

Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents

Blake LeBaron and Leigh Tesfatsion*, *AER Papers & Proceedings*, 98(2), 246-250

“All models are wrong, but some are useful.” G.E.P. Box (1979, p. 202)

Macroeconomists seek to understand the structure and performance of economies at a national or regional level and the manner in which government policy makers attempt to influence this structure and performance over time. Such understanding would seem to require a systematic exploration of the intricate feedback loops connecting micro behaviors, interaction patterns, and macro regularities as observed in real-world economies.

In fact, however, mainstream macroeconomic theory remains firmly rooted in general equilibrium microfoundations (David Colander, 2006). Emphasis is on the isolated optimal choice behaviors of utility-maximizing households and profit-maximizing firms subject to budget and technological feasibility constraints, and on the equilibrium states attained through external imposition of conditions requiring fulfilled expectations and market clearing.

Potentially important real-world factors such as subsistence needs, incomplete markets, imperfect competition, inside money, strategic behavioral interactions, and open-ended learning that tremendously complicate analytical formulations are typically not incorporated.

Starting around the mid-1980s, various researchers have sought to develop agent-based computational economics tools able to capture in useful terms the complexity of real-world

*Blake LeBaron, International Business School, Brandeis University, Mailstop 32, Waltham, MA 02454-9110. Tel: 781- 731-2258. FAX: 781-736-2267. Email: blebaron@brandeis.edu; Leigh Tesfatsion (corresponding author), Economics Department, Iowa State University, Ames, IA 50011-1070. Tel: 515-294-7318. FAX: 515-294-0221. Email: tesfatsi@iastate.edu

economic phenomena. Could the application of such tools facilitate a more empirically-based approach to macroeconomic modeling?

As elaborated in Joshua M. Epstein and Robert L. Axtell (1996) and Leigh Tesfatsion and Kenneth L. Judd (2006), Agent-based Computational Economics (ACE) is the computational study of economic processes modeled as dynamic systems of interacting agents.¹ Here “agent” refers broadly to an encapsulated collection of data and methods representing an entity residing in a computationally constructed world. Individual biological life forms, social groupings, institutions, and physical entities can all be represented as agents.

ACE is a culture-dish approach to the study of economic worlds. Once initial conditions have been specified by the modeler, all subsequent world events are driven by agent interactions. These interactions – the attempts by agents to express actions within their worlds -- are determined dynamically in “run-time” by the agents’ internal structures, informational states, beliefs, motivations, and data-processing methods. A crucial point is that modelers do not need to constrain agent interactions *a priori* by the imposition of equilibrium conditions, homogeneity assumptions, or other external coordination devices that have no real-world referents. Ideally, the agents in ACE models should be as free to act within their computational worlds as their empirical counterparts are within the real world.

In order for an ACE model to facilitate the understanding of a real-world macroeconomy, however, three criteria must be met. First, the model must include an appropriate empirically-based taxonomy of agents. Second, the scale of the model must be suitable for the particular

¹See <http://www.econ.iastate.edu/tesfatsi/ace.htm> for annotated pointers to ACE tutorials, readings, software, teaching resources, research groups, individual researchers, and research resource sites, including a site devoted to ACE studies of decentralized market economies.

purpose at hand. Third, model specifications must be subject to empirical validation in an attempt to provide genuine insight into proximate and ultimate causal mechanisms. The following sections address each of these criteria in turn.

I. Taxonomy: What Types of Agents for Macroeconomic Models?

Taxonomy, the classification of phenomena into ordered groups or categories, is indispensable for scientific research. It facilitates information retrieval and establishes the foundations for comparative research (Ernst Mayr, 1997, Chapter 7).

Empirically grounded taxonomic classification should routinely be undertaken as part of the process of macroeconomic theorizing and model building. What types of human needs and desires are relevant for understanding particular types of macroeconomic phenomena? What types of goods and services meet or could meet these human needs and desires? What types of facilities exist or could exist to produce these goods and services, and who participates in these production activities? What kinds of institutions exist or could exist to distribute these goods and services, and who participates in these distribution activities? And what types of players (if any) oversee the design and/or operation of these institutions, and for what purposes?

ACE modeling per se provides no answers to these questions; it is a methodological approach, not a theory. Rather, ACE modeling provides a systematic way to incorporate whatever taxonomic classification a researcher believes is useful for the exploratory study of a particular economic phenomenon (Teshatsion, 2007). The researcher is freed from the constrictive binds of analytical tractability and from the need to rely on narrow fragmented taxonomies arising from artificial disciplinary boundaries.

More precisely, agents in ACE models can span all the way from passive features of the world with no cognitive function to decision makers with sophisticated cognitive abilities who

actively gather and process data. For example, as illustrated in Tesfatsion and Judd (2006, Chapter 16), an ACE macroeconomic model might include structural agents (e.g. a spatial world), institutional agents (e.g. a legal system, corporations, markets), and cognitive agents (e.g. entrepreneurs, consumers, stock brokers, and government policy makers).

