Econometrica, Vol. 57, No. 4 (July, 1989), 889–910

LEARNING RATIONAL EXPECTATIONS UNDER COMPUTABILITY CONSTRAINTS

BY STEPHEN E. SPEAR[1]

In this paper we consider how boundedly rational agents learn rational expectations. The assumption that agents are boundedly rational is made operational by imposing computability constraints on the economy: all equilibrium price functions or forecasts of future equilibrium prices are required to be computable. Computable functions are defined, as in the computer science literature, as functions whose values can be calculated using some finite algorithm.

The paper examines two learning environments. In the first, agents have perfect information about the state of nature. In this case, the theory of machine inference can be applied to show that there is a broad class of computable economies whose rational expectations equilibria can be learned by inductive inference.

In the second environment, agents do not have perfect information about the state of nature. In this case, a version of Gödel's incompleteness theorem applicable to the theory of computable functions yields the conclusion that rational expectations equilibria cannot be learned.

KEYWORDS: Rational expectations, inductive inference, recursive function, bounded rationality, effectively computable.

1. INTRODUCTION

IN A RATIONAL EXPECTATIONS EQUILIBRIUM (REE), agents correctly forecast future payoff relevant variables conditional on current information. Out of equilibrium, they make systematic errors in forecasting. If agents recognize error, then learning should occur and agents will modify their behavior until the economy attains a REE. This fundamental argument makes a compelling case in favor of REE as the appropriate notion of equilibrium for dynamic economies. Of course the ability of agents to learn by recognizing error is crucial to this argument. In this paper we examine the possibilities for learning REE when the environment facing agents is sufficiently complex that recognizing errors is not a simple task.

In taking this view, we are concerned with the question of how boundedly rational agents learn REE, and will focus, in particular, on the question of whether learning can be procedural, in the sense that agents have algorithms—procedures defined by finite sets of rules—for systematically comparing forecasts with realizations.

Procedural approaches to defining rationality are not new, of course. Simon's (1976) work has been seminal in stressing procedures and algorithms in defining rationality. This theme is echoed in Arrow's (1987) discussion of the information and data processing demands imposed on agents in models which deviate from the static, perfectly competitive, complete markets world of textbook economics.

[1] I am grateful to Larry Blume, Margaret Bray, David Kreps, Bart Lipman, Sanjay Srivastava, Michael Woodford, and two anonymous referees for many useful comments and suggestions on this work.

Similar ideas can be found in Lucas' (1987) characterization of economic agents as collections of decision rules which are systematically reviewed and revised in light of new information.

Our specific approach to modeling boundedly rational agents is one of imposing computability constraints on the model. By computability constraints, we mean that all functions or functionals involved—forecasts of future prices, temporary equilibrium (T.E.) prices, or mappings from forecasts to T.E. price functions—must be machine computable. Given that the REE of a dynamic economy can be represented as the fixed point of a functional equation taking forecast functions to current price functions, computability constraints provide a natural and reasonably simple way of imposing bounded rationality. These assumptions will be formalized and discussed more fully below.

The main analytic tools we apply are drawn from formal computer science. In particular, we make use of the theory of inductive inference and of recursive function theory (see references in the following section). The paper itself is less concerned with examining learning mechanisms and convergence results than with the basic issue of whether, in the presence of computability constraints, learning rational expectations is even possible.

Our interest in algorithmic learning is motivated by two observations. First, most of the approaches to boundedly rational learning in the literature are, in fact, procedural. Agents in these models use well-known procedures from statistical decision theory (all of which can be implemented algorithmically) to learn about the parameters of the world they inhabit, and to refine their guesses about the future. A second, more fundamental observation, is based on the famous undecidability theorems of Kurt Gödel in formal logic. Loosely speaking, these results state that there exist propositions which can be logically formulated whose truth or falsity cannot be logically determined. In the context of learning rational expectations, we are concerned, then, with the possibility that the problem of determining the truth of the statement "my forecast is an REE forecast" may be undecidable. In this light, the notion of procedural rationality simply means that agents cannot decide the undecidable. In terms of the computability formalism, the bounded rationality assumption requires that we examine the question of whether the learning process ever requires that agents compute noncomputable functions.

One can, of course, object to this characterization on the grounds that agents only act "as if" they were solving decision problems beyond the reach of mere mathematicians, in much the same way that a trained dog only acts as if it were solving a differential equation while chasing down a fly ball. While this argument has merit, it is not without its own ambiguities. The dog's brain does, in fact, solve a differential equation, at least to a close enough approximation to allow the dog to catch the ball. Qualitatively, this is no different from having a computer solve the relevant equation numerically and display the result pictorially on its terminal screen. Of course, in the case of the computer, the programmer knows how the computer is solving the problem, since he wrote the algorithm. The question of how the dog's brain does it is not well understood yet.

Ultimately, the "as if" question is one of whether human mental activity is machine simulable or not.

Even if one believes that the human mind is not machine simulable (and that it can do things computers never will), the consideration of algorithmic learning procedures serves a useful purpose. In particular, it focuses attention on those aspects of the problem of economic learning which may require metamathematical capabilities.

On the other side of this issue, if one views the various constructs of economic theory as attempts to model realistically (rather than as if) the actual decision processes of agents, then the requirement that agents not be capable of deciding logically undecidable propositions is actually quite weak, since it leaves available the vast array of results in logic and mathematics which are, in fact, decidable. From this perspective, inquiry into the algorithmic foundations of economic decision-making provides an outer measure of the extent to which current economic theory makes unreasonable demands on the capabilities of real agents.

The importance of the issue of error recognition for learning leads us to focus on inductive learning procedures. Inductive learning is characterized in psychology as the process whereby new information from the environment is organized into cognitive patterns which then serve to either reinforce an individual's cognitive models of the environment, or to alter these models to conform with newly recognized patterns. In the context of economic learning, inductive procedures seem to correspond quite closely to Lucas' notion that agents continuously review and revise their decision rules in response to signals from their environment. We will be particularly interested in whether there are algorithmic induction procedures which allow agents to recognize when their forecasts are wrong.

In the literature on learning rational expectations, two main approaches to the issue of learning have been pursued. The first has been characterized by Blume, Bray, and Easly (1982) in their survey as Bayesian or rational learning. In this mode of learning, agents may hold different forecasts due to differing information structures. The resulting equilibrium corresponding to these forecasts is then required to be a REE, from which agents may attempt to infer other agents' information. The process of rational learning is then one of moving to a REE (by way of a sequence of REE) in which nothing further can be inferred about other agents' information (including possibly their expectations). Analysis of this type of learning includes work by Arrow and Green (1973), Blume and Easly (1981), Bray and Kreps (1987), Cyert and DeGroot (1974), Feldman (1986), and Townsend (1982, 1983).

