FACULTY WORKING PAPER 91-0183

Continuity of Posterior Revision and Bayesian Dynamic Programming

Mark Feldman
Department of Economics
University of Illinois

Bureau of Economic and Business Research
College of Commerce and Business Administration
University of Illinois at Urbana-Champaign

BEBR

FACULTY WORKING PAPER NO. 91-0183

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

December 1991

Continuity of Posterior Revision and Bayesian Dynamic Programming

Mark Feldman

Department of Economics

ABSTRACT

The study of the Bayesian learning dynamics with optimizing agents is facilitated by embedding the individual decision problems within a standard dynamic programming environment. This may be accomplished by augmenting the state space to include the set of possible beliefs over a parameter space Θ representing the set of a priori possible model specifications. But to apply the standard dynamic programming results on the existence of an optimal policy, it is necessary to establish that the distribution of the posterior beliefs is a continuous function (in the topology of weak convergence of probability measures) of the current state variable and chosen action. We develop necessary and sufficient conditions for the continuity of the posterior distribution map.

1. INTRODUCTION

The study of the Bayesian learning dynamics with optimizing agents is facilitated by embedding the individual decision problems in a standard dynamic programming environment. As first demonstrated by Hinderer (1970) and Reider (1975), this may accomplished by augmenting the state space to include the set of possible beliefs over a parameter space Θ representing the set of a priori possible specifications. But to apply the standard dynamic programming results on the existence of an optimal policy as in Blackwell (1965) or Maitra (1968) it is necessary to establish that the distribution of the next period state variable is a continuous function (in the topology of weak convergence of probability measures) of the current state variable and chosen action. In the context of a Bayesian dynamic programming problem this requires that the Bayesian's probability distribution over next period beliefs vary continuously with her current beliefs and chosen action. The object of this paper is to find a minimal set of sufficient conditions which yield the requisite continuity.

There has been considerable recent literature in which agent(s) are modelled as solving Bayesian dynamic programming problems, including work by Aghion, Bolton, Harris and Julien (1991), Easley and Kiefer (1988), Easley and Kiefer (1989), Feldman and McLennan (1989), Feldman and Spagat (1991), Kiefer and Nyarko (1988), Kiefer and Nyarko (1989), McLennan (1984), and Nyarko (1991). Invariably there is a parameter space Θ, an action (or action/state) space X, and an outcome space Y. Let $P(Y)$, $P(\Theta)$ and $P^2(\Theta)$ denote respectively the space of probability measures on Y, Θ and $P(\Theta)$. Suppose that for parameter $\theta \in \Theta$ and action $x \in X$, the distribution of outcomes is $\Psi(\theta, x) \in P(Y)$ and that the distribution of posterior beliefs for prior $\mu \in P(\Theta)$ is $\varphi(\mu, x) \in P^2(\Theta)$. To attain the requisite continuity of the map φ, a frequent assumption is that with respect to some reference measure ν on Y, $\Psi(\theta, x)$ has a Radon-Nikodym derivative $f(\theta, x, \cdot)$ with f jointly continuous in θ, x, and y.

Unfortunately, requiring that the map Ψ has a representation by a jointly continuous density f is highly restrictive. The objective of this paper is to determine minimal conditions on Ψ which are sufficient for the continuity of φ. Theorem 3.2, the main technical result of this paper, establishes that if Θ, X and Y are separable metric spaces, then φ is continuous if Ψ is continuous when $P(\Theta)$ is endowed with the total variation (i.e., norm) topology. Examples are provided to demonstrate that if Ψ is merely setwise continuous, then φ may fail to be continuous. However, as discussed in Section 5, modulo some relabelling, there is no intrinsic difficulty if some of the components of the outcome space are deterministic functions of θ and x.

The organization of the paper is as follows. Definitions and notational conventions are provided in Section 2. The assumptions regarding the spaces Θ, X, and Y, and a statement of Theorem 3.2 are provided in Section 3, along with counterexamples to the conjecture that setwise convergence suffices for the continuity of φ. Section 4 contains needed lemmas which are proved in the Appendix, and the proof of Theorem 3.2. In Section 5 we conclude with some remarks on application and interpretation of the results.

2. DEFINITIONS AND NOTATION

The set of real numbers is denoted by R. If $A \subset X$, the indicator function of A is I_A. Let (S, d) be a metric space. The Borel σ-field is denoted by $B(S)$, the set of Borel probability measures is $P(S)$, and the set of Borel measures is $M(S)$. For $s \in S$, the Dirac measure $\delta_s \in P(S)$ is defined by $\delta_s(A) = I_A(s)$ for $A \in B(S)$. f: $S \to R$ is a *Lipschitz* function if for some $k < \infty$, $\sup_{s \neq t} \frac{|f(s) - f(t)|}{d(s,t)} < k$. The Lipschitz seminorm $\|f\|_L$ is defined by $\|f\|_L = \sup_{s \neq t} \frac{|f(s) - f(t)|}{d(s,t)}$. If f is a bounded Lipschitz function, the *bounded Lipschitz* norm is $\|f\|_{BL} = \|f\|_L + \|f\|_\infty$ where $\|f\|_\infty$ denotes the usual sup norm. BL(S, d) is the set of all real-valued, bounded Lipschitz functions on (S, d). Endowed with the bounded Lipschitz norm, BL(S, d) is a Banach space (see e.g. [Dudley, 1989, Section 11.2 #54]).

