

Learning Rational Expectations: The Finite State Case

J. S. JORDAN*

Department of Economics, University of Minnesota, Minneapolis, Minnesota 55455

Received September 21, 1982; revised October 29, 1984

This paper is devoted to the question of whether traders can learn rational expectations from repeated observations of market data in a stationary environment with finitely many exogenous states of the world. The learning problem is placed in the context of an iterative adjustment process which achieves equilibrium if traders have rational expectations. The main result is that even if traders begin with no knowledge of their environment, there exists an estimation procedure which converges to rational expectations when the environment satisfies a certain regularity condition. The regularity condition is shown to be generic. *Journal of Economic Literature* Classification Numbers: 021, 022, 026, 213. © 1985 Academic Press, Inc.

1. INTRODUCTION

The rational expectations hypothesis requires traders to anticipate the future according to the objective probability distribution of future events conditioned on all of their current information. This information consists of private information, which differs across traders and is usually taken to be exogenous, and market information, such as prices. Since prices are sensitive to the private information of all traders, prices may reveal to each trader some or all of the decision-relevant private information of others. However, forming the correct conditional distribution of future events requires some knowledge of the relation between private information and prices, which in turn depends on the characteristics of the various traders. Such knowledge is not likely to be directly available to each trader, so it must be gained by experience. Hence the plausibility of the rational expectations hypothesis hinges on the ability of traders to learn the correct conditional distributions from repeated observations of market data.

The problem of learning rational expectations is greatly complicated by the well-known dependence of the "correct" conditional distribution on

* The support of National Science Foundation Grant SOC-77-07852 is gratefully acknowledged. I would also like to thank the referee and associate editor for helpful suggestions.

trader's beliefs. Indeed, a rational expectations equilibrium is essentially a fixed point in a space of conditional distributions. Thus the problem of learning rational expectations is partly a problem of consistent estimation and partly a problem of the stability of equilibrium. Because of the latter difficulty one might expect any positive results to be of limited generality.

This expectation is borne out by the results in the literature thus far. Possibly the most favorable result has been obtained by Bray [4] for a market for a financial asset with an uncertain future value. All traders are assumed to have constant absolute risk aversion and all random variables are assumed normal. In [4, Proposition 5] Bray demonstrates that for a large region in the space of parameters which specify the market, the linear regression of eventual asset values on previous prices converges to rational expectations. A more limited result in a more general model has been obtained by Blume and Easley [3]. They study a general class of stochastic exchange environments with finitely many "states of information." Each trader tries to infer the current state of information, which describes the private information of every trader, from the current price. Blume and Easley specify a family of learning processes by which traders adjust their beliefs about the relation between prices and states of information in response to their observations as the environment is repeated over time. Under certain assumptions, rational expectations are shown to be "locally stable" under learning [3, Corollary 1], but non rational expectations may also be locally stable [3, Corollary 2 and Sect. 4]. Studies previous to these two papers are surveyed in [2].¹

The papers mentioned above use the conventional model of rational expectations equilibrium in which each trader's expectations are conditioned on his own private information and the equilibrium price. As the environment is repeated over time, each trader revises his estimate of the correct probability distribution conditioned on private information and prices. In the opinion of the present writer this approach severely and unnecessarily restricts the study of the learning problem. First, the model must be restricted to exclude the well-known examples of the nonexistence of rational expectations equilibrium. Second, trader's beliefs must be restricted to ensure that a market clearing price will exist in each sample period, and the learning process must preserve these restrictions on beliefs. These limitations prevent a general convergence theorem from being stated much less proved.

The present paper uses a different model of rational expectations equilibrium to obtain a fairly general convergence theorem. We place the

¹ I have of course omitted to mention those learning models, termed "rational learning" by Bray and Kreps [5], in which the learning process is itself essentially a rational expectations equilibrium in a larger space of conditional distributions.

learning problem in the context of the dynamic adjustment process constructed in [10] to implement rational expectations equilibria. To give a brief description of this process we will suppose for the moment that the learning problem has been solved, so that traders know the correct conditional distributions conditioned on any random variables they observe. Initially, each trader conditions his expectations only on his own private information, which is represented by a signal s^i . This results in a market clearing price $f_1(s)$ for each joint-signal (state of information) $s = (s^i)_{i=1}^N$. However, before the trades $(y_1^i(s))_i$ associated with the price $f_1(s)$ are consummated, each trader i revises his expectations by conditioning on $f_1(s)$ as well as s^i , leading to a new market clearing price $f_2(s)$. Again, before the new trades are consummated, each trader revises his expectations by conditioning on s^i , $f_1(s)$, and $f_2(s)$, leading to a new price $f_3(s)$, and so on. This adjustment process continues, with traders retaining the information they acquire at each step, until it reaches a price $f_A(s)$ which reveals no new information to any trader. We will assume there are only finitely many joint-signals s , so the process will terminate in a finite number of steps. The terminal price function f_A is called a *dynamic equilibrium* price function.

A dynamic equilibrium f_A is a rational expectations equilibrium in that markets clear and traders condition their expectations on all of the information available to them, including previous prices. If the earlier prices reveal decision-relevant information which is not revealed by the dynamic equilibrium prices, then f_A may not satisfy the more common definition of rational expectations equilibrium which would require expectations to be conditioned on only private information and the final price $f_A(s)$ in each state s . In fact, a conventional rational expectations equilibrium may fail to exist, while a dynamic equilibrium always exists.

This process was suggested to the present author by Reiter [13], who formulated it for general informationally decentralized message processes, and was developed independently by Kobayashi [11] for a financial asset market with normality and constant absolute risk aversion. Kobayashi permits the preequilibrium trades $(y_a^i(s))_i$ to be consummated but assumes that traders do not anticipate the resulting capital gains and losses during the adjustment process. It is shown in [10, Sect. 4] that if preequilibrium trades are consummated and capital gains and losses are rationally anticipated, the preequilibrium market clearing prices $f_a(s)$ may fail to exist.

Placing the learning problem in this context has the immediate effect of giving traders more to learn. Traders must learn the correct distribution of future events conditioned on private information and a sequence of prices rather than a single price. However, this difficulty is offset by the fact that the adjustment process removes the simultaneous determination of prices and expectations. The price $f_a(s)$ influences traders' expectations at stage

$a + 1$ of the process, but not at stage a . This feature of the process leads to the general existence of dynamic equilibria² and also brings the learning problem much closer to the problem of consistent statistical estimation.

Section 2 below gives the formal definitions of the adjustment process and the learning problem. We suppose that the environment is repeated over time, and in each sample period t the adjustment process takes place, giving trader i the data s_t^i, p_{at} for each iteration $a \geq 1$, and the realized future state σ_t . For each adjustment stage a , the i^{th} trader must use the data $(s_t^i; p_{1t}, \dots, p_{at}; \sigma_t)_{t=1}^T$ to form, in sample period $T + 1$, an estimate of the distribution of the future state σ conditioned on $(s_{T+1}^i; p_{1(T+1)}, \dots, p_{a(T+1)})$. These estimates for all traders determine $p_{(a+1)(T+1)}$. A rule for estimating this conditional distribution is called an *estimation procedure*, and the entire array of estimation procedures for all traders and all adjustment stages is an *estimation scheme*. An estimation scheme is *successful (a.s.)* if, for almost every infinite sample, the estimates converge to the conditional distributions mentioned in our initial description of the adjustment process given above.

