

2003

Agent-based computational economics: modeling economies as complex adaptive systems

Leigh Tesfatsion

Iowa State University, tesfatsi@iastate.edu

Follow this and additional works at: http://lib.dr.iastate.edu/econ_las_pubs

 Part of the [Growth and Development Commons](#), [Industrial Organization Commons](#), and the [Other Economics Commons](#)

The complete bibliographic information for this item can be found at http://lib.dr.iastate.edu/econ_las_pubs/424. For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

This Article is brought to you for free and open access by the Economics at Iowa State University Digital Repository. It has been accepted for inclusion in Economics Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Agent-Based Computational Economics: Modelling Economies as Complex Adaptive Systems^{*}

Leigh Tesfatsion

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

Abstract

Agent-based computational economics (ACE) is the computational study of economies modelled as evolving systems of autonomous interacting agents. Thus, ACE is a specialization to economics of the basic complex adaptive systems paradigm. This paper outlines the main objectives and defining characteristics of the ACE methodology, and discusses several active research areas.

Key words: Agent-based computational economics, complex adaptive systems

1 Introduction

Decentralized market economies are complex adaptive systems, consisting of large numbers of adaptive agents involved in parallel local interactions. These local interactions give rise to macroeconomic regularities such as shared market protocols and behavioral norms which in turn feed back into the determination of local interactions. The result is a complicated dynamic system of recurrent causal chains connecting individual behaviors, interaction networks, and social welfare outcomes.

This intricate two-way feedback between microstructure and macrostructure has been recognized within economics for a very long time. Nevertheless, for much of this time economists have lacked the means to model this feedback quantitatively in its full dynamic complexity. The most salient characteristic of

^{*} To appear in *Information Sciences*. This review is an abbreviated version of [20].

traditional quantitative economic models supported by microfoundations has been their top-down construction. Heavy reliance is placed on extraneous coordination devices such as fixed decision rules, common knowledge assumptions, representative agents, and imposed market equilibrium constraints. Face-to-face personal interactions typically play no role or appear in the form of tightly constrained and stylized game interactions. In short, agents in these models have had little room to breathe.

Slowly but surely, however, advances in modeling tools have been enlarging the possibility set for economists (e.g., [1,8]). Researchers can now quantitatively model a wide variety of complex phenomena associated with decentralized market economies, such as inductive learning, imperfect competition, trade network formation, and the open-ended co-evolution of individual behaviors and economic institutions.

One branch of this work has come to be known as *agent-based computational economics (ACE)*, the computational study of economies modeled as evolving systems of autonomous interacting agents.¹ ACE researchers generally rely on computational laboratories to study the evolution of decentralized market economies under controlled experimental conditions.

As in a culture-dish laboratory experiment, the ACE modeller starts by constructing an economy with an initial population of agents. These agents include both economic agents and agents representing various other social and environmental phenomena. The ACE modeler specifies the initial state of the economy by specifying the initial attributes of the agents. The initial attributes of an agent might include type characteristics, internalized behavioral norms, internal modes of behavior (including modes of communication and learning), and internally stored information about itself and other agents. The economy then evolves over time without further intervention from the modeller. All events that subsequently occur must arise from the historical time-line of agent-agent interactions. No extraneous coordination devices are permitted. For example, no resort can be made to the off-line determination and imposition of market-clearing prices through fixed point calculations.

2 ACE Research Areas

Three special ACE journal issues have recently appeared that include a diverse sampling of current ACE research [17–19]. The topics addressed in these issues roughly divide into eight research areas: (i) Learning and the embodied

¹ See <http://www.econ.iastate.edu/tesfatsi/ace.htm> for extensive resources related to the ACE methodology.

mind; (ii) evolution of behavioral norms; (iii) bottom-up modeling of market processes; (iv) formation of economic networks; (v) modelling of organizations; (vi) design of computational agents for automated markets; (vii) parallel experiments with real and computational agents; and (viii) building ACE computational laboratories.²

The major issue driving research area (i) is how to model the minds of the computational agents who populate ACE frameworks. Should these minds be viewed as logic machines with appended data filing cabinets, the traditional artificial intelligence viewpoint? Or should they instead be viewed as controllers for embodied activity, as advocated by evolutionary psychologists? If the focus of an ACE study is the design of a fully automated market, there is no particular reason why the minds of the computational agents should have to mimic those of real people – indeed, this could be positively detrimental to good market performance. On the other hand, if the focus is on the modeling of an economic process with human participants, then mimicry might be essential to ensure predictive content.

Aware of these concerns, ACE researchers are increasingly moving away from the unconsidered adoption of off-the-shelf learning algorithms. Some ACE researchers are systematically investigating the performance of alternative learning algorithms in various economic decision contexts (e.g., [3,21]). Others are attempting to calibrate their learning algorithms to empirical decision-making data (e.g., [13,14]). Another interesting development is the use of human-subject experimental data to calibrate the learning of computational agents (e.g., [6]).

