CHAPTER 6

Control of a linear regression process with unknown parameters

Nicholas M. Kiefer and Yaw Nyarko

1 Introduction

Applications of forms of control theory to economic policy making have been studied by Theil (1958), Chow (1975, 1981), and Prescott (1972). Many of the applications are approximations to the optimal policy - suggestions of how to improve existing practice using quantitative methods rather than development of fully optimal policies. Chow (1975) obtains the fully optimal feedback control policy for linear systems with known coefficients for a quadratic loss function and a finite time horizon. Chow (1981) argues that the use of control technique for the evaluation of economic policies is possible and essential under rational expectations. The use of optimal control for microeconomic planning is fully established. An early analysis with many practical suggestions is Theil (1958). Optimal control theory has also been useful in economic theory, in analyzing the growth of economies as well as the behavior over time of economic agents.

The problem of control of a stochastic economic system with unknown parameters is far less well studied. Zellner (1971, Chapter 11) studied the two-period control problem for a normal regression process with a conjugate prior and quadratic loss function. He obtained an approximate solution to the problem and compared it with two other solutions - the "here and now" solution, in which the agent chooses a fixed action for both periods at the beginning of the first period, and a "sequential updating" solution, in which the agent chooses the first-period policy without regard to its information value and then updates his beliefs on the basis of first-period experience before calculating the second-period policy. The sequential updating solution was studied further and compared with the

We are indebted to many conference participants, to our discussant Ingmar Prucha, and to two referees for useful comments. The revision of this chapter was supported by the National Science Foundation.

certainty equivalence rule by Harkema (1975). The optimal policy, as expected, yields a higher expected reward than the alternative policies. Prescott (1972) developed an approximation to the optimal control solution for a linear dynamic system with unknown coefficients, a finite horizon, and quadratic loss. Chow (1975, Chapter 10) developed a control policy for a linear dynamic equation system with unknown parameters, quadratic loss, and a finite horizon but no learning about the parameters over the period of control. Chow (1975, Chapter 11) provided an approximate solution for the control problem with learning. Taylor (1974) examined a simple regression case and considered asymptotic properties of estimates and policies when an agent uses least-squares control policies (i.e., in each period the agent chooses a policy as if unknown parameters were equal to their current least-squares estimates). He found that the sequence of slope estimates is consistent when the intercept is known. Anderson and Taylor (1976), in work closely related to an example we carry through here, presented simulation results for the least-squares control rule when both the slope and intercept are unknown. They found apparent convergence of the policy variable. Jordan (1985) followed up on this work by establishing conditions under which the sequence of parameter estimates obtained by following the least-squares control rule are consistent. The related problem of maximizing an unknown function has been studied (see Kushner and Clark 1978) without the parametric representation favored in economics, where parameters are sometimes directly interpretable.

We study the problem of controlling a linear regression process with unknown parameters. We use a stochastic control framework that makes explicit the trade-off between current expected reward and the information value of an action. An agent's information about unknown parameters is represented by a probability distribution. The model is set up in Section 2; Section 3 shows the existence of optimal policies. This chapter concentrates attention on the sequence of "beliefs" held by the agent. Section 4 shows that the agent's beliefs converge to some limit, which may or may not be concentrated on the true parameter value. Proofs are given in Section 5. Related results on the sequence of actions are given by Kiefer and Nyarko (1986).

2 The model

In this section we sketch the general framework we wish to study. Let Ω' be a complete and separable metric space, F' its Borel field, and (Ω', F', P') a probability space. On (Ω', F', P') define the stochastic process $\{\epsilon_t\}_1^\infty$, the *shock process,* which is unobserved by the agent. The shock process is assumed to be independent and identically distributed, with the common marginal distribution $p(\epsilon_t \mid \phi)$ depending on some parameter ϕ in R^h that is unknown to the agent. Assume that the set of probability measures $\{p(\cdot \mid \phi)\}$ is continuous in the parameter ϕ (in the weak topology of measures) and that for any ϕ, $\int \epsilon p(d\epsilon \mid \phi) = 0$. Let $\bar{X}$, the *action space,* be a compact subset of R^k. Define $\Theta = R^1 \times R^k \times R^h$ be the parameter space. If the "true parameter" is $\theta = (\alpha, \beta, \phi) \in \Theta$ and the agent chooses an action $x_t \in \bar{X}$ at date t, the agent observes y_t, where

$$y_t = \alpha + \beta x_t + \epsilon_t \tag{2.1}$$

and where ϵ_t is chosen according to $p(\cdot \mid \phi)$.

Let $\tilde{\Theta}$ be the Borel field of Θ, and let $P(\Theta)$ be the set of all probability measures on $(\Theta, \tilde{\Theta})$. Endow $P(\Theta)$ with its weak topology, and note that $P(\Theta)$ is then a complete and separable metric space (see, e.g., Parthasarathy 1967, Chapter II, Theorems 6.2 and 6.5). Let $\mu_0 \in P(\Theta)$ be the *prior probability* on the parameter space with a finite first moment.

The agent is assumed to use Bayes's rule to update the prior probability at each date after any observation of (x_t, y_t). For example, in the initial period, date 1, the prior distribution is updated after the agent chooses an action x_1 and observes the value of y_1. The updated prior, that is, the posterior, is then $\mu_1 = \Gamma(x_1, y_1, \mu_0)$, where $\Gamma: \bar{X} \times R^1 \times P(\Theta) \to P(\Theta)$ represents the Bayes rule operator. If the prior μ_0 has a density function, the posterior may be easily computed. In general, the Bayes rule operator may be defined by appealing to the existence of certain conditional probabilities (see Appendix). Under standard conditions the operator Γ is continuous in its arguments, and we assume this throughout. Any $\{x_t, y_t\}$ process will therefore result in a posterior process $\{\mu_t\}$, where, for all $t = 1, 2, \ldots,$

$$\mu_t = \Gamma(x_t, y_t, \mu_{t-1}). \tag{2.2}$$

Let $\bar{H}_n = P(\Theta)\prod_{i=1}^{n-1}[\bar{X} \times R^1 \times P(\Theta)]$. A *partial history* h_n at date n is any element $h_n = (\mu_0, (x_1, y_1, \mu_1), \ldots, (x_{n-1}, y_{n-1}, \mu_{n-1})) \in \bar{H}_n$; h_n is said to be admissible if equation (2.2) holds for all $t = 1, 2, \ldots, n-1$. Let H_n be the subset of $\bar{H}_n$ consisting of all admissible partial histories at date n.

