Computational Economics **8**: 205–231, 1995.

Self-Organization of Markets: An Example of a Computational Approach

NICOLAAS J. VRIEND
Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA, <vriend@santafe.edu>

Abstract. A model of decentralized trade is simulated with firms that produce a given commodity, and consumers who repeatedly wish to purchase one unit of that commodity. Consumers '*shop around*', while firms may attract the attention of potential customers by sending information signals and offering good service. The main objective of this paper is to present an example of a computational approach to address the following question: How do self-organized markets emerge in the economy, and what are their characteristics?

1. Introduction

One of the main tasks of economic theory is to explain the outcomes of a decentralized economy, and to understand the fundamental mechanisms of trade in such an economy. It seems fair to say that the central problems with which economic theory is concerned remained the same since Smith [1776]: How, why, and when does the '*invisible hand*' work? How is it possible that a group of individual agents, each pursuing his self-interest, leads to order rather than chaos? And also the opposite: When would such a group of individuals give rise to chaos rather than order?[1]

As a decentralized economy consists of locally interacting rational agents who are all continuously pursuing advantageous opportunities, such an economy may very well be studied in the framework of complex adaptive system theory (see Anderson et al. [1988]). A '*complex system*' is a system consisting of a large number of agents that interact with each other in various ways. Such a system is '*adaptive*' if these agents change their actions as a result of the events in the process of interaction. Some examples of formal analyses of economies as locally interacting systems are Föllmer [1974], Durlauf [1990], Bak et al. [1993], Blume [1993], and Ellison [1993], who make frequent use of the analytical apparatus of graph theory, statistical mechanics, or the theory of interacting particle systems.

This paper distinguishes itself from that literature in the following way. The question we want to address concerns the emergence of self-organized markets in a decentralized economy. Transactions do not take place in Walrasian central markets, or through anonymous random matching devices, but instead, market interactions depend in a crucial way on local knowledge of the identity of some potential trading partners. A market, then, is not a central place where a certain good is exchanged, nor is it the aggregate supply and demand of a good. In general, markets emerge as the result of locally interacting individual agents who are themselves actively *pursuing* those interactions that are the most advantageous ones,[2] i.e., they are self-organized.

The fact that the interactions between economic agents that we want to analyze are not determined by their given, fixed position in a grid, graph or lattice, or by some kind of anonymous matching device, implies that a formal analysis seems to encounter many technical difficulties. Therefore, we will consider a simple model of a decentralized economy, using a computer simulation. Each rational economic agent is modeled separately with machine learning techniques recently emerged in the field of Artificial Intelligence. Modeling the *homo oeconomicus* as a '*machine*' does not pose particular conceptual difficulties to economic theory.[3] After all, as Lucas puts it, doing economics means "*programming robot imitations of people*" (in Klamer [1984], p. 49). Although Lucas' statement was only meant as a metaphor, current-day computational potentialities make it interesting to take it literally, and to consider its usefulness for economic theory.

What Lucas was referring to, of course, was that the *homo oeconomicus* is a rather mechanical representation of individual agents. The fundamental characteristic of the *homo oeconomicus* is that he simply chooses (one of) the most preferred option(s) in his perceived opportunity set. In fact, the *homo oeconomicus* is an '*opportunist*'; always doing the best he can. This implies that the question of the modeling of the perceived opportunity sets of the individual agents becomes a primary issue. During the process of interaction between the individual agents in a decentralized economy, perceived opportunities evolve. Such changes may be due either to a change in underlying circumstances or to a change in the perception of these circumstances. The latter is called '*learning*'. Usually, these changes will not only occur simultaneously, but there will also be interaction between the dynamics of learning and the dynamics of economic forces as such.

The problem for economists is, that they are not in a position to tell how a set of given physical stimuli maps onto a set of perceived opportunities (see e.g., Tversky & Kahneman [1986]). One way to solve this problem, is assuming that each agent simply observes the objective truth. Basically, this is what the Rational Expectations Hypothesis is about. Here, we will follow an alternative route to abstract from all psychological issues concerning the perception of opportunities, assuming that the perception of opportunities is an *endogenous* process. That is, the set of perceived opportunities depends strictly upon the preceding sequence of actions and outcomes. We know that the *homo oeconomicus'* actions, given his preferences, depend strictly upon his perceived opportunities. Hence, actions are a function of perceived opportunities, and perceived opportunities a function of earlier actions. As a result we get a sequence analysis of actions and outcomes in which perceptions or expectations do not appear explicitly, but only "*between the lines*" (Hart [1951], p. viii). Hence, formally we can model each agent's actions as a function of previous actions and outcomes.

The problem, then, is how to model this latter mapping. One solution would be to apply simple fixed rules-of-thumb. The drawback of this is not only that they are arbitrary, but that the results will be arbitrary as well. We will give an example of this later. Therefore, we do not tie down the set of functions a priori in an arbitrary way. With the use of Artificial Intelligence techniques, it is possible to keep the relations between actions/outcomes and previous actions/outcomes completely flexible. Hence, we can analyze in how far the market provides sufficient structure to tie down the set of possibly perceived opportunities, thus constraining the behavior of the individual agents. Therefore, we simulate a decentralized economy, and look for the emergence of regularities in actions and outcomes during the process of creating and trading away of opportunities by rational agents. The emergence of such regularities is usually related to the metaphor of the '*invisible hand*'. While the individual agents take care only about their own self-interest, it is the '*invisible hand*' that is thought to perform an ordering function, bringing about coordination of economic activities. The questions to be examined are: *(i)* To what extent do the individual agents create opportunities to trade? *(ii)* How does the information concerning these opportunities spread through the economy? *(iii)* To what extent are these opportunities trade away? Or, to put it another way: How do self-organized markets emerge in the economy, and what are their characteristics?

It should be stressed that this paper's scope is modest. Its objective is not to solve all '*invisible hand*' puzzles, and it is not to understand all the ins and outs of the specific economic model considered. Instead, this paper is an example of a method, i.e., a computational approach, in order to study emergent behavior and self-organization in a decentralized economy. Also, the scope of this paper is not to analyze the method per se, i.e., the specific algorithms applied, but to illustrate the usefulness of these tools with respect to the issues of emergent behavior and self-organization.

In section 2 we will delineate the decentralized economy to be simulated. In section 3 we will describe how we model rational consumers and firms. Section 4 will contain the results of the simulations, which are analyzed and discussed, while section 5 will conclude.

2. The Economy

We consider a closed, decentralized economy, in which there are two types of agents: firms and consumers. Time is divided into a sequence of basic periods. From here on, we will call each single time period a '*day*'. During each day the following sequence of actions takes place:

(i) At the beginning of the day firms produce a given homogeneous consumption good, which has a given price. This price is fixed and equal for all agents. It is also constant through time, and known to all agents.[4] All firms are identical in that they use the same production technology, which exhibits constant returns to scale. Production decided upon at the beginning of the day is immediately for sale, while

unsold stocks perish at the end.[5] Firms may send information signals to some other agents in the population at the beginning of each day. Each signal communicates to its receiver that the firm offers the commodity for sale on that day. Spreading these signals is costly, the marginal costs being constant. Receiving information, on the other hand, is costless. These signals are the only means of communication between individual agents. Thus, a firm, preferring more profits to less profits, makes available for sale a number of units of the commodity, spreads information about this in order to attract the attention of potential customers, and then just waits to receive customers to sell to.

(ii) During the day consumers are '*shopping*'. Preferences and endowments are such that, given its price, each consumer wishes to buy and consume exactly one unit of the produced commodity on each day.[6] During each day consumers shop around in search of agents to visit on a given day. Each consumer can visit only one agent on each day.[7] Trade takes place on a '*first-come first-served*' basis.

(iii) At the end of the day all agents evaluate their own market experience. A consumer who has not found a unit of the commodity available, will turn home unsatisfied. Each firm observes the demand it faced during the day.[8] Demand above the firm's available supply is simply forgone. Then, knowing also its production and communication costs, a firm can computer its profits ex post. These profits might turn out to be negative. In our analysis we discard all kinds of liquidity or bankruptcy problems. Firms simply prefer higher to lower profits. Thus we are implicitly assuming some complementary model of a loan market in which lenders allow the firms to make negative profits on unpleasant conditions. Another possibility that may occur is that a firm will perceive no profitable opportunity to produce a positive amount of output and will therefore choose not to produce or signal at all. Whether such firms are definitely out of business or may re-enter will be determined endogenously. In any case, the number of firms is fixed, and we do not consider the possibility of new entrants.

Each firm's objective is to maximize profits. A consumer's only objective is finding a unit of the commodity on each day. Besides the price of the commodity and their own preferences, endowments, and technology, the individual agents are, at the beginning of day 1, without any quantitative or structural knowledge of the economy. Their situation is one of moving around in a crowd of individuals who are all swarming around. They do perceive that there exist other agents in the economy, but they do not know their characteristics. In particular, a priori they have not identified the agents that are firms or those that are consumers interested in this commodity, possibly offering them a trading opportunity. They are even unaware of the numbers of other agents, and they are ignorant about the actions of others. They do not know anything about the market experience of other agents, nor are they informed what future market developments will be. Their perceptions of trading opportunities will be based on their own experiences in the market.[9] But when markets are not orderly, i.e., when there may be simultaneously rationed sellers and rationed buyers, these individual market experiences are not merely a reflection of the overall state of the economy. Moreover, '*shopping*' is anonymous. The idea is here, that one of the distinguishing characteristics of monetary shopping is that as long as a customer has money in his pocket and agrees to pay cash, while the firm can hand over the commodity immediately, nobody is going to ask any further questions (see e.g., Anderlini & Sabourian [1988] or Shubik [1988]).

As firms are ignorant as to the characteristics of the agents to whom they may send information, they are indifferent in this respect, and therefore they choose the agents to whom they send a signal at random. Hence, the remaining decision problem for a firm is the quantity to produce and the number of signals to send on each day. A firm's objective function can be denoted as:

$$\begin{aligned} \text{profit} \; = \; & \text{price} \cdot \text{minimum[production, demand faced]} \\ & -\text{production costs} - \text{signaling costs} \end{aligned}$$

The result of the ignorance described is that a firm is not in a position to specify the demand function facing it. Hence, a firm is not able to maximize its (expected) profits directly with standard techniques, as it would need to know the complete, possibly time-varying, distribution function of demands directed to it. Note that knowledge or an estimation of the expected demand would *not* be sufficient to maximize expected profits.

