since this expression goes to zero as $q \to \infty$, it follows that state s *appears only a finite number of times with probability one.* We call s a *transient* state.

Because state s in the previous argument was an arbitrary non-absorbing state, it follows that *all non-absorbing states are transient.* It is also clear that with probability one there will be some period t such that no transient state reappears after period t. This means that with probability one the system is absorbed into one of the k absorbing states, from which it never emerges. Inevitably in this system one good emerges as the money good for the economy.

We can in fact often calculate the probability that a system starting out with w_r agents choosing good r as money, is absorbed by state s_*^r. Let us think of the Markov process as that of k gamblers, each of whom starts out with an integral number of k coins, there being n coins in total. The gamblers represent the goods and their coins are the agents who choose that good for money. We have shown that in the long run, one of the gamblers with have all of the coins, with probability one. Suppose the game is fair in the sense that in any period a gambler with a positive number of coins has an equal chance to increase or decrease his wealth by one coin. Then the expected wealth of a gambler in period $t+1$ is just his wealth in period t. Similarly, the expected wealth $E[w^{t'}|w^t]$ in period $t' > t$ of a gambler whose wealth in period t is w^t is $E[w^{t'}|w^t] = w^t$. This means that if a gambler starts out with wealth $w > 0$ and he wins all the coins with probability q_w, then $w = q_w n$, so the probability of being the winner is just $q_w = w/n$.

We now can say that this Markov process, despite its enormous size, can be easily described as follows. Suppose the process starts with $w_r = w$. Then in a finite number of time periods, the process will be absorbed into one of the states $s_*^1, \ldots, s_*^k$, and the probability of being absorbed into state s_*^r is w/n. As it turns out, this description of a finite Markov process is completely general, except for a few technical points. In general, however, the process will not be absorbed into a single state, but rather into what is called an *irreducible Markov subprocess*, making transitions among a number of states, called *communicating recurrent states*. If the process is aperiodic, the fraction of time it is in each of these recurrent states state forms a probability distribution called the *stationary distribution* of the "absorbing" Markov subprocess. If the process is periodic, it can be conveniently subdivided into a number of aperiodic subprocesses, each with a stationary distribution.

3 From Differential Equations to Markov Processes

A plausible model of market dynamics should reflect two fundamental aspects of market competition. First, trades must be bilateral with separate budget equations for each transaction (Starr 1972). The second is that in a decentralized market economy out of equilibrium, *there is no price vector for the economy at all*. The assumption that there is a system of in prices that are common knowledge to all participants (we call these *public prices*) is plausible in equilibrium, because all agents can observe these prices in the marketplace. However, out of equilibrium there is no vector of prices determined by market exchange. Rather, assuming Bayesian rational agents, every agent has a subjective prior concerning prices, based on personal experience, that the agent uses to formulate and execute trading plans.

Consider, for instance, the wage rate for a particular labor service. In equilibrium this price may be common knowledge, but out of equilibrium, every supplier of this service must have an estimate of the probability of selling his service as a function of his offer price. The supplier, if Bayesian rational, will have a subjective prior representing the shape of the demand function for the service. This prior will determine whether the supplier accepts a particular wage offer, or rejects the offer and continues searching for a better offer. The information needed to form this prior includes the distribution of subjective priors of demanders for the labor service, while the subjective prior of each demander will depend in a similar manner on the distribution of subjective priors of suppliers of the service. Thus in this case information is not simply asymmetrically distributed, but rather is effectively indeterminate, since the supply schedule depends on suppliers' assessment of demand conditions, and the demand schedule depends on demanders' assessment of supply schedules. The conditions under which each agent's choice is a best response to the others', even assuming common knowledge of rationality and common knowledge of the Markov process (which in this case is the game in which the players are engaged), are quite stringent and normally not present (Aumann and

Brandenburger 1995).

In analyzing market disequilibrium, we must thus assume that each agent's subjective prior includes a vector of *private* prices that is modified adaptively through the exchange experience. The admissible forms of experience in a decentralized market economy are those that result from observing the behavior of trading partners. This experience is the sole basis for a trader's updating his private price vector, and equilibrium can be achieved only if plausible models of inference and updating lead private prices to converge to equilibrium prices (Howitt and Clower 2000).

In the interest of simplicity in dealing with a daunting problem that has defied solution for more than a century, we will assume that there are no institutions other than markets where individuals congregate to exchange their wares, there are no forms of wealth other than agents' production goods, all transactions take place in the current period, so there is no intertemporal planning, and there is no arbitrage beyond that which can be executed by an agent engaged in a series of personal trades using no information save that acquired through personal trading experience.

If we can show stability under these conditions (which we can) we then know that trader information concerning macro aspects of the economy plays no essential role in achieving market equilibrium, and indeed may well have a destabilizing effect (Gintis 2007).

It follows logically that in so simple a market system with rational actors but no public institutions, *expectations are purely adaptive*. The "rational expectations" notion that agents know the global structure of the economy and use macroeconomic information to form accurate expectations is not plausible in the decentralized context. This conclusion may of course require revision in a model with an institutional structure that creates public information, such as a credible government or national bank.

An appropriate candidate for modeling the Walrasian system in disequilibrium is a Markov process. The states of the process are vectors whose components are the states of individual agents. The state of each agent includes his holding of each good, an array of parameters representing his search strategies for buying and selling, parameters representing his linkage to others in a network of traders, and finally his vector of private prices, which the agent uses to evaluate trading offers.

If state s_i has a positive probability of making a transition to state s_j in a finite number of periods (i.e., $p_{ij}^{(t)} > 0$ for some positive integer t), we say s_i *communicates* with s_j. If all states in a Markov process communicate with each other, we say the process is *irreducible*. We cannot assume a Markov model of a Walrasian system is irreducible because, as in our elementary example above, where a good inevitably emerges as money, the states of the system with one good as money will not communicate with the states where a different good is money.

The Markov model of a Walrasian economy is finite if we assume there are a finite number of agents, a finite number of goods, a minimum discernible quantity of each good, and a finite inventory capacity for each good. A strictly positive probability of remaining in the same state for an agent then ensures that the Markov process is finite it aperiodic. Assuming for the moment that the Markov process is irreducible, being both and aperiodic implies the Markov process is *ergodic* (Feller 1950), which means has a stationary distribution expressing the long-run probability of being in each state of the system, irrespective of its initial state.

We may not care about individual states of the process, but rather about certain aggregate properties of the system, including the mean and standard deviation of prices, and the aggregate pattern of excess demand. As we shall see, the ergodic theorem ensures that under appropriate conditions these aggregates have determinate long-term stationary distributions.

THE ERGODIC THEOREM: Consider an n-state aperiodic and irreducible Markov process with transition matrix P, so the t-period transition probabilities are given by $P^{(t)} = P^t$. Then there is probability distribution $u = (u_1, \ldots, u_n)$ over the states of the Markov process with strictly positive entries that has the following properties for $j = 1, \ldots, n$:

$$u_j = \lim_{t \to \infty} p_{ij}^{(t)} \quad \text{for } i = 1, \ldots, n \tag{2}$$

$$u_j = \sum_{i=1}^{n} u_i p_{ij} \tag{3}$$

We call u the *stationary distribution* of the Markov process.

Equation (2) says that with probability one, u_j is the long-run frequency of state s_j in a realization $\{s^t\} = \{s^0, s^1, \ldots\}$ of the Markov process. Also, this frequency is strictly positive and independent from the starting state $s^0 = s_i$. By a well-known property of convergent sequences, (2) implies that u_j is also with probability one the limit of the average frequency of s_j from period t onwards, for any t. This is in accord with the general notion in a dynamical system that is ergodic, the equilibrium state of the system can be estimated as an historical average over a sufficiently long time period (Hofbauer and Sigmund 1998).

Equation (3) is the *renewal equation* governing stationary distribution u. It asserts that in the long run, the probability of being in state s_j is the sum over i of the probability that it was in some state s_i in the previous period, multiplied by the probability of a one-period transition from state s_i to state s_j, independent from the initial state of the realization $\{s^t\} = \{s^0, s^1, \ldots\}$.

A Markov process thus has only a one period "memory." However, we can consider a finite sequence of states $\{s^{t-l}, s^{t-l+1}, \dots, s^t\}$ of the Markov process of fixed length l as a single state, the process remains Markov and has as an l-period "memory." Because any physically realized memory system, including the human brain, has finite capacity, the finiteness assumption imposes no constraint on modeling systems that are subject to physical law.