In Section 4 we construct an estimation scheme which is successful for all stochastic environments satisfying a regularity condition we call *Axiom L*. Put somewhat loosely, Axiom L requires that for each iteration a , the market clearing price p_a is a Lipschitzian function of traders' conditional distributions when the latter lie in some neighborhood of the correct conditional distributions. Our result is based on the following reasoning. At the first stage of the price adjustment process, the learning problem is trivial. Trader i conditions his expectations of σ on the exogenous information s^i alone, so the conditional frequency distribution provides a convergent estimation procedure. Moreover, the law of the iterated logarithm provides a rate of convergence, which, under Axiom L, extends to a rate of convergence on the joint distribution of s_t^i, σ_t , and the first stage price p_{1t} . At the second stage we have a nontrivial learning problem because trader i conditions his expectations of σ on s^i and p_1 , and the true joint distribution of s_t^i, p_{1t} , and σ_t is changing over time due to learning at the first stage. However, we have a rate of convergence for this distribution, which we use to construct an appropriate modification of the conditional frequency distribution of σ_t given s_t^i and p_{1t} . Again, the law of the iterated logarithm gives a rate of convergence for the second stage estimation procedure; and so on.

The recursive nature of the price adjustment process is critical to this argument. At each stage, learning affects the joint distribution being lear-

² If there are infinitely many joint-signals, the process will generally not terminate after finitely many iterations. However, there is a natural alternative definition of equilibrium for this case, and a general existence theorem is obtained in [10, Sect. 5].

ned at the next stage, but there is no feedback from later to earlier stages. For this reason, the learning problem can be solved sequentially beginning with the trivial first stage problem. Indeed this recursive structure may be the only crucial feature of the model. At the present level of generality, no qualitative restrictions apply to the map from expectations to market-clearing prices, except the Lipschitz condition which we impose by assumption. Therefore, the reasoning used in this paper may apply to other models of endogenous expectations in which the informational variables are generated recursively.

In Section 3, the estimation scheme is introduced in the context of a class of two-trader, two-state stochastic environments. The general result, in Section 4, is a fairly straightforward extension of the example in Section 3. Axiom L is interpreted in Section 4.2 and further motivated by Proposition 5.5, which states that Axiom L is, in a certain precise sense, a typical property of smooth stochastic environments. It may be worthwhile to emphasize that Section 5 serves only to motivate Axiom L, and is not otherwise related to the learning problem. Some remaining open questions are mentioned in Section 6.

2. THE MODEL

2.1. DEFINITIONS. There are N traders, indexed by superscript i , and M commodities. For each i , let S^i be a finite set with generic element $s^i = (z^i, \omega^i)$, where ω^i is an endowment bundle in R_{++}^M ,³ and z^i is an abstract signal. The set S^i is trader i 's *private information set*, and elements s^i will be called *signals*. Let Σ be a finite set of *future states*, with generic element σ , and for each i let $v^i: R_+^M \times \Sigma \rightarrow R$ be i 'th trader's state-dependent utility function. We will assume that for each i and each $\sigma \in \Sigma$,

(A.1) $v^i(\cdot, \sigma)$ is continuous, strictly increasing, and strictly concave.

Let $S = \prod_i S^i$, with generic element $s = (s^i)_i$, and let π be a probability measure on $S \times \Sigma$. We will assume for convenience that $\pi(s, \sigma) > 0$ for all $(s, \sigma) \in S \times \Sigma$, so that the phrases "every (s, σ) " and "almost every (s, σ) " are identical. A *stochastic environment* is completely described by π and an N -tuple of utility functions $(v^i)_i$. For convenience in exposition we will keep the number of traders and the set $S \times \Sigma$ fixed. The set of stochastic environments is the set \mathcal{S} of pairs $(\pi, (v^i)_i)$ with $\pi(s, \sigma) > 0$ for all $(s, \sigma) \in S \times \Sigma$ and v^i satisfying (A.1) for each i .

2.2. *Remarks.* In modelling the i 'th trader's private information as a signal s^i we are departing slightly from the model in [10]. In [10] each

³ $R_{++}^M = \{x \in R^M: x_j > 0 \text{ for all } j\}$, and $R_+^M = \{x \in R^M: x_j \geq 0 \text{ for all } j\}$.

trader's private information is modelled as a partition or subfield of a common probability space on which all random variables, present and future, are defined. Formally the above representation is a special case of the model in [10]. We have used the present definition to support the interpretation that the future state σ is eventually observed by all traders, but the private signal s^i is not observed by trader j at any time. Hence in the present model, the learning problem is complicated by the presence of unobservable variables.

We now define the sequence of informational temporary equilibria and the dynamic equilibrium which would result if traders could form the correct expectations conditional on all random variables they observe.

2.3. DEFINITIONS. Let $\Omega = \{(\omega^i)_i \in (R^M_{++})^N: \text{for some } (z^i, \omega^i)_i \in S, \omega'^i = \omega^i \text{ for all } i\}$, and let Φ denote the set of probability distributions on Σ . Let Δ denote the strictly positive price simplex, $\Delta = \{p \in R^M: p_j > 0 \text{ for all } j \text{ and } \sum_j p_j = 1\}$. Given a stochastic environment $(\pi, (v^i)_i)$, let $w: \Omega \times \Phi^N \rightarrow \Delta$ be a function such that for each $((\omega^i)_i, (\phi^i)_i) \in \Omega \times \Phi^N$, $w((\omega^i)_i, (\phi^i)_i)$ is a Walrasian equilibrium price for the (nonstochastic) exchange economy with endowments $(\omega^i)_i$ and utility functions $(\sum_\sigma \phi^i(\sigma) v^i(\cdot, \sigma))_i$.

For each $s \in S$ and each i , let $\phi_0^i(s) \in \Phi$ denote the conditional distribution on Σ determined by π given s^i . That is,

$$\phi_0^i(s)(\sigma) = \Sigma \{ \pi(s', \sigma): s'^i = s^i \} / \Sigma \{ \pi\{s', \sigma'\}: s'^i = s^i \}.$$

Define $f_1: S \rightarrow \Delta$ by $f_1(s) = w((\omega^i)_i, (\phi_0^i(s))_i)$, where the endowment N -tuple $(\omega^i)_i$ is of course specified by $s = (z^i, \omega^i)_i$. For $a > 1$, we will define $(\phi_{a-1}^i(s))$ and $f_a(s)$ by induction. For each i , let $\phi_{a-1}^i(s)$ denote the conditional distribution on Σ given s^i and $f_1(s), \dots, f_{a-1}(s)$. Let $f_a(s) = w((\omega^i)_i, (\phi_{a-1}^i(s))_i)$. Finally, let $A = \min \{a: f_{a+1} = f_a\}$. The function f_A is a *dynamic equilibrium price function*, and the functions f_a for $a < A$ are *informational temporary equilibria*.

2.4. Remarks. We use the subscript a to denote "adjustment time." The index t will be used below for sampling time. The function w depends on the utility N -tuple and so will vary with the environment. Unfortunately the definition of w is arbitrary to the extent that a nonstochastic exchange environment may have multiple equilibrium prices. To define conditional expectations it is necessary to assume that a particular equilibrium price is consistently selected. We could of course permit a randomized selection, but only by imposing a particular randomization. The definition of w is intended to represent the outcome of the equilibrating forces of the market, and implicitly assumes the market to be sufficiently systematic that a single equilibrium price is selected. The definition of the informational temporary

equilibrium f_a embodies the assumption, mentioned in the Introduction, that the trades at stage $a-1$ were not consummated.

A dynamic equilibrium f_A is a "rational expectations equilibrium" provided that traders condition their expectations on their dynamic equilibrium trades as well as f_A . More precisely, let $(\pi, (v^i)_i)$ be a stochastic environment with the dynamic equilibrium f_A , and for each i and each s , let $y_i(s)$ maximize $\sum_{\sigma} \phi_A^i(s)(\sigma) v^i(\omega^i + y, \sigma)$ subject to $f_A(s) y \leq 0$. For each i , let $\phi_*^i(s)$ denote the conditional distribution on Σ given $s^i, f_A(s)$, and $y^i(s)$ (or merely $f_A(s)$ and $y^i(s)$). Then it is proved in [10, Theorem 3.6] that $(f_A(s), y^1(s), \dots, y^N(s))$ is a Walrasian equilibrium for the exchange economy with endowments $(\omega^i)_i$ (given by s) and utility functions $(\sum_{\sigma} \phi_*^i(s)(\sigma) v^i(\cdot, \sigma))_i$.

