

James B. Bullard

James B. Bullard is an economist at the Federal Reserve Bank of St. Louis. David H. Kelly provided research assistance.

Learning, Rational Expectations and Policy: A Summary of Recent Research

IN THE THREE DECADES since the publication of the seminal work on rational expectations, a steely consensus has been forged in the economics profession regarding acceptable modeling procedures.¹ Simply stated, the consensus is that economic actors do not persist in making foolish mistakes in forecasting over time. People are presumed to be able to both detect past patterns in their prediction errors and base their behavior on the "best possible" forecast of future economic conditions. In the classic phrase of Robert Lucas, the predictions of individuals must be "... free of systematic and easily correctable biases."² The current wide acceptability of this notion is testament to the success of the rational expectations revolution.

Unfortunately, the consensus that in equilibrium systematic forecast errors should be eliminated has been insufficient to end the debate in macroeconomics over expectational assumptions. In particular, some current research examines the idea that *how* systematic forecast errors are

eliminated may have important implications for macroeconomic policy. Researchers who focus on this question are studying what is called "learning," because any method of expectations formation is known as a learning mechanism.³

This paper provides a synopsis of some of the recent research on learning. Three important points are emphasized within the context of the survey. The most salient point is the close relationship between learning issues and macroeconomic policy. In fact, the topic attracts attention precisely because of its perceived policy implications. The second point is more subtle: learning is implicitly an integral part of rational expectations models, and current research only makes this fact explicit. There is little prospect that one can avoid the study of learning by assuming rational expectations. The third point is that including learning in macroeconomic models is unlikely to either confirm or overturn completely the results from rational expectations macro-models. Instead, the concept of rational expectations equilibrium seems to provide the appropri-

¹The seminal work is Muth (1960, 1961).

²Lucas (1977), p. 224.

³The phrase "learning mechanism" is used in this paper in a broad way, even though it conjures up different images

for different readers. This terminology developed because rational expectations is justified by the notion that people eliminate systematic forecast errors over time, and the dynamic elimination of errors is a definition of learning.

ate benchmark for the study of learning in that, when systems converge under learning, they typically converge to stationary rational expectations equilibria.

The next section provides a non-technical introduction to learning and rational expectations. The subsequent section looks at the effects of introducing learning through a simple example attributable to Albert Marcet and Thomas Sargent. Some interpretation and discussion of the example is offered in the third section, along with a review of other attempts to introduce learning into macroeconomic models. The final section provides summary comments.

THE LEARNING IMPLICIT IN RATIONAL EXPECTATIONS

Why Study Learning?

Since decisions made today affecting productive behavior are presumed to be based, in part, on individual assessments of the future, macroeconomic theories and models generally have provided a role for expectations. Around 1970, however, researchers began to realize that the policy implications of their models were often quite sensitive to the choice of expectational assumptions.⁴ This failure of robustness has been increasingly apparent in the last 20 years and has drawn increased attention to the problem of how expectations are formed. This line of research constitutes what has been called the rational expectations revolution.

A capsule characterization of rational expectations contains the following themes: (1) In equilibrium (a steady state in dynamic terms), expectations are "correct" in the sense that individuals make no systematic forecast errors; (2) Individuals use all available information (as defined by the researcher) in forming forecasts; (3) Expectations vary with changes in government policy; and (4) Individuals know "the

model" and thus can predict as well as the economist manipulating the model.⁵ These are the tenets of the theory and are often espoused by its advocates.

The first tenet, the heuristic notion that individuals eliminate systematic forecast errors, is the one most responsible for the rise of the rational expectations hypothesis. In a deterministic environment, this idea implies that, once learning is complete, people have perfect foresight. In a stochastic environment, it means that the remaining forecast errors are white noise.

Since the consensus is that the elimination of systematic forecast errors is a sensible postulate, all reasonable long-run equilibria must be rational expectations equilibria. Macroeconomists generally are interested in the effects of changes in policy parameters at these steady-state equilibrium points.⁶ Given such equilibria, there is at least one reason why the explicit specification of learning could matter: in a model with multiple rational expectations equilibria, the learning mechanism may serve to select the actual outcome.⁷ The next section contains an example of learning as a selection mechanism.

Representing Learning Via Econometric Techniques

Macroeconomists generally have avoided specifying explicitly the optimal learning mechanism underlying rational expectations for a number of reasons. First and foremost, they considered rational expectations a shortcut in expectations modeling that made explicit specification unnecessary.⁸ Further, ascertaining the full implications of the rational expectations hypothesis turned out to be a difficult problem; presumably, these implications must be understood before the issue of learning can be investigated.

There was, however, at least one additional reason why learning was essentially ignored—explicit specifications of how agents form expect-

⁴Sargent (1987) has documented some of this early research in rational expectations.

⁵A more formal approach to defining rational expectations is pursued in the next section.

⁶For a discussion of policy along the transition path, see Taylor (1975).

⁷This is the same as saying that learning provides a stability theory for rational expectations equilibria. In the best case, given initial conditions and some parameter values, statements claiming that the dynamic evolution of the economy always leads to a particular stationary equilibrium

can be made. This is called global stability. Alternatively, one can define local stability, where a particular stationary equilibrium is stable only if the initial conditions are near that equilibrium.

