2

Savage–Bayesian Models of Economics

Nicholas M. Kiefer and Yaw Nyarko

1 Introduction

The "state of the art" in learning models in economics is highly unsettled. On the one hand, we have the optimizing models in which learning occurs as a byproduct of maximizing expected utillty under uncertainty. This includes the work of Rothschild (1974) and Prescott (1972) among the earlier papers and more recently McLennan (1984), Mirman et al., (1984), Easley and Kiefer (1989), and Kiefer and Nyarko (1989), and many others. Since these models have as their motivation expected utility maximization, which in turn follows from the axioms of Savage and implies Bayesian updating behavior, we shall refer to them as Savage–Bayesian models. On the other hand, we have a large number of models in which agents follow various *ad hoc* learning rules. These are often easier to analyze (under the specific assumptions made) and are sometimes advertised as being more realistic as well. We shall call these "*ad hoc*" models since there is no appeal to a universal principle in their construction.

In this chapter we first review the foundational questions bearing on the choice of approach. This turns out to be easy since the debate has in all essential features been seen before in economics. We shall be considering primarily the foundational issues relevant to economic questions. Indeed, the debates between Bayesian and non-Bayesian appoaches in statistics have a long history. Lindley (1972) gives a very nice exposition of some of the issues in this debate of primary interest to statisticians.

In section 2 we review the axioms and the basic argument for the use of Savage–Bayesian models in economics. We illustrate this in section 3 with a simple economic model, that of a monopolist with incomplete information about some key parameters of the stochastic demand curve it faces. Section 4 discusses some common objections to the Savage–Bayesian approach: that it requires agents to form very complicated models, that it breaks down on "probability 0" events, and that it ignores the information and computation processing limitations of agents.

Section 5 discusses whether, in the context of an economic model with heterogeneous agents, Bayesian or non-Bayesians make more money. Section 6 addresses whether Bayesian or non-Bayesian are better at "learning." The answer to both questions is "it depends." The last section is made up of miscellaneous comments and concluding remarks comparing the Savage–Bayesian models with other models used in the economics literature today.

2 The Axiomatic Foundations and Expected Utility Maximization

2.1 The basic argument

One rarely sees an economic theory article in which a static behavioral model is analyzed without some notion of optimization. Indeed, one of the principles of economic theory is that static behavioral models should involve optimization. This was not always the case, as the famous Lester–Machlup debate on the usefulness of marginal utility theory illustrate most sharply. Machlup's view prevailed. Lester's objection that marginal utility theory was fanciful, difficult, and did not realistically model the behavior of agents has been dismissed and would now invite ridicule.

The modern theory of choice under uncertainty, the expected utility approach, similarly has widespread acceptance in economics and statistics. The axiom systems that can be used to produce a utility representation of preferences are quite compelling and widely accepted. The axioms that produce subjective probability distributions over events are nearly as compelling and also widely accepted. The key reference is of course Savage (1954). DeGroot (1970) provides an insightful textbook discussion of the construction of utility functions and probability distributions, giving an alternative but related set of axioms, and Pratt et al. (1964) provide a simple but completely rigorous axiomatic system leading to expected utility maximizing behavior in the finite case.

Early objections to expected utility theory – that it is too difficult or too unrealistic to have any hope of being useful as a behavioral model – have disappeared. Modern refinements, e.g. by Machina and Schmeidler (1992), modify the theory in increasingly technical directions. Now, it is possible for the contribution of this highly technical work to be regarded as making the theory more realistic.

The position we advocate is simply that in modelling economic agents we must first and foremost realize that such agents seek to choose a "best" or "optimal" action. Depending on the context this would be profit maximization or utility or preference maximization. Learning is not a goal or objective in its own right. Learning only occurs as a byproduct, reflecting the informational constraints faced by agents.

2.2 The axioms

The following is taken from Savage (1954). This is only a summary and we encourage the reader to consult Savage (1954) for a very enjoyable exposition. A *state* s is a complete description of all items over which there is uncertainty. The *world* or universe is the set of all states, S. An *event* is a collection of states. The *true state* is the state which pertains. There is a set of consequences, C. One should think of this as the set of all "final consumptions." An *act* f is a mapping from the set of states into the set of consequences, $f: S \to C$. In particular, an act specifies the consequence in each state. Let A denote the set of all acts.

Agents have preferences, $\leq$, over acts. The strict preference $<$ is defined in the usual manner: $f < g$ if $f \leq g$ but not $g \leq f$. The axioms placed on the preferences are as follows.

P1 f is complete (for all f and $g \in A$ either $f \leq g$ or $g \leq f$) and transitive (for all $f, g, h \in A$, $f \leq g$ and $g \leq h$ implies $f \leq h$).

P2 $\leq$ obeys the sure-thing principle. That is, let f, f', g, and g' be acts and let B be a subset of S such that

1 on $\sim B$ (the complement of B in S), $f = g$ and $f' = g'$;

2 on B, $f = f'$ and $g = g'$; and

3 $f \leq g$.

Then $f' \leq g'$. (In words: Suppose acts g and f agree outside of the set B, and g is (weakly) preferred to f. Modify f and g outside of B but ensure that they are still the same outside of B, and maintain their values on B. Then the modified g is (weakly) preferred to the modified f.)

We say the act g is weakly preferred to the act f *given the event* B if, when g and f are modified outside of B so that they are the same outside of B, then the modified g is weakly preferred to the modified f, regardless of how the modification outside of B is done. Under the sure-thing principle the manner of modification outside of B is irrelevant. An event B is said to be a *null* event if, for all acts f and g, $f \leq g$ given B.

P3 Let f and g be constant acts (i.e. independent of s) and let B be an event which is not null. Then $f \leq g$ given B if and only if $f \leq g$.

P4 Let f, f', g, and g' be constant acts such that $f' < f$ and $g' < g$ (so that we may think of f and g as "good" constant acts and f' and g' as bad constant acts). Let A and B be any two events. Define the act (f_A) to be equal to f on A and f' outside of A, and define g_A to be g on A and g' outside of A.

Define f_B and g_B analogously with B replacing A. Then $f_A \leq f_B$ implies $g_A \leq g_B$.

P5 There exists a pair of consequences or constant acts f and f' such that $f' < f$.

P6 *A continuity axiom*: Let g and h be two acts with $g < h$ and let f be any consequence (or constant act). Then there exists a partition of S such that if g or h is modified to become g' and h' with the modification only taking place on one element of the partition and is equal to f there with the other values being unchanged, then $g' < h'$. (This condition requires that the agent has a randomizing device, e.g. the toss of a coin; the agent is required to be able to distinguish consequences contingent on say 10 "heads" as opposed to on 11 "heads.")

P7 Let f and g be any two acts and let B be an event. Let $g(s)$ denote the constant act equal to the consequence $g(s)$ regardless of the state. Then $f \leq g(s)$ given B for all $s \in B$ implies $f \leq g$ given B.

The result of Savage (1954) is as follows.

