

Agent-Based Models as Markov Processes: Theory and Application to the Dynamics of Market Exchange

Herbert Gintis*

March 28, 2012

Abstract

Traditional differential equation approaches to modeling the dynamics of decentralized market exchange have made little progress despite more than sixty years of effort on the part of mathematical economists. Agent-based models have recently been shown to shed light on market dynamics, but their theoretical status is unclear. This paper shows that such models can be treated as finite Markov processes, and their long-run properties are then given by an expanded version of the ergodic theorem for Markov processes. A Markov process model of a simplified market economy shows the fruitfulness of this approach, and highlights important unsolved problems concerning finite Markov processes with large state spaces.

1 Introduction

The scientist often possesses powerful mathematical models, but the equations do not admit of closed form analytical solutions. Computer simulations illuminate the behavior of such processes. However, the social sciences lack the analytical tools to model the complex collective behavior of human agents. In this situation, agent-based models often illuminate system dynamics (Tesfatsion and Judd 2006, Gintis 2007). However, the results of such modeling exercises have an ambiguous scientific status, because they are not solutions to analytically specified but intractable equations.

What is the analytical status of an agent-based model? In this paper I propose an answer, and apply it to a particularly intractable problem in economic theory: the dynamics of decentralized market production and exchange. An agent-based model is a finite Markov process, and the application to market exchange proves the

*Santa Fe Institute and Central European University. For publication in ACM Transactions on Intelligent Systems and Technology

existence of a stationary distribution of market models. Uniqueness of the solution translates into the ergodicity of the Markov process.

Economists are also interested in whether a stationary distribution is close to market equilibrium (equality of supply and demand in all markets), and in the stability of such an equilibrium. The conditions under which the stationary distribution of an ergodic Markov model of the economy is a market equilibrium is currently unknown.

2 A Markov Process Primer

A finite Markov process \mathcal{M} consists of a finite number of *states* $S = \{1, \dots, n\}$, and an n -dimensional square matrix $P = \{p_{ij}\}$ such that p_{ij} represents the probability making a transition from state i to state j . A *path* $\{i_1, i_2, \dots\}$ determined by Markov process \mathcal{M} consists of the choice of an initial state $i_1 \in S$, and if the process is in state i in period $t = 1, \dots$, then it is in state j in period $t + 1$ with probability p_{ij} . Despite the simplicity of this definition, finite Markov processes are remarkably flexible in modeling dynamical systems, and characterizing their long-run properties becomes highly challenging for systems with more than a few states.

To illustrate, consider a rudimentary economy in which agents in each period produce goods and sellers are randomly paired with buyers who offer money in exchange for the seller's good. Suppose there are g distinct types of money, each agent being willing to accept only one type of money in trade. Trade will be efficient if all agents accept the same good as money, but in general inefficiencies result from the fact that a seller may not accept a buyer's money.

What is the long run distribution of the fraction of the population holding each of the types as money, assuming that one agent in each period switches to the money type of another randomly encountered agent? Let the state of the economy be a g -vector $(w_1 \dots w_g)$, where w_i is the number of agents who accept type 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 g distinguishable boxes (the g types), which is $C(n + g - 1, g - 1)$, where $C(n, g) = n!/(n - g)!g!$ is the number of ways to choose g 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 type 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 type i as money, and the number of x 's after the final A is

the number agents choosing type k as money. The total number of x 's is equal to n , and the total number of A 's is $g - 1$, so the length of s is $n + g - 1$. Every placement of the $g - 1$ A 's represents particular state of the system, so there are $C(n + g - 1, g - 1)$ states of the system. For instance, if $n = 100$ and $g = 10$, 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 type as his own money. 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 clearly a finite Markov process.

What is the long-run behavior of this Markov process? Note first that if we start in state i at time $t = 1$, the probability $p_{ij}^{(2)}$ of being in state j in period $t = 2$ is simply

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

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

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

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 these methods 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

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

where all n agents choose type 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 g 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 $q \rightarrow \infty$, it follows that state s appears only a finite number of times with probability one. We call s a *transient* state.

We can often calculate the probability that a system starting out with w_r agents choosing type r as money, $r = 1, \dots, g$ is absorbed by state r . Let us think of the Markov process as that of g gamblers, each of whom starts out with an integral number of coins, there being n coins in total. The gamblers represent the types and their coins are the agents who choose that type for money, there being n agents in total. We have shown that in the long run, one of the gamblers will have all 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 agents holding type r . Then in a finite number of time periods, the process will be absorbed into one of the states $1, \dots, g$, and the probability of being absorbed into state r is w_r/n .

3 Long-run Behavior of a Finite Markov Process

An n -state Markov process \mathcal{M} has a *stationary distribution* $u = (u_1, \dots, u_n)$ if u is a probability distribution satisfying

$$u_j = \sum_{i=1}^n u_i p_{ij}, \quad (2)$$

or more simply $uP = u$. This says that for all states i , the process spends fraction u_i of the time in state i , so the fraction of time in state j is the probability it was in some state i in the previous period, times the probability of transiting from i to j , summed over all states i of \mathcal{M} . We show below that every finite Markov process has at least one stationary distribution.

If state $i \in S$ in \mathcal{M} has a positive probability of making a transition to state j in a finite number of periods (that is, $p_{ij}^{(t)} > 0$ for some nonnegative t), we say j is *accessible* from i . If states i and j are mutually accessible, we say that the two states *communicate*. A state always communicates with itself because for all i , we have $p_{ii}^{(0)} = 1 > 0$. If all states in a Markov process mutually communicate, we say the process is *irreducible*. More generally, if A is any set of states, we say A is *communicating* if all states in A mutually communicate, and no state in A communicates with a state not in A . A Markov process \mathcal{M} with at least two states, one of them absorbing, as was the case in our previous example, cannot be irreducible, because no other state is accessible from an absorbing state. By definition a set consisting of a single absorbing state is communicating.

The communication relation is an equivalence relation of S . To see this, note that if i communicates with j , and j communicates with k , then clearly i communicates with k . Therefore the communication relation is transitive. The relation is symmetric and reflexive by definition. As an equivalence relation, the communication relation partitions S into communicating sets S_1, \dots, S_k .

We say a set of communicating states A is *leaky* if some state $j \notin A$ is accessible from a state in A , in which case j is of course accessible from any state in A . If a communicating set the partition, say S_r , is leaky, then with probability one \mathcal{M} will eventually transit to an accessible state outside S_r and \mathcal{M} will never return to S_r . To see this suppose the contrary, and let $i \in S_r$ and $j \notin S_r$ with $p_{ij} > 0$. Suppose from state j , eventually \mathcal{M} makes a transition to a state $j' \notin S_r$ and $p_{j'i'} > 0$ for some state $i' \in S_r$. Then necessarily j' is accessible from j , and i' is accessible from j' , so i is accessible from j . But then i and j communicate, which is a contradiction. Therefore a leaky closed set S_r consists wholly of transient states, and with probability one there is a time $t_r > 0$ such that no state in S_r appears after time t_r . Because there are only a finite number of closed sets in the partition of S , there is a time $t^* > 0$ such that no member of a leaky closed set appears after time t^* . This proves that there must be at least one non-leaky closed set. We call a non-leaky closed set an *irreducible set*. An irreducible set of states in \mathcal{M} is obviously an irreducible Markov *subprocess* of \mathcal{M} , meaning that the states of the set themselves form an irreducible Markov process with the same transition probabilities as defined by \mathcal{M} .

