JOURNAL OF ECONOMIC THEORY 55, 416–427 (1991)

Notes, Comments, and Letters to the Editor

Learning in Mis-specified Models and the Possibility of Cycles

YAW NYARKO*

New York University, New York, New York 10003

Received January 7, 1990; revised April 9, 1991

I study the problem of a monopolist maximizing a sum of discounted profits facing a linear demand curve whose slope and intercept are unknown. I show that if the monopolist has a mis-specified model, i.e., if the true slope and intercept lie outside of the support of the monopolist's prior beliefs, then actions and beliefs may cycle on every sample path. This behavior is shown to be robust to perturbations in the prior, true parameter, and actions. Such behavior is not possible if the agent's model is correctly specified; instead actions and beliefs necessarily converge. *Journal of Economic Literature* Classification Numbers: 026, 211. © 1991 Academic Press, Inc.

I. INTRODUCTION

Bayesian learning models have been subjected to the criticism that such models assume the agents know too much when learning. Bayesian learning models assume that agents with imperfect information form a prior belief—i.e., a probability distribution—over the set of possible parameter vectors which includes the true parameter vector. The objection raised is this: What if the true parameter lies outside of the set of possible parameters entertained by the agent? The real world, it is claimed by the critics, is so complex that the agent will believe (or behave as if it is believed) that the true parameter vector lies in some small restricted set which will in general not contain the true parameter vector.

In this paper I study the problem of a profit-maximizing monopolist facing a linear demand curve. The monopolist does not know the slope and intercept of the demand curve it faces. The monopolist chooses actions (prices) taking into account both the current period profit and future information value resulting from such an action; hence, the monopolist learns

* I thank Professors Jess Benhabib and Andrew Schotter for useful comments. I also thank an anonymous referee and an associate editor for very helpful remarks. I am grateful to the C.V. Starr Center at New York University for its support.

0022-0531/91 $3.00

"actively." I show that when the monopolist has a mis-specified model (i.e., the "true" parameters of the demand curve lie outside of the support of the monopolist's prior beliefs) then the monopolist's beliefs and actions may cycle ad infinitum. Such behavior is not possible if the monopolist's model is correctly specified.

To isolate the effect of learning, the problem I study is one for which, given any fixed value of the unknown parameter, the agent is solving the same problem at each date. Hence, the only things that change from period to period are the agent's beliefs. This formulation allows the role of learning to be disentangled from any other intertemporal effects.

The result on cyclical behavior is therefore different from that in endogenous business cycle literature ([1] and others) which obtains cycles by critically exploiting the intertemporal links in models. Further, this cyclical behavior is obtained despite the fact that the agent is using the laws of probability, i.e., Bayes' Rule, to revise beliefs.

Most authors who have studied Bayesian problems have used correctly specified models. This includes [21, 19] among earlier papers and more recently [4, 16, 8, 7, 14]. All of these papers study the correctly specified model where the true parameter lies in the support of the agents' prior distribution.

The results obtained in that literature can be characterized by the word "convergence": Both the beliefs and the optimal action converge somewhere over time. Such convergence results have typically been proved using martingale convergence theorems which implicitly require the agent's model to be correctly specified. As clearly demonstrated in [21], the limiting beliefs may or may not be concentrated on the true parameter and the limit action need not be the "correct action," i.e., the action that would be chosen if the true parameter vector were known with certainty. However, actions and beliefs do not wander around forever, but eventually settle down, with the limiting values depending upon the particular problem (and possibly even on the sample path).

In this paper I show that a monopolist's beliefs and actions may cycle despite the fact that the monopolist is using all the laws of probability (Bayes' Rule) and is solving an infinite-horizon problem (so is "actively learning"). One may ask: If the monopolist's beliefs are cycling won't the monopolist realize this and throw away those beliefs? The answer is no. The monopolist takes actions at finite dates, $t = 1, 2, 3, \ldots$. At any finite date, even though the past history has exhibited cycles the monopolist will attribute this to some low, but positive, probability event.

