16 January 2000

Agent-Based Computational Economics: A Brief Guide to the Literature

Leigh Tesfatsion
Department of Economics
Iowa State University
Ames, Iowa 50011-1070
tesfatsi@iastate.edu
http://www.econ.iastate.edu/tesfatsi/

1 Introduction

The newly developing field of agent-based computational economics (ACE) is roughly defined by its practitioners as the computational study of economies modelled as evolving systems of autonomous interacting agents.[1] A principal concern of ACE researchers is to understand the apparently spontaneous formation of global regularities in economic processes, such as the unplanned coordination of trading activities in decentralized market economies that economists associate with Adam Smith's invisible hand. The challenge is to explain how these global regularities arise from the bottom up, through the repeated local interactions of autonomous agents channeled through socioeconomic institutions, rather than from the top-down imposition of fictitious coordination mechanisms such as market clearing constraints or an assumption of single representative agents.

The study of evolutionary economics is by no means new, of course. Even before Darwin, attempts were made to apply evolutionary ideas to socioeconomic behavior.

Although this early work is now largely ignored by economists, economic textbooks still typically include at least some mention of the ideas of J. Schumpeter regarding the evolution of economic institutions. Moreover, Schumpeter's work, together with the seminal work by A. A. Alchian on uncertainty and evolution in economic systems, appears to have strongly influenced the subsequent well-known work by R. Nelson, S. Winter and various of their collaborators on evolutionary theories of economic change. In addition, one has the work of W. B. Arthur on economies incorporating positive feedbacks, the work by Richard Day on dynamic economies characterized by complex phase transitions, the work by J. Foster on an evolutionary approach to macroeconomics, R. Heiner's work

[1] For extensive on-line ACE resources, including surveys, an annotated syllabus of readings, software, and pointers to individual researchers and research groups, see the ACE Web site at http://www.econ.iastate.edu/tesfatsi/ace.htm.

on the origins of predictable behavior, J. Hirshleifer's work on evolutionary models in economics and law, and U. Witt's work on economic natural selection. These and numerous other interesting studies on evolutionary economics are reviewed by Nelson (1995).

More recently, as detailed in Samuelson (1997), a number of researchers have been focusing on the potential economic applicability of evolutionary game theory with replicator dynamics. In these studies, game strategies are distributed over a fixed number of strategy types, and the strategies reproduce over time in direct proportion to their relative fitness.

Exploiting the recent advent of more powerful computational tools, such as object-oriented programming, ACE researchers such as Arifovic (1994), Arthur et al. (1997), Axelrod (1997), Epstein & Axtell (1996), Kirman (1997), Kollman et al. (1997), Marks (1992), McFadzean and Tesfatsion (1999), Miller (1996), Tesfatsion (1997a,b), Vriend (1995), and Young (1998) have been able to extend this earlier work on evolutionary economics in four key ways.

First, agents in ACE frameworks are typically modelled as heterogeneous entities that determine their interactions with other agents and with their environment on the basis of internalized data and behavioral rules. These agents thus tend to have a great deal more internal cognitive structure and autonomy than conventionally modelled economic agents. Second, a broader range of agent interactions is typically permitted in ACE frameworks, with predatory and cooperative associations taking center stage along with price and quantity relationships. Third, the evolutionary process is generally represented as natural selection pressures acting directly on agent characteristics rather than as population-level laws of motion. These natural selection pressures result in the continual creation of new modes of agent behavior and an ever-changing network of agent interactions. Fourth, ACE frameworks are computer implemented as virtual economic worlds that grow themselves along a real time-line, much like a culture grows in a petri dish. In principle, once initial conditions are set, all subsequent events in these virtual economic worlds are initiated and driven by agent-agent and agent-environment interactions; no further outside interventions by the modeler (e.g., off-line fixed point calculations) are permitted.

In brief, then, ACE is a blend of concepts and tools from evolutionary economics, cognitive science, and computer science. It represents a methodological approach that may ultimately permit two important developments: (a) the rigorous testing, refinement, and extension of theories developed in the earlier literature on evolutionary economics that were found to be analytically intractable; and (b) the rigorous formulation and testing of conceptually integrated socioeconomic

```
int main () {
  Init();                          // Construct the initial trader generation
                                   //    with random trade strategies.
  For (G = 1,...,GMax){            // Enter the generation cycle loop.
                                   // Generation Cycle:
    InitGen();                     //    Configure traders with user-supplied
                                   //       parameter values (initial expected
                                   //       utility levels, capacity quotas,...).
    For (I = 1,...,IMax) {         //    Enter the trade cycle loop.
                                   //    Trade Cycle:
      MatchTraders();              //       Determine trade partners,
                                   //          given expected utilities,
                                   //          and record job search and
                                   //          inactivity costs.
      Trade();                     //       Implement trades and
                                   //          record trade payoffs.
      UpdateExp();                 //       Update expected utilities using
    }                              //          newly recorded costs and payoffs.
                                   //    Environment Step:
    AssessFitness();               //       Assess trader fitness scores.
    Output();                      //       Output trader information.
                                   //    Evolution Step:
    EvolveGen();                   //       Evolve a new trader generation.
  }
  Return 0;
}
```

Table I: Pseudo-Code for the Trade Network Game (TNG)

theories compatible with theory and data from many different relevant fields currently separated by artificial disciplinary boundaries.