Theorem (Savage) If $\leq$ obeys P1–P7, there exists a (*utility*) function $u: C \to \mathbb{R}$ and a (*prior*) probability measure P over S such that, for all acts f and g,

$$f \leq g \text{ if and only if } \int u[f(s)]P(\mathrm{d}s) \leq \int u[g(s)]P(\mathrm{d}s) \quad (2.1)$$

2.3 The Savage axioms imply Bayesian updating

Let u and P be the utility function and prior probability obtained from the axioms of Savage as explained above. Let $P(. \mid B)$ denote the probability P conditional on event B. Then an immediate corollary of the main existence result of (2.1) is the following. Let B be any event with $P(B) > 0$; then

$$f \leq g \text{ given } B \text{ if and only if } \int u(f)P(\mathrm{d}s \mid B) \leq \int u(g)P(\mathrm{d}s \mid B) \quad (2.2)$$

We may interpret the preferences "$\leq$ given B" to be the preferences an agent has given knowledge that the event B has occurred. The right-hand side of (2.2) requires the agent to evaluate acts by expected utility but using the conditional probability of the prior given knowledge of B. In particular, conditional on B the agent behaves as if it "updates" its prior by Bayes's rule (i.e. the conditional probability) and then proceeds to maximize expected utility with respect to the updated prior (which is the posterior distribution).

2.4 Where do the priors come from?

This question is widely raised as a criticism of the Bayesian approach. Answer: The priors come from the same place the utility function comes from! Note that from the axioms above we obtained the prior and utility function simultaneously given the preferences. Indeed, what is a prior anyway? We use prior probabilities to represent agents' beliefs about the relative likelihoods of the various states. But what are beliefs?

In particular, let E be the event

$$E \equiv \{\text{the next president of the USA will be a Democrat}\}$$

What is the meaning of the statement "an agent assigns probability 0.2 to the event E above"? We have chosen the event so that it does not allow a frequentist interpretation, i.e. one cannot interpret the above to mean that "in infinitely repeated trials the fraction of next presidents of the USA being Democrats is 0.2." The interpretation in our opinion necessarily has to do with preferences and the amounts the individual is willing to bet for or against the occurrence of the event E.

Indeed, consider a bet which gives the agent \$1 if the event E occurs and \$0 otherwise. Suppose the agent announces that \$$M$ is the maximum amount it is willing to pay to take part in the bet. If we suppose that we can approximate the agent's utility function of wealth by a linear function then we may set M equal to the (subjective) expected winnings. In particular, $M = 0.2$. Hence the probability an agent assigns on an event occurring is the maximum amount it is willing to bet on that event. Beliefs are obtained from preferences over the associated bets!

One of the surprising observations is that some economists who are perfectly comfortable with the use of utility functions in deterministic problems or problems with objective risk are uncomfortable with the use of prior probabilities because they are "arbitrary." Priors are merely as arbitrary as the utility functions themselves.

We therefore disagree with the assertions often made that priors are very different from utility functions and should somehow be treated differently. The two are obtained simultaneously.

3 An Example: The Monopolist Problem

An example which has seen much use in the literature and which serves nicely to illustrate the setting is that of a monopolist producing a good to meet a stochastic demand with unknown mean function. Suppose at each date t a monopolist must choose an output level q_t. This results in a market clearing price

$$P_t = \alpha - \beta q_t + \varepsilon_t \tag{2.3}$$

where $\{\varepsilon_t\}_{t=1}^{\infty}$ is an independent and identically distributed process with zero mean and finite variance. The monopolist does not know the parameters α and β of the demand curve. The monopolist does not observe the shock process $\varepsilon_{t'}$ but knows its distribution.

At date t the monopolist would have observed the date t history of prices and outputs $h_t \equiv \{(q_1, p_1), \ldots, (q_{t-1}, p_{t-1})\}$; the monopolist then chooses a date t output level q_t. This will result in a market-clearing price of p_t via (2.3). Let H_t be the set of all date t partial histories. A policy is a collection of functions $\pi = \{\pi_t\}_{t=1}^{\infty}$ such that, for each t, π_t: $H_t \rightarrow \mathbb{R}$ prescibes the date t output level as a function $q_t = \pi_t(h_t)$ of the date t partial history.

We model the monopolist as a Savage–Bayesian. An act for the monopolist is the choice of a policy π. A consequence is a sample path of prices and outputs, $\{q_t, p_t\}_{t=1}^{\infty}$. A state s is a parameter vector and sample path of shock terms, $s = ((\alpha, \beta), \{\varepsilon_t\}_{t=1}^{\infty})$. It should be clear that the act π results in a consequence as a function of the state. Indeed, we shall let $f_\pi(s)$ be the sample path of outputs and prices $\{q_t, p_t\}_{t=1}^{\infty}$ induced by the policy π when the parameter vector and sample path of shock terms are those specified by the state $s = ((\alpha, \beta), \{\varepsilon_t\}_{t=1}^{\infty})$.

The monopolist has preferences over acts induced by preferences over the consequences, $\{q_t, p_t\}_{t=1}^{\infty}$; the monopolist could for example be maximizing a discounted sum of period-by-period profits. Suppose those preferences obey the axioms P1–P7 stated earlier. Then from the Savage theorem there will exist a prior ϕ over states and a utility function $U(\{q_t, p_t\})$ over consequences such that the monopolist's preferences over acts, π, may be determined by the expected utility, $\int U[f_\pi(s)]\,\mathrm{d}\phi$. In particular; the monopolist will seek a policy which solves the problem $\max_\pi \int U[f_\pi(s)]\,\mathrm{d}\phi$.

Most of the literature on the monopolist problem studies the model where the utility function takes the special additively separable form $U(\{q_t, p_t\}) = \sum_{t=1}^{\infty} \delta^{t-1} r(q_t, p_t)$, where r is the within-period profit or some function of profit. Let us now assume this additively separable structure.

Recall that we assumed that the monopolist knows the distribution of shock terms. Hence any prior probability distribution μ over the parameter vector (α, β) induces a prior probability, denoted by ϕ_μ, over the set of states $s = \big((\alpha, \beta), \{\varepsilon_t\}_{t=1}^{\infty}\big)$ and vice versa. Recall also that $f_\pi(s)$ denotes the sample path of outputs and prices resulting from the policy π in the state $s = \big((\alpha, \beta), \{\varepsilon_t\}_{t=1}^{\infty}\big)$. A policy π and prior probability μ over (α, β) therefore results in an expected utility $V_\pi(\mu) \equiv \int U[f_\pi(s)]\,\mathrm{d}\phi_\mu$. If we let E denote the expectations operator over sample paths of outputs and prices, $f_\pi(s)$, induced by the measure ϕ_μ, then under the additively separable structure

$$V_\pi(\mu) = E\sum_{t=1}^{\infty} \delta^{t-1} r(q_t, p_t) \tag{2.4}$$

The preferences of the monopolist with initial prior μ can then equivalently be represented by $V_\pi(\mu)$, and the monopolist's problem is the optimization problem

$$\sup_\pi V_\pi(\mu) \equiv V(\mu) \tag{2.5}$$

The function V in (2.5) is the value function. (In the literature the measure μ over Θ is what is typically referred to as the prior, and not ϕ_μ.)

Let $\phi_\mu[\cdot|(q, p)]$ denote the probability ϕ_μ conditional on the observation at date 1 of prices $p_1 = p$ from output choice $q_1 = q$. Define $B(p, q, \mu) \equiv \text{marg}_\Theta\, \phi_\mu[\cdot|(p, q)]$ as the marginal distribution of $\phi_\mu[\cdot|(q, p)]$ on the parameter space Θ. $B(p, q, \mu)$ is the monopolist's posterior distribution over the parameter space after observing a price p from output level q and updating the initial prior μ. It represents the Bayesian updating rule to obtain the conditional probability of θ given observation (p, q). (Of course, from the discussion of section 2.3 this additive separable structure is *not* necessary for Bayesian updating; the latter follows from the axioms of Savage (1954).)

