

The Dynamics of Walrasian General Equilibrium: Theory and Application

Herbert Gintis and Antoine Mandel*

December 10, 2012

Abstract

This paper reviews recent progress in treating Walrasian general equilibrium as a complex dynamical system that can be analyzed using evolutionary game theory. We model the general equilibrium economy as the stage game of an evolutionary game in which the strategy of each agent is a vector of prices determining the agent's demand and supply, as well as the set of trades the agent deems acceptable. Gintis and Mandel (2012) have shown that, under mild conditions, the stable equilibria of the resulting dynamical system under a replicator dynamic are the same as the equilibria of the Walrasian economy. For sufficiently large population size, the resulting dynamical system can be arbitrarily closely approximated by a finite Markov process. The stationary distribution of this Markov process thus approximates a Walrasian equilibrium. The dynamics of this model for particular ranges of parameter values can be studied through computer simulation.

1 Introduction

Walras (1954 [1874]) developed a general model of competitive market exchange. A *Walrasian equilibrium* for this model is a set of prices such that supply is determined by all firms maximizing profits, demand is determined by all households maximizing utility subject to a budget constraint given by the value of their endowments, and excess demand for all goods is zero. Walras provided an informal argument for the existence of a Walrasian equilibrium. Wald (1951 [1936]) provided an existence proof for a simplified version of Walras' model. Inspired by John Nash's (1950) proof of the existence of Nash equilibrium for finite games

*Gintis: Santa Fe Institute and Central European University; Mandel: University Paris 1 Panthéon-Sorbonne. For publication in *Complexity Economics*

(Feiwel 1987), Debreu (1952) and Arrow and Debreu (1954) supplied a more general proof, with further contributions by Gale (1955), Nikaido (1956), McKenzie (1959), Negishi (1960), and others. Arrow and Hahn (1971), Mas-Colell et al. (1995), and Florenzano (2005) offer a summary of this and related work on the equilibrium properties of the general equilibrium model.

The stability of the Walrasian economy was a central research focus in the years following the existence proofs (Arrow and Hurwicz 1958, 1959, 1960; Arrow, Block and Hurwicz 1959; Nikaido 1959; McKenzie 1960; Nikaido and Uzawa 1960). Inspired by Walras' tâtonnement process, these models assumed that there is no production or trade until equilibrium prices are attained, and out of equilibrium, there is a price vector shared by all agents, the time rate of change of which is a function of excess demand. These efforts at proving stability were unsuccessful (Fisher 1983). Indeed, Scarf (1960) and Gale (1963) provided simple examples of unstable Walrasian equilibria under a tâtonnement dynamic. Moreover, Sonnenschein (1973), Mantel (1974, 1976), and Debreu (1974) showed that any continuous function, homogeneous of degree zero in prices, and satisfying Walras' Law, is the excess demand function for some Walrasian economy. These results showed that no general stability theorem could be obtained based on the tâtonnement process. Subsequent analysis showed that chaos in price movements is the generic case for the tâtonnement adjustment processes (Saari 1985, Bala and Majumdar 1992).

Several researchers in the early 1960's explored the possibility that allowing trading out of equilibrium could sharpen stability theorems (Uzawa 1959, 1961, 1962; Negishi 1961; Hahn 1962; Hahn and Negishi 1962), and especially Fisher (1970, 1972, 1973, 1983). These efforts, however, did not result in any general stability results relevant to a decentralized market economy.

A novel approach to the dynamics of large-scale social systems, evolutionary game theory, was initiated by Maynard Smith and Price (1973), and adapted to dynamical systems theory in subsequent years (Taylor and Jonker 1978, Friedman 1991, Weibull 1995). The application of these models to economics involved the shift from biological reproduction to behavioral imitation as the criterion for the replication of successful agents.

In Gintis and Mandel (2012), we applied this framework by treating the Walrasian economy as the stage game of an evolutionary game. The stage game assumes each agent is endowed with one unit of a good that he must trade to obtain the various goods he consumes. We will call this good the agent's *production good*, and we will refer to the agent as the *producer* of this good. An agent's trade strategy consists of a set of *private prices* for his endowment and the goods he consumes, such that, according to these private prices, demand and supply are determined by utility maximization, and a trade is acceptable if the value of goods received is at

least as great as the value of the goods offered in exchange. The stage game includes an *exchange process* mapping initial endowments and agent strategy profiles into final allocations. . We assume that the strategies of relatively successful agents in the stage game are occasionally copied by less successful agents with the same production good. The Walrasian economy thus becomes a *multipopulation game*, where all producers of the same good form a single population (Weibull 1995). With rather mild assumptions, the stability of equilibrium is then guaranteed.

