ELSEVIER

Journal of Economic Dynamics & Control
25 (2001) 281–293

JOURNAL OF
Economic
Dynamics
& Control

www.elsevier.nl/locate/econbase

# Introduction to the special issue on agent-based computational economics

Leigh Tesfatsion[1,*]

*Department of Economics, Iowa State University, Ames, Heady Hall 375, IA 50011-1070, USA*

Accepted 22 February 2000

**Abstract**

A brief overview of agent-based computational economics (ACE) is given, followed by a synopsis of the articles included in this special issue on ACE and in a companion special issue on ACE scheduled to appear in *Computational Economics*. © 2001 Elsevier Science B.V. All rights reserved.

*JEL classification:* B4; C6

*Keywords:* Agent-based computational economics

## 1. Agent-based computational economics

Agent-based computational economics (ACE) is the computational study of economies modelled as evolving systems of autonomous interacting agents. ACE is thus a specialization to economics of the basic complex adaptive systems paradigm (Holland, 1992).

One principal concern of ACE researchers is to understand why certain global regularities have been observed to evolve and persist in decentralized market economies despite the absence of top-down planning and control: for example,

[1] http://www.econ.iastate.edu/tesfatsi/

* Tel.: + 515-294-0138; fax: 515-294-0221.

*E-mail address:* tesfatsi@iastate.edu (L. Tesfatsion).


PII: S0165-1889(00)00027-0

trade networks, socially accepted monies, and market protocols. The challenge is to demonstrate *constructively* how these global regularities might arise from the bottom up, through the repeated local interactions of autonomous agents acting in their own perceived self-interest.

A second principal concern of ACE researchers is to use ACE frameworks normatively, as computational laboratories within which alternative socio-economic structures can be studied and tested with regard to their effects on individual behavior and social welfare. This normative concern complements a descriptive concern with actually observed global regularities by seeking deeper possible explanations not only for why certain global regularities have been observed to evolve but also why others have not.

The ACE focus on economies as self-organizing systems is not new, of course; it clearly follows in the tradition of Smith (1937) and Hayek (1948). Moreover, the ACE concern with constructive demonstration reflects the strong influence of researchers such as Schelling (1978), Axelrod (1984), and Arthur (1994) who use simple but carefully crafted choice problems to study the specific processes by which social order can emerge from self-interested micro behavior.

The ACE view of economies as evolving systems is by no means new either. Even before Darwin, attempts were made to apply evolutionary ideas to socio-economic behavior. Although this earliest work is now rarely cited by economists, most have had at least some exposure to the seminal work by Schumpeter (1942) on the evolution of economic institutions and by Alchian (1950) on uncertainty and evolution in economic systems. Moreover, as detailed in Nelson (1995), this early work on evolutionary economics appears to have strongly influenced subsequent work on evolutionary theories of economic change by such well-known researchers as Richard Day, Jack Hirshleifer, Richard Nelson, Sidney Winter, and Ulrich Witt. In addition, some researchers have been exploring the potential economic applicability of evolutionary game theory with replicator dynamics (see Samuelson, 1997).

What is new about ACE is its exploitation of powerful new computational tools, most notably object-oriented programming. These tools permit ACE researchers to extend previous work on economic self-organization and evolution in four key ways.[2]

[2] The best single introduction to the ACE methodology to date remains Epstein and Axtell (1996). In this ground-breaking monograph the authors construct a simple agent-based world, Sugarscape, within which they are able to grow economies, epidemics, tribal formations, large-scale agent migrations, and other social phenomena reminiscent of human societies. Interested readers can also visit the ACE web site at http://www.econ.iastate.edu/tesfatsi/ace.htm. Resources available at this continually updated site include surveys, an annotated syllabus of readings, software, teaching materials, and pointers to individual researchers and research groups.

