Research Paper #616

RATIONAL LEARNING AND RATIONAL EXPECTATIONS

Margaret Bray* and David M. Kreps**

June 1981

Abstract. A general formulation of a rational (Bayesian) learning model with rational expectations is given. The martingale convergence theorem is used to prove that in any rational learning model, each agent's posterior assessments converge as time passes. An example is developed to show that, with sufficient "smoothness," this implies convergence of beliefs to "accurate" beliefs. The usefulness of rational learning models in investigating how agents "learn" about the relation between prices and states in a rational expectations equilibrium is discussed.

*Faculty of Economics and Politics, University of Cambridge, and Graduate School of Business, Stanford University

**Graduate School of Business, Stanford University

Acknowledgements. Conversation with Paul Milgrom proved very helpful. This research was supported in part by National Science Foundation Grant SES80-06407 to the Graduate School of Business, Stanford University, and by the Graduate School of Business, Stanford University.

RATIONAL LEARNING AND RATIONAL EXPECTATIONS

Margaret Bray* and David M. Kreps**

June 1981

A number of recent papers have addressed the question: Will economic agents be able to "learn" the relationship between prices and states of the world so as to achieve the sorts of rational expectations equilibria studied by Grossman [1977] et. al.? This is a very complex question, because as agents learn, they change their behavior, and thus they change the relationship between prices and states of the world. This learning problem is not one of observing a stationary sequence, so standard results in statistical theory (as applied by Kihlstrom and Mirman [1975], for example) do not apply.

Many of the recent studies have concerned "irrational learning." By this we mean: An agent who "understands" the economy in the sense that this agent knows the functional relationship between states of nature and prices (at the various dates) would not rationally learn in the fashion predicated by the model. Papers in this vein include Blume and Easley [1980], Bray [1980], and Radner [1980]. These papers show that with irrational learning, the evolution of agents' beliefs and behaviors can exhibit many different qualities: There can be convergence to correct beliefs; there can be convergence to incorrect beliefs; there can be divergence of beliefs; beliefs can "cycle". In contrast, there have been a few papers which deal with rational learning (often under very strong assumptions or for very special examples); among these are Arrow and Green [1973], Blume and Easley [1981], Feldman [1980], Frydman [1981], Lewis [1981], and (especially) Townsend [1978]. In these papers, one invariably

*Faculty of Economics and Politics, University of Cambridge, and Graduate School of Business, Stanford University.

**Graduate School of Business, Stanford University.

finds convergence of beliefs to correct beliefs. In this paper, we give a general formulation of rational learning (following Townsend), we show that rational learning must entail convergence of beliefs, and we argue (by example) that convergence of beliefs to correct beliefs will follow if the model is sufficiently regular.

The paper is organized as follows: In section 1, we give a genral formulation of a model with rational learning. Then we develop a special case of this formulation, for economies as in Grossman [1977] (among many others), where there is a sequence of disconnnected, "i.i.d." periods. In section 2, a simple example in the style of Grossman is given. In this example, there is one informed and one uninformed agent, with the added complication that the uninformed agent is unsure (at the outset) how risk averse the informed agent is. We show how an equilibrium with rational learning can be defined in this case (although we do not obtain a closed form specification of the equilibrium). This should be compared with the model of Bray [1980] in particular: Bray gives an "irrational" learning model for precisely this economy. Section 3 contains the main result of this paper: By simple application of the martingale convergence theorem, we see that "convergence of beliefs" <u>must</u> take place in any model with rational learning. We sharpen this result in section 4, by showing that, for the example of section 2, this implies that the uninformed agent will (almost surely) learn how risk averse the informed agent is, yielding (in the limit) the usual stationary rational expectations equilibrium for this economy. We also make some informal remarks concerning the extent to which this sharpening can be expected to extend to other particular models. In section 5, we compare our results with the very interesting example of Blume and Easley [1980], wherein (with positive probability) there can be convergence of beliefs to an incorrect model. We see how the "irrationality" present in the Blume and Easley model makes this possible, and why this could never happen if there is rational learning. We close with some general discussion: We argue that rational learning models have limited value for economic analysis, as they presume much too much on the part of agents. But they can teach us some useful things about the amount of "irrationality" that one ought to permit in an irrational learning model.

1. Formulation.

We imagine an economy which meets at a sequence of dates $t = 0,1,\ldots$ At each date t, certain markets (spot and/or future) open, in which agents trade. Agents begin with some information about the economy, and they obtain further information as the economy evolves. In particular, agents always observe the equilibrium prices at which they trade, and they condition current trades on the information contained in current equilibrium prices.

Formally, there is an underlying measure space (Ω,F). Agents are indexed by $n = 1,\ldots,N$, and each agent begins with some prior probability assessment P^n on (Ω,F). At date t in state ω, a set of markets $M_t(\omega)$ opens up, and agents trade to equilibrium in this market, with the equilibrium price in market m beging denoted $p_t(m,\omega)$. We write p_t for the random vector $\{p_t(m,\omega);\ \omega \in \Omega,\ m \in M_t(\omega)\}$; note that at this level of generality, the number of components of p_t may be random.