Using standard dynamic programming arguments we may show that

$$V(\mu) = \max_q \{Er(q, p) + \delta EV[B(q, p, \mu)]\} \tag{2.6}$$

where $p \equiv \alpha - \beta q + \varepsilon_1$ and E is the expectations operator over $p = \alpha - \beta q + \varepsilon_t$ induced by letting (α, β) have distribution μ and ε_1 its known distribution. With the above equation we see immediately the *trade-off* the monopolist faces; the choice of an optimal output level must trade current period return (the first term in (2.6)) against the expected future information value of that choice of output level (the second term in (2.6)). Note that "learning" is in no way treated as an objective of the agent. It is a byproduct of the infinite horizon expected utility maximization.

We will return to this example to illustrate a number of points below. Here we stress that the axioms leading to Savage–Bayesian behavior need no modification to apply to this dynamic case. The economist who finds these compelling in the static case need make no further assumptions to deal equally satisfactorily with dynamic models. We note in section 6 some of the literature on this "monopolist" problem.

4 Mis-specification, Probability 0 Sets, and Information Constraints

This section will address three common but unfounded complaints about the use of the Savage–Bayesian paradigm in economics.

A Agents form simple models. The world is so complicated that agents cannot conceive of the real world, and do not consider it a possibility.

B Bayesian updating breaks down if it is possible for probability 0 events to be observed.

C The information and calculation requirements of the Savage–Bayesian approach are so strong that the model is too unrealistic to be useful.

4.1 Agents form simple models

Objection A reflects a serious confusion and misunderstanding of the role of models. Models are not reality, and a model is not necessarily to be rejected because there is some aspect of reality not captured by the model. Dreze (1972) makes the point nicely in his Presidential Address to the Econometric Society. He argues that models play the same role in decision theory as in fashion, i.e. they provide a framework for experimenting with new ideas, and they provide a frame for showing off one's work to its best advantage.

Consider flipping a coin to observe whether it is "heads" or "tails." Suppose one uses a prior probability model that heads or tails occur with probability of $\frac{1}{2}$ each. Is this how the toss of a coin "really" occurs? Of course not! The toss depends upon such things as the force and direction in which it is tossed, the wind pressure and velocity, etc. "Reality" would require a knowledge of all these variables, and a respectable knowledge of the laws of physics. This would enable us to predict accurately the outcome of the toss of the coin. Since we lack such knowledge, we require a model. The probability model has become an almost universally unchallenged model of the outcome of the toss of a coin. Is it a realistic model? We leave that to the reader to judge!

Now, there is a Savage–Bayesian way of modeling agents who for some reason or the other entertain beliefs which do not include the "truth." We now study some examples, and we shall indicate that this is very much related to the issue to be discussed under objection B.

Consider the monopolist problem above. There we suppose that the monopolist has a prior belief over the vector $\theta = (\alpha, \beta) \in \mathbb{R}^2$ representing the unknown intercept and slope of the demand curve. Suppose that the monopolist has prior beliefs μ_0 over θ with support in some set S (say $S = (0, 1) \times (0, 1)$). Suppose that the "true" value of the parameter is some θ^* in $\mathbb{R}^2$ which does not lie in S (e.g. $\theta^* = (2, 2)$). Such a model has been studied by Nyarko (1991a) and in the language of that paper the agent (monopolist) is said to have a mis-specified model.

To further embellish this example suppose instead that the demand curve in the monopolist example is given by $p_t = \alpha - \beta q_t + \Gamma G(q_t, z_t) + \varepsilon_t$, where $G(q_t, z_t)$ may be some "complicated" and perhaps nonlinear function of output q_t and some other variable z_t. Suppose that the monopolist does not entertain for whatever reason the possibility of G influencing the demand. This is of course equivalent to saying that the agent has beliefs over the tuple (α, β, Γ) with support in some subset of $\mathbb{R}^2 \times \{0\}$ while the true vector may be some $(\alpha^*, \beta^*, \Gamma^*)$ with Γ^* non-zero. Again such an agent has a mis-specified model.

Are such priors inconsistent with the Savage–Bayesian approach? The answer of course is no. Recall that we argued in section 2 that priors and utility functions are obtained from agents having preferences over lotteries. If agents never make bets in favor of some event E occurring, regardless of the odds, then such an agent assigns prior probability 0 to event E. In the language of Savage (1954), the agent considers such events to be null. An outside observer may know that the event E will occur and so can conclude that the agent has a mis-specified model. However, as long as the agent's preferences over lotteries obey the axioms of Savage, that agent will necessarily have a prior and utility function and maximize expected utility. Mis-specified models in this sense are not ruled out by the Savage axioms.

4.2 Bayes's rule breaks down with probability 0 observations

The discussion of mis-specified models above leads us naturally to objection B. What happens in a Savage–Bayesian model when an agent observes data which were assigned prior probability 0? This is a distinct possibility when the agent's prior probability is mis-specified.

Indeed, consider the monopolist example, and suppose that the slope term is $\beta = 1$ and the monopolist knows this. Suppose that the monopolist's prior beliefs over the intercept term α are uniform on [2, 3], but the true intercept term is $\alpha^* = 10$, so that the monopolist has a mis-specified model. Let the error term ε_t be uniformly distributed on $[-1, 1]$.

Suppose that at date 1 the agent chooses the output level $q = 1$. Then under the monopolist's prior beliefs the monopolist expects to observe prices anywhere in the set given below (and excuse the obvious abuse of notation):

$$p = \alpha - \beta q + \varepsilon \text{ lies in the set } [2, 3] - (1)(1) + [-1, 1] = [0, 3]$$

However, the true prices are governed by the true intercept value $\alpha^* = 10$. The true observed prices will therefore lie in the set $10 - (1)(1) + [-1, 1] = [8, 10]$. Hence the monopolist is sure to observe data that are assigned prior probability 0.

Does this model contradict the Savage–Bayesian paradigms? Let us introduce some detail to make our point. Let θ be a set of unknown parameters and let Z denote the set of possible observations of this agent. Let P denote the agent's prior beliefs over the product space $\Omega \equiv \Theta \times Z$. The agent's posterior distribution given the observation of the data z is then the conditional probability $P(\cdot \mid z)$. (Note that what is typically called the prior is actually $\text{marg}_\Theta P$, the marginal of P on Θ; and what is typically called the posterior is $\text{marg}_\Theta P(\cdot \mid Z)$.)

The question we seek to answer is then the following: is $P(. \mid z)$ well defined for all z in Z, and in particular is it defined for z values which may be assigned a prior (i.e. P) probability of zero?

Well, what is a conditional probability $P(. \mid z)$? It is a function mapping observed values of z into the set possible of probability measures on $\Theta \times Z$ and should satisfy the following conditions.

(a) It is a function of z (or more formally, for any subset S of $\Theta \times Z$, $P(S \mid z)$ should be a random variable measurable with respect to $\sigma(\{z\})$, the σ-field generated by z).

(b) It integrates properly (i.e. $\int_Z P(S \mid z)\mathrm{d}P = P(S)$ for each subset S of $\Theta \times Z$).