We thus know that a finite Markov process can be partitioned into a set S^{tr} of transient states, plus irreducible Markov subprocesses S_1, \dots, S_m . S^{tr} is then the union of all the leaky closed sets in \mathcal{M} . The Markov process may start in a transient state, but eventually transits to one of the irreducible subprocesses S_1, \dots, S_m . The long-term behavior of \mathcal{M} depends only on the nature of these irreducible subprocesses.

It is desirable to have the long run frequency of state i of the Markov process be the historical average of the fraction of time the process spends in state i , because in this case we can estimate the frequency of a state in the stationary distribution by its historical frequency in a series of runs of the process. When this is the case, we say the Markov process is *ergodic*. We must add one condition to irreducibility to ensure the ergodicity of the Markov process. We say a state i has *period* $k > 1$ if $p_{ii}^{(k)} > 0$ and whenever $p_{ii}^{(m)} > 0$, then m is a multiple of k . If there are no periodic states in the Markov process, we say it is *aperiodic*. It usually is possible, as we will see, to ensure the aperiodicity of a Markov process representing complex dynamical phenomena. We have the following theorem (Feller 1950).

Theorem 1. ERGODIC THEOREM *Let \mathcal{M} be an n -state irreducible aperiodic Markov process with probability transition matrix P . Then \mathcal{M} has a unique stationary distribution $u = (u_1, \dots, u_n)$, where $u_i > 0$ for all i and $\sum_i u_i = 1$. For any initial state i , we have*

$$u_j = \lim_{t \rightarrow \infty} p_{ij}^{(t)} \quad \text{for } i = 1, \dots, n. \quad (3)$$

Equation (3) implies that u_j is the long-run frequency of state s_j in a realization $\{s^t\} = \{s^0, s^1, \dots\}$ of the Markov process. By a well-known property of absolutely convergent sequences, (3) implies that u_j is also 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).

If a Markov process \mathcal{M} is aperiodic but not reducible, we know that it has a set of transient states S^{tr} and a number of irreducible aperiodic subprocesses S_1, \dots, S_k . Each of these subprocesses S_r is an ergodic Markov process derived from \mathcal{M} by eliminating all the states not in S_r , and so has a strictly positive stationary distribution u_r over its states. If we expand u_r by adding zero entries for the states in \mathcal{M} but not in S_r , this clearly gives us a stationary distribution for \mathcal{M} . Because there is always at least one ergodic subprocess for any finite aperiodic Markov process, this proves that every aperiodic Markov process has a stationary distribution. Moreover, it is clear that there are as many stationary distributions as there are ergodic subprocesses.

The ergodic theorem and the above remarks allow us to fill out the general picture of behavior of the finite aperiodic Markov process. Such a process may start in a transient state, but ultimately it will enter one of the ergodic subprocesses, where it will spend the rest of its time, the relative frequency of different states being given by the stationary distribution of the subprocess. We thus have

Theorem 2. EXPANDED ERGODIC THEOREM *Let \mathcal{M} be a finite aperiodic Markov process. Then there is a probability transition matrix $P = \{p_{ij}\}$ of \mathcal{M} such that*

$$u_{ij} = \lim_{t \rightarrow \infty} P_{ij}^{(t)}. \quad (4)$$

Moreover, there exists a unique partition $\{S^{\text{tr}}, S_1, \dots, S_k\}$ of the states S of \mathcal{M} , a probability distribution u^r over S_r for $r = 1, \dots, k$, such that $u_i^r > 0$ for all $i \in S_r$, and for each $i \in S^{\text{tr}}$, there is a probability distribution q^i over $\{1, \dots, k\}$ such that for all $i, j = 1, \dots, n$ and all $r = 1, \dots, k$, we have

$$u_j^r = u_{ij} \quad \text{if } i, j \in S_r; \quad (5)$$

$$u_j^r = \sum_{i \in S_r} u_i^r p_{ij} \quad \text{for } j \in S_r; \quad (6)$$

$$u_{ij} = q_r^i u_j^r \quad \text{if } s_i \in S^{\text{tr}} \text{ and } s_j \in S_r. \quad (7)$$

$$u_{ij} = 0 \quad \text{if } s_j \in S^{\text{tr}}. \quad (8)$$

$$\sum_j u_{ij} = 1, u_{ij} \geq 0 \quad \text{for all } i = 1, \dots, n. \quad (9)$$

Equation (4) asserts that u_{ij} is the long-run probability of being in state j when starting from state i . Equation (5) and (6) assert that for the states j belonging to an ergodic subprocess S_r of \mathcal{M} , $u_j^r = u_{ij}$ do not depend on i and represent the stationary distribution of S_r . Equations (7) and (8) assert that transient states eventually transit to an ergodic subprocess of \mathcal{M} .

4 Extensions of Markov Processes

A Markov process by construction has only a one period memory, meaning that the probability distribution over states in period t depend only on the state of the process in period $t - 1$. However, we will show that, if we consider a finite sequence of states $\{i_{t-k}, i_{t-k+1}, \dots, i_{t-1}\}$ of the Markov process of fixed length k to be a single state, then the process remains a finite Markov process and is ergodic if the original process was ergodic. In this way we can deal with stochastic processes with any finite 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.

A *stochastic process* \mathcal{S} consists of a state space S and probability transition function P with the following properties. Let $H_{\mathcal{S}}$ be the set of all finite sequences of elements of S . We interpret $(i_t, i_{t-1}, \dots, i_0) \in H_{\mathcal{S}}$ as the process being in state

i_τ in time period $\tau \leq t$. The *history* h_t of the stochastic process at time $t > 1$ is defined to be $h_t = (i_{t-1}, \dots, i_0)$. We also define $h_t^k = (i_{t-1}, \dots, i_{t-k})$. We write the set of histories of \mathcal{S} at time t as H_S^t . The probability transition function for the stochastic process has entries of the form $p_i(h_t)$, which is the probability of being in state i in time t if the history up to time t is h_t . We say \mathcal{S} has *k-period memory* if, for all $i \in S$, $p_i(h_t) = p_i(h_t^k)$ for $t > k$. We write the set of histories h_t^k as \mathcal{S} as H_S^k .

We say a stochastic process \mathcal{S} with state space S and transition probability function P with k -period memory and stochastic process \mathcal{T} with state space S' and transition probability function P' with k' -period memory are *isomorphic* if there is a bijection $\phi : S \rightarrow S'$ and a bijection $\psi : H_S^k \rightarrow H_{S'}^{k'}$ such that for all $i \in S$, $p_i(h_t^k) = p'_{\phi(i)}(\psi(h_t^k))$.