Each agent n, at each date t, formulates his demand based on information that this agent possesses. There are two types of information. Agent n comes endowed with certain prive information; we let G_t^n denote the σ-field generated by n's private information at date t (which we assume is nondecreasing). We assume that M_t is G_t^n measurable for all t and n;

each agent knows which markets are open. (In the spirit of Radner [1968], we could forbid certain agents to trade in markets where they lack the information necessary to verify or consummate the trades, but it is easier notationally to assume that markets are open only for opportunities in which all agents can participate.) In addition, agent n observes certain random variables generated by his own and other agents' behavior. To keep matters simple, we will assume here that each agent learns equilibrium prices and only those (in addition to the information in G_t^n); we shall write $H_t^n(p)$ for the σ-field generated by G_t^n and $\{p_0, p_1, \ldots, p_t\}$.

An equilibrium (of plans, prices and price expectations) is an array $(p,x) = \{p_t(m,\omega),\ x_t^n(m,\omega);\ t = 0,1,\ldots,\ \omega \in \Omega,\ \ m \in M_t(\omega),\ \ n = 1,\ldots,N\}$ where each $p_t(m,\omega) \in [0,\infty)$, each $x_t^n(m,\omega) \in R$, any requisite measurability conditions are met (for example, p_t must be $\bigvee_n G_t^n$ - measurable),

$$(1) \qquad \sum_{n=1}^{N} x_t^n(m,\omega) = 0 \quad \text{for each } t,\ m \text{ and } \omega,$$

and $\{x_t^n(m,\omega);\ t = 0,1,\ldots,\ \omega \in \Omega,\ m \in M_t(\omega)\}$ is an optimal net trade for agent n among all feasible net trades for this agent, given that prices follow p, and given that the agent possesses information $H_t^n(p)$ at date t.

We shall not attempt to be more precise about this at such a high level of generality. Instead, we will develop a much more concrete example, in the spirit of Grossman [1977]. Imagine that this economy consists of a sequence of economically disjoint periods. That is, at each date t, agents start afresh, with no wealth or preference carryover from previous periods. There is some uncertainty at each date t, given by a random element ϕ_t from some space Φ -- we assume that $\{\phi_t\}$ forms an i.i.d. sequence. We can imagine ϕ_t representing such things as agents' endowments, performance of securities, size of harvests, etc. A fixed set of spot markets opens in each period. Agent n has a von Neumann-Morgenstern utility function for each period t, $u^n(x^n_t,\phi_t)$, where x^t_n is the agent's net trade in the period t spot markets. Agent n receives some private information about $(\phi_0,\ldots,\phi_t)$ between the dates t - 1 and t (before date t trading takes place), which we represent by a function I^n_t defined on the space Φ^{t+1}. That is, agent n learns the realization of $I^n_t(\phi_0,\ldots,\phi_t)$ prior to date t trading.

So far we have precisely the very stationary and disconnected sequence economy of Grossman [1977]. To this we add some further uncertainty about the "parameters" of the economy. Formally, there is a random variable θ that is independent of the ϕ_t and that enters agents' utility functions; the utility functions have the form $u^n(x^n_t,\phi_t;\theta)$. Think of θ as a specification of agents' coefficients of risk aversion or of their endowments. Or, with minor modifications in our notation, one can think of θ as determining the number of agents in the economy. We assume that agent n is endowed at the outset with information $J^n(\theta)$ about θ, and that he obtains no further information except what he can glean from equilibrium prices. (This is easily relaxed.)

In terms of our general formulation, in this example $\Omega = \Theta \times \Phi^\infty$, and G^n_t is the σ-field generated by $(J^n,I^n_0,\ldots,I^n_t)$. For simplicity, we assume that agents share prior assessment P over Ω.

An equilibrium in this context is a sequence of prices $p_0(\phi_0,\theta)$, $p_1(\phi_0,\phi_1,\theta)$, ... and net trades $(x^n_0(\phi_0,\theta))^N_{n=1}$, $(x^n_1(\phi_0,\phi_1,\theta))^N_{n=1},\ldots$ such that

(2) $$\sum_{n=1}^{N} x^n_t = 0 \quad \text{P-a.s.} \quad \text{for each } t = 0,1,\ldots,$$

(3) $$x^n_t(\phi_0,\phi_1,\ldots,\phi_t,\theta) \in \underset{x:p_t x\le 0}{\arg\max}\, E(u^n(x,\phi_t;\theta)\,|\,J^n,I^n_0,\ldots,I^n_t,p_0,\ldots,p_t),$$

(4) x_t^n is $(J^n, I_0^n, \ldots, I_t^n, p_0, \ldots, p_t)$ - measurable, and

(5) p_t is $((J^n, I_0^n, \ldots, I_t^n, p_0, \ldots, p_t))_{n=1}^N$ - measurable.

In (3), the arg max should be over x that are feasible net trades for agent i, and $E(\cdot|\cdot)$ denotes conditional expectation with respect to P. (For (3) to make sense, some attention should be paid to fixing a regular version of conditional probability for each t; cf. Kreps [1977].)

In both this example and in our general formulation, the reader will see that what we have in mind is simply a grand rational expectations equilibrium as in Radner [1972], except that we allow differential information, and agents learn from equilibrium prices. This seems the natural extension of the static rational expectations equilibrium definition to sequence economies; see, for example, Futia [1979] or Townsend [1978].