A *policy* is a sequence $\pi = \{\pi_t\}_{t=1}^{\infty}$, where, for each $t \geq 1$, the policy function $\pi_t: H_t \to \bar{X}$ specifies the date t action $x_t = \pi_t(h_t)$ as a Borel function of the partial history h_t in H_t at that date. A policy function is *stationary* if $\pi_t(h_t) = g(\mu_t)$ for each t, where the function $g(\cdot)$ maps $P(\Theta)$ into $\bar{X}$. Note that μ_t can be regarded as a state variable at date t, containing all relevant information about the parameters provided by the partial history h_t.

Define $(\Omega, F, P) = (\Theta, \tilde{\Theta}, \mu_0)(\Omega', F', P')$. Any policy π then generates a sequence of random variables $\{(x_t(\omega), y_t(\omega), \mu_t(\omega)\}_{t=1}^{\infty}$ on (Ω, F, P), as described above, using (2.1) and (2.2) (the technical details are stated in the Appendix).

For any $n=1,2,\ldots$, let F_n be the subfield of F generated by the random variables (h_n, x_n). Notice that $x_n \in F_n$ but y_n and $\mu_n \in F_{n+1}$. Next define $F_\infty = V_{n=0}^{\infty} F_n$.

We now discuss the utility and reward functions and define the optimality criterion. Let $u: \bar{X} \times R^1 \to R^1$ be the utility function, and in particular, $u(x_t, y_t)$ is the utility to the agent when action x_t is chosen at date t and the observation y_t is made. We assume:

(A.1) u is bounded above and continuous.

The reward function $r: \bar{X} \times P(\Theta) \to R^1$ is defined by

$$r(x_t, \mu_{t-1}) = \int_{\Theta} \int_{R} u(x_t, y_t) p(d\epsilon_t \mid \phi) \mu_{t-1}(d\alpha\, d\beta\, d\phi), \tag{2.3}$$

where $y_t = \alpha + \beta x_t + \epsilon_t$.

Let δ in $(0,1)$ be the discount factor. Any policy π generates a sum of expected discounted rewards equal to

$$V_\pi(\mu_0) = \int \sum_{t=1}^{\infty} \delta^t r(x_t(\omega), \mu_{t-1}(\omega)) P(d\omega), \tag{2.4}$$

where the (x_t, μ_t) processes are those obtained using the policy π. A policy π^* is said to be an *optimal policy* if, for all policies π and all priors μ_0 in $P(\Theta)$,

$$V_{\pi^*}(\mu_0) \ge V_\pi(\mu_0). \tag{2.5}$$

The processes (x_t, y_t) corresponding to an optimal policy are called *optimal processes*. Even though the optimal policy π^* (when it exists) may not be unique, the function $V(\mu_0) = V_{\pi^*}(\mu_0)$ is always well defined and will be referred to as the *value function*.

3 The existence theorem

We now indicate that stationary optimal policies exist and that the value function is continuous.

Theorem 3.1. *A stationary optimal policy $g: P(\Theta) \to \bar{X}$ exists. The value function V is continuous on $P(\Theta)$, and the following functional equation holds:*

$$V(\mu_{t-1}) = r(x_1, \mu_{t-1}) + \delta \int V(\mu_t) p(d\epsilon_t \mid \phi) \mu_{t-1}(d\alpha\, d\beta\, d\phi), \tag{3.1}$$

where $\mu_t = \Gamma(x_t, y_t, \mu_{t-1})$ and $y_t = \alpha + \beta x_t + \epsilon_t$ and the integral is taken over $\Theta \times R^1$.

Proof: Let $S = \{f: P(\Theta) \to R \mid f \text{ is continuous and bounded}\}$. Define T: $S \to S$ by

$$Tw(\mu_0) = \max_{x \in \bar{X}} r(x, \mu_0) + \delta \int V(\mu_1) p(d\epsilon_1 \mid \phi) \mu_0(d\alpha\, d\beta\, d\phi), \qquad (3.2)$$

where $\mu_1 = \Gamma(x, y_1, \mu_0)$ and $y_1 = \alpha + \beta x + \epsilon_1$. One can easily show that for $w \in S$, $Tw \in S$; also, T is a contraction mapping. Hence, there exists a $v \in S$ such that $v = Tv$. Replacing w with v in (3.2) then results in (3.1); and since $v \in S$, v is continuous. Finally, it is immediate that the solution to the maximization exercise in (3.2) (replacing w with v) results in a stationary optimal-policy function [one should consult Blackwell (1965) or Maitra (1968) and Schal (1979) for the details of the preceding arguments]. Q.E.D.

4 Convergence properties of posterior process

In this section the convergence properties of the posterior process $\{\mu_t\}$ for arbitrary (i.e, not necessarily optimal) policies are studied.

The main results of this section may be described as follows. Proposition 4.1 shows that the posterior process always converges (in the weak topology of measures) with probability 1. However, the limiting probability μ_∞ may or may not be concentrated on the true parameter. Proposition 4.2 indicates that if there exists a strongly consistent estimator – be this the ordinary least-squares (OLS) estimator, the maximum-likelihood estimator, or some other – the posterior process necessarily converges to the true parameter.