Observe therefore that a conditional probability is unique only up to probability 0 sets. In particular, let $P(. \mid z)$ be a conditional probability and define $P'(. \mid z)$ such that, for each z in Z, $P'(. \mid z)$ is a probability measure on $\Theta \times Z$ satisfying (a) and such that $P(\{z \text{ in } Z : P(. \mid z) = P'(. \mid z)\}) = 1$. Then $P'(. \mid z)$ is also a conditional probability. In particular we are free to change arbitrarily $P(. \mid z)$ on a set of z values with P probability 0. $P(. \mid z)$ and $P(. \mid z)$ are referred to as different *versions* of the conditional probability. When describing agents who are using possibly mis-specified models not only must we state their priors (in addition to utility functions, discount factor and other primitives) but we must also fix a version of the conditional probability which will provide us with a conditional probability even on probability 0 events.

Let us now turn to the monopolist example with mis-specified model. Recall that the monopolist expects observed price to lie in the interval [0, 3]. However, the prices will always lie in the interval [8, 10]. So what does the monopolist do when a price, in [8, 10], is observed?

As argued above we must fix a version of the monopolist's prior beliefs. We illustrate the range of possibilities with two examples.

1 Whenever prices outside of [0, 3] are observed the monopolist's beliefs about α are represented by some distribution, say normal, with support on all of the real line. This of course implies that the monopolist's beliefs about the set of possible prices have support on all of the real line. Hence any subsequent possible price will be entertained as a possibility.

2 Whenever prices outside of [0, 3] are observed the monopolist ignores the data. The monopolist's posterior distribution given any such observation will then be the same as the prior distribution. The two different versions will of course in general lead to different evolution of beliefs and hence actions over time. Indeed, the monopolist that reverts to the normal distribution prior when faced with "strange" data will over time "learn" the true value of α. The monopolist who throws away the data will of course throw away the data in each period and will never "learn" anything; its beliefs stay the same at each date.

The stark difference between these possibilities masks the fact that whenever we do Bayesian updating we typically have a version in mind. For example suppose we have a prior over a parameter θ which is N(0, 1), i.e. normally distributed with mean 0 and variance 1. Suppose we observe the data $z = \theta + \varepsilon$, with ε independent of θ with distribution N(0, 1). The Bayes's rule type of formula is then applied to show that the posterior distribution on θ conditional on any observation is N(z/2, 1/2), the normal distribution with mean $z/2$ and variance 1/2. This is assumed to be true for *all* z. This of course fixes a version.

To see this, note that we could have defined the posterior distribution on θ as N(z/2, 1/2) for all z values which are irrational numbers and N(0, 1), or any arbitrary distribution, for all z values which are rational numbers. This is another version of the conditional probability. Why is this any better or worse than the previous one?

One may be tempted to argue that the difference between these two possibilities in our normal updating example are irrelevant since an "outside observer" would conclude *ex ante* (i.e. before the realizations) that the two versions will be the same (with probability 1). However, this appeals to a special arbitrarily defined "outside observer." If instead we choose an outside observer who knows the true values of θ and ε and knows that z will be a rational number then that observer will conclude that the versions will differ with probability 1! When we use the formula N(z/2, 1/2) for *all* z as the posterior distribution we are stating our preference for that particular version of the conditional probability.

Some restriction on the set of possible versions is possible, and in the context of our monopolist example may indeed be useful. Suppose an event occurs which was assigned prior probability 0. If we have fixed a version of the prior probability the conditional probability will be defined even on this event. One may then wish to require the conditional probability to assign probability 1 to the event upon which it has been conditioned. This requirement on the version leads to the concept of a *proper* conditional probability, as in Blackwell and Dubins (1975).

In many cases the requirement of properness will result in a significant reduction in the set of possible versions. For example consider the monopolist problem above where the monopolist has a mis-specified model. Let us suppose that the mis-specification is only in the monopolist's beliefs about the parameter θ and not in the distribution of the data, given θ. Suppose the monopolist observes the data (prices) p^*. Let us require the monopolist to have a proper version of the conditional probability. Then the monopolist's conditional probability must assign probability 1 to the set of θ values which could possibly generate p^*. In the monopolist example no θ in the support of the monopolist's prior (the set of α in [2, 3]) will generate any observed p^* (in the set [8, 10]). Hence properness implies that the monopolist cannot have a posterior distribution with support in the set {α in [2, 3]}. In particular the monopolist cannot have a posterior distribution which is equal to its prior and so cannot "throw away" the data.

The requirement of properness has been used in the recent game theory literature (see, for example, Brandenburger and Dekel, 1987; Blume et al., 1991). However, even when we require versions to be proper we are still left with many versions. One may therefore object to the use of any one particular version of conditional distribution (even when proper) because they are too *ad hoc*. One may ask: where do the arbitrarily specified updating rules on the probability zero sets come from? Our answer: Yes, as you may have predicted, they come from the same place as the utility functions! The versions of the conditional distributions are just as *ad hoc* as the prior distributions and the utility functions. Indeed, they should be considered an integral part of the prior beliefs. When the model is being specified, one typically states the prior probability and the utility function. We ask that the prior probability should be stated with its associated version of the conditional probabilities.

From Savage (1954), we know that an agent's prior beliefs may be determined by asking the agent to state preferences among different lotteries. It is assumed the agent can do that. Similarly to obtain conditional probabilities on probability 0 events we ask the agent to choose among different lotteries – conditional on these probability 0 events occurring. If we assume that agents can do this then we arrive at a conditional probability on those events. (Of course there remains some subtleties of interpretation of the Savage axioms in this case. This is further discussed in forthcoming work. However, one should consult Blume et al. (1991) for more on this.)

In summary we stress that the axioms of Savage do not exclude the possibility of agents observing zero probability events. Since these are not thought to be terribly relevant to statistics, they have not received much attention in the statistical literature. Nevertheless they are of considerable interest in economic modeling (and perhaps should receive more attention in statistics as well, as illustrated by the "model revision" problem). One such economic model has been studied by Nyarko (1991a).

4.3 Information processing and computational constraints

We turn now to objection C. We first note that this is related to objection A, and hence some resolution to objection C may be obtained by simply supposing that the agent's beliefs assign probability 1 to a set of "simple" models; we then proceed as discussed in section 4.1 above.

However, objection C should really be interpreted as information processing constraints involved in the use of Bayesian modeling. This objection to Bayesian methods also has a history. Good (1983) suggests the concept of type II rationality. Type II rational Bayesians maximize expected utility when costs of calculation are taken into account. This straightforward idea has not generated adequate development in economics. Our position is clear: information processing and computational constraints should be modeled explicitly as part of the constraints (Savage–Bayesian) that agents face when maximizing their expected utility.

Kiefer and Rao-Sahib (1991) have taken up the topic and looked at simple models in which agents cannot process all the relevant data generated in the economy. For example, data may be generated at the rate of K bits per minute but the agent can absorb only $K/2$ bits per minute. These models use coding theorems to bound expected losses relative to full information decisions. Another promising but largely unexplored approach is to approximate arbitrary posterior distributions by, for example, orthogonal polynomials and just keep finitely many terms, thus reducing the dimension of the state space. This is relevant in important practical cases.