⁸Lucas' comment, ". . . take the rational expectations equilibrium . . . as the model to be tested and view [learning] as . . . an adjunct to the theory that serves to lend it plausibility," hints at this pragmatism. See Lucas (1987a), p. 231.

tations have often been attacked. In particular, one key criticism of the adaptive expectations formulation was that it allowed systematic errors to persist over time.⁹

The notion of adaptive expectations is essentially that people predict the future value of a variable via a geometric distributed lag (or Koyck lag) of its past values. One of Muth's results was that adaptive expectations is an optimal prediction method if the variable being forecast follows a random walk.¹⁰ Muth's result makes it clear that, from the beginning, rational expectations theorists had econometric techniques in mind when thinking about ways to model optimal forecasting.

The idea of looking to econometrics to model learning obviously can be extended, since the Koyck lag is only one of many available econometric techniques. In principle, it should be possible to take advantage of the developments in econometric theory to shed light on the problem of how expectations are formed. Furthermore, if econometric methods are to be applied to solve the inference problem faced by individuals, a good place to start would be some variant of least squares regression. This is the logic behind the recent work on least squares learning, an example of which is presented below.

The application of econometric theory to the expectations problem was vigorously pursued recently because the mathematical knowledge to analyze the problem became available.¹¹ This technology was developed in the engineering literature and has been extended to economics by Albert Marcet and Thomas Sargent.¹² Economists can now determine, therefore, whether it makes any difference for policy implications if the learning mechanism behind rational expectations is explicitly specified. Viewed this way, the work on learning can be considered part of a continuing effort to understand the microfoundations of expectations formation.

The impetus for a detailed analysis of learning has come from both advances in research technology and one pressing problem: many rational expectations models are characterized by multiple equilibria. Moreover, these different equilibria can have different policy implications, as the next section illustrates.¹³

LEARNING AND THE UNPLEASANT MONETARIST ARITHMETIC

This section examines a simple expository example of Marcet and Sargent;¹⁴ it is a simplified version of the "unpleasant monetarist arithmetic" of Sargent and Neil Wallace.¹⁵ The example is meant only to illustrate the types of issues that can arise; it is not intended as a definitive statement of the effects of including learning in economic models. As the next section will discuss, these effects are still uncertain.

The model's most important feature is that there are two steady states (high inflation and low inflation) with differing policy implications. At the low-inflation steady state, a permanent increase in the government deficit implies a permanent increase in the stationary inflation rate. However, increasing the government deficit leads to lower steady-state inflation if the economy is in the high-inflation stationary equilibrium. Because the policy implications of the model hinge on the choice, it is desirable to name one of the stationary states as the more likely outcome. The policy conclusions turn out to be different under learning as opposed to perfect foresight, even though systematic forecast errors are eventually eliminated under learning and the only possible limit points of the system are the two perfect foresight stationary states. This occurs because the stability properties of the two steady states are altered under learning as opposed to perfect foresight.

In Marcet and Sargent's model, fiat currency is the only money. The model consists of two

⁹See Cagan (1956) and Nerlove (1958) for expositions of the adaptive expectations hypothesis.

¹⁰See Muth (1960).

¹¹Formally, the problem is one of stability in a recursive stochastic environment. This means that the economist wants to take uncertainty, time and feedback (beliefs affect outcomes and outcomes affect beliefs) into account in the same model to find out the conditions under which convergence will occur.

¹²See Marcet and Sargent (1988, 1989a,b,c).

¹³The selection of an actual outcome among these potential outcomes is the subject of a spirited debate in the

macroeconomics literature which is not considered in detail here. A number of authors have suggested methods of choosing a unique equilibrium through *deeper theorizing*, that is, by applying certain restrictions to the model to eliminate unsavory equilibria. See, for instance, McCallum (1983), Taylor (1977) and Evans (1985, 1986). Some recent evaluations of these selection criteria suggest they are less than satisfactory. See, for instance, Pesaran (1988) and Boyd and Dotsey (1990).

¹⁴See Marcet and Sargent (1989b).

¹⁵See Sargent and Wallace (1981, 1987).

equations relating the price level to the stock of currency. Let p_t be the price level and c_t be the stock of currency per capita at time t . The model is given by

$$(1) p_t = \lambda F_t p_{t+1} + \gamma c_t$$

$$(2) c_t = c_{t-1} + \delta p_t$$

$$(3) F_t p_{t+1} = \beta_t p_t$$

where $0 < \lambda < 1$; $\gamma, \delta > 0$; $p_t, c_t > 0$; and $c_0 > 0$ is given. The term $F_t p_{t+1}$ is the forecast at time t of the price at time $t+1$, where β_t is the expected (gross) inflation rate at time t . Rearranging equation 1 with c_t on the left-hand side yields a linear money demand equation. Equation 2 can be interpreted as a government budget constraint, where the government chooses c_t to maintain a fixed real deficit δ . The model has some microfoundations, since it can be derived in an overlapping generations framework with utility-maximizing individuals. The model generates time paths for c_t and p_t given an additional assumption on how expectations are formed. The remainder of the section will develop and compare the results under two alternative assumptions: rational expectations (or perfect foresight) and learning as described by a least squares autoregression.