In Section 4.1 the model is simplified somewhat (in particular, assume that the distribution of shocks is known and further that $k = 1$ so as to have a simple regression equation $y = \alpha + \beta x + \epsilon$). Under this simplification, some characterization of the limiting distribution can be provided. In particular, if for some ω in Ω, $x_t(\omega)$ does not converge to a limit, the limiting posterior distribution for that ω in Ω is concentrated on the "true" parameter value. Alternatively, if $x_t(\omega)$ does converge to some $x(\omega)$, say, the posterior process converges to a limiting probability with support a subset of the set $\{\alpha', \beta': \alpha' + \beta' x(\omega) = \alpha + \beta x(\omega)\}$, where α, β represent the "true" parameter values.

4.1 *Convergence of* $\{\mu_t\}$

First we prove that under the very general conditions of Sections 2 and 3, the posterior process converges for P almost everywhere (a.e.) ω in Ω to a well-defined probability measure (with the convergence taking place in the weak topology).

Note that for any Borel subset D of the parameter space Θ, if we suppress the ω's and let, for some fixed ω, $\mu_t(D)$ represent the mass that measure $\mu_t(\omega)$ assigns to the set D,

$$\mu_t(D) = E[1_{\{\theta \in D\}} \mid F_t]. \tag{4.1}$$

Define a measure μ_∞ on Θ by setting, for each Borel set D in Θ,

$$\mu_\infty(D) = E[1_{\{\theta \in D\}} \mid F_\infty]. \tag{4.2}$$

The proposition that follows shows that μ_∞ is the limiting posterior distribution and is indeed a well-defined probability measure.

Proposition 4.1. *The posterior process $\{\mu_t\}$ converges, for P-a.e. ω in Ω, in the weak topology, to the probability measure μ_∞.*

Recall P is the probability on Ω. Define P_θ to be the conditional distribution of P on Ω given the value θ in Θ. Here, P_θ should be interpreted as the distribution of histories - sequences $\{x_t, y_t\}$ - given values of the parameters of the regression equation and of the shock process θ. The proof of the following proposition is due to Schwartz (1965, Theorem 3.5, p. 14). Let 1_θ be the point mass at θ.

Proposition 4.2. *Suppose there exists an F_∞-measurable function g such that for μ_0-a.e. θ in Θ, $g(\omega) = \theta$, P_θ-a.e. Then for μ_0-a.e. θ in Θ, $\mu_\infty(\omega) = 1_\theta$, P_θ-a.e.*

The existence of a strongly consistent estimator is equivalent to the existence of a function g with the properties stated in Proposition 4.2.

4.2 *Simple regression equation model*

In this section a few simplifying assumptions are introduced to enable some rather strong characterizations of the convergence properties of the posterior process. These assumptions reduce the model to the situation of a simple regression equation. In particular, suppose that condition (S) holds:

Condition (S): The shock process has a distribution that is known to the agent and possesses finite second moment; $k = 1$, so that the action space $\bar{X}$ is a subset of R^1.

Proposition 4.3 shows that if the x_t process does not converge, the posterior process converges to the point mass on the true parameter value.

Note, however, that nonconvergence of the x_t process is not necessary for convergence of μ_t to point mass.

Let $B = \{\omega : x_t(\omega)$ does not converge$\}$, and recall that 1_θ is the point mass at θ.

Proposition 4.3. *For μ_0-a.e. θ in Θ, the posterior process $\mu_t(\omega)$ converges to 1_θ for P_θ-a.e. ω in B.*

Define on B^C, the set where $x_t(\omega)$ converges, $x(\omega) = \lim_{t \to \infty} x_t(\omega)$. In Proposition 4.4, it is shown that if the x_t process does converge to $x(\omega)$, the posterior process converges to a limiting probability with support a subset of the set $\{(\alpha', \beta') : \alpha' + \beta' x(\omega) = \alpha + \beta x(\omega)\}$, where α, β represent the true parameter values.

Proposition 4.4. *For μ_0-a.e. $\theta = (\alpha, \beta)$ in Θ, the posterior process $\mu_t(\omega)$ converges to a limiting distribution $\mu_\infty(\omega)$ with support a subset of the set $\{(\alpha', \beta') : \alpha' + \beta' x(\omega) = \alpha + \beta x(\omega)\}$ for P_θ-a.e. ω in B^C.*

5 Proofs

Proof of Proposition 4.1

The proof may be summarized as follows: We use equation (4.1) to show that for any Borel set D in Θ, $\mu_t(D)$ is a martingale sequence and apply the martingale convergence theorem to show that μ_t converges weakly to μ_∞. This argument does not assure us that the limit is a probability measure. However, the sequence of probability measures $\mu_t(\omega)$ for fixed ω is tight, and Prohorov's theorem can be applied to deduce that μ_∞ is a probability measure.

A sequence of probability measures v_n on Θ is said to be tight if, for all $\epsilon > 0$, there exists a compact set K^ϵ such that $v_n(K^\epsilon) \geq 1 - \epsilon$ for all n. Claim 5.1 establishes the tightness of $\{\mu_t\}$.

Claim 5.1. *For P-a.e. ω in Ω, the sequence of probability measures $\{\mu_t(\omega)\}$ is tight.*

Proof: Let K_r be the closed (compact) ball with center the origin and radius r. It suffices to show that for P-a.e. ω in Ω,

$$\lim_{r \to \infty} \left[\inf_t \mu_t(K_r)(\omega) \right] = 1. \tag{5.1}$$

However, using Chebyshev's inequality,

$$\mu_t(\Theta - K_r) = P(\|\theta\| > r \mid F_t) < E[\|\theta\| \mid F_t]/r,$$

so

$$\mu_t(K_r) \geq 1 - E[\|\theta\| \mid F_t]/r. \tag{5.2}$$

One can check that $\{E[\|\theta\| \mid F_t]\}$ is a positive martingale sequence and so converges to $E[\|\theta\| \mid F_\infty]$ (see, e.g., Chung 1974, Theorem 9.4.8, p. 340). We assumed that μ_0 has finite first moment, which implies that $E[\|\theta\|] < \infty$, which in turn implies that $E[\|\theta\| \mid F_\infty] < \infty$, P-a.e. Hence, $\sup_t E[\|\theta\| \mid F_t] = L < \infty$, P-a.e. Using this in (5.2) results in

$$\inf_t \mu_t(K_r) \geq 1 - \sup_t E[\|\theta\| \mid F_\infty]/r = 1 - L/r. \tag{5.3}$$

Taking limits as $r \to \infty$ then results in (5.1). This concludes the proof of Claim 5.1. Q.E.D.