2 The Stability of Walrasian Equilibrium

We consider an economy \mathcal{W} with a finite number of goods and a finite number of agents. Agents have continuous, strictly concave, and locally non-saturated utility functions defined over some subset of goods, and an initial endowment consisting of one unit of one of the goods, which we term his *production good*. We assume an agent can achieve a positive level of utility only by trading this good for other goods that he consumes. Each agent also has a set of private prices of the goods in the economy. He uses these prices to maximize utility, thus determining his demand and supply of various goods.

We say an allocation of goods to the agents in the economy \mathcal{W} is *feasible* if it is a reallocation of the initial endowments among the agents. An *exchange mechanism* ξ associates a feasible allocation with each profile of private prices of agents in the economy \mathcal{W} . We interpret this allocation as the result of trades among agents, where a bilateral trade is acceptable to both parties only if the trade is (weakly) profitable according to both agents' private prices. The payoff to a player subject to a particular exchange mechanism is the utility of the agent's allocation as determined by this exchange mechanism and a particular profile of private prices.

We say an exchange mechanism ξ is *proper* if first, if all agents choose the same Walrasian equilibrium price vector, then ξ assigns each agent his desired allocation given this price vector, and second, if a good is in excess supply for some profile of private prices, at least one producer of the good can lower his price and (weakly) increase his utility, given exchange mechanism ξ .

A proper exchange mechanism ξ with these properties thus defines a finite game \mathcal{G} in which each agent's strategy is his private price vector, and he payoff is the utility derived from consuming the goods allocated to him by ξ . A profile of private prices for all agents is a *Nash equilibrium* of \mathcal{G} if each agent's private price vector is a best response to the price vectors of the other agents. A Nash equilibrium of \mathcal{G} is *strict* if any alteration in relative prices by any agent of the goods he consumes renders the agent strictly worse off. We show in Gintis and Mandel (2012) that for a proper exchange mechanism ξ , the only strict Nash equilibria of

the game \mathcal{G} are the Walrasian equilibria of the economy \mathcal{W} .

We now consider an evolutionary game \mathcal{E} for which \mathcal{G} is the stage game and the evolution of agent strategies is governed by a replicator dynamic (Weibull 1995, Hofbauer and Sigmund 1998). We interpret this dynamic as agents periodically adopting the strategy of an agent with the same production good who has had more trading success in the recent past. The evolutionary game \mathcal{E} is thus a multipopulation game in the sense of Weibull (1995), where each population is the set of agents producing a particular good. Multipopulation games have the property that the only stable strategy profiles for the replicator dynamic are strict equilibria of the stage game \mathcal{G} (Weibull 1995). Because the strict equilibria of \mathcal{G} are the Walrasian equilibria of the economy \mathcal{W} , it follows that the Walrasian equilibria of the economy \mathcal{W} are stable fixed points of the evolutionary dynamic \mathcal{E} . The converse also holds: the only stable equilibria of \mathcal{E} are the Walrasian equilibria of the economy \mathcal{E} (Gintis and Mandel 2012).

3 Evolutionary Dynamics and Finite Markov Processes

than replicator equations are somehow aggregated and to check their are consistent with some micro representation, one needs stochastic Markov process models.

The equations of our dynamical system \mathcal{E} do not express macroeconomic dynamics in any convenient form. However, we can construct a discrete version of \mathcal{E} as a finite Markov process. This formulation is conducive to macroeconomic aggregation. The link between stochastic Markov process models and deterministic replicator dynamics is well documented in the literature. Helbing (1996) shows, in a fairly general setting, that mean-field approximations of stochastic population processes based on imitation and mutation lead to the replicator dynamic. Moreover, Benaim and Weibull (2003) show that large population Markov process implementations of the stage game have approximately the same behavior as the deterministic dynamical system implementations based on the replicator dynamic.