Our point is that our agents may well be constrained in processing information or in computations. These constraints should be incorporated into the model. They will result in a restriction on the set of feasible acts that can be chosen by the Savage–Bayesian agent. Incorporating constraints into optimization problems is an exercise in which economists are well trained. It is surprising that so little work has been done on information processing and computation constraints.

5 Do Bayesians Make More Money Than Non-Bayesians?

(We have no information about the salaries of economists so we cannot answer the question about Bayesian versus non-Bayesian economists!) What we are concerned with here is the question of whether Bayesian agents make more money than non-Bayesian agents in economic models where both exist. The general answer is "maybe"; we illustrate below why a sharper general answer cannot be expected. We then turn to an example in which a single Savage–Bayesian agent amidst many non-Bayesians makes more money.

5.1 It depends upon the prior and the truth

For simplicity suppose that the common goal is to make money, i.e. bets are small or utility is linear. Even here, the question is not focused sharply enough to admit a clear answer. For example, suppose the Bayesian is betting "heads" at even odds on coin flips, and is certain that the coin falls "heads" with probability 0.9. The Bayesian will expect to make money. However, if we take the Bayesian's strategy as given and evaluate the winnings under that strategy and with the knowledge that the coin actually falls "heads" with probability 0.3 say, we see that the Bayesian will lose money (and breaks even if the coin is fair). A non-Bayesian who always bets on "tails" (perhaps because it is July) will make money, thus beating the Bayesian.

This example illustrates the difficulty in reaching any firm result on who makes more money. Of course, given common information, the Savage–Bayesian expects to make at least as much money as anyone else. However, this is an expectation with respect to the Savage–Bayesian's prior beliefs. Arguments very similar to the above have been made by Blume and Easley (1992) in a model of trading in a financial market.

5.2 Example where Bayesians win (against OLS agents)

One easy way to make Savage–Bayesian agents make more money is via the interpretation that such agents take into account much more information, and in particular take into account the way other agents are behaving. This is an often heard interpretation of the definition of a Bayesian agent in the economics literature. This is of course not necessarily our definition, which is one which follows only from the axioms of Savage (1954). However, we shall study the implication of such an interpretation via the following example.

Consider the following model of the interaction of a large number of competitive firm. Suppose there is a continuum of firms indexed by the set $I = [0, 1]$. Each agent (i.e. firm) i in I must choose an output level x_{it} at each date t. Define x_t to be the average output level of the agents at date t; i.e $x_t = \int x_{it}\mathrm{d}i$. The equilibrium price of the good depends upon the aggregate output and is obtained via a linear demand curve $p_t = \alpha - \beta x_t + \varepsilon_t$ where (α, β) are constants *known* to each agent. $\{\varepsilon_t\}_{t=1}^{\infty}$ is the shock process and is an unobserved independent and identically distributed process with zero mean and finite variance, with a distribution known to all agents.

If each agent has a cost function $c(x_{it}) = x_{it}^2$, the expected profits of each agent will be $x_{it}E_i p_t - 0.5x_{it}^2/2$, where $E_i p_t$ is the expectation operator of agent i. The profit-maximizing output can then be easily shown to be $x_{it} = E_i p_t$. When agents are choosing their date t actions, they know the history of past prices $\{p_\tau\}_{\tau=1}^{t-1}$ but they do know what the date t price will be. Let us suppose that that agents do *not* observe the past aggregate output levels.

Suppose that each agent i in I predicts the end of period price to be the average of all previous prices:

$$p_t^{\text{ols}} = \sum_{\tau=1}^{t-1} \frac{p_\tau}{t-1} \tag{2.7}$$

We shall call such agents OLS agents since their predicted price is that which would be predicted by the ordinary least squares algorithm.

Each such agent will then choose an output level equal to $x_{it} = p_t^{\text{ols}}$. This results in an end of date t price $p_t = \alpha - \beta p_t^{\text{ols}} + \varepsilon_t$. Define $P_t^* \equiv Ep_t = \alpha - \beta p_t^{\text{ols}}$ to be the expected price in the economy with OLS agents. The expected profit of each such agent will then be

$$E\pi_t^{\text{ols}} = E\left[p_t x_{it} - 0.5x_{it}^2\right] = p_t^* p_t^{\text{ols}} - 0.5\left(p_t^{\text{ols}}\right)^2 \tag{2.8}$$

Now suppose that there is one agent, a "Bayesian," who is by definition one that knows that other agents are choosing actions via the OLS learning rule but is otherwise the same as all other agents in the economy. Suppose that that agent is so small that it does not affect the prices. Then such an agent will know that the end of period expected price will be p_t^* and will choose an output level $x_t^* = p_t^*$ resulting in a profit of

$$E\pi_t^* = Ep_t p_t^* - 0.5\left(p_t^*\right)^2 = 0.5\left(p_t^*\right)^2 \tag{2.9}$$

It is therefore very easy to see from (2.8) and (2.9) that

$$E\left(\pi_t^* - \pi_t^{\text{ols}}\right) = 0.5\left(P_t^* - P_t^{\text{ols}}\right)^2 = 0.5\left[\alpha - (1+\beta)p_t^{\text{ols}}\right]^2 \tag{2.10}$$

In particular, the following observations are clear. First, at each date the Bayesian will be making profits no less than that of the OLS agent. This profit is strictly positive as long as the OLS estimate is different from the rational expectations value $\alpha/(1+\beta)$. When the noise term ε_t is nondegenerate this will in general be the case at each date.

The profits will of course tend to zero in the limit if the OLS estimate converges to the rational expectations level $\alpha/(1+\beta)$. One can actually show that if $1+\beta \geq 0$ and the first moment of the error term is finite then indeed the OLS prices converge to $\alpha/(1+\beta)$. It is easy to show, however, that when any of these two hypotheses are violated we may indeed obtain situations where the OLS estimator does not converge to $\alpha/(1+\beta)$ and so the Bayesian agent makes strictly more profits than the OLS agent even in the limit. In summary we have the following:

Proposition 1 (Bayesian is infinitely richer infinitely often)

(i) The Bayesian cannot make less money than the OLS agent, i.e. $\pi_t^* \geq p_t^{\text{ols}}$ for each t.

(ii) If either (a) $1+\beta < 0$, $\varepsilon_t \equiv 0$ for all t, and the initial price $p_0 \geq \alpha/(1+\beta)$; *or* (b) the shock distribution has fat tails, i.e. $E(|\,\varepsilon\,|) = \infty$ (e.g. a Cauchy distribution), then $\limsup_{t\to\infty} E\left(\pi_t^* - \pi_t^{\text{ols}}\right) = +\infty$.

Remarks We could also consider a model where there are many Bayesian agents interacting with many OLS agents. Those Bayesians agents will then have to make predictions about how other Bayesians are behaving. This results in an infinite hierarchy of beliefs about beliefs about beliefs, etc. For more on this problem see Nyarko (1990, 1991c).

6 On the Question of Learning in the Limit

6.1 It depends on the prior and induced sample path

As argued earlier, this is an inappropriate question in economics since it is not the objective of agents to "learn." Agents seek to choose the best actions given their informational constraints. Recall in the monopolist problem the choice of an optimal action involved a trade-off between actions that yeild high current period returns and which results in high future information. A monopolist may therefore optimally decide to choose actions which do not reveal the true parameter values.