Theorem 3. FINITE HISTORY THEOREM *Consider a stochastic process \mathcal{S} with finite state space S and transition probability function P with k -period memory, where $k > 1$. Then \mathcal{S} is isomorphic to a Markov process \mathcal{M} with state space S' , where $i' \in S'$ is a k -dimensional vector $(i_1, \dots, i_k) \in S^k$. The function ϕ is the identity on S , and ψ is the canonical isomorphism from H_S^k to S^k .*

Proof: We define the transition probability of going from $(i_1, \dots, i_k) \in S^i$ to $(j_1, \dots, j_k) \in S^k$ in \mathcal{M} as

$$p_{(i_1, \dots, i_k), (j_1, \dots, j_k)} = \begin{cases} p_{j_k}(i_1, \dots, i_k) & j_1 = i_2, \dots, j_{k-1} = i_k \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

This equation says that (i_1, \dots, i_k) represents “state i_k in the current period and state i_τ in period $\tau < k$.” It is easy to check that with this definition the matrix $\{p_{ij,kl}\}$ is a probability transition matrix for \mathcal{M} , and \mathcal{M} is isomorphic to \mathcal{S} , proving the theorem.

Note that we can similarly transform a finite Markov process \mathcal{M} with state space S into a finite Markov process \mathcal{M}^k with a k -dimensional state space S^k for $k > 1$, and if \mathcal{M} is ergodic, \mathcal{M}^k will also be ergodic. For ease of exposition, let us assume $k = 2$ and write $(i, j) \in S^2$ as ij , where i is the state of \mathcal{M} in the current period and j is its state in the previous period. Then if $\{u_1, \dots, u_n\}$ is the stationary distribution associated with \mathcal{M} , then

$$u_{ij} = u_i p_{ij} \quad (11)$$

defines a stationary distribution $\{u_{ij}\}$ for \mathcal{M}^2 . Indeed, we have

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

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

$$u_{ij} = u_i p_{i,j} = \sum_k u_k p_{k,i} p_{i,j} = \sum_k u_{ki} p_{i,j} = \sum_{kl} u_{kl} p_{kl,ij}. \quad (12)$$

It is straightforward to show that pairs of states of \mathcal{M} correspond to single states of \mathcal{M}^2 . These two equations imply the ergodic theorem for $\{p_{ij,kl}\}$ because equation 11 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 ,

$$u_{kl} = \lim_{t \rightarrow \infty} p_{ij,kl}^{(t)} \quad (13)$$

$$u_{ij} = \sum_{kl} u_{kl} p_{kl,ij}. \quad (14)$$

Let \mathcal{M} be a Markov process with state space S , transition probability matrix $P = \{p_{ij}\}$, and initialization probability distribution q on S , so the probability that the first state assumed by the system is i is given by q_i . The probability that a sequence $(i_0, i_1, \dots, i_\tau)$ occurs when \mathcal{M} is run for τ periods is then given by

$$q_{i_0} p_{i_0 i_1} \times \dots \times p_{i_{\tau-1} i_\tau}.$$

Theorem 4. STACKING THEOREM *Consider a finite Markov process \mathcal{M} with state space S , transition matrix $P = \{p_{ij} | i, j \in S\}$, and an initialization probability distribution q on S . Let \mathcal{M}^* be the stochastic process given by running \mathcal{M} for τ^* periods, reinitializing the process, running \mathcal{M} again for τ^* periods, and so on. Then \mathcal{M}^* is a finite Markov process. If $q_i > 0$ for all $i \in S$, the \mathcal{M}^* is ergodic.*

Proof: We set the state space for \mathcal{M}^* to be $S \times \mathbb{Z}_{\tau^*}$, where $\mathbb{Z}_k = \{1, \dots, k\}$. The transition probabilities are given by

$$p^{(i,t)(j,t')} = \begin{cases} p_{ij} & t' = t + 1 \leq \tau^* \\ q_j & t = \tau^*, t' = 1 \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

Then \mathcal{M}^* is clearly a finite Markov process, and because q is strictly positive on S , \mathcal{M}^* has no transient states. Because there is a positive probability of remaining in the same state for two periods, \mathcal{M}^* is aperiodic, and since it is also irreducible, by the ergodic theorem (Theorem 1, \mathcal{M}^* is ergodic.

An important question is the nature of aggregations 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. We have

Theorem 5. AGGREGATION THEOREM *Suppose a finite Markov process with state space S has a set of states $A \subset S$ with all $j \in A$ identically situated in the sense that $p_{ji} = p_{ki}$ for all states $i \in S$. Then there is a Markov process with the same transition probabilities as \mathcal{M} , except the states in A are replaced by a single state.*

Proof: 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

$$u_m = \lim_{t \rightarrow \infty} P_{im}^t \quad (16)$$

$$u_m = \sum_i u_i p_{im} \quad (17)$$

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

$$P_{mi} = u_j p_{ji} + u_k p_{ki}. \quad (18)$$

If j and k are identically situated, then (18) implies

$$u_i = \sum_r u_r p_{ri}, \quad (19)$$

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 identically situated states can be aggregated into a single in this manner.

More generally, we may be able to partition the states of \mathcal{M} into cells m_1, \dots, m_l such that, for any $r = 1, \dots, l$ and any states i and j of \mathcal{M} , i and j are identically situated with respect to each m_k . When this is possible, then m_1, \dots, m_l are the states of a derived Markov process, which will be ergodic if \mathcal{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. Note that in many Markov models of market interaction, the states of two agents can be interchanged without changing the transition probabilities, so such agents are identically situated. In this situation, 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 too many to permit an analytical derivation of the stationary distribution.

5 Markov Processes and Agent-based Models

For a Markov process with a small number of 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, as is typically the case in modeling social processes, these methods are impractical. Rather, we must construct an accurate computer model of the Markov process, and ascertain empirically the dynamical properties of the irreducible Markov subprocesses. We are in fact often 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 Expanded 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, \dots, 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 \mathcal{M} with transient states S^{tr} and ergodic subprocesses S_1, \dots, S_k enters a subprocess S_r after τ^* periods with high probability, and suppose the historical average over states from $t = 0$ to $t = t^*$ is a close approximation to the stationary distribution of S_r . The existence of τ^* is guaranteed by the expanded ergodic theorem (Theorem 2). Consider the Markov process \mathcal{M}^+ consisting of running \mathcal{M} , reinitializing \mathcal{M} every τ^* periods. By the stacking theorem (Theorem 4), \mathcal{M}^+ is ergodic, so a sufficiently large sample of historical averages starting τ^* periods after reinitialization and continuing until the next initialization will reveal the stationary distribution of \mathcal{M}^+ . This is the methodology we will use in estimating the aggregate properties of a Markov model of a market economy.

6 A Mathematical Model of Decentralized Market Exchange

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 has become the standard model of market exchange in economic theory. Such economies are particularly attractive because they capture the basic structural characteristics of market economies, and because a Walrasian equilibrium always entails an efficient allocation of resources, abstracting from certain conditions that require non-market regulation, such as externalities (e.g., air pollution), natural monopolies (such as the supply of water in a city), and public goods (such as national defense). Bator (1957), Katzner (1992), and Ellickson (1994) provide accessible introductions to the Walrasian model for mathematically sophisticated readers.

