The Dynamics of Pure Market Exchange

Herbert Gintis*

April 11, 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.

Abstract

This paper investigates the out-of-equilibrium price and quantity adjustment process for a decentralized market economy with individual production. The economy is modeled as an ergodic 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 a system 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 Markov 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.

1 Introduction

Adam Smith (2000[1759]) envisioned a decentralized economy that sustains an efficient allocation of resources through the "invisible hand" of market competition.

*Santa Fe Institute and Central European University

Smith's vision was formalized by Léon Walras in 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 provided a rather complete proof of the existence of equilibrium in decentralized market economies. Such economies are that a particularly attractive because of the so-called first fundamental theorem of welfare economics, which asserts Walrasian equilibrium is Pareto-efficient (Arrow 1951, Debreu 1951, Debreu 1954, Hurwicz 1960).

The question of stability of the Walrasian economy was a central concern of mathematical economists in the years immediately following the proof of existence of equilibrium (Arrow and Hurwicz, 1958, 1959, 1960; Arrow, Block and Hurwicz, 1959; McKenzie, 1960; Morishima 1960; Nikaido 1959: Nikaido and Uzawa 1960; Uzawa 1960a,b). The models studied by these the researchers 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 assumption of public prices was implemented by positing 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."

This quest for a general stability theorem was untracked in 1959, however, by Herbert Scarf"s simple examples of unstable Walrasian equilibria (Scarf 1960). Indeed, more recent studies have shown that the tâtonnement dynamic resulting from public prices adjusted by an auctioneer who knows only the structure of excess demand is stable only under extremely stringent and implausible conditions (Fisher 1983, Kirman 1992), and chaos in price movements is the generic case for the tâtonnement adjustment processes (Saari 1985, Bala and Majumdar 1992).

There were attempts soon after Scarf"s counterexamples to continue the analysis of tâtonnement by adding out-of-equilibrium trading (Uzawa 1959, 1961, 1962; Negishi 1961; Hahn 1962; Hahn and Negishi 1962), but with only limited success. Fisher (1970, 1972) later explored tâtonnement processes with refined structures of out-of-equilibrium trading, but these too proved less than successful. Surveying the state of the art some quarter century after Arrow and Debreu's seminal existence theorems, Fisher (1983) concluded that virtually no progress had been made towards a cogent model of Walrasian market dynamics.

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, 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, satisfying Walras' Law, and even requiring that utility functions be homothetic, is the excess demand function for some Walrasian economy, when mild technical conditions hold.

More recently, 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 rather disappointing assessment, but his conclusion remains valid. Despite the centrality of the general equilibrium model to economic theory, we know nothing systematic about Walrasian market dynamics.

2 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, or by anything else. 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. 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 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 sole admissible forms of experience in a decentralized market economy result from producing, consuming, trading, and observing the behavior of agents with whom one interacts. Recall of one's own trading experience and knowledge of the trading strategies of exchange partners alone can be the basis for updating the system of private prices, and equilibrium can be achieved only if plausible models of inference and updating lead private prices to converge to public prices through market exchange (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.

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 analytical tool 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, although there may be some additional purely macro states as well (e.g., the supply of a public good). We may assume agents have utility functions and perhaps other characteristics, such as acquired skills and talents, that do not change over the time period under investigation. The state of each agent, then, includes his holding of each good, service, and resource, 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. When a state variable is a price or quantity, it can be represented using scientific notation as an ordered pair consisting of an integral mantissa and a signed integral exponent, where both mantissa and exponent lied in bounded intervals. Thus there are a finite number of states, but even a very small economy therefore has an astronomical number of states.

Nevertheless, the Markov process 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 a capacity for each good. An small but strictly positive probability of mutation among "nearby" states then ensures that the Markov process is aperiodic. Assuming the Markov process is irreducible, being both finite and aperiodic implies the Markov process is ergodic (Feller 1950), which means it has a stationary distribution expressing the long-run probability of being in each state of the system, irrespective of its initial state. Of course, we do 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. The ergodic theorem ensures that these aggregates have determinate long-term stationary distributions.