To illustrate this possibility suppose that in the monopolist problem there are only two vectors of the unknown parameter, (α', β') and (α'', β''). This results in two mean demand curves as, for example, in figure 2.1 with point of intersection $(\bar{q}, \bar{p})$. If the monopolist chooses the action $\bar{q}$, the monopolist will receive no information on the unknown demand parameters since with that output level both demand parameters yield the same distribution of prices. Suppose that $\bar{q}$ maximizes the current period expected profit. Any output level different from $\bar{q}$ results in a loss in current period profits which may or may not be compensated for in the future information value of that different output level. If the agent is risk averse and the discount factor is low, it may indeed be possible that the future information value never compensates for the loss in current period profit. In that case the action $\bar{q}$ is optimal and the Bayesian agent does no learning over time.

The potential for non-learning has been shown for low discount factors under differentiability assumptions by Kiefer and Nyarko (1989); it has been shown in numerical simulations by Kiefer (1989). This question has also been studied under genericity conditions by McLennan (1987) and under normality assumptions by Feldman and McLennan (1989).

The monopolist problem with actions which are possibly vector-valued and where other exogenous variables may be present has been studied by Nyarko (1991b). Huffman and Kiefer (1990) and El-Gamal and Sundaram (1989) study a variant of the monopolist problem (and, in particular, optimal growth) where capital is an additional state variable. Nyarko and Olson (1991) study the optimal growth model where the unknown parameter is not a fixed parameter θ but an endogenously moving quantity (the stock of resources). The conclusions from all the above variants of the monopolist problem remain the same as before: the answer to the question of whether the Savage–Bayesian agent learns the unknown parameters over time is "it depends." The answer can be either yes or no depending on the particular model, prior beliefs and utility function.

6.2 On the same sample path the Bayesians cannot be beaten

As we illustrated in the previous section a Savage–Bayesian agent need not learn the unknown parameter vector in the limit. Other agents using different behavior rules may therefore do better than the Savage–Bayesian agent when it comes to learning. Indeed, consider the monopolist example outlined above where figure 2.1 applies. Consider a non-Bayesian agent who chooses the action $\bar{q} + e$ at even dates and $\bar{q} - e$ at odd dates where $e > 0$ is some "small" number. It can easily be shown that such an agent will learn the unknown parameter vector over time. (This is similar to the concept and discussion of ε-equilibria in Easley and Kiefer (1988).) Such an agent will therefore do better than the Savage–Bayesian agent as far as learning is concerned. There should of course be nothing too surprising about this. The objective of the Savage–Bayesian is not to learn, but to maximize utility! Utility maximization may lead to decisions and hence a sample path of observations which do not maximize the ability to learn the unknown parameter.

However, if we *fix* the sample path of observations, the Savage–Bayesian agent cannot be beaten when it comes to the question of learning. In particular if some other agent is able to learn the unknown parameter vector then so will the Bayesian

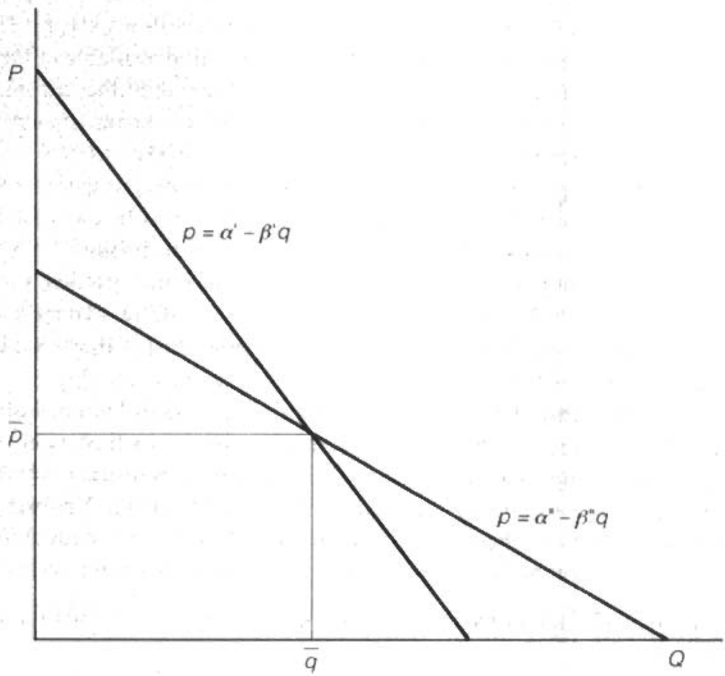

Figure 2.1 Possibility of non-learning monopolist

agent. More precisely, the following theorem can be proved. If an estimator can be constructed which is consistent along a given sample path (e.g. if the OLS estimator is consistent on some sample path), then the Bayesian posterior process (with sufficiently diffuse priors) will converge to the point mass on the true parameter along that sample path. In particular, along a *given* sample path as far as *learning* is concerned Bayesian updating "cannot be beaten."

The rest of this section will be devoted to formal statement and proof of this assertion. Those uninterested in the details may wish to proceed to the next section.

Suppose there is some parameter θ in a parameter space Θ which is unknown to the agent. Let μ_0 be the probability over Θ representing the agent's prior beliefs. There is some observation process $\{Z_t\}_{t=1}^{\infty}$ taking values in the set Z whose probability[1] law P_θ depends upon the parameter θ. Define $\Omega = \Theta \times Z^{\infty}$. Any $\omega \in \Omega$ specifies a parameter value $\theta(\omega)$ and a sample path of observations $\{z_t(\omega)\}_{t=1}^{\infty}$. We then define P to be probability measure over Ω induced by μ_0 and P_0; i.e. for any subset S of Ω,

$$P(S) \equiv \int P_\theta(S)\, \mathrm{d}\mu_0 \tag{2.11}$$

Let $\mathcal{F}_t$ be the σ-field generated by the data $\{z_1, \ldots, z_t\}$. Let $\mathcal{F}_\infty$ denote the information generated by the entire sample path of observations $\{z_1, z_2, \ldots\}$. Fix any $t = 1, 2, \ldots$ or $t = \infty$. $\mathcal{F}_t$ represents the information available at the end of date t. Let $P(. \mid \mathcal{F}_t)$ denote the probability P conditional on the information in $\mathcal{F}_t$. Let μ_t denote the marginal distribution of $P(. \mid \mathcal{F}_t)$ on the parameter space θ. μ_t is the agent's posterior distribution over θ at the end of date t.

Suppose the true parameter is θ. Let 1_θ be the probability measure on θ which assigns probability one to θ. The Bayesian agent learns the true parameter vector in the limit if μ_t converges to 1_θ in some metric with P_θ-probability 1. The metric we shall use is that of the topology of weak convergence; in particular a sequence of measures on θ, $\{\mu^n\}_{n=1}^{\infty}$, converges to the measure μ^∞ if, for all continuous and bounded real-valued functions on θ, $\int f\, \mathrm{d}\mu^n$ converges to $\int f\, \mathrm{d}\mu^\infty$; we let wlim denote the operation of taking limits of measure in the weak topology.

First we show that the Bayesian agent's posterior distribution will always converge. It converges to the probability distribution μ^∞ which is conditional on the complete sample path of observations, $\mathcal{F}_\infty$. The conclusion of this result is that the beliefs do not wander around but eventually settle down. Nothing is being said at this stage about whether or not the limiting belief, μ_∞, is equal to 1_θ and hence whether or not the Bayesian agent learns the true parameter vector.

