Macroeconomic Dynamics, **5**, 2001, 143–147. Printed in the United States of America.

INTRODUCTION TO THE SPECIAL ISSUE: NEW APPROACHES TO LEARNING IN MACROECONOMIC MODELS

JASMINA ARIFOVIC
Simon Fraser University

JAMES BULLARD
Federal Reserve Bank of St. Louis

The research questions addressed by the literature on learning in macroeconomics can be classified into four categories: First, there are issues related to the convergence and stability under learning in models with unique rational expectations equilibria. Authors here are concerned mainly with the learnability of a rational expectations equilibrium, as a measure of that equilibrium's plausibility as an observed outcome in an actual economy. Second, there are issues related to convergence and stability under learning in models with multiple rational expectations equilibria. In this case, learnability serves as an equilibrium selection device, helping economists decide which equilibria are the more likely to be actually observed among the many that exist under rational expectations. A third set of issues involves the examination of transitional dynamics that accompanies the equilibrium selection process. Following some type of unexpected strcutural change or change in policy regime, for instance, economies necessarily must follow temporary transitional paths to a rational expectations equilibrium associated with the new reality. Learning is sometimes used to help model such transitional dynamics. Finally, there are issues related to the examination of learning dynamics that are intrinsically different, even asymptotically, from the dynamics of the rational expectations versions of the models. In these cases, the learning dynamics do not converge to the rational expectations fixed points, and (unexploitable) expectational errors persist indefinitely. Some authors have tried to make use of this possibility in order to build explanations of otherwise puzzling macroeconomic phenomena based on constantly changing expectations.

The papers in this volume use a variety of approaches to address these issues. Although the specific economic problems being worked on come from many different sources, in a broader sense all authors are grappling with the same core questions of how economic and financial theory should think about expectations formation in formal models. All authors are exploring the degree to which departures from

Address correspondence to: Jasmina Arifovic, Department of Economics, Simon Fraser University, Burnaby, British Columbia V5A 1S6, Canada; e-mail: arifovic@sfu.ca.

© 2001 Cambridge University Press 1365-1005/01 $9.50

https://doi.org/10.1017/S1365100501019010 Published online by Cambridge University Press

rational expectations can influence economic outcomes in ways that better describe the data, while maintaining respect for the tenets of the rational expectations revolution, especially that there be no easily correctable forecast errors made by agents, and that rational expectations fixed points serve as important benchmarks. At the same time, authors wish to explore and emphasize the importance of the interaction between beliefs and outcomes—beliefs affect outcomes, but outcomes also affect beliefs—which is implicit in all models with an important role for expectations.

Many of the authors in this volume have gone to great lengths in an effort to show that their models have properties that match actual data generated in observed markets. We think that this is a healthy trend in the macroeconomics learning literature, which has, by necessity, often been focused mainly on technical, theoretical issues. The ability to move to models that can match otherwise puzzling data is the hallmark of growth in this area of economics and finance. There is also an impressive array of models employed in this volume with regard to how agents (or groups of agents) might learn. Several papers employ genetic or other evolutionary algorithms, whereas others have agents that use more conventional adaptive rules including least squares.

The issues covered in the special issue are indeed broad, which is perhaps testament to the pervasiveness of expectations in economic life. Two papers take up topics in finance: Blake LeBaron in "Evolution and Time Horizons in an Agent-Based Stock Market," and Bryan Routledge in "Genetic Algorithm Learning to Choose and Use Information." LeBaron develops a partial equilibrium model with two assets, a stock paying a stochastic dividend and a risk-free asset with a fixed interest rate. Agents have different time horizons about how far they should look into the past to evaluate the performance of different trading strategies that evolve over time. There is a separate pool of trading strategies that are represented as neural networks, the inputs of which are variables that are commonly used in real markets, such as past dividends, returns, the price/dividend ratio, and trends following technical trading indicators. Agents choose those strategies that yield the highest returns when evaluated at the agents' time horizons. Trading strategies are evolved as well. Simulations with populations of long-horizon agents resulted in the sequences of prices that correspond to the rational expectations equilibrium. On the other hand, simulations with populations inhabited by both short-horizon and long-horizon agents exhibit different behavior characterized by persistence in volatility of returns and in trading volume, with a strong contemporaneous connection between volume and volatility. These features are observed in actual markets. Another interesting result is that short-horizon agents are not driven out of the market by long-horizon agents. The volatility generated by the behavior of short-horizon agents prevents the long-horizon agents from gaining enough market strength to affect price movements.

Routledge examines the dynamics of learning in a version of the Grossman–Stiglitz model, a repeated, one-period economy with risk-free and risky assets. There is a costly signal about the risky asset's dividend. If traders choose to

https://doi.org/10.1017/S1365100501019010 Published online by Cambridge University Press

acquire it, they become informed. Otherwise, they remain uninformed and make an inference about the signal from the market-clearing asset price. Agent learning is modeled using a genetic algorithm. Agents trade for a fixed number of periods, *R*, without updating their decision rules, whereupon they update using a process modeled by a genetic algorithm. The choice of *R* influences the level of noise in the genetic-algorithm reproduction operator, and, in particular, smaller values of *R* are associated with increased levels of noise. Whether or not convergence to the rational expectations equilibrium occurs depends on the level of noise in the genetic updating procedure relative to that in the economy.

