On the Parametric Continuity of Dynamic Programming Problems

Dutta, Prajit, Mukul Majumdar, and Rangarajan Sundaram

Working Paper No. 269
March 1991

On the Parametric Continuity of Dynamic Programming Problems

Prajit K. Dutta, Mukul Majumdar and Rangarajan Sundaram

Rochester Center for Economic Research
Working Paper No. 269

On the Parametric Continuity of Dynamic Programming Problems

Prajit Dutta
Columbia University

Mukul Majumdar
Cornell University

Rangarajan Sundaram
University of Rochester

Rochester Center for Economic Research
Working Paper No. 269

First Draft: December 1990
Revised February 1991

Abstract We provide several alternative sets of conditions under which the solutions to parametric families of dynamic programming problems are continuous in the parameters. The applicability of these results is illustrated using frequently studied classes of economic models.

1. Introduction

A number of recent papers in economics have been concerned with the effect of changes in underlying parameters (preferences, technology, government policy, etc.) on the solutions to dynamic economic models[1]. This paper represents an attempt to provide a general theory of variational analysis in stochastic dynamic programming problems. Specifically, we investigate alternative conditions under which ***parametric continuity*** obtains, *viz.,* conditions under which the solutions to such problems vary continuously with underlying parameters[2]. A primary objective of this enterprise being the identification of broad classes of ***economic*** models in which such continuity obtains, the conditions we examine are often motivated by economic considerations. It is our hope that this paper will also be of use in addressing further questions of interest such as the parametric differentiability or monotonicity of solutions.

We consider families of general stochastic dynamic programming problems in which the one–period reward function, the transition probabilities, and the discount factor are all indexed by an (arbitrary dimensional) parameter φ. Clearly for each fixed value of φ, the solutions to these problems will depend on φ. Evidently also, ***some*** basic continuity assumption will have to be made on the structure of this family of problems if one is to obtain continuity of the solutions in φ. We begin with the minimal and natural requirements that the one–period reward function and the transition probabilities are both continuous[3] in φ for each fixed value of the state

---

[1]See, e.g., Amir, et al (1990), Araujo–Scheinkman (1977), Becker (1985a), Boldrin–Montrucchio (1986), or Dutta (1988), in the context of optimal growth models; Santos (1989a), Boldrin–Montrucchio (1990), or Kehoe, et al (1990) on dynamic general equilibrium models; Becker (1985b) and Boyd (1989) on changes in the rate of capital taxation; and McLennan–Feldman (1989), in optimal learning models.

[2]By ***solutions*** to dynamic programming problems, we mean the value function of the problem, and the correspondence of maximizers of the corresponding Bellman Equation. By ***continuity*** of these solutions in the parameter, we mean the continuity of the value function, and the upper–semicontinuity of the correspondence of maximizers of the Bellman equation, in the parameter. For detailed definitions, see section 2.1 below.

[3]Throughout, continuity of the transition probabilities refers to ***weak***–continuity

variable[4]; and that the discount factor varies continuously with φ. These assumptions, termed the ***separate continuity*** requirements, are maintained throughout the paper.

The usual dynamic programming arguments establish that under the separate continuity assumptions, separate continuity of the solutions in the ***state*** variable obtains. Namely, for each fixed parameter value, the value function is continuous and the correspondence of maximizers of the Bellman equation is upper-semicontinuous in the state. It appears a reasonable conjecture that these assumptions will also suffice to obtain separate continuity of the solutions in the ***parameters.*** An example in section 3 shows that this conjecture is, surprisingly, false. Indeed, in this example, the value function is discontinuous in the parameter for ***each*** fixed state[5].

We therefore turn to an investigation of supplementary conditions under which parametric continuity obtains. We examine several alternatives. Our positive findings may be summarized under three headings, as presented below. Section 7 of this paper shows that in many economic models verification of at least ***one*** of these conditions is immediate either directly from the primitives or from the Bellman Equation.

A) Primitive Continuity:

Our first positive result shows that parametric continuity obtains if the continuity assumptions on the primitives are considerably strengthened. In section 4, we prove (Theorem 1) that if the primitives of the problem are assumed to be ***jointly*** continuous in states and parameters, then the solutions are also ***jointly*** continuous in these

---

unless otherwise specified.