The second type of learning model has been characterized as "boundedly rational" learning. In these models, agents update their forecasts systematically based on observations of state variables determined by the economy. Agents learn to be rational if the sequence of updated forecasts converges to a REE. During the updating process, however, none of the forecasts need be rational. Within this class of learning models, there is an additional dichotomy between models in which agents hold a fixed forecast and collect sequences of observations of state variables, and models in which agents update their forecasts after every new observation. We will refer to the first type of model as a two-stage learning model, and to the second type as an incremental learning model. Analysis of boundedly rational learning includes work by Blume and Easly (1982), Blanchard (1976), Bray (1982), Brock (1972), Cyert and DeGroot (1974), DeCanio (1979), and Marcet and Sargent (1986, 1987), and Woodford (1987). Viewed in the context of the existing literature on learning, this paper is clearly about boundedly rational learning.

The main results developed are as follows. In two-stage learning models where agents have full information about the state of the economy, the theory of inductive inference can be applied to show that it is possible for agents to learn the rational expectations equilibrium. When agents have incomplete information (i.e. an incomplete signal about the current state), however, the result breaks down since agents are required to infer not a function, but a correspondence. This inference problem is undecidable since it involves the determination of a nontrivial set of admissible (recursive) functions.

We then examine models of incremental learning in which agents use a specific (but fixed) updating procedure to generate new forecasts given old forecasts and current information. For this analysis, it is assumed that there is a *model consistent* equilibrium (in the sense that agents' forecasting procedures converge to a forecast which is never controverted by the data). For these models, we show that the problem of determining whether a given updating procedure in fact yields a REE (for a fixed economy) is undecidable. Hence, while agents may hit upon a procedure which does yield a REE, there is no computationally feasible way of *choosing* a procedure which will yield rational expectations.

The paper is structured as follows. In Section 2, we present the basic economic model and our fundamental computability assumptions. We also briefly discuss aspects of the theory of recursive functions which will be useful in developing our results. In Section 3, we consider inductive learning. This section contains a brief discussion of the theory of inductive inference as well. In Section 4, we consider incremental learning models. Section 5 contains concluding observations. The main technical material used in the paper is presented in the Appendix, along with appropriate citations to the technical literature.

2. THE MODEL

In modeling inductive learning of rational expectations, the key feature of the process is the idea that when agents make forecasts of future prices, the temporary equilibrium of the economy determines a mapping from price functions to price functions whose fixed points are the rational expectations equilibria of the economy. Agents attempt to learn this mapping and its fixed points by observing the temporary equilibria which result and updating their forecasts of future prices using the observed data. This formulation is standard in the literature on economic dynamics and can be found in a variety of models.

For specificity, we will consider a stylized version of a simple overlapping generations models in which agents live for two periods, though the results developed here apply more generally to any model in which the REE can be determined as the fixed points of a mapping taking forecasts into temporary equilibrium prices. Time is discrete, with periods denoted $t = 1, 2, \ldots$. The economy then consists of a collection of agents (indexed by $h \in H$, where H may be finite or infinite) who trade a finite number of commodities in each period on a sequence of period-by-period spot markets. Agents are characterized by their endowments, which are stochastic, and by their demand functions. Demand functions depend on current and lagged prices, endowments, and on each agent's forecast of future prices when young. Since our interest is not in questions of existence or uniqueness of the REE, we will forego detailed (and unneeded) specifications of these characteristics.

To characterize the relationship between agents' forecasts of future prices and the resulting temporary equilibrium generated by these forecasts, let S denote a set of state variables for the economy. Elements of S may include any finite number of exogenous or lagged endogenous variables. Agents forecast prices as *functions* ϕ: $S \to P$, where P denotes the set of spot market prices for the economy. Let Φ denote a set of admissible forecast functions. We then make the following assumption.

ASSUMPTION 2.1: *The economy maps admissible forecasts $\phi_0 \in \Phi$ into T.E. price functions $\phi_1 \in \Phi$. This mapping, denoted $\hat{g}$: $\Phi \to \Phi$, has a fixed point.*

Two points about this assumption are worth noting. First, while the assumption is restrictive in precluding the specification of forecasts and T.E. as conditional probability measures or correspondences, it follows closely the assumptions typically made in models of learning. Secondly, the assumption on the mapping $\hat{g}$ can be easily modified to admit the possibility that different agents make different forecasts of the future. At a REE, agents necessarily forecast with the same price function (although they may form expectations differently if there are informational asymmetries present). To simplify the exposition, we assume that all agents forecast using the same forecasting function.

We turn next to our rationality assumptions. Given our focus on procedural rationality, we define admissible price functions as those which are computable. While we will discuss this assumption in detail below, the essential reason for imposing the restriction that admissible price functions be computable is the fact that computable functions are precisely those which can be characterized by finite algorithms. To formalize these restrictions as assumptions on the model, however, we will need to digress briefly and discuss some of the central ideas of the theory of computable functions.

As is standard in the literature on computability, we characterize computation in terms of Turing machines. For a formal description of such machines, see, e.g., Hopcroft and Ullman (1979). Less formally, a Turing machine is a computer which operates in discrete time. At any point in time, the computer is in one of a *finite* number of internal states. A read-write head scans letters drawn from a finite alphabet and recorded in a potentially infinite memory (the *tape*). A pair

(q, a) consisting of an internal state and a letter read from the machine's tape determines a triple (q', a', m) consisting of a new state, a new letter recorded over the existing letter (or possibly on a different tape), and a "move" of the read-write head (left or right). The transition from state to state is determined by the machine's finite control, which gives the transition rule for changing the machine's internal state as a function of the existing state and the letter being scanned. The finite control also determines what output is produced. A Turing machine begins a computation by scanning the leftmost letter of a given input string written on its tape while in a given initial internal state. Transitions then take place according to the instructions in the finite control. The computation *halts* if after a finite number of steps the machine enters a designated final state. Otherwise, the computation *loops*.

The key feature of Turing machines is that they represent the general model of computation: every finite algorithm can be realized as a Turing machine. Finite algorithms, in turn, are important because they determine completely the set of feasible computations that a human being can undertake. This is the content of the so-called Church-Turing thesis in formal logic (see Cutland (1980) for details). The Church-Turing thesis can be equivalently stated as the principle of noncomputability: if a function cannot be calculated by a Turing machine, it cannot be calculated by a human being.

Because Turing machines operate on strings of letters (*words*) with letters drawn from a finite alphabet, the functions calculated by such machines have countable domains and ranges. Hence, we may view all such functions as mappings from $N \to N$, where N denotes the set of natural numbers. Given a specific Turing machine, we define the function calculated by the machine as that f: $N \to N$ such that when the machine begins computation with input $j \in N$, it either produces output $f(j)$ or fails to halt. If the machine produces an output, we write $f(j)\downarrow$. Otherwise, we write $f(j)\uparrow$ and say that the function is undefined at $j \in N$. The set of functions computed by some Turing machine is called the set of *partial recursive* functions, where the adjective "partial" refers to the fact that the function may not be defined for certain values $j \in N$. A function which is defined for all $j \in N$ is said to be *total recursive*.