2.5. Estimation Procedures

For each trader i , let $X_0^i = S^i \times \Sigma$ and for $a > 0$ let $X_a^i = S^i \times \Delta^a \times \Sigma$. An *estimation procedure* for trader i and adjustment stage a is a sequence of functions $\{e_{at}^i\}_{t=2}^{\infty}$ with, for each t ,

$$e_{0t}^i: X_0^{i(t-1)} \times S^i \rightarrow \Phi$$

and

$$e_{at}^i: X_a^{i(t-1)} \times (S^i \times \Delta^a) \rightarrow \Phi \quad \text{for each } a > 0.$$

The entire array of estimation procedures $\{e_{at}^i\}_{i=1}^N, a=0, \infty, t=2, \infty$ is called an *estimation scheme*.

2.6. Remarks. The estimator e_{at}^i associates with a sample $(s_{\tau}^i; p_{1\tau}, \dots, p_{a\tau}; \sigma_{\tau})_{\tau=1}^{t-1}$ a function on $S^i \times \Delta^a$ to Φ which acts as a conditional expectation. Hence e_{at}^i might be constructed as an estimator of the joint distribution on $S^i \times \Delta^a \times \Sigma$ which is then conditioned on $S^i \times \Delta^a$. However, the estimation scheme constructed in Sections 3 and 4 will estimate the "conditional distributions" directly, without regard to the joint distribution from which they might be derived.

The definition of e_{at}^i also embodies the assumption that trader i never observes s_j , even after the future state σ is realized. It might appear from the definitions in Section 2.3 that learning the conditional expectation ϕ_a^i requires learning f_x for each $\alpha \leq a$, which is impossible since $(s_j)_{j \neq i}$ is unobservable. However it will be shown in Sections 3 and 4 that the conditional expectations can be estimated directly.

Given a stochastic environment and an estimation scheme we can now describe how the data for estimation are generated. Essentially all that needs to be specified is that in each sampling period, a pair (s, σ) is drawn independently from the distribution π . Then at each stage of the adjustment process, a traders' expectations are determined by his estimation procedure and the data generated in previous sampling periods.

2.7. The Generation of the Data

Let $(\pi, (v^i)_i)$ be a stochastic environment and let $\{e_{at}^i\}_{t,a,i}$ be an estimation scheme. In each sample period t , a pair (s, σ) is drawn independently from the distribution π on $S \times \Sigma$. For each a , the first period a^{th} stage price p_{a1} is chosen arbitrarily in \mathcal{A} . For each $t > 1$, the first stage price p_{1t} is defined by

$$p_{1t} = w((\omega_t^i)_i, (e_{0t}^i[(s_\tau^i, \sigma_\tau)_{\tau=1}^{t-1}, s_t^i])_i).$$

Similarly for each $t > 1$ and $a > 1$,

$$p_{at} = w((\omega_t^i)_i, (e_{(a-1)t}^i[(s_\tau^i; p_{1\tau}, \dots, p_{(a-1)\tau}, \sigma_\tau)_{\tau=1}^{t-1}, (s_t^i; p_{1t}, \dots, p_{(a-1)t})])_i).$$

These definitions determine a probability distribution π_0^∞ on the infinite sample space $\prod_{t=1}^\infty (S \times \Sigma)$ and, for each $a > 0$, a distribution π_a^∞ on $\prod_{t=1}^\infty (S \times \mathcal{A}^a \times \Sigma)$. Of course, π_0^∞ is simply the product distribution induced by π . For each $a > 0$, the distribution π_a^∞ depends on π , on $(v^i)_i$ (through the definition of w), the arbitrarily selected prices $p_{11}, \dots, p_{(a-1)1}$, and the estimation procedures $\{e_{at}^i\}_{t=1}^\infty$ for each i and each $a < a$.

Given a stochastic environment $(\pi, (v^i)_i)$, an estimation scheme $\{e_{at}^i\}_{t,a,i}$ is said to be *successful almost surely* (a.s.) if for each i , each a , and almost every infinite sample $(s_t; p_{1t}, \dots, p_{at}; \sigma_t)_{t=1}^\infty$,

$$\lim_{t \rightarrow \infty} \|e_{at}^i[(s_\tau^i; p_{1\tau}, \dots, p_{a\tau}; \sigma_\tau)_{\tau=1}^{t-1}, (s_t^i; p_{1t}, \dots, p_{at})] - \phi_a^i(s_t)\| = 0, \quad (*)$$

where “almost every” refers to the distribution π_a^∞ , and $\|\phi - \phi'\| = \sup\{|\phi(\sigma) - \phi'(\sigma)| : \sigma \in \Sigma\}$ for each $\phi, \phi' \in \Sigma$. Recall that $\phi_a^i(\cdot)$ denotes trader i 's “rational expectations” at stage a .

2.8. Remarks. The above definitions do not require the adjustment process to terminate at any stage in any sample period. Loosely speaking, we are permitting time to be doubly infinite, with an infinite number of adjustment stages occurring within each sample period. It is apparent from the definition of the dynamic equilibrium f_A in Section 2.3 that $A < N^\# S$ for every stochastic environment in \mathcal{S} . This follows from the fact that the conditional expectations $\{\phi_a^i\}_{a \geq 0}$ are obtained by conditioning on successively finer partitions of the finite set S . If $f_{a+1} = f_a$, no trader's partition is refined by observing p_{a+1} , so $\phi_{a+1}^i = \phi_a^i$ for all i , and $f_{\alpha} = f_a$ for all $\alpha \geq a$. However, we have placed no restrictions on the estimators $\{e_{at}^i\}_{a \geq 0}$ to ensure that the estimated conditional distributions share this successive refinement property. Hence there is no guarantee that in any sample period t there will exist some adjustment stage a with $p_{a+1} = p_a$. It might be natural to require the adjustment process to terminate with the first stage in sample period one, the second stage in period two, and so on. That is,

traders stop revising their (estimated) conditional expectations when the adjustment process reaches a stage they have not previously encountered. However this convention would considerably complicate the notation since the first sample point for the estimation of ϕ_a^i would not occur until sample period a . None of the succeeding results would be otherwise affected by imposing this convention.

As noted above, the estimation scheme influences the true price distribution, which makes the estimation problem conceptually more complex than conventional econometric problems. Prices at stage a are influenced by the estimated expectations at earlier stages, but are *not* influenced by the stage a estimation procedures $\{e_{at}^i\}_{t,i}$, which estimate expectations conditioned on stage a prices. The recursive nature of our dynamic model enables us to break the estimation problem into stages so that, at each stage, estimation itself does not influence the data.