[4]More accurately, we assume that for each fixed value of the state (resp. parameter), the one period reward function and the transition probabilities are continuous in actions and parameters (resp. actions and states).

[5]This example points to a fundamental difference between static and dynamic models. In the former, the application of the Berge Maximum Theorem separately to states and parameters demonstrates that separate continuity requirements on the primitives do suffice to obtain separate continuity of the solutions.

variables. Unfortunately, joint continuity is, in general, a rather severe restriction to place on the model. For instance, if the underlying spaces (i.e., the state space and the parameter space) are compact, as is often the case in applications, this results in an ***equicontinuous*** family of dynamic programming problems.

B) ***Monotonicity:***

In section 5, we supplement the separate continuity assumptions on the primitives with monotonicity restrictions on the problem. We consider two situations: value monotonicity, and monotonicity of the primitives.

(i) ***Value Monotonicity:*** A great variety of economic problems possess the property that the value function is a ***non–decreasing*** function of the state for each given parametrization[6]. In subsection 5.1, we supplement the separate continuity assumptions with this condition of value monotonicity. We prove two results. First, with no further assumptions, we show that a partial restoration of continuity occurs: namely, the value function is now upper–semicontinuous on the parameter space (Theorem 2(i)). Second, we show that if the problem also possesses ***atomless*** transition probabilities, the full power of Theorem 1 is restored: the solutions are now ***jointly*** continuous in states and parameters (Theorem 2(ii)).

(ii) ***Monotonicity of the Primitives:*** In subsection 5.2, we place conditions on the primitives of the problem that guarantee monotonicity of the value function. We prove the somewhat surprising result (Theorem 3) that full ***joint*** continuity of the solutions in states and parameters now obtains.

---

[6]See, e.g., Burdett and Mortensen (1980) in the context of search models; Brock and Mirman (1972), Dechert and Nishimura (1983), or Majumdar, Mitra, and Nyarko (1989) in the context of growth models; Lucas and Prescott (1971) on investment under uncertainty; Scheinkman and Schechtman (1983) on inventory models; or Reed (1974) on renewable resources. It should be emphasized that while all these papers deal with a one–dimensional state space, the state space in our paper is n–dimensional Euclidean space, for arbitrary finite n. All references to monotonicity should be taken as referring to weak monotonicity with respect to the usual partial ordering on this space.

C) Strongly Stochastic Models:

In section 6, we drop all monotonicity assumptions, maintaining only the separate continuity restrictions. However, we strengthen the continuity requirement on the transition probabilities, by requiring them to be ***setwise continuous*** on the parameter space[7]. We prove (Theorem 4) that the solutions are now ***separately*** continuous in states and parameters. An immediate corollary of this result is that weak continuity and atomlessness of the transition probabilities suffice to obtain ***separate*** continuity of the solutions in the state and parameter[8].

Before proceeding to the main body of the paper, we briefly indicate the related literature. In simultaneous and independent work, Feldman and McLennan (1990) have also studied this issue of parametric continuity in dynamic programming problems. The assumptions they place on the primitives of the model are more restrictive than ours, and their results correspondingly stronger. In particular, their analysis is conducted for the case where the primitives are jointly many times continuously differentiable; they also take a first–order approach to the problem which necessiates the use of convexity restrictions on the primitives, restrictions which are not present in our framework. On the other hand, they obtain under these conditions not only joint continuity of the solutions, but also the many times continuous (joint) differentiability of the value function and the optimal policy function in states and parameters[9].

---

[7]Setwise convergence, while less restrictive than strong convergence, is much stronger than weak convergence, and is not, in general, satisfied by deterministic transitions.

[8]It is, perhaps, worth emphasizing two features of these results at this stage. First, all results barring the ones in which we assume atomlessness of the transition presuppose only ***weak–continuity*** of the transition probabilities, and consequently apply to ***deterministic*** models as well. Second in a variety of economic applications one is interested in the continuity of the solutions in the discount factor alone. In this special case, all ***joint*** continuity assumptions on the primitives are vacuously satisfied, so that parametric continuity immediately obtains by Theorem 1.