Specifically, Benaim and Weibull (2003) show that if the evolutionary dynamic at some point in time enters the basin of attraction of a stable equilibrium of the stage game, then the corresponding Markov process with the same initial state will enter any given neighborhood of the equilibrium with a probability that exponentially approaches one as the population size goes to infinity. The Markov process will then remain in this neighborhood for a period of time that is an exponential function of the population size.

These results allow us to study market dynamics quite rigorously. While analytical solutions for the discrete system exist (Kemeny and Snell 1960, Gintis 2009), they also cannot be practically implemented. However, the dynamics of the

Markov process model can be studied for various parameter values by computer simulation (Gintis 2007, 2012).

We shall begin by reviewing basic facts concerning finite Markov processes.

4 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 construction, 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

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

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 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 .

5 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} . 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 in 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 observed historical frequency. 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 irreducible, 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} .

6 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 $H_{\mathcal{S}}^t$. 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 S as H_S^k .

We say a stochastic process 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 $\psi : H_S^k \rightarrow H_{\mathcal{T}}^{k'}$ 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 S with finite state space S and transition probability function P with k -period memory, where $k > 1$. Then 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 $(i_1, \dots, i_k) \in S^k$ 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 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

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

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 4 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

$$u_m = \lim_{t \rightarrow \infty} P_{im}^n \quad (15)$$

$$u_m = \sum_i u_i p_{im} \quad (16)$$

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

$$P_{mi} = u_j p_{ji} + u_k p_{ki}. \quad (17)$$

If j and k are identically situated, then (17) implies

$$u_i = \sum_r u_r p_{ri}, \quad (18)$$

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.

7 Estimating Markov Processes

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 market dynamics, 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 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.

8 The Dynamics of a Multi-Good Markov Economy

We assume there are goods $k = 1, \dots, n$. Each agent consumes a subset of goods other than his goods. The Markov process is initialized by creating N producers for each production good g^k , so there are Nn agents in the economy. Each agent A is assigned a private price vector $p^A = (p_1^A, \dots, p_n^A)$ by choosing each price from a uniform distribution on $(0, 1)$, then normalizing so that the price of the production good is unity. Each g^k producer is then randomly assigned a set $H \subseteq G$, $g^k \notin H$ of consumption goods.

We create highly heterogeneous utility functions to ensure that our results are not the result of assuming an excessively narrow set of consumer characteristics. The high degree of randomness involved in creating a large number of agents ensures that all goods will have approximately the same aggregate demand characteristics. If we add to this that all goods have the same supply characteristics, we can conclude that the Walrasian equilibrium will occur when all prices are equal.

The utility function u^A of each agent A 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 $1 \dots r/2$) of randomly chosen sizes m_1, \dots, m_k , $m_j > 1$ for all j , and $\sum_j m_j = r$. We randomly assign goods to the various segments, and for each segment, we generate a CES utility function 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 agents 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

$$u(x_1, \dots, x_m) = \left(\sum_{l=1}^m \alpha_l x_l^\gamma \right)^{1/\gamma}, \quad (19)$$

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 $\gamma \neq 0$ (i.e., the utility function segment is not Cobb-Douglas), this gives

$$x_i = \frac{M f_i}{\sum_{l=1}^m p_l \phi_{il}^{1/(1-\gamma)}}, \quad (20)$$

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

$$x_i = \frac{M f_i \alpha_i}{p_i}. \quad (21)$$

In each period, each agent produces a unit of his production good g^h . Producers of the same good congregate in a marketplace that agents who would like to acquire g^h visit with the aim of exchanging some of their production good for a quantity of g^h . Visitor A, who produces g^k , offers an exchange with producer B at a price ratio given by A's private price vector p^A . A's demand for g^h is given by the maximization of his utility function u^A , using his private price vector p^A . Agent B accepts this offer if it is weakly profitable given p^B , if B consumes g^k , if B has a positive amount of g^h in inventory, and if B has not already satisfied his demand for g^k .