The *dual bounded Lipschitz* metric on $P(S)$, denoted by β_S (or by β is there is no ambiguity), is defined by $\beta_S(P, Q) = \sup \{|\int f\, dP - \int f\, dQ|: \|f\|_{BL} \leq 1\}$, for $P, Q \in P(X)$. If S is separable, β_S metrizes the topology of weak convergence on $P(X)$. Further details on the properties of the dual bounded Lipschitz metric can be found in Dudley (1966) and Dudley (1989). The *total variation norm* on $M(S)$, denoted by τ_S (or by τ), is defined for $\eta, \gamma \in M(S)$ by $\tau_S(\eta, \gamma) = \sup_{A \in B(S)} |\eta(A) - \gamma(A)|$. The restriction of τ_S to $P(S)$ is also denoted by τ_S. The sequence $\{\mu_n\}$ converges *setwise* to $\mu \in M(S)$ if $\mu_n(A) \to \mu(A)$ for all $A \in B(S)$.

3. STATEMENT OF THEOREM AND INADEQUACY OF SETWISE CONVERGENCE

3.1. *Statement of Theorem*

Consider a Bayesian decision-maker with prior belief μ_0 on a parameter space Θ who must choose an action $x \in X$. Given a parameter $\theta \in \Theta$ and her action $x \in X$, the decision-maker will observe an outcome $y \in Y$, the realization of a random element with (unknown) law $\Psi(\theta, x) \in P(Y)$. We make the following assumptions:

ASSUMPTION 1: (Θ, d_Θ), (X, d_X), and (Y, d_Y) are separable metric spaces with respective Borel σ-fields $B(\Theta)$, $B(X)$, and $B(Y)$.

ASSUMPTION 2: The mapping $\Psi: \Theta \times X \to (P(Y), \tau_Y)$ is continuous where τ_Y is the total variation metric on $P(Y)$.

Given a prior μ and action x, the induced probability measure on the outcome space $(Y, B(Y))$ is $\phi(\mu, x)$ defined by $\phi(\mu, x)(A) = \int_\Theta \Psi(\theta, x)(A)\, \mu(d\theta)$ for $A \in B(Y)$. The corresponding probability on $(\Theta \times Y, B(\Theta) \times B(Y))$ is $\Phi(\mu, x)$ defined by $\Phi(\mu, x)(A \times B) = \int_A \Psi(\theta, x)(B)\, \mu(d\theta)$, for $A \in B(\Theta)$ and $B \in B(Y)$. The posterior belief given prior μ, action x, and outcome y is denoted as $\Gamma(\mu, x, y)$. Dynkin and Yushkevich (1979) establish the joint measurability of Γ by extending a standard proof of the existence of a regular version of conditional probability. The formal statement of their result is:

LEMMA 3.1. There is a measurable function $\Gamma: P(\Theta) \times X \times Y \to P(\Theta)$ such that for all $(\mu, x) \in P(\Theta) \times X$, $\Gamma(\mu, x, \cdot)$ is a regular version of $\Phi(\mu, x)(\cdot \,\|\, B(Y))$.

Proof. Follows directly from discussion on p. 263 of Dynkin and Yushkevich (1979). ■

For given $(\mu, x) \in P(\Theta) \times X$ and Borel set $D \subset P(\Theta)$, the Bayesian's prior probability of $\{\Gamma(\mu, x, y) \in D\}$ is $\varphi(\mu, x)$ where the mapping $\varphi: P(\Theta) \times X \to P^2(\Theta)$ is defined by $\varphi(\mu, x)(D) = \int_Y I_D(\Gamma(\mu, x, y))\, \phi(\mu, x)(dy)$. The main technical result of this paper is:

THEOREM 3.2. Endowing $P(\Theta)$ and $P^2(\Theta)$ with the topology of weak convergence, $\varphi: P(\Theta) \times X \to P^2(\Theta)$ is continuous.

Proof. See Section 4.

3.2. *CounterExamples To Theorem* 3.2 *When* Ψ *is Setwise But Not Total-Variation Continuous*

CounterExample 1: We first provide an example where Ψ is only setwise continuous in x (and jointly setwise continuous) and φ is discontinuous. Suppose $\Theta = \{\theta_0, \theta_1\}$, $X = [0, 1]$, $Y = [0, 1)$ and m is Lebesgue measure on $(Y, B(Y))$. Let δ_i denotes the Dirac measure with mass on θ_i. Suppose that $\Psi(\theta_1, 0) = m$, and $\Psi(\theta_0, x) = m$ for all $x \in X$, (So no learning ocurs if $x = 0$). Let g be the density function corresponding to m.

We now proceed to specify $\Psi(\theta_1, x)$ for $x > 0$. For $x \in (0, 1]$, define $k(x) = 2\cdot[x^{-1}]$ where $[c]$ is the largest integer less than or equal to c. For $j = 1, 2, \ldots, k(x)$, define the interval $I_j(x) = [\backslash F(j - 1, k(x)), \frac{j}{k(x)})$; and for $x > 0$, define $\Psi(\theta_1, x)$ as the probability measure with density $g_x(y) = 2$ for $y \in I_j(x)$ when j is even and $g_x(y) = 0$ otherwise. It is well-known that as $x \to 0$, g_x converges weakly to g in $L^1(Y, B(Y), m)$. It follows that $\Psi(\theta_1, \cdot)$ is continuous at $(\theta_1, 0)$ when $P(Y)$ is endowed with the topology of setwise convergence.