Agents can also be composed of more elementary agents in various forms of hierarchical organization. For example, an ACE macroeconomic model might include the following hierarchy of nested agent refinements: national economy → {financial sector, business sector, household sector, government sector, foreign sector}; financial sector → {commercial banks, insurance companies, stock brokers, bond dealers}; commercial banks → {employees, shareholders}; employees → {salaried workers, wage workers}; and so forth.

ACE modeling thus provides macroeconomists with tremendous flexibility to tailor the breadth and depth of the “representative agents” in their models to particular applications at hand. The taxonomy can be adjusted to the application rather than the application to the taxonomy, surely an essential prerequisite for sound scientific research.

Once a taxonomy is specified, the data and methods of each agent type can be initialized using available evidence from field studies, econometric studies, human-subject laboratory experiments, surveys, and interviews. ACE models are typically implemented using programming languages with “object-oriented” capabilities that permit highly modular and extensible model formulations. It is therefore relatively easy to successively refine a model’s taxonomy as experience with the model is gained.

II. Scale Robustness: How Many Agents for Macroeconomic Models?

As in all sciences, macroeconomics strives to construct models that are “simple but not too simple” for whatever purpose is at hand. Scaling is a critical aspect of this simplification. How

many households with diverse needs and wants should be considered? How many goods and services to meet these needs and wants should be represented? How many institutional aspects of production and distribution should be explicitly incorporated? And so forth.

Early in their training all economists are taught the potentially powerful effects of scale on the fundamental nature of economic activity. At one end of the spectrum, when market participants are sufficiently numerous, “perfect competition” among buyers and sellers can result in effective price-taking behaviors with prices and quantities determined at efficient market-clearing levels. At the other end of the spectrum, when a market consists of a single seller facing numerous buyers (or vice versa), monopoly pricing can occur in which the single seller (or single buyer) sets the market price to extract as much own-profit as possible.

In between these two extremes lies a vast, largely unexplored region in which rivalry among buyers and sellers leads to “imperfect competition.” Within this region, economists have studied several special cases in some depth, including Bertrand or Cournot duopoly, particular types of oligopoly, and monopolistic competition. However, measured against the full range of possibilities, these special cases constitute a set of measure zero.

Modern macroeconomic theory is largely founded on assumptions of perfect competition, driven to this modeling strategy not so much by empirical evidence as by considerations of analytical tractability. Coordination issues are commonly sidestepped by one of two means. One approach is to postulate market equilibrium in terms of high-level aggregate constructs, e.g., aggregates Y , K , and L satisfying an aggregate production relation $Y=AF(K,L)$, despite the incredibly stringent conditions required for the existence of such constructs (Franklin M. Fisher, 1993). The other approach is to postulate market equilibrium in terms of single representative agents for the household and firm sectors assumed to behave as competitive price takers despite

the absence of contestable markets or large numbers of rival buyers and sellers, and despite the many logical difficulties associated with the presumption that single agents can adequately represent collectives of decision makers (Alan P. Kirman, 1992).

One of the greatest potential contributions that ACE could make to macroeconomic theory is permitting the constructive exploration of scale effects without the external imposition of artificial coordination devices. What does it matter if an economy has 10,000 versus 300 million participants? What macroeconomic purposes are served by small-scale models, and which require a scale closer to empirical reality? Do macroeconomies exhibit important regularities that simply cannot be generated using small-scale models?

ACE models implemented on modern computational platforms can include millions of heterogeneous interacting agents (Robert L. Axtell, 2001). The question is not whether this can be done but whether it should be done, and for what purposes.

III. Empirical Validation: Connecting to Data

Empirical validation is obviously important for more traditional economic models as well as for ACE models. Nevertheless, ACE researchers and critical observers both acknowledge that certain validation problems facing ACE researchers are special to the ACE methodology.

One problem involves degrees of freedom. ACE models often contain many parameters, and the claim is that the clever researcher can match any desired empirical feature using these degrees of freedom. This problem is compounded by the fact that functional forms and entire learning algorithms are at the disposal of the ACE researcher. Should a model use genetic algorithm learning or gradient ascent? Should information be stored in a neural network or in a linear forecasting model? This design flexibility suggests that ACE modeling tools might

almost be too rich in terms of fitting data. Another problem is that the properties of many ACE models are currently not well understood and not well motivated by observed human behavior.

Despite these problems, there are substantive reasons to hope that ACE researchers will eventually be able to push the bar of empirical validation very high. Consider how ACE models are currently validated.

One commonly-used method is to connect agent-level behavior to experiments with real people. In an early example of this type of research, Jasmina Arifovic (1996) found that the learning behavior of agents in an ACE foreign exchange model aligned well with participant behaviors in parallel human-subject laboratory experiments. A wide-ranging survey of more recent research along these lines can be found in Tesfatsion and Judd (2006, Chapter 19). It is clear that laboratory experiments will provide a crucial foundation for ACE modeling, since we still have relatively little information regarding how people learn in various field situations.

In addition to laboratory data comparisons, another direct and obvious empirical validation test for an ACE model is to replicate empirical features at many levels and at multiple time scales. For example, a reasonable test for a financial market would be to fit the equity premium. However, this alone would not be a very convincing test. Fitting a wide array of features ranging from the non-normality of returns at daily frequencies to the long-range correlations of volatility and dividend-price ratios would be a much more impressive test; see Blake LeBaron (2006).