Perfect Foresight Dynamics

First, close the model by assuming individuals have rational expectations, which in this case is equivalent to perfect foresight because the model is deterministic. Then

$$(4) \beta_t = \frac{p_{t+1}}{p_t}$$

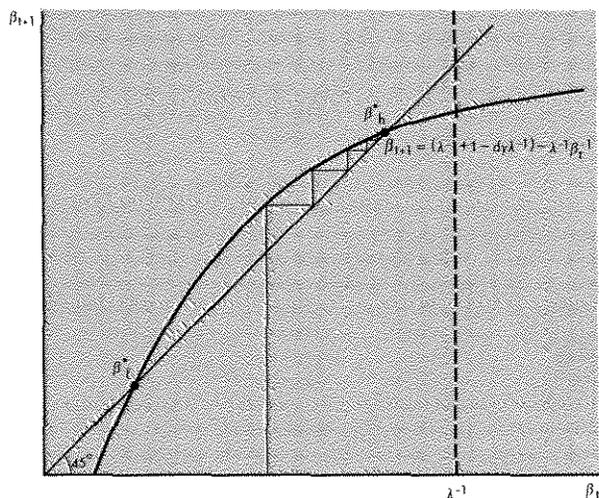
Equations 1-4 can be rearranged to yield:¹⁶

$$(5) \beta_{t+1} = [\lambda^{-1} + 1 - \delta\gamma\lambda^{-1}] - \lambda^{-1}\beta_t^{-1}$$

A rational expectations equilibrium is a sequence $\{\beta_t\}_{t=0}^{\infty}$ that satisfies (5), and it determines the equilibrium sequences for c_t and p_t . The difference equation 5 has two roots, β_t^* and β_n^* , where $1 < \beta_t^* < \beta_n^* < \lambda^{-1}$; these roots are the stationary states of the model. The real government deficit δ is a shift parameter in equation 5.

The differing policy implications at the two steady states are illustrated in figure 1, which is borrowed from Marcet and Sargent's paper. A

Figure 1
Perfect Foresight



permanent increase in the government deficit δ shifts the entire curve down in the figure. At the low-inflation steady state, β_t^* , this parameter change raises the stationary inflation rate. But at the high-inflation steady state, β_n^* , a permanent increase in δ lowers the stationary inflation rate. In the first case, the comparative statics are "classical," while in the second they are "perverse." Any policy advice about inflation based on this model differs depending on which asymptotic outcome is considered more likely to be observed.

The high-inflation stationary state, β_n^* , is the attractor for all initial values β_0 between β_t^* and λ^{-1} .¹⁷ The low-inflation stationary state is attainable only if $\beta_0 = \beta_t^*$. If the initial conditions are between zero and β_t^* no equilibrium sequence exists. Altogether, there are many possible equilibrium sequences in the model, and they can be indexed by β_0 . Of these, all but one converge to the high-inflation stationary state β_n^* , where the comparative statics are perverse.

In related work, Sargent and Wallace have used similar arguments to claim that monetarist models can yield "unpleasant" results. Under rational expectations, the "bad" stationary state with the perverse comparative statics is the

¹⁶See the appendix for the detailed derivation of this and subsequent statements in this section.

¹⁷The appendix also describes how to determine stability properties in figures 1 and 2. Only initial values between zero and λ^{-1} are feasible.

eventual outcome for virtually all initial conditions, provided an equilibrium exists.

Perfect foresight is a strong assumption, but it can be made palatable by arguing that people eliminate systematic errors in their forecasts over time. Therefore, perfect foresight may provide a good approximation once the steady state has been attained and learning is complete.

Marcet and Sargent have developed techniques to analyze this argument in detail. Operationally speaking, people can be viewed as using some type of statistical technique to infer a future price from available data. One widely known technique is ordinary least squares (OLS). The next portion of the paper analyzes the model using the assumption of least squares learning.

Least Squares Learning Dynamics

To develop results analogous to the perfect foresight case under least squares learning, Marcet and Sargent replace equation 4 with

$$(4') \beta_t = \left[\sum_{s=1}^{t-1} p_s^2 \right]^{-1} \left[\sum_{s=1}^{t-1} p_s p_{s-1} \right].$$

People form their forecasts of future inflation by calculating β_t , which is found via a first order autoregression using available data through time $t-1$. The difference equation that describes the evolution of β_t is in this case given by:

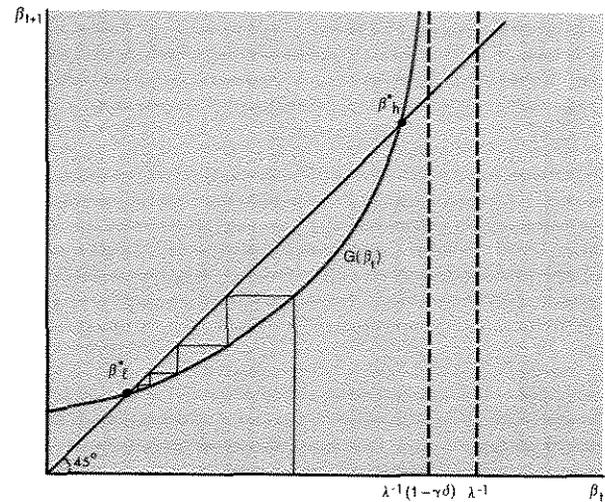
$$(5') \beta_t = \beta_{t-1} + \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_s^2} \left[\frac{(1 - \lambda\beta_{t-2})}{(1 - \lambda\beta_{t-1})} G(\beta_{t-1}) - \beta_{t-1} \right],$$

where

$$(5'') G(\beta_{t-1}) = \frac{1 - \lambda\beta_{t-1}}{1 - \lambda\beta_{t-1} - \gamma\delta}.$$