We now demonstrate that φ is not continuous. Choose μ_0 such that $\mu_0(\{\theta_0\}) = a \in (0, 1)$. Then $\Gamma(\mu_0, 0, y) = \mu_0$ for all $y \in Y$ and $\varphi(\mu_0, 0) = \delta_{\mu_0}$. But for $x > 0$, $\Gamma(\mu_0, x, y) = \delta_1$ for $y \in I_j(x)$ with j odd. So $\varphi(\mu_0, x)(\{\delta_0\}) = \frac{a}{2}$ and as $x \to 0$, and $\varphi(\mu_0, x) \not\Rightarrow \varphi(\mu_0, 0)$.

CounterExample 2: An example is now provided where map $\theta \to \Psi(\theta, x)$ is setwise continuous but $\mu_n \Rightarrow \mu_0$ does not imply $\varphi(\mu_n, x) \Rightarrow \varphi(\mu_n, x)$. Let $X = \{x_0\}$, $\Theta = [0, 1/2] \cup \{1\}$ and $Y = [0, 1]$. Suppose $\Psi(0, x_0)$ and $\Psi(1, x_0)$ are defined by $\Psi(0, x_0) = \delta_0$, $\Psi(1, x_0)(\{0\}) = 1/2$ and $\Psi(1, x_0)(\{1\}) = 1/2$. For $\theta \in (0, 1/2]$, define $\Psi(\theta, x_0)$ as having uniform density on $[0, \theta]$. So for any sequence $\theta_n \to \theta_0$, $\Psi(\theta_n, x_0) \Rightarrow \Psi(\theta_n, x_0)$; in particular, if $\theta_n \to 0$ then $\Psi(\theta_n, x_0) \to \delta_0 = \Psi(0, x_0)$.

Now define $\mu_0 \in P^2(\Theta)$ by $\mu_0(\{0\}) = \mu_0(\{1\} = 1/2$ and consider the sequence of priors $\{\mu_n\}$ where $\mu_n(\{1\}) = 1/2$ and $\mu_n(\{n^{-1}\}) = 1/2$. Observe that $\phi(\mu_0, x_0)(\{0\}) > 0$ and $\Gamma(\mu_0, x_0, 0)(\{0\}) = 2/3$, and so $\Gamma(\mu_0, x_0, 0) \notin \{\delta_0, \delta_1\}$. But for $n > 0$, the outcome is completely informative; i.e., $\phi(\mu_n, x_0)$ a.s. $\Gamma(\mu_n, x_0, y) \in \{\delta_0, \delta_1\}$. So $\{\varphi(\mu_n, x_0)\}$ doesn't converge to $\varphi(\mu_0, x_0)$.

4. PROOF OF THEOREM

4.1. *Some Lemmas for Probability Measures on Metric Spaces*

To prove the Theorem, several intermediate results need to be established. The first, LEMMA 4.1, is an approximation theorem for probability measures. LEMMA 4.2 is a sort of converse to Scheffe's Theorem. LEMMA 4.3, taken from Dudley, 1989 #54], asserts that if μ is a probability measure on a separable metric space S, then S can be decomposed into a countable collection of disjoint μ-continuity sets of arbitrarily small diameter.

LEMMA 4.1. Suppose (S, d) is a separable metric space, $\varepsilon > 0$, and $\{A_1, A_2, \ldots\}$ is a disjoint family of Borel subsets of S with $\cup A_j = S$ and diam $A_j < \delta = \frac{\varepsilon}{5}$ for all j. Define $B_n = \cup_{j=1}^{n} A_j$. Suppose also that for probability measures μ and η on $(S, B(S))$ there exists an integer k such that $\mu(B_k) > 1 - \delta$ and $\sup_{j \leq k} |\mu(A_j) - \eta(A_j)| < \delta \cdot k^{-1}$. Then $\beta(\mu, \eta) < \varepsilon$.

Proof. See Appendix. ■

LEMMA 4.2. Suppose (S, d) is a metric space and $\{\nu_j\}$ and $\{\eta_j\}$ are sequences of finite Borel measures on $(S, B(S))$ such that: (i) $\nu_j \to \nu_0$ in total variation, (ii) $\eta_j \to \eta_0$ in total variation, and (iii) $\eta_j << \nu_j$ for $j = 0, 1, \ldots$. Suppose also that $\exists$ a sequence $\{f_j\}$ of versions of $\frac{d\eta_j}{d\nu_j}$ which are uniformly bounded in $L^\infty(S, B(S), \nu_0)$ and let f be a version of $\frac{d\eta_0}{d\nu_0}$. Then $f_j \to f$ in $L^1(S, B(S), \nu_0)$ and in ν_0 measure.

Proof. See Appendix. ■

LEMMA 4.3. For any separable metric space (S, d), $\varepsilon > 0$, and $P \in P(S)$, there exists a sequence $\{A_j\}$ of disjoint P-continuity sets with $\cup A_j = S$ and diam $A_j < \varepsilon$ for all j.