One may also perhaps suggest that the monopolist "choose" at the very beginning beliefs which are diffuse and have support over the set of all possible parameter vectors. However, we cannot just fiddle around with the monopolist's beliefs, just as we cannot fiddle around with utility functions.

In a consumer optimization problem where a consumer has low propensity to save, resulting in long-run poverty, is it valid to ask the consumer to "throw away" its utility function and "choose" a different utility function? Of course the answer is no! The objective of the consumer is not long-term wealth maximization but instead maximization of expected discounted utility of consumption. Similarly, in the monopolist problem the objective is not to learn the true parameter vector per se, but instead to maximize the sum of expected discounted profits.

Utility functions and beliefs both originate from agents' preferences: From [22] prior beliefs may be constructed by asking the agent to make choices over uncertain lotteries; utility functions are constructed by asking agents to make choices over certain prizes (or consumption bundles). Beliefs and utility functions are not choice variables; agents are born with them!

Requiring agents to throw away or change beliefs becomes problematic when the beliefs of the agents represent the approximations to the real world they are willing to entertain. For example the agent may be prepared to consider only linear relationships between variables as approximations to the "truth." The "truth" may be that the variables of concern are related by some non-algebraic function. To ask that the agent have beliefs with support over all functional relationships may be asking too much of the agent.

The possibility of non-convergent behavior in economic models has been studied in [4] and also mentioned in [6]. In [4] the model of [20] is used to give an illustration of the possibility of cyclical behavior in a general equilibrium exchange economy with asymetric information where agents use incorrect models to predict the underlying state of the economy from prices. (This possibility is made more precise in the very illuminating piece by Foster and Friedman [10].) The exact source of the mis-specification in agents' models is however not stated in [4]. In contrast, in this paper the source of mis-specification in agents' models is explicit: The true model lies outside of the support of the agent's prior distribution.

[11] studies the case where the true parameter lies in an infinite dimensional set and shows that the set of true parameters and prior probabilities that leads to non-convergent behavior is in a topological sense large (in particular it is a residual set—the complement of a set of first category). [9] skillfully shows that the same results hold for the two-armed bandit problem with arms that have an outcome space equal to the set of integers. A potential problem with the notion of residual sets is of course that in a finite dimensional Euclidean space a residual set may have Lebesgue measure zero. In this paper, in contrast, the true parameter lies in a finite dimensional Euclidean space. Further, I am able to index the true parameter vector and prior by finitely many positive numbers and to show

that for such numbers in a set of positive Lebesgue measure there is non-convergent behavior.

In this paper learning is done by Bayesian updating. Many authors (e.g., [5, 12, 13, 15]) have studied models where learning is done by using ordinary least squares point estimates. Those models are inherently mis-specified since while agents are learning all agents are using the wrong statistical model. Most of these authors have focussed on conditions under which there is convergence to rational expectations, while [12] discusses the inherent impossibility of learning the rational expectations model. [24] studies a mis-specified overlapping generations model where agents choose optimal actions via a recursive Robins–Monroe scheme, and shows that agents' optimal actions may converge to a sunspot equilibrium.

How pervasive is the phenomenon of mis-specified models? Well, the standard case is where agents use linear models for reasons of analytical tractability (just as do their economist advisors?) while the true model may be non-linear. And how relevant are the conclusions obtained about these mis-specified models? Dare I suggest that the cyclical behavior of economic aggregates observed in times series data is due to governments making policy with mis-specified models?

II. The Monopolist's Decision Problem

The monopolist must decide at each date t a price, p_t, to charge for a good. The chosen price leads to sales of q_t, given by a demand curve

$$q_t = a - bp_t + \varepsilon_t \tag{2.1}$$

with $\{\varepsilon_t\}_{t=1}^{\infty}$ an i.i.d. normally distributed process with mean zero and unit variance. (We choose ε to be normally distributed for simplicity despite the fact that it leads to sales q_t which may be negative with positive probability!)

For simplicity we shall suppose that there are exactly two prices the monopolist can charge; a "high" price, $p = 10$, and a "low" price, $p = 2$. After choosing a price p_t the monopolist observes the quantity sold, q_t, and receives a profit $p_t q_t$ (hence, we normalize costs to zero). The monopolist does not observe the shocks, ε_t, but knows their distribution.