If the transition matrix of an n-state Markov process is P, and the t-period transition matrix is $P^{(t)} = P^t$, the stationary distribution $u = (u_1, \ldots, u_n)$ of the Markov process is a probability distribution with strictly positive entries that has the following properties for $j = 1, \ldots, n$.

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

$$u_j = \sum_i u_i p_{ij} \tag{2}$$

Equation (1) says that u_j the long-run frequency of state s_j in any realization $\{s^t\}$ of the Markov process with probability 1, and that this probability is strictly positive and independent from the starting state $s^0 = s_i$ of the realization. By a well-known property of convergent sequences, (1) 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 (2) 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 $\{s^t\}$. This means that a Markov process has only a one period "memory." However, we can treat a finite sequence of historical states of the Markov process of fixed maximum 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 Markovian assumption imposes no substantive constraint on modeling systems that are subject to the laws of physics.

In case the latter claim is not obvious, suppose a Markov process has transition matrix P, with entries $p_{i,j}$ and consider two-period states of the form ij. We define the transition probability of going from ij to kl as

$$p_{ij,kl} = \begin{cases} p_{j,l} & j = k \\ 0 & j \neq k, \end{cases} \tag{3}$$

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 $\{u_1, \ldots, u_n\}$ is the stationary distribution associated with P, then

$$u_{ij} = u_i p_{i,j} \tag{4}$$

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

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

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

$$u_{ij} = u_i p_{i,j} = \sum_k u_k p_{k,i} p_{i,j} = \sum_k u_{ki} p_{i,j} = \sum_{kl} u_{kl} p_{kl,ij}. \tag{5}$$

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 $\{p_{ij,kl}\}$. Equation 4 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,

$$u_{kl} = \lim_{t\to\infty} p^{(t)}_{ij,kl} \tag{6}$$

$$u_{ij} = \sum_{kl} u_{kl} p_{kl,ij}. \tag{7}$$

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

If the 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 $t' = 1, 2, \ldots$. 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 necessarily reappears infinitely often with probability one in any realization of the Markov process.

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 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 and 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, which can be proved using a similar argument.

It follows that the communicate relation 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 the 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_r^i 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 *quasi-ergodic theorem*: 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 $r = 1,\ldots,k$, and for each $i \in S^{\mathrm{tr}}$, 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

$$u_{ij} = \lim_{t\to\infty} P_{ij}^{(t)}; \tag{8}$$

$$u_j^r = u_{ij} \quad \text{if } i, j \in S_r; \tag{9}$$

$$u_j^r = \sum_{i\in S_r} u_i^r p_{ij} \quad \text{for } j \in S_r; \tag{10}$$

$$u_{ij} = q_r^i u_j^r \quad \text{if } s_i \in S^{\mathrm{tr}} \text{ and } s_j \in S_r. \tag{11}$$

$$u_{ij} = 0 \quad \text{if } s_j \in S^{\mathrm{tr}}. \tag{12}$$

$$\sum_j u_{ij} = 1 \quad \text{for all } i = 1,\ldots,n. \tag{13}$$

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 ergodic Markov process as it unfolds. We are in fact only interested in measuring certain aggregate properties of the process rather than its stationary distribution. 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 ergodic 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.

3 Scarfian Instability Revisited

To assess the effect of passing from public to private prices, this section analyzes Herbert Scarf's seminal example of Walrasian instability, first using standard public prices and then Markovian private prices. Scarf constructed a three-good economy in which each agent produces one good and consumes some of his production 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 X-producer consumes X and Y in proportion $x = y$, so his utility is given by

$$u_x(x, y, z) = \min\{x, y\}.$$

Similarly, for the other two agents, we have

$$u_y(x, y, z) = \min\{y, z\},$$

and

$$u_z(x, y, z) = \min\{z, x\}.$$

It is straightforward to show that utility maximization for the X-producer, where p_x and p_y are the prices of X and Y and x_X^d, y_X^d are the X-producer's final demand for X and Y, gives

$$x_X^d = y_X^d \qquad \text{and } p_x(1 - x_X^d) = p_y y_X^d,$$

which give the agents final demand (x_X^d, y_X^d):

$$x_X^d = y_X^d = \frac{p_x}{p_x + p_y}.$$

These equations allows us to calculate total excess demand E_g for each good $g \in$ {X,Y,Z} as a function of the prices of the three goods. It is easy to check that the and market-clearing prices, normalizing $p_z^* = 1$, are given by $p_x^* = p_y^* = 1$.