The economy sketched seems to capture the fundamental aspects of real decentralized economies that most commodities are produced in advance, while they are purchased on a repeated basis by shopping consumers.

3. Modeling the Individual Actions

As explained in section 2, an individual firm's problem is ill-defined, and therefore it has to work inductively. Each day, a firm chooses a point in its (production, signaling)-space, or $\mathbb{N}^2$, and then observes the outcome generated by that action. Hence, a firm's trading opportunities can be represented by a profit landscape, in which a firm continuously tries to find higher points. Clearly, this search space is large. We will use a combination of a Classifier System (CS) and a Genetic Algorithm (GA) to model each firm's actions (see e.g., Holland [1986] and [1992], or Machine Learning [1988]).[10]

The CS, first, predicts in parallel for a set of points in the (production, signaling)-space the eventual payoff that might be generated by these points. At the beginning of day 1, a firm does not have any information as to what the most valuable actions are. Therefore, the initial set of points for each firm's CS consists of randomly chosen actions in the firm's search domain, and the initial predictions are equal for all points. These predicted payoffs are updated on the basis of the firm's own experience, each prediction being some weighted average of past payoffs. Second, the CS decides on each day which of these points is chosen as the current action, where the probability of a point being activated depends on its predicted payoff. This choice of actions is a stochastic function of predicted payoffs, i.e., it is not simply the strongest point that is activated because a rational firm will search to balance the exploitation and the exploration of perceived opportunities.

A rational firm will choose his actions in those regions of the (production, signaling)-space that it perceives to be most advantageous. Therefore, a GA is used to direct the set of actions towards the regions of the (production, signaling)-space in which the performance is likely to improve. Through the application of the genetic operators reproduction, crossover and mutation some new action strings are created every now and then to replace weak existing strings. The frequency at which this is done is determined by a parameter called the '*GA rate*'. Note that a too high GA rate would mean that the CS does not get enough time to predict the value of the newly created strings, while a too low GA rate would lead to lack of exploration of new regions. Furthermore, a firm knows that when a stock-out occurs, its profit would have been higher if it had produced more, and the same applies if it had produced less in case there is some stock left. Therefore, every day a mutation operator is applied to the most recently activated point, such that the number for production is increased or decreased slightly in the apparently right direction.

On each day, a consumer is only interested in obtaining one unit of the commodity. Therefore, he has to find a firm, and this firm must have at least one unit of the commodity available. A priori, however, a consumer does not know which agents are firms, and he has no knowledge as to the inventory policies of the individual firms. As to the various potential suppliers of the commodity, as perceived by a consumer on a given day, we model these into the following three categories. First, the agents the consumers visited on their most recent trip to the market. Returning to such an agent, we will call '*patronage*'. We define '*patronage*' as '*return to the last visited agent*', encompassing any possible motive to adhere to any type of agent. A second category consists of the agents the consumers have identified as firms selling the commodity, through the information signals received on that day. Consumers do not use the information signals received on previous days, as these signals communicate only the firms' willingness to sell the commodity on the day concerned. In case a consumer decides to visit a firm in this category, he chooses randomly one firm among the signals he received. Third, any other, unidentified agent; to be chosen at random.

There are also three possible market experiences for a consumer. First, he succeeds in finding one unit of the commodity and turns home satisfied. Second, the consumer visits a firm but arrives too late and finds only empty shelves. Third, the consumer does not even succeed in tracking down a firm and is '*lost in the mist*'. With respect to the information a consumer received from the firms, there are two possibilities. Either he received some information signals or he did not get any signal at all. These possible states and actions for a consumer on a given day can very well be modeled by a CS. In table I, the complete CS is presented. Note that given that this set of rules covers the consumer's whole action space, there is no need to generate new rules, by applying a GA.

TABLE I. Classifier System consumer[11]

1)	if	Sat	and	Info	then	Patr
2)	if	Sat	and	Info	then	Known
3)	if	Sat	and	Info	then	Rand
4)	if	Sat	and	¬Info	then	Patr
5)	if	Sat	and	¬Info	then	Rand
6)	if	Late	and	Info	then	Patr
7)	if	Late	and	Info	then	Known
8)	if	Late	and	Info	then	Rand
9)	if	Late	and	¬Info	then	Patr
10)	if	Late	and	¬Info	then	Rand
11)	if	Mist	and	Info	then	Patr
12)	if	Mist	and	Info	then	Known
13)	if	Mist	and	Info	then	Rand
14)	if	Mist	and	¬Info	then	Patr
15)	if	Mist	and	¬Info	then	Rand

At the beginning of day 1 a consumer does not have any information as to what the most valuable shopping rules are. Therefore, the predicted payoffs are equal for all rules. At the end of each day, each consumer evaluates his market experience. There are two possible outcomes for a consumer's shopping efforts. Either he has succeeded in buying one unit of the commodity, or he has not. While the former is a positive outcome, the latter is a negative one. The CS governs the reinforcement learning process, updating the predicted payoffs according to experienced payoffs, and determines the actions analogously to the procedures described above for the firms. Note that here, only those rules for which the '*if*...' part is fulfilled may be activated.

GAs are search procedures based on the mechanics of natural selection and natural genetics. The key feature of GAs is their ability to exploit accumulating information about an initially unknown search space, in order to bias subsequent search efforts into promising regions. GAs are especially appropriate when, for one reason or another, analytical tools are inadequate, and when '*point-for-point*' search is unfeasible because of the amount of possibilities to process, which may be aggravated by the occurrence of non-stationarity. But the most attractive feature of GAs is that they do not need a supervisor. That is, no knowledge about the optimal, or '*correct*', action, or a measure of the distance between the actions of the CS/GAa and the '*correct*' action, is needed in order to adjust the set of actions of the CS/GA. The *only* information needed are the (predicted) outcomes that would be generated by each action. This information is supplied by the CS. In this sense CS/GAs exploit the local character of information, and no further knowledge about the underlying outcome generating mechanisms is needed, e.g., the derivatives of certain functions.

CS/GAs are not the only possible algorithms in this context. They are examples of reinforcement learning algorithms. *"Reinforcement learning is the learning of a mapping from situations to actions so as to maximize a scalar reward or reinforcement signal. The learner is not told which actions to take, as in most forms of machine learning, but instead must discover which actions yield the highest reward by trying them"* (Sutton [1992], p. 225). Agents must experiment to try new actions, and actions that led in the past to more satisfactory outcomes are more likely to be chosen again in the future. Machine Learning [1992] presents a survey of reinforcement learning. Closely related stochastic learning models can be found in Cross [1983], based on Bush & Mosteller [1955], and Roth & Erev [forthcoming], based on Harley [1981].[12]

Other recent models of dynamics with experimentation in the economics literature are, for example, Ellison [1993], Kandori et al. [1993], and Young [1993]. There is, however, a difference between those models and reinforcement learning. In the evolutionary dynamic models mentioned, adaptive behavior is basically a one-step error correction mechanism. The agents have a well-specified model of the game, they can reason what the optimal action would be, given the actions of the other players, completely independent from any payoff actually experienced, and they play a best-response strategy against the frequency distribution of a given (sub-)population of other players. The evolutionary dynamics consist of a co-evolutionary adaptive process, players adapting to each others' adaptation to each other..., plus experimentation in the form of trembling. The very first task for a reinforcement learning algorithm, on the other hand, is to *learn* what the optimal action would be. The agents do not have a well-specified model of their environment, and they do not know which action would be the best response.[13]

4. Simulation Results

In this section, we analyze the simulated economy. We run the model five times for 3000 basic periods, or days, with 50 firms and 5000 consumers. The other parameter values of the economic model are: price of the commodity 1.00, marginal cost of signaling 0.08, and marginal cost of production 0.25. At the end of day 3000, of an average run, more than 16 million units of the commodity have been produced, almost 80 million information signals have been sent, and 15 million times a consumer has made a trip to the market. This has resulted in more than 14 million transactions. We will characterize the history of actions and outcomes by looking for regularities in the data set. In particular, we are interested in those regularities that cannot be deduced directly from the built-in properties of the individual agents or some other microeconomic aspect of the model; at least not by any argument which is substantially shorter than producing that regularity by running the simulation itself (see Lane [1992]).[14] We will search for such regularities not only at the

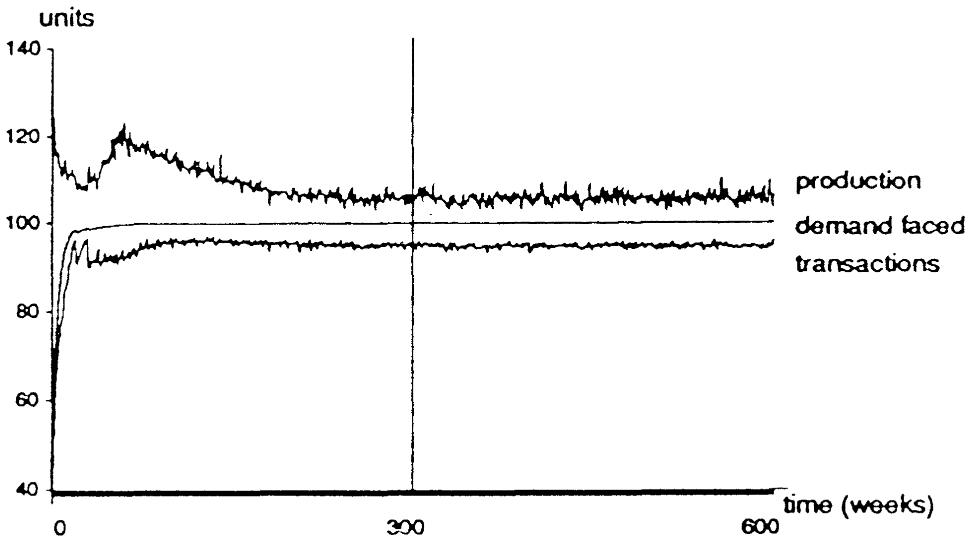

Fig. 1. Time series average production, demand faced, and transactions.