Let H be the subset of R^2 denoting the set of all possible values of the parameter vector or "model," $\theta = (a, b)$, representing the intercept and slope of the demand curve. The monopolist does not know the true model, θ, but believes it lies in some subset H_0 of H, which may or may not equal H. The beliefs of the monopolist are represented by the prior probability μ_0 with support equal to H_0. If the true parameter lies in H_0 (i.e., if $\theta \in H$) then we say that the monopolist has a correctly specified model. If the true parameter lies outside of H_0 (i.e., if $\theta \in H - H_0$) then we say that the monopolist has a *mis-specified* model.

At the beginning of any date n, the monopolist would have observed the history of prices and quantities, $h_n = \{p_t, q_t\}_{t=1}^{n-1}$; the monopolist chooses a date n price, p_n, conditional on the history of observations up to date n.

Let $f_n(h_n \mid \theta)$ be the density function of the joint distribution of the date n history, h_n, conditional upon any "model" $\theta = (a, b)$. (f_n is sometimes called the likelihood function). Since the shocks to the demand, ε_t, are normally distributed it is easy to check that if $\theta = (a, b)$,

$$\log f_n(h_n \mid \theta) = \frac{-1}{2} \sum_{t=1}^{n-1} [q_t - a + bp_t]^2$$

$$+ \text{terms independent of } (\theta, p_t, q_t). \tag{2.2}$$

The beliefs the monopolist has at the beginning of date n are represented by the posterior probability μ_{n-1} obtained in the usual way via Bayes' rule: For any subset A of the parameter space H,

$$\mu_{n-1}(A) = \int_A f_n(h_n, \theta)\, \mu_0(d\theta) \Big/ \int_H f_n(h_n, \theta)\, \mu_0(d\theta). \tag{2.3}$$

The objective of the monopolist is to choose a sequence of prices, $\{p_t\}_{t=1}^{\infty}$, with the date t price a function of the date t observed history h_t only (and independent of future unobserved variables!) to maximize the sum of discounted profits, given initial prior μ_0 and discount factor δ in (0, 1),

$$\operatorname*{Max}_{\{p_t\}_{t=1}^{\infty}} E_0 \sum_{t=1}^{\infty} \delta^{t-1} p_t q_t = V(\mu_0), \tag{2.4}$$

where E_0 is the beginning of date 1 expectation given initial prior μ_0.

It is very easy to state the above problem in stationary dynamic programming terms to prove the existence of an optimal policy. Indeed, the state variable at each date is the beginning of period beliefs, μ_t; the action is the chosen price. An application of the Blackwell [2] dynamic programming techniques (and the modifications to handle this particular problem) implies the existence of a stationary policy function, $g(\mu_{t-1})$, which provides the optimal price at each date t, p_t, as a function of the beginning of period beliefs μ_{t-1}.

Let $R(p_t, \mu_{t-1})$ be the date t expected profits of a monopolist that chooses the price p_t at date t when the beginning of period beliefs over θ is μ_{t-1}:

$$R(p_t, \mu_{t-1}) = \int_H \int_{-\infty}^{\infty} p_t(a - bp_t + \varepsilon_t)\, N(d\varepsilon_t)\, \mu(d\theta) \tag{2.5}$$

where $\Theta = (a, b)$ and $N(d\varepsilon)$ denotes integration over ε with respect to the Normal distribution. Also let $B((p, q), \mu)$ denote the posterior distribution of the monopolist at the end of any period, with beginning of period beliefs μ and observed price and quantity, (p, q), as defined for example as in (2.3). If $V(\mu)$ denotes the value function as in (2.4) then the Bellman equation for the monopolist's problem is

$$V(\mu_{t-1}) = \operatorname*{Max}_{p_t} p_t q_t + \delta E_{t-1} V(\mu_t), \tag{2.6}$$

where $\mu_t = B((p_t, q_t), \mu_t)$ and where E_{t-1} denotes expectations given beginning of date t beliefs μ_{t-1}. (For more on technical probabilistic details on the existence of optimal processes for these types of problems see [18] or [14].)