Proof. This is Lemma 11.7.3 of Dudley (1989). ■

To motivate LEMMA 4.4 suppose that S is a metric space, $A \in B(S)$ and $\mu_n \Rightarrow \mu_0 \in P(S)$ with $\mu_0(A) > 0$. We can define a sequence of conditional measures η_k on (S, B(S)) for k = 0, 1, 2, ..., by $\eta_k(C) = \frac{\mu_k(C \cap A)}{\mu_k(A)}$. In general, η_k doesn't necessarily converge, and if it does converge, the limit need not be η_0. But if A is a μ_0 contiinuity set, then $\eta_k \Rightarrow \eta_0$.

LEMMA 4.4. Suppose S is a metric space, $\mu_n \Rightarrow \mu_0 \in P(S)$, and A is a μ_0-continuity set with $\mu_0(A) > 0$. For k = 0, 1, ..., define the probability measure η_k on (S, B(S)) by $\eta_k(C) = \frac{\mu_k(C \cap A)}{\mu_k(A)}$. Then $\eta_k \Rightarrow \eta_0$.

Proof. Let $F \subset S$ be closed; it suffices to show that limsup $\eta_k(F) \le \eta_0(F)$. Since $\mu_k \Rightarrow \mu_0$ and $F \cap \text{cl } A$ is closed, limsup $\mu_k(F \cap \text{cl } A) \le \mu(F \cap \text{cl } A) = \mu(F \cap A)$, where the equality follows from A being a μ_0-continuity set. Since $\mu_k(A) \to \mu(A) > 0$, we have:

$$\text{limsup } \eta_k(F) \le \text{limsup } [\frac{\mu_k(F \cap \text{cl } A)}{\mu_k(A)}] \le \frac{\mu(F \cap \text{cl } A)}{\mu(A)} = \frac{\mu(F \cap A)}{\mu(A)} = \eta_0(F).$$ ■

Lemmas 4.5 and 4.6 are respectively variants of Scheffe's Theorem [e.g., [Billingsley (1986), Theorem 16. 11]] and a generalized Dominated Convergence Theorem in Royden (1988, Proposition 11.18). Since I was unable to locate a reference containing needed versions of these theorems, for completeness proofs are provided.

LEMMA 4.5. Let (S, F, ν) be a measure space. Suppose: (i) for all $A \in F$, $\lambda_n(A) = \int_A f_n(s)\, \nu(ds)$ and $\lambda(A) = \int_A f(s)\, \nu(ds)$ for densities f_n and f, and (ii) for n = 1,2, ..., $\lambda_n(S) = \lambda(S) < \infty$. Then $f_n \to f$ in ν measure is a necessary and sufficient condition for $\tau_S(\lambda_n, \lambda) \to 0$.

Proof. See Appendix. ■

LEMMA 4.6. Let (S, F) be a measurable space and $\{\nu_n\}$ a sequence of measures that converge setwise to a finite measure ν. Suppose $\{f_n\}$ is a sequence of uniformly bounded, real-valued functions on S with $f_n \to f$ in ν measure. Then: $\int f_n(s)\, \nu_n(ds) \to \int f(s)\, \nu(ds)$.

Proof. See Appendix. ■

4.2. *Proof of Theorem and Auxiliary Results*

PROPOSITION 4.7: If A is a μ_0-continuity set, then the map $\phi_A: P(\Theta) \times X \to (M(Y), \tau_Y)$ defined by $\phi_A(\mu, x)(B) = \Phi(\mu, x)(A \times B) = \int_A \Psi(\theta, x)(B)\, \mu(d\theta)$ is continuous at (μ_0, x) for all $x \in X$.

Proof. See Appendix. ■

COROLLARY 4.8: Endowing $P(Y)$ with the total variation topology, the map $\phi: P(\Theta) \times X \to P(Y)$ is continuous.

Proof. For all $\mu \in P(\Theta)$, Θ is a μ-continuity set and so ϕ_Θ is total variation continuous at all $(\mu, x) \in P(\Theta) \times X$. Since $\phi_\Theta(\mu, x) = \phi(\mu, x)$, the result follows. ■

Proof of THEOREM 3.2. Let $g: P(\Theta) \to R$ be a bounded, continuous function. It is necessary to verify that for any sequence $\{(\mu_n, x_n)\}$ with $\mu_n \Rightarrow \mu_0$ and $x_n \to x_0$, that $\int g(\mu)\, \varphi(\mu_n, x_n)(d\mu) \to \int g(\mu)\, \varphi(\mu_0, x_0)(d\mu)$. Since $\int g(\mu)\, \varphi(\mu_n, x_n)(d\mu) = \int g(\Gamma(\mu_n, x_n, y))\, \phi(\mu_n, x_n)(dy)$ and $\int g(\mu)\, \varphi(\mu_0, x_0)(d\mu) = \int g(\Gamma(\mu_0, x_0, y))\, \phi(\mu_0, x_0)(dy)$, it suffices to demonstrate that for $(\mu_n, x_n) \to (\mu_0, x_0)$, $\int g(\Gamma(\mu_n, x_n, y))\, \phi(\mu_n, x_n)(dy) \to \int g(\Gamma(\mu_0, x_0, y))\, \phi(\mu_0, x_0)(dy)$.