The key to this formulation is that we directly model agents' uncertainty about what is the "true economy" by expanding the state space. This is in the spirit of Harsanyi's [1967-8] models of games with incomplete information: There is some shared underlying state space -- where agents differ initially is in their initial information about the state space and (perhaps) in thier prior assessments over that space. And an equilibrium must give equilibrium prices and trades for all the states, just as in Harsanyi's "Bayesian Nash equilibria." This approach to rational learning, therefore, has antecedents both in the literature on rational expectations equilibria (see Townsend [1978]) and in the literature of game theory.

The apparent "flaw" in this formulation of the problem is that it really avoids the main question: How do agents "learn" the relationship between prices and states of the world? Instead, it takes that problem for relatively simple economies (such as those of Grossman) and "resolves" it by positing that agents know the rational expectations equilibrium for a far more complex economy. The point of view is that if agents do not "know" the equilibrium, it is because they are uncertain about some parameters of the economy, such as the information possessed by or the risk tolerence of their fellows. However, given a realization of those parameters, the agents know what equilibrium prices will prevail. In effect, we imagine agents who are possessed of quite substantial computational power -- the power to "figure out" how the economy will evolve in any of its possible manifestations. (Compare this with the context of economies with a sequence of markets but with no differential information. In this context, the rational expectations notion is that of Radner [1972], where agents are presumed to "work out" what future prices and net trades will be.)

Is so much rationality necessary? For example, must agents agree on the underlying state space Ω? This question is motivated in part by the observation that agents will never see "disproved" their hypotheses for what will happen for states $\omega \in \Omega$ that do not in fact pertain. We can perhaps get by with something less than the above: We could imagine that each agent n has a personal model of what will ensue (given by a personal state space (Ω^n, F^n, P^n), and price and net trade mappings $\{p_t^n\}$ and $\{x_t^n\}$ on this space). To maintain the "rationality" of these diverse models, we would assume that they "intersect" for the true state of the world -- what does in fact happen is consistent with each agent's private model. (This weaker form of "rationality" is sometimes called "fulfilled expectations" in the literature.) We shall not pursue this sort of model further; the reader will see that while our later results do not apply to it directly (as they consist of a.s. statements, and thus have no particular validity for any single point ω), some salvaging may be possible.

A final comment: We use a model where the equilibrium concept is nothing more than the usual rational expectations equilibrium. Therefore, besides the question of how agents "figure out" this equilibrium, we must face the problems of existence of equilibrium and implementation of equilibrium. We have nothing to add to these questions -- we simply assume that equilibria exist (or, rather, prove statements about them when they do exist), and we assume that they can be implemented without further "informational leakage" to agents.

2. An example.

We now consider a particular example of the special case of section 1. This is an infinitely repeated version of the model of Grossman and Stiglitz [1980]; Bray [1980] uses this model to study an irrational learning process based on ordinary least squares estimation.

In this model there are two types of agents -- informed and uninformed. Each agent is endowed at date t with one unit of a one period risky asset yielding a random gross return r_t at date $t + 1/2$. There is also a safe asset, whose price and gross return are normalized to unity. Agent n has a utility function for period t

$$\exp(-(x_t^n r_t + y_t^n)/\theta^n),$$

where x_t^n and y_t^n are his holdings (between dates t and $t + 1/2$) of the risky and safe assets, and θ^n is a parameter measuring risk tolerence. (That is, $1/\theta^n$ is the agent's coefficient of absolute risk aversion.)

Both types of agents observe last period's return on the risky asset r_{t-1} prior to date t trading. In addition, the informed agents observe an unbiased predictor ρ_t of r_t. We assume that $\{\rho_t\}$ is an i.i.d. sequence of Normal random variables, and that $r_t = \rho_t + \varepsilon_t$, where $\{\varepsilon_t\}$ is also a sequence of i.i.d. Normal random variables and where the full collection of random variables $\{\rho_t, \varepsilon_t;\ t = 0,1,\ldots\}$ is independent. The error terms ε_t have zero means and variances σ^2. In the notation of section 1, the underlying uncertainty at date t concerns $\phi_t = (\rho_t, \varepsilon_t)$.

If all the agents knew ρ_t at t, they would have no reason to attempt to draw inferences about ρ_t from the price. If the price were p, agent n would demand $(\theta^n/\sigma^2)(\rho_t - p)$, and if there were N agents and so a total endowment of N, the equilibrium price would be $p = \rho_t - (N\sigma^2/\sum_n \theta^n)$. Following Radner [1979], we term this the full communication equilibrium price. If $(N\sigma^2/\sum_n \theta^n)$ is known, then ρ_t can be inferred from p_t. Thus even if initially some of the agents were uninformed of ρ_t, but these agents knew $(N\sigma^2/\sum_n \theta^n)$, the full communication equilibrium would be a rational expectations equilibrium in which the price p is completely informative.

Suppose however that there is uncertainty about the θ^n. To keep matters as simple as possible, we assume that it is common knowledge that there are only two agents, one informed, the other uninformed. The values of σ^2 (the variance of the error term) and the uninformed agent's coefficient of risk tolerence θ^U are also common knowledge. The uninformed agent is however initially uncertain about θ^I. (The informed agent of course knows the value of θ^I.) We suppose that at the outset, the uninformed agent has a prior assessment concerning θ^I that is given by a distribution function on some interval $[a,b] \subseteq (0,\infty)$ and that has a density.