First, economic worlds can be computationally constructed that are populated with heterogenous agents who determine their interactions with other agents and with their environment on the basis of internalized social norms, internal behavioral rules, and data acquired on the basis of experience. Consequently, these agents have a richer internal cognitive structure and more autonomy than conventionally modelled economic agents.

Second, a broad range of agent behaviors and interactions can be permitted in these economic worlds, with predatory and cooperative associations taking center stage along with price and quantity relationships. Agents continually adapt their behavior in response to agent–agent and agent–environment interactions in an attempt to satisfy their needs and wants. That is, behavioral rules are state conditioned, and agents co-adapt their behavior in an intricate dance of interactions. The economic worlds can therefore exhibit self-organization.

Third, the evolutionary process can be represented as natural selection pressures acting directly on agent behavioral attributes rather than as population-level laws of motion. These natural selection pressures induce agents to engage in continual open-ended experimentation with new rules of behavior. That is, agents in the economic worlds co-evolve.

Fourth, the economic worlds can be grown along a real time-line, observed but not disturbed, much like a culture grows in a petri dish. Once initial conditions are set by the modeller, all subsequent events in these worlds can be initiated and driven by agent-agent and agent-environment interactions without further outside intervention.

In brief, then, ACE is a methodology that blends concepts and tools from evolutionary economics, cognitive science, and computer science in a manner that may ultimately permit three important developments: (a) The constructive grounding of economic theories in the thinking and interactions of autonomous agents; (b) the testing, refinement, and extension of these theories through careful computational experiments, statistical analysis of findings, and appropriate comparisons with analytical studies, econometric studies, field studies, and human-subject laboratory studies; and (c) the formulation and testing of conceptually integrated socioeconomic theories compatible with theory and data from many different relevant fields of social science currently separated by artificial disciplinary boundaries.

As with any new methodology, initial excitement over possibilities must give way to careful research that demonstrates more concretely both its advantages and its limitations. This special issue on ACE, together with the companion special issue on ACE scheduled to appear in *Computational Economics*, will give readers a chance to judge for themselves the potential of the ACE methodology and the extent to which convincing results have been achieved to date. To help guide the reader through these special issues, the following section provides an article synopsis.

## 2. Article synopsis

Six of the 18 articles accepted for the two special ACE issues use stylized problem contexts to explore general economic concerns such as the evolution of norms, self-organization, and agent learning. The remaining 12 articles focus on efficiency and welfare concerns for particular types of markets: namely, financial markets, labor markets, retail fish markets, business-to-business markets, electricity markets, entertainment markets, automated markets, e-commerce, and cattle pasture access markets in North Cameroon. The six articles in the first category are included in the special issue scheduled to appear in *Computational Economics* while the remaining 12 articles are included in the special issue at hand.

A more detailed overview of the articles included in the special ACE issue at hand will now be given.

Understanding the causal connections relating structure, behavior, and welfare outcomes in markets comprised of boundedly rational agents who learn imperfectly from the past has resurfaced as a major focus of economic research in recent years, spurred on by influential studies such as Kreps (1990), Sargent (1993), Krugman (1996), Arthur et al. (1997), and Young (1998). To date, however, definitive results have been difficult to achieve due to the complexity of these connections. The authors writing for this special ACE issue attempt to shed some light on these connections for various specific types of markets by constructing agent-based computational laboratories within which alternative hypotheses can be experimentally studied.

In the lead-off article, John Duffy focuses on the well-known search model of money by Kiyotaki and Wright (1989). Previous experimental studies of this model using both real and artificial agents have found that convergence to the speculative equilibrium is by no means ensured, even when the parameters of the model are set to values for which the speculative equilibrium is the unique Nash equilibrium of the model. Since the designs of the experiments with real and artificial agents differ considerably from each other, however, it is difficult to make a direct comparison of their findings.