Suppose $(\mu_n, x_n) \to (\mu_0, x_0)$ and pick $\varepsilon > 0$. Defining $\delta = \frac{\varepsilon}{5}$, by Lemma 4.3 there exists a disjoint cover $A_1, A_2, \ldots$ of Θ, with diam $A_j < \delta$, consisting of μ_0-continuity sets. Define $\eta_{j,n}$ and ν_n on $(Y, B(Y))$ by $\eta_{j,n}(C) = \Phi(\mu_n, x_n)(A_j \times C) = \phi_{A_j}(\mu_n, x_n)(C)$ and $\nu_n(C) = \Phi(\mu_n, x_n)(\Theta \times C) = \phi(\mu_n, x_n)(C)$. By Proposition 4.7 and Corollary 4.8, $\eta_{j,n} \to \eta_{j,0}$ and $\nu_n \to \nu_0$ in total variation. Define $f_{j,n}$ as the version of $\frac{d\eta_{j,n}}{d\nu_{j,n}}$ such that $f_{j,n}(y) = \Gamma(\mu_n, x_n, y)(A_j)$ for all $y \in Y$. Define $B_m = \cup_{j=1}^{m} A_j$, and choose k such that: $\phi(\mu_0, x_0)(\{y: \Gamma(\mu_0, x_0, y)(B_k) > 1 - \delta\}) > 1 - \delta/2$. Defining $D = \{y: \Gamma(\mu_0, x_0, y)(B_k) > 1 - \delta\}$, $\phi(\mu_0, x_0)(D) > 1 - \delta/2$.

We now define a set $E_n \subset Y$ such that $\beta(\Gamma(\mu_n, x_n, y), \Gamma(\mu_0, x_0, y)) < \varepsilon$ for $y \in D \cap E_n$. Let $E_n = \{y: \sup_{1 \le j \le k} |\Gamma(\mu_0, x_0, y)(A_j) - \Gamma(\mu_n, x_n, y)(A_j)| < \delta \cdot k^{-1}\}$. If $y \in D \cap E_n$, then $\beta(\Gamma(\mu_n, x_n, y), \Gamma(\mu_0, x_0, y)) < \varepsilon$ by Lemma 4.1. By Lemma 4.2, $f_{j,n} \to f_{j,0}$ in ν_0 measure for all j (and in particular for $j \le k$). So for n sufficiently large, $\nu_0(E_n) > 1 - \delta/2$, $\nu_0(D \cap E_n) > 1 - \delta$, and $\nu_0(\{y: \beta(\Gamma(\mu_n, x_n, y), \Gamma(\mu_0, x_0, y)) < \varepsilon\}) > 1 - \varepsilon$.

Define the functions $h_0: Y \to P(\Theta)$ and $h_n: Y \to P(\Theta)$ by $h_0(y) = \Gamma(\mu_0, x_0, y)$ and $h_n(y) = \Gamma(\mu_n, x_n, y)$. From the above paragraph and Theorem 9.2.2 of Dudley (1989), $h_n \to h_0$ in $\phi(\mu_0, x_0)$ probability (when $P(\Theta)$ is endowed with the topology of weak convergence). Since g is continuous,

$g(h_n) \to g(h_0)$ in $\phi(\mu_0, x_0)$ probability. By Corollary 4.8, $\phi(\mu_n, x_n) \to \phi(\mu_0, x_0)$ in total variation and hence setwise. So by Lemma 4.6:

$$\int g(h_n(y))\, \phi(\mu_n, x_n)(dy) \to \int g(h_0(y))\, \phi(\mu_0, x_0)(dy),$$

or

$$\int_Y g(\Gamma(\mu_n, x_n, y))\, \phi(\mu_n, x_n)(dy) \to \int_Y g(\Gamma(\mu_0, x_0, y))\, \phi(\mu_0, x_0)(dy). \blacksquare$$

5. SOME REMARKS ON APPLICATIONS

In many dynamic programming problems it is natural for some components of the state space be deterministic functions of previous actions and state variables. But in such settings, if the set Y in Assumptions 1 and 2 is identified with the state space, Assumption 2 will not be satisfied and Theorem 3.2 will not be applicable. Fortunately, such difficulties are easily surmounted in a manner which will be briefly sketched.

Consider a Bayesian dynamic programming problem with parameter space Θ, state space S, action space A and transition map $\xi: \Theta \times S \times A \to P(S)$. The generalized or Bayesian state space is $\Sigma = P(\Theta) \times S$. With this generalized state space, the Bayesian optimization problem is a conventional dynamic programming problem. Let $\zeta: \Sigma \times A \to P(\Sigma)$ be the induced generalized transition map. To apply standard results on the existence of optimal policies, it is necessary to be able to verify that ζ is continuous.

In many dynamic programming problems, including those alluded to above, the space S will have a product representation $S = Y \times Z$. Suppose this is the case and let $S_t = (Y_t, Z_t)$ represent the time t state variable with Y_t and Z_t respectively taking values in Y and Z. Suppose additionally that Y_t is a sufficient statistic; that is, conditional upon the past state variable action pair (s_{t-1}, a_{t-1}) and outcome y_t, the distribution of Z_t is independent of Θ. Then all that is required for the distribution of posterior beliefs to be continuous is that the marginal distribution $\xi_Y: \Theta \times S \times A \to P(Y)$ defined by $\xi_Y(\theta, s, a)(B) = \xi(\theta, s, a)(B \times Z)$ be continuous when $P(Y)$ is endowed with the total variation topology. Summarizing, ζ will be continuous if ξ is continuous and the marginal map ξ_Y is total variation continuous.