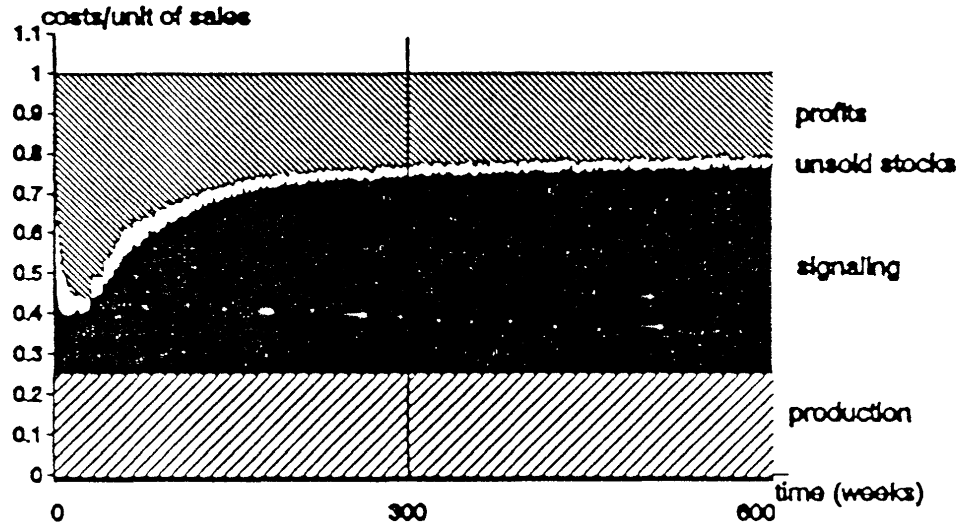

Fig. 2. Time series average costs per unit of sales.

level of aggregate or macroeconomic activity, but we will also analyze whether the microeconomic distributions, concerning the experiences of the individual agents underlying those aggregates, shows some regularities.

4.1. Macroeconomics: Time Series of Aggregates

The most commonly considered macroeconomic variables are production, demand, sales, advertising, profits, profits per unit of sales, and unsold stocks.[15] We also consider the relative number of consumers choosing to patronize, the service offered by the firms, i.e., the probability that a firm can satisfy a given client, and, in order to analyze how successful the coordination of economic activities is, we constructed a measure of the efficiency of the economy. Two factors have to be taken into

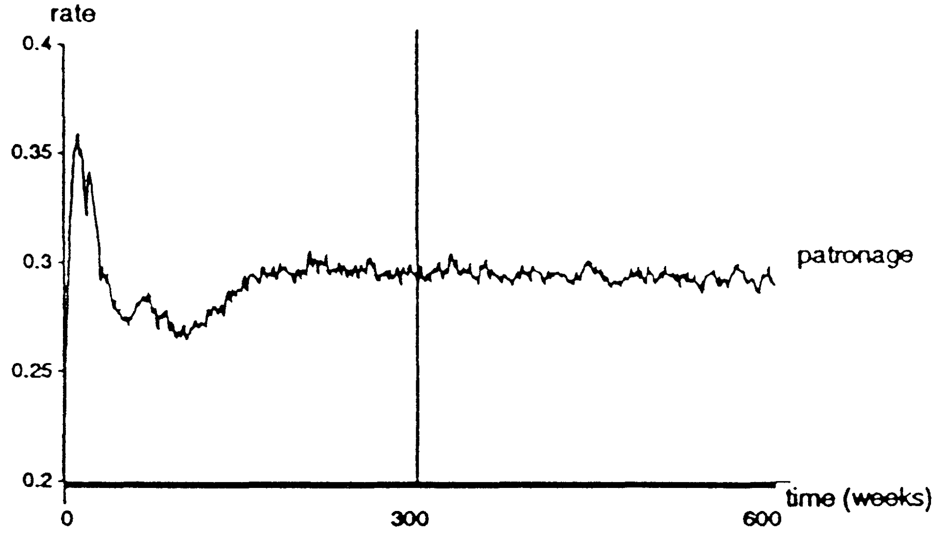

Fig. 3. Time series patronage as proportion of consumer population.

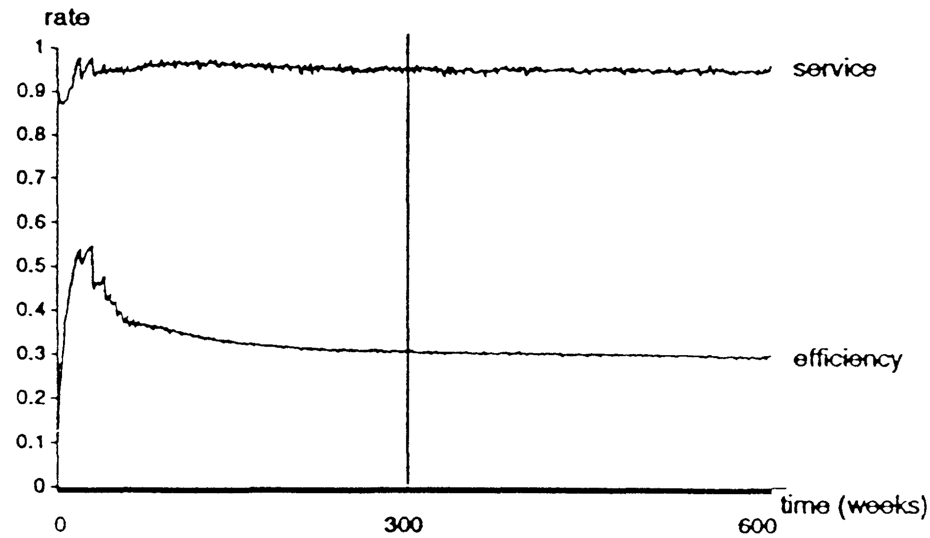

Fig. 4. Time series average service rate and macroeconomic efficiency.

account. On the one hand, the number of actual transactions relative to aggregate demand, and on the other hand, the signaling and production costs incurred to realize these transactions compared with the resources that were technically absolutely indispensable to create these transaction opportunities, i.e., the production costs. Hence, efficiency = (actual transactions/aggregate demand) · (indispensable resources/aggregate costs), where $0 \leqslant$ efficiency $\leqslant 1$. The series, averaged over the five runs, are presented in figures 1 to 4.[16] [17]

These graphs show four features of the macroeconomic time series. First, we see strong movements in most series at the very beginning of the history. As all agents were initially ignorant as to the relative values of their possible actions, and to the objectively given, overall economic opportunities, the first phase of the time series appears to be dominated by an overall learning effect. Second,

TABLE II. Macroeconomic statistics, days 1501–3000[18]

Variable	Avg.	(s.d.)	st. dev.	(s.d.)	min.	(s.d.)	max.	(s.d.)
Transactions	95	(0.1)	1	(0.1)	91	(2.0)	98	(0.1)
Production	106	(0.4)	3	(0.1)	98	(2.1)	115	(0.8)
Signals	597	(1.8)	18	(1.5)	535	(17.3)	673	(11.3)
Demand faced	100	(0.0)	0	(0.0)	100	(0.0)	100	(0.0)
Unsold stocks	11	(0.4)	2	(0.1)	5	(0.1)	19	(1.6)
Profits	20	(0.2)	1	(0.1)	15	(0.6)	24	(0.3)
Patronage	0.29	(0.003)	0.00	(0.000)	0.28	(0.006)	0.31	(0.003)
Prob. service	0.95	(0.001)	0.01	(0.001)	0.92	(0.013)	0.98	(0.002)
Efficiency	0.3	(0.001)	0.00	(0.000)	0.28	(0.004)	0.32	(0.002)

the macroeconomic efficiency coefficient increases fast right from the start, and after about 20 weeks the economy reaches a high performance level. The average demand faced by the firms is almost 100, the maximum attainable, implying that all consumers have discovered their way to the market, while production and sales are close to that level. Moreover, profits per unit, efficiency, and patronage reach their historical maximum also after 20 weeks. Hence, '*on average*' agents learn fast about the overall opportunities of the economic environment; and that while each individual agent observes only his own actions and outcomes. Third, the economy does not settle down at that high performance level. Although the average demand faced by the firms remains constant, the average production shows a prolonged upward sweep, whereas the average transactions move in the opposite direction. Evidently, consumers have some difficulties in finding the right firms. This seems to be confirmed by the fluctuations in the patronage series. Furthermore, profits per units appear to be squeezed steadily by increasing signaling costs, and the system efficiency decreases accordingly. It seems that in this second phase, roughly from week 20 to 300, the economic interaction between the individual agents, who are all continuously learning about their opportunities, forcefully sways the economy, although the overall economic environment is constant. Fourth, after week 300 all series appear to become comparatively stationary. Table II gives some macroeconomic statistics of this phase, based on the daily observations, and averaged over the five runs.

