Introduction to the CE Special Issue on Agent-Based Computational Economics

Leigh Tesfatsion
Department of Economics, Iowa State University, Ames, IA 50011-1070

Revised July 8, 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 the *Journal of Economic Dynamics and Control*.

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, 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 socioeconomic 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 socioeconomic 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.[1]

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 the *Journal of Economic Dynamics and Control*, 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 eighteen 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 twelve 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 the North Cameroon. The special issue at hand includes the six articles in the first category. The remaining twelve articles are included in the special ACE issue scheduled to appear in the *Journal of Economic Dynamics and Control*.

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

In the lead-off article, Joshua Epstein uses an agent-based computational model to experimentally study an important observed aspect of social norm evolution: namely, that the amount of time an individual devotes to thinking about a behavior tends to be inversely related to the strength of the social norms that relate to this behavior. In the limit, once a norm is firmly entrenched in a society, individuals tend to conform their behavior to the norm without explicit thought. Epstein's model permits agents to learn how to behave (what norm to adopt), but it also permits agents to learn how much to think about how to behave.

Epstein's model consists of a ring of interacting agents characterized by two attributes: a binary-valued norm (either 0 or 1); and a sampling radius which agents use to determine current social opinion. Each agent updates his sampling radius according to a simple rule of thumb and updates his norm to keep it in conformity with the norm currently adopted by the majority of the agents within his sampling radius. Epstein reports findings from computational experiments conducted with this model that are generally supportive of various stylized facts regarding the evolution of social norms: namely, local conformity, global diversity, and punctuated equilibrium. Moreover, he is also able to show that the amount of individual computation (sampling) rises and falls as social norms dissolve and become locked in, respectively.

Scott Page focuses on self-organization. He constructs a minimalist model of temporal task selection to study whether a collection of agents can self-organize in time and space without a central planner. Each agent in the collection must select a route among a finite set of locations, one location per period. One possible property of the resulting route collection assumed to be socially desirable is "organization," meaning that the number of agents in attendance at each location in each period is minimal. For example, given eight agents and four locations, a route collection exhibits organization if only two agents are at each location per period. Attainment of organization requires that individual routes display some dissimilarity. A second possible property assumed to be socially desirable is "robust organization" with regard to structural shocks (location closings), in the sense that a route collection remains organized after such a shock.

Page first establishes a number of analytical results regarding organization, diversity, and robust organization. He then experimentally demonstrates that the ability of agents to evolve an organized route collection starting from an unorganized route collection, using different rules of thumb for the choice of routes, becomes easier as the number of agents increases and harder as the number of locations increases, all else equal. Another key finding is that the route collections that agents evolve tend to be highly diverse, with no two agents selecting the same route. This tremendous diversity might appear to be costly, but in fact it has an unintended benefit — robustness. The diverse collections of routes evolved by the self-organizing agents proved to be more robust to shocks than the minimally diverse routes a central planner might choose to achieve organization, despite the fact that the agents did not explicitly consider robustness when choosing their routes.

A small-world network is a connected network in which each node is linked to a relatively small number of other nodes but the presence of short-cut connections between some nodes makes the average minimum path length between nodes small as well (Watts and Strogatz, 1998). Such networks have local connectivity and global reach.

Allen Wilhite uses an agent-based computational model of a bilateral exchange economy to explore the micro-level and aggregate-level consequences of restricting trade to small-world trade networks. His experiments focus on the trade-off between global market efficiency and transactions costs under four alternative types of trade networks: (a) completely connected trade networks (every agent can trade with every other trader); (b) locally disconnected trade networks consisting of disjoint trade groups; (c) locally connected trade networks consisting of trade groups aligned around a ring with a one-trader overlap at each meeting point; and (d) small-world trade networks constructed from the locally connected trade networks by permitting from one to five randomly specified short-cut trade links between members of non-neighboring trade groups. Given each type of trade network, traders endowed with stocks of two goods seek out feasible partners, negotiate prices, and then trade with those who offer the best deals.