Duffy first describes a set of experiments with real and artificial agents that were both conducted within similar simplified versions of the Kiyotaki–Wright model. The behavioral rules used by the artificial agents to conduct their trades are modeled on the basis of evidence obtained from the human subject experiments. The artificial agents adaptively select among their feasible behavioral rules by means of a simple form of reinforcement learning. Duffy reports that the findings for the artificial agent experiments match basic features of the findings for the real agent experiments.

Duffy then uses the findings from the artificial agent experiments to predict what might happen in two modified versions of the Kiyotaki–Wright model that are designed to encourage greater speculative behavior by certain player types.

A key prediction is that both modifications of the Kiyotaki–Wright model will serve to increase the speed with which players learn to adopt speculative strategies and hence the likelihood of convergence to the speculative equilibrium. Actual experiments are then run for the two new versions of the Kiyotaki–Wright model using real agents, with encouraging results: the findings from the experiments with real agents are roughly similar to those predicted by the artificial agent experiments.

Various regularities observed in actual common stock returns (e.g., kurtosis) are not well explained by existing asset market models. Nicholas Tay and Scott Linn conjecture that better explanatory power might be obtained by allowing agents in otherwise traditional asset market models to form their expectations in accordance with the way investors form their expectations in real life: namely, in fuzzy terms using inductive reasoning. They argue that these features can be faithfully captured by a genetic-fuzzy classifier system, a modification of the classifier system developed by Holland (1992).

To test their claim, they modify the Santa Fe artificial stock market model studied by LeBaron et al. (1998) by permitting traders to form their expectations inductively using a genetic-fuzzy classifier system and by modifying the manner in which traders decide which prediction rules to rely on when making demand decisions. They report experimental findings that show that the asset prices and returns generated by their model exhibit characteristics, including measures of kurtosis, that are very similar to actual data.

Shu-Heng Chen and Chia-Hsuan Yeh argue that social learning in the form of imitation of strategies is important in stock market contexts, along with individual learning, but that standard modellings of stock markets do not include the mechanisms by which such social learning actually takes place. They construct an agent-based computational framework for the analysis of artificial stock markets that includes an additional social learning mechanism, referred to as a *school*.

Roughly, the school consists of a group of agents (e.g., business school faculty members) who are competing with each other to supply publicly the best possible models for the forecasting of stock returns. The success (fitness) of school members is measured by the current forecasting accuracy of their models, whereas the success of traders is measured in terms of their wealth. Traders continually choose between trading in the market and taking time out to visit the school. Once at school, a trader tests a sample of the forecasting models currently proposed by school members in an attempt to discover a model that is superior to the one he is currently using. The school and the traders co-evolve over time in an intricate feedback loop.

To test the implications of their stock market model, Chen and Yeh conduct an experiment consisting of 14,000 successive trading periods. One key finding is that market behavior never settles down; initially successful forecasting models quickly become obsolete as they are adopted by increasing numbers of agents.

Another key finding is that individual traders do not act as if they believe in the efficient market hypothesis even though aggregate market statistics suggest that the stock market is efficient.

Jasmina Arifovic uses a computational two-country overlapping generations model with a flexible exchange rate and no restrictions on foreign currency holdings to explore what happens when two countries issue competing currencies. The government of each country finances its deficit via seignorage, and the private citizens of each country use a genetic algorithm to continuously update their savings and portfolio decision rules. One issue is what happens to the exchange rate. A second related issue is whether the currency competition will impose fiscal and monetary discipline on government policy makers by favoring the currency of the government with the least inflationary policy.

Arifovic presents findings from computational experiments that show that a stationary rational expectations equilibrium in which both currencies are valued is unstable in her model economy. More precisely, the dynamics of the genetic algorithm updating always move the economy away from this equilibrium and towards a stationary single-currency equilibrium in which the currency used to finance the larger of the two deficits does not survive.