III. The Correctly Specified Model, ($\theta^* \in H_0$)

When the true parameter vector lies in the support of the monopolist's prior beliefs ($\theta^* \in H_0$), so that the monopolist has a correctly specified model, then it can be shown that both the beliefs and the chosen prices of the monopolist converge. The proof of that assertion involves noting first that the posterior distribution at date t, μ_t, is a conditional probability over θ conditional upon information up to date t. Conditional probabilities of a fixed set can be shown to form a martingale. One then invokes the martingale convergence theorem to obtain the required convergence of beliefs and hence actions. (See [14] or [18] for details.)

As mentioned in the introduction, the above convergence result says nothing about the limits: in particular, beliefs need not converge to the true parameter vector and actions need not converge to the actions which would be optimal if the true parameter vector were known.

IV. A Mis-specified Model ($\theta^* \in H - H_0$)

Let $a' = 20$, $a'' = 16$, $b' = 1$, and $b'' = 4$, and define $\theta' = (a', b')$ and $\theta'' = (a'', b'')$. Suppose that the prior probability μ_0 is uniformly distributed on the rectangle with vertices at the points θ', θ'', (a', b'') and (a'', b'). Suppose that the true parameter vector is $\theta^* = (a^*, b^*) = (28.5, 5.25)$. Observe that θ^* lies outside of the support of μ_0, so that the monopolist's model is mis-specified (see Fig. 1).

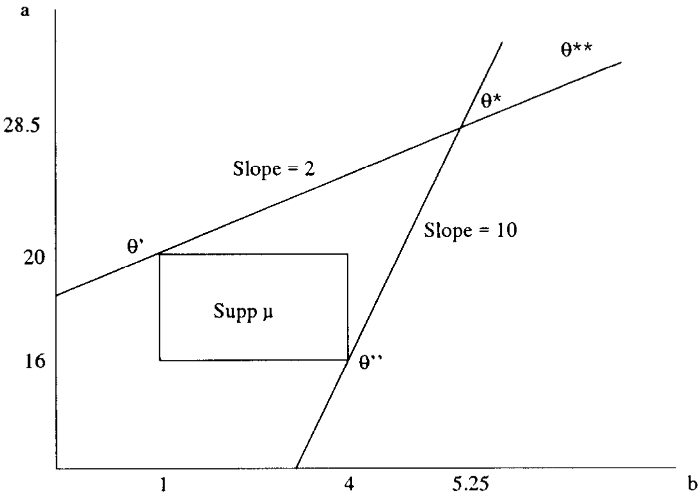

FIG. 1. A mis-specified model.

We proceed to show that on every sample path the monopolist's optimal prices oscillate ad infinitum between the "high" and the "low" prices. We begin with an informal explanation, and then state and prove the results more formally in a proposition.

Any variable with a date subscript and a bar "$^-$" on top of it will denote a time average (e.g., $\bar{p}_n = \sum_{t=1}^{n} p_t/n$). From the linear demand curve relation (2.1) we obtain

$$\bar{q}_n = a - b\bar{p}_n + \bar{\varepsilon}_n. \tag{4.1}$$

From the strong law of large numbers we know $\bar{\varepsilon}_n$ is approximately zero for large n. Suppose the monopolist has chosen the low price $p = 2$ for a large number of consecutive periods. Then from (4.1) we conclude that the monopolist would be observing sales quantities which on average equal $\bar{q}_n = a^* - b^*(2) = 28.5 - (5.25)(2) = 18$. The monopolist, upon observing sales quantities which average 18, will conclude from (2.1) or (4.1) that the true slope and intercept lie in the set $\{(a, b) \mid 18 = a - b(2)\}$. This set is precisely the line drawn in Fig. 1 with the smaller slope equal to 2 and passing through the point $\theta^* = (a^*, b^*) = (28.5, 5.25)$. The only point in the support of the monopolist's beliefs which lies on that line is the point $\theta' = (a', b') = (20, 1)$. Hence over time, if the low price $p = 2$ is chosen for many consecutive periods, the monopolist's beliefs will assign a larger and larger weight to θ' being the true parameter vector.