A key result in the theory of computation is that not every function f: $N \to N$ can be computed by a Turing machine. In particular, this means that there are functions defined on N for which there is no algorithm that will calculate the values of the functions. To prove this, we need to introduce the concept of *coding* or *indexing* the set of Turing machines. First, every Turing machine can be completely described by listing the (finite) symbols of the alphabet on which it operates, the (finite) internal states of the machine, and the (finite) set of transition rules of its control. By coding each element of this finite set suitably (i.e. assigning each element a unique integer), every Turing machine TM can be described by a string $w(TM)$ of finite length. Such strings can be further collapsed using the fact that there exist mappings from $\bigcup_k N^k$ (where N^k is the k-fold product of N) into N which are one-to-one. One such mapping, due to Gödel, can be realized by assigning to each string of integers $i = [i_1, \ldots, i_n]$ the integer

$$j(i) = p_i^{i_1} \dots p_n^{i_n}$$

where p_i is the ith prime.

Since the set of recursive functions (partial and total) is countable, we can identify each such function with the index number $i \in N$ of the Turing machine that calculates the function. Given such an indexing, it follows immediately that the set of Turing machines is at most denumerably infinite. The set of all functions from N to N, however, is uncountable. Hence, not every function is recursive. Equivalently, there exist functions for which there are no algorithms that will systematically calculate the values of the function in question.

Having introduced the key elements of the theory of recursive functions, we can now proceed to state and discuss the computability assumptions we impose on the economic model.

ASSUMPTION 2.2: *The set S is countable. The set Φ consists of total recursive functions on S.*

The assumption that S is discrete departs significantly from the continuity assumptions typically made in economic analysis (although see Mas-Colell (1975) and Laitner (1985)). Continuity assumptions, while extremely useful for the mathematical results they make possible, are nonetheless only idealizations of the actual discrete computations performed in real economies peopled with real agents who always round off real numbers. Hence, the assumption that states are discrete hardly needs justification. What does require further exploration (though we do not undertake it here) are the issues of existence and uniqueness of equilibria associated with economies in which variables take on discrete values. Some work on this issue is done in Laitner (1985) in the context of an overlapping generations economy, though the REE he obtains take the form of probability measures rather than functions. For our purposes, we will view the discreteness assumption as one of approximating (on rational numbers, for example) the real-valued results delivered by continuous economies.

The assumption that the set of admissible prices functions Φ consists of total recursive functions can be justified as follows. That Φ should be a subset of the recursive functions (total or partial) is based on the fact that forecasts *must* be computable (by the Church-Turing thesis). At any REE, therefore, the equilibrium price function must also be recursive. That nonequilibrium temporary equilibrium price functions should be computable can be defended by viewing these prices as computable approximations of the T.E. price functions associated with a given forecast. That such approximations exist (to any degree of precision) can be shown formally (though we do not undertake a proof here) by working with computable approximations to agents' excess demand functions (given the current state and each agent's forecast) and applying the generalized Sturm algorithm (see, e.g., Jacobson (1964)) to compute the zeros of the approximation to the aggregate excess demand function. Note also that the need for considering such approximations is also raised by recent results of Lewis (1985), where it is shown, in the context of the consumer's optimization problem, that *exact* optimization may be computationally impossible.

The assumption that Φ consists of total functions can be justified by again viewing elements of Φ as computable approximations of T.E. equilibrium prices, and then appealing to standard results on the existence of temporary equilibrium prices (see, for example, Grandmont and Hildenbrand (1972)). It must, of course, be noted that these results cannot be applied directly to the model because of the discreteness assumption on S. Hence, we cannot derive these properties of Φ as theorems, but must impose them as assumptions. Clearly, it would be interesting to know whether these properties can be formally derived.

Given Assumption 3.2, the fact that the set of recursive functions is countable implies that we may unambiguously determine a function $\phi \in \Phi$ by its index under some acceptable indexing of the recursive functions. It then follows that the set Φ is isomorphic to N and we may define the mapping $\hat{g}$ of Assumption 3.1 as g: $N \to N$, where the temporary equilibrium price function associated with a forecast ϕ_i is $\phi_{g[i]}$. We make the following assumptions on g.

ASSUMPTION 2.3: *The mapping g: $N \to N$ is total recursive. The mapping g has a fixed point in the sense that there exist functions ϕ_i and $\phi_{g[i]}$ such that $\phi_i = \phi_{g[i]}$.*

As with Assumption 3.2, the requirement that g be total reflects the general results on the existence of temporary equilibria (again, see Grandmont and Hildenbrand (1972)). The second part of the assumption states that there exists a rational expectations equilibrium in the sense that g has a fixed point in the space of total recursive price functions. Note that $g[i]$ need not equal i at the REE; it is possible that the algorithms used by agents to forecast prices may be different from those used by the market to generate actual prices, even though both procedures yield the same price functions.

Finally, note that if we relax Assumption 3.2 to require only that forecasts and T.E. price functions be partial recursive (i.e. for some states, forecasts, or temporary equilibrium prices may be undefined) then we can guarantee the existence of REE under Assumption 3.3 by appealing to the so-called recursion theorem (Theorem A.3 in the Appendix). This possibility has been examined in some detail by McAfee (1984), who interprets the possibility that the computation of a T.E. price for some state of nature may not halt as a case of some market(s) failing to open. Under this interpretation, the REE will specify indefinite delay in committing to a forecast for those markets which fail to open.

While it is certainly not inconceivable that there are states of nature which occur with positive probability for which markets fail to open (many risks are not insurable), it is not clear that inclusion of such states is reasonable for models designed to consider the pricing of systematic risk, and for which states involving nonpriced risks are ignored. Since these models are more typical of the kind considered in the literature on rational expectations, we impose Assumption 3.2. McAfee's results are, nevertheless, worth noting since they show that none of the restrictions imposed by consideration of recursive economies necessarily precludes the existence of REE. We turn next to the question of learning.

3. INDUCTIVE LEARNING

To begin our analysis of learning, we first introduce several ideas and results from the theory of inductive inference.

The theory of inductive inference addresses the question of characterizing when it is possible to construct algorithms that can infer which function among a class of recursive functions has generated an observed sequence of ordered pairs of numbers of the form $(j, f[j])$. To state the main results of this theory, the following notation is useful. Let

$$\Sigma = \bigcup_j \left[\bigtimes_1^j N \right].$$

Then $\sigma \in \Sigma$ is a sequence of natural numbers. Let $\sigma_n = \pi_n(\sigma)$ denote the projection of the sequence σ on its first n components. Let $\langle j, m \rangle$ denote the *number* obtained from the ordered pair by some coding of $N \times N$ into N. If a sequence σ consists of numbers of the form $\langle j, f(j) \rangle$ (i.e. coding for the ordered pairs in the graph of the function f) and all such pairs for which the function f is defined appear (possibly with repetitions) in the sequence σ, then σ is called a *text* for f. A learning function is a mapping ϕ: $\Sigma \to N$ which takes text for a function f into a conjecture about the index of the TM which computes f. We then make the following definition.