An important aspect of imperfectly competitive markets with strategically interacting agents is the manner in which agents determine their transaction partners and the form of the transaction networks that tend to evolve and persist over time.[3] Transaction networks are now frequently analyzed by means of transaction cost economics (Williamson and Masten, 1999). To date, however, this literature has not stressed the dynamics of learning, adaptation, and innovation, nor the development of trust. Instead, it is assumed that optimal forms of organization or governance will arise that are suited to the particular characteristics of agent transactions, such as the need for transaction-specific investments.

The next four articles focus on the endogenous formation of transaction networks for, respectively, a labor market, a retail fish market, an intermediate goods market, and a pasture access problem arising for nomadic cattle herdsmen and farmers in North Cameroon. In each of these studies, trust and the establishment of durable relations among agents play key roles.

Leigh Tesfatsion experimentally studies the relationship between structure, behavior, and market power in an agent-based computational labor market framework. Workers and employers repeatedly participate in costly searches for preferred worksite partners on the basis of continually updated expected utility, engage in efficiency-wage worksite interactions modelled as prisoner's dilemma games, and evolve their worksite strategies over time on the basis of the earnings

[3] For pointers to individual researchers and research groups who are currently studying this issue, visit the web site titled 'Formation of Social and Economic Networks' at http://www.econ.iastate.edu/tesfatsi/netgroup.htm.

secured by these strategies in past worksite interactions. Any dissatisfied worker can quit working for an employer by directing his future work offers elsewhere, and any dissatisfied employer can fire a worker by refusing to accept future work offers from this worker.

Tesfatsion constructs and uses various descriptive statistics to study experimentally determined correlations between market structural specifications and worker–employer network formations, and between network formations and the types of worksite behaviors and market power outcomes that these networks support. Two key market structural specifications of interest are job concentration and job capacity. She finds that job capacity consistently outperforms job concentration as a predictor of the relative market power of workers and employers. Surprisingly, holding job capacity fixed, changes in job concentration have only small and unsystematic effects on attained market power levels.

Alan Kirman and Nicolaas Vriend construct an agent-based computational model of the wholesale fish market in Marseilles that captures in simplified form the structural aspects of the actual fish market. Their objective is to understand two persistently observed features of the actual fish market: price dispersion; and widespread buyer loyalty to sellers in the form of repeat business. Each buyer and seller must make multiple decisions during each trading day regarding price, quantity, choice of trading partner, and treatment of trading partner (e.g., should a seller offer better deals to his more loyal buyers). Each of these decisions is separately modelled for each individual agent by means of a classifier system, a particular type of reinforcement learning algorithm developed by Holland (1992).

Kirman and Vriend report that, in experimental runs with their model, price dispersion and loyalty emerge as a result of the co-evolution of buyer and seller decision rules. For example, regarding loyalty, buyers learn to become loyal as sellers learn to offer a higher payoff to loyal buyers, while these sellers, in turn, learn to offer a higher payoff to loyal buyers as they happen to realize a higher payoff from loyal buyers. The authors provide a detailed discussion of the dynamic processes which underly the emergence of price dispersion and loyalty in their reinforcement learning framework.

Tomas Klos and Bart Nooteboom use an agent-based computational model to explore how transaction networks develop among buyer and supplier firms who repeatedly choose and refuse their transaction partners on the basis of continually updated anticipations of future returns. These anticipations depend in part on trust, where trust increases with the duration of a relationship, and in part on profitability. Buyer firms face a ‘buy or make’ decision: they can search for suppliers to obtain components for the production of differentiated products to be sold in a final goods market, or they can choose to produce these components themselves. Supplier firms engage in both specific and general-purpose asset investment tailored to the collection of buyer firms with whom they are transacting. Buyer firms can increase revenues by selling more differentiated products, and supplier firms can reduce input costs for buyer firms by generating learning-by-doing efficiencies for the buyer firms with whom they are in longer-term relationships.