Proof of Proposition 4.1 (continued)

Let U be the subclass of F made up of sets of the following kind: First, since Θ is separable, let $\{s_1, s_2, s_3, \dots\}$ be a separant; let B_n^k be the ball of radius $1/n$ and center s_k; then define U as the set of all finite intersections of the balls B_n^k, where $k = 1, 2, \dots$ and $n = 1, 2, \dots$. One may check that U is countable.

Next, for any fixed set D, $\mu_t(D) = E[1_{\{\theta \in D\}} \mid F_t]$, so using Chung (1974, Theorem 9.4.8, p. 340), the sequence $\{\mu_t(D)\}$ can be shown to be a positive martingale, and so the martingale convergence theorem applies, and we conclude that $\mu_t(D)$ converges with P probability 1 to $\mu_\infty(D)$. Since the set U is countable, convergence holds on all of U, simultaneously, with P probability 1. Then we check that U satisfies conditions (i) and (ii) of Billingsley (1968, Theorem 2.2, p. 14), so, from that Theorem, μ_t converges weakly to μ_∞ with P probability 1.

Finally, from Claim 5.1, for P-a.e. ω in Ω, the sequence of posterior distributions is tight. Hence, using Prohorov's theorem (see, e.g., Billingsley 1968, Theorem 6.1, p. 37), we may conclude that μ_∞ is a probability measure (P-a.e.). Q.E.D.

Proof of Proposition 4.2

A stronger version of this proposition is stated and proved later in Lemma 5.4.

Comment on Proof of Proposition 4.3

The idea behind the proof of Proposition 4.3 is the following. Suppose first that $x_t(\omega) = x'$ for all t and for all ω. Then $y_t(\omega) = \alpha + \beta x' + \epsilon_t(\omega)$, and $\Sigma_{t=1}^{n} y_t(\omega)/n = \alpha + \beta x' + \Sigma_{t=1}^{n} \epsilon_t/n$. However, by the strong law of large numbers, $\lim_{n\to\infty} \Sigma_{t=1}^{n} \epsilon_t/n = 0$, P-a.e., so $y' = \alpha + \beta x'$ if we define $y' = \lim_{n\to\infty} \Sigma_{t=1}^{n} y_t(\omega)/n$; in particular, the agent will learn that the true parameter will satisfy this relation in the limit. Next, if $x_t(\omega)$ does not converge but alternates between two numbers x' and x'', it is obvious that applying the preceding argument first to the even sequence $\{x_{2t}\}_{t=1}^{\infty}$ and then to the odd sequence $\{x_{2t-1}\}_{t=1}^{\infty}$, the two equations $y' = \alpha + \beta x'$ and $y'' = \alpha + \beta x''$ will be obtained, where $y' = \Sigma_{t=1}^{\infty} y_{2t}$ and $y'' = \Sigma_{t=1}^{\infty} y_{2t-1}$, from which one may compute the true parameters α and β. In this situation the agent will learn the true parameters in the limit. This idea is behind the proof of Proposition 4.3.

In the preceding example, the law of large numbers had to be applied first to the even time subsequence and then to the odd subsequence. In Lemma 5.1, it is shown that the law of large numbers may be applied to a very large set of time subsequences. As indicated in Lemma 5.3, using the result of Lemma 5.1, one can compute the true parameter in a manner very similar to that explained in the preceding example (i.e., by solving two simultaneous equations involving the true parameters α and β). Lemma 5.4 states that if the true value can be computed, beliefs converge to point mass at the true value.

Proof of Proposition 4.3

Define $1_{\{\omega\in K\}}$ equal to 1 if $\{\omega \in K\}$ and equal to zero otherwise, where K is any subset of Ω. For any $1 < m$, define

$$N_k(\omega) = \inf\left\{n: \sum_{t=1}^{n} 1_{\{1\le x_t(\omega)\le m\}} = k\right\}$$

if the set in brackets, $\{\cdot\}$, is nonempty and equal to infinity otherwise. Notice that for each k, $N_k(\omega)$ is a stopping time, that is, $\{\omega: N_k(\omega) = t\} \in F_t$ for each t.

Lemma 5.1. *There exists a set A in F with $P(A) = 1$ such that for all rational numbers ℓ and m, with $\ell < m$, and for all $\omega \in A$,*

$$\lim_{k\to\infty} \frac{1}{k} \sum_{t=1}^{N_k(\omega)} \epsilon_t(\omega) 1_{\{\ell\le x_t(\omega)\le m\}} = 0, \tag{5.4}$$

where the N_k are those corresponding to the ℓ and m.

Proof: Fix an ℓ and m with $\ell < m$. For ease of notation, drop the ω's in the random variables $\epsilon_t(\omega)$ and $x_t(\omega)$, and define, for fixed $\ell \ge m$, $\ell_t = 1_{\{\ell < x_t < m\}}$.