[9]Twice–continuous differentiability of the value function and continuous

The recent work of Dutta (1989) is also closely related to our paper. The analyses are complementary. While our paper studies the effect of parametric variation in ***discounted*** dynamic programming problems, Dutta concenterates on the case when the discount factor is the sole parameter. He examines the continuity of solutions to stochastic dynamic programming problems as the discount factor approaches unity. In particular, his focus is on identifying conditions under which solutions under "large" discount factors approximate the solutions of the undiscounted model under various alternative undiscounted criteria.

2. The Framework

This section is divided into three parts. Section 2.0 gathers notation, definitions, and some preliminary results. Section 2.1 describes dynamic programming problems. The separate continuity assumptions and the formal questions of interest are the subject of section 2.3.

2.0 Definitions and Preliminaries

We denote k–dimensional Euclidean space by $\mathbb{R}^k$. Given 2 vectors x and y in $\mathbb{R}^k$, we write $x \geq y$ (resp. $x >> y$) if $x_i \geq y_i$ for $i = 1,\ldots,k$ (resp. $x_i > y_i$ for $i = 1,\ldots,k$).

Functions and Correspondences:

A real valued function f defined on $\mathbb{R}^k$ is said to be: (a) ***non–decreasing*** if for all $x,y \in \mathbb{R}^k$, $x \geq y$ implies $f(x) \geq f(y)$; (b) ***right–continuous*** at a point x, if for all sequences $x_n \downarrow x$ (i.e., $x_n \geq x$ for all n, and $x_n \to x$), we have $f(x_n) \to f(x)$; (c) ***upper–semicontinuous*** or usc at x, if for all sequences $x_n \to x$, we have $\limsup_n f(x_n) \leq$

differentiability of the policy function jointly in states and parameters has also been shown by Santos (1989b), in the context of the (deterministic) multisector growth problem under convexity and smoothness assumptions on the primitives. See also Boldrin and Montrucchio (1990) in this regard.

f(x); (d) ***lower–semicontinuous*** or lsc at x if –f is usc at x, and (e) ***continuous*** at x if f is both usc and lsc at x. We note that a non–decreasing function is upper–semicontinuous iff it is right continuous.

A correspondence G from X to Y (where X and Y are metric spaces and Y is compact) is said to be: (a) ***upper–semicontinuous*** or usc at $x \in X$, if $\forall\, x_n \to x$, $y_n \in G(x_n)$, and $y_n \to y$, we have $y \in G(x)$; (b) ***lower–semicontinuous*** or lsc at x if for $\forall\, x_n \to x$ and $y \in G(x)$, there is $y_n \in G(x_n)$ such that $y_n \to y$; (c) ***continuous*** at x if it is usc and lsc at x.

If a function f is right–continuous (resp. usc, continuous) at each x, then we simply say that f is right continuous (resp. usc, continuous). A similar statement holds for correspondences. Appendix I provides a formal statement of the Maximum Theorem of Berge (1963) based on these definitions.

Convergence Concepts:

By analogy with probability distribution functions, a sequence of non–decreasing right–continuous functions F_n is said to ***converge weakly*** to a limit F, if F is also non–decreasing and right continuous, and $F_n(x) \to F(x)$ at all x where F is continuous.

A sequence of probability measures μ_n on some probability space $(\Omega,\mathcal{F})$ is said to ***converge weakly*** to a limit probability measure μ, if any of the following (equivalent) conditions hold: (i) $\mu_n(A) \to \mu(A)$ at all $A \in \mathcal{F}$ which satisfy $\mu(\partial A) = 0$, where ∂A is the boundary of the set A; (ii) for all bounded continuous functions $f:\Omega \to \mathbb{R}$, $\int f d\mu_n \to \int f d\mu$.; or (iii) F_n converges weakly to F, where F_n and F are the (non–decreasing, right continuous) distribution functions corresponding to μ_n and μ respectively.

A sequence of probability measures μ_n on some probability space $(\Omega,\mathcal{F})$ is said to ***converge setwise*** to a limit μ if $\mu_n(A) \to \mu(A)\ \forall\, A \in \mathcal{F}$. Observe that setwise convergence implies weak–convergence, but not *vice versa*.

Appendix I provides various integration–to–the–limit theorems based on these definitions, that we use in this paper.

Finally, a probability measure μ_1 on $(\mathbb{R}^n, \mathcal{F})$ is said to ***stochastically dominate*** a probability measure μ_2 on that space if for all non–decreasing functions $f:\mathbb{R}^n \to \mathbb{R}$, we have $\int f d\mu_1 \geq \int f d\mu_2$.