The Klos and Nooteboom model permits an assessment of the efficiency of resulting profit outcomes as a function of trust and market conditions. The authors report illustrative computational experiments with alternative settings for the degree of differentiation among the buyers' products. As predicted by transaction cost economics, more product differentiation favors 'make' relative to 'buy' decisions due to higher switching costs and scale effects. Nevertheless, the path dependencies and uncertainties that arise for firms due to the ability to make and break relationships on the basis of past experience result in profit outcomes that are not always efficient.

Rouchier et al. are motivated by a field study focusing on seasonal mobility ('transhumance') among nomadic cattle herdsmen in North Cameroon. The field study explores the conditions that determine the access that nomadic herdsmen have to pasture lands. A key finding is that the grazing patterns and individual relationships established among herdsmen, village leaders, and village farmers tend to be very regular.

In an attempt to better understand these observed regularities, Rouchier et al. construct and use an agent-based computational framework to simulate the dynamics of the relationships among three agent types: nomadic herdsmen who need both water and grass for their cattle and who seek access to these resources from village leaders and farmers in return for access fees; village leaders who provide herdsmen with either good or poor access to water depending on their order of arrival; and village farmers who own pasture land that they may or may not permit the herdsmen to use for cattle grazing. Herd sizes evolve as a function of the agreements that are reached. The authors test two different models of reasoning for their agents: a 'cost priority' model based on ideas from transaction cost economics under which agents care only about minimizing their costs; and a 'friend priority' model based on ideas from institutional theory (North, 1990) under which agents also care directly about the stability of their relationships.

The authors run experiments in which the land in some villages randomly becomes unavailable for use as pasture for short periods of time, so that the farmers in these villages refuse all access requests from herdsmen during these periods. They show that the cost-priority and friend-priority models of agent reasoning result in dramatically different experimental outcomes. In particular, the global efficiency of the cost priority model is surprisingly low relative to the friend priority model, leading in some cases to the disappearance of herds. In explanation, the authors note that the cost priority model tends to result in less flexible agent behavior, and this in turn results in less robustness to land disruption shocks and more overgrazing of pasture lands. In reality, nomadic herdsmen are careful to sustain an extended social network of friends across a wide variety of villages through repeated interactions, and only the friend priority model produced such a pattern.

To date, research into alternative designs for auctions has largely focused on symmetric risk-neutral agents bidding competitively for a unit of an item in a one-period auction. In reality, however, many auctions involve oligopolistic sellers, asymmetric in size, who meet repeatedly and frequently and who determine their price and quantity offers strategically in an effort to exploit market power opportunities. As a case in point, electricity market auctions typically involve price and quantity offers for the sale of large amounts of bulk electricity by small numbers of electricity generators, some of whom have relatively large market shares.

John Bower and Derek Bunn construct an agent-based computational framework to study the following issue for the England and Wales wholesale electricity market: How would the government-proposed change from a uniform price auction to a discriminatory price auction affect prices for bulk electricity? The market is modelled as a repeated game among electricity generators with profit and market share objectives. In each trading period each generator submits to the auction a supply function expressing its price and quantity offers. Each generator is represented, at the level of its individual power plants, by a separate autonomous adaptive agent capable of evolving its supply strategy by means of a simple reinforcement learning algorithm. In contrast, agents on the demand side of the market are assumed to be passive price takers; their behavior is modelled by a fixed aggregate demand curve reflecting a standardised daily load profile corresponding to a typical winter day.

The main experimental finding of the authors is that the proposed change from a uniform to a discriminatory auction design significantly increases the ability of generators to exercise market power. In particular, when supply function offers are not publicly available, large generators gain a significant informational advantage over small generators under the discriminatory auction because they submit more offers and therefore can learn more precisely about the current state of the market. The uniform price auction mitigates this advantage by letting small generators share in the industry's collective learning by receiving the same market price for their electricity as any other generator. The authors conclude that, under certain circumstances, the choice of the auction design may actually be less important than simply ensuring that all auction participants have equal access to information, regardless of their size.