DEFINITION: Let ϕ be a learning function and σ a text for the function f.

1. ϕ is said to be defined on σ if $\phi(\sigma_n)\downarrow$ for all $n \in N$.
2. Let $i \in N$. ϕ is said to converge on σ to i if (a) ϕ is defined on σ and (b) for all but finitely many $n \in N$, $\phi(\sigma_n) = i$.
3. ϕ is said to identify σ if there is $i \in N$ such that (a) ϕ converges on σ to i, and (b) the TM of index i calculates f.

Intuitively, the learning function processes the information in the sequence σ and, for each new observation $\langle j, f(j) \rangle$, offers a conjecture about the index of the recursive function which might have generated the observed data sequence. If these conjectures converge to a fixed index i and the data were in fact generated by the function f_i, then ϕ identifies f_i.

The identification of any given recursive function is trivial; if j is the index of the function in question, then define $\phi(\sigma) = j$ for all σ. The more interesting (and less trivial) problem is that of identifying collections of recursive functions. For our purposes, the following theorems, due to Gold (1967), are most important.

THEOREM 3.1: *The set of total, primitive recursive functions is identifiable. Furthermore, identification is effective; the learning function ϕ is recursive.*

(A proof of this theorem can be found in Gold (1967), and in Blum and Blum (1975). Primitive recursive functions are defined in the Appendix.)

THEOREM 3.2: *The set of total recursive functions is not identifiable by any recursive learning function.*

(Again, formal proof of this theorem can be found in Gold (1967), or in Osherson et. al. (1986).)

A corollary to Theorem 3.2 (see Blum and Blum (1975)) states that the class of total recursive functions can be effectively identified if the text from which the identification is made is generated by primitive recursive enumeration (see Appendix for definitions). This in turn requires that there exist a primitive total recursive function p: $N \rightarrow N$ such that $p(n) = \langle j_n, f(j_n) \rangle$. Since every recursive function has a primitive recursive enumeration (see Cutland (1980)), identification by primitive recursive enumeration is always possible.

While Gold's theorem is quite powerful in asserting our ability to infer recursive functions from primitive recursive enumerations of their graphs (at least in the limit), the fact that not all recursive functions can be identified if the graph is enumerated in an arbitrary way poses problems for learning when there is no *a priori* reason to expect that nature will oblige us with a primitive recursive enumeration of the evidence. In the Appendix, we provide a brief sketch of the proof of Theorem 3.2 since the constructions in this proof illustrate why the order in which the graph of a function is enumerated is important for identification.

Having briefly outlined the theory of machine inference, we now consider the problem of learning rational expectations. In examining inductive learning of rational expectations, we consider three broad models of boundedly rational learning. The first two correspond to models of learning in which agents hold fixed expectations and collect samples of equilibrium prices. These models have been examined by Blanchard (1976), Bray (1982), Brock (1972), and DeCanio (1979). The third model deals with incremental learning in which new observations are immediately incorporated into forecasts. These models have been considered by Blume and Easly (1982), Bray (1982), Cyert and DeGroot (1974), and Marcet and Sargent (1986). In developing our results, we first consider the possibility of inductive learning of rational expectations in two-stage models with full and incomplete information. We then turn to models of incremental learning.

The first model we consider is one of two-stage learning under full information. In two-stage learning, agents hold (common) fixed forecasts and observe pairs (p_t, s_t) of prices and states generated by the economy as temporary equilibria corresponding to the fixed forecast. Agents then attempt to infer the temporary equilibrium price function corresponding to the fixed forecast. Assuming this is possible, agents can then learn the function g by varying their forecasts and learning the corresponding T.E. price function. Once agents have learned the economy, they can presumably find the REE. The assumption that agents have full information simply means that agents observe all prices and state variables at each point in time.

The second model we consider is a variation on the first in which agents learn by stages, but in this model, we assume that information is not complete. In particular, each agent observes only a signal about the current state s_t. This signal is given by a mapping $\eta_n(s)$ which is *not* one-to-one. We then ask whether agents can infer the T.E. price functions and, as in the case of full information, the mapping g.

We consider first the simplest case of two-stage inference with full information. In this model, agents hold a fixed, common forecast ϕ_i and observe a sequence of temporary equilibrium pairs (p_t, s_t). Hence, when the forecast is ϕ_i, agents observe s_t and

$$p_t = \phi_{g[i]}(s_t).$$

Applying Gold's theorems then yields the following results.

PROPOSITION 3.3: *If the T.E. price function is primitive recursive, agents can identify it. If $\phi_{g[i]}$ is not primitive recursive, identification may not be possible.*

PROOF: Apply Theorems 3.1 and 3.2. *Q.E.D.*

Several comments about this result are in order. First, it should be kept in mind that Gold's theorem states that identification occurs in the limit. Specifically, this means that there is some finite time at which the learning function has converged, but no *a priori* bound can be put on when convergence occurs. Furthermore, it is not possible for the computer to announce that convergence has occurred, since this depends on the enumeration of the data. Thus, the identification result obtained here is similar in nature to results in statistical inference which depend on the law of large numbers. In the problem of parameter estimation, for example, the law of large numbers guarantees consistent estimates and delivers confidence intervals for the estimates, but there is no way to determine how close the estimates obtained from any given sample are to the true parameter values, except in the limit. We will return to this point below.

The assumption that the temporary equilibrium price function be primitive recursive is, of course, restrictive. The assumption can be avoided if we can guarantee that the data from which agents learn is generated by primitive recursive enumeration, but this would require that agents be able to arrange the data they observe according to a primitive recursive enumeration of the T.E. function's graph. As is noted in the Appendix, this generally requires knowledge about the program generating the values of the function being inferred. Of course, if a program for the function $\phi_{g[i]}$ is known, learning is irrelevant. If such a program is not known *a priori*, then learning an arbitrary T.E. function requires that the stochastic process generating the states of nature s_t do so in a way that yields a primitive recursive enumeration of the graph of the function.

On the other hand, the assumption that the T.E. price function is primitive recursive can be defended on the grounds that the primitive recursive functions include all functions which can be *practically* computed in the sense of computational complexity.

Given these observations and caveats, it is natural to ask whether agents who can learn the T.E. price function of an economy can, by varying their forecasts, learn the functional mapping g determined by the economy. We have the following result.