The Walrasian economy consists of households and firms. Firms acquire production inputs (labor, raw materials, and capital) at given market prices and combine them to produce outputs (goods and services) which they sell at given market prices to the households. Labor and capital inputs are rented from the households that own them, so the firm itself owns nothing—all wealth is in the hands of households. The inputs also include raw materials and the outputs of other firms, which firms also purchase on markets. The labor and raw material inputs, as well as shares of the net profit of the firms, are owned by the households, and form their wealth.

Households buy the output of the various firms, some of which they consume, and some of which they add to their stock of wealth. The economy is in equilibrium when the vector of prices p for all goods and services is set so that the supply equals the demand in each market.

Consider the following highly simplified Walrasian economy. There are two goods, which we will call apples (a) and nuts (n), two households (x and y), and two factor inputs, labor (l) and capital (k). Suppose l_x and l_y are the amounts of labor owned by households x and y , l_a and l_n are the amounts of labor used in producing apples and nuts, k_x and k_y are the amounts of capital owned by households x and y , and k_a and k_n are the amounts of capital used in producing apples and nuts. Then we have the equations

$$l_a + l_n = l_x + l_y \quad (20)$$

$$k_a + k_n = k_x + k_y \quad (21)$$

The equations say that the total amount of labor and capital demanded by firms to use in production equals the total amount supplied by households.

Now suppose the wage rate is w and the interest rate (which is the rental price of capital) is r . Also, suppose the price of apples is p_a , the price of nuts is p_n , and household x consumes a_x of apples and n_x of nuts, while household y consumes a_y of apples and n_y of nuts. Finally, we assume household x owns a share α_x^a of the profits of the apple firm and α_x^n of the nut firm, while household y owns a share

α_y^a of the profits of the apple firm and α_y^n of the nut firm, where

$$\alpha_x^a + \alpha_y^a = \alpha_x^n + \alpha_y^n = 1 \quad (22)$$

because all profits are distributed to the owners.

Then if m_x and m_y are the incomes of households x and y , we have the following two equations.

$$m_x = \alpha_x^a \pi_a + \alpha_x^n \pi_n + wl_x + rk_x \quad (23)$$

$$m_y = \alpha_y^a \pi_a + \alpha_y^n \pi_n + wl_y + rk_y. \quad (24)$$

where π_a and π_n are the profits generated in the two firms.

The next two equations are production functions for the apples and nuts firms. They say that each good is produced by using capital and labor.

$$f(l_a, k_a) = a_x + a_y \quad (25)$$

$$g(l_n, k_n) = n_x + n_y \quad (26)$$

We assume the firm maximizes profits, given by

$$\pi_a = p_a f(l_a, k_a) - (wl_a + rk_a) \quad (27)$$

$$\pi_n = p_n g(l_n, k_n) - (wl_n + rk_n) \quad (28)$$

Choosing inputs to maximize profits (that is, setting the partial derivatives of π_a and π_n to zero and solving the resulting equations), we find

$$f_{l_a} = \frac{w}{p_a} \quad f_{k_a} = \frac{r}{p_a} \quad (29)$$

$$g_{l_n} = \frac{w}{p_n} \quad g_{k_n} = \frac{r}{p_n} \quad (30)$$

where subscripts on the production functions f and g refer to partial derivatives.

We assume households have utility function $u^x(a_x, n_x)$ and $u^y(a_y, n_y)$, which they maximize subject to their income constraints. The income constraints are given by

$$m_x = p_a a_x + p_n n_x \quad (31)$$

$$m_y = p_a a_y + p_n n_y, \quad (32)$$

and households maximize $u^x(a_x, n_x)$ and $u^y(a_y, n_y)$ subject to these constraints. Maximizing utility given these income constraints gives

$$\frac{u_a^x}{p_a} = \frac{u_n^x}{p_n} \quad (33)$$

$$\frac{u_a^y}{p_a} = \frac{u_n^y}{p_n}, \quad (34)$$

where subscripts on u^x and u^y represent partial derivatives. Finally, we can take one price as numeraire, since both firms and households care only about price ratios (there is no money or monetary wealth in this model).

$$p_a = 1. \quad (35)$$

In this system, $\alpha_x^a, \alpha_x^n, l_x, l_y, k_x,$ and k_y are independent parameters representing the structure of ownership in the economy. There remain eighteen variables to be determined: $\alpha_y^a, \alpha_y^n, l_a, l_n, k_a, k_n, a_x, a_y, n_x, n_y, \pi_a, \pi_n, m_x, m_y, p_a, p_n, w,$ and r . There are ostensibly nineteen equations (equations (29), (30), and (22) are each two separate equations). However, if we add up (27) and (28), then substitute in (25) and (26), we get the sum of (23) and (24). This means one equation is superfluous, giving us a system of sixteen equations in sixteen unknowns.

A *market equilibrium* is a solution to these equations with non-negative prices and quantities. If these equations were linear, this system would then generically have exactly one solution, although even then the non-negativity conditions would not be guaranteed. But of course these equations are not linear, so there could be zero or many solutions. Thus proving the existence of a market equilibrium is non-trivial. An existence proof 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), Arrow and Hahn (1971) and others provided rather complete proofs of the existence of equilibrium in Walrasian economies very similar to the model above, but with arbitrary finite numbers of goods, households, firms, capital goods, types of labor, and raw materials. The conditions that make this possible, roughly speaking, are that consumers have concave preferences (declining marginal utility) and firms have convex production functions (declining marginal productivity).

7 Walrasian Tâtonnement Price Dynamics

We live in a period of financial fragility exacerbated by increasingly interdependent national economies, thus transforming local into global instability. Standard macroeconomic models cannot address these problems because they are highly aggregated “representative agent” versions of the Walrasian model that do not treat the interaction of heterogeneous agents (Kirman 1992, Guilmi et al. 2012). We must thus go beyond the currently static Walrasian model showing that under a plausible dynamic exhibiting the conditions under which the general equilibrium model is locally and globally stable. The only dynamics that have been proposed follow Walras’ tâtonnement analysis, which assumes that the adjustment process out of equilibrium is carried out on a centralized institutional level beyond the control of economic agents. Walras assumed there is an “auctioneer” (like the famous

crieur of the Paris *Bourse*), who calls out prices, calculates excess demand or supply for each good, based on the reaction of firms and households, and adjusts the price by lowering the price of a good in excess supply and raising the price of a good in excess demand. The auctioneer repeats this process of “tâtonnement” until prices clear all markets, at which time the auctioneer broadcasts this equilibrium price vector, and permits firms and household to make plans based on it. Of course, Walras could not show that this process converges, despite its intuitive plausibility.

Twentieth century economists were not satisfied with Walras’ nineteenth century argument. The question of stability of the Walrasian economy thus became 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 conformed to Walras’ tâtonnement framework, assuming that out of equilibrium, there is a system of public prices shared by all agents. This researchers implemented an ordinary differential equation dynamic, in which the time rate of change of price for each good is a function of excess demand for that good. They had only extremely limited success, even in this highly implausible stability framework.