There still remains the issue of determining if ξ_Y is total variation continuous. This task is simplified if there exists a σ-finite measure λ on $(Y, B(Y))$ such that for all $(\theta, s, a) \in \Theta \times S \times A$, $\xi(\theta, s, a) << \lambda$. Then by Lemma 4.6, a necessary and sufficient condition is that for $(\theta_n, s_n, a_n) \to (\theta_0, s_0, a_0)$ that $\dfrac{d\xi_Y(\theta_n, s_n, a_n)}{d\nu} \to \dfrac{d\xi_Y(\theta_0, s_0, a_0)}{d\nu}$ in ν measure.

In particular, suppose $Y \subset R^m$, λ is Lebesgue measure, and $\xi_Y(\theta, s, a)$ can be represented by a density $f(y, \theta, s, a)$. Then ξ_Y is total variation continuous if and only $(\theta_n, s_n, a_n) \to (\theta_0, s_0, a_0)$ implies that $f(\cdot,\theta_n, s_n, a_n)$ converges in Lebesgue measure to $f(\cdot,\theta_0, s_0, a_0)$.

APPENDIX

Proof of LEMMA 4.1. Select $f \in BL(S, d)$ with $\|f\|_{BL} \le 1$. It suffices to verify that $|\int_S f(s)\,\mu(ds) - \int_S f(s)\,\eta(ds)| < \varepsilon$. Define $a_n = \inf_{s \in A_n} f(s)$ and $b_n = \sup_{s \in A_n} f(s)$. Since $\|f\|_{BL} \le 1$, it follows that: (i) $\sup\{|f(s) - f(t)|: s, t \in A_n\} \le \operatorname{diam} A_n < \delta$, (ii) $b_n - a_n < \delta$, and (iii) $|a_n|, |b_n| \le 1$.

Since,

$$\int_{A_n} f(s)\,\eta(ds) \ge \int_{A_n} a_n\,\eta(ds) = \eta(A_n)\cdot a_n > [\mu(A_n) - \delta\cdot k^{-1}]\cdot a_n,$$

and

$$\int_{A_n} f(s)\,\mu(ds) \le \mu(A_n)\cdot b_n,$$

we have

$$\int_{A_n} f(s)\,\mu(ds) - \int_{A_n} f(s)\,\eta(ds) < \mu(A_n)\cdot b_n - [\mu(A_n) - \delta\cdot k^{-1}]\cdot a_n < (\mu(A_n) + k^{-1})\cdot\delta.$$

Similarly, $\int_{A_n} f(s)\,\eta(ds) - \int_{A_n} f(s)\,\mu(ds) < (\mu(A_n) + k^{-1})\cdot\delta$,

so

$$\left|\int_{A_n} f(s)\,\eta(ds) - \int_{A_n} f(s)\,\mu(ds)\right| < (\mu(A_n) + k^{-1})\cdot\delta,$$

and

$$\left|\int_{B_n} f(s)\,\eta(ds) - \int_{B_n} f(s)\,\mu(ds)\right| < [\mu(B_n) + 1]\cdot\delta \le 2\cdot\delta.$$

The remaining task is to bound $\left|\int_{\sim B_n} f(s)\,\eta(ds) - \int_{\sim B_n} f(s)\,\mu(ds)\right|$. Observe that

$$\eta(\sim B_n) = 1 - \eta(B_n) < 1 - \mu(B_n) + \delta < 2\delta.$$

Since $|f(s)| \le 1$,

$$\left|\int_{\sim B_n} f(s)\,\eta(ds)\right| < 2\delta,\ \left|\int_{\sim B_n} f(s)\,\mu(ds)\right| < \delta, \text{ and } \left|\int_{\sim B_n} f(s)\,\eta(ds) - \int_{\sim B_n} f(s)\,\mu(ds)\right| < 3\delta.$$

Combining these results,

$$\left|\int_S f(s)\,\eta(ds) - \int_S f(s)\,\mu(ds)\right|$$

$$\le \left|\int_{\sim B_n} f(s)\,\eta(ds) - \int_{\sim B_n} f(s)\,\mu(ds)\right| + \left|\int_{B_n} f(s)\,\eta(ds) - \int_{B_n} f(s)\,\mu(ds)\right| < 5\cdot\delta = \varepsilon. \ \blacksquare$$