A key finding of Wilhite's experiments is that small-world trade networks provide most of the market-efficiency advantages of the completely connected trade networks while retaining almost all of the transaction cost economies of the locally connected trade networks. His findings also suggest that there exist micro-level incentives for the formation of small-world trade networks, since the traders who use this type of network tend to do well relative to the traders who do not.

In Wilhite's study, the trade network in each experiment is taken as given. A natural extension of this work is to consider how networks among trade partners initially form and subsequently evolve.

In recent years a number of researchers have focused on this issue. Some have studied random matching of partners, some have studied probabilistic matching of partners with optional exit, and still others have studied the effects of permitting agents to preferentially choose and refuse potential partners on the basis of continually updated expected utility. In the latter work, agents have generally been assumed to have a fixed tolerance level: if the expected utility assigned to a potential partner ever falls below this level, no further offers are directed to this agent and all further offers from this agent are refused. Thus, once an agent becomes intolerable to another agent, all interactions with that other agent permanently cease; there are no second chances.

Esther Hauk introduces an important new element to the work on choice and refusal of partners: namely, endogenous pickiness. She studies an iterated prisoner's dilemma (IPD) game with choice and refusal of partners in which each agent adaptively adjusts his tolerance level to equal the average payoff he has attained from past interactions. After a certain number of periods of play, each agent can either keep his current IPD strategy or switch to a more successful strategy, after which he re-initializes all of his expected utilities for potential partners back to optimistic levels. The combined effect of having both endogenous tolerance levels and reconstituted optimism is that cooperative agents can punish defectors by refusal in a manner that induces the defectors to change towards more cooperative behavior in the future. This occurs because potential partners who are refused are eventually given a second chance to clean up their act.

Hauk first derives a number of analytical results for the special case in which agents are restricted to just two strategies: always defect (AllD); and always cooperative (AllC). She then uses computational experiments to investigate what happens when agents adaptively select their IPD strategies from among a set of five strategies observed in human-subject experiments that range from the purely nasty AllD to the purely cooperative AllC. She shows that the effect of having endogenous tolerance levels and reconstituted optimism in such a context is to ensure that, over time, cooperators exclusively interact with each other and are more successful than exploiters.

In Hauk's study, all agents use the same fixed learning algorithm to implement their strategy revision. In contrast, Ann Bell focuses on the comparative performance of alternative learning algorithms.

Specifically, Bell compares the performance of three different reinforcement learning (RL) algorithms in a repeated multi-agent congestion game similar in structure to the well known *El Farol* problem studied by (Arthur, 1994). The rules of the game are stationary, but the choice environment continuously evolves as agents learn to associate state and action pairs on the basis of past rewards. The performance of a standard RL algorithm from the computer science literature is contrasted with that of two simpler RL algorithms used by Alvin Roth and Ido Erev to explain experimental data from human subject experiments.

Bell first examines a single adaptive learning agent facing a stationary stochastic environment in which all other agents choose their actions randomly. She shows that the performance of the Roth and Erev RL algorithms in such a setting tends to be extremely sensitive to apparently innocuous changes in experimental specifications, such as the scaling of rewards, whereas the RL algorithm from computer science results in rapid and robust convergence to stable aggregate system behavior. She then increases the proportion of adaptive agents, which results in a nonstationary endogenously evolving choice environment. A key finding in the latter environment is that her former findings are reversed: the Roth and Erev RL algorithms tend to induce rapid convergence to stable aggregate system behavior despite the slow and erratic behavior of individual learners, while the computer science RL algorithm does not. She concludes that much more study is needed to understand the dynamics and performance of learning algorithms under different choice conditions.

In a similar vein, Bruce Edmonds criticizes the unexamined use of off-the-shelf computer algorithms to represent learning processes in agent-based modelling. Rather, he argues that modellers should strive for descriptive accuracy, making use of specific findings in human subject experiments as well as current developments in cognitive science in general.