2.1 Dynamic Programming Problems: A Description

Since dynamic programming problems are well understood, our description of them will be relatively terse. For greater detail, the reader is referred to the papers by Blackwell (1965) or Maitra (1968), or the books by Ross (1983) or Stokey, et al (1989).

A Dynamic Programming problem is described by a quintuple $<S, A, r, q, \delta>$ with the usual interpretation that S is the set of ***states*** of some system; A is the set of ***actions*** available to the decision maker; $r:S\text{x}A \to \mathbb{R}$, a bounded measurable function, is the instantaneous ***reward function;*** $q(.|.,.)$ is the ***law of motion*** for the system that associates Borel measurably with each $(s,a) \in S\text{x}A$, a probability measure $q(.|s,a)$ over the Borel sets of S; and $\delta \in [0,1)$ is the ***discount factor*** used by the decision maker. Throughout this paper, we assume that $S = \mathbb{R}^m$ for some m[10]; and that A is a compact Borel subset of some metric space.

A ***t–history,*** or a history upto t, for this problem, is a list $(s_0, a_0, s_1, a_1, \ldots, s_t)$ of states and actions upto period t–1, and the period–t state. A generic t–history will be denoted by h_t. A ***policy*** π for the decision maker is a sequence of measurable maps $\{\pi_t\}$, such that for each t, π_t specifies an action to be taken by the decision maker in period t as a (measurable) function of the history h_t upto t. A policy is ***Markovian*** if for each t π_t depends only on the period t state s_t. Thus, a Markovian policy can be represesnted by a sequence $\{g_1, g_2, \ldots\}$ where, for each t, g_t is a measurable map from S to A. A ***stationary Markovian*** policy π (henceforth, simply stationary policy) is a

[10]More generally, S could be any partially ordered linear metric space.

Markovian policy for which $g_t \equiv g$ for all t, where g is a measurable function from S to A. We denote such a policy by $g^{(\infty)}$.

The decision maker is assumed to discount future rewards by some factor $\delta \in [0,1)$. Each policy π defines, in the obvious manner, from each initial state s, and for each t, a period–t expected reward for the decision maker denoted $r_t(\pi)(s)$. Hence, each policy defines, from each initial state s, a total discounted reward for the decsion maker over the infinite horizon, denoted $W(\pi)(s)$ defined by $W(\pi)(s) = \Sigma_t \delta^t r_t(\pi)(s)$. The decision maker's objective is to find a policy π^* such that $W(\pi)(s) \leq W(\pi^*)(s)$ for all π, and for all $s \in S$. When such a π^* exists it will be termed an ***optimal policy,*** and the associated total payoff function $W(\pi^*)$ will be referred to as the ***value function.*** Note that if π^* and π' are both optimal then $W(\pi^*) \equiv W(\pi')$. An optimal policy which is also stationary will be termed a ***stationary optimal policy.***

The following result on the existence of optimal policies is well–known:

Theorem 0 (Maitra, 1968):

Suppose (i) r *is continuous on* SxA, *and* (ii) q *is weakly continuous on* SxA, *i.e., if* $(s_n,a_n) \to (s,a)$, *then the sequence of probability measures* $q(.|s_n,a_n)$ *converges weakly to* $q(.|s,a)$. *Then there is an optimal policy for the decision maker. The associated value function, denoted* V, *is continuous on* S, *and is the unique bounded function that satisfies the following functional equation* ***(Bellman's Equation),*** *from each* $s \in S$:

$$V(s) = \max_{a \in A} \{r(s,a) + \delta \int V(s')dq(s'|s,a)\}.$$

Let G(s) *denote the set of maximizers in the above equation at* $s \in S$. *Then,* G *is an upper–semicontinuous correspondence from* S *to* A, *and admits a measurable selection; furthermore, the policy* $g^{(\infty)}$ *defined through any measurable selection* g *from* G *is a stationary optimal policy.*

2.2 Parametric Families of Dynamic Programming Problems

A parametric family of dynamic programming problems is defined by the sextuple $<\Phi, S, A, r, q, \delta>$, where Φ, assumed to be a Borel subset of some metric space, is a set of ***parameters*** indexing δ, r and q, i.e., each φ in Φ defines a dynamic programming problem <S, A, $r(.,.,\varphi)$, $q(.|.,.,\varphi)$, $\delta(\varphi)$>. We make the following weak initial assumptions of ***separate continuity*** on this structure, that are maintained throughout the paper:

Assumption 1. $r: S \times A \times \Phi \to \mathbb{R}$, *is separately continuous on* SxA *and* AxΦ, *i.e., for each fixed* s, $r(s,.,.)$ *is continuous on* AxΦ, *and for each fixed* φ, $r(.,.,\varphi)$ *is continuous on* SxA.