Proof of LEMMA 4.2. Pick $\varepsilon > 0$, let τ_S denote the total variation metric on $P(S)$, and let $\|\cdot\|_\infty$ denote the $L^\infty(S, B(S), \nu_0)$ norm. Select $M > 0$ such that $\|f_j\|_\infty < M$ for all j, and choose k such that for $n > k$, $\tau_S(\nu_n, \nu_0) < \frac{\varepsilon}{8\cdot M}$ and $\tau(\eta_n, \eta_0) < \varepsilon/4$. Since

$$\int_S |f_n(s) - f(s)|\ \nu_0(ds) \le 2\cdot \sup_{A\in \mathcal{B}(S)} |\int_A f_n(s)\ \nu_0(ds) - \int_A f(s)\ \nu_0(ds)|,$$

and

$$|\int_A f_n(s)\ \nu_0(ds) - \int_A f(s)\ \nu_0(ds)|$$

$$\le |\int_A f_n(s)\ \nu_0(ds) - \int_A f_n(s)\ \nu_n(ds)| + |\int_A f_n(s)\ \nu_n(ds) - \int_A f(s)\ \nu_0(ds)|$$

$$\le 2\cdot M\cdot \tau_S(\nu_0, \nu_n) + |\eta_n(A) - \eta_0(A)|$$

$$\le 2\cdot M\cdot \tau_S(\nu_0, \nu_n) + \tau_S(\eta_n, \eta_0)$$

$$< \varepsilon/2,$$

we conclude that $\int_S |f_n(s) - f(s)|\ \nu_0(ds) < \varepsilon$. Since ε is arbitrary, $f_n \to f$ in $L^1(S, \mathcal{B}(S), \nu_0)$ and hence in ν_0 measure. ■

Proof of LEMMA 4.5. To prove sufficiency, suppose $f_n \to f$ in ν-measure but $\tau_S(\lambda_n, \lambda) \not\to 0$. Then $\exists\ \varepsilon > 0$ and a subsequence $\{f_{n_k}\}$ with $\sup_{A\in F} |\lambda_{n_k}(A) - \lambda(A)| > \varepsilon$. But since $f_{n_k} \to f$ in ν measure, $\exists$ a further subsequence $f_{n_{k_j}} \to f$ ν a.e. [see e.g., Theorem 2.5.3 of Ash (1972)]. From the standard version of Scheffe's Theorem (as in Billingsley (1986)), $\sup_{A\in F} |\lambda_{n_{k_j}}(A) - \lambda(A)| \to 0$, contradicting $\sup_{A\in F} |\lambda_{n_k}(A) - \lambda(A)| > \varepsilon$ for all k.

To prove necessity, suppose $\tau_S(\lambda_n, \lambda) \to 0$. By Theorem 1.1 of Devroye and Gyorfi (1985), $2\cdot\sup_{A\in F} |\lambda_n(A) - \lambda(A)| = \int |f_n(s) - f(s)|\ \nu(ds)$; and by definition, $\tau_S(\lambda_n, \lambda) = \sup_{A\in F} |\lambda_n(A) - \lambda(A)|$. So $f_n \to f$ in $L^1(S, F, \nu)$ and hence in ν measure. ■

Proof of LEMMA 4.6. We first establish that if $f_n \to f$ ν-a.e., then $\int f_n(s)\ \nu_n(ds) \to \int f(s)\ \nu(ds)$. Choose K such that $|f_n(s)| \le K$ for all s. Define $E = \{s\colon f_n(s) \not\to f(s)\}$. Define $g_n, g\colon S \to R$ by $g_n(s) = f_n(s)$ for $s \notin E$ and $g(s) = f(s)$ for $s \notin E$, while $g_n(s) = g(s) = 0$ for $s \in E$. Since $g_n(s) \to g(s)$ for all $s \in S$, by Proposition 11.18 of Royden (1988), $\int_S g_n(s)\ \nu_n(ds) \to \int_S g(s)\ \nu(ds)$. Since $\nu(\{s\colon g(s) \ne f(s)\}) = 0$, $\int_S g(s)\ \nu(ds) = \int_S f(s)\ \nu(ds)$. Furthermore, since $\int_{\sim E} |g_n(s) - f_n(s)|\ \nu_n(ds) = 0$, and $\int_E |g_n(s) - f_n(s)|\ \nu_n(ds) < 2K\cdot\nu_n(\sim E)$,

$$|\int_S g_n(s)\ \nu_n(ds) - \int_S f_n(s)\ \nu_n(ds)| < 0 + 2K\cdot\nu_n(\sim E) \to 0.$$

So,

$$|\int_S f(s)\ \nu(ds) - \int_S f_n(s)\ \nu_n(ds)|$$

$$\le |\int_S g(s)\ \nu(ds) - \int_S g_n(s)\ \nu_n(ds)| + |\int_S g_n(s)\ \nu_n(ds) - \int_S f_n(s)\ \nu_n(ds)| \to 0,$$

verifying that $\int f_n(s)\, \nu_n(ds) \to \int f(s)\, \nu(ds)$ when $f_n \to f$ ν almost everywhere.

Extending the proof to the case where $f_n \to f$ in ν measure, now follows the technique used in the proof of Lemma 6. Suppose that for some $\varepsilon > 0$, $\exists$ a subsequence $\{f_{n_k}\}$ with $\left|\int f_{n_k}(s)\, \nu_{n_k}(ds) - \int f(s)\, \nu(ds)\right| > \varepsilon$. But then $\exists$ a subsubsequence $f_{n_{k_j}} \to f$ ν a.e.. But by the above paragraph, $\left|\int f_{n_{k_j}}(s)\, \nu_{n_{k_j}}(ds) - \int f(s)\, \nu(ds)\right| \to 0$, contradicting $\left|\int f_{n_k}(s)\, \nu_{n_k}(ds) - \int f(s)\, \nu(ds)\right| > \varepsilon$.