3. AN EXAMPLE

This section is devoted to a simple example illustrating the estimation scheme which is proved to be successful a.s. in Section 4. The set of stochastic environments is drawn from the traditional example of rational expectations theory. There are two traders, $i = 1, 2$, and two future states, $\Sigma = \{\sigma_\alpha, \sigma_\beta\}$. Trader 1's private information set contains two signals, $S^1 = \{s_a, s_b\}$, and trader 2 receives no private information, that is, S^2 is a one-point set, so we will mildly abuse the notation of Section 2 above to let $S = S^1$. We assume, in this example, that each trader's endowment ω^i is nonstochastic. As above, each trader i has a state-dependent utility function $v^i: R_+^M \times \Sigma \rightarrow R$. The space of stochastic environments is thus the set \mathcal{S} of pairs $(\pi, (\omega^i, v^i)_{i=1}^2)$, where π is a probability distribution on the four-point set $S \times \Sigma$. Recall that $\mathcal{A} = \{p \in R^M: p_j > 0 \text{ for all } j \text{ and } \sum_j p_j = 1\}$, and that Φ is the set of probability distributions on the two point set Σ . As in Section 2, given $(\omega^i, v^i)_{i=1}^2$, let $w: \Phi^2 \rightarrow \mathcal{A}$ be a function such that for each $(\phi^1, \phi^2) \in \Phi^2$, $w(\phi^1, \phi^2)$ is a Walrasian equilibrium price for the two person economy $(\omega^i, \phi^i(\sigma_\alpha) v^i(\cdot, \sigma_\alpha) + \phi^i(\sigma_\beta) v^i(\cdot, \sigma_\beta))_{i=1}^2$.

Since S is a two-point set, the dynamic process described in Section 2.3 above, when traders form the correct conditional expectations, will require only two iterations. Given π , let $\phi_0^1(s_a)$ and $\phi_0^1(s_b) \in \Phi$ denote the conditional distribution on Σ given s_a and s_b , respectively. Also, let ϕ_0^2 denote the unconditional distribution on Σ determined by π . Initially, trader 1 observes s and trader 2 observes nothing, so the initial expectations of the two traders are given by $\phi_0^1(s)$ and ϕ_0^2 , respectively. The first stage price will be given by $w(\phi_0^1(s), \phi_0^2)$, so define $f_1: S \rightarrow \mathcal{A}$ by $f_1(s) = w(\phi_0^1(s), \phi_0^2)$. At the second stage, trader 2 conditions his expectations on the first stage price, so

let $\phi_1^2(s)$ denote the conditional distribution on Σ given $f_1(s)$, for each $s \in S$. Trader 1 already knows s , so let $\phi_1^1(\cdot) = \phi_0^1(\cdot)$. Finally, define $f_2: S \rightarrow \mathcal{A}$ by $f_2(s) = w(\phi_1^1(s), \phi_1^2(s))$, which is the second and final stage price function.

If $f_1(s_a) \neq f_1(s_b)$, then trader 2 infers the signal, so both traders know the signal in equilibrium. If $f_1(s_a) = f_1(s_b)$, then trader 2 remains uninformed and $f_2(\cdot) = f_1(\cdot)$. Note that, according to the conventional definition, a rational expectations equilibrium fails to exist if $f_1(s_a) \neq f_1(s_b)$ and $f_2(s_a) = f_2(s_b)$. However, the price function f_2 is a *dynamic equilibrium* because we assume that trader 2 does not forget the information he inferred from f_1 even when that information is not also revealed by f_2 .

Now suppose that trader 1 knows only $S \times \Sigma$ and (w^1, v^1) , and trader 2 knows only (w^2, v^2) . Neither trader knows π or the characteristics of the other trader. Trader 1 needs to learn $\phi_0^1(\cdot)$, and trader 2 needs to learn ϕ_0^2 and $\phi_1^2(\cdot)$. (Note that $\phi_1^2(s)$ can be expressed as a function of the first stage price $f_1(s)$.) In each sample period t , (s_t, σ_t) is drawn independently according to π . The traders estimate their initial expectations, and these estimates determine the first stage price $f_{1t}(s_t)$. Since the initial expectations are not conditioned on any endogenous variables, their estimation is statistically trivial. For each t , define $e_{0t}^2: \Sigma^{t-1} \rightarrow \Phi$ as the frequency distribution determined by each sample $(\sigma_1, \dots, \sigma_{t-1})$ of previously drawn states. Define $e_{0t}^1: (S \times \Sigma)^{t-1} \times S \rightarrow \Phi$ by setting $e_{0t}^1(s_1, \sigma_1, \dots, s_{t-1}, \sigma_{t-1}; s_t)(\sigma)$, the estimated probability of σ , equal to the conditional frequency $\# \{ \tau < t: s_\tau = s_t \text{ and } \sigma_\tau = \sigma \} / \# \{ \tau < t: s_\tau = s_t \}$. These estimators clearly converge a.s. to ϕ_0^2 and $\phi_0^1(\cdot)$, respectively.

Trader 2's estimation of $\phi_1^2(\cdot)$, the conditional distribution on Σ given the first-stage price, is much more involved. In each sample period t , the first-stage price, which we will denote p_t , is generated as

$$p_t = w(e_{0t}^1(s_1, \sigma_1, \dots, s_{t-1}, \sigma_{t-1}; s_t), e_{0t}^2(\sigma_1, \dots, \sigma_{t-1})).$$

Hence the true conditional distribution of σ given p_t is changing over time, and of course p_t will generally differ from the "correct" price, $f_1(s_t)$. However, the arguments of w approximate $(\phi_0^1(s_t), \phi_0^2)$ as $t \rightarrow \infty$, so if w is well behaved, p_t will be near $f_1(s_t)$ for large t . We will assume that the function w , which is in the nature of a selection from the Walrasian equilibrium correspondence, is Lipschitzian near $(\phi_0^1(s_a), \phi_0^2)$ and $(\phi_0^1(s_b), \phi_0^2)$. This is essentially an assumption on the class of stochastic environments, and is discussed at more length in Sections 4.2 and 5. For example, w can be chosen to satisfy the assumption if the two exchange economies determined by the expectations $(\phi_0^1(s_a), \phi_0^2)$ and $(\phi_0^1(s_b), \phi_0^2)$, respectively, are "regular" in the sense of Debreu [7]. Under this assumption, p_t approximates $f_1(s_t)$ sufficiently closely to permit the construction of a convergent estimator of $\phi_1^2(\cdot)$.

In sample period t , trader 2 observes the first stage price p_t and must estimate the conditional distribution of σ_t from the data $(p_1, \sigma_1, \dots, p_{t-1}, \sigma_{t-1})$. The construction of a conditional frequency requires a criterion for deciding whether a previous price p_τ is close enough to p_t to include the succeeding state σ_τ . The law of the iterated logarithm implies the existence of a constant $K > 0$ such that

$$\limsup_{t \rightarrow \infty} \|(e_{0t}^1[(s_\tau, \sigma_\tau)_{\tau=1}^{t-1}, s_t], e_{0t}^2[(\sigma_\tau)_{\tau=1}^{t-1}]) - (\phi_0^1(s_t), \phi_0^2)\| / K[(\log \log t)/t]^{1/2} < 1 \quad \text{a.s.}$$

Therefore, by the Lipschitzian property of w , there is some $\lambda > 0$ such that, for almost every sample sequence $(s_t, \sigma_t)_{t=1}^\infty$,

$$\|p_t - f_1(s_t)\| < \lambda K[(\log \log t)/t]^{1/2}$$

for large t . For each t , let $\varepsilon_t = [(\log \log t)/t]^{1/3}$. Note that $\varepsilon_t \rightarrow 0$ and, for any λ and K , $\varepsilon_t > \lambda K[(\log \log t)/t]^{1/2}$ for large t . Given $((p_\tau, \sigma_\tau)_{\tau=1}^{t-1}, p_t)$, construct the estimated conditional probability of a state σ given p_t as a conditional frequency, counting the previous observation of σ_τ if and only if $\|p_\tau - p_t\| < \varepsilon_\tau$. More explicitly,

$$\begin{aligned} & e_{1t}^2[(p_\tau, \sigma_\tau)_{\tau=1}^{t-1}, p_t](\sigma) \\ &= \# \{ \tau < t : \|p_\tau - p_t\| < \varepsilon_\tau \text{ and } \sigma_\tau = \sigma \} / \# \{ \tau < t : \|p_\tau - p_t\| < \varepsilon_\tau \}. \end{aligned}$$