This quest for a general stability theorem was derailed by Herbert Scarf’s 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 imposed on the aggregate dynamics of the system. These conditions cannot be derived from the underlying market structure or the behavior of producers and consumers (Kirman 1992), and chaos in price movements is the generic case for the tâtonnement adjustment processes (Saari 1985, Bala and Majumdar 1992).

General equilibrium theorists in the early 1970’s harbored some expectation that plausible restrictions on utility functions might entail stability. However, Sonnenschein (1973), Mantel (1974, 1976), and Debreu (1974) showed that any virtually any continuous function is the excess demand function for some Walrasian economy. 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.

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. This remains the case today. Despite the centrality of the general equilibrium model to economic

theory, little about real-world market dynamics in realistic settings can be inferred from the differentiable manifold approach.

Even if a well-behaved centralized institutional price adjustment had been found, this would not have solve the problem of real-world economic dynamics. This is true even in principle and even accepting a highly idealized framework. This is because in the approach taken to stability in twentieth century general equilibrium theory, a single price vector is assumed to govern all agents. We call this the *public prices* assumption. This assumption precludes the attainment of stability through a price-adjustment processes in the hands of households and firms. No Walrasian agent can have any affect whatsoever on any price. Indeed, this is the very definition of a “competitive” Walrasian equilibrium. Therefore the is no possible revision of the tâtonnement process that operates a the level of decentralized individual decision-making.

Only a dynamic based on agents adjusting prices to maximize profits or utility while producing and consuming out of equilibrium can be a plausible candidate for a general Walrasian stability theory. I will show that a Markov process model of market dynamics can succeed where a differential equation approach has failed, although because of the large size of the resulting Markov model, in our current, and highly incomplete, state of knowledge of Markov processes, its properties can be ascertained only through computer simulation.

There is surprisingly little precedent for this research strategy. There are several agent-based models described in the economics literature (Tesfatsion and Judd 2006), but none deals with several sectors, disaggregated to the level of heterogeneous agents. Chen et al. (2000) use Swarm software (Stefansson 1997) to model a one-sector economy using labor, the only endogenous variable being the single product price. Epstein and Axtell (1997) model a highly complex multi-agent system, but have one good and no firms. Lian and Plott (1998) implement a laboratory experiment with human subjects, with one good, labor, and *fiat* money, and Anderson et al. (2004) implement a similar experiment with three goods and three consumers with Leontief consumption functions, based on Scarf (1960). Weddepohl (1997) simulates a global tâtonnement process in an economy with two goods plus labor, one firm producing each good, and three consumers who also supply a fixed amount of labor. Other than these papers, I have been able to find only examples of sector-level models that assume aggregate demand function. such as Taylor and Uhlig (1990). There are several quite elegant models of bilateral trade in durable goods (that is, there is no consumption until the trade process is completed), without assuming public prices, leading to market-clearing Pareto-optimal allocations (Smale 1976, Friedman 1979, Foley 1994, Axtell 2005). These models, however, assume that agents trade only in goods that they consume plus perhaps money. An acceptable model of real-world market exchange must assume that agents consume

out of equilibrium and trade in goods that they do not consume.

My goal here is to present Markov process models of two highly simplified market economies. The first, a highly simplified Markov process model of market dynamics with private prices, implements the Scarf (1960) Walrasian model. While this model exhibits limit cycles under the tâtonnement dynamic, the Markov process implementation is ergodic, and hence has a stationary distribution. Estimates of this stationary distribution under a variety of system parameters indicate that this stationary distribution is close to the market-clearing Walrasian equilibrium. I hypothesize that the stationary distribution is exactly the market-clearing Walrasian equilibrium, although I cannot prove this. Moreover, this result does not depend on Scarf's peculiar choice of Leontieff utility functions, but applies widely to constant elasticity of substitution utility functions as well. Similarly, the result does not depend on Scarf's particular choice of what goods individuals consume, but rather applies to all situations in which individuals produce one good and consume both this good and one other good.

A notable characteristic of this model is that there are continual short-run excursions from market-clearing prices, on the order of 20% or more from equilibrium prices. These are not generated by macro-level shocks because no such shocks are included in system dynamics. Rather, they are produced by the completely disaggregated process whereby agents periodically copy the behavior of more successful agents, under conditions of random matching of agents into trading partners.

The second model, which is a generalization of the above to an economy with many goods, agents who consume several goods, and the potential for the emergence of one or more goods serving as a universal medium of exchange (money), exhibits similar long-run dynamics. We thus find that in general, ergodic Markov process economies with simple production technologies have stationary distributions close to market-clearing Walrasian equilibrium, with significant short-run excursions away from equilibrium generated by the inherent stochasticity of the Markov process. Proving that the stationary equilibrium is exactly the Walrasian equilibrium, or finding a counterexample, remains to be accomplished.

8 A Decentralized Price Dynamic for the Scarfian Economy

This section revisits Herbert Scarf's (1960) seminal example of Walrasian instability assuming public prices and tâtonnement dynamics, showing that instability holds in the Markov process version of the model as well. 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

$$u_X(x, y, z) = \min\{x, y\}, \quad (36)$$

$$u_Y(x, y, z) = \min\{y, z\}, \quad (37)$$

$$u_Z(x, y, z) = \min\{z, x\}. \quad (38)$$

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

$$x_X^d = y_X^d = \frac{p_x}{p_x + p_y} \quad (39)$$

$$y_Y^d = z_Y^d = \frac{p_y}{p_y + p_z} \quad (40)$$

$$z_Z^d = x_Z^d = \frac{p_z}{p_z + p_x} \quad (41)$$

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

$$E_x = x_X^d + x_Z^d - 1 = \frac{-p_y}{p_x + p_y} + \frac{p_z}{p_z + p_x} \quad (42)$$

$$E_y = y_X^d + y_Y^d - 1 = \frac{-p_z}{p_y + p_z} + \frac{p_x}{p_x + p_y} \quad (43)$$

$$E_z = z_Y^d + z_Z^d - 1 = \frac{-p_x}{p_z + p_x} + \frac{p_y}{p_y + p_z}. \quad (44)$$

The tâtonnement price adjustment process is given by

$$\dot{p}_i = E_i(p_x, p_y, p_z), \quad \text{where } i = x, y, z. \quad (45)$$

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.

For the Markov process version of the Scarf economy, we maintain the above assumptions, except now we assume 500 agents of each of the three types. We drop the assumption that out of equilibrium there is a common price system accepted by

all agents. Instead, we assume each agent has his personal subjective prior over the appropriate price structure, which we call the agent's *private prices*. These prices are, except occasionally when one agent copies the price vector of another agent, the private information of their owner. The agent uses his private price vector to set his consumption goals by maximizing utility subject to these prices, and to determine his trading strategy, according to which he offers exchanges at rates given by his price vector, and he accepts offers only if the value of what he receives is at least as great as the value of what he gives up, all evaluated at his private price vector. Each agent is endowed at time $t = 0$ with a set of private prices randomly drawn from a uniform distribution.