An information cascade is a sequence of decisions for which it is optimal for agents to imitate the selections of the agents who select before them. Two key questions arise for information cascades: can they occur; and are they fragile? For the special case of binary decisions the answer to each question is yes, but for more general types of decisions the answer to these questions is currently unknown.

Arthur De Vany and Cassey Lee construct an agent-based computational framework within which they explore the existence and fragility of information cascades under a variety of alternative structural specifications. This framework differs from standard information cascade models in two basic respects. First, each decision can involve a selection from among more than two options. Second, agents can receive local quality signals from neighboring agents in addition to global quantity information about the proportion of agents who have selected each option to date.

For concreteness, De Vany and Lee apply their model to the study of the dynamics of motion picture box office revenues. The experimental findings they report for this application indicate that the resulting information cascade dynamics can be quite complex. The authors' main finding is that multiple cascades can coexist in an intermittent pattern in which two or more intertwined cascades are observed to repeatedly alternate over time as the dominant cascade pattern. This intermittence makes it difficult to isolate individual cascades and to predict which if any of the competing cascades will ultimately win out. The authors argue that the complex dynamical patterns observed in their computational experiments resemble the irregular dynamics observed in actual time-series data for movie picture box office revenues.

In addition to saving labor time, automated contracting through artificial agents can increase search efficiency because artificial agents are often more effective at finding beneficial contractual arrangements in strategically complex settings with large strategy domains. To date, however, automated contracts have generally been binding contracts that limit the ability of the artificial agents to react to unforeseen events. Recently, the concept of a 'leveled commitment contract' has been proposed that permits agents to decommit from contracts by paying a monetary penalty to the contracting partner, but the efficiency of the resulting contracts depends heavily on the structuring of the penalties.

Martin Andersson and Tuomas Sandholm use an agent-based computational model of an automated negotiation system to experimentally study the sensitivity of leveled commitment contractual outcomes to changes in penalty structurings and to changes in the design of the artificial agent negotiators. Four types of penalties are considered: fixed; percentage of contract price; increasing based on contract start date, and increasing based on contract breach date. Agents differ by amount of look-ahead and by degree of self-interested behavior. Multiple task allocation problem instances are tested, with five negotiation rounds permitted for each instance. In all tested settings, the authors find that choosing relatively low but positive decommitment penalties works best. Surprisingly, however, the authors also find that self-interested myopic agents achieve a higher social welfare level, and more rapidly, than cooperative myopic agents when decommitment penalties are low. While a look-ahead capability improves agent performance, over short ranges of penalty parameters myopic agents perform almost as well.

Internet congestion leads to a large variability in the time required to complete electronic transactions. Given the increasing use of Internet transmissions for financial and other forms of e-commerce, the development of mechanisms that can ensure timely and reliable exchanges of data on the Internet is an important outstanding issue.

In earlier studies, Sebastian Maurer and Bernardo Huberman proposed a 'restart strategy' for carrying out Internet transactions that exploits the analogy between management of Internet delay-time mean and variance and the risk-return management of portfolios. They demonstrated that the use of their restart strategy by a single agent resulted in a faster average execution time with a smaller variance, given current Internet usage patterns. A potential problem with this strategy, however, is that it might be rendered useless if adopted and used by larger groups of agents.

In their current study, Maurer and Huberman construct an agent-based computational framework to explore the effects of restart strategy use by a group of agents deciding asynchronously whether to use the Internet. Each agent repeatedly faces a simple binary decision: to download a web page or not. The agents base their decisions on knowledge of congestion statistics for a past window of time. If the perceived current waiting time exceeds a certain threshold, the agent refrains from using the Internet; otherwise the agent attempts to download. The threshold values are determined optimally for each agent as a function of the congestion statistics and general problem constraints. The loads created by agents who decide to use the Internet determine the congestion statistics for subsequent use decisions.