If the denominator is zero, the estimate is defined arbitrarily. The fact that ε_τ does not depend on λ or K is essential, since these constants are not known by the traders. For large τ and $t > \tau$, $\|p_\tau - f_1(s_\tau)\| < \varepsilon_\tau/2$ and $\|p_t - f_1(s_t)\| < \varepsilon_t/2 < \varepsilon_\tau/2$ so $\|p_\tau - p_t\| < \varepsilon_\tau$ if $f_1(s_\tau) = f_1(s_t)$. Also, $\|f_1(s_\tau) - f_1(s_t)\| \leq \|p_\tau - f_1(s_\tau)\| + \|p_t - f_1(s_t)\| + \|p_t - p_\tau\|$, so for large t , τ , if $\|p_t - p_\tau\| < \varepsilon_\tau$ then $\|f_1(s_\tau) - f_1(s_t)\| < 2\varepsilon_\tau$. Since S is a finite set and $\varepsilon_\tau \rightarrow 0$, the last inequality will be satisfied for large τ , t only if $f_1(s_\tau) = f_1(s_t)$. Hence, for large τ , t , $\|p_\tau - p_t\| < \varepsilon_\tau$ if and only if $f_1(s_\tau) = f_1(s_t)$. Therefore this estimator behaves asymptotically like the true conditional frequency of σ given $f_1(s_t)$, and thus converges to $\phi_1^2(s_t)$ a.s.

Since the estimator of $\phi_1^2(\cdot)$ is asymptotically a conditional frequency, it converges at the same rate as the estimators of $(\phi_0^1(\cdot), \phi_0^2)$. Therefore, in the general model where S is an arbitrary finite set and the price adjustment process requires more than two stages, the estimators of $\phi_2^i(\cdot)$, $\phi_3^i(\cdot)$; etc., can be constructed in the same way. This is done in the next section.

4. THE MAIN RESULT

This section generalizes the estimation scheme constructed in the previous example. We first formalize a general version of the Lipschitz condition used above.

4.1. *Axiom L*

Let $(\pi, (v^i)_i) \in \mathcal{S}$ and let $B(\pi, (v^i)_i) = \{((\omega^i)_i, (\phi^i)_i) \in (R_{++}^M)^N \times \Phi^N: \text{for some } s = (z^i, \omega^i)_i \in S \text{ and some } a \geq 0, (\omega^i)_i = (\omega^i)_i \text{ and } (\phi_a^i(s)) = (\phi^i)_i\}$. A stochastic environment $(\pi, (v^i)_i)$ satisfies *Axiom L* if for each $((\omega^i)_i, (\phi^i)_i) \in B(\pi, (v^i)_i)$ there is a neighborhood \mathcal{N} of $(\phi^i)_i$ in Φ^N and a constant $\lambda > 0$ with

$$\|w((\omega^i)_i, (\phi^i)_i) - w((\omega^i)_i, (\phi'^i)_i)\| \leq \lambda \|(\phi^i)_i - (\phi'^i)_i\|$$

for all $(\phi'^i)_i \in \mathcal{N}$, where $\|(\phi^i)_i - (\phi'^i)_i\| = \max \{|\phi^i(\sigma) - \phi'^i(\sigma)|: \sigma \in \Sigma \text{ and } 1 \leq i \leq N\}$.

Let \mathcal{S}_L denote the set of stochastic environments satisfying *Axiom L*.

4.2. *Remarks.* The interpretation of *Axiom L* may be clarified somewhat by the following example. Suppose that for each i and each $\sigma \in \Sigma$, $v^i(\cdot, \sigma)$ has the form

$$v^i(x, \sigma) = \alpha^i(\sigma)x + \sum_{j=1}^M \beta_j^i(\sigma) \ln x_j,$$

where $\alpha^i(\sigma) \in R_+^M$ and $\beta_j^i(\sigma) > 0$ for each j . Then $v^i(\cdot, \sigma)$ belongs to a convex family of utility functions which yield excess demand functions with the gross substitutes property (we can overlook the fact that these utility functions are not defined on the boundary of R_+^M , since they lead to interior equilibria). Hence for every $(\phi^i)_i \in \Phi^N$, the nonstochastic exchange environment $((\omega^i)_i, (\Sigma_\sigma \phi^i(\sigma) v^i(\cdot, \sigma))_i)$ has a unique equilibrium, and the function $w: \Omega \times \Phi^N \rightarrow \mathcal{A}$ is uniquely defined. Also w is C^∞ , since the gross substitutes property implies that the aggregate excess demand function has a nonsingular derivative. Hence $(\pi, (v^i)_i)$ satisfies *Axiom L* for any π . This example indicates that *Axiom L* can be interpreted as a regularity condition on the nonstochastic environments determined by a stochastic environment over the course of the adjustment process. More specifically, let $(\pi, (v^i)_i) \in \mathcal{S}$ and for each $s \in S$ and each $a \geq 0$, let $u_a^i(s)$ denote the utility function $\Sigma_\sigma \phi_a^i(s)(\sigma) v^i(\cdot, \sigma)$. Suppose that for each a, s , the nonstochastic environment $((\omega^i)_i, (u_a^i(s))_i)$, where $(\omega^i)_i$ is specified by s , is smooth and regular in the sense that the aggregate excess demand function has a nonsingular derivative at each of the Walrasian equilibrium prices. It may be worthwhile to emphasize that prices enter this excess demand function only in the usual way, through each trader's budget constraint.

Then it would be natural to expect the function w to be differentiable on a neighborhood of $((\omega^i)_i, (\phi_a^i(s))_i)$ for each a and s . This amounts to assuming that w is as smooth a selection from the Walrasian equilibrium correspondence as this correspondence permits. We will show in Section 5 below that the regularity of $((\omega^i)_i, (u_a^i(s))_i)$ for each a and s is generic for reasonable topologies on the set of smooth stochastic environments.

The conditions mentioned in the preceding paragraph lead to a property somewhat stronger than Axiom L, namely the continuous differentiability of $w((\omega^i)_i, \cdot)$ on \mathcal{N} . Axiom L is more in the nature of a Lipschitz condition, although slightly weaker than a conventional Lipschitz condition in that the inequality is not required for all pairs of N -tuples $(\phi''^i)_i, (\phi'''^i)_i \in \mathcal{N}$.

4.3. PROPOSITION. *There exists an estimation scheme which is successful a.s. for every stochastic environment in \mathcal{S}_L .*

Proof. For each i, t define e_{0t}^i to be a conditional frequency distribution. That is, for any event $E \subset \Sigma$,

$$e_{0t}^i[\{(s_\tau^i, \sigma_\tau)\}_{\tau=1}^{t-1}, s_t^i](E) = \# \{(s_\tau^i, \sigma_\tau): \tau < t, s_\tau^i = s_t^i, \sigma_\tau \in E\} / \# \{(s_\tau^i, \sigma_\tau): \tau < t, \text{ and } s_\tau^i = s_t^i\}^4$$

For $a > 0$, we define e_{at}^i as follows: For each t , let $\varepsilon_t = [(\log \log t)/t]^{1/3}$ and for any event $E \subset \Sigma$, let

$$e_{at}^i[\{(s_\tau^i; p_{1\tau}, \dots, p_{a\tau}; \sigma_\tau)\}_{\tau=1}^{t-1}, (s_t^i; p_{1t}, \dots, p_{at})](E) = \# \{(s_\tau^i; p_{1\tau}, \dots, p_{a\tau}; \sigma_\tau): \tau < t, s_\tau^i = s_t^i, \|p_{a\tau} - p_{at}\| < \varepsilon_t$$

for each $1 \leq \alpha \leq a$, and $\sigma_\tau \in E\} / \# \{(s_\tau^i; p_{1\tau}, \dots, p_{a\tau}; \sigma_\tau): \tau < t, s_\tau^i = s_t^i, \text{ and } \|p_{a\tau} - p_{at}\| < \varepsilon_t \text{ for each } 1 \leq \alpha \leq a\}$. This defines an estimation scheme $\{e_{at}^i\}_{t,a,i}$.