While equation 5' is quite complicated, it can be interpreted without too much difficulty. Near a stationary equilibrium, the term $(1 - \lambda\beta_{t-2})/(1 - \lambda\beta_{t-1})$ is close to unity, while the term multiplying the brackets is always between zero and one. Therefore, near a steady state, equation 5' states that β_t is approximately a convex combination of β_{t-1} and $G(\beta_{t-1})$. That is, near a stationary equilibrium, the projected gross inflation rate is a weighted average of last period's projected gross inflation rate and a certain function

Figure 2
Least Squares Learning



$G(\beta_{t-1})$. Alternatively, one can view (5') as last period's prediction "updated" by a certain small amount.

The possible steady states under learning, β_t^* and β_t^* , are the same steady states possible under perfect foresight. This captures the notion that systematic forecast errors are eliminated over time and suggests that stationary rational expectations equilibria provide a good benchmark for the study of learning. Marcet and Sargent claim, in fact, that when least squares learning mechanisms converge, they always converge to rational expectations equilibria.¹⁸

Now that a method of expectations formation has been added to the model, however, some aspects have changed. Marcet and Sargent have completed an analysis of the complicated difference equation given by (5'). They show that equation 5' describing the dynamics under least squares learning is closely related to the following simpler equation:

$$(6) \beta_t = G(\beta_{t-1}).$$

For the purposes of this exposition, it will suffice to analyze (6).¹⁹ The graph that Marcet and Sargent use to describe the dynamics of (6) is given in figure 2.

¹⁸See Marcet and Sargent (1989a,b,c).

¹⁹Readers interested in the actual dynamics under least squares learning are referred to Marcet and Sargent (1989b).

The approximate dynamics under least squares learning as described by (6) indicate that the low-inflation steady state β_l^* is the attractor for all initial conditions β_0 between zero and β_h^* . The high-inflation stationary state is attained only if $\beta_0 = \beta_h^*$. If β_0 is between β_h^* and λ^{-1} , no equilibrium exists under least squares learning. The policy advice emanating from this analysis is therefore approximately the opposite from that offered in the perfect foresight case. Under least squares learning, the low-inflation stationary equilibrium with the classical comparative statics is the asymptotic outcome for virtually all initial conditions, provided that an equilibrium exists.

It is important to emphasize at this point that no generic result links learning with "good" equilibria and rational expectations with "bad" equilibria. In some other model, rational expectations may be associated with good outcomes and learning with bad outcomes. Nevertheless, some may take comfort in the fact that, at least for the present case, the comparative statics are once again "classical" under a plausible assumption about how expectations are formed.

Experimental Evidence

So far, the example has illustrated that, while the potential stationary equilibrium inflation values are exactly the same under perfect foresight and least squares learning, the stability properties are approximately reversed, leading to opposite policy implications. Nothing has been said about whether perfect foresight or least squares learning is the better description of human behavior in this context. Recently, however, Ramon Marimon and Shyam Sunder have gathered some experimental evidence that bears on this issue.²⁰

Marimon and Sunder summarize the results from a set of controlled experiments that use students as subjects. Their model is an overlapping generations version of the one used by Marcet and Sargent. The students were oriented to the context of the model and asked to forecast inflation, with monetary rewards for more accurate predictions. The authors were especially interested in characterizing the actual outcomes for inflation in the model where the decisions are made by humans.

The results of these experiments indicate that actual inflation tends to cluster around the low-

inflation stationary equilibrium. Marimon and Sunder never observed a tendency of the inflation rate to converge to the high inflation stationary equilibrium in any of their experiments. Based on the forecasts made by their subjects, the authors conclude that least squares learning provides a good approximation to observed behavior. While the work has some caveats and is open to interpretation, it at least provides preliminary evidence on the viability of assuming least squares learning.

LEARNING IN OTHER MACROECONOMIC CONTEXTS

The notion of a rational expectations revolution stresses the much greater emphasis that macroeconomists have placed on the plans of individuals in macroeconomic models since 1970. One way to think about current research in the area is to categorize models according to their treatment of expectations formation. While most authors want to analyze models where expectations are "rational," several different approaches have been taken. These different approaches *within* rational expectations macroeconomics reflect unresolved issues in the theory of expectations formation. Various views have been espoused about the specific information that individuals take into account when they are planning for the future; such issues become salient partially because rational expectations theory does not specify a method of learning. Thus, at least three different categories of macromodels that attempt to analyze individuals with rational expectations can be identified.