However, if the monopolist is sure that θ' is the true parameter vector, the low price $p=2$ results in per period expected profits of $2(a'-2b')=2(20-2)=36$, while the choice of the high price $p=10$ results in per period expected profits of $10(a'-10b')=10(20-10)=100$. In particular, the action $p=2$ is not optimal. In summary, we see that if the monopolist chooses the low price $p=2$ for a long time, the monopolist's beliefs will assign more and more weight to θ' being the true parameter vector, in which case the monopolist would switch to the high price $p=10$.

Similarly, one can check that if the monopolist chooses the high price $p=10$ for a large number of consecutive periods, the monopolist will receive data leading the monopolist to believe that the true parameter vector lies on the steeper line in Fig. 1 with slope equal to 10 passing through the point θ^*. Since the only point in the support of the monopolist's beliefs that lies on this line is the point θ'', the monopolist will assign more and more weight to θ'' being the true parameter vector. However, if the true parameter vector is θ'' the monopolist would rather choose the low price $p=2$. In particular, if the high price 10 is chosen for many periods the monopolist receives data which makes the monopolist conclude it is optimal to switch prices to the lower price $p=2$.

More formally we have the following (with the details of the proof relegated to the appendix):

PROPOSITION 4.1. *For the Monopolist Problem* (2.4) *with the parameter values as described above, if* $\{p_t\}_{t=1}^{\infty}$ *is an optimal price sequence then on each sample path both the high price and the low price are chosen infinitely often. In particular, on each sample path,* $\liminf_{t\to\infty} p_t = 2$ *and* $\limsup_{t\to\infty} p_t = 10$.

4.2. ROBUSTNESS OF CYCLES RESULT. *It should be clear for example from Fig.* 1 *that the conclusion of Proposition* 4.1 *is robust to perturbations in most directions. If the true parameter vector is not* $\theta^*=(28.5, 5.25)$ *but is some point like* θ^{**} *in Fig.* 1, *then exactly the same result holds, via a similar argument.*

For example, if the lower price $p=2$ is chosen for many periods, the monopolist will receive data leading the monopolist to believe that the true slope and intercept lie approximately on the line through θ^{**} with slope 2. Since the closest point in the support of the monopolist's prior beliefs to that line is the point θ' the monopolist will assign more and more weight to that point (this being a consequence of Theorem 4.1.1 in the appendix). The rest of the argument proceeds just as before.

The initial prior, μ_0, need not be uniformly distributed on the rectangle described in Fig. 1. It can be any distribution on that rectangle so long as the points θ' and θ'' lie in the support of that distribution. Given the prior and true parameter vector θ^*, as in Fig. 1, if the "low" price is any number less than 2 and the "high" price is any number greater than 10, then again it is easy to see that the same result holds via the same argument.

If we parameterize the monopolist problem by the vector $m = (a', a'', b', b'', a^*, b^*, \bar{p}, \underline{p}) \in R^8_+$ where $\bar{p} > \underline{p}$ are the high and low prices, resp., then it is easy to show that the set of such vectors for which the cyclical behavior of the previous section holds has strictly positive Lebesgue measure.

V. Conclusion

I have studied the problem of a profit maximizing monopolist who does not know the slope and intercept of the demand curve it faces. The only change from the standard Bayesian Learning models introduced here is that the monopolist has a mis-specified model. I have shown explicitly that this may result in the monopolist's beliefs and actions cycling on each and every sample.

This result is interesting first because such behavior is not possible in a correctly specified model, and second because the monopolist has beliefs cycling despite the fact that the model is stationary in all respects other than the learning, and the monopolist behaves "rationally" in most senses of the word: The monopolist maximizes sum of discounted profits, taking into account the future information value of current period actions (so that learning is "active"), and the monopolist learns or updates beliefs using Bayes' rule (or the laws of probability).

The monopolist problem above has therefore shown that updating via the laws of probability by "rational" economic agents is not enough to ensure the convergence of beliefs and actions. The agents must also have a correctly specified model.