One question to ask is, how '*reasonable*' the actions of the agents are, on average, during this second half of the history. As a strategy for an agent would be an action for the first day plus an action for each day conditional upon the responses experienced in market, these responses depending upon the actions of both the consumers and the firms, deriving a proper equilibrium strategy would be complicated. Therefore, we make the following approximation. Take the average patronage rate of the satisfied consumers during the last 4 weeks, assume there is no stochasticity, i.e., the demand equals the expected value of the stochastic

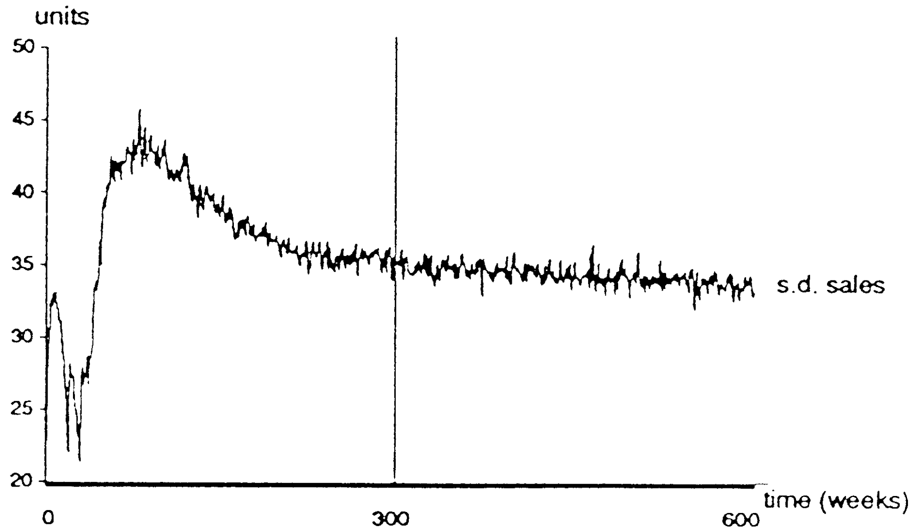

Fig. 5. Time series standard deviation in sales.

demand function, and that the firms know the demand function. Then, the only decision variable for a firm is the number of signals to send. We look for the one-period optimal strategy for the individual firm that is symmetric and stationary. The analysis can be found in appendix B. This equilibrium would be: signaling 652, while production and sales are 100. During the last 4 weeks, we observe, on average: signaling 611, production 106, and sales 95. The differences may be partly explained as follows. In our model, demand is stochastic, which means that each given signal is less profitable, because it is uncertain whether there will be a unit available when a consumer arrives, implying that less signaling will be optimal. Moreover, with demand being stochastic, in general over-production is profitable, while sales must be less than production.

4.2. Microeconomics: Differences between Agents

In this section we examine the differences between the individual firms and between the individual consumers in their actions and outcomes. Individual experiences may differ in the sense that the distribution over time of an individual agent's experiences merely reflects the, possibly changing, cross-sectional distribution of experiences in the economy. But it may also be that individual experiences do not '*average out*' over time.

First, we consider the firms. In figure 5 the development of a regularity in the differences in sales between the individual firms is traced. After two strong sweeps, the standard deviation of the sales seems to be relatively constant, while the spread in sales among the firms in any given week is considerable.[19] A first question to ask is: Do these differences average out over time? In table III, we give some statistics, averaged over the five runs, of the realization of the action/outcome variables, for each firm averaged over the days 1–3000. We see that the spread in actions

TABLE III. Statistics individual firms, averaged days 1–3000.

Variable	Avg.	(s.d.)	st. dev.	(s.d.)	min.	(s.d.)	max.	(s.d.)
Sales	94	(0.1)	13	(0.7)	59	(4.2)	122	(5.4)
Production	108	(0.3)	14	(0.9)	69	(4.4)	137	(7.7)
Signals	530	(1.6)	72	(3.7)	329	(28.2)	676	(25.5)
Demand faced	99	(0.0)	13	(0.7)	63	(4.8)	128	(5.3)
Unsold stocks	14	(0.3)	2	(0.2)	10	(0.5)	18	(1.5)
Profits	25	(0.2)	4	(0.2)	15	(1.0)	34	(1.5)

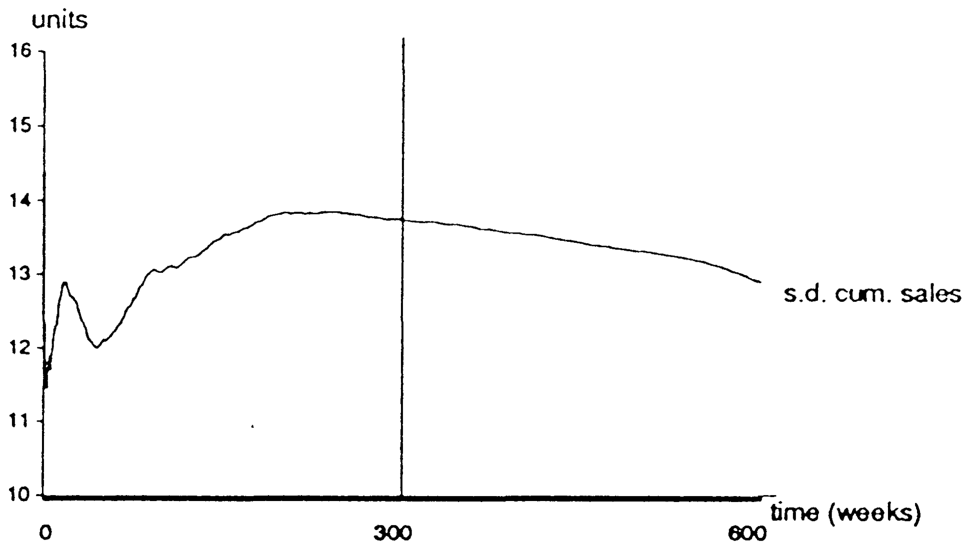

Fig. 6. Time series standard deviation cumulative sales.

and in outcomes is considerable, and that there are systematic and significant differences between the firms' experiences, although all firms were identical at the start. Accumulated over the history of 3000 days, with respect to each of the variables listed in table III, the highest value is about double the lowest.[20] This leaves the question unanswered in how far the firms would have differed among each other after, e.g., 2 million days. Although we cannot answer that question, we can show what the trend is during the simulated history of 3000 days.

Figure 6 shows that, at the end of the history, the cumulative sales of the firms tend to become more equal to each other.[21] Therefore, a second question is: Suppose these differences do average out over time, although within each period the firms are not equal, how would a stationary distribution of firm sizes look like? Therefore, we take the last 1500 days, divide sales in 10 classes, assume a discrete Markov process, calculate the Markov matrix of transition probabilities, determine the fixed point of the transformation, which is independent of the initial state, and obtain the following stationary sales distribution. Figure 7 also shows the distribution of firm sizes at the final day of the simulations, averaged over the five runs.[22]

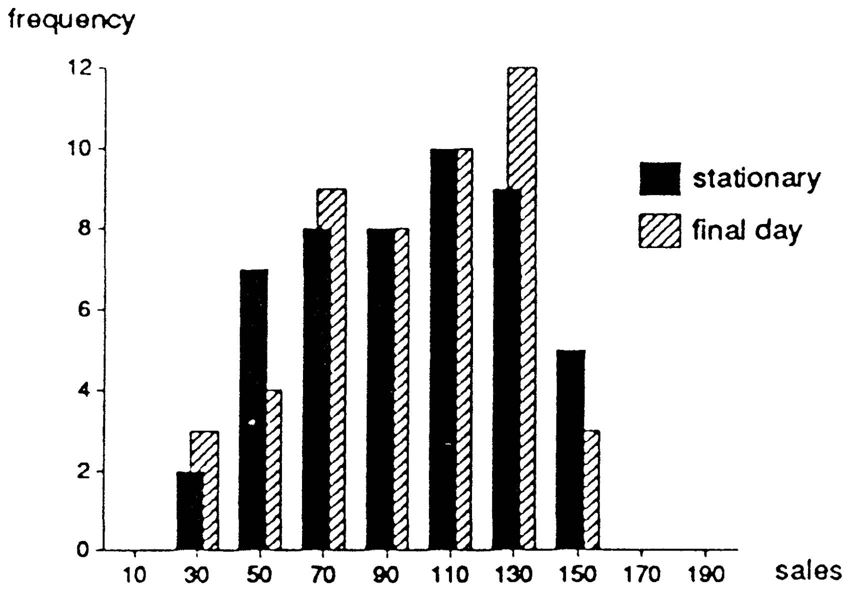

Fig. 7. Stationary and final firm-size distribution.

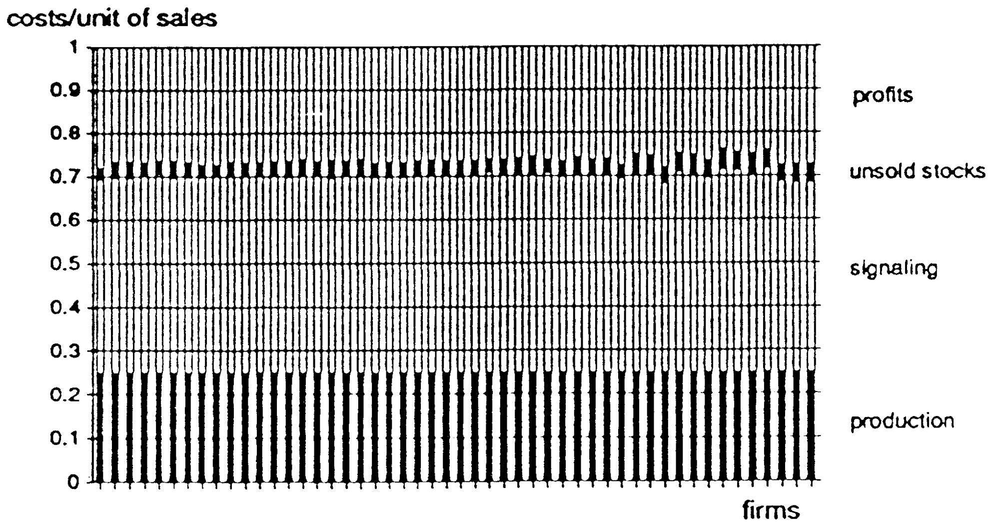

Fig. 8. For each firm: average costs/unit of sales days 1–3000.