This is a convenient place to explain why we used complete random initialization of all our Markov models. According to the Stacking Theorem (Theorem 4), if we repeatedly run an instance of a finite Markov chain with random initialization, or indeed any initialization that places positive probability on all states, then the stochastic process consisting of a finite series of such instances forms an ergodic Markov process. It follows that this Markov process has a unique stationary distribution. Therefore the frequency of particular states in the history of the process represents the probability of this state in the stationary distribution.

We allow 50000 generations and ten 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.

Note that this model is indeed a finite Markov process, where a state of the system is a vector of states of individual agents, the state of an agent being his private price structure, and his current utility. Each price, as well as utility, is a floating point number, which has a finite range on any computer platform. Thus there is a finite, albeit extremely large, state space for the process. Note that whether the random generators used in the computer realization of the model may be either generate true random number (such as the timing of a Beta decay) or may use a deterministic algorithm with an approximation to the behavior of a uniform random variable. The second case is clearly a finite process, and the first case has finite range because the output of the random number generator is measured as a computer floating point number.

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 fraction of his inventory 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 he gets exceeds the value of what he gives up, according to his private prices, then he agrees to trade. If the responder has less of the second good than the initiator wants, the trade is scaled down proportionally. Agents are thus rational maximizers, where an agent's subjective priors is his vector of private prices.

Which good an initiator 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 but can exchange in trade the T-good. Note that an agent must be willing to acquire his T-good despite the fact that it does not enter into his utility function. This is because X-producers want Y, but Y-producers do not want X. Only Z-producers want X. Therefore there can be no direct barter trade between consumption and production goods in this economy. Because 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 his 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 terms are dictated by the initiator and the amount is the maximum compatible with the inventories of the initiator and responder.

Figure 1 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 quite small. Quasi-public prices are the closest the Markov process comes to approximating the public prices of standard Walrasian general equilibrium theory. Such prices are "almost public," but an agent is free to change the prices he uses at will.

The Markov dynamic in this case is a stationary distribution depicted in Figure 2. It is clear that mean long-run behavior of prices in the system is close to

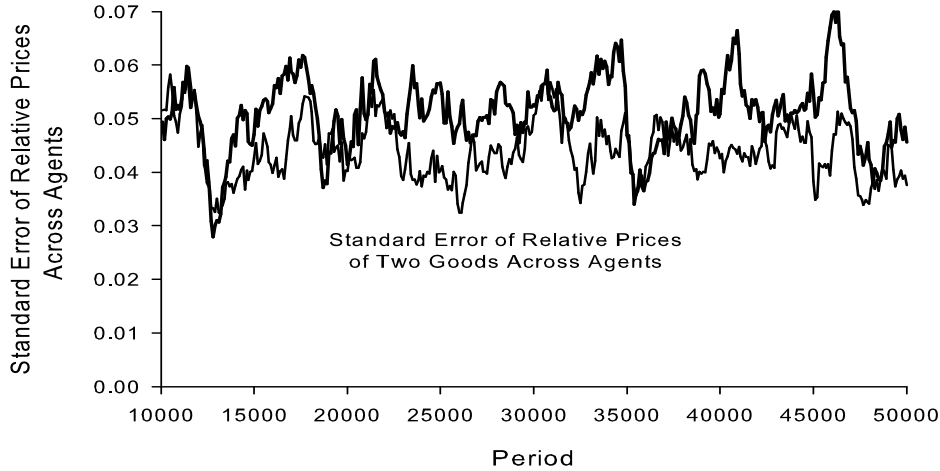


Figure 1: The Markov process version of the Scarf economy initialized with random prices quickly transitions to quasi-public prices (private prices with low standard error across agents).

the market-clearing Walrasian equilibrium. This is the only behavior of the stationary distribution observed, independent of the initial state of the system, so it is apparently a characteristic of the stationary distribution.

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.

The assumptions behind this and other Markov models of market interaction are rather generic in the following sense. The assumption that each agent has a set of private prices is quite general. Of course, in complex institutional economics, some prices might be set through an extra-market process, as when the central bank sets an interest rate, or a government agency sets the price of a service. In these cases, all agents will have the same entry in their private price vectors for these prices. The assumption that individuals use their private prices to form individual demand and supply for each good by maximizing a utility function follows from the assumption that agents are rational maximizers. The assumption that individuals will accept a trade if they gain more than they give up according to their private prices is reasonable, but can be proven from a few more basic assumptions, the

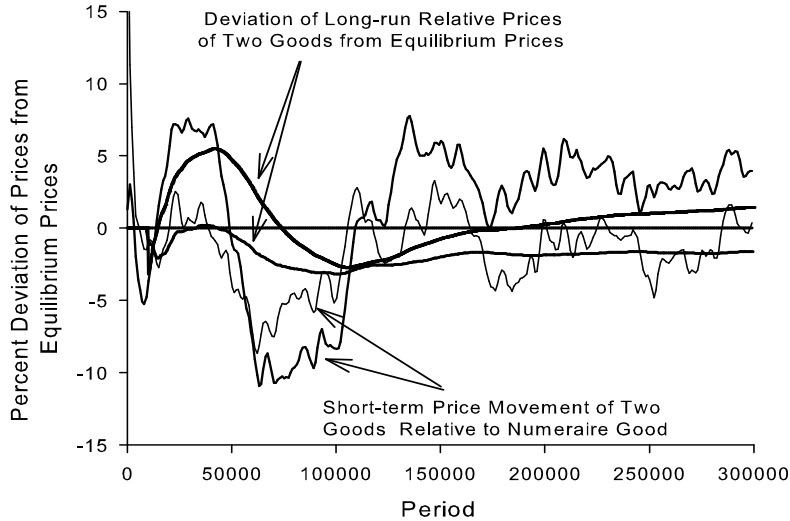


Figure 2: The Markov process Scarf economy with private prices exhibits a stationary distribution close to Walrasian market clearing.

most important being (a) if one set of terms of trade is acceptable, then any other terms of trade that improves the benefit-cost ratio for the agent is acceptable; and (b) if two exchanges are acceptable, then any linear combination of exchanges is also acceptable. The production conditions for our economies are extremely primitive (there is no labor or capital), but this assumption can be replaced by more realistic ones. The result is a more complex dynamic with many of the same properties, and some others as well. See Gintis (2007) for details.

The equilibrium of the Markov process is a stationary distribution that can be analytically specified in principle, but in practice is many 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. The extended ergodic theorem (Theorem 2), the stacking theorem (Theorem 4) and the aggregation theorem (Theorem 5) together ensure that this process has a stationary distribution that can be empirically estimated. 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 cautious conclusion we can offer is that we have proven, for the first time, that there exist Scarfian economies with private prices and a decentralized dynamical system of price adjustment without tâtonnement that in the long run approximates the unique equilibrium price vector.

9 A Multi-Good Market Economy with Simple Production

We will now shift our attention to a pure market economy with many goods, many agents and only one institution—the marketplace. We assume each agent produces one good, in fixed amount, using only personal labor, but consumes a variety of goods. Agents are endowed with private prices, and they have no information about the economy other than that gathered from private experience in trade. These assumptions lack realism, but each can be replaced by more realistic assumptions when necessary. For instance, we could add a labor market, a capital market, firms, and a central bank without difficulty (Gintis 2007).