■

Proof of PROPOSITION 4.7: Let A be a μ_0-continuity set, $\mu_n \Rightarrow \mu_0$, and $x_n \to x_0 \in X$. We must verify that for $B \in \mathcal{B}(Y)$, $\phi_A(\mu_n, x_n)(B) \to \phi_A(\mu_0, x_0)(B)$ uniformly in B. If $\mu_0(A) = 0$, then $\phi_A(\mu_n, x_n)(B) \le \mu_n(A) \to 0 = \phi_A(\mu_0, x_0)$. So we may restrict attention to $\mu_0(A) > 0$.

For $\mu_0(A) > 0$ and n sufficiently large, define γ_n and γ_0 on $(\Theta, \mathcal{B}(\Theta))$ by $\gamma_n(B) = \dfrac{\mu_n(A \cap B)}{\mu_n(A)}$ and $\gamma_0(B) = \dfrac{\mu_0(A \cap B)}{\mu_0(A)}$. By Lemma 4, $\gamma_n \Rightarrow \gamma_0$. Let $\delta_n, \delta_0 \in \mathcal{P}(X)$ denote respectively the Dirac measures with mass at x_n and x_0. Define the probability measures η_n, η_0 on $(\Theta \times X, \mathcal{B}(\Theta) \times \mathcal{B}(X))$ by $\eta_n(B \times C) = \dfrac{\mu_n(A \cap B)}{\mu_n(A)} \cdot \gamma_n(C)$ and $\eta_0(B \times C) = \dfrac{\mu_0(A \cap B)}{\mu_0(A)} \cdot \gamma_0(C)$. By Theorem 4.4 of Billingsley [1968], $\eta_n \Rightarrow \eta_0$.

For $B \in \mathcal{B}(Y)$, define $\Psi_B\colon \Theta \times X \to R$ by $\Psi_B(\theta, x) = \Psi(\theta, x)(B)$, and note that $\int_\Theta \Psi_B(\theta,x)\eta_n(d(\theta, x)) = \phi_A(\mu_n, x_n)(B)$ and $\int_\Theta \Psi_B(\theta,x)\eta_0(d(\theta, x)) = \phi_A(\mu_0, x_0)(B)$. A direct consequence of Assumption 2 is the equicontinuity of the family of functions $\{\Psi_B\colon B \in \mathcal{B}(Y)\}$. Since $\{\Psi_B\colon B \in \mathcal{B}(Y)\}$ are equicontinuous and $\eta_n \Rightarrow \eta_0$, by Exercise 8, p. 17 of Billingsley [1968],

$\phi_A(\mu_n, x_n)(B) = \int_\Theta \Psi_B(\theta, x)\, \eta_n(d(\theta, x)) \to \int_\Theta \Psi_B(\theta, x)\, \eta(d(\theta, x)) = \phi_A(\mu_0, x_0)(B)$ uniformly in B. ■

REFERENCES

Ash, R. (1972). "Real Analysis and Probability."New York, Academic Press.

Billingsley, P. (1986). "Probability and Measure." New York, Wiley.

Blackwell, D. (1965). "Discounted Dynamic Programming." *Annals of Mathematical Statistics* 36, 226-235.

Devroye, L. and L. Gyorfi. (1985). "Nonparametric Density Estimation: The L_1 View."

Dudley, R. M. (1966). "Convergence of Baire Measures." *Studia Mathematica* 27, 251-268.

Dudley, R. M. (1989). "Real Analysis and Probability." Belmont, CA, Wadsworth & Brooks/Cole.

Dynkin, E. B. and A. A. Yushkevich. (1979). "Controlled Markov Processes." New York, Springer-Verlag.

Easley, D. and N. Kiefer. (1988). "Controlling a Stochastic Process with Unknown Parameters." *Econometrica* (56), 1045-1064.

Easley, D. and N. Kiefer. (1989). "Optimal Learning with Endogenous Data." *International Economic Review* 30, 963-978.

Feldman, M. and A. McLennan. (1989). "Learning in a Repeated Statistical Decision Problem with Normal Disturbances."

Feldman, M. and M. Spagat (1991), "A Bayesian Model of Planning in Real Time with Costly Adjustment," unpublished manuscript, Brown University.

Hinderer, K. (1970). "Foundations of Non-stationary Dynamic Programming with Discrete Time Parameter." Springer-Verlag.

Kiefer, N. and Y. Nyarko. (1988). Control of a Linear Regression Process with Unknown Parameters. "Third International Symposium in Economic Theory and Econometrics." Cambridge University Press.

Kiefer, N. and Y. Nyarko. (1989). "Optimal Control of an Unknown Linear Process with Learning." *International Economic Review* 30, 571-588.

Maitra, A. (1968). "Discounted Dynamic Programming in Compact Metric Spaces." *Sankhya, Ser A* 30, 211-216.

McLennan, A. (1984). "Price Dispersion and Incomplete Learning in the Long Run." 7, 331-347.

Nyarko, Y. (1991). "On the Convergence of Bayesian posterior processes in linear economic models: Counting equations and unknowns." *Journal of Economic Dynamics and Control* 15, 687-714.

Reider, U. (1975). "Bayesian Dynamic Programming." *Advances in Applied Probability* 7, 330-348.

Royden, H. L. (1988). "Real Analysis." New York, Macmillan.