If, for some $s<\infty$, $N_k(\omega)=\infty$ for all $k \geq s$, then

$$\sum_{t=1}^{N_k(\omega)} \epsilon_t 1_t = \sum_{t=1}^{s} \epsilon_t 1_t < \infty \quad \text{for all } k \geq s,$$

and (5.4) follows immediately. If $N_k(\omega)<\infty$ for all k, then $\sum_{t=1}^T 1_t \to \infty$ and

$$\lim_{k\to\infty} \frac{1}{k} \sum_{t=1}^{N_k(\omega)} \epsilon_t 1_t = \lim_{T\to\infty} \frac{\sum_{t=1}^T \epsilon_t 1_t}{\sum_{t=1}^T 1_t}. \tag{5.5}$$

Since for $t' \geq t$, $\epsilon_{t'}$ is independent of $\{1_1, \ldots, 1_t\}$, (5.4) follows from (5.5) and Lemma 5.2 due to Taylor (1974).

Hence, if for fixed ℓ and m with $\ell<m$, $A(\ell, m)$ denotes the set on which (5.4) holds, $P(A(\ell,m))=1$. Define A to be the intersection over all rational numbers $\ell<m$ of the sets $A(\ell,m)$; then $P(A)=1$, and A satisfies the conclusion of the lemma. Q.E.D.

Lemma 5.2. *Let $\{v_t\}$ be a sequence of independent random variables with mean zero and uniformly bounded variance. Let $\{z_t\}$ be a sequence of random variables such that for each t, t' with $t'>t$, $y_{t'}$ is independent of $\{z_1, \ldots, z_t\}$; then for almost every realization with $\sum_{t=1}^T z_t^2 \to \infty$,*

$$\lim_{T\to\infty} \frac{\sum_{t=1}^T z_t v_t}{\sum_{t=1}^T z_t^2} = 0. \tag{5.6}$$

Proof: One applies Taylor (1974, Lemmas 1-3) with minor modifications. Q.E.D.

In Lemma 5.3, on the set where x_t does not converge, there exists a consistent estimator for the true parameter.

Lemma 5.3. *There exists an F_∞-measurable function g such that $g(\omega)=\theta$, P_θ-a.s. on the set where x_t does not converge, that is, such that if B is the set where x_t does not converge,*

$$P_\theta(\{\omega : g(\omega)=\theta\} \cap B) = P_\theta(B). \tag{5.7}$$

Proof: Construct such a function g. To ease the exposition, assume that $\bar{X}=[0,1]$. One may check that, since $\bar{X}$ is assumed compact, this simplifying assumption is without loss of generality.

Let $\bar{Q}$ be the set of rational numbers in $\bar{X}$, and let $\bar{x}(\omega)=\limsup x_t(\omega)$ and $\underline{x}(\omega)=\liminf x_t(\omega)$. Proceed to define two random variables $h(\omega)$ and $h'(\omega)$ taking values in $\bar{Q}$ and such that $\underline{x}(\omega)<h'(\omega)<h(\omega)<\bar{x}(\omega)$. Define the function $h: \bar{X}\times\bar{X} \to Q$ as follows. First, any integer $k=1,2,\ldots$ can be uniquely written as $k=2^{n+1}+p$, where $n=1,2,\ldots$ and $0 \leq p \leq 2^{n+1}-1$; so define $s_k=(2p+1)/2^n$, where $k=2^{n-1}+p$. The sequence $\{s_k\}$

is therefore a sequence of rational numbers in $\bar{X}=[0,1]$. Define $t(x,x')=\inf\{k: s_k\in[x,x']\}$ if $x<x'$ and $t(x,x')=0$ if $x\geq x'$; and $h(x,x')=s_{t(x,x')}$ with $s_0=0$. Hence, h takes values in $\bar{Q}$, and one can check that h is Borel measurable. [In fact, to prove the measurability of h, note that $t(x,x')=1_{\{x<x'\}}\sum_{t=1}^{\infty} r_t$, where r_t is the indicator function that equals 1 when $s_k\in[x,x']^c$ for all integers $k<t$ and $s_t\in[x,x']$ and zero otherwise (with $r_1=1$). Because $1_{\{x>x'\}}$ and r_t (for each t) are Borel measurable, we obtain the measurability of $t(x,x')$; the measurability of h then follows from $h(x,x')=s_{t(x,x')}=\sum_{k=1}^{\infty} s_k 1_{\{t(x,x')=k\}}$.]

Next, we define the random variable $h(\omega)=h(\bar{x}(\omega),\underline{x}(\omega))$ (note the abuse of notation!). Since $\bar{x}$ and $\underline{x}$ are both F_∞-measurable and $h(x,x')$ is Borel measurable, we obtain that $h(\omega)$ is F_∞-measurable. We have therefore constructed an F_∞-measurable random variable $h(\omega)$ taking values in $\bar{Q}$ and such that on B, $\underline{x}(\omega)<h(\omega)<\bar{x}(\omega)$. By replacing $\bar{x}(\omega)$ with $h(\omega)$ and repeating the preceding construction, we obtain an F_∞-measurable random variable $h'(\omega)$ taking values in $\bar{Q}$ and such that on B, $\underline{x}(\omega)<h'(\omega)<h(\omega)$.

We now need some notation and a few definitions to construct the function g. Let

$$\bar{1}_t=1_{\{h\leq x_t\leq 1\}};\qquad \bar{N}_k=\begin{cases}\inf\{n:\sum_{t=1}^{n}\bar{1}_t=k\} & \text{if well-defined,}\\ \infty & \text{otherwise;}\end{cases}$$

$$\bar{X}_k=\frac{1}{k}\sum_{t=1}^{\bar{N}_k} x_t\bar{1}_t;\qquad \bar{Y}_k=\frac{1}{k}\sum_{t=1}^{\bar{N}_k} y_t\bar{1}_t;\qquad \bar{\epsilon}_k=\frac{1}{k}\sum_{t=1}^{\bar{N}_k}\bar{\epsilon}_t\bar{1}_t. \tag{5.8}$$