Formally, for each good g^k there is a market M_k consisting of all agents who produce good g^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^h producer A is the current active trader, for each good g^k for which A has positive demand, A is assigned a random member $B \in M_k$ who consumes g^h and has a positive amount of g^k in inventory. A then offers B the quantity y_h of g^h in exchange for the quantity x_k^A of good g^k , subject to the constraints $y_h \leq \mathbf{i}_h^A$, where \mathbf{i}_h^A represents A's current inventory of good g^h , and $y_h = p_k^A x_k^A / p_h^A$. Thus agent A must receive in value at least as much as he gives up, according to his private prices p^A . Agent B accepts this offer provided the exchange is weakly profitable at B's private prices; i.e., provided $p_k^B x_k^A \leq p_h^B y_h$. 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^k . If A fails to trade with this agent, he still might secure a trade giving him g^k , because $A \in M_k$ may also be on the receiving-end of trade offers from g^k producers at some point during the period. If a g^k producer exhausts his supply of g^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 trader updates his private price vector on the basis of his trading experience over the period, raising the price of a consumption or lowering the price of his production good by 0.05% if he

failed to purchase any of the consumption good or sell all of his production good, and lowering price by 0.05% if he succeeded in obtaining his consumption good or raising the price of his production good by 0.05% if sold all his production good.

9 Networking

We implement in the Markov process an imitation mechanism that corresponds to the replicator dynamic in the evolutionary dynamic. The information structure of this mechanism consists of a network in which two linked agents have access to each other's strategies and a noisy measure of each other's trading success in previous periods. We create a random graph with a variable degree, ranging from 0.1 to 25. In addition, there are random and infrequent (a small fraction per period) disappearance of linkages and the forging of an equal number of new linkages. We will report on a network degree of five, which means that each agent has a mean of five links to other agents. In fact, similar results hold for a network degree of 0.1 (only one in ten agents has a link to another), with the mean payoff being a decreasing, and the standard error of prices across agents being an increasing function of the graph degree (results not reported here).

After a number of trading periods, the population of traders is updated using the following process. For each market M_k and for each g^k -trader A, let f^A be the accumulated payoff 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_* be the minimum over f^B for all g^k producers B to whom A is linked in the network. For each such linked g^k producer B, let $p^B = (f^B - f_*) / \sum (f^B - f_*)$, where the sum is taken over all agents linked to A. Then $\{p^B\}$ is a probability distribution over the agents linked to 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 A in M_k . Next, among the agents to which A is linked, choose agent B with probability p^B . If B's payoff over the past reproduction interval is greater than A's, then A copies B's price vector.

The dynamic specified by the above conditions determines the evolution of the distribution of private prices from period to period. We find 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.

We will illustrate this dynamic assuming nine goods ($n = 9$) and three hundred producers per good ($N = 300$). The complexity of the utility functions do not allow us to calculate equilibrium properties of the system perfectly, but we know

that Walrasian equilibrium prices are approximately equal because the initial endowment consists of equal quantities of all goods, and the randomization process in specified parameters of utility functions guarantees that aggregate demand for all goods is approximately equal when all relative prices are unity.

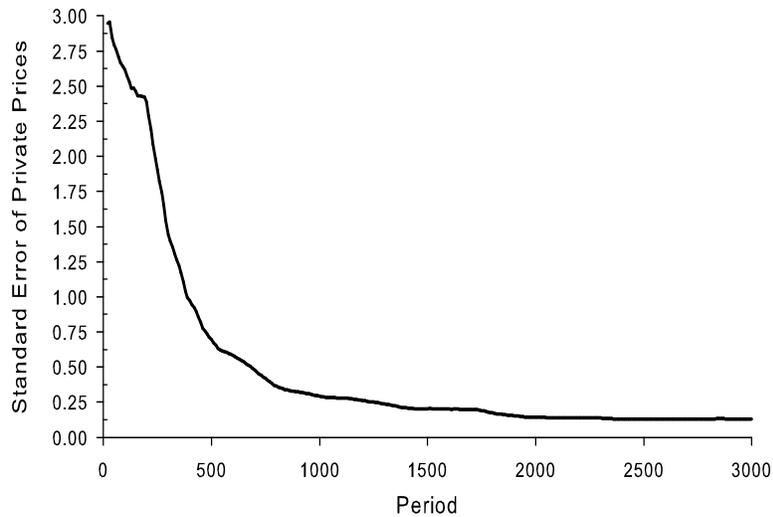


Figure 1: The standard error of prices across agents becomes small, a situation we term *quasi-public prices*.