Assumption 2. q *is separately weakly-continuous on* SxA *and* AxΦ, *i.e., for each* φ $q(.|.,.,\varphi)$ *is weakly continuous on* SxA, *and for each fixed* s, $q(.|s,.,.)$ *is weakly continuous on* AxΦ.

Assumption 3. $\delta(.)$ *is continuous on* Φ, *and satisfies* $\delta(\varphi) \in [0,\alpha]$ *for some* $\alpha \in [0,1)$, *for all* $\varphi \in \Phi$.

Under these assumptions, Theorem 0 applies for each fixed φ, yielding a value function $V(.,\varphi)$ that is continuous on S, and a correspondence of maximizers of the Bellman Equation $G(.,\varphi)$ that is usc on S. The problem we wish to study in this paper may now be precisely stated:

(i) When will $V(.,.)$ be ***jointly continuous*** on SxΦ, or at least,

(i') When will V be ***separately continuous*** on Φ, for each fixed s?

(ii) When will G be ***upper–semicontinuous*** on SxΦ, or, at least,

(ii') When will G be ***upper–semicontinuous*** on Φ for each fixed s?

Remarks: (i) These questions are, of course, just the analogs for the dynamic programming problem of the Maximum Theorem. Note that the separate continuity in φ of V is ***not*** sufficient to guarantee separate upper–semicontinuity of G in φ through the maximum theorem, since the RHS of the Bellman Equation (in particular, $\int Vdq$) may fail to be separately continuous.

(ii) If continuity of solutions in the discount factor is the sole question of interest, all ***joint*** continuity assumptions are vacuuously satisfied (for Assumption 3, let $\phi = [0,\alpha]$ and $\delta(.)$ be the identity function). Theorem 1, thus, applies.

3. An Example

We now provide an example to show that while the separate continuity assumptions suffice to obtain separate continuity in the state, they will ***not*** suffice to obtain separate continuity of the solutions in the parameter. Indeed, the value function in the example is discontinuous on ϕ for ***each*** fixed value of s.

The example is constructed as follows: let $S = \mathbb{R}$, $A = \{0\}$, and $\phi = [0,1]$. Let $\delta(.) \equiv \delta \in (0,1)$. To ease notation, we suppress dependence of r and q on the single action available. Define r as follows: $r \equiv 0$ if $\varphi = 0$, and if $\varphi \neq 0$, then

$$\begin{aligned} r(s,\varphi) &= 0, && \text{if } s < 0 \text{ or } s > 2\varphi \\ &= s/\varphi, && \text{if } s \in [0,\varphi) \\ &= 2 - s/\varphi, && \text{if } s \in [\varphi,2\varphi]; \end{aligned}$$

and let q be defined for any φ by:

$$\begin{aligned} q(H|s,\varphi) &= 1, && \text{if } \varphi \in H \\ &= 0, && \text{otherwise} \end{aligned}$$

It is readily verified that r and q satisfy Assumptions 1 and 2. Note also that r is not jointly continuous on $S\mathrm{x}\phi$. Now pick any initial state $s \in \mathbb{R}$, and any sequence $\varphi_n > 0$, $\varphi_n \to 0$. For a fixed $\delta \in (0,1)$, $V(s,\varphi_n) = r(s,\varphi_n) + \delta/1-\delta \geq \delta/1-\delta > 0$. But $V(s,0) \equiv 0$. Thus, V fails to be continuous on ϕ for any value of $s \neq 0$.

4. Joint Continuity

This section presents our first positive answer to the parametric continuity question. We show that if r and q are both ***jointly*** continuous on their entire domains, then the solutions are also jointly continuous on $S\times\Phi$.