To illustrate his arguments in more concrete terms, Edmonds studies the behavior exhibited by human subjects participating in a previously conducted one-sided auction experiment. Each human subject had one unit of good to sell, received a sequence of bids for this unit at some cost per bid, and had to decide when to stop the bidding process and accept the last bid made. At the end of the experiment the types of stopping rule strategies used by the human subjects were extracted in a form amenable to computational representation. Edmonds attempts to reproduce the human subjects' observed behavior in an analogous computational auction experiment with artificial agents by endowing the artificial agents with these observed types of stopping rule strategies together with a learning mechanism based on the genetic programming (GP) paradigm developed by (Koza, 1992).

Edmonds' artificial agent experiments closely approximate several important aspects of the human subject experiments: for example, the mean and spread of attained earnings. Nevertheless, some aspects are not captured well. For example, the GP strategies evolved by the artificial agents frequently involved complicated expressions that humans would automatically simplify, and certain conditioning events (predicates) that appeared frequently in the strategies used by the human subjects appeared only infrequently in the evolved GP strategies. He concludes that the major problem currently facing researchers interested in incorporating descriptively accurate cognitive processes in artificial agents is a lack of empirical data regarding how humans learn strategies. Examining the artificial agent run generating the closest match to the human subject experiments, he is able to suggest a number of hypotheses regarding this learning as a guide for further study.

In conclusion, all of the articles included in this special ACE issue attempt to learn about some poorly understood aspect of human behavior by reproducing and studying this behavior in an agent-based computational framework. Since the chains of causality in such frameworks tend to be extremely complicated, the specific global regularities that arise tend not to be directly predictable consequences of the initial agent and environmental specifications. This loss of analytical tractability is understandably viewed by some as a drawback of ACE research. Nevertheless, watching ACE worlds evolve over time does provide the continual pleasure of discovery; and, as stressed by (Feynman, 1999), this is exactly how scientific research ought to be.

Notes

[1] 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.

References

Armen A. Alchian. Uncertainty, Evolution, and Economic Theory. *Journal of Political Economy* 58: 211–222, 1950.

W. Brian Arthur. Inductive Reasoning and Bounded Rationality: The *El Farol* Problem. *American Economic Association Papers and Proceedings* 84: 406–411, 1994.

Robert Axelrod. *The Evolution of Cooperation*. Basic Books, New York, N.Y., 1984.

Ann Maria Bell. Reinforcement Learning Rules in a Repeated Game. *Computational Economics*, this issue.

Bruce Edmonds. Towards a Descriptive Model of Agent Strategy Search. *Computational Economics*, this issue.

Joshua M. Epstein. Learning to Be Thoughtless: Social Norms and Individual Computation. *Computational Economics*, this issue.

Joshua M. Epstein and Robert Axtell. *Growing Artificial Societies: Social Science from the Bottom Up*. The MIT Press, Cambridge, MA, 1996.

Richard Feynman. *The Pleasure of Finding Things Out*. Perseus Books, Cambridge, 1999.

Esther Hauk. Leaving the Prison: Permitting Partner Choice and Refusal in Prisoner's Dilemma Games. *Computational Economics*, this issue.

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

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

John Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. The MIT Press, Cambridge, MA, 1992.

Richard Nelson. Recent Evolutionary Theorizing About Economic Change. *Journal of Economic Literature* 33: 48–90, 1995.

Scott E. Page. Self-Organization and Coordination. *Computational Economics*, this issue.

Lawrence Samuelson. *Evolutionary Games and Equilibrium Selection* The MIT Press, Cambridge, MA, 1997.

Thomas C. Schelling. *Micromotives and Macrobehavior* W. W. Norton & Company, New York, N.Y., 1978.

Joseph A. Schumpeter. *Capitalism, Socialism, and Democracy* Harper & Row Publishers, Inc., New York, N.Y., 1942.

Adam Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Cannan Edition, American Modern Library Series, N.Y., 1937.

Duncan Watts and Steven H. Strogatz. Collective Dynamics of ‘Small-World’ Networks. *Nature* 393: 440–442, June 1998.

Allen Wilhite. Bilateral Trade and ‘Small-World’ Networks. *Computational Economics*, this issue.