To see this, suppose a Markov process has transition matrix $P = \{p_{ij}\}$ and consider two-period states of the form ij. We define the transition probability of going from ij to kl as

$$p_{ij,kl} = \begin{cases} p_{j,l} & j = k \\ 0 & j \neq k, \end{cases} \tag{4}$$

This equation says that ij represents "state s_i in the previous period and state s_j in the current period." It is easy to check that with this definition the matrix $\{p_{ij,kl}\}$ is a probability transition matrix, and if $\{u_1, \dots, u_n\}$ is the stationary distribution associated with P, then

$$u_{ij} = u_i p_{i,j} \tag{5}$$

defines the stationary distribution $\{u_{ij}\}$ for $\{p_{ij,kl}\}$. Indeed, we have

$$\lim_{t\to\infty} p_{ij,kl}^{(t)} = \lim_{t\to\infty} p_{j,k}^{(t-1)} p_{k,l} = u_k p_{k,l} = u_{kl}$$

for any pair-state kl, independent from ij. We also have, for any ij,

$$u_{ij} = u_i p_{i,j} = \sum_k u_k p_{k,i} p_{i,j} = \sum_k u_{ki} p_{i,j} = \sum_{kl} u_{kl} p_{kl,ij}. \tag{6}$$

It is straightforward to show that pairs of states of P correspond to single states of $\{p_{ij,kl}\}$. These two equations imply the ergodic theorem for $\{p_{ij,kl}\}$ because equation 5 implies $\{u_{ij}\}$ is a probability distribution with strictly positive entries, and we have the defining equations of a stationary distribution; for any pair-state ij,

$$u_{kl} = \lim_{t\to\infty} p_{ij,kl}^{(t)} \tag{7}$$

$$u_{ij} = \sum_{kl} u_{kl} p_{kl,ij}. \tag{8}$$

An argument by induction extends this analysis to any finite number of sequential states of P.

An important question is the nature of collections of states of a finite Markov process. For instance, we may be interested in total excess demand for a good without caring how this breaks down among individual agents. From the case of two states j and k it will be clear how to generalize to any finite number. Let us make being in either state j or in state k into a new macro-state m. If P is the transition matrix for the Markov process, the probability of moving from state i to state m is just $P_{im} = P_{ij} + P_{ik}$. If the process is ergodic with stationary distribution u, then the frequency of m in the stationary distribution is just $u_m = u_j + u_k$. Then we have

$$u_m = \lim_{t\to\infty} P^n_{im} \tag{9}$$

$$u_m = \sum_i u_i p_{im} \tag{10}$$

However, the probability of a transition from m to a state i is given by

$$P_{mi} = u_j p_{ji} + u_k p_{ki}. \tag{11}$$

Now suppose states j and k are *interchangable* in the sense that $p_{ji} = p_{ki}$ for all states i. Then (11) implies

$$u_i = \sum_r u_r p_{ri}, \tag{12}$$

where r ranges over all states except j and k, plus the macro state m. In other words, if we replace states j and k by the single macro-state m, the resulting Markov process has one fewer state, but remains ergodic with the same stationary distribution, except that $u_m = u_j + u_k$. A simple argument by induction shows that any number of interchangeable states can be aggregated into a single macro-state in this manner.

More generally, we may be able to partition the states of M into cells $m_1, \ldots, m_l$ such that, for any $r = 1, \ldots, l$ and any states i and j of M, i and j are interchangeable with respect to each m_k. When this is possible, then $m_1, \ldots, m_l$ are the states of a derived Markov process, which will be ergodic if M is ergodic.

For instance, in a particular market model represented by an ergodic Markov process, we might be able to use a symmetry argument to conclude that all states with the same aggregate demand for a particular good are interchangeable. If so, we can aggregate all states with the same total excess demand for this good into a single macro-state, and the resulting system will be an ergodic Markov process with a stationary distribution. In general this Markov process will have many fewer states, but still far to many to permit an analytical derivation of the stationary distribution.

4 The Structure of Finite Aperiodic Markov Processes

If a Markov process is finite and aperiodic but is not irreducible, its states can be partitioned into subsets S^{tr}, $S_1, \ldots S_k$, where every state $s \in S^{\mathrm{tr}}$ is *transient*, meaning that for any realization $\{s^t\}$ of the Markov process, with probability one there is a time t such that $s \neq s^{t+t'}$ for all $t' = 1, 2, \ldots$; i.e., s does not reappear in $\{s^t\}$ after time t. It follows that also with probability 1 there is a time t such that no member of S^{tr} appears after time t. A non-transient state is called *recurrent*, for it reappears infinitely often with probability one in a realization of the Markov process.

If s_i is recurrent and communicates with s_j, then s_j is itself recurrent and communicates with s_i. For if j does not communicate with s_i, then every time s_i appears, there is a strictly positive probability, say $q > 0$ that it will never reappear. The probability that s_i appears k times is thus at most $(1-q)^k$, so s_i reappears an infinite number of times with probability zero, and hence is not recurrent. If s_j communicates with s_i, then s_j must be recurrent, as can be proved using a similar argument.

It follows that communication of states is an equivalence relation over the recurrent states of the Markov process. We define $S_1, \ldots S_k$ to be the equivalence classes of the recurrent states of the Markov process with respect to this equivalence relation. It is clear that the restriction of Markov process to any one of the S_r, $r = 1, \ldots, k$ is an ergodic Markov process with a stationary distribution. Moreover, if $s_i \in S^{\mathrm{tr}}$, there is a probability distribution q^i over $\{1, \ldots, k\}$ such that q^i_r is the probability, starting in s_i, the Markov process will eventually enter S_r, from which it will of course never leave. Thus for an arbitrary finite, aperiodic Markov process with transition matrix $P = \{p_{ij}\}$, we have the following.

Extended Ergodic Theorem: Let M be a finite aperiodic Markov process. There exists a unique partition $\{S^{\mathrm{tr}}, S_1, \ldots, S_k\}$ of the states S of M, a probability distribution u^r over S_r for $r = 1, \ldots, k$, such that $u^r_i > 0$ for all $i \in S_r$, and for each $i \in S^{\mathrm{tr}}$, there is a probability distribution q^i over $\{1, \ldots, k\}$

such that for all $i, j = 1, \ldots, n$ and all $r = 1, \ldots k$, we have

$$u_{ij} = \lim_{t \to \infty} P_{ij}^{(t)}; \tag{13}$$

$$u_j^r = u_{ij} \quad \text{if } i, j \in S_r; \tag{14}$$

$$u_j^r = \sum_{i \in S_r} u_i^r p_{ij} \quad \text{for } j \in S_r; \tag{15}$$

$$u_{ij} = q_r^i u_j^r \quad \text{if } s_i \in S^{\text{tr}} \text{ and } s_j \in S_r. \tag{16}$$

$$u_{ij} = 0 \quad \text{if } s_j \in S^{\text{tr}}. \tag{17}$$

$$\sum_j u_{ij} = 1 \quad \text{for all } i = 1, \ldots, n. \tag{18}$$

For a Markov process with few states, there are well-known methods for solving for the stationary distribution (Gintis 2009, Ch. 13). However, for systems with a large number of states, these methods are impractical. Rather, we here create a computer model of the Markov process, and ascertain empirically the dynamical properties of the irreducible Markov subprocesses. We are in fact only interested in measuring certain aggregate properties of the subprocess rather than their stationary distributions. These properties are the long-run average price and quantity structure of the economy, as well as the short-run volatility of prices and quantities and the efficiency of the process's search and trade algorithms. It is clear from the quasi-ergodic theorem that the long-term behavior of an any realization of aperiodic Markov process is governed by the stationary distribution of one or another of the stationary distributions of the irreducible subprocesses $S_1, \ldots, S_k$. Generating a sufficient number of the sample paths $\{s^t\}$, each observed from the point at which the process has entered some S_r, will reveal the long-run behavior of the dynamical system.

Suppose an aperiodic Markov process M with transient states S^{tr} and ergodic subprocesses $S_1, \ldots, S_k$ enters a subprocess S_r after t_0 periods with high probability, and suppose the historical average over states from t_0 to t_1 is a close approximation to the stationary distribution of S_r. Consider the Markov process M^+ consisting of reinitializing M every t_1 periods. Then M^+ is ergodic, and a sufficiently large sample of historical averages starting t_0 periods after reinitialization and continuing until the next initialization will reveal the stationary distribution of M^+. This is the methodology we will used in estimating the aggregate properties of a Markov model of a market economy.