Theorem 1: Suppose r *is continuous and* q *is weakly continuous on* $S\times A\times\Phi$, *and Assumption 3 holds. Then,* V *is continuous on* $S\times\Phi$, *and* G *is usc on* $S\times\Phi$.

Proof. Let $Z = S\times\Phi$. For any set X, let B(X) define the Borel sets of X. Define the family of (conditional) probability measures $Q(.|z)$ on B(Z), where $z = (s,\varphi)\in Z$, as follows: for sets of the form $H\times I \in B(Z)$ where $H \in B(S)$ and $I \in B(\Phi)$, let $Q^*(H\times I|z) = q(H|s,\varphi)$ if $\varphi \in I$, and $Q^*(H\times I|z) = 0$, otherwise. Then, by the Caratheodary extension theorem, Q^* has an extension to B(Z), denoted Q. We show that Q so defined is weakly continuous on Z. Let $f:Z \to \mathbb{R}$ be any continuous bounded function, and suppose $z_n = (s_n,\varphi_n) \to z = (s,\varphi)$. We are required to show that

$$\int f(z')Q(dz'|z_n) \;\to\; \int f(z')Q(dz'|z).$$

Note that for each φ, $Q(.|s,\varphi)$ places full mass on $S \times \{\varphi\}$. Thus, defining $k_n(s) = f(s,\varphi_n)$ and $k(s) = f(s,\varphi)$, this is the same as showing

$$\int k_n(s')dq(s'|s_n,\varphi_n) \;\to\; \int k(s')dq(s'|s,\varphi)$$

(where now the integrals are being taken over S). But, since f is continuous on $S\times\Phi$, so $\forall\ s'_n \to s'$, we have $k_n(s'_n) \to k(s')$. Therefore, the set D, defined as $\{s'\in S|\ \exists\ s'_n \to s'$ but $\lim k_n(s'_n) \neq k(s')\}$, is empty, and hence has $q(.|s,\varphi)$–measure zero. Also by hypothesis, $q(.|s_n,\varphi_n)$ converges weakly to $q(.|s,\varphi)$. The desired result now follows from Billingsley (1968, Theorem 5.5; see Theorem A in Appendix I).

Next, let $C(Z; \mathbb{R})$ be the space of all continuous, bounded functions from Z to $\mathbb{R}$, endowed with the sup–norm topology. Define a map T on $C(Z; \mathbb{R})$ by

$$Tw(s,\varphi) = \max_{a\in A} \{r(s,a,\varphi) + \delta(\varphi) \int w(s',\varphi')dQ(s',\varphi'|s,a,\varphi)\}$$

It is routine to check that the hypotheses of the Theorem combined with the continuity of, and the upper–bound on, $\delta(.)$ imply that T maps $C(Z; \mathbb{R})$ into itself, and is a contraction. By the contraction mapping theorem, T has a unique fixed point, denoted say, V. Standard arguments, as those used in the proof of Theorem 0, show that $V(.,\varphi)$ is the value function of the problem pertaining to the parameter φ. By construction V is continuous on Z. The upper–semicontinuity of G now obtains from this by applying the Berge maximum theorem to the Bellman Equation. □

5. Monotonicity

This section is divided into two parts. In subsection 5.1, we examine the impact of supplementing the separate continuity assumptions with the assumption of value monotonicity. Subsection 5.2 then explores the effect of imposing conditions directly on the primitives that will guarantee value monotonicity.

5.1 Value Monotonicity

In this subsection, we supplement the separate continuity assumptions first with a requirement that for each fixed value of the parameter, the value function be a non–decreasing function of the state (condition M below), and then also with an atomlessness requirement (condition A below), and show that strong implications for parametric continuity result. The formal conditions are:

Condition M: *For each fixed* φ, $V(.,\varphi)$ *is monotone non–decreasing on* S, *i.e.,* $s,s' \in S$, *and* $s \geq s'$ *implies* $V(s,\varphi) \geq V(s',\varphi)$.

Condition A: *For all* (s,a,φ), *the measures* $q(.|s,a,\varphi)$ *are atomless.*

With Assumptions 1 through 3 maintained, we now have:

Theorem 2:(i) *Under Condition M,* $V(s,.)$ *is upper–semicontinuous on* Φ *for each* s.