Imagine that at some date t, after observing ρ_{t-1}, the uninformed agent's posterior assessment of θ^I is given by a distribution with density function f on the interval $[a,b]$. The demand by the informed agent, as a function of the price p, the information ρ_t, and θ^I, is

$$(6) \quad X^I(p,\rho_t,\theta^I) = \theta^I(\rho_t - p)/\sigma^2.$$

In equilibrium, $X^I + X^U = 2$, so the uninformed agent can infer X^I from his own equilibrium holding X^U. Equation (6) can be used to compute ρ_t as a function of θ^I and the other known variables (σ^2, X^I and p):

$$(7) \quad \rho_t = (\sigma^2 X^I/\theta^I) + p.$$

The uninformed agent's current assessment of θ^I, together with his prior on ρ^t, thus yields a posterior on r_t (which will not be Normal in general). We now show how this yields a demand function for the uninformed agent which is downward sloping and depends only on p.

Fix some p in (7) for the moment, and let X^I vary. As X^I increases, the uninformed agent's assessment over ρ_t and thus over r_t shifts (stochastically) upward. At any price p, this increases the desirability of the risky good for the uninformed agent: We can compute the uninformed agent's demand given the equilibrium data X^I and p, obtaining a schedule $x^U(p,X^I)$ (which depends, implicitly, on the density f) -- an equilibrium condition is that $X^I + x^U(p,X^I) = 2$, and since x^U increases with X^I, for each p there is a <u>unique</u> X^I satisfying this equation. Calling this value $x^I(p)$, we have that the demand of the uninformed agent, as a function of p alone, is $X^U(p) = 2 - x^I(p)$. Differentiating $x^I(p) + x^U(p,x^I(p))$ yields

$$\frac{dx^I}{dp} = -\left(\frac{\partial x^U}{\partial p}\right)/\left(1 + \frac{\partial x^U}{\partial x^I}\right) > 0.$$

And since $dx^I/dp = -dX^U/dp$, $X^U(p)$ is downward sloping. Recalling that $X^I(p,\rho_t,\theta^I)$ is a decreasing function of p, this implies that there is a unique and well-defined equilibrium price p_t solving

(8) $X^I(p,\rho_t,\theta^I) + X^U(p) = 2.$

This price p_t depends upon θ_t, ρ^I, and the density function f. In order to specify the entire equilibrium, we must finally specify how this density f evolves through time. Suppose that the agent, after learning r_{t-1} at time t, has an assessment given by f, and then acts as above. At time t + 1 he has three new pieces of data: p_t -- the equilibrium price last period; X_t^U -- his allocation; and r_t. As $X_t^I = 2 - X_t^U$, he can infer X_t^I. From (7), the likelihood of observing p_t, X_t^U and r_t given θ^I is the likelihood that $\rho_t = (\sigma^2(2 - X_t^U)/\theta^I) + p_t$ given r_t, which as $r_t = \rho_t + \varepsilon_t$, and (ρ_t,ε_t) is Normal, is the Normal density function $g(\cdot|r_t)$ of ρ_t given r_t. Thus using Bayes' rule, the revised posterior assessment for θ^I is:

$$(9) \quad P(\theta^I \leq z|p_t, X_t^U, r_t) = \frac{\int_a^z f(\theta)g((\sigma^2(2-X_t^U)/\theta) + p_t|r_t)\, d\theta}{\int_a^b f(\theta)g((\sigma^2(2-X_t^U)/\theta) + p_t|r_t)\, d\theta} .$$

Note that as the uninformed agent begins with a prior over θ^I that has a density function on the interval [a,b], (9) ensures that the posterior will also be of this form. As we assume that agents begin with such a prior, all subsequent posterior densities will have this form.

It is worth noting that while we have not computed the equilibrium prices, we could, in principle, do so. (All that would be needed is a large computer budget and, almost certainly, a good graduate student.) The major difficulty is in computing $x^U(p,X^I)$ -- since the uninformed agent's posterior assessment for r_t is not Normal, finding his optimal net trade is quite hard.

If this can be done efficiently, then solving numerically $x^I(p) + x^U(p,x^I(p)) = 2$ to get $X^U(p)$, and then (8) and (9) will not be difficult. But closed form solutions are out of the question.

This example would be no more difficult if the informed agent were uncertain about θ^U -- the informed agent doesn't care at all about θ^U in forming his demand. But things would be much more complex is, say, we had two agents, each of whom got a different noisy signal about r_t, neither of which was strictly superior information, and if neither knew the other's coefficient of risk tolerence. As long as each side can learn something of value to it from the equilibrium price, finding non-trivial equilibria seems impossible.

It is, on the other hand, quite easy to construct trivial examples of "fully revealing" equilibria. Following Radner [1979], if Θ and Φ are both finite sets, then "generically" p_0 will be one-to-one and will reveal J instantly; and thereafter each p_t will reveal I_t completely. Similar results, following the analysis of Allen [1979], where Θ and Φ have low dimension when compared with the number of coordinates in price vectors, are also possible. And if one is content with "almost equilibria," then Jordan's [198] results apply to give full revelation. The point is that we have here nothing more than a "complicated" rational expectations equilibrium problem, and the literature that concerns these equilibria applies directly.