An interesting related learning issue being addressed by ACE researchers is the extent to which the learning processes of real-world market participants are mal-adapted to market institutions, leaving room for improvement from the application of optimization tools [11]. Conversely, to what extent have existing market protocols evolved or been designed to avoid the need for any great rationality on the part of market participants [10]?

An important issue for the evolution of behavioral norms is how mutual cooperation manages to evolve among economic agents even when cheating reaps immediate gains and binding commitments are not possible [2]. What roles do reputation, trust, reciprocity, retaliation, spitefulness, and punishment play? More generally, how do exchange customs and other behavioral norms important for economic processes come to be established, and how stable are these norms over time? Are these behavioral norms diffusing across traditional political and cultural boundaries, resulting in an increasingly homogeneous global economy?

² See <http://www.econ.iastate.edu/tesfatsi/aapplic.htm> for pointers to resources for each research area.

As detailed in [9], the evolution of behavioral norms has been studied using classical game theory. In the latter, the approach has been to explain this evolution on the basis of individual rationality considerations, such as anticipations of future reciprocity. In contrast, many ACE researchers (e.g., [7,9]) have tended to place equal or greater stress on peer emulation, parental mimicry, and other socialization forces thought to underly the transmission of culture.

A key issue driving ACE market studies is the evolution or design of market protocols and other market-related institutions. Many of the articles included in the special ACE issues [18] and [19] lie in this research area. Types of markets studied include financial, electricity, labor, agricultural, entertainment, retail, business-to-business, and automated Internet auctions.

A challenging issue motivating ACE research in the area of economic network formation is the manner in which economic interaction networks are determined through deliberative choice of partners as well as by chance. A key concern has been the emergence of trade networks among collections of buyers and sellers who determine their trade partners adaptively, on the basis of past experiences with these partners. Recent research on the endogenous formation of trade networks has tended to focus more concretely on specific types of markets (e.g., [12,16]).

Another important issue driving this research area is the extent to which interaction networks are important for predicting market outcomes. If interaction effects are weak, as in some types of auction markets [10], then the structural aspects of the market (e.g., numbers of buyers and sellers, costs, capacities) will be the primary determinants of market outcomes. In this case, given a particular market structure treatment, multiple trial runs should result in a relatively simple *central-tendency* outcome distribution. If interaction effects are strong, as in labor markets [16], then each market structure treatment might instead map under multiple trial runs into a *spectral* outcome distribution with outcomes clustered around two or more distinct “attractors” corresponding to distinct interaction networks.

The main questions traditionally driving the study of organizations have been normative: How to determine the optimal form of organization for achieving specified goals? More generally, what is the relationship between environmental properties, organization structure, and organization performance? As illustrated by [4], the increased use of ACE modelling in this research area might ultimately permit a significant widening of this traditional scope by permitting the quantitative study of multiple organizations within broader economic settings, e.g., the study of intra-firm organization for multiple firms participating within a market.

To date, much of the work on computational agent design for automated markets has been driven by the quest for optimal performance. Nevertheless, some studies are taking a broader tact, focusing instead on the increasing ubiquity of artificial life forms that this trend toward automation entails and on some of the concerns this trend raises (e.g., [11]).

The attempts by ACE researchers to conduct parallel experiments with real and computational agents has raised a number of challenging issues. Chief among these is the need to make the parallel experiments truly parallel, so that comparisons are meaningful and lead to robust insights.

One major hurdle is the need to ensure that the salient aspects of an experimental design as perceived by the human participants are captured in the initial conditions specified for the computational agents. Another major hurdle is that experiments run with human participants generally have to be kept short, both to prevent boredom among the participants and to prevent the bankruptcy of the investigators who provide the participants with monetary payments. Thus, the “shadow of the past” might be strongly affecting experimental outcomes for individual human participants in ways not understood and controlled for by investigators. In contrast, experiments with computational agents can be run for many generations to diminish dependence on initial conditions.

An important question, then, is which type of horizon – short run or long run – provides the best approximation for real-world economic processes. Do real-world economic agents essentially move from one new economic situation to the next, never having a chance to settle into long-run behavior? Or do these agents participate repeatedly in economic situations with enough similarity that they are able to use long-run learned (or inherited) behaviors to deal effectively with these situations?

A *computational laboratory (CL)* is a framework that permits the study of systems of multiple interacting agents by means of controlled and replicable computational experiments [5,15].³ A number of issues arise for ACE researchers wishing to build computational laboratories. For example, should a separate CL be constructed for each application, or should researchers strive for general multi-faceted platforms? How can experimental findings be effectively communicated to other researchers by means of descriptive statistics and graphical visualizations without information overload? How might these findings be validated by comparisons with data obtained from other sources? A particularly important unresolved issue is the need to ensure that findings from ACE experiments reflect fundamental aspects of a considered problem

³ See <http://www.econ.iastate.edu/tesfatsi/acecode.htm> for pointers to a wide variety of CLs, authoring tools, and general programming languages that are currently being used to design and/or test multi-agent systems.

and not simply the peculiarities of the particular CL or programming language used to implement the experiments.