VI. Appendix

In the proof of Proposition 4.1 we will require the following theorem:

Theorem 4.1.1 (Yamada). *On each sample path (or on a set with probability one) the following is true: Suppose that there is some function $w(\theta)$ (which may depend upon the sample path), such that*

$$\operatorname*{Sup}_{\theta \in H_0} \left| \left(\frac{1}{n}\right) \log f_n(h_n \mid \theta) - w(\theta) \right| \to 0, \qquad (6.1)$$

then for any open set A in H_0 which contains the asymptotic carrier $A_0 = \{\theta \in H_0 \mid w(\theta) = \operatorname{Sup}_{\theta' \in H_0} w(\theta')\}$,

$$\lim_{n \to \infty} \mu_n(A) \to 1. \tag{6.2}$$

Proof. The theorem above extends the theorem of [25, Theorem 1], stating the result for each sample path; it follows from almost trivial modifications of the proof in [25].

Proof of Proposition 4.1. Fix any sample path and suppose that p_t converges to some p_∞. Then it is easy to show that

$$\lim_{n \to \infty} \sum_{t=1}^{n} p_t/n = p_\infty \qquad \text{and} \qquad \lim_{n \to \infty} \sum_{t=1}^{n} p_t^2/n = p_\infty^2. \tag{6.3}$$

Now using the fact that $q_t = a^* - b^* p_t + \varepsilon_t$, we obtain upon simplification

$$\begin{aligned}
&\frac{1}{n} \sum_{t=1}^{n} [q_t - a + bp_t]^2 \\
&\quad = \frac{1}{n} \sum_{t=1}^{n} [(a^* - a) - (b^* - b) p_t + \varepsilon_t]^2 \\
&\quad = (a^* - a)^2 + (b^* - b)^2 \left(\sum_{t=1}^{n} p_t^2/n \right) \\
&\qquad - 2(a^* - a)(b^* - b) \left(\sum_{t=1}^{n} p_t/n \right) \\
&\qquad + \left(\sum_{t=1}^{n} \varepsilon_t^2/n \right) + 2(a^* - a) \left(\sum_{t=1}^{n} \varepsilon_t/n \right) \\
&\qquad - 2(b^* - b) \left(\sum_{t=1}^{n} p_t \varepsilon_t/n \right).
\end{aligned} \tag{6.4}$$

Now from the strong law of large numbers $\sum_{t=1}^{n} \varepsilon_t^2/n \to E\varepsilon_1^2 = 1$ and $\sum_{t=1}^{n} \varepsilon_t/n \to E\varepsilon_1 = 0$; a modification of the strong law of large numbers shows that $\sum_{t=1}^{n} p_t \varepsilon_t/n \to 0$ (see, e.g., [23, p. 476, Lemmas 2 and 3]). Hence taking limits as $n \to \infty$ in (6.4), using (6.3), and simplifying results in

$$\lim_{n \to \infty} \frac{1}{n} \sum_{t=1}^{n} [q_t - a + bp_t]^2 = [(a^* - b^* p_\infty) - (a - bp_\infty)]^2 + 1. \tag{6.5}$$

If $\theta = (a, b)$, (2.2) and (6.5) imply

$$\lim_{n \to \infty} \frac{1}{n} \log f_n(h_n \mid \theta) = \frac{-1}{2} [(a^* - b^* p_\infty) - (a - b p_\infty)]^2 + K, \tag{6.6}$$

where K is a constant independent of $\theta = (a, b)$ or p_∞.