(ii) *Under Conditions M and A, the value function* V *is continuous on* $S\times\Phi$, *and the correspondence of maximizers* G *is usc on* $S\times\Phi$.

Proof: See Appendix II. □

5.2. Monotonicity in the Primitives

This subsection proves a somewhat surprising result in the light of Theorem 2(i): under the weakest conditions on the primitives that are sufficient to ***guarantee*** monotonicity of V on S, the value function turns out to be ***jointly continuous*** on $S \times \Phi$. No restrictions beyond Assumption 2 are required on the transition probabilities.

Assumption 4 $r(.,a,\varphi)$ *is monotone non–decreasing in* s *for each pair* (a,φ).

Assumption 5 *For all* (a,φ), $q(.|s,a,\varphi)$ *stochastically dominates* $q(.|s',a,\varphi)$ *if* $s \geq s'$.

Theorem 3: *Under Assumptions 1 to 5,* V *is continuous on* $S \times \Phi$, *and G is usc on* $S \times \Phi$.

Proof: Let $C^*(S \times \Phi)$ be the space of all separately continuous bounded real–valued functions on $S \times \Phi$, that are also monotone nondecreasing in s. Endowed with the topology of sup norm convergence this is a complete metric space. Define the operator T on $C^*(S \times \Phi)$ by

$$Tw(s,\varphi) \quad = \quad \max_{a \in A} \{r(s,a,\varphi) \; + \; \delta(\varphi) \int w(s',\varphi)dq(s'|s,a,\varphi)\}.$$

We will show that Tw is also in $C^*(S \times \Phi)$. For notational ease, let $Lw(s,a,\varphi) = \int w(s',\varphi)dq(s'|s,a,\varphi)$. First, note that the separate continuity of Lw on SxA follows immediately from Assumption 2 and the hypothesis that w is continuous on S for each φ. To see separate continuity of Lw on $A \times \Phi$, let $(a_n,\varphi_n) \to (a,\varphi)$, and denote by $w_n(.)$ and $w^*(.)$ respectively the functions $w(.,\varphi_n)$ and $w(.,\varphi)$. By hypothesis, w_n and w^* are both monotone nondecreasing functions on S that are continuous. Further, $w_n(s) \to w^*(s)$ for each $s \in S$. Therefore, by lemma 1 (see Appendix II, proof of Theorem 2(i)) the set $\{s \in S | \; \exists \; s_n \to s, \text{ but } \lim w_n(s_n) \neq w^*(s)\}$ is empty. Since $q(.|s,a_n,\varphi_n) \equiv q_n$ converges weakly to $q(.|s,a,\varphi) \equiv q$, it now follows from Billingsley (1968; see Theorem

A, Appendix I) that $\int w_n dq_n$ converges to $\int w^*(.)dq$, i.e., that Lw is separately continuous on A$\times\phi$ also as required. That Tw is separately continuous now follows easily from the separate continuity of Lw and r by applying the Maximum Theorem separately to s and φ. Finally, since w is nondecreasing on S, Assumption 5 implies that Lw also has this property. So, therefore, does Tw, which then maps $C^*(S\times\phi)$ into itself.

Routine arguments show that Tw is a contraction, and hence has a unique fixed–point, denoted say V. It is evident by construction that $V(.,\varphi)$ is the value function of the problem with parameter φ. To complete the proof requires a lemma:

Lemma 3.1 *Let* f:S$\times$A$\times\phi \to \mathbb{R}$ *be separately continuous in* (s,a) *and* (a,φ). *If f is monotone non–decreasing in* s, *then f is jointly continuous on* S$\times$A$\times\phi$.

Proof: Let $(s_n,a_n,\varphi_n) \to (s',a',\varphi')$. Let $f_n(s) = f(s,a_n,\varphi_n)$ for all s and $f(s) = f(s,a',\varphi')$ for all (a,φ). By assumption, each f_n (as also f) is monotone and continuous on S, and $f_n(s) \to f(s)$ for each s. Since each $s \in S$ is also a continuity point of f, it follows that f is the weak limit of the sequence f_n. Lemma 1 now implies that for any sequence $s_n \to s$, $f_n(s_n)$ converges to f(s), completing the proof. □

By this lemma V is continuous on S$\times\phi$, establishing one part of Theorem 3. To see the other part, note that r and Lw as separately continuous and monotone functions are also jointly continuous on S$\times$A$\times\phi$. Hence, so is $r + \delta LV$, and the joint upper–semicontinuity of G now obtains by the Maximum Theorem applied to the Bellman Equation. □

6. Setwise Convergence and Parametric Continuity

We now consider a stronger concept of convergence for q – namely, that for each s, q is ***setwise continuous*** on A$\times\phi$ – and show, without any further assumptions, that now V and G will both exhibit ***separate*** continuity.