Define $\underline{N}_k$, $\underline{X}_k$, $\underline{Y}_k$, and $\underline{\epsilon}_k$ in a manner similar to $\bar{N}_k$, $\bar{X}_k$, $\bar{Y}_k$, and $\bar{\epsilon}_k$ but replace $\bar{1}_t$ with $\underline{1}_t=1_{\{0\leq x_t\leq h'\}}$. Finally, define $g:\Omega\to\Theta$ by $g(\omega)=(0,0)$ on B^c and $g(\omega)=(g^\alpha(\omega),g^\beta(\omega))$ on B, where (dropping the ω's for clarity)

$$g^\alpha=\lim_{k\to\infty}\frac{\bar{X}_k\cdot\underline{Y}_k-\underline{X}_k\cdot\bar{Y}_k}{\bar{X}_k-\underline{X}_k}\quad\text{and}\quad g^\beta=\lim_{k\to\infty}\frac{\bar{Y}_k-\underline{Y}_k}{\bar{X}_k-\underline{X}_k}. \tag{5.9}$$

The remainder of the proof is devoted to showing that the random variable g just constructed is well defined and satisfies the conclusions of the lemma.

Recall that $B=\{\omega:\underline{x}(\omega)<\bar{x}(\omega)\}$. On B, notice that $\underline{X}_k<h'<h<\bar{X}_k$, $\bar{N}_k(\omega)<\infty$, and $\underline{N}_k(\omega)<\infty$ for all k sufficiently large. Because $y_t=\alpha+\beta x_t+\epsilon_t$,

$$\frac{1}{k}\sum_{t=1}^{\bar{N}_k} y_t\bar{1}_t=\alpha\frac{1}{k}\sum_{t=1}^{\bar{N}_k}\bar{1}_t+\beta\frac{1}{k}\sum_{t=1}^{\bar{N}_k} x_t\bar{1}_t+\frac{1}{k}\sum_{t=1}^{\bar{N}_k}\epsilon_t\bar{1}_t. \tag{5.10}$$

One can check that $(1/k)\sum_{t=1}^{\bar{N}_k}\bar{1}_t=1$ for all k. Hence, (5.10) becomes

$$\bar{Y}_k=\alpha+\beta\bar{X}_k+\bar{\epsilon}_k. \tag{5.11}$$

Similarly, one can show that

$$\underline{Y}_k = \alpha + \beta \underline{X}_k + \underline{\epsilon}_k. \tag{5.12}$$

Solving (5.11) and (5.12) for β yields

$$\beta = \frac{(\bar{Y}_k - \underline{Y}_k) - (\bar{\epsilon}_k - \underline{\epsilon}_k)}{\bar{X}_k - \underline{X}_k}. \tag{5.13}$$

Let A be the set where the conclusion of Lemma 5.1 holds, so that $P(A) = 1$. For any fixed $\theta = (\alpha, \beta)$, let A_θ be the set of ω's in A whose first coordinate is θ (recall $\Omega = \Theta\Omega'$). On the set A, since h and h' are both rational numbers, both $\bar{\epsilon}_k$ and $\underline{\epsilon}_k$ tend to 0 as $k \to \infty$. Hence, on $A_\theta \cap B$, taking limits on (5.13) leads to

$$\beta = \lim_{k \to \infty} \frac{\bar{Y}_k - \underline{Y}_k}{\bar{X}_k - \underline{X}_k}. \tag{5.14}$$

By a procedure similar to that used in deriving (5.14), one obtains, on $A_\theta \cap B$,

$$\alpha = \lim_{k \to \infty} \frac{\bar{X}_k \cdot \underline{Y}_k - \underline{X}_k \cdot \bar{Y}_k}{\bar{X}_k - \underline{X}_k}. \tag{5.15}$$

From (5.14) and (5.15), for all ω in $A_\theta \cap B$, $g(\omega) = \theta$. But $P_\theta(A_\theta) = 1$. Hence, $P_\theta(\{\omega : g(\omega) = \theta) \cap B) = P_\theta(\{\omega : g(\omega) = \theta) \cap A_\theta \cap B) = P_\theta(A_\theta \cap B) = P_\theta(B)$. Since, clearly, $g \in F_\infty$, this completes the proof of Lemma 5.3.

Q.E.D.

The final step in the proof of Proposition 4.3 involves showing that on the set where there exists a consistent estimator for the true parameter, the posterior distribution will converge to point mass on the true parameter. Since Lemma 5.3 implied the existence of a consistent estimator on the set B, Lemma 5.4 concludes the proof of Proposition 4.3.

Lemma 5.4. *Suppose there exists an F_∞-measurable function g and a set $B \in F_\infty$ such that for μ_0-a.e. θ in Θ,*

$$P_\theta(\{\omega : g(\omega) = \theta) \cap B) = P_\theta(B).$$

Then for μ_θ-a.e. θ in Θ,

$$P_\theta(\{\omega : \mu_\infty(\omega) = 1_\theta\} \cap B) = P_\theta(B), \tag{5.16}$$

where 1_θ is the point mass at θ.

Proof: To make things precise, in particular to indicate that the true parameter can be considered a random variable (with distribution μ_0), let

$\pi(\omega) \in \Theta$ be the projection of $\omega \in \Omega = \Theta\Omega'$ onto its first coordinate, Θ. Define $C = \{\omega: \mu_t(\omega) \to 1_{\pi(\omega)}$ in the weak topology$\}$ and $C_\theta = \{\omega: \mu_t(\omega) \to 1_\theta$ in the weak topology$\}$.

We seek to show that for μ_0-a.e. θ in Θ,

$$P_\theta(C_\theta \cap B) = P_\theta(B). \tag{5.17}$$

Now,

$$P(C \cap B) = \int_\Theta P_\theta(C_\theta \cap B)\mu_0(d\theta) \quad \text{and} \quad P(B) = \int_\Theta P_\theta(B)\mu_0(d\theta). \tag{5.18}$$

To prove (5.17), it suffices to prove that

$$P(C \cap B) = P(B). \tag{5.19}$$

For if (5.19) holds, using (5.18) yields

$$\int_\Theta P_\theta(C_\theta \cap B)\mu_0(d\theta) = \int_\Theta P_\theta(B)\mu_0(d\theta), \tag{5.20}$$

which implies, since $P_\theta(C_\theta \cap B) < P_\theta(B)$ for all θ, that (5.17) holds.