PROPOSITION 3.4: *If the function g is primitive recursive, it can be identified in the limit.*

PROOF: Assuming that agents know the index of the forecast ϕ_i, the assumption that they can learn $\phi_{g[i]}$ implies that (in the limit) agents observe pairs $(i, g[i])$. Since g is primitive recursive, Gold's theorem implies that the function g can be identified. *Q.E.D.*

Note that the question of whether or not agents can learn g *in finite time* is not addressed by this result, since Gold's theorem provides only for identification in the limit. If we take the two-stage learning paradigm literally, it seems highly implausible that agents will ever be able to learn g and hence the REE since they can never know when they have, in fact, identified $\phi_{g[i]}$ and hence, when to change their conjectures. On the other hand, two-stage procedures which rely on the law of large numbers can be criticized on the same grounds if one insists on certain knowledge of the parameters being estimated. It seems at least intuitively plausible that one could modify the inductive inference paradigm to avoid this problem along the same lines as models which rely on "sampling time" assumptions. Such a modification would allow agents to terminate the production of conjectures once they appeared to have converged, but at some risk of being wrong. While a formal analysis of this is beyond the scope of this paper, it seems natural to conjecture that one would obtain reasonable "approximate" learning results from this procedure, with the risk of error going to zero as the arbitrary termination time is pushed further and further into the future. Nevertheless, it is clear that, like other two-stage learning models, the need for substantial sampling is a serious drawback of the procedure.[2]

As with Proposition 3.3, the assumption that g is primitive recursive does restrict the economy significantly. Once again, the assumption can be weakened by assuming that agents are capable of varying their forecasts in a way that generates a primitive recursive enumeration of the pairs $[i, g(i)]$. As was the case before, though, if g is not primitive recursive, such an enumeration of the graph of g may require knowledge of a program for g itself, in which case there is nothing to be learned. The assumption that g is primitive recursive can be defended (as with the T.E. price function) on the grounds that practical computability requires primitive recursion.

[2] I am indebted to an anonymous referee for pointing out this complication.

We will call an economy primitive recursive if all the relevant functions ϕ_i, $\phi_{g[i]}$, and g are primitive recursive. Taking Propositions 3.3 and 3.4 together, then, we obtain the positive result that for primitive recursive economies in which information is complete, agents may be able to inductively infer the mapping g and hence (by calculating the fixed point of g) learn the rational expectations equilibrium price function(s).

We turn next to the question of inductive learning in the two-stage model when agents do not have complete information, but instead observe only a signal $\eta_h(s)$ correlated with the current state. In this case, when agents again hold a fixed, common forecast ϕ_i, agent h observes a T.E. pair $(p_t, \eta_h[s_t])$, where

$$p_t = \phi_{g[i]}(s_t).$$

When the mapping η_h is many-to-one, the best agent h can hope to infer from observations of the temporary equilibrium is a correspondence of possible T.E. price functions consistent with his signal. Let

$$N(\bar{\eta}_h) = \{s \mid \eta_h(s) = \bar{\eta}_h\}.$$

Agent h must infer the set

$$\Psi_h(i) = \{j \mid \text{range } \phi_j|_{N(\bar{\eta}_h)} = \text{range } \phi_{g[i]}|_{N(\bar{\eta}_h)} \text{ for all } \bar{\eta}_h\}.$$

If this set can be inferred, then h can hope to learn the *correspondence* induced by g and given explicitly by $\Psi_h(i)$. Since g has a fixed point, this correspondence has a fixed point. Unfortunately, because the set $\Psi_h(i)$ is a nonempty, nontrivial set of recursive functions, we immediately obtain the following negative result.

PROPOSITION 3.5: *There is no effective procedure for determining the set* $\Psi_h(i)$.

PROOF: Apply Rice's theorem (see Appendix). *Q.E.D.*

COMMENT 1: While it might seem that inductive inference with full information is just a special case of inference with incomplete information, it should be noted that the inference results with full information do not contradict Rice's theorem since those results do not require a characterization of the *set* of indices consistent with the T.E. price function $\phi_{g[i]}$. The identification theorems only imply the existence of a learning function which in the limit correctly conjectures some index j such that $\phi_j = \phi_{g[i]}$ given the observations of independent and dependent variables. Identification fails in the case of incomplete information because there are simply too many functions in the set $\Psi_h(i)$ for any finite state Turing machine to keep track of.

COMMENT 2: While Rice's theorem implies that we cannot effectively determine membership of the set $\Psi_h(i)$, if the set were recursively enumerable there would be some hope that by systematically listing the elements of the set, the REE could be inferred. Unfortunately, the requirement that the pricing functions be total immediately implies (by the Rice-Shapiro Theorem) that if the set S is infinite, $\Psi_h(i)$ is not recursively enumerable. If we relax Assumption 3.2 and work with partial recursive price functions, then the set $\Psi_h(i)$ may be recursively enumerable. But in this case, we will lose the positive result on learning in environments of complete information, since Gold's theorem allows the identification of at best a total recursive function which extends the given partial recursive function. In the context of a stochastic economy, extending a forecast function defined on some subset of states to all possible states amounts to opening new markets for trades contingent on the added states (see the discussion at the end of Section 2). In general, one would expect this to alter the fixed points of the functional map taking forecasts into T.E. price functions. Hence, for agents to find the fixed point $\phi_{g[i]} = \phi_i$ requires that they determine the fixed point of the correspondence

$$\Xi(\phi_i) = \{\phi \mid \phi = \phi_{g[j]} \text{ and } j \text{ is an index for some restriction of the total recursive function } \phi_i\}.$$

Generating this correspondence, in turn, requires that agents know the set of total recursive functions. By the Rice-Shapiro theorem, determining this set is impossible. One alternative which avoids both of these problems is to assume that there are only finitely many states of nature. We will return to this point below.

COMMENT 3: There are results in the inductive inference literature showing that identification is possible under some circumstance when text (i.e. the observed (p, s) pairs) is incomplete in the sense that some finite number of pairs is missing. There are no results to date, however, on learning when infinitely many observations are missing.

Are there alternatives to characterizing the REE which avoid these problems? Since agents are ultimately interested only in the question of whether they are at an REE, one possibility is for them to try simply to determine whether $\phi_i = \phi_{g[i]}$ without trying to determine the function g. But this amounts to asking whether agents can compute the characteristic function of the set of fixed points of g, i.e.,

$$\{i \mid \phi_i = \phi_{g[i]}\}.$$

Again, by Rice's theorem, this function is not computable. For total recursive price functions ϕ_i, this set is not recursively enumerable, so that there is no effective way for agents to determine that they are *not* at a fixed point of g, either. Note also that the fact that this set is not recursive also implies that its complement is not recursive, so that there is no effective procedure for determining when a forecast is wrong.

One approach which avoids all of the problems discussed so far is to assume that there are only finitely many states of nature. We will return to this point below.

The negative implications of Proposition 3.5 together with the observation that in most realistic situations of economic inference, information will not be complete (or will be excessively costly to collect), suggests that we examine alternative models of learning which may be robust to incompleteness of information. This leads us to consider models of incremental learning. These models are considered in the following section.

4. INCREMENTAL LEARNING

In this section, we consider models of incremental learning, in which agents use a fixed (perhaps agent specific) procedure for updating forecasts in each period based on the current (possibly incomplete) information about prices and states. These models are intuitively appealing because they avoid the need for the "sampling time" assumptions required of the two-stage inference models.

In looking at incremental learning, we will be less concerned with the question of whether or not a given updating procedure converges, than with the question of deciding whether or not the procedure can yield a REE given that convergence occurs. This in turn leads naturally to a consideration of how agents actually choose their updating procedures and whether this process can be rationalized in any way.