2 Illustration: An ACE Trade Network Framework

To illustrate the ACE approach in more concrete terms, consider the Trade Network Game (TNG) developed by Tesfatsion (1997b) and McFadzean and Tesfatsion (1999) for studying the formation and evolution of trade networks. As depicted in Table I, the TNG consists of successive generations of traders who choose and refuse trade partners on the basis of continually updated expected payoffs, engage in risky trades modelled as two person games, and evolve their trade strategies over time.

The TNG framework facilitates the computational study of markets from an agent-based perspective in three key ways.

First, as seen in Table I, the TNG framework is modular in design. This means that experimentation with alternative specifications for market structure, search and matching among traders, trade site interactions, expectation formation and updating, and evolution of trade site strategies

```
class TradeBot
{
    Internalized Social Norms:
        Market protocols for communicating with other traders;
        Market protocols for trade partner search and matching;
        Market protocols for trade interactions.
    Internally Stored State Information:
        My attributes;
        My endowments;
        My beliefs and preferences;
        Addresses I have for myself and for other traders;
        Additional data I have about other traders.
    Internal Behavioral Rules:
        My rules for gathering and processing new information;
        My rules for determining my trade behavior;
        My rules for updating my beliefs and preferences;
        My rules for measuring my utility (fitness) level;
        My rules for modifying my rules.
};
```

Table II: A TNG Trader as a Software Agent

can easily be undertaken — much like changing a lightbulb in a multi-bulb lamp — as long as the interfaces (inputs and outputs) for the modules implementing these specifications remain unchanged. Moreover, each of these modules can potentially be grounded in trader-initiated actions in the sense that the module is implemented via behavioral rules internal to the traders.

Second, as seen in Table II, each TNG trader is an autonomous software agent ("tradebot") with internalized social norms, internally stored state information, and internal behavioral rules. The traders can therefore engage in anticipatory behavior. Moreover, using stored addresses together with internalized communication protocols, the traders can communicate with each other at event-triggered times, a feature not present in standard economic models.

Third, the TNG permits the rigorous and routine experimental study of non-steady state market dynamics at three different levels: individual characteristics of traders; interactions among traders (network formation); and social welfare as measured by descriptive statistics such as average trader welfare. For example, the TNG framework is currently being applied to the study of worker-employer market power relationships in evolutionary labor markets with adaptive search; see Tesfatsion (1999).

3 Conclusion

In summary, the hallmark of the ACE approach to the study of economic processes is a bottom up perspective, in the sense that global behavior is grounded in local agent interactions. The TNG framework briefly outlined in Section 2 illustrates how the ACE approach might be specialized to the study of evolutionary trade networks. More generally, however, the ACE approach permits economists to begin the difficult task of studying the self-organizing capabilities of decentralized market economies.

References

Arifovic, Jasmina (1994), "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control* 18, 3–28.

Arthur, W. Brian, John H. Holland, Blake LeBaron, Richard Palmer, and Paul Taylor (1997), "Asset Pricing Under Endogenous Expectations in an Artificial Stock Market", in *The Economy as an Evolving Complex System, II*, edited by W. Brian Arthur, Steven Durlauf, and David Lane, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XXVII, Reading, MA: Addison-Wesley.

Axelrod, Robert (1997), *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*, Princeton: Princeton University Press.

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

Kirman, Alan (1997), "The Economy as an Interactive System", in *The Economy as an Evolving Complex System, II*, *op. cit.*.

Kollman, Ken, John H. Miller, and Scott E. Page (1997), "Computational Political Economy", in *The Economy as an Evolving Complex System, II*, *op. cit.*.

Marks, Robert (1992), "Breeding Hybrid Strategies: Optimal Behavior for Oligopolists", *Journal of Evolutionary Economics* 2, 17–38.

McFadzean, David, and Leigh Tesfatsion (1999), "A C++ Platform for the Evolution of Trade Networks," *Computational Economics* 14, 109–134.

Miller, John (1996), "The Coevolution of Automata in the Repeated Prisoner's Dilemma", *Journal of Economic Behavior and Organization* 29 (1996), 87–112.

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

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

Tesfatsion, Leigh (1997a), "How Economists Can Get Alife", in *The Economy as an Evolving Complex System, II*, *op. cit.*.

Tesfatsion, Leigh (1997b), "A Trade Network Game with Endogenous Partner Selection", pp. 249–269 in *Computational Approaches to Economic Problems*, edited by Hans Amman, Berc Rustem, and Andrew Whinston, Dordrecht, the Netherlands: Kluwer Academic Publishers.

Tesfatsion, Leigh (1999), "Structure, Behavior, and Market Power an in Evolutionary Labor Market with Adaptive Search", Economic Report 51, Iowa State University, to appear in the *Journal of Economic Dynamics and Control*. [A preprint of this paper is available online at http://www.econ.iastate.edu/tesfatsi/mpevlab.ps]

Vriend, Nicolaas (1995), "Self-Organization of Markets: An Example of a Computational Approach", *Computational Economics* 8, 205–231.

Young, Peyton, *Individual and Social Structure: An Evolutionary Theory of Institutions*, Princeton: Princeton University Press, Princeton, 1998.