Let $(\pi, (v^i)_i) \in \mathcal{S}_L$ and for each i let π_0^i be the marginal distribution on $S^i \times \Sigma$. The law of the iterated logarithm (e.g., [6, p. 64]) implies that for each $(s^i, \sigma) \in S^i \times \Sigma$,

$$\limsup_{t \rightarrow \infty} \# \{(s_\tau, \sigma_\tau): \tau \leq t \text{ and } (s_\tau^i, \sigma_\tau) = (s^i, \sigma)\} / t - \pi_0^i(s^i, \sigma) / K_0 [(\log \log t)/t]^{1/2} < 1$$

a.s. with respect to π_0^∞ , where $K_0 > [2\pi_0^i(s^i, \sigma)(1 - \pi_0^i(s^i, \sigma))]^{1/2}$ for all (s^i, σ) . It follows that for some K'_0 depending on π ,

$$(i) \quad \pi_0^\infty \{(s_t, \sigma_t)_{t=1}^\infty : \text{there is some } T \text{ such that for all } i \text{ and all } t > T, \|e_{0t}^i(\{(s_\tau^i, \sigma_\tau)\}_{\tau=1}^{t-1}, s_t^i) - \phi_0^i(s_t)\| < K'_0 [(\log \log t)/t]^{1/2}\} = 1.$$

⁴ We maintain the convention that the estimated conditional distribution is arbitrary when the denominator is zero.

In particular, $\{e_{0t}^i\}_{t=2}^\infty$ satisfies the convergence condition 2.7(*) a.s. with respect to π_0^∞ for each i . Axiom L, together with (i), implies that

(ii) $\pi_1^\infty\{(s_t, p_{1t}, \sigma_t)_{t=1}^\infty: \text{there is some } T \text{ such that for all } t > T,$

$$\|p_{1t} - f_1(s_t)\| < \varepsilon_t/2\} = 1.$$

Hence

(iii) $\pi_1^\infty\{(s_t, p_{1t}, \sigma_t)_{t=1}^\infty: \text{there is some } T \text{ such that if } \tau, t > T, \text{ then}$

$$\|p_{1\tau} - p_{1t}\| < \varepsilon_\tau \quad \text{if and only if} \quad f_1(s_\tau) = f_1(s_t)\} = 1.$$

For each i , let π_1^i denote the distribution on $S^i \times \mathcal{A} \times \Sigma$ determined by π and f_1 . That is,

$$\pi_1^i(s^i, p_1, \sigma) = \Sigma\{\pi(s', \sigma): s'^i = s^i \text{ and } f(s') = p_1\}.$$

By the law of the iterated logarithm, for each $\sigma \in \Sigma$,

$$(iv) \quad \limsup_{t \rightarrow \infty} |\#\{(s_\tau, p_{1\tau}, \sigma_\tau): \tau < t, s_\tau^i = s_t^i, f_1(s_\tau) = f_1(s_t), \\ \sigma_\tau = \sigma\}|/(t-1) - \pi_1^i(s_t^i, f_1(s_t), \sigma)|/K_0[(\log \log t)/t]^{1/2} < 1$$

a.s. with respect to π_1^∞ , since the probability of (iv) is actually given by π_0^∞ which is the marginal distribution of π_1^∞ on $\Pi_{t=1}^\infty(S \times \Sigma)$. Then (iii) and (iv) imply that for each $\sigma \in \Sigma$,

$$\limsup_{t \rightarrow \infty} |\#\{(s_\tau, p_{1\tau}, \sigma_\tau): \tau < t, s_\tau^i = s_t^i, \|p_{1\tau} - p_{1t}\| < \varepsilon_\tau \text{ and} \\ \sigma_\tau = \sigma\}|/(t-1) - \pi_1^i(s_t^i, f_1(s_t), \sigma)|/K_0[(\log \log t)/t]^{1/2} < 1$$

a.s. with respect to π_1^∞ . Hence there is some K'_1 with

(v) $\pi_1^\infty\{(s_t, p_{1t}, \sigma_t)_{t=1}^\infty: \text{there is some } T \text{ such that for all } i \text{ and all } t > T,$

$$\|e_{1t}^i(\{(s_\tau^i, p_{1\tau}, \sigma_\tau)_{\tau=1}^{t-1}, s_t^i, p_{1t}) - \phi_1^i(s_t)\| < K'_1[(\log \log t)/t]^{1/2}\} = 1.$$

In particular, $\{e_{1t}^i\}_{t=2}^\infty$ satisfies the convergence condition 2.7(*) a.s. with respect to π_1^∞ for each i . Since (v) is analogous to (i), one can again apply Axiom L to obtain an expression analogous to (ii) involving π_2^∞ , and so on to complete the proof.

4.4. Remarks. Axiom L plays a crucial role in the construction of the estimation scheme $\{e_{at}^i\}_{t,a,i}$. The law of the iterated logarithm determines a sequence ε_t converging to zero with the property that for almost every infinite sample $(s_t, \sigma_t)_{t \geq 1}$ there is some T with

$$\|e_{0t}^i((s_\tau^i, \sigma_\tau)_{\tau=1}^{t-1}, s_t^i) - \phi_0^i(s_t)\| < \varepsilon_t$$

for all $t > T$ and all i . Axiom L then implies the existence of some $\lambda > 0$ and some $T' \geq T$ with

$$\|p_{1t} - f_1(s_t)\| < \lambda \varepsilon_t \quad (1)$$

for all $t > T'$. This enables the stage 1 conditional expectations ϕ_1^i to be estimated in a straightforward manner. In the absence of Axiom L, the distance $\|p_{1t} - f_1(s_t)\|$ need not converge to zero, and if Axiom L is weakened to the continuity of $w((\omega^i)_i, \cdot)$ at $(\phi_0^i(s))$ for each s , then the distance $\|p_{1t} - f_1(s_t)\|$ may converge to zero arbitrarily slowly.

The inequality (1) depends on every trader i using the estimation procedure $\{e_{0t}^i\}_{t \geq 2}$. If some trader departs from this procedure, the estimation procedures $\{e_{1t}^i\}_{t \geq 2}$ may fail to converge.

5. THE GENERICITY OF AXIOM L

The purpose of this section is to motivate Axiom L by demonstrating a sense in which it can be regarded as a typical rather than an exceptional property of stochastic environments. We will do this by choosing a topology for stochastic environments and showing that the regularity property mentioned in Section 4.2 above is satisfied on an open and dense set of environments. Unfortunately there is no single most natural topology for finite-state stochastic environments. Some topologies fix the support of π , as we have done above, or at least fix the cardinality of the support, and topologize other characteristics of the environment in a natural way. This approach yields results such as the well-known theorem of Radner [11] that fully revealing rational expectations equilibria exist generically (see also [1]). However, if the topology is weakened to permit perturbations which add to the support of π a "nearby" point with small probability, the result is lost [9].

In order to avoid such ambiguities we will use a topology which does not fix the cardinality of the support of π . Unfortunately this requires abandoning the convenience of the fixed support $S \times \Sigma$ we have employed above. Several other natural modifications will also be made. The space of joint signals S becomes the product $\prod_i (Z^i \times R_{++}^M)$, where Z^i is no longer required to be finite. Since the set of future states Σ is used merely to parameterize the dependence of trader's utilities on the future, we will define Σ to be the space of N -tuples of utility functions themselves.