A third paper, "Learning and Excess Volatility" by James Bullard and John Duffy, straddles macroeconomics and finance by taking up the topic of stock-market volatility in an otherwise conventional macroeconomic model. Bullard and Duffy present a general equilibrium model with two assets, unbacked government liabilities and capital. Their analysis of the stability of the steady state with positive aggregate holdings of both assets under least-squares learning shows that *learning equilibria* may exist with limiting dynamics that are cyclical or possibly more complicated. Bullard and Duffy calibrate their model to match U.S. time series on asset returns over the past 100 years, and find artificial economies that partially capture excess volatility phenomena, namely the much greater variation in real returns to capital compared with underlying fundamentals as captured by per-capita consumption growth rates. This research provides support for the hypothesis that much of the observed volatility in capital asset returns may be due to expectations that are continually being revised, whereas fundamental factors are not changing in quantitatively important ways.

The implicit idea in all of these papers, that observed financial market volatility may, at least in part, be endogenously generated by the financial markets themselves (as opposed to being strictly due to the effects of exogenous "shocks" hitting the market) has tremendous intuitive appeal, and is often casually discussed by market participants, policymakers, and academics. Formal modeling of such phenomena is difficult, however, and these three papers contain some ideas about how one might go about constructing such models. In terms of the four issue areas listed at the outset, all three of these papers provide details on the convergence properties of the learning dynamics vis-à-vis a rational expectations benchmark (issue 1) and on the character of the learning dynamics that persistently depart from those under rational expectations (issue 4), often in great quantitative detail.

Jasmina Arifovic, in "Performance of Rational and Boundedly Rational Agents in a Model with Persistent Exchange-Rate Volatility," is also concerned with learning-induced endogenous volatility, but this time in the markets for foreign exchange. Arifovic studies a model of exchange rates in which agents learn using a version of the stochastic replicator dynamic. Under rational expectations, exchange-rate indeterminacy holds, so that learning plays a role in selecting among equilibria in the model (issue 2). However, Arifovic shows that the stationary rational expectations equilibria of the model are unstable under evolutionary adaptation.

https://doi.org/10.1017/S1365100501019010 Published online by Cambridge University Press

The resulting persistent learning dynamics exhibit volatility in the exchange rates that is robust to changes in the parameters of the economic model as well as to the details of the learning algorithm itself. Application of evolutionary algorithms in macroeconomic models are often simulation-based because the environment is usually too complex to be characterized analytically, but Arifovic has formulated a more tractable model with the stochastic replicator dynamic and has been able to derive one- and two-period-ahead forecasts of the model's endogenous variables. Importantly, she shows how an omniscient, rational agent would fare relative to the boundedly rational agents who inhabit the economy she studies.

Topics closer to traditional macroeconomics are examined in three additional papers. Cars Hommes and J. Barkley Rosser, Jr., in "Consistent Expectations Equilibria and Complex Dynamics in Renewable Resource Markets," study the idea of consistent expectations equilibrium in the context of the economics of the fishery. Fishery dynamics often are studied as a prelude to capital theory because the dynamic structures have much in common across the two applications. A consistent expectations equilibrium means, roughly, that although the agents in the model are adaptive learners, and the dynamics under learning do not converge to a rational expectations equilibrium, the agents would not observe any obvious information that would cause them to depart from their linear forecast rules. In their model, Hommes and Rosser are able to provide conditions under which a chaotic or even noisy chaotic consistent expectations equilibrium exists. Again, this paper is concerned mainly with the fourth issue in the taxonomy listed at the outset.

Todd Allen and Christopher Carroll, in "Individual Learning About Consumption," examine learning in a stochastic, intertemporal optimization problem. The solutions to this class of problems for plausible preferences and a realistic specification of stochastic shocks have only recently been found using computationally intensive numerical methods. The solutions usually involve complicated, nonlinear consumption rules, which actually appear to be quite consistent with empirically observed consumption behavior. How, then, could actual households find such a rule when economists only now have the computational ability to uncover it? Allen and Carroll show that the exactly correct nonlinear consumption policy rule can be approximated very closely, in utility terms, by an intuitive linear form that should be much simpler to learn. However, the results in the paper, using a form of trial-and-error learning, show that even the simplified linear consumption function is very difficult for individuals to learn. Allen and Carroll suggest that some type of "social learning," in which consumers exchange information, is a potential mechanism to explain how households learned the optimal consumption rule.

Gian-Italo Bischi and Ramon Marimon, in "Global Stability of Inflation Target Policies with Adaptive Agents," are the authors in this volume who most explicitly address issues of how policymakers might use results on the stability of equilibria under adaptive learning in order to assess the costs and benefits of pursuing certain policies. Theirs is a model of monetary policy rules interacting with a government budget constraint. They provocatively label rational expectations a

https://doi.org/10.1017/S1365100501019010 Published online by Cambridge University Press

misspecification, which can be misleading because it leaves policymakers without a good sense of the stability properties of the equilibria they are trying to implement. Thus, policies designed under the assumption of rational expectations can lead to undesirable outcomes.

We hope these interesting papers provide some stimulus for additional research on the effects of learning in macroeconomic contexts.

TABLE OF CONTENTS FOR THE SPECIAL ISSUE

https://doi.org/10.1017/S1365100501019010 Published online by Cambridge University Press