5 Scarfian Instability Revisited

To assess the effect of passing from differential equation to Markov process models, this section revisits Herbert Scarf's seminal example of Walrasian instability. Scarf's is a three-good economy in which each agent produces one good and consumes some of this good plus some of one other good, in fixed proportions. Labeling the goods X, Y, and Z, following Scarf, we assume X-producers consume X and Y, Y-producers consume Y and Z, and Z-producers consume Z and X, where the conditions of production are identical for all three goods. The utility functions for the three agents are assumed to be

$$u_X(x, y, z) = \min\{x, y\}, \tag{19}$$
$$u_Y(x, y, z) = \min\{y, z\}, \tag{20}$$
$$u_Z(x, y, z) = \min\{z, x\}. \tag{21}$$

It is straightforward to show that utility maximization, where p_x, p_y, and p_z are the prices of the three goods and x^d, y^d, and z^d are the final demands for the three goods, gives

$$x_X^d = y_X^d = \frac{p_x}{p_x + p_y} \tag{22}$$
$$y_Y^d = z_Y^d = \frac{p_y}{p_y + p_z} \tag{23}$$
$$z_Z^d = x_Z^d = \frac{p_z}{p_z + p_x} \tag{24}$$

These equations allows us to calculate total excess demand for each good as a function of the prices of the three goods. It is easy to check that the market-clearing prices, normalizing $p_z^* = 1$, are given by $p_x^* = p_y^* = 1$. The excess demand functions for the economy are then given by

$$E_x = x_X^d + x_Z^d - 1 = \frac{-p_y}{p_x + p_y} + \frac{p_z}{p_z + p_x} \tag{25}$$
$$E_y = y_X^d + y_Y^d - 1 = \frac{-p_z}{p_y + p_z} + \frac{p_x}{p_x + p_y} \tag{26}$$
$$E_z = z_Y^d + z_Z^d - 1 = \frac{-p_x}{p_z + p_x} + \frac{p_y}{p_y + p_z}. \tag{27}$$

The tâtonnement price adjustment process is given by

$$\dot{p}_i = E_i(p_x, p_y, p_z), \quad \text{where } i = x, y, z. \tag{28}$$

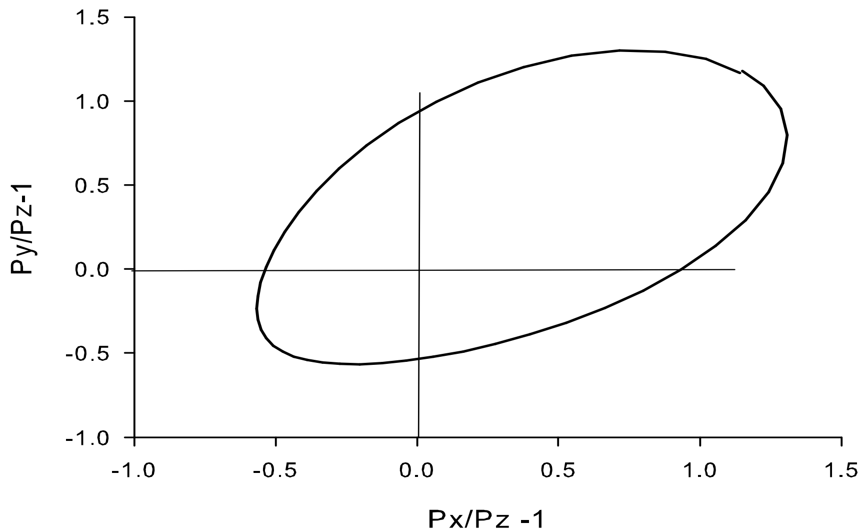

Figure 1: Neutral stability of three-good Scarf Economy, modeled as a Markov process with public prices and tâtonnement price adjustment.

It is easy to show that the expression $p_x p_y p_z$ is constant on paths of the dynamical system, which implies that the equilibrium is neutrally stable, the system moving in closed paths about the equilibrium at every non-equilibrium point.

Before moving to the non-tâtonnement version of the Scarf economy, we will implement Scarf's differential equation solution as a Markov process in which time t becomes discrete, $t = 1, 2, \ldots$, and the differential equations (28) are replaced by difference equations

$$p_x^{t+1} = p_x^t + E_x^t/\Delta \tag{29}$$
$$p_y^{t+1} = p_y^t + E_y^t/\Delta, \tag{30}$$
$$p_z^{t+1} = p_z^t + E_z^t/\Delta, \tag{31}$$

where prices are restricted to a bounded interval of rational numbers, and Δ is an integer chosen so that period-to-period price changes are small. Note that Δ affects the speed of adjustment of the system, but not the path of adjustment, provided it is not so small as to lead the system to violate the price bounds.

This system is a deterministic Markov process in which the current state is the vector of current prices. Suppose we start with disequilibrium prices $p_x = p_x^* + \delta_x$, $p_y = p_y^* + \delta_y$, $p_z = p_z^* = 1$, and in each period prices are updated according to the tâtonnement equations (29). The resulting path of deviations of prices p_x and p_y from equilibrium, with $p_z = 1$ as numeraire, after 5200 periods with $\Delta = 100$, $\delta_X = 3$, and $\delta_Y = -2$ is shown in Figure 1, and perfectly replicates the analytical results of Scarf (1960).

6 The Scarf Economy Without Tâtonnement

For the Markov process version of the Scarf economy without tâtonnement, we maintain the above assumptions, except now we assume 1000 traders of each of the three types, each trader endowed at time $t = 0$ with a set of private prices randomly drawn from a uniform distribution. We allow 50000 generations and 10 periods per generation. At the start of each period, each agent's inventory is re-initialized to one unit of his production good and zero units of the other goods. Each agent in turn is then designated a trade initiator and is paired with a randomly chosen responder, who can either accept or reject the proposed trade. Each agent is thus an initiator exactly once and responder on average once per period. After a successful trade, agents consume whatever is feasible from their updated inventory.

In the reproduction stage, which occurs every ten periods, 5% of agents are randomly chosen either to copy a more successful agent or to be copied by a less successful agent, where success is measured by total undiscounted utility of consumption over the previous ten periods. Such an agent is chosen randomly and assigned a randomly chosen partner with the same production and consumption parameters. The less successful of the pair then copies the private prices of the more successful. In addition, after the reproduction stage, each price of each agent is mutated with 1% probability, the new price either increasing or decreasing by 10%.

The trade procedure is as follows. The initiator offers a certain quantity of one good in exchange for a certain quantity of a second good. If the responder has some of the second good, and if the value of what get gets exceeds the value of what he gives up, according to his private prices, then he agrees to trade. If he has less of the second good than the initiator wants, the trade is scaled down proportionally. Traders are thus rational maximizers, where their subjective priors are their vectors of private prices, and each is ignorant of the other's subjective prior.

Which good he offers to trade for which other good is determined as follows. Let us call an agent's production good his P-good, the additional good he consumes his C-good, and the good which he neither produces nor consumes he T-good. Note that agents must be willing to acquire their T-good despite the fact that it does not enter their utility function. This is because X-producers want Y, but Y-producers do not want X. Only Z-producers want X. Since a similar situation holds with Y-producers and Z-producers, consumption ultimately depends on at least one type of producer accepting the T-good in trade, and then using the T-good to purchase their C-good.

If the initiator has his T-good in inventory, he offers to trade this for his C-good. If this offer is rejected, he offers to trade his T-good for his P-good. which will be a net gain in the value of his inventory provided his subjective terms of trade are favorable. If the initiator does not have his T-good but has his P-good, he offers this in trade for his C-good. If this is rejected, he offers to trade half his P-good for his T-good. If the trade initiator had neither his T-good nor his P-good, he offers his C-good in trade for his P-good, and if this fails he offers to trade for his T-good. In all cases, when a trade is carried out, the term are dictated by the initiator and the amount is the maximum compatible with the inventories of the initiator and responder.

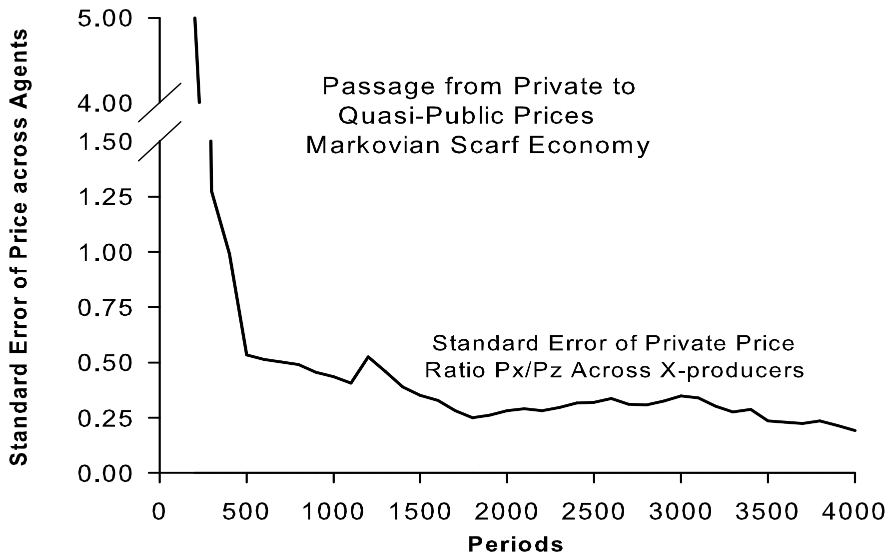

Figure 2: The Markov process version of the Scarf economy initialized with random prices quickly transitions to quasi-public prices (private prices with low standard error across traders).