Maurer and Huberman show that, for a range of problem parameters, all agents are at least as well off when they use the restart strategy as when they do not. They caution, however, that the good performance of their restart strategy with multiple agent use depends on their assumption that a server can efficiently detect multiple requests from the same user and cancel the superfluous ones in order to avoid sending the same data to a user multiple times.

## 3. Concluding remarks

Participants in real-world markets base their current behavior in part on their own past experiences and in part on perceived aggregate market characteristics that their past individual behaviors have conjointly helped to determine. This feedback loop can lead to intricate causal chains connecting structure, behavior, and welfare outcomes that are difficult to understand and predict by standard analytical and statistical tools.

The twelve articles included in this special ACE issue use agent-based computational frameworks as computational laboratories within which to study these market connections through controlled experimentation. In attempting to address previously unresolved questions, they also raise interesting new questions that to date remain relatively unexplored. How should agent adaptation, learning, and evolution be constructively represented in these artificial economic worlds? How should computational experiments be designed to test well-posed hypotheses so that findings are informative, whether in support or in contradiction? How can findings from computational experiments be reported in a way that combines efficiency with informativeness? And how can findings from computational experiments be validated by appropriate comparisons with findings obtained by other means?

If these questions can be successfully tackled, then agent-based computational modelling should one day join analytical modelling, econometrics, field study, and human-subject laboratory study as a standard tool in every economist's toolkit.

## References

Alchian, A.A., 1950. Uncertainty, evolution, and economic theory. Journal of Political Economy 58, 211–222.

Arthur, W.B., 1994. Inductive reasoning and bounded rationality: the El Farol problem. American Economic Association Papers and Proceedings 84, 406–411.

Arthur, W.B., Durlauf, S.N., Lane, D.A. (Eds.), 1997. The Economy as an Evolving Complex System II, Proceedings Volume XXVII, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, Reading, MA.

Axelrod, R., 1984. The Evolution of Cooperation. Basic Books, New York, NY.

Epstein, J.M., Axtell, R., 1996. Growing Artificial Societies: Social Science from the Bottom Up. MIT Press, Cambridge, MA.

Hayek, F.A., 1948. Individualism and Economic Order. University of Chicago Press, Chicago.

Holland, J., 1992. Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence, 2nd Edition. MIT Press/Bradford Books, Cambridge, MA.

Kiyotaki, N., Wright, R., 1989. On money as a medium of exchange. Journal of Political Economy 97, 924–954.

Kreps, D.M., 1990. Game Theory and Economic Modelling. Clarendon Press, Oxford, UK.

Krugman, P., 1996. The Self-Organizing Economy. Blackwell Publishers, Oxford, UK.

LeBaron, B., Arthur, W.B., Palmer, R., 1998. Time series properties of an artificial stock market. Journal of Economic Dynamics and Control 23, 1487–1516.

Nelson, R., 1995. Recent evolutionary theorizing about economic change. Journal of Economic Literature 33, 48–90.

North, D.C., 1990. Institutions, Institutional Change, and Economic Performance. Cambridge University Press, Cambridge, UK.

Samuelson, L., 1997. Evolutionary Games and Equilibrium Selection. MIT Press, Cambridge, MA.

Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Clarendon Press, Oxford, UK.

Schumpeter, J.A., 1942. Capitalism, Socialism, and Democracy. Harper & Row, New York, NY.

Schelling, T.C., 1978. Micromotives and Macrobehavior. W.W. Norton & Company, New York, NY.

Smith, A., 1937. An Inquiry into the Nature and Causes of the Wealth of Nations. Cannan Edition, American Modern Library Series, NY.

Williamson, O.E., Masten, S.E. (Eds.), 1999. The Economics of Transaction Costs. Edward Elgar Publishing, Inc., MA.

Young, H.P., 1998. Individual Strategy and Social Structure: An Evolutionary Theory of Institutions. Princeton University Press, Princeton, NJ.