Forecast Functions That Use Only Historical Data

The first group contains models in which individuals forecast using only the history of the economy as a guide. In the previous example, for instance, learning was modeled as being based on the historical price sequence that people observed. The approach could be extended relatively easily to include historical time series on other variables in the individuals' forecast functions. While authors assuming rational expectations normally do not make statements about the learning mechanism implicitly underlying their models, it seems that an appropriate

²⁰See Marimon and Sunder (1990).

use of available history is what many have in mind.

Considering the use of historical data only, one might be tempted to conclude (say, from looking at the results of the model from the last section) that the only possible long-run outcomes are stationary states. In fact, most discussion of rational expectations applies to dynamic steady states. This notion of long-run equilibrium as a stationary state, which dominates most macroeconomic thought, was challenged recently by an argument that the dynamic equilibrium of an economy might be periodic or chaotic, even under rational expectations and perfect competition.²¹ Therefore, the artificial economy may not converge to any of the multiple rational expectations steady states, instead remaining in a permanent periodic equilibrium. Cyclical equilibria are important because, presumably, policy implications are altered if long-run equilibria are periodic.²²

Forecast Functions That Include the Beliefs of Others

A second group of macromodels contains forecast functions where, in addition to historical time series, the beliefs of others also play a role.²³ Consider, for instance, the view of Cass and Shell: "In seeking to optimize his own actions, an economic actor must attempt to predict the moves of all other economic actors."²⁴ Such consideration adds a new element to the inference process.²⁵ Of course, the notion of interaction among individuals (especially individual firms) has a long history in economics; it is eliminated in the Arrow-Debreu competitive equilibrium framework by the assumption of a large population.

Still, a participant contemplating a forecast may well be concerned with the aggregate expectations of the remaining players, or merely with the beliefs of one other player, such as the government in a monetary policy game.²⁶ The

game—theoretic macromodels that take this latter view are structurally very different from the traditional models, even though they are both based on rational expectations. Most importantly, they are different in terms of policy implications: these models often yield "... an equilibrium that is extremely sensitive to the public's beliefs about the monetary authority's preferences."²⁷ This conclusion does not hold in a model in which individual forecasts are a function of historical data only, such as the model of section two, because it excludes the preferences of the government and the public's beliefs about them.

Forecast Functions That Include Frivolous Variables

When individuals forecast based only on relevant past history, they are sometimes said to be basing expectations on "fundamentals." When deciding what is relevant, however, some people may rationally take into account what seems to an observer to be irrelevant information in the form of a frivolous variable (typically called the sunspot variable). This variable acquires importance in determining the actual outcomes of the economy only because people think it is important. The frivolous variable serves only to signal changes in expectations. Once some individuals take the frivolous variable into account in forming expectations, it becomes rational for all others to do so, since the variable actually does influence outcomes.²⁸

Predictions in models with this type of individual forecast function are based not only on relevant and objectively irrelevant historical information, but also on the expectations of others, since the beliefs of others are taken into account when the frivolous variable is assessed. The literature on sunspot equilibria provides a third but distinctly different strain of rational expectations macroeconomics, with distinctly different policy implications.²⁹ In particular, "... a con-

²¹See, for instance, Grandmont (1985). Chaos means that the equilibrium sequence is aperiodic but bounded and displays sensitive dependence on initial conditions. For a general discussion of chaos, see Butler (1990).

²²In the overlapping generations model, for example, the existence of periodic equilibria is disturbing because it implies that welfare varies from generation to generation.

²³John Maynard Keynes, for instance, discussed this type of forecast function in some detail. See Keynes (1936), p. 156.

²⁴Cass and Shell (1989).

²⁵But not necessarily a strategic element. See Rogoff (1989).

²⁶This theme is outlined in detail in the volume edited by Frydman and Phelps (1983).

²⁷Rogoff (1989).

²⁸See, for instance, Azariadis (1981).

²⁹See Cass and Shell (1989).

sideration of the complete set of possible equilibria, [including the sunspot equilibria], associated with a given policy regime may alter one's evaluation of the relative desirability of alternative policies, relative to the conclusion that one might reach if one considered only a single possible equilibrium . . ."³⁰

The existence of sunspot equilibria raises the question of whether models with learning have dynamics converging to them. One author, Michael Woodford, has shown that exactly this sort of dynamics is possible.³¹ Also, in general, the extent of the sunspot phenomenon is wide-ranging since there are no limits to the number of possible frivolous variables.

These three approaches to macroeconomics differ according to their alternative assumptions about what it means to assume rational expectations. Since the theory provides no method of expectations formation (that is, no learning process), researchers are free to provide their own: perhaps individuals base their expectations on sunspots, or the expectations of others, or a straightforward application of classical or Bayesian inference to historical time series. As part of the legacy of the rational expectations revolution, all three approaches place heavy emphasis on the role of individuals' views of the future in influencing current macroeconomic equilibria.

Unfortunately, there is little prospect that econometric analysis will decide which version of rational expectations is correct. Because the theory is not well-defined, the empirical tests are unconvincing.³² While economists want to assume that expectations are rational, the implications of this consensus for modeling and for policy are in doubt.