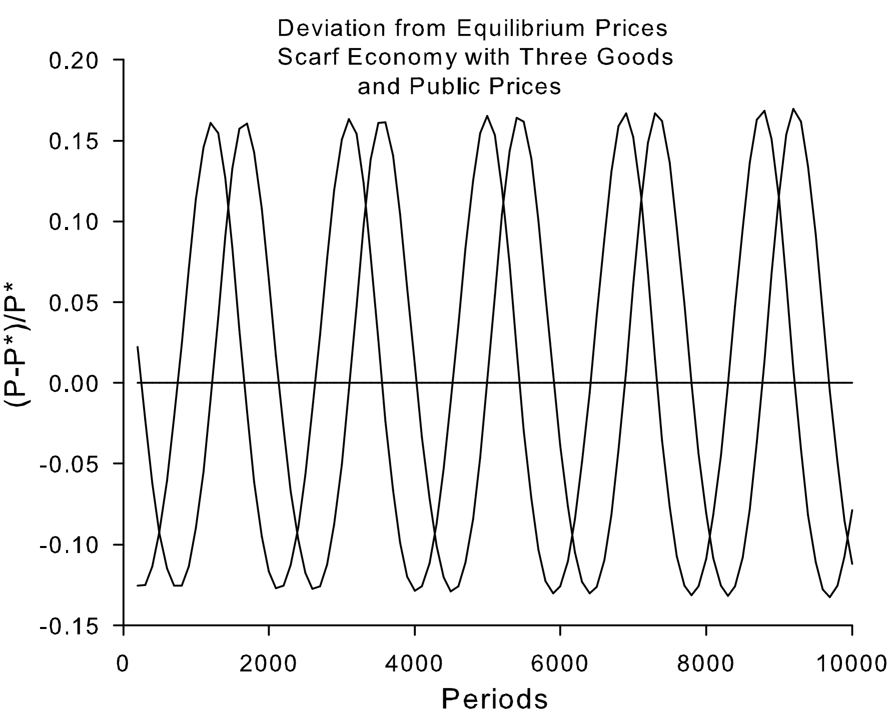

Figure 1: Non-convergence of price to steady state in a three-good Scarf Economy with public prices.

Suppose we start with disequilibrium prices $p_x = p_x^* + 3$, $p_y = p_y^* - 2$,

$p_z = p_z^* = 1$. The excess demands E_x and E_y are then used to update prices:

$$p'_x = p_x + E_x/100 \tag{14}$$
$$p'_y = p_y + E_y/100 \tag{15}$$

Note that the denominator on the right hand side 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 be unstable. The process of calculating excess demand and updating prices is then repeated a large number of times. We are thus treating the Scarf system as a Markov process, but by using public prices, the dynamics of the system are virtually identical to the differential equation dynamics that Scarf analyzed. The result after 10,000 periods is shown in Figure 1, and perfectly replicates the analytical results of Scarf (1960).

4 The Scarf Economy with Private Prices

For the private price version of the Scarf economy, we maintain the above assumptions, except now there are 1000 traders of each of the three types and each trader is endowed at the beginning of a run a 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 show 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 with a period roughly ten times as long as in the case of the Scarfian economy with public prices. This is actually the unique limit cycle of the model, which can be verified by reinitializing the process from many different initial states. In all cases, the stationary distribution has approximately the same period and amplitude.d

In sum, we have developed a perfectly specified 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 barter algorithm, for which alternative algorithms are plausible, all modeling choices are uniquely determined by the standard conception of the Walrasian general equilibrium model. The equi-