Most standard macro models would stop here, but ACE models offer the possibility of further testing because they generate a rich micro economy underneath the macro data. For example, subject to data availability, one could test whether simulated trader behaviors aligned well with actual trader behaviors in data sets such as used by Terrance Odean (1999).

It is also critical to realize that, beyond simple time series, ACE models generate a complete distributional dynamics for a modeled economy. One can check features such as simulated wealth distributions and firm sizes and compare these with corresponding distributions from real economies; see, e.g., Domenico Delli Gatti et al. (2006). This micro-level distributional approach to empirical validation requires information that might not always be available. However, modern macroeconomies with their increasing reliance on electronic transactions continue to put more data into machine-readable form. Moreover, inclusive data sets at the micro level are not required. Limited but representative samples of real-world micro data provide an important check on the empirical plausibility of simulated micro-level distributions.

Since ACE macro models represent a fully functioning economy, they can also be empirically validated using procedures that are far from our traditional empirical toolbox. Human players can be allowed to interact with the models, and validation can take the form of testing the impact of human or machine players in a given situation. Also, simulated agents can be allowed to live with an actual data flow, as in the case of financial markets where one can observe the evolution of an order book with actual order flow augmented by simulated agent trades. A measure of validation would be how closely the simulated order book tracks the actual order book in the real market. See Michael Kearns and Luis Ortiz (2003) for an example of this type of work.

Certainly many ACE models could be validated in more traditional identify/estimate/test cycles as with standard models. However, they do raise some practical complications for the applied econometrician. Most importantly, their computational nature makes them costly to estimate. Analytics are most likely impossible, and computational methods such as simulated method of moments, while fine in theory, might be too computationally costly to undertake. Also, ACE models could display complicated dynamics related to nonlinearities in their

interconnections, raising difficult questions about stationarity and ergodicity. In short, the very properties that make ACE models so interesting to study can cause empirical headaches when estimating them. All is not lost for traditional econometrics, but researchers will need to be more creative in how they apply statistical techniques.

Finally, it is important to conjecture how policy makers could potentially make use of ACE macro models. Researchers at central banks might never decide to fit giant ACE macro models to data. It is quite possible that the usual reduced-form models will never be beaten at this task. Instead, policy makers could turn to ACE models to try to expand their current thinking.

For example, policy makers could use ACE models to explore major policy changes that diverge far from current policy settings. An ACE macro model with learning and adapting agents provides a kind of living version of a policy experiment, exploring the importance of behavioral adjustments in a given situation. Such models might also be well suited for analyzing an economy in extreme situations, e.g., for evaluating the probability of a financial crash and recommending appropriate recovery policies. In short, ACE macro models could thrive in the tails of distributions where standard empirical models are likely to fail.

References

- Arifovic, Jasmina 1996. "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies." *Journal of Political Economy*, 104(3):510-541.
- Axtell, Robert L. 2001. "Zipf Distribution of U.S. Firm Sizes." *Science* 293(5536), 1818-1820.
- Box, George E.P. 1979. "Robustness in the Strategy of Scientific Model Building," pp. 201-236 in *Robustness in Statistics*, eds. R.L. Launer and G.N. Wilkinson. NY:Academic Press.
- Colander, David. ed. 2006. *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*. NY: Cambridge University Press.

- Delli Gatti, Domenico, Edoardo Gaffeo, Mauro Gallegati, Gianfranco Giulioni, and Antonio Palestrini. 2006. *Emergent Macroeconomics: An Agent-Based Approach to Business Fluctuations*, Springer, 2008, forthcoming.
- Epstein, Joshua M., and Robert L. Axtell. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, MA: MIT Press.
- Fisher, Franklin M. 1993. *Aggregation: Aggregate Production Functions and Related Topics*. Cambridge, MA: MIT Press.
- Kearns, Michael, and Luis Ortiz. 2003. "The Penn-Lehman Automated Trading Project." *IEEE Intelligent Systems*, 18(6):22-31.
- Kirman, Alan P. 1992. "Whom or What does the Representative Agent Represent." *Journal of Economic Perspectives*, 6(2):117-136.
- LeBaron, Blake. 2006. "Agent-Based Financial Markets: Matching Stylized Facts with Style," pp. 221-235 in *Post Walrasian Macroeconomics*, ed. David Colander, NY:Cambridge University Press.
- Mayr, Ernst. 1997. *This is Biology: The Science of the Living World*. Cambridge, MA:Harvard University Press.
- Odean, Terrance. 1999. "Do Investors Trade Too Much?." *American Economic Review*, 89(5):1279-1298.
- Tesfatsion, Leigh. 2007. "Agents Come to Bits: Towards a Constructive Comprehensive Taxonomy of Economic Entities." *Journal of Economic Behavior and Organization*, 63(2):333-346.
- Tesfatsion, Leigh, and Kenneth L. Judd. eds. 2006. *Handbook of Computational Economics: Volume 2, Agent-Based Computational Economics*. Amsterdam, NL: Elsevier.