For this analysis, we let

$$\left(\phi_{g[i_0]}(s), \eta(s)\right) = \xi(s, i_0)$$

where, as before, $\eta(s)$ denotes the signal agents receive about the current state, and i_0 is the index of the common initial forecast. For this analysis, we will assume that agents receive the same signal, and that agents share a common means of updating forecasts based on observations of ξ. These commonality assumptions simplify the exposition of the results to follow and, as was the case in the previous models, the assumptions can be relaxed without altering any of the results. (Indeed, by reinterpreting the indices as codings of vectors of indices, the assumption that signals are common can be dropped.)

We assume that the updating scheme is defined by a (total) recursive function f which takes forecasts and observations ξ into new forecasts, so that

$$i_1 = f(i_0, \xi[s_1, i_0])$$

where the updating process maps the index i_0 into the index i_1 based on observation ξ at s_1 for given forecasts ϕ_{i_0}. We will say that a forecast ϕ_j is consistent with the updating scheme if

$$\phi_j = \phi_{f(j, \xi[s, j])} \quad \text{for all } s.$$

Given the updating scheme f, any forecast consistent with f will be called a model consistent equilibrium. If, in addition, $\phi_j = \phi_{g[j]}$, then ϕ_j is a REE forecast.

COMMENT 1: The model developed so far can also be applied to the models analyzed by Anderson and Sonnenschein (1982, 1985) by retaining the updating framework but allowing agents to collect many observations of the state s_t (so that, in particular, $\eta(s) = s$). In the models they consider, agents update using statistical methods (OLS estimation and convolutional smoothing of empirical distributions) to incorporate information about market outcomes in the revision of forecasts. As the Anderson-Sonnenschein analysis shows, there may exist consistent equilibria (in which no agent's model is controverted) which are *not* REE. We note that in general specifications of the updating procedure, market information need not be used, as when the updating function is constant.[3]

COMMENT 2: By relaxing the requirement that the price functions ϕ be total, applications of the *s-m-n* Theorem and Recursion Theorems (see Appendix) can be used to show the existence of equilibrium forecasts which are not controverted by the data generated by the economy when agents hold these forecasts. As was the case with the two-stage model, however, the equilibrium price functions delivered by these theorems can only be guaranteed to be partial recursive, allowing the possibility that equilibrium prices may not be defined for some states which occur with positive probability. Since we have argued before that this does not seem reasonable, we will simply make the following assumption.

ASSUMPTION 4.1: *There exist consistent equilibrium price functions. These functions are total recursive.*

Given an updating scheme which leads to an agent's model not being controverted, we now consider the consistency of the model relative to the actual economy. In particular, we are interested in the question of characterizing when the model consistent equilibrium is in fact a REE, and whether it is possible for agents to *choose* an updating scheme which yields the REE as a model consistent equilibrium.

As before, we fix the function g defining the economy and ask for a characterization of the set of updating schemes which yield REE. Since updating functions are recursive, we let $f=f_i$ for some i, and identify updating functions by their indices. Let j be an index such that $\phi_j=\phi_{g[j]}$ and suppose that $\phi_n=\phi_{f_i[n,\xi]}$ for all ξ. We are then interested in the set of what might be called rational updating schemes

$$R_g=\{i\,|\,\phi_{f_i[j,\xi]}=\phi_{g[j]}\}$$

and in the question of whether there is any effective procedure (i.e. an algorithm) which will allow agents to determine the indices in this set. We have the following result.

PROPOSITION 4.2: *There is no effective procedure for determining when a given, model consistent updating scheme yields a REE, unless R_g is empty.*

[3] The fact that the assumption of model consistency tells us as much about an agent's updating rule as it does about his forecasts was also pointed out by a referee.

PROOF: By Rice's theorem, there is no effective procedure for determining membership in the set R_g (i.e., the characteristic function of the set R_g is not computable). *Q.E.D.*

COMMENT: As in the case of Proposition 3.5, the requirement that the functions involved be total immediately implies (by the Rice-Shapiro Theorem) that R_g is not recursively enumerable. If the totality assumption is relaxed, this set may be recursively enumerable, in which case it would be possible to systematically list the rational updating schemes. The recursive enumerability of this set remains an interesting open question. Note, however, that even in this case, agents must know which economy they inhabit before they can hope to become rational. Of course, to know the economy is to know the function g, so that, once again, the problem of choosing a rational updating scheme becomes one of learning how the economy transforms forecasts into temporary equilibrium outcomes.

5. CONCLUSIONS

In this paper, we have considered in some detail the question of learning rational expectations equilibria under computability constraints. The analysis presented here suggests two major conclusions about learning and REE.

When agents have perfect information, inductive learning is feasible and limit identification of the REE is possible. Furthermore, the identification paradigm provides effective procedures for realizing the REE. In the process, agents learn the economy in the sense that they learn how the economy maps forecasts into market outcomes.

When information is incomplete, however, attempts to inductively infer the T.E. correspondence become infeasible, since there are no effective procedures for computing the set of T.E. price functions corresponding to a given forecast and signal realization. When we attempt to circumvent this obstacle by considering incremental adjustment processes which yield model consistent equilibria, we find that attempts to characterize adjustment processes yielding REE are impossible. The question of when a given updating procedure yields a REE (for a given computable economy) is undecidable. Hence, if agents are required to choose an updating scheme which yields model consistent equilibria which are also REE, they are facing an undecidable problem.

While there are alternative specifications of the computability assumptions which ameliorate the decidability problems posed by incomplete information, none of these alternatives is unambiguously better than the maintained assumptions of the paper. We can, for example, relax the assumption that forecasts and T.E. prices must be total recursive. In this case, the identification of some nontrivial sets of functions may become feasible if these sets can be shown to be recursively enumerable. On the other hand, relaxing the totality assumption means that the positive result on learning with complete information may no longer be true. In addition, allowing the price functions to be partially recursive opens up the possibility that the economy may have REE in which risks which are routinely priced by observed economies cannot be priced by a model of such an economy. Indeed, by considering only states of the world involving infinitely many actuarial risks, this problem can easily arise. A third alternative is that suggested by McAfee, of arbitrarily imposing time bounds on the length of computations which agents are allowed to make. While this eliminates the possibility of equilibria involving infinite delay, the resulting equilibrium allocations will typically fail to be optimal. A fourth alternative, which eliminates all problems of computability, is to assume that there are only finitely many states of nature. While this may be a reasonable approximation of the number of risks actually priced by the economy, the assumption becomes problematic when we include lagged values of endogenous variables in our description of the state since it requires that there be only finitely many different prices at which exchange may take place. Since this is precisely the context in which many of the most interesting and important results on rational expectations equilibrium have been obtained, the question of whether one can uniformly approximate an economy with infinitely many state variables by one with finitely many is most interesting. One possibility for dealing with this issue when it occurs because of the inclusion of lagged endogenous variables would be to show a version of the so-called Shadowing Lemma for compact, stochastic dynamic systems. This lemma states that for a nonstochastic, hyperbolic dynamic system, any approximation of a given trajectory is uniformly "shadowed" by an exact trajectory. This result can be applied, for example, to maps having chaotic trajectories. For such maps, no computer simulation of a trajectory can ever be precise because unavoidable round-off errors in calculating iterations of the map get magnified by the overall expansive action of the mapping. (This is the source of the sensitive dependence on initial conditions exhibited by such maps.) Nevertheless, the Shadowing Lemma asserts the existence of an exact trajectory which approximates the computed trajectory uniformly closely (see Gukenheimer and Holmes (1983) for details). If such a result could be shown for stochastic dynamic systems, it would justify approximating price processes on finitely many ε-balls.