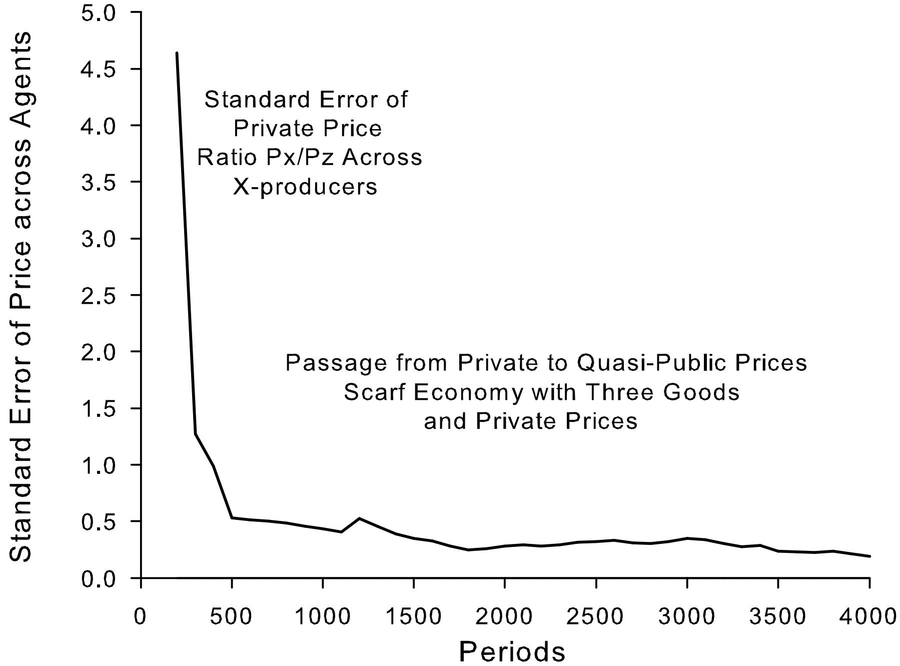

Figure 2: The Markov version of the Scarf economy cycles initialized with random prices quickly transitions to quasi-public prices, which are public prices with relatively low standard error across individuals.

librium 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 calculated, even with the fastest and most powerful conceivable computer aids. Thus, as in the natural sciences, we are obliged to investigate the stationary distribution by simulating the process 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 exhibit unique limit cycles around the equilibrium price vector, with private prices and a completely decentralized dynamical system of price adjustment.

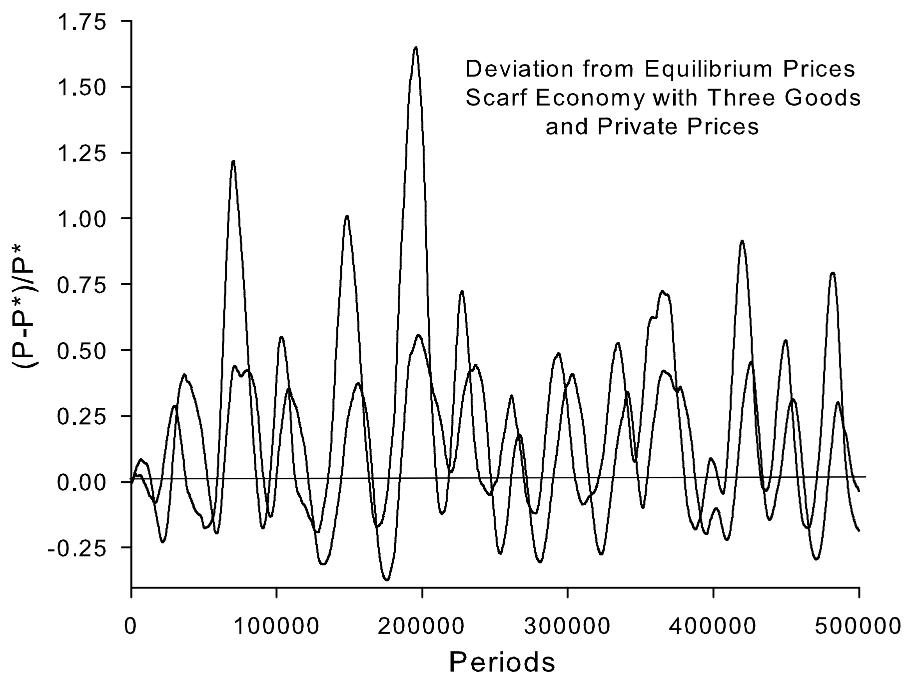

Figure 3: The Scarf economy cycles without the presence of a centralized price-making and adjusting agent when modeled as a Markov process with private prices.