Figure 2 shows that within a relatively few periods, the randomly initialized private prices move to *quasi-public* prices, in which the standard error of prices for the same good across individuals is relatively small. Quasi-public prices are the closest the Markov process comes to approximating the public prices of standard Walrasian general equilibrium theory.

The Markov dynamic in this case is a stationary distribution depicted in Figure 3. It is clear that by the time quasi-public prices have become established, the Markov process has attained its stationary distribution, which is a cycle around the equilibrium. This is the only behavior of the stationary distribution observed, independent of the initial state of the system, so it is the stationary distribution version of a limit cycle. In all observed cases, the stationary distribution has approximately the same period and amplitude.

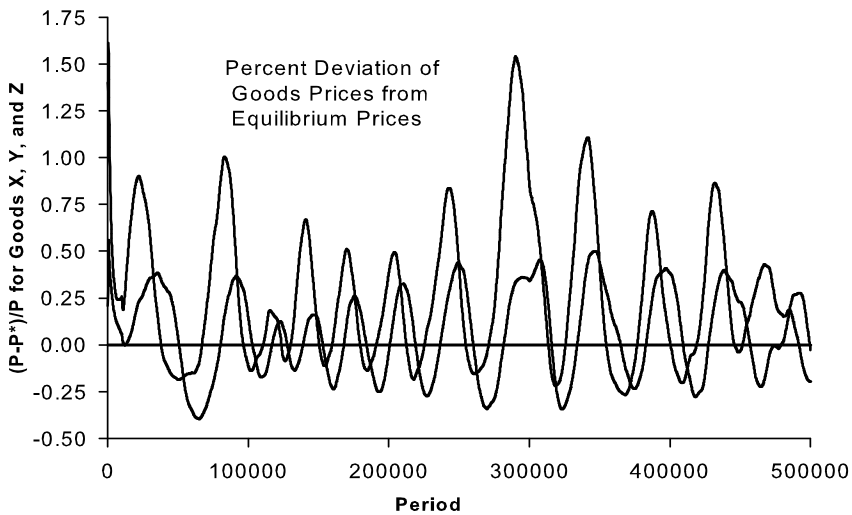

Figure 3: The Markov process Scarf economy with private prices exhibits a stationary distribution akin to a limit cycle. The average excursion from equilibrium and the approximate period of the stationary distribution are independent from initial conditions.

In sum, we have developed a dynamic mathematical model of Scarfian exchange in the form of a Markov process. The transition probabilities of the Markov process are specified implicitly by the algorithms for agent pairing, trading, updating, and reproduction. Except for the trade algorithm, for which alternative algorithms are plausible, all modeling choices are uniquely determined by the standard conception of the Walrasian general equilibrium model.

The equilibrium of the Markov process is a stationary distribution that can be analytically specified in principle, but in practice is orders of magnitude too large to calculate, even with the fastest and most powerful conceivable computational aids. Thus, as in the natural sciences, we are obliged to investigate the stationary distribution by running the process on a computer with a variety of choices of numerical parameters. Such calculations are subject to statistical error, but our results are so robust that we are virtually certain to have captured the characteristics of system equilibrium almost perfectly. The most caution conclusion we can offer is that we have proven, for the first time, that there exist Scarfian economies that without tâtonnement that exhibit limit cycles around the equilibrium price vector, with private prices and a completely decentralized dynamical system of price adjustment.

7 A Multi-Good Market Economy with Simple Production

Despite its considerable historical value, the Scarf economy's extreme assumptions are uncharacteristic of a normal market economy, and as we shall see, it is the extreme assumptions that account for the Scarfian economies lack of stability. The remainder of this paper will deal with a canonical case of a pure market economy with only one institution—the marketplace, but with highly heterogeneous agents. We assume each agent produces one good, in fixed amount, using only personal labor, but consumes a variety of goods produced by others. Agents are endowed with subjective priors concerning the value of all goods, which we call their *private prices*. They have no information about the economy other than that gathered from private experience in trade, including periodically discovering the private price vector of another agent they have encountered and copying it, with some possible mutation, if that agent appears to be more successful than himself. The only serious design decision is that of the trade algorithm which, while much more straightforward than in the case of the Scarf economy, is still in principle somewhat underspecified by the logic of Walrasian exchange. Happily, the details of the trade protocol do not affect the dynamical movement to Walrasian equilibrium as far as we can ascertain.

We assume there are n sectors. Sector $k = 1, \dots, n$ produces good k in "styles" $s = 1, \dots, m$ (we use "styles" to enrich the heterogeneity of goods in the model without seriously increasing the computational resources needed to estimate the stationary distribution of the resulting Markov process). Each agent consumes a subset of non-production goods, but only a single style of any good. In effect, then, there are nm distinct goods g_s^k, but only n production processes and correspondingly n equilibrium prices, since goods g_s^k and g_t^k with styles s and t respectively, have the same production costs and hence the same price in equilibrium. We write the set of goods as $G = \{g_s^k | k = 1, \dots, n, s = 1, \dots m\}$. We also write $g = g^k$ when $g = g_s^k$ for some style s.

A producer of good g_s^k, termed a g_s^k*-agent*, produces with no inputs other than personal labor an amount q_k of good g_s^k which depreciates to zero at the end of a trading period. In a non-monetary economy, only the production good is carried in inventory, but when individuals are permitted to acquire non-consumption goods, as in later sections of the paper, a trade inventory includes all goods that are not the agent's consumption goods.

The Markov process is initialized by creating N agents, each of whom is randomly assigned a production good g_s^k. Thus, in an economy with goods in m styles, there are Nnm traders. Each of these traders is assigned a private price vector by choosing each price from a uniform distribution on (0, 1), then normalizing so that the price of the n^{th} good is unity. Each g_s^k-agent is then randomly assigned a set $H \subseteq G$, $g_s^k \notin H$ of consumption at most one style of a given good.

The utility function of each agent is the product of powers of CES utility functions of the following form. Suppose an agent consumes r goods. We partition the r goods into k segments (k is chosen randomly from $1 \ldots r/2$) of randomly chosen sizes $m_1, \ldots, m_k$, $m_j > 1$ for all j, and $\sum_j m_j = n$. We randomly assign goods to the various segments, and for each segment, we generate a CES consumption with random weights and an elasticity randomly drawn from the uniform distribution on an interval $[\epsilon_*, \epsilon^*]$. Total utility is the product of the k CES utility functions to random powers f_j such that $\sum_j f_j = 1$. In effect, no two agent have the same utility function.

For example, consider a segment using goods $x_1, \ldots, x_m$ with prices $p_1, \ldots, p_m$ and (constant) elasticity of substitution s, and suppose the power of this segment in the overall utility function is f. It is straightforward to show that the agent spends a fraction f of his income M on goods in this segment, whatever prices he faces. The utility function associated with this segment is then

$$u(x_1, \ldots, x_n) = \left(\sum_{l=1}^{m} \alpha_l x_l^{\gamma} \right)^{1/\gamma}, \tag{32}$$

where $\gamma = (s-1)/s$, and $\alpha_1, \ldots, \alpha_m > 0$ satisfy $\sum_l \alpha_l = 1$. The income constraint is $\sum_{l=1}^{m} p_l x_l = f_i M$. Solving the resulting first order conditions for utility maximization, and assuming $\gamma \neq 0$ (i.e., the utility function segment is not Cobb-Douglas), this gives

$$x_i = \frac{M f_i}{\sum_{l=1}^{m} p_l \phi_{il}^{1/(1-\gamma)}}, \tag{33}$$

where

$$\phi_{il} = \frac{p_i \alpha_l}{p_l \alpha_i} \quad \text{for } i, l = 1, \ldots, m.$$

When $\gamma = 0$ (which occurs with almost zero probability), we have a Cobb-Douglas utility function with exponents α_l, so the solution becomes

$$x_i = \frac{M f_i \alpha_i}{p_i}. \tag{34}$$

By creating such a complex array of utility functions, we ensure that our results are not the result of assuming an excessively narrow set of consumer characteristics. However, the high degree of randomness involved in creating a large number of agents ensures that all goods will have approximately the same aggregate demand characteristics. If we add to this that all goods have the same supply characteristics, we can conclude that the market-clearing Walrasian equilibrium will occur when all prices are equal. This in fact turns out to be the case. If we assume heterogeneous production conditions, then we cannot calculate equilibrium prices, but we can still judge that the dynamical system is asymptotically stable by the long-run standard error of the absolute value of excess demand, which will be very small in equilibrium.