We assume there are n sectors. Sector $k = 1, \dots, n$ produces good k in “styles” $s = 1, \dots, m$. We use “styles” to ensure that our Markov process economies have large numbers of goods, but only a limited number of prices. This is for the sake of computational convenience only, and does not affect the dynamics of the model. 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 produced in sector k with style s , 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 $G = \{g_s^k | k = 1, \dots, n, s = 1, \dots, m\}$. 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 personal labor and no other inputs an amount q_k of good g_s^k which depreciates to zero if it remains in inventory at the close of a trading period. A trade inventory for an agent can include any good acquired through production or trade.

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 agents. Each of these agents is assigned a private price vector by choosing each price from a uniform distribution on the open unit interval, 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 rendered unique by randomly setting several parameters of a hybrid CES utility function, explained in Appendix A. These utility functions, which are generalizations of a functional form widely used in economic models, do not satisfy the gross substitutability assumption (Arrow, Block and Hurwitz 1959), so stability in the tâtonnement dynamic generally does not hold.

For each good $g_s^k \in G$ there is a *market* $m[k, s]$ of agents who sell good g_s^k . In each period, the agents in the economy are randomly ordered and are permitted

one-by-one to engage in active trading. When the g_s^h -agent A is the current active agent, for each good g_t^h for which A has positive demand, 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 $y_k \leq i_k^A$, where 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 ; that is, he offers B an equivalent value of good g_s^k , the valuation being at A's prices. B accepts this offer provided the exchange is weakly profitable at B's private prices; that is, 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 . The code for this trading algorithm is rather complex, but the interested reader can download it from my website: <http://people.umass.edu/gintis>.¹ 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.

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 agent 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 (that is, if he failed to purchase any of the consumption good or sell all of his production good), and lowering price by 0.05% otherwise (that is, 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 agents is updated using a standard replicator dynamic, in which agents who have high scores for trading and consuming have a high probability of reproducing, while unsuccessful trades are eliminated from the economy. In all cases, the new agents inherit the price vector of its parent, perhaps mutated a bit. The resulting updating process is a discrete approximation of a monotonic dynamic in evolutionary game theory. In differential equation systems, all monotonic dynamics have the same properties as the simplest, which is the replicator dynamic (Taylor and Jonker 1978, Samuelson and Zhang 1992). Other monotonic approximations, including choosing a pair of

¹For related analyses of trading algorithms in a multiple market setting, see Rubinstein and Wolinsky (1985), Gale (1987), and Binmore and Herrero (1988).

agents in $m[k, s]$ and let the lower-scoring agent copy the higher-scoring agent, produce similar dynamical results.

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.

10 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, \dots, g_6^9 , and 16200 agents in the economy. There are then nine distinct prices p_1^A, \dots, p_9^A for each agent A. We treat g^9 as the numeraire good for each agent, so $p_9^A = 1$ for all agents 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 $\epsilon = 0.01$.

The results of a typical run of this model is illustrated in Figures 3 and 4. Figure 3 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 3 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 its initial value of 6.7.

Figure 4 shows the movement 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




Figure 3: Convergence of private prices to quasi-public prices in a typical run with nine goods in 6 styles each (fifty-four goods).

another order of magnitude by the end of the run.

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.

11 Conclusion

Agent-based models have been interpreted either as a way to study the behavior of mathematical models that have no closed-form solution, or as an alternative to mathematical modeling altogether (Tsefatson and Judd 2006). This paper develops a third interpretation, and applies it to a long-standing problem in economic theory: characterizing the dynamical property of general market exchange.




Figure 4: The path of aggregate excess demand over 50,000 periods.

In this alternative, an agent-based model is a well-known mathematical structure, the finite Markov process. I develop several theorems, including an expanded ergodic theorem (Theorem 1), a finite history theorem (Theorem 3), a stacking theorem (Theorem 4), and an aggregation theorem (Theorem 5), all of which relate to the construction and estimation of finite Markov process, and when appropriately applied, ensure that the resulting Markov process is ergodic. This implies that the process has a stationary distribution, and that we can empirically estimate this distribution.

I then showed that agent-based models of social processes can generally be treated as finite Markov processes with analytically determinate stationary distributions that be estimated using empirical sampling techniques. I applied this specifically to a highly simplified generalized exchange system that Scarf (1960) showed had a unique but unstable market clearing equilibrium under the tâtonnement public price dynamic with limit cycles. I showed that the Markov process translation with private prices and many heterogeneous agents has a stationary distribution with quasi-public prices that approximates Walrasian market-clearing in the long run. I then turned to a decentralized system of barter exchange with rather generic properties and heterogeneous agents with hybrid constant elasticity of substitution (CES) utility functions (see Appendix A), so the classical conditions for stabil-

ity of equilibrium (gross substitutability) do not obtain (Arrow et al. 1959). The Markov process dynamic for this economy swiftly converges to quasi-public prices and a quasi-Walrasian equilibrium. These results are not due to the simplicity of the economy, because similar results obtain with a considerably more institutionalized economy with firms, labor and credit markets, and a government sector that taxes and provides unemployment insurance (Gintis 2007). My experience with constructing generalized competitive exchange models in different contexts leads me to believe this is the general case, although I cannot prove this (Gintis 2012)

Like many mathematical constructs, the Markov process is simple to define, but the complexities of Markovian dynamics are little understood. Markov processes with analytically determinate and very small transition matrices can be solved using well-known techniques (Gintis 2009, Ch.13). But the Markov processes arising modeling complex biological and social phenomena have immense state spaces and their transition matrices are only implicitly defined by the detailed behavior of the model. This gives rise to two questions that we would like to see addressed in the future.

First, what are the statistical properties of very large Markov processes? How can we ascertain when enough data has been collected to estimate with some confidence the stationary distribution of the process. In practice, we must settle for running the Markov process until it appears to have steadied down to some constant or cyclically varying behavior. The theoretical question is: when is this practice justified?

Second, can we make strategic simplifications of a very large Markov process that is analytical tractable and closely approximates the behavior of the process? A possible strategy is to use a statistical mechanics approach. An analogy with the statistical mechanics of ideal gases might prove useful. First, a continuous process is approximated by a finite process (this step is unnecessary in our case), then intractable binomial distributions are approximated by exponential distributions using the Stirling approximation. The resulting system is amenable to analytical treatment. Whether such a strategy can be successful in the case of large Markov processes remains to

12 Appendix A: Constructing Utility Functions

The utility function of each agent is the product of powers of CES utility functions of the following form. For each consumer, we partition the n consumer goods into k segments (k is chosen randomly from $1 \dots n/2$) of randomly chosen sizes m_1, \dots, m_k , with $m_j > 1$ 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 elasticity. 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 consumers have the same utility function.