5.1. DEFINITIONS. Let U denote the space of utility functions $u: R_{++}^M \rightarrow R$ satisfying

(1) u is C^2 and for each $x \in R_{++}^M$, $Du(x) \in R_{++}^M$, and $D^2u(x)$ is negative definite; and

(2) for each $x \in R_{++}^M$, $\text{cl } \{x' \in R_{++}^M : u(x') \geq u(x)\} \subset R_{++}^M$, where cl denotes closure in R^M .

Let U have the topology of uniform C^2 convergence on compact subsets of R_{++}^M , and let $\Sigma = U^N$. For each i , define $v^i: R_{++}^M \times \Sigma \rightarrow R$ by $v^i(x, \sigma) = v^i(x; u^1, \dots, u^N) = u^i(x)$.

For each i , let Z^i be a separable metric space, and let $S^i = Z^i \times R_{++}^M$. Let $S = \prod_i S^i$, and let $\mathcal{M}_F(S \times \Sigma)$ denote the set of probability measures on $S \times \Sigma$ with *finite support*. Each $\pi \in \mathcal{M}_F(S \times \Sigma)$, together with the N -tuple of state-dependent utilities $(v^i)_i$ defined above, determines a stochastic environment $(\pi, (v^i)_i)$ which conforms closely enough to the definition of a stochastic environment in Section 2.1, that all results in the preceding sections apply. The principal distinction is that the support of π is no longer equal to $S \times \Sigma$, which is now an infinite set, but this merely requires replacing the term “every (s, σ) ” with “almost every (s, σ) ” where necessary. Also, the utility functions v^i defined above do not satisfy assumption (A.1) of 2.1 since $v^i(\cdot, \sigma)$ is not defined on the boundary of R_{++}^M , but condition (2) above ensures that all equilibria are interior, so this difference can be ignored.

By identifying each $\pi \in \mathcal{M}_F(S \times \Sigma)$ with $(\pi, (v^i)_i)$ as defined above, we will refer to elements of $\mathcal{M}_F(S \times \Sigma)$ as stochastic environments. Let H denote the space of all compact subsets of $S \times \Sigma$ with the Hausdorff metric, and define $h: \mathcal{M}_F(S \times \Sigma) \rightarrow H$ by $h(\pi) = \text{supp } \pi$, where $\text{supp } \pi$ denotes the support of π . The topology on $\mathcal{M}_F(S \times \Sigma)$, which we need not specify completely, is assumed to satisfy

(3) h is a continuous open map.

5.2. *Remarks.* Statement (3) requires the topology on $\mathcal{M}_F(S \times \Sigma)$ to be weak enough to allow small perturbations of $\text{supp } \pi$, and strong enough so that small perturbations of π lead to small perturbations of $\text{supp } \pi$. That is, h behaves like a projection map. An example is the metric space obtained by defining the metric

$$d(\pi, \pi') = d_w(\pi, \pi') + d_H(\text{supp } \pi, \text{supp } \pi'),$$

where d_w is a metric for the topology of weak convergence and d_H is the Hausdorff metric. This topology has appeared in the core convergence literature, of course without the finite support requirement (e.g., [8, Theorem 2, p. 192]).

With the above definition of Σ , it is no longer reasonable to suppose that the i^{th} trader observes the complete description of the realized state $\sigma = (u^j)_j$ in each sample period. Rather, it becomes more natural to assume that in each sample period t , the i^{th} trader observes only the i^{th} coordinate $\sigma_t^i = u_t^i$. However, this modification requires only the obvious notational changes in the proof of Proposition 4.3.

We now give a precise definition of the regularity condition described in Section 4.2.

5.3. Regular Stochastic Environments

A (nonstochastic) exchange environment $(\omega^i, u^i)_{i=1}^N \in (R_{++}^M \times U)^N$ is regular if for each Walrasian equilibrium price p , the Jacobian of the aggregate demand function is nonzero at p [14].⁵ For each $\pi \in \mathcal{M}_F(S \times \Sigma)$, let π_S denote the marginal distribution on S . For each $s \in \text{supp } \pi_S$, let $\Phi(s) = \{(\phi^i)_i \in \Phi^N: \text{for each } i \text{ there is some subset } E^i \subset S \text{ satisfying}$

- (1) $s \in E^i$ and for each $s' \in E^i$, $s'^i = s^i$; and
- (2) for each event $F \subset \Sigma$, $\phi^i(F) = \pi(E^i \times F) / \pi(E^i \times \Sigma)\}$.

A stochastic environment $\pi \in \mathcal{M}_F(S \times \Sigma)$ is regular if for each $s = (z^i, \omega^i)_i \in \text{supp } \pi_S$ and each $(\phi^i)_i \in \Phi(s)$, the exchange environment

$$(\omega^i, \sum_{\sigma} \phi^i(\sigma) v^i(\cdot, \sigma))_i$$

is regular. Let \mathcal{S}_R denote the set of regular stochastic environments in $\mathcal{M}_F(S \times \Sigma)$.

5.4. Remarks. Given a stochastic environment π , the N -tuple of conditional distributions $(\phi_a^i(s))_{i=1}^N$, defined in Section 2.3 above, is in $\Phi(s)$ for each a . This fact implies directly that if π is regular, the function w can be chosen so that Axiom L is satisfied.

5.5. PROPOSITION. \mathcal{S}_R contains an open and dense subset of $\mathcal{M}_F(S \times \Sigma)$.

Proof. We will construct an open dense set $G \subset \mathcal{S}_R$. Let $G_1 = \{\pi \in \mathcal{M}_F(S \times \Sigma):$

- (1) if $(s, \sigma) = ((z^i, \omega^i)_i, \sigma)$ and $(s', \sigma') = ((z'^i, \omega'^i)_i, \sigma')$ are distinct elements of $\text{supp } \pi$ then $\omega^i \neq \omega'^i$ for all i ; and
- (2) for each $((z^i, \omega^i), (u^i)_i) \in \text{supp } \pi$, the exchange environment $(\omega^i, u^i)_i$ is regular.

It follows directly from [14, Theorem 2] that G_1 is dense. Also, (1) implies that for each $(s, \sigma) \in \text{supp } \pi$, $\Phi(s)$ contains only the N -tuple $(\phi^i)_i$, where ϕ^i is the degenerate distribution at σ for each i . Hence the regularity of π follows from (2).

Let $\pi \in G_1$ and for each $(s, \sigma) = (s, (u^i)_i) \in \text{supp } \pi$, let $\mathcal{N}_S(s, \sigma) \times (\Pi_i \mathcal{N}_U^i(s, \sigma))$ be a neighborhood of $(s, (u^i)_i)$ in $S \times U^N = S \times \Sigma$ with the properties that

- (3) $\mathcal{N}_U^i(s, \sigma)$ is convex for each i ;

⁵ That is, the matrix representing the derivative of the aggregate demand for the first $M-1$ commodities with respect to the first $M-1$ prices has a nonzero determinant. This definition is equivalent to [14, Definition, p. 7].

(4) for each $(u^i)_i \in \Pi_i \mathcal{N}_U^i(s, \sigma)$ and each $(z'^i, \omega'^i)_i \in \mathcal{N}_S(s, \sigma)$, $(\omega'^i, u^i)_i$ is a regular exchange environment; and

(5) if $(s, \sigma) \neq (s', \sigma')$, then for each $(z_1^i, \omega_1^i)_i \in \mathcal{N}_S(s, \sigma)$, $(z_2^i, \omega_2^i)_i \in \mathcal{N}_S(s', \sigma')$, $\omega_1^i \neq \omega_2^i$ for all i .