Notwithstanding the differences between the experiences of the individual firms on the market, and the resulting firm-size distribution, there are also some correspondences among the firms. Splitting the gross revenue for each unit sold between the various costs incurred by each individual firm shows an identical distribution for each firm. See figure 8, where the firms are ordered on their cumulative sales.[23]

That the production costs per unit sold are the same for each firm is determined directly by the production technology, which is identical for all firms. But the firms, though they must have such different perceptions with respect to the extent of their trading opportunities, turn out to spend all the same amount on signaling per unit sold, they all let perish the same amount of unsold stocks per unit sold, and they all make the same profit per unit sold. Note that all firms use the same

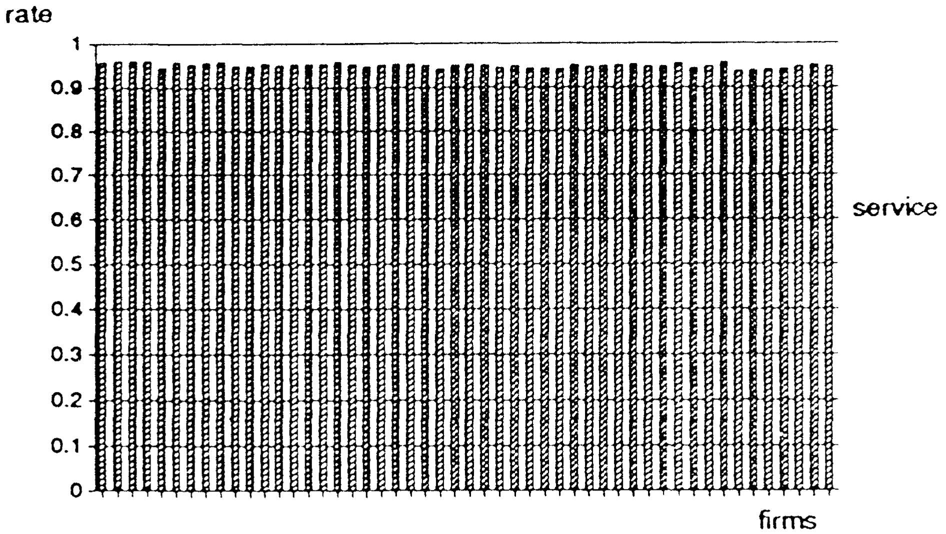

Fig. 9. For each firm: average probability to serve a given client days 1–3000.

TABLE IV. Statistics individual firms, averaged days 1–3000.

Variable	Avg.	(s.d.)	st. dev.	(s.d.)	min.	(s.d.)	max.	(s.d.)
Profit/unit	0.26	(0.002)	0.009	(0.008)	0.24	(0.006)	0.28	(0.002)
Unsold stocks/unit	0.04	(0.001)	0.006	(0.004)	0.03	(0.001)	0.05	(0.002)
Signaling/unit	0.45	(0.001)	0.007	(0.006)	0.43	(0.003)	0.46	(0.004)
Prob. service	0.95	(0.001)	0.005	(0.004)	0.94	(0.002)	0.96	(0.003)

constant-returns-to-scale technology. But that does not yet imply that they should follow the same signaling and production policy per unit sold. A second regularity is presented in figure 9, where we see that, on average, also the probability to be able to serve a given client turns out to be equal for all firms. Table IV gives some statistics concerning these uniformities among the firms.

Next, we consider the consumers. In table V we give some descriptive statistics, where for each consumer each variable has been averaged over the days 1–3000. Reading table V from the bottom row to the top, the following interesting points emerge. First, we see that consumers have discovered that agents who never send signals and are never observed selling the commodity, are very unlikely to be firms

TABLE V. Statistics individual consumers, averaged days 1–3000.

Variable	Avg.	(s.d.)	st. dev.	(s.d.)	min.	(s.d.)	max.	(s.d.)
Satisfaction	0.94	(0.001)	0.004	(0.000)	0.92	(0.001)	0.96	(0.002)
Patronage	0.29	(0.002)	0.055	(0.005)	0.13	(0.006)	0.52	(0.018)
Visits known signaler	0.70	(0.002)	0.055	(0.005)	0.47	(0.018)	0.86	(0.007)
Visits random agent	0.01	(0.000)	0.001	(0.000)	0.00	(0.000)	0.01	(0.001)

offering the commodity for sale. Second, the information signals sent by the firms do matter to the consumers, as in 70% of the cases they use these to select a firm. Third, there is a systematic difference in the consumers' shopping behavior, as some consumers patronize four times as often than some others. But, fourth, there is only a small variance between the consumers' market outcomes.[24]

4.3. Competition, Patronage and Arbitrage

In section 4.1 we observed that the firms, although reaching a high performance level quickly, continued to increase their signaling activity steadily, thus eroding their profits. In the end, the firms differed strongly in their perceptions of and experiences on the market, but on average they all had the same cost profile per unit sold. Moreover, on average they all offered the same opportunities to their clients. This does not mean that differences in service rates do never occur, but just that they do not persist. This all suggests that competition works, and leads to regularities. Note that firms are not aware of the existence of competitors, let alone the actions of those competitors. Competition works via the market.

To explain these observations, it is important to remember that in the simulated economy, the price of the commodity is given, known and equal for all agents and constant through time, that the commodity itself is homogeneous, and that all firms and all consumers were identical at the start. The systematic differences between the individual firms *emerge* in the economic process. Hence, all references to presumed differences in skills or attitudes, like e.g., '*aggressivity*', and other psychological factors would be out of place. Individual firms may become different in the following two sense, as perceived by consumers. First, their identity may be better or worse known by means of the information signals. Second, they may differ in the reliability of their service. Consumers are only interested in obtaining a unit of the commodity. Hence, whereas in models with price-setting firms, consumers are looking for bargains, in our economy they are pursuing high service rates. Therefore, firms compete using signaling and a good service as weapons.

In this section we will focus upon the role played by the consumers in the market process, and in particular upon the phenomenon of patronage. In section 3 we defined '*patronage*' by '*returning to the last visited agent*', i.e., independent of the state the consumer finds himself in.[25] An important question is whether consumers who were satisfied on a given day do more often return to the same firm on the next day than consumers who were disappointed by the firm they visited. This is a crucial issue, because it is this type of behavior, which we might call '*strict patronage*', that leads to arbitrage of trading opportunities. For suppose some firms offer higher service rates than other firms. Strict patronage would imply that a firm not able to satisfy its clients is likely to loose some of its customers. Given its level of production, that would mean a higher coefficient of customer satisfaction on the next day. On the other hand, a firm satisfying its customers is likely to enlarge its clientele, thus lowering its service rate. Hence, ceteris paribus, strict patronage

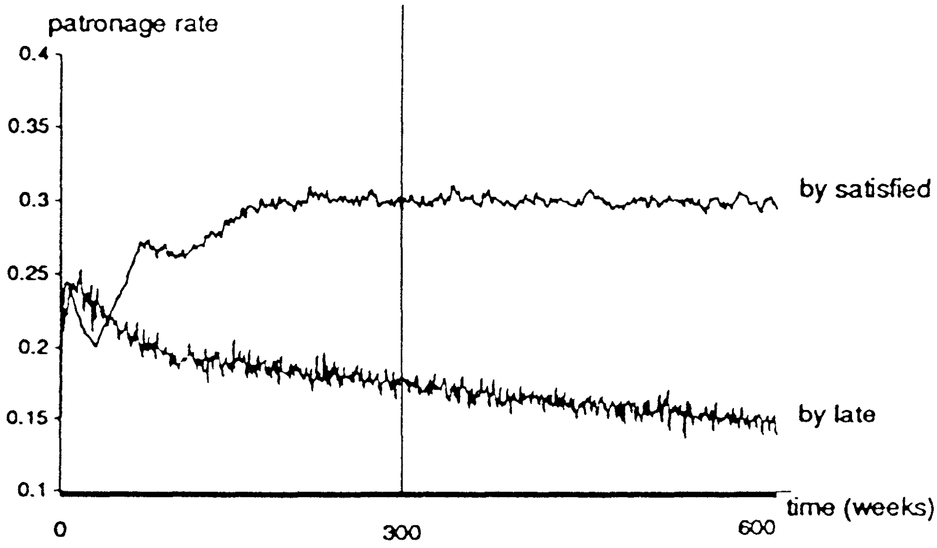

Fig. 10. Time series patronage for two categories of informed consumers.

TABLE VI. Average satisfaction rate for informed consumers, averaged days 1–3000.

	Patronizing	(s.d.)	Switching	(s.d.)
Previous day satisfied	0.94	(0.001)	0.95	(0.001)
Previous day late	0.87	(0.005)	0.93	(0.003)

directly implies arbitrage of trading opportunities, in the sense of the equalization of service rates across firms.[26]

Figure 10 shows that consumers soon perceive the distinct opportunities offered by strict patronage.[27] On each day, informed satisfied consumers patronize about twice as often as informed consumers that had arrived late on the previous day. In appendix A we explain how the patronage rate is some way biased. As the rules that imply visiting a random, unidentified agent are chosen with very low probabilities, the probability of activating one of the rules that imply patronage is biased towards slightly less than 1/3. This implies that strict patronage does turn out to develop, but mainly in the negative sense that disappointed consumers perceive it to be advantageous to avoid their failing supplier.[28]

In table VI we compare the average satisfaction rate of previously satisfied consumers who decided to patronize with those who switched to another firm, visiting one of the firms they received a signal from, and of previously disappointed consumers who adhered to their failing supplier with those who changed firm. We see that, averaging over all days, all consumers, and the five runs, the perception of disappointed consumers that it is advantageous to change firm is confirmed by these averages.

TABLE VII. Average actions and outcomes two variants, 1–3000.

Variable	Standard	(s.d.)	'fixed' patr.	(s.d.)
Transactions	94	(0.1)	96	(0.1)
Production	108	(0.3)	103	(0.1)
Signals	529	(1.6)	40	(3.0)
Demand faced	99	(0.0)	99	(0.0)
Unsold stocks	14	(0.3)	7	(0.2)
Profits	25	(0.2)	67	(0.2)
Patronage	0.29	(0.002)	0.98	(0.000)
Satisfaction	0.95	(0.001)	0.97	(0.000)
Efficiency	0.32	(0.001)	0.81	(0.005)

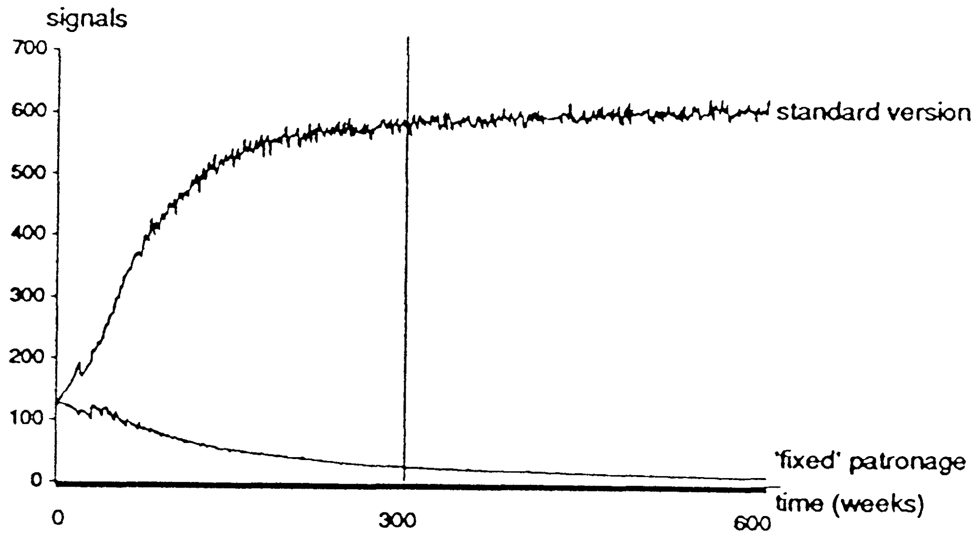

Fig. 11. Time series signaling in two variants.