For each good $g_s^k \in G$ there is a *market* $m[k,s]$ of traders who sell good g_s^k. In each period, the traders in the economy are randomly ordered and are permitted one-by-one to engage in active trading. When the g_t^h-agent A is the current active trader, for each good g_t^h for which A has positive demand (i.e, $x_h^{A*} > 0$), A is assigned a random member B $\in m[h,t]$ who consumes g_s^k. A then offers B the maximum quantity y_k of g_s^k, subject to the constraints $y_k \leq \mathbf{i}_k^A$, where $\mathbf{i}_k^A$ represents A's current inventory of good g_s^k, and $y_k \leq p_h^A x_h^A / p_k^A$, where x_h^A is A's current demand for g_t^h. This means that if A's offer is accepted, A will receive in value at least as much as he gives up, according to A's private prices. A then offers to exchange y_k for an amount $y_h = p_k^A y_k / p_h^A$ of good g_t^h; i.e., he offers B an equivalent value of good g_t^h, the valuation being at A's prices. B accepts this offer provided the exchange is weakly profitable at B's private prices; i.e., provided $p_k^B y_k \geq p_h^B y_h$. However, B adjusts the amount of each good traded downward if necessary, while preserving their ratio, to ensure that what he receives does not exceed his demand, and what he gives is compatible with his inventory of g_t^h. If A fails to trade with this agent, he still might secure a trade giving him g_s^k, because A $\in m[k,s]$ may also be on the receiving-end of trade offers from g_t^h-agents at some point during the period. If a g_s^k-agent exhausts his supply of g_s^k, he leaves the market for the remainder of the period.

The assumption that each trading encounter is between agents each of whom produces a good that the other consumes could be replaced by the assumption is that each g_s^k-producer A can locate the producers of his consumption goods, but that finding such a producer who also consumes g_s^k will require a separate search. We simply collapse these two stages, noting that when a second search is required and its outcome costly or subject to failure, the relative inefficiency of the non-monetary economy, by comparison with the monetary economies described below, is magnified. Note, however, that while A's partner is a consumer of g_s^k, he may have fulfilled his demand for g_s^k for this period by the time A makes his offer, in which case no trade will take place.

The trade algorithm involves only one substantive design choice, that of allowing A to make a single take-it-or-leave-it relative price offer, while obliging A to accept quantity terms that are set by B, when it is feasible to do so. Such alternatives as allowing B to make the take-it-or-leave-it offer, and choosing the mean of the two offers provided that each is acceptable to the other, or using a Nash bargaining solution, do not alter the market dynamics.

After each trading period, agents consume their inventories provided they have a positive amount of each good that they consume, and agents replenish the amount of their production good in inventory. Moreover, each trader updates his private price vector on the basis of his trading experience over the period, raising the price of a consumption or production good by 0.05% if his inventory is empty (i.e., if he failed to purchase any of the consumption good or sell all of his production good), and lowering price by 0.05% otherwise (i.e., if he succeeded in obtaining his consumption good or sold all his production inventory). We allow this adjustment strategy to evolve endogenously according to an imitation processes.

After a number of trading periods, the population of traders is updated using the following process. For each market $m[k,s]$ and for each g_s^k-trader A, let f^A be the accumulated utility of agent A since the last updating period (or since the most recent initialization of the Markov process if this is the first updating period). Let f_* and f^* be the minimum and maximum, respectively, over f^A for all g_s^k-agents A. For each g_s^k-agent A, let $p^A = (f^A - f_*)/(f^* - f_*)$, so p^A is a probability for each A. If r agents are to be updated, we repeat the following process by r times. First, choose an agent for reproducing as follows. Identify a random agent in $m[k,s]$ and choose this agent for reproduction with probability p^A. If A is not chosen, repeat the process until one agent is eventually chosen. Note that a relatively successful trader is more likely to be chosen to reproduce than an unsuccessful trader. Next, choose an agent B to copy A's private prices as follows. Identify a random agent B in $m[k,s]$ and choose this agent with probability $1-p^B$. If B is not chosen, this process is repeated until B is chosen. Clearly, a less successful trader is likely to be chosen this criterion. Repeat until an agent B is chosen. Finally, endow B with A's private price vector, except for each such price, with a small probability $\mu =$ randomly increase or decrease its value by a small percentage ϵ. The resulting updating process is a discrete approximation of a monotonic dynamic in evolutionary game theory, and in differential equation systems, all monotonic dynamics have the same dynamical properties (Taylor and Jonker 1978, Samuelson and Zhang 1992). Other monotonic approximations, including the simplest, which is repeatedly to choose a pair of agents in $m[k,s]$ and let the lower-scoring agent copy the higher-scoring agent, produce similar dynamical results.

Using utility as the imitation criterion is quite noisy, because utility functions are heterogeneous and individuals who prefer goods with low prices do better than agents who prefer high-priced goods of independent of the trading prowess. Using alternative criteria, such as the frequency and/or volume trading success, with results similar to those reported herein.

The result of the dynamic specified by the above conditions is the change over time in the distribution of private prices. The general result is that the system of private prices, which at the outset are randomly generated, in rather short time evolves to a set of *quasi-public* prices with very low inter-agent variance. Over the long term, these quasi-public prices move toward their equilibrium, market-clearing levels.

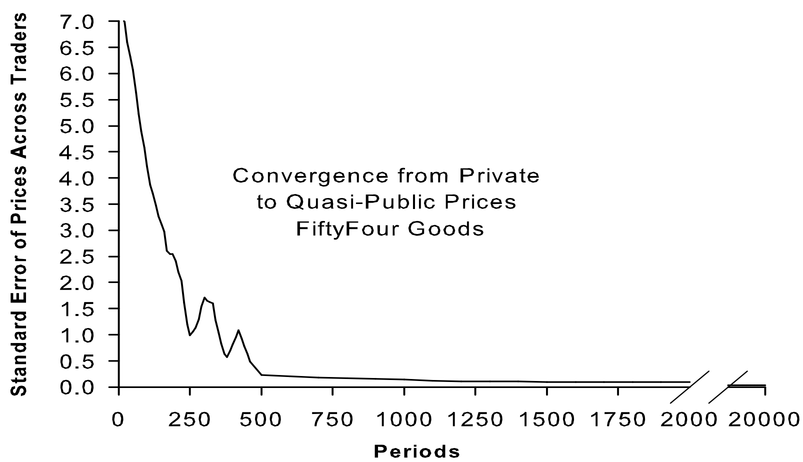

Figure 4: Convergence of private prices to quasi-public prices in a typical run with nine goods in 6 styles each (fifty-four goods).

8 Estimating the Stationary Distribution

I will illustrate this dynamic assuming $n = 9$, $m = 6$, and $N = 300$, so there are fifty-four distinct goods which we write as $g_1^1, \ldots, g_6^9$, and 16200 traders in the economy. There are then nine distinct prices $p_1^A, \ldots p_9^A$ for each agent A. We treat g^9 as the numeraire good for each trader, so $p_9^A = 1$ for all traders A. A g^k-agent produces one unit of good k per period. We assume that there are equal numbers of producers of each good from the outset, although we allow migration from less profitable to more profitable sectors, so in the long run profit rates are close to equal in all sectors. The complexity of the utility functions do not allow us to calculate equilibrium properties of the system perfectly, but we will assume that market-clearing prices are approximately equal to unit costs, given that unit costs are fixed, agents can migrate from less to more profitable sectors, and utility functions do not favor one good or style over another, on average. Population updating occurs every ten periods, and the number of encounters per sector is 10% of the number of agents in the sector. The mutation rate is $\mu = 0.01$ and the error correction is $\epsilon = 0.01$.