A second line of alternatives worth considering is the relaxation of the very strong consistency requirements imposed by the rational expectations hypothesis. One possible alternative is the model consistent equilibrium concept developed and analyzed by Anderson and Sonnenschein (1982, 1985). These models are attractive in imposing some consistency on agents' forecasts relative to the actual prices delivered by the economy, but the notion of rationality embodied in these models is procedural rather than substantive. The fact that an agent's forecast of the future must be consistent with his available information (when viewed through the lens of his own model of the world) avoids the criticism of temporary equilibrium models in which no consistency is required, without forcing agents to decide undecidable problems.

While these comments in no way exhaust the possible ways of dealing with issues of computability in economic models, they (and the results developed in this paper) are meant to suggest that the problems of decidability imposed by our fundamental equilibrium concepts are serious, and that attempts to deal with the questions of how rational (but human) beings deal with economic calculations should be taken seriously.

Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, PA 15213, U.S.A.

Manuscript received July, 1987; final revision received September, 1988.

APPENDIX

In this Appendix, we cite and discuss several results from the theory of recursive functions which are used in the text. We also present a sketch of the proof of Theorem 3.2.

1. The set of recursive functions was defined in the text as those function computed by some Turing machine. We give here an alternative definition of the recursive functions.

Consider the set of functions on N containing the constant function equal to 1, the successor function $\text{Succ}(N) = N + 1$, and the characteristic function for the relationship of equality, i.e. the function which is 1 on the graph $[x, f(x)]$ of the function $f(x) = x$ and 0 elsewhere. Next, let S denote the operation of composition of functions, and define the operation of recursion, denoted R, as follows. Given a function $f(x)$ and a function $g(x, y, z)$ defined on $N \times N \times N$, define a new function $h(x, y)$ by the conditions

$$h(x,0) = f(x)$$

and

$$h(x, y+1) = g(x, y, h[x, y]).$$

Note that while the functions g and h are defined, respectively, on $N \times N \times N$ and $N \times N$, this is mostly a matter of convenience since, by using the coding described in the text, we can convert these into functions defined on N. With this in mind, the procedure above gives a way of mapping functions defined on N into functions defined on N. In particular, we denote the function h as $R(f, g)$, and h is said to be defined from f and g by recursion.

The *primitive recursive* functions are then defined as the set of functions which is closed under the operations S and R, i.e. the set of functions having the property that, given any two functions in the set, any function obtained by composition and recursion based on the given functions is again in the set. To obtain the full class of recursive functions, we define a third operation known as *minimalization*. Given a function $f(x, y)$, define $\mu y[f(x, y) = 0]$ as the least y such that: (i) $f(x, z)$ is defined for all $z \leqslant y$, and (ii) $f(x, y) = 0$ if such a y exists. Otherwise, the expression is undefined. We can then define a new function $h(x) = \mu y[f(x, y) = 0]$. The set of recursive functions (partial and total) is then obtained as the class of functions closed under composition, recursion, and minimalization. Primitive recursive functions are a strictly proper subset of the recursive functions. The Ackermann function (see, e.g., Cutland (1980, p. 46) for details) is an example of a recursive function which is not primitive recursive.

2. The fact that not all functions can be effectively calculated is closely related to the famous undecidability results of Gödel, which state that there are propositions in formal logic which cannot be proved either true or false. The analogue of Gödel's Theorem in the formal theory of computation is the following result, known as Rice's Theorem.

THEOREM A.1: *Let F be a nonempty proper subset of the set of partial recursive functions of one variable. The characteristic function of the set*

$$S_F = \{ i \mid f_i \in F \}$$

is not recursive.

A proof of this theorem can be found, for example, in Cutland (1980), Hopcroft and Ullman (1979), or Salomaa (1985). Intuitively, the theorem states that every nontrivial property of partial recursive functions is undecidable in the sense that given a property shared by some, but not all, partial recursive functions, there is no algorithm capable of deciding which functions share the property and which do not.

While Rice's theorem imposes constraints on our ability to determine whether a given function is an element of some nontrivial set of recursive functions, it does not necessarily imply that we can have no knowledge about the elements of such a set. Indeed, it is sometimes possible to find an algorithm which will list the elements of such a set. A nontrivial set of recursive functions having the property that its elements can be listed by a Turing machine is said to be recursively (or machine) enumerable. While recursive enumerability of an infinite set allows us to determine that functions appearing in the list of elements of the set are in fact in the set, it does not allow us to determine which elements are not in the set since the listing of elements in the set necessarily takes forever. The following proposition, known as the Rice-Shapiro Theorem, provides a characterization result for recursively enumerable sets of recursive functions.

THEOREM A.2 (Rice-Shapiro): *Suppose that A is a set of recursive functions such that the set $\{j \mid \phi_j \in A\}$ is recursively enumerable. Then for any recursive function f, $f \in A$ if and only if there is a finite function $\Theta \subset f$ with $\Theta \in A$.*

Here, a function is *finite* if its domain is finite. The notation $\Theta \subset f$ means that Θ is a restriction of f. A proof of this result can be found in Cutland (1980). We will make extensive use of these two theorems below in our discussion of the decidability problems associated with learning rational expectations when information is incomplete.

In addition to Rice's theorem, the following two results are also useful.

THEOREM A.3 (Recursion Theorem): *For every total recursive function g: $N \to N$, there is a natural number n called a fixed point of g such that*

$$f_n = f_{g[n]}.$$

PROOF: Salomaa (1985), Theorem 4.6

The proof of the recursion theorem makes use of the following result.

THEOREM A.4 (Kleene's *s-m-n* Theorem): *For all positive integers m and n, there is a total recursive function s (depending on m and n) such that*

$$f_i(y_1, \dots, y_m, z_1, \dots, z_n) = f_{s[i, z_1, \dots, z_n]}(y_1, \dots, y_m)$$

where the functions f_i and f_s are recursive.

PROOF: Salomaa (1985), Theorem 4.1.

3. SKETCH OF PROOF OF THEOREM 3.2: The set of total recursive functions is not identifiable by any recursive learning function.