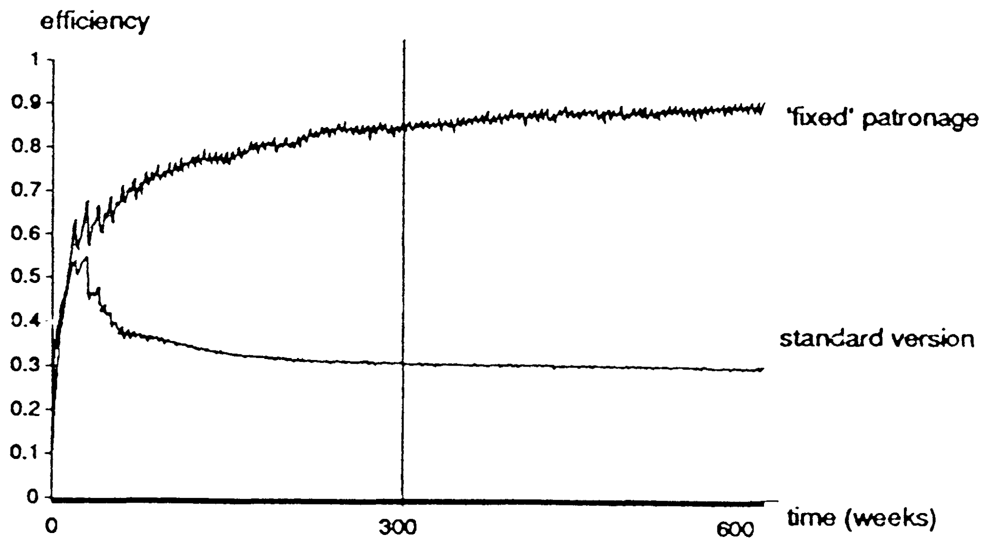

Fig. 12. Time series efficiency in two variants.

4.4. A VARIANT: 'FIXED' PATRONAGE

To illustrate the significance of the occurrence of self-organized markets, i.e. of the fact that all actions and outcomes in the simulated economy emerge purely *endogenously* as the result of locally interacting and learning individual agents who are all continuously looking for advantageous opportunities, we did five runs of a simulation in which we exogenously fixed one aspect of the behavior of the individual agents. That is, we impose that if a consumer has been satisfied by a firm then he will patronize next day surely. Table VII gives the average market experiences of firms and consumers for both the standard model and the variant with '*fixed*' patronage, averaged over the five runs.

We see that the apparently small, and intuitively reasonable, modification in the choice menu of the consumers leads to very different market outcomes. In the variant with '*fixed*' patronage, firms realize higher sales with a lower average production level, implying that less unsold stocks perish, and spend much less resources on signaling. As a result, the firms obtain a large profit increase. Figure 11 shows that firms perceive almost immediately that it is advantageous to limit their signaling activity.[29] Remember that firms do not know anything about the choice procedures followed by the consumers, and that they do not know which of their consumers are patronizing and which are new clients caught by an information signal. It is merely that they perceive sending more signals to be a waste of resources. This will be related to the fact that almost every consumer patronizes almost always. As the consumers are on average, slightly so, better off as well, the overall efficiency of the economy is much higher, although in figure 12, we observe that during the first phase of about 20 weeks the efficiency curves for the standard version and the variant are almost inextricable.[30] The main source of this improvement is the substantial saving of communication expenditures in the variant with '*fixed*' patronage.[31]

Note that the only difference with the standard version of the model is a *restriction* on the consumers' behavior. That is, all the actions and favorable outcomes of the variant with '*fixed*' patronage are also feasible in the unrestricted standard version. The reason that an individual consumer does not patronize always when he is satisfied in the standard version, is that it would not be rational to do so. Only when patronage is '*fixed*', and both the firms and all the other consumers change their behavior too, patronage in case of satisfaction will make the consumers better off on average. Furthermore, in the standard version without '*fixed*' patronage, firms could decide as well to signal only very scarcely. Then, consumers would be more or less forced to patronize, making everybody better off, and in particular the firms. Again, the point is that for an individual firm that would not be rational.[32]

5. Conclusions

The main conclusion to draw from this paper, is that this kind of approach, assuming that the agents' perception of opportunities is an endogenous process, and applying

Artificial Intelligence techniques to model the agents' actions, is promising. The data generated by the simulation contain regularities, which can be analyzed almost without limits of data availability.

After an initial phase of '*overall*' learning, the macroeconomic situation is characterized by comparatively steady aggregates. Competition appears to lead to coordination of economic activities, communication by firms and patronage by consumers play an important role herein, and the high communication expenditures are the main source of macroeconomic inefficiency. The microeconomic distributions underlying those aggregates show strong differences between the market shares of firms and between the shopping behavior of consumers. However, all firms offer an identical service rate, and the costs incurred per unit sold are also identical for all firms, while the rate of success is equal for all consumers. The emerging regularities in the agents' actions and outcomes show that the market process does tie down the set of possible beliefs of the agents, and does constrain their actions. Moreover, we have not only shown how self-organized markets may emerge in a decentralized economy, but we have also illustrated the essential difference between self-organized markets and organized markets.

Acknowledgements

I wish to thank Masanao Aoki, Brian Arthur, Chris Birchenhall, Thomas Brenner, Tim Cason, Raja Das, Pierre Dehez, John Duffy, Steven Durlauf, Dan Friedman, Maarten Janssen, Alan Kirman, Blake LeBaron, Reinhard Selten, Robert Waldmann, and seminar and conference participants in Amsterdam, Austin, Bonn, Florence, Los Angeles, Maastricht, Nafplio, Pittsburgh, Quebec City, Santa Fe and Tilburg for their critical comments and helpful discussions with respect to previous versions of this paper. The criticisms of two anonymous referees led to considerable improvements. All errors and responsibilities are mine. Much of the work for this paper was done at the European University Institute in Florence. The Santa Fe Institute's stimulating environment, a grant from the Niels Stensen Foundation, and financial support from the Deutsche Forschungsgemeinschaft through Sonderforschungsbereich 303 at the University of Bonn, are also gratefully acknowledged.

Appendix A. Outline of Classifier System and Genetic Algorithm

A.1 A CONSUMER

As the set of rules for a consumer is complete, only a Classifier System is used, in order to determine which rule is active on each day, and to update the strengths of the rules. Initially, the strength of each rule is 0.50. Eventually, all strengths are a number between 0 and 1. On each day, all rules for which the '*if...*' part is fulfilled participate in a '*stochastic auction*'. Using two coefficients, $b_1 = 0.10$ and $b_2 = 0.10$, each valid rule makes the bid $:= b_1 \cdot \text{strength} + \epsilon$, with $\epsilon \simeq$

$N(\mu = 0, \sigma = 0.00875)$. A given bid is disregarded with probability 0.025. The highest bidder wins the right to be active, and has to '*pay*' ($b_1 \cdot$ strength), to be subtracted from its strength. The active rule is reinforced according to: strength := strength + $b_1 \cdot (1 - b_2) \cdot$ income, where the income is 1 if the consumer, using that rule, succeeds in buying a unit of the commodity, and 0 otherwise. The rule that was active on the preceding day is reinforced receiving part of the currently active rule's raw bid: strength[previously active rule] := strength[previously active rule] + $b_2 \cdot (b_1 \cdot$ strength[active rule]). This system of strength transfers implies that the strengths converge towards the income they generate, each strength being a weighted average of past payoffs.

One issue of arbitrariness in modeling the consumers' actions has to be faced. Suppose all firms are equally advantageous for a consumer. That is, it does not matter whether he patronizes his previous supplier, or visits any other firm. A problem, then, is that we as economists know nothing about the decision procedures used by the consumers. One possibility would be, that such a consumer puts the names of all firms, including his last supplier, in an urn, and draws one name from this urn. However, this is not necessarily the right representation of the case in which all sellers are perceived to be equally attractive. The following two-stage decision procedure might be equally obvious to be followed. First, the consumer decides whether to patronize or to change firm, and second, he chooses a different firm at random when the option of changing had been chosen. Then, in the absence of any reason to patronize, this decision procedure would give a patronage rate of 50%. We implement some mediation between these two extremes, by letting the rules of a consumer's CS that imply visiting a signaling firm making two bids. Clearly, this means that one has to be careful when attaching a meaning to the absolute values of the probabilities with which the various rules are chosen.

A.2 A Firm

A firm uses a Classifier System, in order to determine which rule is active on each day, and to update the strengths of the rules, and a Genetic Algorithm in order to generate new rules.