3. Convergence of Posterior Assessments.

It is trivial to prove that posterior assessments "converge" almost surely. First consider any agent's posterior assessments of a fixed measurable set $A \subseteq \Omega$. At date t, the agent's assessment that A obtains is given by $E^n[1_A|H^n_t(p)]$, where $E^n[\cdot]$ denotes expectation taken with respect to P^n.

Proposition 1. For arbitrary measurable $A \subseteq \Omega$, and for each agent n,

$$\lim_{t\to\infty} E^n[1_A | H_t^n(p)] = E^n[1_A | H_\infty^n(p)] \quad P^n\text{-a.s.},$$

where $H_\infty^n(p) = \bigvee_{t=1} H_t^n(p)$.

The proof is a trivial consequence of the fact that $E^n[1_A | H_t^n(p)]$ forms a bounded martingale -- see Chung [1974, Theorem 9.4.8]. It is worthwhile to note something this result does not say: It does not say that posteriors converge for all agents "almost surely," where "almost surely" refers to some objective probability measure. If, for example, P^n and $P^{n'}$ are mutually singular assessments for some agents n and n', then it is possible that there is no point in Ω for which convergence holds for both n and n' simultaneously. But, of course, in cases where the P^n all have identical null sets, we get simultaneous convergence of posteriors (except possible on one of these null sets).

Now we wish to sharpen this convergence result, by considering how agent's "entire posterior" evolve. To do this requires some topological assumptions on the "probability space of interest," so to illustrate matters we will use the concrete example of section 1, where $\Omega = \Theta \times \Phi^\infty$, and we will focus on agents' marginal posteriors on the space Θ. (This restriction of attention is natural in the "disconnected periods" models of Grossman.)

Thinking of Θ as recording agents' risk tolerences and other similar parameters, it is natural to suppose that Θ has "nice" topological structure. We shall assume that Θ is a complete separable metric space and that, with the σ-field on Θ generated by its open sets, it is a Borel space. (So, for example, Θ is some subset of R^k for finite or even countably infinite k.) Thus we can fix a regular version of conditional probability P^n_t on Θ that represents agent n's posterior at date t over Θ; that is, for A a measurable subset of Θ,

$$P^n_t(A) \text{ is a version of } E^n[1_A | H^n_t(p)], \text{ and}$$

$$P^n_t(\cdot) \text{ is a probability measure on } \Theta, \; P^n\text{-a.s.}$$

Similarly, we can fix regular versions P^n_∞ of $E^n[1_{\{\cdot\}} | H^n_\infty(p)]$.

Proposition 2. Under the conditions above (Θ is a complete and separable metric space, and it forms a Borel space), for fixed regular versions of conditional probability P^n_t and P^n_∞, we have that P^n_t converges weakly to P^n_∞, P^n-a.s.

Proof. Since Θ is a complete and separable metric space, we can find for each integer k a countable collection $\{A^k_m ; m = 1,2,\ldots\}$ of disjoint measurable subsets of Θ, such that each A^k_m is contained within a $1/k$-ball and $\bigcup_{m=1}^{\infty} A^k_m = \Theta$. By proposition 1, $\lim_t P^n_t(A^k_m) = P^n_\infty(A^k_m)$ for every k and m, except possibly on a P^n-null set. (Since there are countably many k and m, this is true for all k and m simultaneously.) Fix some $\omega \in \Omega$ not from that null set.

We prove the weak convergence of P^n_t to P^n_∞ (at the point ω) using criterion (ii) of the portmanteau theorem of Billingsley [1968, page 12]. Let f be a bounded, uniformly continuous, real valued function on Θ. Let Z denote the bound on f ($Z \geq |f(\phi)|$ for all $\theta \in \Theta$) and let $\delta(\varepsilon)$ denote the modulus of continuity of f. Fix some $\varepsilon^\circ > 0$ and let $\delta^\circ = \delta(\varepsilon^\circ)$.

Let L° be an integer larger than $1/\delta^\circ$. Pick K° large enough so that $\sum_{k=1}^{K^\circ} P^n(A^{L^\circ}_k) > 1 - \varepsilon^\circ/Z$. Pick T° large enough so that for all $t > T^\circ$

and for all $k = 1,\ldots,K°$, $|P_t^n(A_k^{L°}) - P_\infty^n(A_k^{L°})| < \varepsilon°/(K°Z)$. This implies that $\sum_{k=K°+1}^\infty P_t^n(A_k^{L°}) \leq 2\varepsilon°/Z$ for all $t > T°$. And therefore, for all $t > T°$,

$$|\int f\, dP_t^n - \int f\, dP_\infty^n| \leq$$

$$\sum_{k=1}^{K°} [|\int_{A_k^{L°}} f\, dP_t^n - f(a_k^{L°})P_t^n(A_k^{L°})| + |f(a_k^{L°})(P_t^n(A_k^{L°}) - P_\infty^n(A_k^{L°}))|$$

$$+ |\int_{A_k^{L°}} f\, dP_\infty^n - f(a_k^{L°})P_\infty^n(A_k^{L°})|] +$$

$$\sum_{k=K°+1}^{\infty} [|\int_{A_k^{L°}} f\, dP_t^n| + |\int_{A_k^{L°}} f\, dP_\infty^n|] \leq \varepsilon° + \varepsilon° + \varepsilon° + 2\varepsilon° + \varepsilon°,$$

where $a_k^{L°}$ is an arbitrarily selected point from $A_k^{L°}$. []