SUMMARY: THE DIFFICULTY OF DEFINING OPTIMAL BELIEFS

Several lines of research, each in its own way, are attempting to extend the rational expectations hypothesis to include learning. The first and most obvious is a direct attempt to find mild assumptions showing that most reasonable methods of expectations formation converge to particular types of rational expectations equilibria. If this can be done, the concept of rational expectations equilibrium can be said to provide

a good approximation to the concept of long-run competitive equilibrium. This literature, however, generally ignores the problem of the expectations of others and of frivolous variables.

Even simple learning mechanisms often have not yielded the outcomes that intuition suggests might have occurred. The notion that a reasonable method of expectations formation must give rise to a dynamic system that always converges to a certain *a priori* plausible rational expectations equilibrium has been gradually eroded. The general conclusion so far seems to be that explicitly introducing learning into macroeconomic models is unlikely to provide a widely applicable selection criterion for rational expectations equilibria. That is, in a rational expectations model with multiple equilibria, introducing a learning mechanism does not appear to reduce the set of potential outcomes in any meaningful way.

One reason for this disappointing result is that it is difficult to define optimal learning. Not only is the class of plausible mechanisms quite large, it is also hard to limit the learning techniques under study to one that can be justified by some optimality argument. One of the biggest problems is that the usual statistical techniques are, strictly speaking, inapplicable to the problem of individual inference in the context of macroeconomic models.³³ The source of difficulty is that, in models with expectations, there is an aspect of simultaneity in the sense that beliefs affect outcomes and outcomes affect beliefs. In order to apply standard inference techniques, people must be unaware of the effects of beliefs on outcomes. Making this assumption is unsatisfactory, however, because it means that individuals ignore relevant and potentially useful information when forming their forecasts.

Nevertheless, work continues on ways to explicitly model learning; Marcet and Sargent provide one example. Some attempt is made to choose the learning mechanism via an optimization criterion, and the asymptotic properties of the implied systems are then analyzed. This research agenda is difficult and relies to a large extent on mathematical machinery only recently developed to study such systems.

The policy implications of including learning have been emphasized in this summary of the recent research. In models with multiple ra-

³⁰Woodford (1988).

³¹See Woodford (1990).

³²Webb (1988).

³³See Marcet and Sargent (1989a,b,c).

tional expectations equilibria, the consensus opinion that people eliminate systematic forecast errors is generally not enough to determine the actual outcome. Moreover, because different rational expectations equilibria have different implications for policy, merely stating that the economy will converge to one of the possible stationary equilibria, without saying which one, is insufficient for useful policy advice. Work on how expectations are formed has only just begun; one hopes that it will lead economists eventually to understand how equilibrium is achieved and what their policy advice is likely to produce.