The results of a typical run of this model is illustrated in Figures 1 to 3. Figure 1 shows the passage from private to quasi-public prices over the first 3,000 trading periods of a typical run. The mean standard error of prices across agents falls from 2.95 to below 0.13 over this period. Thus quasi-equilibrium prices are extremely close to a uniform public price, but with the added attraction that each individual agent is free to change his private price vector at will.

Figure 2 shows the movement of the mean absolute value of excess supply, as a fraction of supply, over 3,000 periods. This falls from 8.50 at the start of the process to below 0.12 after 1,000 periods.

Figure 3 shows that 86% of full allocative efficiency is achieved after 2,000 periods.

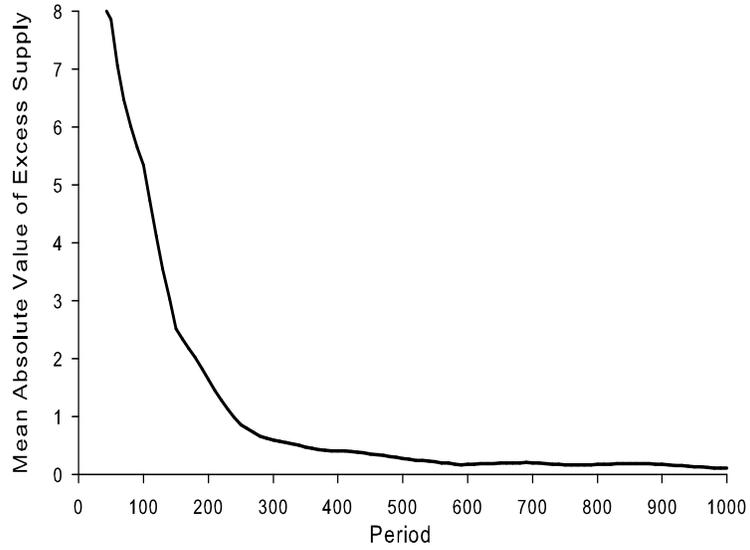


Figure 2: The decline in mean absolute excess supply is rapid, leading to a Walrasian *quasi-equilibrium*.

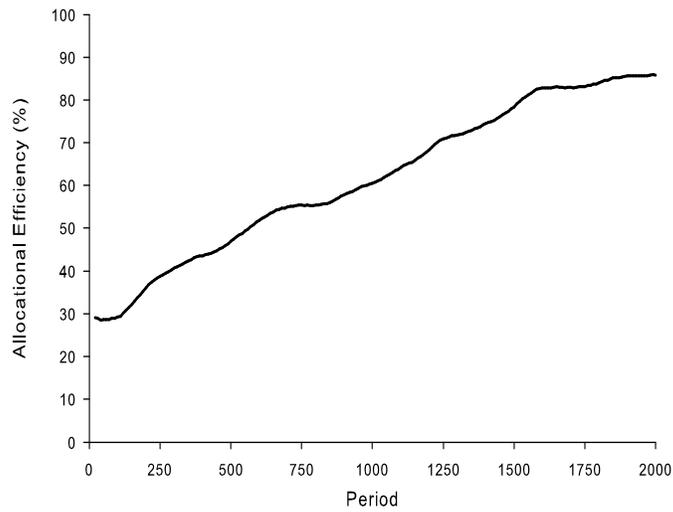


Figure 3: Allocational Efficiency of the Walrasian economy.

10 Conclusion

The traditional approach to modeling Walrasian dynamics was to assume that all agents shared a common vector of prices (*public prices*). Given agent endowments and utility functions, this price vector determined excess demand and supply for all goods. The assumption of perfect competition, according to which all agents are price-takers, precludes endogenous price dynamics, so these models posited an extra-market *auctioneer* that updates the public price vector as a function of the pattern of excess demand and supply. Unfortunately, this approach is doomed to failure because price dynamics based on excess supply and demand are generically chaotic, and 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 (Bala and Majumdar 1992, Saari 1995).

Even had this approach succeeded, it is unclear that these models would have led to a plausible model of price dynamics in a decentralized market system, in which no extra-market price-making entity exists. Perhaps even more problematic is the assumption of public prices itself, for there is no plausible mechanism leading large numbers of independent agents to adopt a common price vector.