We have shown in Proposition 3.1 that $P(\{\omega: \mu_t(\omega) \to \mu_\infty(\omega)$ in the weak topology$\}) = 1$. Hence, to prove (5.19), we need to show that if D is any Borel subset of Θ, and we denote by $\mu_\infty(D)_{(\omega)}$ the mass that the measure $\mu_\infty(\omega)$ assigns to the set D, then

$$\mu_\infty(D)_{(\omega)} 1_{\{\omega \in B\}} = 1_{\{\pi(\omega) \in D\}} 1_{\{\omega \in B\}} \ P\text{-a.e.} \tag{5.21}$$

Using the definition of g and B,

$$\begin{aligned} P(\{\omega: g(\omega) = \pi(\omega)\} \cap B) &= \int P_\theta(\{\omega \mid g(\omega) = \theta\} \cap B)\mu_0(d\theta) \\ &= \int P_\theta(B)\mu_0(d\theta) = P(B). \end{aligned} \tag{5.22}$$

Noting that $B \in F_\infty$ and $g \in F_\infty$, if D is any Borel subset of Θ (by definition of μ_∞, dropping the ω's for ease of exposition),

$$\begin{aligned} \mu_\infty(D)_{(\omega)} 1_{\{\omega \in B\}} &= E[1_{\{\pi \in D\}} \mid F_\infty] 1_B \\ &= E[1_{\{\pi \in D\}} 1_B \mid F_\infty] \quad [\text{since } B \in F_\infty] \\ &= E[1_{\{g \in D\}} 1_B \mid F_\infty] \quad [\text{using } (5.22)] \\ &= 1_{\{g \in D\}} 1_B \quad [\text{since } g \text{ and } B \in F_\infty] \\ &= 1_{\{\pi \in D\}} \cdot 1_B \ P\text{-a.e.} \quad [\text{using } (5.22)]. \end{aligned}$$

This proves (5.19) and completes the proof. Q.E.D.

Proof of Proposition 4.4

Define on B^c,

$$x(\omega)=\lim_{t\to\infty} x_t(\omega).$$

We will indicate below that for P-a.e. ω in B^c,

$$y(\omega)=\lim_{n\to\infty}\frac{1}{n}\sum_{t=1}^{n} y_t(\omega)$$

is well defined. Let $M(\omega)=\{(\alpha',\beta')\in\Theta:\alpha'+\beta'x(\omega)=y(\omega)\}$. The proof of the proposition is complete if we show that for P-a.e. ω in B^c,

$$\pi(\omega)\in M(\omega) \quad\text{and}\quad \mu_\infty(M(\omega))_{(\omega)}=1, \tag{5.23}$$

where for any Borel set D in Θ, we define $\mu_\infty(D)_{(\omega)}$ to be the mass assigned to the set D by the probability measure $\mu_\infty(\omega)$.

Remark: First note that for each ω, $M(\omega)$ is a closed subset of Θ and is hence a Borel subset of Θ. Next,

$$\{\omega\in B^c:\pi(\omega)\in M(\omega)\}=\{\omega\in B^c:\alpha(\omega)+\beta(\omega)x(\omega)=y(\omega)\}$$

is clearly in F since B^c, $\pi(\omega)$, $x(\omega)$, and $y(\omega)$ are in F; hence, the random variable $\mu_\infty(M(\omega))_{(\omega)}=E[1_{\{\pi(\omega)\in M(\omega)\}}\mid F_\infty\}$ is F-measurable. The expressions in (5.23) are therefore all well defined.

Since $y_t(\omega)=\alpha(\omega)+\beta(\omega)x_t(\omega)+\epsilon_t(\omega)$, where $(\alpha(\omega),\beta(\omega))=\pi(\omega)$,

$$\frac{1}{n}\sum_{t=1}^{n} y_t(\omega)=\alpha(\omega)+\beta(\omega)\frac{1}{n}\sum_{t=1}^{n} x_t(\omega)+\frac{1}{n}\sum_{t=1}^{n}\epsilon_t(\omega). \tag{5.24}$$

From the strong law of large numbers,

$$\lim_{n\to\infty}\frac{1}{n}\sum_{t=1}^{n}\epsilon_t=0 \;\; P\text{-a.e.}$$

(see, e.g., Chung 1974, Theorem 5.4.2, p. 126). Hence, taking limits in (5.24), we obtain, for P-a.e. ω in B^c,

$$y(\omega)=\alpha(\omega)+\beta(\omega)x(\omega) \;\; P\text{-a.e.} \tag{5.25}$$

[Notice that this implies that P-a.e. $y(\omega)$ is well-defined whenever $x(\omega)$ is.] From (5.25), $\pi(\omega)\in M(\omega)$ for P-a.e. ω in B^c.

Next, $\pi(\omega)\in M(\omega)$ for P-a.e. ω in B^c implies that $1_{\{\pi\in M(\omega)\}}\cdot 1_{B^c}=1_{B^c}$ P-a.e.; hence, noting that $1_{B^c}\in F_\infty$,

$$\mu_\infty(M(\omega))_{(\omega)}\cdot 1_{B^c}=E[1_{\{\pi(\omega)\in M(\omega)\}}\,|\,F_\infty]\cdot 1_{B^c}=E[1_{\{\pi(\omega)\in M(\omega)\}}\cdot 1_{B^c}\,|\,F_\infty]$$
$$=E[1_{B^c}\,|\,F_\infty]=1_{B^c}\ P\text{-a.e.} \tag{5.26}$$

This shows that (5.23) holds and concludes the proof. Q.E.D.