5 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 other information about the economy other than that gathered from private experience in trade, except that periodically they discover the private price vector of another agent and copy 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 necessarily still in principle somewhat underspecified. 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 $k = 1, \ldots, n$ produces good k in "styles" $s = 1, \ldots, m$ (we use "styles" to enrich the heterogeneity of goods in the model without seriously increasing the computational resources needed to estimate the stationary of the distribution resulting Markov process). Each agent consumes a subset of his 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 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 $G = \{g_s^k | k = 1, \ldots, n, s = 1, \ldots m\}$. 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 personal labor and no other inputs 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

$$u(x_1,\ldots,x_n) = \left(\sum_{l=1}^{m} \alpha_l x_l^{\gamma}\right)^{1/\gamma}, \tag{16}$$

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

$$x_i = \frac{M f_i}{\sum_{l=1}^{m} p_l \phi_{il}^{1/(1-\gamma)}}, \tag{17}$$

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

$$x_i = \frac{M f_i \alpha_i}{p_i}. \tag{18}$$

The reason for creating such a complex array of utility functions is to ensure that our results are not the result of assuming an excessively narrow set of consumer characteristics. However, a 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.

6 Estimating the Stationary Distribution

I will illustrate this dynamic assuming $n = 9$, $m = 5$, and $N = 300$, so there are 40 distinct goods which we write as $g_1^1, \ldots, g_5^9$, and 12000 traders in the economy. There are thus 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 unity, given that costs of production are the same in all industries and the utility functions are randomly generated. 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 $\epsilon = 0.01$.

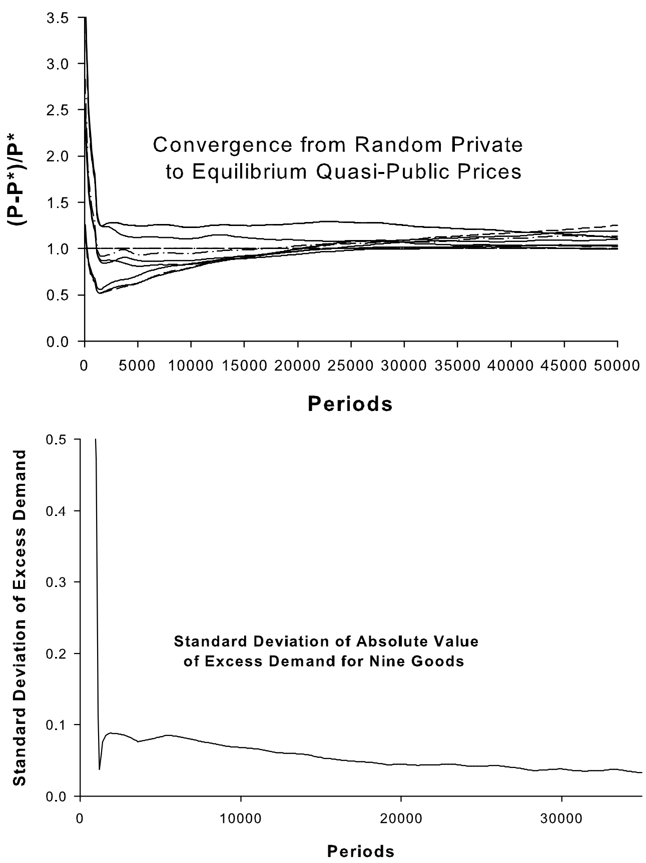

Figure 4: The Market Economy

The results of a typical run of this model is illustrated in Figure 4. The upper pane of the figure shows the passage from private to quasi-public prices over the first 10,000 trading periods. 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 nine 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 1500 periods, and falling another order of magnitude over the next 8500 periods. The final value of this standard error is 0.0223, as compared with it initial value of 6.3992. 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 simulation. 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.