We have shown that treating the market economy as a complex system with a walrasian stage game embedded in an evolutionary game dynamic, with agent strategies being private price vectors, leads to a dynamical system whose stable equilibria correspond to Walrasian market-clearing. We have also seen that such a dynamical system can be closely approximated by a finite Markov process, and that the stationary distribution of this Markov process can be studied through computer simulation.

We have studied only very simple Walrasian economies, but the method holds the promise of extending to systems involving firms, raw materials and intermediate goods markets, money, the holding of inventories across periods, and the presence of capital goods. Similarly, one might explore extensions of our approach in which there is some public information (e.g., central bank interest rates) and contracts across periods are permitted, so that a financial sector can be modeled.

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: Non-linear Price Adjustment and Adaptive Expectations," Technical Report, Of-

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.

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

Benaim, Michel and Jürgen W. Weibull, “Deterministic Approximation of Stochastic Evolution in Games,” *Econometrica* 71,3 (2003):873–903.

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

Debreu, Gérard, “Excess Demand Function,” *Journal of Mathematical Economics* 1 (1974):15–23.

Feiwel, George R., *Arrow and the Ascent of Modern Economic Theory* (New York: New York University Press, 1987).

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 Adjustments without an Auctioneer,” *Review of Economic Studies* 39 (1972):1–15.

——, “Stability and Competitive Equilibrium in Two Models of Search and Individual Price Adjustment,” *Journal of Economic Theory* 6 (1973):446–470.

——, *Disequilibrium Foundations of Equilibrium Economics* (Cambridge: Cambridge University Press, 1983).

- Florenzano, M., *General Equilibrium Analysis: Existence and Optimality Properties of Equilibria* (Berlin: Springer, 2005).
- Friedman, Daniel, "Evolutionary Games in Economics," *Econometrica* 59,3 (May 1991):637–666.
- Gale, David, "The Law of Supply and Demand," *Math. Scand.* 30 (1955):155–169.
- _____, "A Note on Global Instability of Competitive Equilibrium," *Naval Research Log Quarterly* 10 (1963):81–87.
- 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 Masahiko Aoki, Kenneth Binmore, Simon Deakin, and Herbert Gintis (eds.) *Complexity and Institutions: Norms and Corporations* (London: Palgrave, 2012).
- _____ and Antoine Mandel, "The Stability of Walrasian General Equilibrium," *under submission* 0,0 (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.
- Helbing, Dirk, "A Stochastic Behavioral Model and a 'Microscopic' Foundation of Evolutionary Game Theory," *Theory and Decision* 40 (1996):149–179.
- Hofbauer, Josef and Karl Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge: Cambridge University Press, 1998).
- Kemeny, John G. and J. Laurie Snell, *Finite Markov Chains* (Princeton: Van Nostrand, 1960).
- 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.

- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, *Microeconomic Theory* (New York: Oxford University Press, 1995).
- Maynard Smith, John and G. R. Price, "The Logic of Animal Conflict," *Nature* 246 (2 November 1973):15–18.
- 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.
- Nash, John F., "Equilibrium Points in n -Person Games," *Proceedings of the National Academy of Sciences* 36 (1950):48–49.
- 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.
- Saari, Donald G., "Iterative Price Mechanisms," *Econometrica* 53,5 (September 1985):1117–1131.
- _____, "Mathematical Complexity of Simple Economics," *Notices of the American Mathematical Society* 42,2 (February 1995):222–230.
- Scarf, Herbert, "Some Examples of Global Instability of Competitive Equilibrium," *International Economic Review* 1 (1960):157–172.
- Sonnenschein, Hugo, "Do Walras' Identity and Continuity Characterize the Class of Community Excess Demand Functions?," *Journal of Economic Theory* 6 (1973):345–354.
- Taylor, Peter and Leo Jonker, "Evolutionarily Stable Strategies and Game Dynamics," *Mathematical Biosciences* 40 (1978):145–156.

- Uzawa, Hirofumi, “Walras’ Tâtonnement in the Theory of Exchange,” *Review of Economic Studies* 27 (1959):182–194.
- , “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]).
- Weibull, Jörgen W., *Evolutionary Game Theory* (Cambridge, MA: MIT Press, 1995).

c:\papers\Complexity Economics StabilityReview\GESStabilityReview.Tex\ December 10, 2012