REFERENCES

- Azariadis, Costas. "Self-Fulfilling Prophecies," *Journal of Economic Theory* (1981), pp. 380-96.
- Boyd, John H., III, and Michael Dotsey. "Interest Rate Rules and Nominal Determinacy," Working paper No. 90-1, University of Rochester and the Federal Reserve Bank of Richmond (1990).
- Butler, Alison. "A Methodological Approach to Chaos: Are Economists Missing the Point?" this *Review* (March/April 1990), pp. 36-48.
- Cagan, Phillip. "The Monetary Dynamics of Hyperinflation," in Milton Friedman, ed., *Studies in the Quantity Theory of Money* (University of Chicago Press, 1956), pp. 25-120.
- Cass, David, and Karl Shell. "Sunspot Equilibrium in an Overlapping-Generations Economy with an Idealized Contingent Commodities-Market," in William A. Barnett, John Geweke, and Karl Shell, eds., *Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinearity* (Cambridge: Cambridge University Press, 1989), pp. 3-20.
- Evans, George. "Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models," *Quarterly Journal of Economics* (1985), pp. 1217-33.
- Evans, George. "Selection Criteria for Models with Non-uniqueness," *Journal of Monetary Economics* (1986), pp. 147-57.
- Frydman, Roman, and Edmund Phelps. *Individual Forecasting and Aggregate Outcomes: Rational Expectations Examined*, (Cambridge University Press, 1983).
- Grandmont, Jean-Michel. "On Endogenous Competitive Business Cycles," *Econometrica* (1985), pp. 995-1045.
- Harvey, A. C. *The Econometric Analysis of Time Series*, (Philip Allan, 1981).
- Keynes, John Maynard. *The General Theory of Employment, Interest and Money* (MacMillan, 1936).
- Lucas, Robert E., Jr. "Understanding Business Cycles," in Karl Brunner and Allan H. Meltzer, eds., *Stabilization of the Domestic and International Economy*, Carnegie-Rochester Conference Series on Public Policy (Amsterdam: North-Holland, 1977), pp. 215-39.
- _____. "Adaptive Behavior and Economic Theory," in Robin M. Hogarth and Melvin W. Reder, eds., *Rational Choice: The Contrast Between Economics and Psychology* (University of Chicago Press, 1987a), pp. 217-42.
- _____. *Models of Business Cycles*, Yrjo Jahnsson Lectures, (Basil-Blackwell, 1987b).
- Marcet, Albert, and Thomas J. Sargent. "The Fate of Systems with Adaptive Expectations," *American Economic Review* (1988), pp. 168-71.
- _____. "Convergence of Least-Squares Learning Mechanisms in Self Referential, Linear Stochastic Models," *Journal of Economic Theory* (1989a), pp. 337-68.
- _____. "Least-Squares Learning and the Dynamics of Hyperinflation," in William A. Barnett, John Geweke, and Karl Shell, eds., *Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinearity* (Cambridge: Cambridge University Press, 1989b), pp. 119-40.
- _____. "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy* (1989c), pp. 1306-22.
- Marimon, Ramon, and Shyam Sunder. "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence," Working paper, University of Minnesota and Universitat Autònoma de Barcelona, and Carnegie-Mellon University (January 1990).
- McCallum, Bennett T. "On Non-Uniqueness in Rational Expectations Models: An Attempt at Perspective," *Journal of Monetary Economics* (1983), pp. 139-68.
- Muth, John F. "Optimal Properties of Exponentially Weighted Forecasts," *Journal of the American Statistical Association* (1960), pp. 299-306.
- _____. "Rational Expectations and the Theory of Price Movements," *Econometrica* (1961), pp. 315-35.
- Nerlove, Marc. "Adaptive Expectations and Cobweb Phenomena," *Quarterly Journal of Economics* (1958), pp. 227-40.
- Pesaran, M. Hashem. *The Limits to Rational Expectations* (Basil-Blackwell, 1988).
- Rogoff, Kenneth. "Reputation, Coordination, and Monetary Policy," in Robert J. Barro, ed., *Modern Business Cycle Theory* (Harvard University Press, 1989), pp. 236-64.
- Sargent, Thomas J. *Macroeconomic Theory* (Academic Press, 1987).
- Sargent, Thomas J. and Neil Wallace. "Some Unpleasant Monetarist Arithmetic," *Federal Reserve Bank of Minneapolis Quarterly Review* (Fall, 1981), pp. 1-17.
- _____. "Inflation and the Government Budget Constraint," in Assaf Razin and Efraim Sadka, eds., *Economic Policy in Theory and Practice* (MacMillan, 1987), pp. 170-200.
- Taylor, John B. "Monetary Policy during a Transition to Rational Expectations," *Journal of Political Economy* (1975), pp. 1009-21.
- _____. "Conditions for Unique Solutions in Stochastic Macroeconomic Models with Rational Expectations," *Econometrica* (1977), pp. 1377-85.
- Webb, Roy H. "The Irrelevance of Tests for Bias in Series of Macroeconomic Forecasts," *Federal Reserve Bank of Richmond Economic Review* (November/December 1988), pp. 3-9.
- Woodford, Michael. "Monetary Policy and Price Level Indeterminacy in a Cash-in-Advance Economy," Working paper, University of Chicago (December 1988).
- _____. "Learning to Believe in Sunspots," *Econometrica* (1990), pp. 277-308.

Appendix

Details of the Marcet and Sargent Model

This appendix provides the mathematical details for Marcet and Sargent's model outlined in section two. For convenience, the model is reproduced here; it consists of the three equations:

$$(1) p_t = \lambda F_t p_{t+1} + \gamma c_t$$

$$(2) c_t = c_{t-1} + \delta p_t$$

$$(3) F_t p_{t+1} = \beta_t p_t$$

with $0 < \lambda < 1$; $\gamma, \delta > 0$; $p_t, c_t > 0$; and $c_0 > 0$ given. First close the model under perfect foresight:

$$(4) \beta_t = \frac{p_{t+1}}{p_t}$$

Substituting (4) in (1) and rearranging shows that $\beta_0 < \lambda^{-1}$ is required to be compatible with $c_0 > 0$. Substitute (3) in (1) to obtain

$$(A.1) p_t = \lambda \beta_t p_t + \gamma c_t$$

$$(A.2) p_{t-1} = \lambda \beta_{t-1} p_{t-1} + \gamma c_{t-1}$$

Rearranging (A.2),

$$(A.3) c_{t-1} = \gamma^{-1} [p_{t-1} - \lambda \beta_{t-1} p_{t-1}]$$

Substituting (2) into (A.1) gives

$$(A.4) p_t = \lambda \beta_t p_t + \gamma [c_{t-1} + \delta p_t]$$

Now substitute (A.3) into (A.4):

$$(A.5) p_t = \lambda \beta_t p_t + \gamma [\gamma^{-1} (p_{t-1} - \lambda \beta_{t-1} p_{t-1}) + \delta p_t]$$

$$(A.6) \lambda \beta_t = 1 - \beta_{t-1}^{-1} + \lambda - \gamma \delta$$

or, iterating forward and rearranging,

$$(5) \beta_{t+1} = (1 + \lambda^{-1} - \gamma \delta \lambda^{-1}) - \lambda^{-1} \beta_t^{-1}$$

as in the text. A rational expectations equilibrium is defined as a sequence $\{\beta_t\}_{t=1}^{\infty}$ that solves (5). Proceed to analyze (5) as follows:

$$(A.7) \frac{d\beta_{t+1}}{d\beta_t} = \lambda^{-1} \beta_t^{-2} > 0$$

$$(A.8) \frac{d^2\beta_{t+1}}{d\beta_t^2} = -2\lambda^{-1} \beta_t^{-3} < 0$$

The roots of (5) are found by setting $\beta_{t+1} = \beta_t$ and applying the quadratic formula:

$$(A.9) \beta_t^* \beta_n^* =$$

$$\frac{(1 + \lambda^{-1} - \gamma \delta \lambda^{-1}) \pm \sqrt{(1 + \lambda^{-1} - \gamma \delta \lambda^{-1})^2 - 4\lambda^{-1}}}{2}$$