7 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 simulate this economy using the same parameters as in our previous simulations, including zero inventory costs. Figure 5 shows that the use of money increases monotonically over the first 2000 periods, spread almost equally among the six 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"

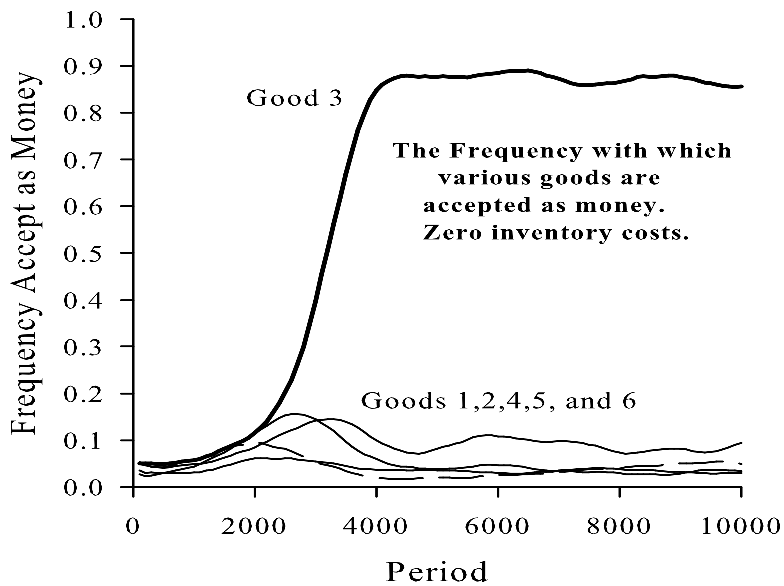

Figure 5: The Emergence of Money in a Market Economy. The parameters of the model are the same as in the baseline case treated previously. Inventory costs are assume absent.

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 6. 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%.

8 The Robustness of the Decentralized Market Economy

The above Markov model is extremely resilient in the face of aggregate shocks. We illustrate this admitting of a single new good, *fiat* money, that is neither produced nor consumed, and enjoys zero inventory storage costs. When such a good is available, it quickly becomes a universal medium of exchange for the economy, accepted by almost 100% of market traders. The nature of market dynamics in this context is not noticeably different from the baseline case of non-monetary exchange, and the efficiency of the economy is heightened in the same manner as depicted in Figure 6. Suppose, however, that every 5000 periods, we impose a shock on the economy consisting of a reduction in the *fiat* money holdings of each

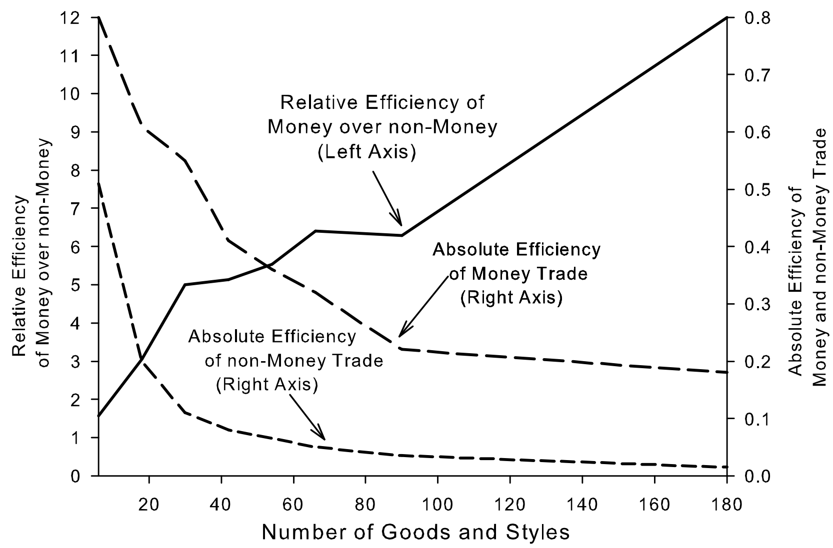

Figure 6: The Relative Efficiency of Money in a Market Economy

trader to 20% of its current level. The reduced holding are maintained for 100 periods, after which the money holdings of each trader is multiplied by five, restoring the money stock for the economy to its initial level. The southwest pane in Figure 7 shows the course of the money supply. The southeast pane of the figure shows that relative efficiency suffers during the shock period, but is perfectly restored upon the resumption of the normal money supply. The northeast and northwest panes show that prices and excess demand are virtually unaffected by the sequence of 100 period shocks.