Let $\pi' \in \mathcal{M}_F(S \times \Sigma)$ with $\text{supp } \pi' \subset \bigcup_{(s, \sigma) \in \text{supp } \pi} \mathcal{N}_S(s, \sigma) \times (\Pi_i \mathcal{N}_U^i(s, \sigma))$. Then (3) and (5) imply that for each $(s', \sigma') \in \text{supp } \pi'$ and each $(\phi^i)_i \in \Phi(s')$, there is some $(s^*, \sigma^*) \in \text{supp } \pi$ such that $s' \in \mathcal{N}_S(s^*, \sigma^*)$, and the utility function $\Sigma_\sigma \phi^i(\sigma) v^i(\cdot, \sigma) \in \mathcal{N}_U^i(s^*, \sigma^*)$ for each i . The regularity of π' now follows from (4).

Hence each $\pi \in G_1$ has an open neighborhood in \mathcal{S}_R . The proof is completed by defining G as the union of these neighborhoods.

5.6. Remarks. Topologies which satisfy (3) of 5.1 may be considered too weak in that small perturbations in $\text{supp } \pi_S$ do not preserve the exogenous information structure. In particular they can lead to drastic changes in the initial conditional distributions $(\phi_0^i)_i$. The simplest way to counter this effect is to fix $\text{supp } \pi_S$ and to use the relative topology on the set $R = \{\pi' \in \mathcal{M}_F(S \times \Sigma) : \text{supp } \pi'_S = \text{supp } \pi_S\}$. This topology leads directly to the generic existence of fully revealing rational expectations equilibria. It seems likely that by further refining this relative topology by permitting only small perturbations in the probabilities of joint signals as well, that is, by requiring the continuity of the function $\pi' \rightarrow \max_{s \in \text{supp } \pi_S} \{|\pi'_S(s) - \pi_S(s)|\}$ on R , the analogue of Proposition 5.5 could be obtained.⁶

The principal reason for using topologies satisfying (3) is that they permit, in certain open regions of $\mathcal{M}_F(S \times \Sigma)$, dense sets of stochastic environments which have no rational expectations equilibria (when traders condition expectations on private information and equilibrium prices alone). This can be shown as follows. Let $\pi \in \mathcal{M}_F(S \times \Sigma)$. Since we are demonstrating density, we can choose π to satisfy 5.5 (1). That is, each trader's endowment identifies the future state. Let $(s, \sigma) = ((z^i, \omega^i)_i, (u^i)_i) \in \text{supp } \pi$ and let $\lambda = \pi(s, \sigma)$. We will assume that

the exchange environment $(\omega^i, u^i)_i$ is regular and has a *unique*
Walrasian equilibrium. (*)

Choose a nearby point $(s', \sigma') = ((z'^i, \omega'^i)_i, (u'^i)_i)$ and a small number $\varepsilon > 0$ with the properties that

(1) $\omega'^1 \neq \omega^1$ and ω'^1 is distinct from trader 1's endowment at all other points in $\text{supp } \pi$;

⁶ In a fully revealing rational expectations equilibrium the probabilities of the joint-signals in the support of π_S are irrelevant. Since these probabilities may influence the conditional distributions $(\phi_0^i(s))_i$ even if f_1 , say, is a 1-1 function, the topology must restrict the extent to which they can be perturbed.

(2) $\omega'^1 = \omega^i$ for all $i > 1$;

(3) the exchange environments

(a) $(\omega^i, u^i)_i$

(b) $(\omega'^i, u'^i)_i$

(c) $((\omega^1, u^1), (\omega^i, (\lambda(1-\varepsilon) + \varepsilon)^{-1}(\lambda(1-\varepsilon) u^i + \varepsilon u'^i)))_{i>1})$

(d) $((\omega'^1, u'^1), (\omega^i, (\lambda(1-\varepsilon) + \varepsilon)^{-1}(\lambda(1-\varepsilon) u^i + \varepsilon u'^i)))_{i>1})$

all have unique Walrasian equilibria (for (a) this is just (*)); and

(4) the equilibrium prices for (a) and (b) are identical, and the equilibrium prices for (c) and (d) are distinct.

Let π' be the stochastic environment described by

$$\begin{aligned}\pi'(s'', \sigma'') &= (1 - \varepsilon) \pi(s'', \sigma'') && \text{if } (s'', \sigma'') \in \text{supp } \pi; \\ &= \varepsilon && \text{if } (s'', \sigma'') = (s', \sigma'); \text{ and} \\ &= 0 && \text{otherwise.}\end{aligned}$$

Then conditions (1)–(4) directly contradict the existence of a rational expectations equilibrium. Condition (*), which makes possible the satisfaction of (3), confines the demonstration of density to open sets of stochastic environments which have such a point in their support.

Proposition 5.5 and the above remarks indicate that a large class of stochastic environments are sufficiently regular to permit w to be chosen to satisfy Axiom L. Furthermore, this class is not dependent, in its elements or the measure of its size, on the existence of rational expectations equilibria.

6. CONCLUDING REMARKS

Our main result, Proposition 4.3, strengthens the intuitive plausibility of expectations equilibrium theory. It asserts that traders who begin with no knowledge of the characteristics of other traders, and no knowledge of the probability distribution governing the uncertainty they face, can learn to form correct expectations from repeated observations of prices and realized states. Also, even though traders use prices to infer the private information of others, a trader's learning procedure does not require any data involving other traders' private signals. Such a result appears to be possible only in recursive dynamic models of expectations equilibrium, such as the one developed in [10] and used in the present paper. Under the conventional definition of price-conditional rational expectations equilibrium, the existence problem, and the discontinuity which causes it preclude any general convergence results.

Proposition 4.3 is stated as an existence result to emphasize that we have not attempted any qualitative comparisons among alternative successful estimation schemes. One especially troublesome feature of the scheme we have constructed is that the convergence of a trader's estimated expectations depends on the use of the same estimation procedure by other traders. This raises the possibility that convergence could be impaired if traders seek to somehow tailor their estimation procedures more closely to their own characteristics.

More work is also needed to enlarge the scope of the existence result. The most immediate question is whether or not the Lipschitz condition, Axiom L can be dropped. The extension to stochastic environments with an infinite joint-signal space S also remains problematic.

REFERENCES

1. B. ALLEN, "Equilibria in Which Prices Convey Information: The Finite Case," CARESS Working Paper No. 80-10, University of Pennsylvania, June 1980.
2. L. BLUME, M. BRAY, AND D. EASLEY, Introduction to the stability of rational expectations equilibrium, *J. Econ. Theory* **26** (1982), 313-317.
3. L. BLUME AND D. EASLEY, Learning to be rational, *J. Econ. Theory* **26** (1982), 340-351.
4. M. BRAY, Learning, estimation, and the stability of rational expectations, *J. Econ. Theory* **26** (1982), 318-339.
5. M. BRAY AND D. KREPS, "Rational Learning and Rational Expectations," Research Paper No. 616, Graduate School of Business, Stanford University.
6. L. BREIMAN, "Probability," Addison-Wesley, Reading, Mass. 1968.
7. G. DEBREU, Economies with a finite set of equilibria, *Econometrica* **38** (1970), 387-392.
8. W. HILDENBRAND, "Core and Equilibria of a Large Economy," Princeton University Press, Princeton, N.J., 1974.
9. J. JORDAN, Comments on "Rational expectations with price information," mimeo, 1976.
10. J. JORDAN, A dynamic model of expectations equilibrium, *J. Econ. Theory*, in press.
11. T. KOBAYASHI, "A Convergence Theorem on Rational Expectations Equilibrium with Price Information," IMSSS Working Paper No. 79, Stanford University, January, 1977.
12. R. RADNER, Rational expectations equilibrium: Generic existence and the information revealed by prices, *Econometrica* **47** (1979), 655-678.
13. S. REITER, On expectations equilibrium, mimeo, August 1976.
14. S. SMALE, Global analysis and economics IIA: Extension of a theorem of Debreu, *J. Math. Econ.* **1** (1974), 1-14.