The level of the real deficit $\delta > 0$ must be chosen such that these roots are real. This requires

$$(A.10) \delta < \lambda \gamma^{-1} [\lambda^{-1} + 1 - (4\lambda^{-1})^{1/2}] = \delta_{\max}$$

The roots also satisfy

$$(A.11) 1 < \beta_t^* \beta_n^* < \lambda^{-1}$$

These facts provide the basis for the qualitative graph of figure 1.

To analyze the dynamics of the model using the graph, consider some initial condition β_0 on the horizontal axis. Find the value β_1 by tracing up from β_0 to the plotted function. The value of β_1 now serves as the input for the next period. To transfer the β_1 value to the horizontal axis, trace horizontally from the plotted function to the 45 degree line, and from the 45 degree line down to the horizontal axis. Now repeat the procedure as though β_1 is the initial condition.

Next, solve the model under least squares learning. Use

$$(4') \beta_t = \left[\sum_{s=1}^{t-1} p_s^2 \right]^{-1} \left[\sum_{s=1}^{t-1} p_s p_{s-1} \right]$$

In order to get to equation 5' in the text, (4') must be written in recursive form. To do this, define temporarily two vectors

$$(A.12) P_{t-1} = [p_0, \dots, p_{t-1}]'$$

$$(A.13) P_{t-2} = [p_1, \dots, p_{t-2}]'$$

Then

$$(A.14) \beta_t = [P'_{t-2} P_{t-2}]^{-1} P'_{t-2} P_{t-1}$$

and

$$(A.15) \beta_{t-1} = [P'_{t-3} P_{t-3}]^{-1} P'_{t-3} P_{t-2}$$

where now

$$(A.16) P_{t-2} = [p_1, \dots, p_{t-2}]'$$

$$(A.17) P_{t-3} = [p_0, \dots, p_{t-3}]'$$

The additional information in β_t is the observation p_{t-1} . The relationship between β_t and β_{t-1} is given by a well-known recursive least squares formula¹ which can be applied as follows:

$$(A.18) \beta_t = \beta_{t-1} + [P'_{t-3}P_{t-3}]^{-1} p_{t-2}(p_{t-1} - p_{t-2}\beta_{t-1})f_t^{-1}$$

where

$$(A.19) f_t = 1 + p_{t-2} [P'_{t-3}P_{t-3}]^{-1} p_{t-2}$$

is a scalar. Since

$$(A.20) [P'_{t-3}P_{t-3}]^{-1} = \left[\sum_{s=1}^{t-2} p_{s-1}^2 \right]^{-1},$$

the scalar f_t can be written as

$$(A.21) f_t = \left[\sum_{s=1}^{t-2} p_{s-1}^2 \right]^{-1} \left[\sum_{s=1}^{t-2} p_{s-1}^2 + p_{t-2}^2 \right].$$

Substituting into (A.18) yields

$$(A.22) \beta_t = \beta_{t-1} + \frac{p_{t-2}}{\sum_{s=1}^{t-1} p_{s-1}^2} [p_{t-1} - p_{t-2}\beta_{t-1}].$$

To obtain equation (5'), use (1), (2) and (3) to find

$$(A.23) p_{t-1} = G(\beta_{t-1}) \left[\frac{1 - \lambda\beta_{t-2}}{1 - \lambda\beta_{t-1}} \right] p_{t-2}$$

Substituting (A.23) into (A.22) gives the equation in the text:

$$(5') \beta_t = \beta_{t-1} + \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} \left[\frac{(1 - \lambda\beta_{t-2})}{(1 - \lambda\beta_{t-1})} G(\beta_{t-1}) - \beta_{t-1} \right].$$

A set of positive sequences $\{\beta_t, c_t, p_t\}_{t=0}^{\infty}$ satisfying (1)-(3) and (4') is an equilibrium under least squares learning. To approximate the dynamics under least squares learning, consider the simpler but closely related difference equation:

$$(6) \beta_t = G(\beta_{t-1}) = \frac{1 - \lambda\beta_{t-1}}{1 - \lambda\beta_{t-1} - \gamma\delta}.$$

The derivatives are given by

$$(A.24) \frac{dG}{d\beta_{t-1}} = \lambda\gamma\delta(1 - \lambda\beta_{t-1} - \gamma\delta)^{-2} > 0$$

$$(A.25) \frac{d^2G}{d\beta_{t-1}^2} = 2\lambda^2\gamma\delta(1 - \lambda\beta_{t-1} - \gamma\delta)^{-3} > 0$$

provided $\beta < \lambda^{-1}(1 - \gamma\delta)$. This provides the information for the second qualitative graph, given by figure 2.

¹See Harvey (1981), p. 54.