Appendix

Bayes rule operator

Let $P(dy_t\,d\theta\,|\,x_t,\mu_{t-1})$ be the joint distribution on $R^1\times\Theta$ obtained as follows: An element θ in Θ is first chosen according to the probability μ_{t-1}; then, given this chosen value of $\theta=(\alpha,\beta,\phi)$, y_t is chosen according to the relation $y_t=\alpha+\beta x_t+\epsilon_t$, where ϵ_t has the distribution $p(\cdot\,|\,\phi)$. Next, define $P(dy_t\,|\,x_t,\mu_{t-1})$ to be the marginal distribution of $P(dy_t\,d\theta\,|\,x_t,\mu_{t-1})$ on R^1. We now apply Parthasarathy (1967, Chapter V, Theorem 8.1) to obtain the existence of a conditional probability measure on Θ, $\Gamma(d\theta\,|\,x_t,y_t,\mu_{t-1})$, which, for fixed (x_t,μ_{t-1}), is measurable in y_t, and where

$$P(dy_t\,d\theta\,|\,x_t,\mu_{t-1})=P(dy_t\,|\,x_t,\mu_{t-1})\cdot\Gamma(d\theta\,|\,x_t,y_t,\mu_{t-1}).$$

The conditional probability $\Gamma(d\theta\,|\,x_t,y_t,\mu_{t-1})$ defines the Bayes rule operator $\Gamma(x_t,y_t,\mu_{t-1})$.

Random variables $\{x_t,y_t,\mu_t\}$

We now provide technical details behind construction of the $\{x_t,y_t,\mu_t\}$ processes. Recall $(\Omega,F,P)=(\Theta,\tilde{\Theta},\mu_0)(\Omega',F',P')$. Any policy π generates a sequence of random variables $\{(x_t,y_t,\mu_t\}_{t=1}^\infty$ on (Ω,F,P) as follows: First consider $\{\epsilon_t\}$ as a stochastic process on (Ω,F,P) rather than on (Ω',F',P') by $\epsilon_t(\omega)=\epsilon_t(\omega')$, where ω' is the second coordinate of ω (recall $\Omega=\Theta\times\Omega'$). Here, μ_0 is given a prior; define $x_1(\omega)=\pi_0(\mu_0)$, $y_1(\omega)=\alpha+\beta x_1(\omega)+\epsilon_1(\omega)$, and $\mu_1(\omega)=\Gamma(x_1(\omega),y_1(\omega),\mu_0)$, where α and β are obtained from the first coordinate of ω (recall $\Omega=\Theta\times\Omega'$). Since both π_0 and Γ are Borel functions (recall Γ is continuous), observe that x_1, y_1, and μ_1 are (Borel measurable) random variables on (Ω,F,P).

Next, suppose that the random variables x_i, y_i, and μ_i have been defined for $i=1,\dots,t-1$; then we may define, inductively, x_t,y_t,μ_t by putting $h_t(\omega)=(\mu_0;\,(x_1(\omega),y_1(\omega),\mu_1(\omega)),\,\dots,\,(x_{t-1}(\omega),y_{t-1}(\omega),\mu_{t-1}(\omega)))$ and

$$x_t(\omega)=\pi_t(h_t(\omega)),\qquad y_t(\omega)=\alpha+\beta x_t(\omega)+\epsilon_t(\omega),$$
$$\mu_t(\omega)=\Gamma(x_t(\omega),y_t(\omega),\mu_{t-1}(\omega)).$$

Since π_t is Borel measurable, x_t,y_t,μ_t are (measurable) random variables.

REFERENCES

Anderson, T. W. and J. Taylor. 1976. "Some Experimental Results on and Statistical Properties of Least Squares Estimates in Control Problems." *Econometrica* 44: 1289–1302.

Billingsley, P. 1968. *Convergence of Probability Measures.* New York: Wiley.

Blackwell, D. 1965. "Discounted Dynamic Programming." *Annals of Mathematical Statistics* 36: 2226–35.

Chow, G. C. 1975. *Analysis and Control of Dynamic Economic Systems.* New York: Wiley.

1981. *Econometric Analysis by Control Methods.* New York: Wiley.

Chung, K. L. 1974. *A Course in Probability Theory,* 2d ed. New York: Academic.

Harkema, R. 1975. "An Analytical Comparison of Certainty Equivalence and Sequential Updating." *Journal of the American Statistical Association* 70: 348–50.

Jordan, J. S. 1985. "The Strong Consistency of the Least Squares Control Rule and Parameter Estimates." Manuscript, University of Minnesota.

Kiefer, N. and Y. Nyarko. 1986. "Optimal Control of a Linear Process with Learning." Manuscript, Cornell University.

Kushner, H. and D. Clark. 1978. *Stochastic Approximation Methods for Constrained and Unconstrained Systems.* New York: Springer-Verlag.

Maitra, A. 1968. "Discounted Dynamic Programming in Compact Metric Spaces." *Sankhya, Series A* 30: 211–16.

Parthasarathy, K. 1967. *Probability Measures on Metric Spaces.* New York: Academic.

Prescott, E. 1972. "The Multiperiod Control Problem Under Uncertainty." *Econometrica* 40: 1043–58.

Schal, M. 1979. "On Dynamic Programming and Statistical Decision Theory." *Annals of Statistics* 7: 432–45.

Schwartz, L. 1965. "On Bayes Procedures." *Z. Wahrscheinlichkeits-theorie* 4: 10–26.

Taylor, J. B. 1974. "Asymptotic Properties of Multiperiod Control Rules in the Linear Regression Model." *International Economic Review* 15: 472–84.

Theil, H. 1958. *Economic Forecasts and Policy.* Amsterdam: North-Holland.

Zellner, A. 1971. *An Introduction to Bayesian Inference in Econometrics.* New York: Wiley.