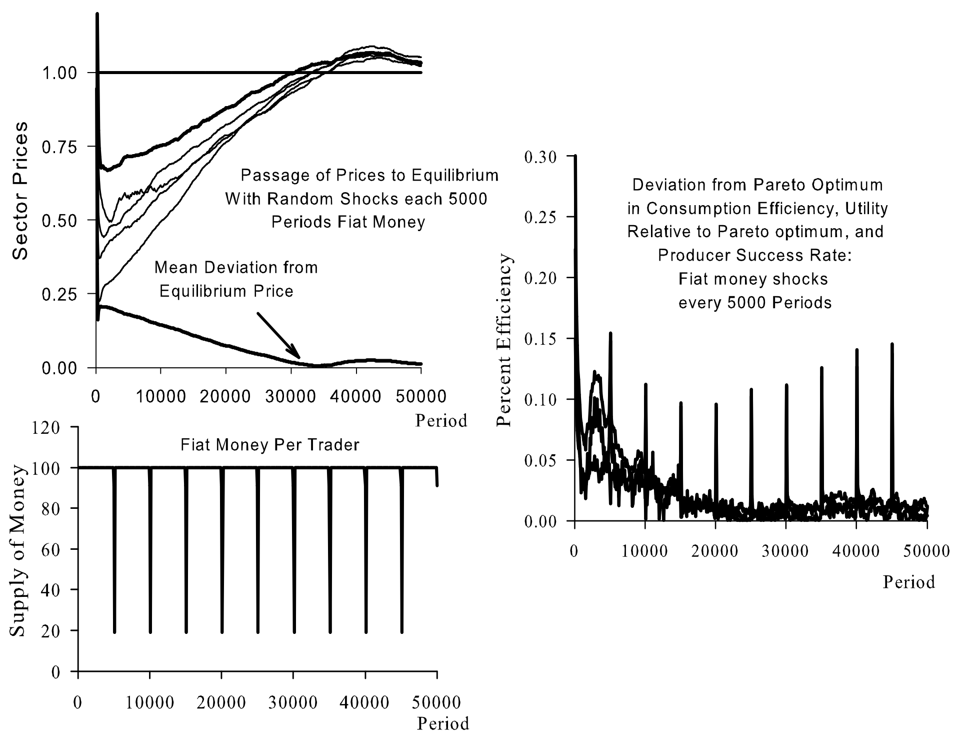

Figure 7: The Resilience of a Market Economy with *Fiat* Money

9 Conclusion

Attempting to explain market dynamics assuming public prices has been an heroic but ultimately disappointing enterprise, Fisher (1983) doubtless being the most ambitious and successful attempt. Even had the quest for an effective dynamic with public prices been successful, however, it would have been of little practical value, because out of equilibrium there is no reasonable sense in which public prices exist in a market economy. By modeling market exchange assuming each agent has as set of *private prices* that is updated through learning and imitation, we have seen than effective dynamics can be analytically modeled. Our Markov model of the Scarf economy produced a limit cycle stationary distribution akin to the neutral stability of the Scarf economy with public prices, and our Markov model of simple decentralized market exchange is extremely stable and resilient in the face of macroeconomic shocks.

Markov models of market exchange are complex dynamical systems with precise mathematical specification in algorithmic form. These models cannot be solved analytically, and we have as yet no formal means of demonstrating properties of the stationary distribution other than by empirical estimation. Yet we know that these Markov models behave in consistent and comprehensible ways. First, it is always the case that starting from a state of pure randomness, under the twin influence of learning and imitation, private prices rapidly converge to *quasi-public prices*, which have the property of differing across individuals, but with a very small standard error. Second, in a general market setting with heterogeneous agents, quasi-public prices adjust to their equilibrium levels in the long run, leading to a *quasi-market equilibrium* with approximately constrained Pareto-optimal allocations, at least in the simple case of market exchange in which each agent produces a single good using only personal labor.

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 Economics," 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., "Quasi-competitive Price Adjustment by Individual Firms: A Preliminary Paper," *Journal of Economic Theory* 2 (1970):195–206.

— , "On Price Adustments without an Auctioneer," *Journal of Economic Theory* 39 (1972):1–15.

— , *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, *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.

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

c\Papers\Stability of Market Exchange\MarkovExchange.tex April 11, 2011