Next, by using Theorem 4.1.1 in the previous section we conclude that if along any sample path p_t converges to p_∞, in the limit the belief process, μ_t, will have all its mass on the asymptotic carrier, which in this case is given by

$$A_0 = \{\theta \in H_0 \mid w(\theta) = \sup_{\theta' \in H_0} w(\theta')\}, \tag{6.7}$$

where (using (6.6))

$$w((a, b)) = -(\tfrac{1}{2})[(a^* - b^* p_\infty) - (a - b p_\infty)]^2$$

Consider a graph with a on the vertical axis and b on the horizontal axis. The set A_0 in (6.7) is the set of points $\theta = (a, b)$ in H_0 which are the closest to the line through $\theta^* = (a^*, b^*)$ with slope equal to p_∞. In the intuitive proof of Proposition 4.1 given in the discussion preceding that proposition, we used the fact that if p_t does converge to some p_∞ then beliefs converge to singleton set A_0 with the above geometric interpretation. We have now established a formal proof of that assertion. The rest of the proof of Proposition 4.1 therefore follows just as in the discussion preceding it.

References

1. J. Benhabib and K. Nishimura, Competitive equilibrium cycles, *J. Econ. Theory* **35** (1985), 284–306.
2. D. Blackwell, Discounted dynamic programming, *Ann. of Math. Stat.* **36** (1965), 226–235.
3. L. Blume and D. Easley, Learning to be rational, *J. Econ. Theory* **26** (1982), 340–351.
4. L. Blume and D. Easley, Rational expectations equilibrium: An alternative approach, *J. Econ. Theory* **34** (1984), 116–129.
5. M. Bray, Learning, estimation and the stability of rational expectations, *J. Econ. Theory* **26** (1982), 318–319.
6. M. Bray and D. Kreps, Rational learning and rational expectations, *in* "Arrow and the Ascent of Modern Economic Theory" (G. Feiwel, Ed.), New York Univ. Press, New York, 1987.
7. D. Easley and N. Kiefer, Controlling a stochastic process with unknown parameters, *Econometrica* **56** (1988), 1045–1064.
8. M. Feldman, An example of convergence to rational expectations, *Internat. Econ. Rev.* **28** (1988), 635–650.
9. M. Feldman, On the generic nonconvergence of Bayesian actions and beliefs, *Econ. Theory*, to appear.

10. J. Foster and M. Friedman, Learning rational expectations: Classical conditions ensure uniqueness and global stability, *Economica* **57** (1990), 439–453.
11. D. Freedman, On the asymptotic properties of Bayes estimates in the discrete case, II, *Ann. of Math. Stat.* **36** (1965), 454–456.
12. R. Frydman, Towards an understanding of market processes: Individual expectations and convergence to rational expectations, *Amer. Econ. Rev.* **72** (1982), 652–668.
13. R. Frydman, "Diversity of Information, Least Squares Learning Rules and Market Behavior," Manuscript, New York University (1987).
14. N. Kiefer and Y. Nyarko, Optimal control of an unknown linear process with learning, *Internat. Econ. Rev.* **30** (1989), 571–586.
15. A. Marcet and T. Sargent, Convergence of least squares learning in self referential linear stochastic models, *J. Econ. Theory* **48** (1989), 337–368.
16. A. McLennan, Price dispersion and incomplete learning in the long-run, *J. Econ. Dynamics and Control* **7** (1984), 331–347.
17. Y. Nyarko, On the convergence of Bayesian posterior processes in linear economic models, *J. Econ. Dynamics and Control,* to appear.
18. Y. Nyarko, "Learning in Mis-specified models and the Possibility of Cycles," C.V. Starr Center Working Paper No. R-R 90-03, New York University, 1990.
19. E. Prescott, The multiperiod control problem under uncertainty, *Econometrica* **40** (1972), 1043–1058.
20. R. Radner, Rational expectations equilibrium: Generic existence and the information revealed by prices, *Econometrica* **47** (1979), 655–678.
21. M. Rothschild, A two-armed bandit theory of market pricing, *J. Econ. Theory* **9** (1974), 185–202.
22. L. Savage, "The Foundations of Statistics," Wiley, New York, 1954.
23. J. Taylor, Asymptotic properties of multiperiod control rules in the linear regression model, *Internat. Econ. Rev.* **15** (1974), 472–484.
24. M. Woodford, Learning to believe in sunspots, *Econometrica* **58** (1990), 277–308.
25. K. Yamada, Asymptotic behavior of posterior distributions for random processes under incorrect models, *J. Math. Anal. Appl.* **56** (1976), 294–308.