4. Sharpening the Convergence Result.

We return to the example of section 2. Let F_t denote the uninformed agent's posterior distribution function for θ^I at date t after learning r_{t-1}, p_{t-1}, and x_{t-1}^I. (Of course, F_t is a random element depending upon the uninformed agent's initial prior on θ^I, the actual θ^I, and the values of ρ_s and ε_s for $s < t$.) Applying proposition 2 (or, more precisely, a small variation on proposition 2), we know that with probability one, F_t converges weakly to some F_∞. What is F_∞? In this example we can show that F_∞ will be a point mass at θ^I. The argument that leads to this conclusion has three parts:

(A) The price at date t, $p_t(\theta^I,\theta^U,\rho_0,\ldots,\rho_t,\varepsilon_0,\ldots,\varepsilon_{t-1})$ will almost surely approach some random price functional $p_\infty(\rho_t)$ as t goes to infinity. To see this, note that all that is germane in $(\theta^I,\theta^U,\rho_0,\ldots,\rho_t,\varepsilon_0,\ldots,\varepsilon_{t-1})$ for computing p_t are the values of θ^I, θ^U, and ρ_t, and the distribution function F_t. Now θ^I and θ^U are the same for all t, and the effect on

$X^U(p)$ of F_t is continuous in the weak topology on F_t. So as F_t approaches F_∞ weakly, X^U will approach the demand the uninformed agent would make if he had posterior F_∞. This, in turn, implies that equilibrium prices approach what they would be if the uninformed agent had this posterior. (We leave details of this argument to the reader.) To show the dependence of this limiting relation p_∞ on F_∞, θ^I, and θ^U, we shall write $p_\infty(\rho_t;F_\infty,\theta^I,\theta^U)$.

(B) The distribution of $p_\infty(\theta_t;F_\infty,\theta^I,\theta^U)$ is $H^U_\infty(p)$ - measurable. To see this, note that as ρ_t forms an i.i.d. sequence, if the uninformed agent simply keeps a frequency distribution of the observed equilibrium prices, part (A) implies that this will (almost surely) approach the distribution function of $p_\infty(\rho_t;F_\infty,\theta^I,\theta^U)$. As the frequency distribution functions are all $H^U_\infty(p)$ - measurable, their limit must also be so. (Again we leave the precise details to the reader.)

(C) Knowledge of the distribution function of $p_\infty(\rho_t;F_\infty,\theta^I,\theta^U)$ and of θ^U and F_∞ is sufficient to compute the true value of θ^I. One more we leave details to the reader, but the key is that p_∞ is stochastically decreasing in θ^I for fixed θ^U and F_∞. This together with (B) implies that θ^I is $H^U_\infty(p)$ - measurable. Thus for any subinterval $[c,d]$ of $[a,b]$,

$$P^U_\infty(\{\theta^I \varepsilon [c,d]\}) = E^U[1_{\{\theta^I\varepsilon[c,d]\}}|H^U(p)] = 1_{\{\theta^I\varepsilon[c,d]\}}.$$

To what extent does this result generalize? Knowing that posteriors converge is something, but can we say anything in general about to what they converge? The answer to these questions is mixed. The arguments above can be generalized to some extent, in contexts where $\Omega = \Theta \times \Phi^\infty$. (For models with less structure, we do not even know how to pose precisely a meaningful sharpening of proposition 1.) But such generalizations can be difficult:

Generalizing step (A) can be quite hard. In our very simple model, it is easy to show that equilibrium prices "settle down" to some stationary relationship if posteriors converge. But in general, this sort of argument depends on a lot of smoothness, especially since small changes in a price functional can lead to enormous changes in the information that those prices communicate.

One step (A) is surmounted, step (B) is easy -- only some sort of ergodicity of $\{\rho_t\}$ is needed.

As for (C), this definitely exploits the special structure of our model. In general, the long-run distribution of prices and other observables (for the limit posteriors) as a function of θ will not be one-to-one. For example, suppose that in the example of section 2 there were to informed agents with coefficients of risk aversion θ^{I1} and θ^{I2}, both unknown by the uninformed agent. Then in the resulting equilibrium, the uninformed agent would learn the value of $\theta^{I1} + \theta^{I2}$, but he could never disentangle θ^{I1} any further than that. (On the other hand, once the uninformed agent knows $\theta^{I1} + \theta^{I2}$, there would be no value to him in knowing θ^{I1} in addition.)