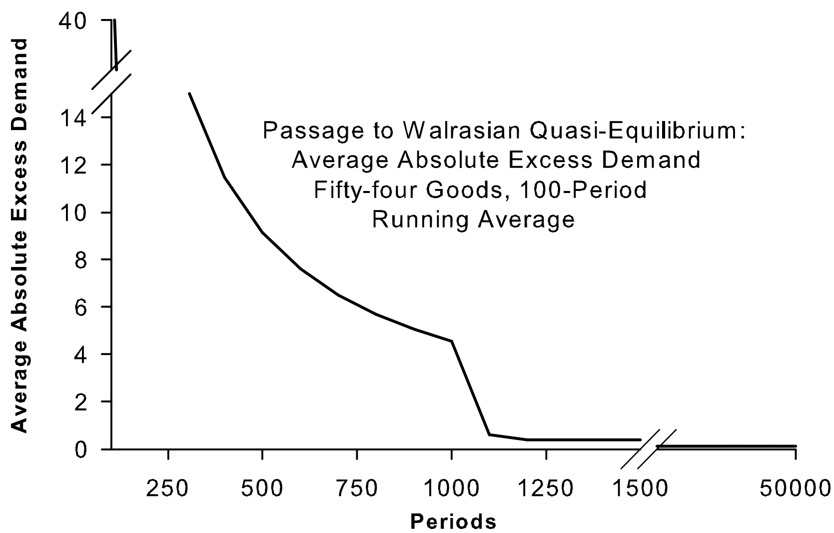

Figure 5: The path of aggregate excess demand over 50,000 periods.

The results of a typical run of this model is illustrated in Figures 4 to 6. Figure 4 shows the passage from private to quasi-public prices over the first 20,000 trading periods of a typical run. The mean standard error of prices is computed as follows. For each good g we measure the standard deviation of the price of g across all g-agents, where for each agent, the price of the numeraire good g_9 unity. Figure 4 shows the average of the standard errors for all goods. The passage from private to quasi-public prices is quite dramatic, the standard error of prices across individuals falling by an order of magnitude within 300 periods, and falling another order of magnitude over the next 8500 periods. The final value of this standard error is 0.029, as compared with it initial value of 6.7.

Figure 8 shows the movement of the average standard error of the absolute value of excess demand over 50,000 periods for nine goods in 6 styles each. Using this measure, after 1500 periods excess demand has decrease by two orders of magnitude, and it decreases another order of magnitude by the end of the run.

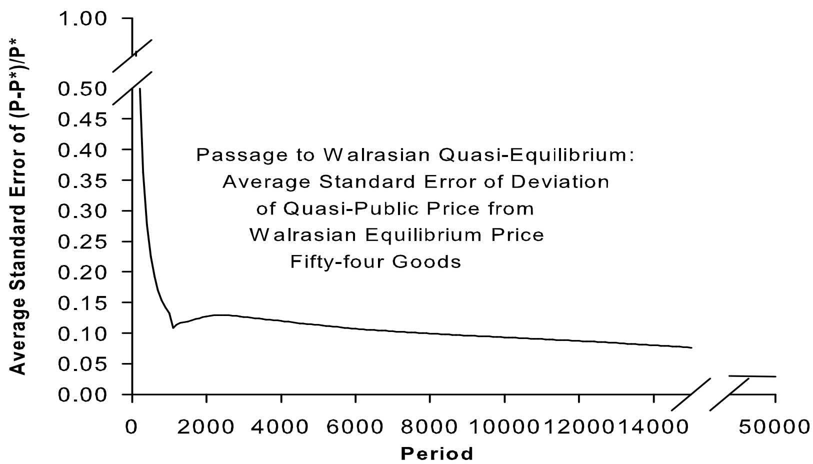

Figure 6: The passage from private to quasi-public prices in the Markovian Scarf economy. The long-run standard error of prices across traders is rather high, due to the fact that the system does not tend to Walrasian equilibrium.

It is not surprising, given the behavior of excess demand for this model, that prices would approach their Walrasian equilibrium values. This process is illustrated in Figure 6. After 50,000 periods, the standard error of the deviation of prices from (our calculated) equilibrium values are about 3% of its starting value.

The distinction between low-variance private prices and true public prices is significant, even when the standard error of prices across agents is extremely small, because stochastic events such as technical changes propagate very slowly when prices are highly correlated private prices, but very rapidly when all agents react in parallel to price movement. In effect, with private prices, a large part of the reaction to a shock is a temporary reduction in the correlation among prices, a reaction that is impossible with public prices, as the latter are always perfectly correlated.

There is nothing special about the parameters used in the above example. Of course adding more goods or styles increases the length of time until quasi-public prices become established, as well as the length of time until market quasi-equilibrium is attained. Increasing the number of agents increases the length of both of these time intervals.

9 The Emergence of Money

There is no role for money in the Walrasian general equilibrium model because all adjustments of ownership are carried out simultaneously when the equilibrium prices are finally set. When there is actual exchange among individual agents in an economy, two major conditions give rise to the demand for money, by which we mean a good that is accepted in exchange not for consumption or production, but rather for resale at a later date against other intrinsically desired goods. The first is the failure of the "double coincidence of wants," (Jevons 1875), explored in recent years in this is and other journals by Starr (1972) and Kiyotaki and Wright (1989,1991,1993). The second condition the existence of transactions costs in exchange, the money good being the lowest transactions-cost at good (Foley 1970, Hahn 1971, Hahn 1973, Kurz 1974b, Kurz 1974a, Ostroy 1973, Ostroy and Starr 1974, Starrett 1974). We show that these conditions interact in giving rise to a monetary economy. When one traded good has very low transactions costs relative to other goods, this good may come to be widely accepted in trade even by agents who do not consume or produce it. Moreover, when an article that is neither produced nor consumed can be traded very low transactions costs, this good, so-called *fiat* money. will emerge as a universal medium of exchange.

We now permit traders to buy and sell at will any good that they neither consumer nor produce. We call such a good a *money good*, and if there is a high frequency of trade in one or more money goods, we say the market economy is a *money economy*. We assume that traders accept all styles of a money good indifferently. We first investigate the emergence of money from market exchange by assuming zero inventory costs, so the sole value of money is to facilitate trade between agents even though the direct exchange of consumption and production goods between a pair of agents might fail because one of the parties is not currently interested in buying the other's production good. The trade algorithm in case agents accept a good that they do not consume is as follows. At the beginning of each period, each agent calculates how much of each consumption good he want to acquire during that period, as follows. The agent calculate the market value of his inventory of production and money goods he holds in inventory, valued at his private prices. This total is the agent's income constraint. The agent then choose an amount of each consumption good to purchase by maximizing utility subject to this income constraint. The trade algorithm is similar to the case of pure market, except that either party to a trade may choose to offer and/or accept a money good in the place of his production good.

We evaluate the performance of this economy using the same parameters as in our previous model, including zero inventory costs. Figure 7 shows that the use of money increases monotonically over the first 2000 periods, spread almost equally

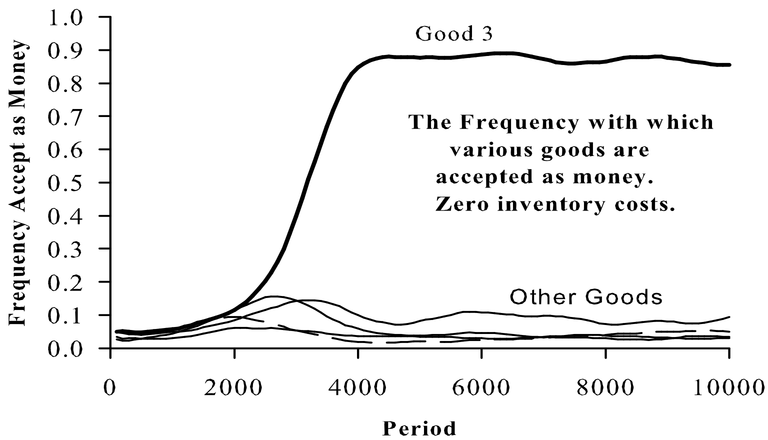

Figure 7: The Emergence of Money in a Market Economy. The parameters of the model are the same as in the baseline case treated previously. Inventory costs are assume absent.

among the remaining goods. From period 2000 to period 4000, one good becomes a virtually universal currency, driving the use of the others to low levels. It is purely random which good becomes the universal medium of exchange, but one does invariable emerge as such after several thousand periods. If we add inventory costs with g^1 being lower cost than the others, g^1 invariably emerges as the medium of exchange after 1000 periods, and the other goods are not used as money at all. I did not include graphs of the passage to quasi-public prices or other aspects of market dynamics because they differ little from the baseline economy described above.