Proposition 2 (Beliefs converge) For μ_0 almost every $\theta \in \Theta$, $P_\theta(\{\mathrm{wlim}_{n\to\infty} \mu_t = \mu_\infty\})$.

Let us now consider a "sophisticated" and not necessarily Bayesian *econometrician.* At the end of date t the econometrician will have an estimate g_t of the parameter θ. This must be a function of the data from date 1 to t (the econometrician may be sophisticated but is not clairvoyant!). In the language of measure theory, g_t must be $\mathcal{F}_t$-measurable.

We define an *estimator* to be a sequence $\{g_t\}_{t=1}^{\infty}$ such that, for all t, $g_t : \Omega \to \Theta$ and is $\mathcal{F}_t$ measurable for all t. Suppose that the true parameter vector is θ. Then we say that the estimator is *consistent* if it converges to θ on a set of sample paths with P_θ-probability 1; i.e. if

$$P_\theta \left(\lim_{n \to \infty} g_t = \theta \right) = 1 \tag{2.12}$$

Given any sequence $\{g_t\}_{t=1}^{\infty}$ of $\mathcal{F}_t$-measurable Θ-valued functions, note that the limit $g \equiv \lim_{n\to\infty} g_t$ if it exists is an $\mathcal{F}_\infty$-measurable function. From proposition 3(i) below we may conclude that if the sophisticated econometrician can construct a consistent estimator then necessarily the Bayesian agent will learn the value of the parameter vector in the limit.

Indeed, more is true. One can show that, on any sample path where the sophisticated econometrician can construct an estimator which converges to the true parameter, the Bayesian agent will necessarily learn the parameter vector on the same sample path. The formal statement of this stronger assertion is proposition 3(ii) below.

Proposition 3 Let g be an F_∞-measurable Θ-valued function.

(i) Suppose that for μ_0 almost everywhere θ, $P_\theta(\{\omega \in \Omega \mid g(\omega) = \theta\}) = 1$. Then for μ almost everywhere θ, $P_\theta(\text{wlim}_{n\to\infty} \mu_n = 1_\theta) = 1$.

(ii) Suppose there exists an F_∞-measurable set B such that for μ_0 almost everywhere θ, $P_\theta(\{\omega \in \Omega | g(\omega) = \theta)\} \cap B) = P_\theta(B)$. Then for μ almost everywhere θ, $P_\theta(\{\text{wlim}_{n\to\infty} \mu_n = 1_\theta\} \cap B) = P_\theta(B)$.

7 Concluding Remarks

Many of the results available in the literature on *ad hoc* learning models are adapted from results available from the study of stochastic approximation methods, or other methods used to study the properties of numerical algorithms in computer science and engineering. The pejorative that the results are often "mechanical" is thus hard to avoid. Nevertheless, it is often the case that the *ad hoc* methods turn out to be exactly or approximately Savage–Bayesian methods for some specification of prior and utility. In these cases the results are useful, although strictly applicable of course only in these particular cases and not in the generality sometimes claimed. That is, an *ad hoc* method may have the implication for a class of models that "learning occurs," i.e. the *ad hoc* estimators are consistent. This is the "lucky learning" result which often seems to be the goal of investigation. However, it is only when the *ad hoc* method coincides with a Savage–Bayesian approach that we can conclude from this that "learning is optimal."

We have discussed in section 4 some common but unfounded criticisms of the Savage–Bayesian paradigm. The Savage–Bayesian model rests upon the foundations of Savage axioms. Many may consider these axioms and their implications too strong. Currently work is being done by Machina and others on relaxing these axioms. We welcome these advances. The main conclusions of this literature is that under some relaxed assumptions, instead of subjective expected utility maximization where the probabilities enter the utility functional linearly, we obtain the maximization of a utility functional where the probabilities enter in some specific nonlinear way. We look forward to seeing applications of these non-expected utility formulations in more applied problems.

Appendix

Proof of proposition 3

(i) To keep the exposition precise, in particular to indicate that the true parameter can be considered a random variable (with distribution μ_0), we let $\theta(\omega) \in \Theta$ be the projection of $\omega \in \Omega \equiv \Theta \times Z^\infty$ onto its Θ coordinate. We also index the posterior distributions $\mu_t(\omega)$ by the sample path, ω. For fixed $\theta \in \Theta$, define $C_\theta \equiv \{\omega \in \Omega \mid \mathrm{wlim}_{t\to\infty} \mu_t = 1_\theta\}$ and $C \equiv \{\omega \in \Omega \mid \mathrm{wlim}_{t\to\infty} \mu_t(\omega) = 1_{\theta(\omega)}\}$. Recall that $P \equiv P_\theta \times \mu_0$. Then

$$P(C) = \int_\theta P_\theta(C_\theta)\mathrm{d}\mu_0 \tag{A.1}$$

We proceed to show that $P(C) = 1$, which then completes the proof since $P(C) = 1$ implies, using (A.1), that $P_\theta(C_\theta) = 1$ for μ_θ almost every θ, which is proposition 3.

From proposition 2, $P_\theta(\{\omega \in \Omega \mid \mathrm{wlim}_{n\to\infty} \mu_t = \mu_\infty\}) = 1$ for μ_0 almost every θ. Hence $P(\{\omega \in \Omega \mid \mathrm{wlim}_{n\to\infty} \mu_t(\omega) = \mu_\infty(\omega)\}) = 1$. To show that $P(C) = 1$, it therefore suffices to show that $\mu_\infty(\omega) = 1_{\theta(\omega)}$, P almost everywhere. However, from the definition of g, $P_\theta(\{\omega \in \Omega | g(\omega) = \theta(\omega)\}) = 1$ for μ_θ almost every θ. This in turn implies $P(\{\omega \in \Omega \mid g(\omega) = \theta(\omega)\}) = 1$. So for any Borel set D in θ, if we let E denote expectations with respect to P, $\mu_\infty(D) = E\left[1_{\{\theta \in D\}} \mid \mathcal{F}_\infty\right] = E\left[1_{\{g \in D\}} \mid \mathcal{F}_\infty\right] = 1_{\{g \in D\}}$ (since g is F_∞-measurable) $= 1_{\{\theta \in D\}}$. Hence $\mu_\infty(\omega) = 1_{\theta(\omega)}$, almost everywhere.

(ii) The proof of part (ii) is very similar to that of part (i) with minor changes. The complete proof may be found in Kiefer and Nyarko (1988, Lemma 5.4, p. 116).

Proof of proposition 1

(a) The following formula is very easily verified:

$$p_{t+1}^{\text{ols}} = p_t^{\text{ols}} + \frac{\alpha - (1+\beta)p_t^{\text{ols}}}{1+t} + \varepsilon_{t+1} \tag{A.2}$$

Suppose that there is no noise term (or, equivalently, that $\varepsilon_t = 0$ almost everywhere for all t). Suppose that we start the OLS process at some $p_0 > \alpha/(1+\beta)$. Then it is easy to check that the date 1 price, $p_1 = \alpha - \beta p_0$, exceeds p_0. Now suppose that $p_t^{\text{ols}} > \alpha/(1+\beta)$ for some t. Then (A.2) with $\varepsilon = 0$ and the fact that $1+\beta < 0$ implies that

$$p_{t+1}^{\text{ols}} > p_t^{\text{ols}} > \frac{\alpha}{1+\beta} \tag{A.3}$$

Hence by induction (A.3) holds for all t. In particular the sequence of OLS price estimate are monotone increasing, and hence converge either to infinity or to a finite limit.