Classifier System

The initial rules are 20 points chosen from a uniform random distribution on the domain [0-255, 0-255] in (production, signaling)-space. Each rule refers simply to a single point in (production, signaling)-space. In fact, each rule can be considered as an '*if...then...*' rule in a standard CS, where the conditional '*if...*' component is always satisfied, such that only the action or '*then...*' component remains. Initially, the strength of each rule is 0.30. Eventually, all strengths are a number between 0 and 1. On each day, all rules participate in a '*stochastic auction*'. Using two coefficients, $b_1 = 0.25$ and $b_2 = 0.40$, each valid rule makes the bid := $b_1 \cdot$ strength + ϵ, with $\epsilon \simeq N(\mu = 0, \sigma)$, where σ decreases during each generation from 0.075 to 0.03. A given bid is disregarded with probability 0.025. The highest bidder wins the right to be active, and has to '*pay*' ($b_1 \cdot$ strength), to be subtracted from its strength. The active rule is reinforced according to: strength := strength + $b_1 \cdot (1 - b_2) \cdot$ income, where the income is a number between 0 and 1 as a function of the external rewards. That is, the firm's actual profits are scaled to [0, 1], taking into account its profits experienced during the last 200 days. The rule that was active on the preceding day is reinforced receiving part of the presently active rule's raw bid: strength[previously active rule] := strength[previously active rule] + $b_2 \cdot (b_1 \cdot$ strength[active rule]). This system of strength transfers implies that the strengths converge towards the income they generate, each strength being a weighted average of past payoffs.

Genetic Algorithm

After each day, the value for production in the last activated rule is adapted taking into account the demand generated on that day by that rule. The production value of that rule is moved towards the demand faced (on average 10% of the distance between the production and the demand faced). The strength of the rule remains the same.

After each 50 days, the main genetic algorithm is applied. This may be triggered earlier (after 25 days) if in the most recent fourth of the current generation cycle the average performance is more than 20% worse than in the corresponding fourth of the preceding cycle. It can also be triggered after 25 days with a probability depending upon the number of rules with a strength smaller than the initial strength of 0.30.

The genetic algorithm generates 1 new rule on the basis of two existing strong rules. The '*parents*' are randomly chosen from the strongest 25% of rules, taking as selection probability for each rule its strength proportional to the total strength of these 25%. The rules to be replaced are chosen among the worst 50% of the rules: more precisely, the one that is most similar to the newly created rule, to avoid crowding of too similar rules. Duplicate rules are not allowed for.

Values are encoded in strings using the binary alphabet, alternating the bits for production and signaling, and ordering both in the same direction from '*high*' to '*low*'. For example, the string 101011 would be (production, signaling) = (7, 1). After each generation cycle, the length of the production and/or signaling sub-string can be increased or decreased with 1 bit, taking into account the production and signaling values of the strongest strings. The minimum length of the complete string is 5, while the maximum length of each sub-string is 10.

The genetic operators applied are reproduction, crossover and mutation. Reproduction and crossover take place as follows: take two strings, place them parallel to each other, and determine two crossing points randomly: $\begin{smallmatrix}\square\square\square\\ \blacksquare\blacksquare\blacksquare\end{smallmatrix} \times \begin{smallmatrix}\square\square\square\square\square\square\\ \blacksquare\blacksquare\blacksquare\blacksquare\blacksquare\blacksquare\end{smallmatrix} \times \begin{smallmatrix}\square\square\square\square\square\\ \blacksquare\blacksquare\blacksquare\blacksquare\blacksquare\end{smallmatrix}$, cross the two strings and choose one of the two '*children*', □□□■■■■■■□□□□□ or ■■■□□□□□□■■■■■, at random. Mutation means that any bit independently in the newly created string can switch from 0 to 1 or from 1 to 0. Depending upon the degree of convergence of the strings in the best 25% of the agent's rules, the probability of mutation of each bit varies between 0.10 and 0.001; higher convergence implying higher mutation rates. The strength of the new rules is the average of the strengths of the 2 '*parents*'.

Appendix B. Formal Analysis

f	patronage rate of satisfied consumers
g	price minus '*marginal*' cost of production
k	'*marginal*' cost of signaling
m	number of firms
n	number of consumers
N	total number of agents
q	demand directed towards firm i in period t
q^-	demand directed towards firm i in period $t-1$
Q	aggregate demand by all firms in period t
Q^-	aggregate demand by all firms in period $t-1$
s	number of signals sent by firm i in period t
S	aggregate number of signals sent by all firms in period t
S_-	aggregate number of signals sent by all other firms in period t
V	profit of firm i in period t
x	transactions of firm i in period t
z	production of firm i in period t

We limit our attention to the question whether there exists a stationary symmetric strategy that would give the highest immediate profits to firm i. With a deterministic demand function, a firm will choose $z = q$. Hence, z is simply a function of s, a strategy for a firm reduces to a number s, and we have $z = q = x$. $V = g \cdot q - k \cdot s$, where $q = (f \cdot q^-) + (s/S) \cdot \{1 - \exp(-S/N)\} \cdot (n - f \cdot Q^-)$.[33]

First, we determine all symmetric strategies that would give a stationary market. Sum the demand over all firms and take $Q = Q^-$: $Q = f \cdot Q + \{1 - \exp(-S/N)\} \cdot (n - f \cdot Q^-) \Rightarrow Q \cdot [1 - f - \{1 - \exp(-S/N)\} \cdot -f] = n \cdot \{1 - \exp(-S/N)\} \Rightarrow Q = [n \cdot \{1 - \exp(-S/N)\}]/[1 - f \cdot \exp(-S/N)]$. This is an increasing function of S.

Second, we determine for firm i the optimal s, and then impose that the strategy is symmetric and stationary. The derivative with respect to s is: $\mathrm{d}V/\mathrm{d}s = g \cdot (n - f \cdot Q^-) \cdot [S_-/S^2 \cdot \{1 - \exp(-S/N)\} + (s/S) \cdot (1/N) \cdot \exp(-S/N)] - k = 0$. Symmetry: $S_- = \{(m-1)/m\} \cdot S$. Stationarity: $Q^- = Q$. Hence, $(n - f \cdot Q) \cdot [(m-1)/(m \cdot S)] \cdot \{1 - \exp(-S/N)\} + (1/m) \cdot (1/N) \cdot \exp(-S/N)] = k/g \Rightarrow Q = (1/f) \cdot \{n - k/g \cdot 1/[(m-1)/(m \cdot S) \cdot \{1 - \exp(-S/N)\} + (1/m) \cdot (1/N) \cdot \exp(-S/N)]\}$. This is a decreasing function of S. Hence, one can solve numerically for which S the first-order condition is fulfilled.

Notes

[1] Clearly, the widespread use of the concept of a '*representative agent*' has eclipsed such questions for some time (see Kirman [1992].

[2] Three recent papers considering this endogenous determination of interactions as well, are Mailath et al. [1993], Durlauf [1994] and Stanley et al. [1994].

[3] See e.g., Marimon et al. [1990], Andreoni & Miller [1991], Rust et al. [1992], or Arthur et al. [1994].

[4] The main reason to assume fixed prices is keep the model as simple as possible. There are, however, also empirical examples of decentralized markets with fixed prices; e.g., newspapers in most countries, or bread in Italy.

[5] Implicitly, this assumes a specific choice of the parameters of the storage technology; infinite costs of storage. Just as with respect to any parameter, it would be interesting to consider the effect of different values.

[6] We abstract from the question under which circumstances this would be the case.

[7] This does not restrict the nature of the problem in any meaningful sense; it simply changes the value of some of its parameters. The more general scenario would be one in which consumers wander continuously about, searching for trading opportunities, making more visits during each day, also on more markets to buy various commodities, while the '*shopping*' behavior would not be synchronized between the consumers. Let us assume that time flows continuously, that the visits and exchanges are discrete events of zero duration, like the arrivals in a Poisson process (see Foley [1975] or Diamond [1982]), associate with each agent a '*random clock*' that rings independently for each agent at the instances of a Poisson process, and let each consumer make a visit when his clock rings (see the theory of interacting particle systems, e.g., Griffeath [1979]). The one-visit-only assumption is just a specific parameter choice with respect to this Poisson process plus a slight simplification.

[8] In reality, firms will often only observe how much they have actually sold. Then, in case a firm has some units in stock left at the end of the day it can calculate the total demand directed towards it. But in case of a stock out it does not even know this for sure and has to estimate it, because a firm with empty shelves simply closes for the rest of the day (see e.g., Alpern & Snower [1988]). We abstract from such estimation procedures.

[9] Including the information consumers might get through the information signals by which the firms try to attract their attention.

[10] For details concerning the CS/GAs we implemented, see appendix A.

[11] Sat = satisfied; Late = arrived late at firm; Mist = not found a firm; Info = information signals received; ¬Info = no information signals received; Patr = patronage; Known = visit firm known from signal; Rand = visit random, unidentified agent.

[12] Roth & Erev's model is basically a CS, updating the propensities to choose from a given set of possible actions, without the experimentation implied by the use of a GA. No new actions need to be created in their model, as the action space could easily be covered by a CS.

[13] See Vriend [1994] for a more extensive discussion of these issues.

[14] Clearly, the emergent behavior and self-organization *are* a function of the underlying configuration. The relevant point is, however, the following. Given a certain model with a certain parametrization, can one reason, i.e., without running a simulation, *which* functions of the parametrization the outcomes are?

[15] See e.g., Zarnowitz [1985].

[16] For presentational reasons, in all time series we aggregated (averaged) the daily observations to weekly observations; one '*week*' covering 5 days. Moreover, as the numbers of firms and consumers are constant, where appropriate, we will express the variables as averages over firms resp. consumers. For each observation we calculated the standard deviation over the five runs. The average of these standard deviations of each series is given in footnotes. The graphs of the individual runs can be requested from the author.

[17] The average standard deviations over the five runs of production, demand face, and transactions are 2.43, 0.04 resp. 0.87. For profits, unsold stocks, and signaling costs per unit sold these average

standard deviations are 0.011, 0.005 resp. 0.008. The average standard deviation over the five runs of patronage is 0.006, of the service rate 0.007, and of efficiency 0.004.

[18] In all tables, we present the statistics of the *average* run. The standard deviations over the five runs of these statistics are, in each table, in parentheses. Additional statistics of the individual runs, for this and subsequent tables, can be requested from the author.

[19] The average standard deviation over the five runs of this variable is 1.47.