As in traditional monetary theory (Menger 1892, Wicksell 1911, Kiyotaki and Wright 1989,1991) money emerges from goods trade both because it is a low transactions cost good and it solves the problem of the "double coincidence of wants" that is required for market exchange (Jevons 1875). The relative efficiency of money over direct goods trade increases with the number of goods, as illustrated in Figure 8. While with six goods and one style the relative efficiency of money is only 150%, for nine goods and twenty styles (180 goods), the relative efficiency is 1200%.

10 The Resilience of the Decentralized Market Economy

We now show that the above Markov model is extremely resilient in the face of aggregate shocks when the number of producers per good is sufficiently large, but becomes unstable when this number fall below a certain (relatively high) threshold. We illustrate this in the context of a *fiat* money economy. We take the non-monetary

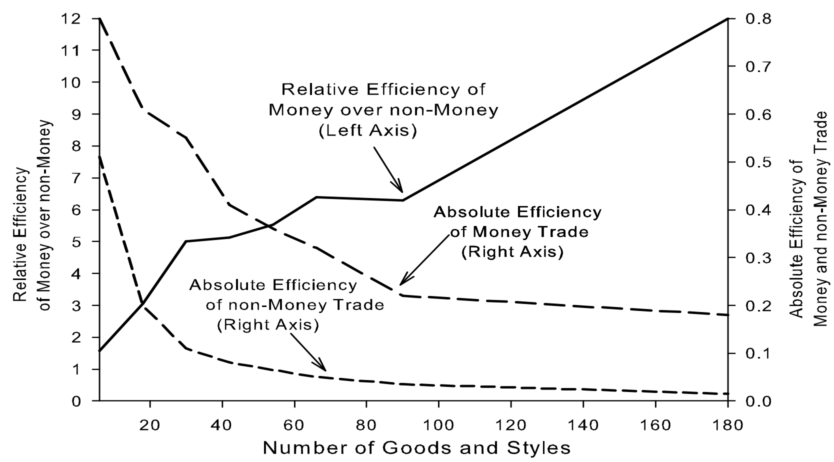

Figure 8: The Relative Efficiency of Money in a Market Economy

economy described above, with nine goods in six styles each, and add a single new good that is neither produced nor consumed, and has zero inventory storage costs. When such a good is available, it quickly becomes universal medium of exchange for the economy, accepted by almost 100% of market traders. The nature of market dynamics in a *fiat* money economy is not noticeably different from the economies described above.

In this *fiat* money economy with 300 producers per good, every 1000 periods, we impose an aggregate shock on the economy consisting of a reduction in the *fiat* money holdings of each trader to 20% of its normal level, as shown in figure 9. The reduced holding are maintained for 100 periods, after which the money holdings of each trader is multiplied by two, restoring the money stock for the economy to its initial level. Figure 10 shows the effect of the period shock on average prices and on the standard error of prices across traders; there is no noticeable effect. Figure 11 shows that the quantity side of the economy is also virtually unaffected by the system of aggregate shocks.

11 Conclusion

The search for stability of market exchange using differential equations with public prices, while producing some brilliant mathematical analyses, was doubly defective. First, no plausible dynamic price and quantity adjustment mechanism was found. Second, such a dynamic, even were it found, would be of doubtful value

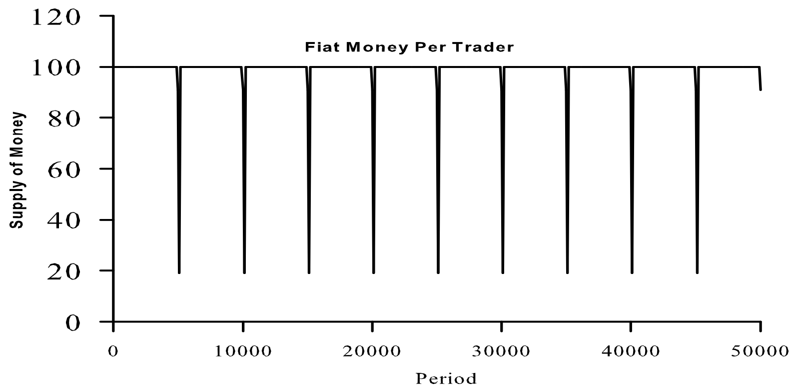

Figure 9: To test the resilience of the Markov economy, we impose a periodic shock, sustained for 100 periods, that reduces the money supply to 20% of its normal level.

because out of equilibrium public prices cannot exist in a decentralized market economy. By modeling market exchange as a Markov process, we have shown that under plausible conditions we get convergence to a quasi-equilibrium. However, under the extreme conditions of the Scarf economy (only three goods, three traders, and fixed coefficients utility functions) we get a stationary distribution akin to a limit cycle in continuous models. With a plain-vanilla Walrasian economy with individual production, we find global stability under a wide range of parameters. Yet we do not have the analytical machinery to ascertain when Scarf instability will hold, when global stability holds, or whether there are other possible dynamic characteristics of a Markovian market system.

On the positive side, we have proved that certain Markov models of market dynamics are globally stable over a wide range of parameters. This is an existence theorem possibly as informative as the existence proofs for general equilibrium. Computational proofs, however, are not as powerful as purely analytical proofs. Those unused to working with complex dynamical systems may object that a computational proof is no proof at all. In fact, a computational proof may not be a mathematical proof, but it is a scientific proof: it is evidential rather than tautological, and depends on induction rather than deduction. The natural sciences, in which complex systems abound, routinely use mathematical to models that admit no closed-form analytical solutions, ascertain their properties through approxima-

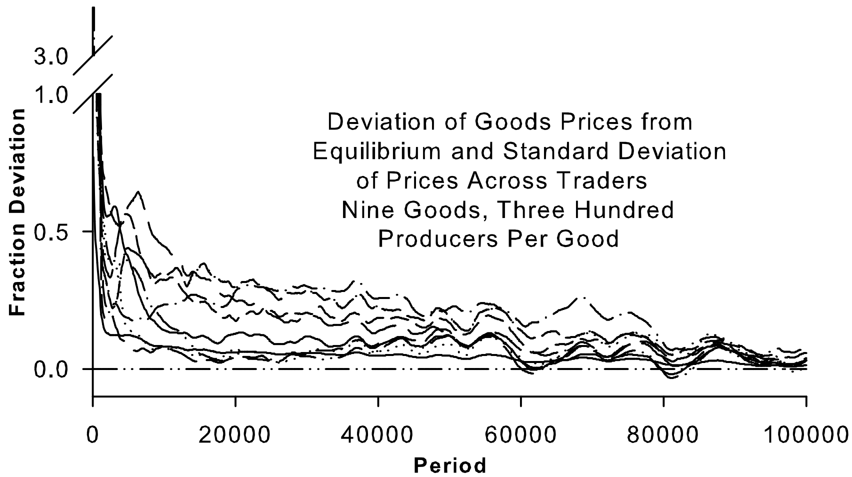

Figure 10: Walrasian equilibrium is resilient to aggregate shocks. With 300 producers per good, we impose an aggregate shock on the Markov process consisting of halving the money supply every 1000 periods, and restoring the money supply after 100 periods have elapsed with the smaller money supply. There is virtually no effect on the passage to a quasi-Walrasian equilibrium.

tion and simulation, and justify these models by virtue of how they conform empirical reality. This appears to be the current state of affairs with respect to Markov processes and general equilibrium theory.

REFERENCES

Arrow, Kenneth and Leonid Hurwicz, "On the Stability of the Competitive Equilibrium, I," *Econometrica* 26 (1958):522–552.

— and — , "Competitive Stability under Weak Gross Substitutability: Nonleinear Price Adjustment and Adaptive Expectations," Technical Report, Office of Naval Research Contract Nonr-255, Department of Economics, Stanford University (1959). Technical Report No. 78.

— and — , "Some Remarks on the Equilibria of Economics Systems," *Econometrica* 28 (1960):640–646.

— , H. D. Block, and Leonid Hurwicz, "On the Stability of the Competitive Equilibrium, II," *Econometrica* 27 (1959):82–109.