Now, $p_t = \alpha - \beta p_t^{\text{ols}}$. If the OLS estimate converge to a finite limit then the prices must converge to a finite limit. The latter limit must necessarily be the same as the limit of the OLS estimates since those estimates are averages of the prices. But then $p_t = \alpha - \beta p_t^{\text{ols}}$ implies that the common limit must equal $\alpha/(1+\beta)$.

In particular if the estimates converge to a finite limit those estimates as well as the date t prices must converge to the rational expectations level $\alpha/(1+\beta)$. However, we just argued that $p_t > p_0 > \alpha/(1+\beta)$ for all t so the prices cannot converge to $\alpha/(1+\beta)$. We therefore conclude that both the prices and the estimates increase monotonically to $+\infty$. This in turn implies that the difference between the profits of the Bayesian and the OLS agent converges to plus infinity.

(b) One can easily show that, for any fixed integer $k > 0$,

$$\sum_{t=1}^{\infty} P\left(|\varepsilon_t| \geq kt\right) \;\geq\; E\left(\left|\frac{\varepsilon_1}{k}\right|\right) - 1 \tag{A.4}$$

(see Chung, 1974, theorem 3.2.1). Since $E(|\varepsilon_1|) = \infty$, we may conclude from the Borel–Cantelli lemma that prob ($|\varepsilon_t| \;\geq\; kt$ infinitely often) $= 1$. Since the integer k is arbitrary this implies that

$$\limsup_{t\to\infty} \left|\frac{\varepsilon_t}{t}\right| = \infty \text{ a.e.} \tag{A.5}$$

Now it is very easy to check that the following recursion formula holds:

$$p_{t+1}^{\text{ols}} = \frac{\alpha}{t} + \left(\frac{t-1-\beta}{t}\right) p_t^{\text{ols}} + \left[\frac{\varepsilon_t}{t}\right] \tag{A.6}$$

Fix any sample path. If $\limsup \varepsilon_t/t = \infty$ then $\limsup p_t^{\text{ols}} < \infty$ results in a contradiction to (A.6). If $\liminf \varepsilon_t/t = -\infty$ then $\liminf p_t^{\text{ols}} > -\infty$ results in a contradiction to (A.6). Hence (A.5) implies that $\limsup |p_t^{\text{ols}}| = \infty$, almost everywhere. Equation (2.10) then proves proposition 1.

Notes

This paper was prepared for the Workshop on Learning and Rationality in Games and Economics, European University Institute, Florence, July 1991. We thank the participants for helpful comments. The paper has benefited from seminar audiences at Cornell and Brown Universities. We thank the National Science Foundation for its assistance. The second author is grateful to both the C.V. Starr Center and the Challenge Fund at New York University for their generosity.

1 Some technical remarks: We assume the spaces Θ and Z are complete and separable metric spaces. $\Theta \times Z^{\infty}$ is endowed with its product topology. Let $B(\Omega)$ denote the set of Borel subsets of Ω. The measure P is then a probability measure over the measure space $(\Omega, B(\Omega))$.

References

Blackwell, D. and Dubins, L. 1975: On existence and non-existence of proper, regular conditional probabilities. *Annals of Probability*, 3, 741–52.

Blume, L. and Easley, D. 1992: Evolution of market behavior. *Journal of Economic Theory*, 58 (1), 9–40.

Blume, L., Brandenburger, A. and Dekel, E. 1991: Lexicographic probabilities and choice under uncertainty. *Econometrica*, 59, 61–80.

Brandenburger, A. and Dekel, E. 1987: Rationalizability and correlated equilibria. *Econometrica*, 6, 1391–1402.

Chung, K.L. 1974: *A Course in Probability Theory*. New York: Academic Press.

DeGroot, M. 1970: *Optimal Statistical Decisions*. New York: McGraw-Hill.

Dreze, J. 1972: Econometrics and decision theory. *Econometrica*, 40 (1), 1–18.

Easley, D. and Kiefer, N. 1989: Controlling a stochastic process with unknown parameters. *Econometrica*, 56 (5), 963–78.

El-Gamal, M.A. and Sundaram, R.K. 1989: Bayesian economists ... Bayesian agents I: an alternative approach to optimal learning. Social Science Working Paper 705, California Institute of Technology.

Feldman, M. and McLennan, A. 1989: Learning in a repeated statistical decision problem with normal disturbances. Manuscript, University of Illinois, Urbana-Champaign.

Good, I.J. 1983: *The Foundations of Statistics and its Applications*. Minneapolis, MN: University of Minnesota Press.

Huffman, G. and Kiefer, N. 1990: Optimal learning in a one-sector growth model. Manuscript, University of California, Riverside.

Kiefer, N. 1989: A value function arising in the economics of information. *Journal of Economic Dynamics and Control*, 13, 201–23.

Kiefer, N. and Nyarko, Y. 1988: Control of a linear regression process with unknown parameters. In W.A. Barnett, E.R. Berndt and H. White (eds), *Dynamic Econometric Modelling, Proceedings of the Third International Symposium in Economic Theory and Econometrics*, Cambridge: Cambridge University Press.

Kiefer, N. and Nyarko, Y. 1989: Control of an unknown linear process with learning. *International Economic Review*, 30 (3), 571–86.

Kiefer, N. and Rao-Sahib 1991: An economist's guide to information theory. Manuscript, Cornell University.

Lindley, D.V. 1972: *Bayesian Statistics: A Review*. Philadelphia, PA: SIAM.

Machina, M. and Schmeidler, D. 1992: *Econometrica*, 60 (4), 745–80.

McLennan, A. 1984: Price dispersion and incomplete learning in the long-run. *Journal of Economic Dynamics and Control*, 7, 331–47.

McLennan, A. 1987: Incomplete learning in a repeated statistical decision problem. Manuscript, University of Minnesota.

Mirman, L., Postlewaite, A. and Kihlstrom, R. 1984: Experimental consumption and the "Rothschild effect". In M. Boyer and R.E. Kihlstrom (eds), *Bayesian Models in Economic Theory*.

Nyarko, Y. 1990: Bayesian rationality and learning without common priors. C.V. Starr Center for Applied Economics Working Paper 90-45, New York University.

Nyarko, Y. 1991a: Learning in mis-specified models and the possibility of cycles. *Journal of Economic Theory*, 51, 416–27.

Nyarko, Y. 1991b: On the convergence of Bayesian posterior processes in linear economic models: counting equations and unknowns. *Journal of Economic Dynamics and Control*, 15, 687–713.

Nyarko, Y. 1991c: The convergence of Bayesian belief hierarchies. C.V. Starr Center for Applied Economics Working Paper 91-50, New York University.

Nyarko, Y. and Olson, L. 1991: Optimal growth with unobservable resources and learning. C.V. Starr Center for Applied Economics Working Paper 91-01, New York University.

Pratt, J.W., Raiffa, H. and Schlaifer, R. 1964: The foundations of decision under uncertainty: an elementary exposition. *Journal of the American Statistical Association*, 59, 353–75.

Prescott, E. 1972: The multiperiod control problem under uncertainty. *Econometrica*, 40, 1043–58.

Rothschild, M. 1974: A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9, 185–202.

Savage, L.J. 1954: *The Foundations of Statistics*. New York, Wiley.