5. An Example of Blume and Easley.

It is interesting to compare the results of section 4 with those of Blume and Easley [1980], in which convergence to an "irrational expectations" equilibrium is possible. In what follows we present a somewhat simplified version of Blume and Easley's model, translated into language consistent with our own. The reader should consult their paper for a better picture of what they are doing. There are two possible values of θ, θ^0 and θ^1, and there are two agents. Agent n observes a signal I_t^n prior to date t trading. After date t trading, agent n learns the entire signal $I_t = (I_t^1, I_t^2)$. The signal I_t can take on only a finite number of different values. If θ^0 obtains and both agents know this, then there is a rational expectations equilibrium price function $p^0(I_t)$ for date t. If θ^1 obtains and this is known, then equilibrium prices are given by $p^1(I_t)$.

Each agent n entertains two models ψ_n^0 and ψ_n^1 about how their world might be structured, corresponding to their initial uncertainty about whether θ^0 or θ^1 obtains. A model is a function $\psi_n(I|I^n,p)$ which gives the probability of $I_t = I$ given that agent n observes I^n and the equilibrium price p. The models ψ_n^0 are consistent with the rational expectations equilibrium which prevails when θ^0 obtains, in the sense that for any p observed in the rational expectations equilibrium p^0, $\psi_n^0(I|I^n,p)$ gives the objective equilibrium distribution for I given I^n and p. However each ψ_n^0 is also defined for values of p which are not observed in equilibrium. Similar remarks apply to the ψ_n^1.

Each agent starts with a prior probability assessment over which model obtains, which yields a probability distribution over I_t given I_t^n and p_t. After observing I_t and p_t, each agent updates his assessment over his two models in a manner that attributes increased posterior probability to whichever model gave greater likelihood to the observed (I_t, p_t). Blume and Easley show that in this world, it is possible that expectations converge to a non-rational expectations equilibrium with positive probability.

This happens as follows. If θ^0 obtains and agents believe that model 1 is correct, prices are given by $p^*(I_t)$ which is equal to neither $p^0(I_t)$ or to $p^1(I_t)$. Moreover, prices "close" to p^* are much more "likely" under model 1 than under model 0: $\psi_n^1(I|I_n,p^*(I)) > \psi_n^0(I|I_n,p^*(I))$ for all possible I, and similarly for all p close to p*. Thus is agents start by believing that model 1 is much more likely to apply than model 0, and if θ^0 does indeed obtain, then they are likely to come to assess probability approaching one to the "incorrect" model 1. This may happen even if model 1 is a very poor description of the actual situation. To give a concrete example, suppose that I can take on only two values, I' and I", and that $p^*(I') = 10$ and $p^*(I'') = 20$, each with probability one. Suppose that agent 1 has no private information, and that his models are $\psi_1^1(10|10) = .01$, $\psi_1^1(20|20) = .01$, $\psi_1^0(10|10) = .001$, and $\psi_1^0(20|20) = .001$. In the equilibrium to which the learning process converges, agent 1 observes repeatedly that $p = 10$ whenever $I = I'$ and $p = 20$ whenever $I = I''$. But he continues to believe that if $I = I'$, $p = 10$ with probability .01, and $p = 20$ with probability .01 if $I = I''$. This is because the only other model that this agent entertains assesses even lower likelihood to the observed events. Like any good detective, this agent rejects the probable impossibility for the improbable possibility, because although improbable, this possibility is the <u>most</u> probable possibility.

In this example, our step (A) from section 4 is satisfied -- prices do settle down into a stable relation. But step (B) is inapplicable. Even though in the long run the agent observes that $I = I'$ <u>always</u> gives $p = 10$, he assesses zero posterior probability to this relationship, because he assessed zero prior probability to it. (It was not one of his two models.) This could never happen in a rational learning model (or, more precisely, it happens with zero prior probability), because agents must assess positive prior probability to pricing relations that have positive prior probability. And, we contend, even if agents do not learn in a wholly rational manner, it is hard to imagine agents so irrational that they don't begin to wonder if some stable relation is emerging that they hadn't counted on.

6. Concluding Remarks.

In this paper we have shown that if agents are rational Bayesians, then their beliefs must converge. Moreover, with "sufficient" regularity, beliefs converge to the "correct beliefs" which characterize a stationary rational expectations equilibrium.

However we do not believe that these results provide a satisfactory answer to the question: "How does a rational expectations equilibrium come about?" The Bayesian learning process that we describe allows for agents' uncertainty about things which are taken to be known parameters in standard models of rational expectations equilibrium (e.g., risk aversion). However we assume that agents have extraordinary insight into the way their world operates given these parameters, and an ability to calculate the probability of events in their world given their prior assessments which greatly exceeds our own. We were able to outline how such a calculation could be performed for an extremely simple example in which only one agent learns, but we not provide a closed-form solution. Other authors have also worked particular examples of models with Bayesian learning (see the introduction). However in general agents would be solving very difficult problems with extremely complex dynamics, particularly when two or more agents are learning from each other.

We find this incredible, and so despite the attractive convergence properties of rational learning, we are led to reject it as a plausible model for investigating the attainment of rational expectations equilibria. One could, of course, accept the point that economic agents, like billiards players, probably do not carry out the appropriate calculations and probability assessments, but follow Friedman [1953] in arguing that, to a good approximation, agents act as if they did so. Friedman argues that the appropriate test of a behavioral assumption is not its superficial plausibility, but its predictive power. Without entering into this methodological debate, we feel safe in saying that we are all in favor of empirical tests of the rational expectations hypothesis.