Arrow, Kenneth J., "An Extension of the Basic Theorems of Classical Welfare Eco-

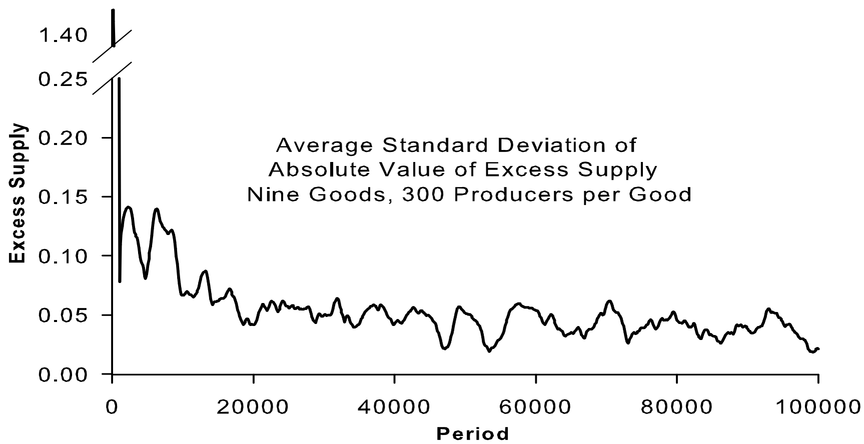

Figure 11: Excess demand is resilient in the face of large macro-level shocks. The parameters are as in the previous figure. Note that there is virtually no effect on aggregate excess demand in any sector of the economy.

nomics," in J. Neyman (ed.) *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (Berkeley: University of California Press, 1951) pp. 507–532.

— and Gérard Debreu, "Existence of an Equilibrium for a Competitive Economy," *Econometrica* 22,3 (1954):265–290.

Aumann, Robert J. and Adam Brandenburger, "Epistemic Conditions for Nash Equilibrium," *Econometrica* 65,5 (September 1995):1161–1180.

Bala, V. and M. Majumdar, "Chaotic Tatonnement," *Economic Theory* 2 (1992):437–445.

Debreu, Gerard, "The Coefficient of Resource Utilization," *Econometrica* 19 (1951):273–292.

Debreu, Gérard, "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences* 38 (1952):886–893.

Debreu, Gerard, "Valuation Equilibrium and Pareto Optimum," *Proceedings of the National Academy of Sciences* 40 (1954):588–592.

— , "Excess Demand Function," *Journal of Mathematical Economics* 1 (1974):15–23.

Feller, William, *An Introduction to Probability Theory and Its Applications* Vol. 1 (New York: John Wiley & Sons, 1950).

Fisher, Franklin M., *Disequilibrium Foundations of Equilibrium Economics* (Cambridge: Cambridge University Press, 1983).

— , *Microeconomics: Essays in Theory and Applications* (Cambridge, UK: Cambridge University Press, 1999). Maarten Pieter Schinkel (ed.).

Foley, Duncan, "Economic Equilibria with Costly Marketing," *Journal of Economic Theory* 2 (1970):276–291.

Gale, David, "The Law of Supply and Demand," *Math. Scand.* 30 (1955):155–169.

Gintis, Herbert, "The Dynamics of General Equilibrium," *Economic Journal* 117 (October 2007):1289–1309.

— , *Game Theory Evolving* Second Edition (Princeton: Princeton University Press, 2009).

Hahn, Frank, "A Stable Adjustment Process for a Competitive Economy," *Review of Economic Studies* 29 (1962):62–65.

— , "Equilibrium with Transactions Costs," *Econometrica* 39 (1971):417–439.

— , "On Transactions Costs, Inessential Sequence Economies and Money," *Review of Economic Studies* 40 (1973):449–461.

— and Takashi Negishi, "A Theorem on Non-Tâtonnement Stability," *Econometrica* 30 (1962):463–469.

Hayek, F. A., "The Use of Knowledge in Society," *American Economic Review* 35,4 (September 1945):519–530.

Hofbauer, Josef and Karl Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge: Cambridge University Press, 1998).

Howitt, Peter and Robert Clower, "The Emergence of Economic Organization," *Journal of Economic Behavior and Organization* 41 (2000):55–84.

Hurwicz, Leonid, "Optimality and Informational Efficiency in Resource Allocation Processes," in *Mathematical Methods in Social Sciences* (Stanford, CA: Stanford University Press, 1960) pp. 27–46.

Jevons, William Stanley, *Money and the Mechnanism of Exchange* (London: D. Appleton & Co., 1875).

Kirman, Alan P., "Whom or What does the Representative Individual Represent?," *Journal of Economic Perspectives* 6 (1992):117–136.

Kiyotaki, Nobuhiro and Randall Wright, "On Money as a Medium of Exchange," *Journal of Political Economy* 94,4 (1989):927–954.

— and — , "A Contribution to a Pure Theory of Money," *Journal of Economic Theory* 53,2 (1991):215–235.

— and — , "A Search-Theoretic Approach to Monetary Economics," *American Economic Review* 83,1 (1993):63–77.

Kurz, Mordecai, "Equilibrium in a Finite Sequence of Markets with Transactions Costs," *Econometrica* 42 (1974):1–20.

— , "Equilibrium with Transactions Cost and Money in a Single Market," *Journal of Economic Theory* 7 (1974):418–452.

Mantel, Rolf, "On the Characterization of Aggregate Excess Demand," *Journal of Economic Theory* 7 (1974):348–53.

— , "Homothetic Preferences and Community Excess Demand Functions," *Journal of Economic Theory* 12 (1976):197–201.

McKenzie, L. W., "On the Existence of a General Equilibrium for a Competitive Market," *Econometrica* 28 (1959):54–71.

— , "Stability of Equilibrium and Value of Positive Excess Demand," *Econometrica* 28 (1960):606–617.

Menger, Karl, "On the Origin of Money," *Economic Journal* 2 (1892):239–255.

Morishima, Michio, "A Reconsideration of the Walras-Cassel-Leontief Model of of General Equilibrium," in *Mathematical Methods in Social Sciences* (Stanford, CA: Stanford University Press, 1960) pp. 63–76.

Negishi, Takashi, "Welfare Economics and the Existence of an Equilibriumfor a Competitive Economy," *Metroeconomica* 12 (1960):92–97.

— , "On the Formation of Prices," *International Economic Review* 2 (1961):122–126.

Nikaido, Hukukaine, "On the Classical Multilateral Exchange Problem," *MetroEconomica* 8 (1956):135–145.

— , "Stability of Equilibrium by the Brown-von Neumann Differential Equation," *MetroEconomica* 27 (1959):645–671.

— and Hirofumi Uzawa, "Stability and Nonnegativity in a Walrasian Tâtonnement Process," *International Economic Review* 1 (1960):50–59.

Ostroy, Joseph, "The Informational Efficiency of Monetary Exchange," *American Economic Review* 63 (1973):597–610.

— and Ross Starr, "Money and the Decentralization of Exchange," *Econometrica* 42 (1974):1093–1113.

Saari, Donald G., "Iterative Price Mechanisms," *Econometrica* 53 (1985):1117–1131.

— , "Mathematical Complexity of Simple Economics," *Notices of the American Mathematical Society* 42,2 (February 1995):222–230.

Samuelson, Larry and Jianbo Zhang, "Evolutionary Stability in Asymmetric Games," *Journal of Economic Theory* 57,2 (1992):363–391.

Scarf, Herbert, "Some Examples of Global Instability of Competitive Equilibrium," *International Economic Review* 1 (1960):157–172.

Smith, Adam, *The Theory of Moral Sentiments* (New York: Prometheus, 2000[1759]).

Sonnenschein, Hugo, "Do Walras' Identity and Continuity Characterizethe Class of Community Excess Demand Functions?," *Journal of Ecomonic Theory* 6 (1973):345–354.

Starr, Ross, "The Structure of Exchange in Barter and Monetary Economies," *Quarterly Journal of Economics* 86,2 (May 1972):290–302.

Starrett, David, "Inefficiency and the Demand for 'Money' in a Sequence Economy," *Review of Economic Studies* 40 (1974):437–448.

Taylor, Peter and Leo Jonker, "Evolutionarily Stable Strategies and Game Dynamics," *Mathematical Biosciences* 40 (1978):145–156.

Uzawa, Hirofumi, "Edgeworth's Barter Process and Walras' Tâtonnement Process," Technical Report, Office of Naval Research Contract NR-047-004, Department of Economics, Stanford University (1959). Technical Report No. 83.

— , "The Stability of Dynamic Processes," *Econometrica* 29,4 (October 1961):617–631.

— , "On the Stability of Edgeworth's Barter Process," *International Economic Review* 3,2 (May 1962):218–231.

Wald, Abraham, "On Some Systems of Equations of Mathematical Economics," *Econometrica* 19,4 (October 1951 [1936]):368–403.

Walras, Leon, *Elements of Pure Economics* (London: George Allen and Unwin, 1954 [1874]).

Wicksell, Knut, *Lectures in Political Economy, Vol. 2. Money* (New York: Kelley, [1911]1967).

c\Papers\Stability of Market Exchange\MarkovExchange.tex December 16, 2011