3 Concluding Remarks

The defining characteristic of ACE model economies is their constructive grounding in the interactions of autonomous adaptive agents, broadly defined to include economic, social, and environmental entities. The state of an ACE model economy at each point in time is given by the internal attributes of the individual agents that currently populate the economy. These kinds of state descriptions will presumably have direct meaning for economists and other social scientists, thus increasing the transparency and clarity of the modelling process.

The use of ACE model economies might also facilitate the development and experimental testing of integrated theories that build on theory and data from many different fields of social science. In particular, ACE frameworks could encourage economists to address growth, distribution, and welfare issues in a more comprehensive manner embracing a variety of economic, social, political, and psychological factors, thus restoring the broad vision of early political economists.

References

- [1] W. B. Arthur, S. N. Durlauf, D. A. Lane (Eds.), *The economy as an evolving complex system II*, Proceedings Vol. XXVII, Santa Fe Institute Studies in the Sciences of Complexity, Addison-Wesley, Reading, MA, 1997.
- [2] R. Axelrod, *The complexity of cooperation: Agent-based models of conflict and cooperation*, The Princeton University Press, Princeton, NJ, 1997.
- [3] H. Dawid, *Adaptive learning by genetic algorithms: Analytical results and applications to economic models*, Revised Second Edition, Springer-Verlag, Berlin, 1999.
- [4] H. Dawid, M. Reimann, B. Bullnheimer, *To innovate or not to innovate?*, IEEE Transactions on Evolutionary Computation 5 (2001) 471–481.
- [5] C. Dibble, *Theory in a complex world: GeoGraph computational laboratories*, Ph.D. Thesis, Geography Department, University of California at Santa Barbara, CA, 2001.
- [6] J. Duffy, *Learning to speculate: Experiments with artificial and real agents*, Journal of Economic Dynamics and Control 25 (2001) 295–319.

- [7] J. M. Epstein, Learning to be thoughtless: Social norms and individual computation, *Computational Economics* 18 (2001) 9–24.
- [8] J. M. Epstein, R. Axtell, *Growing artificial societies: Social science from the bottom up*, The MIT Press, Cambridge, MA, 1996.
- [9] H. Gintis, *Game theory evolving*, The Princeton University Press, Princeton, NJ, 2000.
- [10] D. K. Gode, S. Sunder, Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality, *Journal of Political Economy* 101 (1993) 119–137.
- [11] J. Kephart, Software agents and the information economy, *Proceedings of the National Academy of Sciences U.S.A.*, Volume 99, Supplement 3, 7207-7213, May 14, 2002.
- [12] A. Kirman, N. J. Vriend, Evolving market structure: An ACE model of price dispersion and loyalty, *J. Economic Dynamics and Control* 25 (2001) 459–502.
- [13] B. LeBaron, Empirical regularities from interacting long and short horizon investors in an agent-based stock market model, *IEEE Transactions on Evolutionary Computation* 5 (2001) 442–455.
- [14] R. E. Marks, Breeding hybrid strategies: Optimal behaviour for oligopolists, *Journal of Evolutionary Economics* 2 (1992) 17–38.
- [15] D. McFadzean, D. Stewart, L. Tesfatsion, A computational laboratory for evolutionary trade networks, *IEEE Transactions on Evolutionary Computation* 5 (2001) 546–560. Preprint and source code are available at <http://www.econ.iastate.edu/tesfatsi/tnghome.htm>
- [16] L. Tesfatsion, Structure, behavior, and market power in an evolutionary labor market with adaptive search, *Journal of Economic Dynamics and Control* 25 (2001) 419–457. Preprint at <http://www.econ.iastate.edu/tesfatsi/mpevlab.pdf>
- [17] L. Tesfatsion (Guest Ed.), Special Issue on Agent-Based Computational Economics, *Computational Economics*, Vol. 18, No. 1, October 2001, 1–135.
- [18] L. Tesfatsion (Guest Ed.), Special Issue on Agent-Based Computational Economics, *Journal of Economic Dynamics and Control*, Vol. 25, No. 3-4, March 2001, 281–654.
- [19] L. Tesfatsion (Guest Ed.), Special Issue on the Agent-Based Modeling of Evolutionary Economic Systems, *IEEE Transactions on Evolutionary Computation*, Vol. 5, No. 5, October 2001, 437–560.
- [20] L. Tesfatsion, Agent-based computational economics, ISU Economics Working Paper No. 1, Economics Department, Iowa State University, July 2002. <http://www.econ.iastate.edu/tesfatsi/acewp1.pdf>
- [21] N. J. Vriend, An illustration of the essential difference between individual and social learning, and its consequences for computational analyses, *Journal of Economic Dynamics and Control* 24 (2000) 1–19.