For example, consider a segment using goods x_1, \dots, x_m with private prices p_1, \dots, p_m and 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

$$u(x_1, \dots, x_n) = \left(\sum_{l=1}^m \alpha_l x_l^\gamma \right)^{1/\gamma}, \quad (46)$$

where $\gamma = (s - 1)/s$, and $\alpha_1, \dots, \alpha_m > 0$ 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 $\gamma \neq 0$ (that is, the utility function segment is not Cobb-Douglas), this gives

$$x_i = \frac{M f_i}{\sum_{l=1}^m p_l \phi_{il}^{1/(1-\gamma)}}, \quad (47)$$

where

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

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

$$x_i = \frac{M f_i \alpha_i}{p_i}. \quad (48)$$

The utility function of each agent is the product of powers of CES utility functions of the following form. For each consumer, we partition the n consumer goods into k segments (k is chosen randomly from $1 \dots n/2$) of randomly chosen sizes m_1, \dots, m_k , with $m_j > 1$ 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 a uniform distribution. 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 consumers have the same utility function.

For example, consider a segment using goods x_1, \dots, x_m with prices p_1, \dots, 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

$$u(x_1, \dots, x_n) = \left(\sum_{l=1}^m \alpha_l x_l^\gamma \right)^{1/\gamma}, \quad (49)$$

where $\gamma = (s - 1)/s$, and $\alpha_1, \dots, \alpha_m > 0$ 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 $\gamma \neq 0$ (that is, the utility function segment is not Cobb-Douglas), this gives

$$x_i = \frac{M f_i}{\sum_{l=1}^m p_l \phi_{il}^{1/(1-\gamma)}}, \quad (50)$$

where

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

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

$$x_i = \frac{M f_i \alpha_i}{p_i}. \quad (51)$$

REFERENCES

- Anderson, Christopher M., Charles L. Plott, Ken-Ichi Shimomura, and Sander Granat, "Global Instability in Experimental General Equilibrium: The Scarf Example," *Journal of Economic Theory* 115 (2004):209–249.
- 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: Nonlinear 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. and Frank Hahn, *General Competitive Analysis* (San Francisco: Holden-Day, 1971).

- and Gérard Debreu, “Existence of an Equilibrium for a Competitive Economy,” *Econometrica* 22,3 (1954):265–290.
- Axtell, Robert, “The Complexity of Exchange,” *Economic Journal* 115,504 (2005):F193–F210.
- Bala, V. and M. Majumdar, “Chaotic Tatonnement,” *Economic Theory* 2 (1992):437–445.
- Bator, Francis, “The Simple Analytics of Welfare Maximization,” *American Economic Review* 47 (1957):22–59.
- Binmore, Kenneth and M. J. Herrero, “Matching and Bargaining in Dynamic Markets,” *Review of Economic Studies* LV (1988):17–31.
- Chen, Yu, Xiaoming Su, Jinghai Rao, and Hui Xiong, “Agent-Based Microsimulation of Economy from a Complexity Perspective,” 2000. IFIP World Computer Congress, Beijing, China.
- Debreu, Gérard, “A Social Equilibrium Existence Theorem,” *Proceedings of the National Academy of Sciences* 38 (1952):886–893.
- Debreu, Gerard, “Excess Demand Function,” *Journal of Mathematical Economics* 1 (1974):15–23.
- Ellickson, Brian, *Competitive Equilibrium: Theory and Applications* (Cambridge: Cambridge University Press, 1994).
- Epstein, Joshua and Robert Axtell, *Sugarscape* (Cambridge: MIT Press, 1997).
- 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).
- Foley, Duncan, “A Statistical Equilibrium Theory of Markets,” *Journal of Economic Theory* 62,2 (April 1994):321–345.
- Friedman, Daniel, “Money-mediated Disequilibrium Processes in a Pure Exchange Economy,” *Journal of Mathematical Economics* 6 (1979):149–167.
- Gale, David, “The Law of Supply and Demand,” *Math. Scand.* 30 (1955):155–169.
- Gale, Douglas, “Limit Theorems for Markets with Sequential Bargaining,” *Journal of Economic Theory* 43 (1987):20–54.
- Gintis, Herbert, “The Dynamics of General Equilibrium,” *Economic Journal* 117 (October 2007):1289–1309.
- , *Game Theory Evolving* Second Edition (Princeton: Princeton University Press, 2009).
- , “The Dynamics of Pure Market Exchange,” in Herbert Gintis (ed.) *Approaches to the Evolving World Economy: Complex Dynamics, Norms, and Organizations* (London: Palgrave, 2012).

- Guilmi, Corrado Di, Mauro Gallegati, Simone Landini, and Joseph E. Stiglitz, "Towards an Analytical Solution for Agent Based Models: An Application to a Credit Network Economy," in Herbert Gintis (ed.) *Approaches to the Evolving World Economy: Complex Dynamics, Norms, and Organizations* (London: Palgrave, 2012).
- Hahn, Frank, "A Stable Adjustment Process for a Competitive Economy," *Review of Economic Studies* 29 (1962):62–65.
- and Takashi Negishi, "A Theorem on Non-Tâtonnement Stability," *Econometrica* 30 (1962):463–469.
- Hofbauer, Josef and Karl Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge: Cambridge University Press, 1998).
- Katzner, Donald, *The Walrasian Vision of the Microeconomy : an Elementary Exposition of the Structure of Modern General Equilibrium Theory* (Ann Arbor: University of Michigan Press, 1992).
- Kirman, Alan P., "Whom or What does the Representative Individual Represent?," *Journal of Economic Perspectives* 6 (1992):117–136.
- Lian, Peng and Charles R. Plott, "General Equilibrium, Markets, Macroeconomics and Money in a Laboratory Experimental Environment," *Economic Theory* 12 (1998):21–75.
- 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.
- Morishima, Michio, "A Reconsideration of the Walras-Cassel-Leontief Model 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 Equilibrium for 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.
- Rubinstein, Ariel and Asher Wolinsky, “Equilibrium in a Market with Sequential Bargaining,” *Econometrica* 53 (1985):1133–1150.
- 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.
- Smale, Stephen, “Exchange Processes with Price Adjustment,” *Journal of Mathematical Economics* 3 (1976):211–226.
- Smith, Adam, *The Theory of Moral Sentiments* (New York: Prometheus, 2000[1759]).
- Sonnenschein, Hugo, “Do Walras’ Identity and Continuity Characterize the Class of Community Excess Demand Functions?,” *Journal of Economic Theory* 6 (1973):345–354.
- Stefansson, Benedikt, *Swarm: An Object Oriented Simulation Platform Applied to Markets and Organizations* (New York: Springer-Verlag, 1997).
- Taylor, J. and H. Uhlig, “Solving Nonlinear Stochastic Growth Models: A Comparison of Alternative Methods,” *Journal of Business and Economic Statistics* 8 (1990):1–19.
- Taylor, Peter and Leo Jonker, “Evolutionarily Stable Strategies and Game Dynamics,” *Mathematical Biosciences* 40 (1978):145–156.
- Tesfatsion, Leigh and Kenneth L. Judd, *Handbook of Computational Economics II: Agent-Based Computational Economics* (Amsterdam: Elsevier/North-Holland, 2006).
- 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]).

Weddepohl, Claus, “Simulating Tatonnement in a Production Economy,” *Economic Theory* 10 (1997):551–558.

c:\papers\Stability of Market Exchange\ACM-TIST-Markov.tex March 28, 2012