Nevertheless, we would like to point out that we do in fact know something about how agents form expectations. In some instances people are likely to be somewhat inarticulate about what their expectations are and where they come from. However in many cases economic expectations are derived in whole or in part from the conscious application of techniques of statistical inference to past data. These techniques may or may not be those which our hypothetical brilliant Bayesian would use -- econometricians may lack the insight to know _a priori_ what the appropriate specification of the model to be estimated is, as do we. However, given a model for which a rational expectations equilibrium has been posited, one can pose the question: "Could agents in the model who initially don't know how to form rational expectations learn how to do so by using standard statistical techniques on the data generated by the model?" One of us (Bray) has done precisely this and has obtained, albeit with some difficulty, and for a very simple model, a condition guaranteeing almost sure convergence to the rational expectations equilibrium. General analytic results along these lines may be hard to come by; the simulation approach followed by Cyert and DeGroot [1974] is more tractable, although care is required in interpretation of the results.

On the other hand, considerations of rational statistical inference can be used to place bounds upon the degree of non-rationality that we are prepared to countenance. In particular, we find it implausible that a learning process converges to a stationary equilibrium in which agents maintain beliefs indefinitely which are systematically confounded by the events that they observe. This is reminiscent of the old argument for adopting the rational expectations equilibrium concept, and we accept that in the long run, equilibrium expectations must either be nonstationary or else rational. We merely believe that we have no satisfactory story about the attainment of a long run equilibrium in which expectations matter.

Acknowledgements. Conversation with Paul Milgrom proved very helpful. This research was supported in part by National Science Foundation Grant SES80-06407 to the Graduate School of Business, Stanford University, and by the Graduate School of Business, Stanford University.

References.

Allen, B. [1979], "Generic Existence of Completely Revealing Equilibria for Economies with Uncertainty when Prices Convey Information," Center for Analytic Research in Economics and the Social Sciences, Working Paper #79-07, University of Pennsylvania.

Arrow, K. J. and J. R. Green [1973], "Notes on Expectations Equilibria in Bayesian Settings," Institute for Mathematical Studies in the Social Sciences, Working Paper No. 33, Stanford University.

Blume, L. E. and D. Easley [1980], "Learning to be Rational," mimeo, University of Michigan, forthcoming in the Journal of Economic Theory.

__________________________ [1981], Rational Expectations Equilibrium and the Efficient Markets Hypothesis," mimeo, Cornell University.

Bray, M. M. [1980], "Learning, Estimation, and the Stability of Rational Expectations," mimeo, Stanford University, forthcoming in the Journal of Economic Theory.

Chung, K. L. [1973], A Course in Probability Theory, New York: Academic Press.

Cyert, R. M. and M. H. DeGroot, "Rational Expectations and Bayesian Analysis," Journal of Political Economy, Vol. 82, 521-536.

Feldman, M. [1981], "Structural Stability with Learning of Rational Expectations," Thesis Proposal, University of California at Davis.

Friedman, M. [1953], Essays in Positive Economics, Chicago: University of Chicago Press.

Frydman, R., "Towards an Understanding of Market Processes: Individual Expectations, Market Behavior and Convergence to Rational Expectations Equilibrium," mimeo, C. V. Starr Center for Applied Economics, New York University.

Futia, C. [1981], "Rational Expectations in Stationary Linear Models," Econometrica, Vol. 49, 171-190.

Grossman, S. [1977], "The Existence of Futures Makrets, Noisy Rational Expectations and Informational Externalities," Review of Economic Studies, Vol. 64, 431-449.

Grossman, S. and J. Stiglitz [1980], "On the Impossibility of Informationally Efficient Markets," American Economic Review, Vol. 70, 393-408.

Harsanyi, J. [1967-68], "Games with Incomplete Information Played by Bayesian Players," Parts I, II, and III, Management Science, Vol. 14, 159-182, 320-334, 486-502.

Jordan, J. S. [1981], "On the Generic Existence of Rational Expectations Equilibrium," mimeo, University of Minnesota.

Kihlstrom, R. E. and L. J. Mirman [1975], "Information and Market Equilibrium," Bell Journal of Economics, Vol. 6, 357-376.

Kreps, D. M. [1977], "A Note on 'Fulfilled Expectations' Equilibria," Journal of Economic Theory, Vol. 14, 32-43.

Lewis, G. [1981], "The Phillips Curve and Bayesian Learning," Journal of Economic Theory, Vol. 24, 240-264.

Radner, R. [1968], "Competitive Equilibrium Under Uncertainty," Econometrica, Vol. 36, 31-58.

__________ [1972], "Existence of Equilibrium of Plans, Prices and Price Expectations in a Sequence of Markets," Econometrica, Vol. 40, 289-303.

__________ [1979], "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," Econometrica, Vol. 47, 655-678.

__________ [1980], "Equilibrium under Uncertainty," in Handbook of Mathematical Economics, Vol. 2, (K. J. Arrow and M. D. Intrilligator, eds.), North-Holland.

Townsend, R. M. [1978], "Market Anticipations, Rational Expectations and Bayesian Analysis," International Economic Review, Vol. 19, 481-494.