[20] Clearly, with such large differences and so many observations, these differences form a regularity, in the sense that they are not just caused by some stochastic noise. For example, with respect to the variable profits, an analysis of variance of the matrix of 3000 days $\times$ 50 firms, testing for the significance of column effects (i.e., of the factor '*firms*') yields, averaged over the five runs, $F(49, 146902) = 233.9$, with a standard deviation of 24.95.

[21] The average standard deviation over the five runs of this variable is 1.28.

[22] The average standard deviation over the five runs for the final distributions for the ten classes (0–19, 20–39,...) are 0.4, 1.4, 1.5, 3.4, 2.0, 1.5, 2.0, 1.2, 0.0, and 0.0.

[23] Figure 8, as well as figure 9, concerns one run only, because averaging over the runs would lead to a loss of information. Graphs for the other runs can be requested from the author.

[24] An analysis of variance of the matrix of 50 quarters (on '*quarter*' being 12 weeks, or 60 days) $\times$ 5000 consumers testing for the significance of column effects (i.e., of the factor '*consumers*') yields, averaged over the five runs, $F(4999, 244951) = 0.99$, with a standard deviation of 0.020. Hence, there is no significant difference between the consumers in their success over time. But this depends much on the point of view one takes. It matters which of the two possible market outcomes one takes for granted and considers as '*natural*', and which outcome one considers as a noteworthy deviation. If we express the consumers' outcomes not as a satisfaction rate, but as a disappointment rate, then we get a minimum of 0.04 and a maximum of 0.08. That is, the '*worst*' consumer turns home disappointed twice as often as the '*best*'. As we didn't specify the commodity and the consumers' preferences further, we can't quantify the loss for a disappointed consumer in a given day. It may very well be that from the '*worst*' consumer's point of view this difference is not negligible at all. Having put it this way, it is remarkable, at least as a curiosity, that also the '*worst*' and '*best*' firms, averaged over the whole history, differed a factor 2.

[25] As table I shows, the consumers were indifferent in this respect at the start. For example, if they were disappointed by a firm, or, even more so, if they had visited a randomly chosen, unidentified agent, and by chance contacted another consumer instead of a firm, then they opted to patronize with the same initial probability as any other option.

[26] In a more general setting, the service rates are just one of the possible qualitative aspects of the product that a firm offers. Other aspects might be the price of the commodity, or intrinsic qualitative characteristics of the good. In our model these aspects are identical for all firms. As far as the analysis of patronage is concerned, in our model patronage leads straightaway to arbitrage, whereas with respect to the other mentioned reasons to patronize, this relation will be indirect.

[27] The average standard deviation over the five runs of the patronage rate for satisfied consumers is 0.006, and for disappointed consumers 0.010.

[28] That strict patronage in a positive sense does not occur more frequently is related to the facts that our consumers patronize only when they perceive this to be advantageous, that all firms on average turned out to offer the same service rate, and that patronage does not involve any kind of preferential treatment by the firm.

[29] The average standard deviation over the five runs of signaling in the standard version is 12.93, and of signaling with '*fixed*' patronage 3.72.

[30] The average standard deviation over the five runs of efficiency with '*fixed*' patronage is 0.013.

[31] It might be that this leaves the consumers in the '*fixed*' patronage variant more vulnerable, e.g., to the shock of a firm that exits the market.

[32] The results of this variant show that one has to be very careful in modeling the shopping behavior of consumers, as very small modifications may lead to completely different market outcomes. In particular, it is not innocuous to bias models towards patronage; either by assuming it right away (e.g., Bergmann [1989]) or by imposing ad hoc additional costs on non-patronage (e.g., Sutton [1980]).

[33] See Vriend [1993] for the derivation of such a demand function in a closely related framework.

References

1. Alpern, S., & Snower, D.J. (1988), "High-Low Search" in Product and Labor Markets. *American Economic Review, 78*, 356–362.
2. Anderlini, L., & Sabourian, H. (1988). *Some Notes on the Economics of Barter, Money and Credit* (Economic Theory Discussion Paper No. 128) University of Cambridge.
3. Anderson, P.W., Arrow, K.J., & Pines, D. (Eds.). (1988), *The Economy as an Evolving Complex System. The Proceedings of the Evolutionary Paths of the Global Economy Workshop, Held September, 1987 in Santa Fe, New Mexico* (Vol. V, Santa Fe Institute Studies in the Sciences of Complexity). Redwood City, CA: Addison-Wesley.
4. Andreoni, J., & Miller, J.H. (1991). *Auctions with Adaptive Artificially Intelligent Agents* (Working Paper 91-01-004) Santa Fe Institute.
5. Arthur, W.B., Holland, J., LeBaron, B., Palmer, R., & Taylor, P. (1994). *An Artificial Stock Market* (mimeo).
6. Bak, P., Chen, K., Scheinkman, J., & Woodford, M. (1993). *Aggregate Fluctuations from Independent Sectoral Shocks: Self-Organized Criticality in a Model of Production and Inventory Dynamics* (Working Paper 93-01-004) Santa Fe Institute.
7. Bergmann, B.R. (1989), A Microsimulated Model of Inventories in Interfirm Competition. *Journal of Economic Behavior and Organization, 14*, 65–77.
8. Blume, L.E. (1993). The Statistical Mechanics of Strategic Interaction. *Games and Economic Behavior, 5*, 387–424.
9. Bush, R.R., & Mosteller, F. (1955). *Stochastic Models for Learning*. New York: Wiley.
10. Cross, J.G. (1983). *A Theory of Adaptive Economic Behavior*. Cambridge: Cambridge University Press.
11. Diamond, P.A. (1982). Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy, 90*, 881–894.
12. Durlauf, S.N. (1990). *Locally Interacting Systems, Coordination Failure, and the Behavior of Aggregate Activity* (mimeo Dept. of Economics) Stanford University.
13. Durlauf, S.N. (1994). *Neighborhood Feedbacks, Endogenous Stratification, and Income Inequality* (mimeo).
14. Ellison, G. (1993). Learning, Local Interaction, and Coordination. *Econometrica, 61*, 1047–1071.
15. Föllmer, H. (1974). Random Economies with Many Interacting Agents. *Journal of Mathematical Economics, 1*, 51–62.
16. Foley, D.K. (1975). On Two Specifications of Asset equilibrium in Macroeconomic Models. *Journal of Political Economy, 83*, 303–324.
17. Griffeath, D. (1979). *Additive and Cancellative Interacting Particle Systems. (Lecture Notes in Mathematics 724)*. Berlin: Springer.
18. Harley, C.B. (1981). Learning the Evolutionarily Stable Strategy: *Journal of Theoretical Biology, 89*, 611–633.
19. Hart, A.G. (1951). *Anticipations, Uncertainty and Dynamic Planning* (2nd ed.). New York: Kelley.
20. Holland, J.H. (1986). Escaping Brittleness: The Possibilities of General-Purpose learning Algorithms Applied to parallel Rule-Based Systems. In R.S. Michalski, J.G. Carbonell & T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach (Vol. 2)* (pp. 593–623). Los Altos, CA: Morgan Kaufmann.
21. Holland, J.H. (1992). *Adaptation in Natural and Artificial Systems. An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (2nd ed.). Cambridge, MA: MIT Press.
22. Kandori, M., Mailath, G.J., & Rob, R. (1993). Learning, Mutation, and Long Run Equilibria in Games. *Econometrica, 61*, 29–56.
23. Kirman, A. (1992). Whom or What Does the Representative Individual Represent? *Journal of Economic Perspectives, 6, No. 2*, 117–136.

24. Klamer, A. (1984). *The New Classical Macroeconomics. Conversations with New Classical Economists and their Opponents*. Brighton: Wheatsheaf.
25. Lane, D.A. (1992). *Artificial Worlds and Economics* (Working Paper 92-09-048) Santa Fe Institute.
26. *Machine Learning* (1988). *3, Nos. 2/3, Special Issue on Genetic Algorithms.*
27. *Machine Learning* (1992). *8, Nos. 3/4, Special Issue on Reinforcement Learning.*
28. Mailath, G., Samuelson, L., & Shaked, A. (1993). *Correlated Equilibria as Network Equilibria* (Discussion Paper No. B-244) University of Bonn.
29. Marimon, R., McGrattan, E., & Sargent, T.J. (1990). Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents. *Journal of Economic Dynamics and Control, 14*, 329–373.
30. Roth, A.E., & Erev, I. (forthcoming). Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behavior.*
31. Rust, J., Palmer, R., & Miller, J.H. (1992). *Behavior of Trading Automata in a Computerized Double Auction Market* (Working Paper 92-02-008) Santa Fe Institute.
32. Shubik, M. (1988). *A Game Theoretic Approach to the Theory of Money and Financial Institutions* (Lecture notes) Siena International School of Economic Research.
33. Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations.* London: Strahan and Cadell.
34. Stanley, E.A., Ashlock, D., & Tesfatsion, L. (1994). Iterated Prisoner's Dilemma with Choice and Refusal of Partners. In C.G. Langton (Ed.), *Artificial Life III* (Vol. XVII, Santa Fe Institute Studies in the Sciences of Complexity, pp. 131–175). Redwood, CA: Addison-Wesley.
35. Sutton, J. (1980). A Model of Stochastic Equilibrium in a Quasi-Competitive Industry. *Review of Economic Studies, 47*, 705–722.
36. Sutton, R.S. (1992). Introduction: The Challenge of Reinforcement Learning. *Machine Learning, 8, Nos. 3/4*, 225–227.
37. Tversky, A., & Kahneman, D. (1986). Rational Choice and the Framing of Decisions. *Journal of Business, 59, No. 4, Pt. 2*, S251–S278.
38. Vriend, N.J. (1993). *Essays on Decentralized Trade* (Ph.D. thesis) European University Institute, Florence.
39. Vriend, N.J. (1994). Artificial Intelligence and Economic Theory. In E. Hillebrand & J. Stender (Eds.), *Many-Agent Simulation and Artificial Life* (pp. 31–47). Amsterdam: IOS.
40. Young, H.P. (1993). The Evolution of Conventions. *Econometrica, 61*, 57–84.
41. Zarnowitz, V. (1985). Recent Work on Business Cycles in Historical Perspective: A Review of Theories and Evidence. *Journal of Economic Literature, 23*, 523–580.