Suppose the result of the theorem is not true, and let ϕ identify the total recursive functions. Then ϕ must identify the subset of functions which are eventually zero. Denote this subset R_0. Given this observation, one can construct a text for a total recursive function $\bar{g}$ on which ϕ "changes its mind" infinitely often, and hence does not converge to any index for $\bar{g}$. Let σ_n be given. Construct the sequence

$$\tau = [\sigma_n, \langle n+1, 0\rangle, \dots, \langle n+j, 0\rangle]$$

for some $j > 1$. Since ϕ identifies all total recursive functions, there exists a j such that ϕ will conjecture an index for a function in R_0. Similarly, given the text τ, construct the sequence

$$\tau' = [\tau, \langle n+j+1\rangle, \dots, \langle n+j+k, 1\rangle]$$

for some $k > 1$. Then, since the set of functions in $R_1 = \{$total recursive functions which are eventually 1$\}$ is a subset of the total recursive functions, ϕ must eventually conjecture an index for some function in R_1 given the text τ'. By choosing the function $\bar{g}$ to have values in $\{0,1\}$ with $\bar{g}(j)$

generating sequences of 0's and 1's of ever increasing length (doing so in an algorithmically predictable way), one can confront the learning function ϕ with text sequences of the form τ and τ'. Hence, ϕ changes its mind infinitely often and cannot, therefore, identify $\bar{g}$.

That ϕ can identify $\bar{g}$ from a primitive recursive enumeration of the graph of $\bar{g}$ (avoiding the endless changes of conjectures) follows from the fact that such an enumeration of the graph uses the program which calculates $\bar{g}$ to search out the boundaries of the increasing strings of 0's and 1's. In effect, the evidence is stacked (literally) in favor of identification.

REFERENCES

ANDERSON, R. M., AND H. SONNENSCHEIN (1982): "On the Existence of Rational Expectations Equilibrium," *Journal of Economic Theory*, 26, 261–278.

——— (1985): "Rational Expectations Equilibrium with Econometric Models," *Review of Economic Studies*, 52, 359–369.

ARROW, K. J. (1987): "Rationality of Self and Others in an Economic System," in *Rational Choice: The Contrast Between Economics and Psychology*, ed. by R. Hogarth and M. Reder. Chicago: University of Chicago Press, pp. 201–215.

ARROW, K. J., AND J. R. GREEN (1973): "Notes on Expectations Equilibria in Bayesian Settings," IMSSS Working Paper No. 33, Stanford University.

BINMORE, K. G. (1986): "Remodeled Rational Players," mimeo, London School of Economics.

BLANCHARD, O. (1976): "The Non-transition to Rational Expectations," mimeo, MIT Department of Economics.

BLUM, L., AND M. BLUM (1975): "Toward a Mathematical Theory of Inductive Inference," *Information and Control*, 28, 125–155.

BLUME, L. E., M. M. BRAY, AND D. EASLY (1982): "Introduction to the Stability of Rational Expectations Equilibrium," *Journal of Economic Theory*, 26, 313–317.

BLUME, L. E., AND D. EASLY (1981): "Rational Expectations Equilibrium and the Efficient Markets Hypothesis," mimeo, Cornell University.

——— (1982): "Learning to be Rational," *Journal of Economic Theory*, 26, 340–351.

BRAY, M. M. (1982): "Learning, Estimation, and the Stability of Rational Expectations," *Journal of Economic Theory*, 26, 318–339.

BRAY, M. M., AND D. KREPS (1987): "Rational Learning and Rational Expectations," in *Arrow and the Ascent of Modern Economic Theory*, ed. by George R. Fiewel. London: Macmillan, pp. 597–625.

BROCK, W. A. (1972): "On Models of Expectations that Arise from Maximizing Behavior of Economic Agents Over Time," *Journal of Economic Theory*, 5, 348–376.

CUTLAND, N. J. (1980): *Computability: An Introduction to Recursive Function Theory*. Cambridge, England: Cambridge University Press.

CYERT, R. M., AND M. H. DEGROOT (1974): "Rational Expectations and Bayesian Analysis," *Journal of Political Economy*, 82, 521–536.

DECANIO, S. J. (1979): "Rational Expectations and Learning from Experience," *Quarterly Journal of Economics*, 92, 47–57.

FELDMAN, M. (1986): "An Example of Convergence to Rational Expectations with Heterogeneous Beliefs," mimeo, Cornell University.

GOLD, E. M. (1967): "Language Identification in the Limit," *Information and Control*, 10, 447–474.

GUKENHEIMER, J., AND P. HOLMES (1983): *Nonlinear Oscillations, Dynamic Systems, and Bifurcations of Vector Fields*. New York, NY: Springer-Verlag.

GRANDMONT, J. M., AND W. HILDENBRAND (1972): "Stochastic Processes of Temporary Equilibria," *Journal of Mathematical Economics*, 4, 247–277.

HOPCROFT, J. E., AND J. D. ULLMAN (1979): *Introduction to Automata Theory, Languages and Computation*. Reading, MA: Addison-Wesley.

JACOBSON, N. (1964): *Lectures in Abstract Algebra*. Princeton, NJ: Van Nostrand and Co.

LAITNER, J. (1985): "Stationary Equilibrium Transition Rules for an Overlapping Generations Model with Uncertainty," *Journal of Economic Theory*, 35, 83–108.

LEWIS, A. (1985): "On Effectively Computable Realizations of Choice Functions," *Mathematical Social Science*, Vol. 10.

LUCAS, R. E. (1987): "Adaptive Behavior and Economic Theory," in *Rational Choice: The Contrast Between Economics and Psychology*, ed. by R. Hogarth and M. Reder. Chicago: University of Chicago Press, pp. 217–242.

MARCET, A., AND T. J. SARGENT (1986): "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models," mimeo.

——— (1987): "Convergence of Least Squares Learning in Environments with Hidden State Variables and Private Information," mimeo.

MAS-COLELL, A. (1975): "A Model of Equilibrium with Differentiated Commodities," *Journal of Mathematical Economics*, 2, 263–295.

MCAFEE, R. P. (1984): "Effective Computability in Economic Decisions," mimeo.

OSHERSON, D. N., M. STOB, AND S. WEINSTEIN (1986): *Systems that Learn*. Cambridge, Mass: MIT Press.

SALOMAA, A. (1985): *Computation and Automata*. Cambridge, Eng: Cambridge University Press.

SIMON, H. (1976): "From Substantive to Procedural Rationality," in *Method and Appraisal in Economics*, ed. by S. Latis. Cambridge, Eng: Cambridge University Press.

TOWNSEND, R. M. (1982): "Equilibrium Theory with Learning and Disparate Expectations: Some Issue and Methods," in *Individual Forecasting and Aggregate Outcomes*: "*Rational Expectations*" *Examined*, ed. by R. Frydman and E. S. Phelds. Cambridge, Eng: Cambridge University Press.

——— (1983): "Forecasting the Forecasts of Others," *Journal of Political Economics*, 91, 4, 546–588.

WOODFORD, M. (1987): "Learning to Believe in Sunspots," University of Chicago, mimeo.