Assumption 2': *For each* s, *if* $(a_n,\varphi_n) \to (a,\varphi)$, *then the sequence of probability measures* $q(.|s,a_n,\varphi_n)$ *converges setwise to the probability measure* $q(.|s,a,\varphi)$; *while, for each* φ, *if* $(s_n,a_n) \to (s,a)$ *then* $q(.|s_n,a_n,\varphi)$ *converges weakly to* $q(.|s,a,\varphi)$.

Theorem 4: *Under Assumptions 1, 2' and 3,* V *is continuous on* S *for each* φ, *and on* Φ *for each* s, *while* G *is upper-semicontinuous on* S *for each* φ, *and on* Φ *for each* s.

Proof: This result follows from simple modifications of the proof of Theorem 1, exploiting additionally the setwise convergence assumption and Proposition 18 of Royden (1968, Chapter 11; presented here as Theorem C in Appendix I). For completeness, we have included a formal proof in Appendix III to this paper. □

In the event that the transition probabilities are atomless for each (s,a,φ), weak–convergence implies setwise convergence. This translates into the following obvious corollary of Theorem 5:

Corollary 6: *Under Assumptions 1 through 3 and Condition A,* V *is separately continuous in* s *and* φ, *and* G *is separately upper-semicontinuous in* s *and* φ.

7. Illustrations

We present in this section a series of economic models to illustrate the applicability of our results. For expositional purposes, we choose the simplest member of each class of models, although our results typically apply to a much more general formulation of the respective problems; we also present the unparametrized form of the problems to save on notation. As a final point, we note that in each case we exploit the result whose hypotheses are easiest to check (and, indeed, in each case this is almost immediate). Other results that we have provided may also, of course, apply, but may be harder to check.

7.1 Search Models

In the simplest version of the search model, a worker samples from a wage offer distribution in each period while unemployed. Upon receiving a draw of (say) w, the worker must decide whether to accept the wage. If he accepts it, the problem terminates, and he receives that wage in each period thereafter. If he rejects it, he receives unemployment compensation of c for that period. Denote the wage offer disribution by F(.). If δ denotes the worker's discount factor, then his value function V satisfies the following functional equation:

$$V(w) \quad = \quad \max.\{w/[1-\delta], \quad c + \delta \int V(w')dF(w')\},$$

where w is the wage currently under consideration. It is immediate from the form of the equation that V is a non–decreasing function of w, so that value monotonicity obtains. If F is atomless (in particular, if F admits a density) then ***Theorem 3*** applies immediately: the solutions (the value function and the reservation wage) vary continuously with the parameter φ, provided c(.) and δ(.) are continuous functions of φ, and $F(.|\varphi)$ is weakly continuous in φ.□

7.2 Inventory Models

Here, in each period of an infinite horizon, an inventory of q units of a commodity is used to meet the stochastic demand that period which is the realization of a given distribution F(.). We assume that F satisfies $\lim_{q\to\infty} q(1-F(q)) = 0$. The price at which sales take place is fixed at some level p. Hence, expected one–period revenues from holding an inventory of q in any period are given by $p.[\int \aleph(x\leq q)xdF(x) + q(1-F(q)]$ ($\equiv$ h(q), say), where $\aleph(x\leq q)$ denotes the indicator function which takes on the value 1 if $x \leq q$, and is 0 otherwise. Costs are incurred from two sources. First, there is a holding cost of of b per period per unit of inventory held. Secondly, inventories can be replenished to any desired level by paying a fixed cost[11] of c. Summing up, the one period return r(q,y) from begining with an inventory level of q and reordering a quantity of y is:

$$\begin{aligned} r(q,y) &= h(q) - bq - c, && \text{if } y > 0, \\ &= h(q), && \text{if } y = 0 \\ &\equiv U(q) - c\aleph(y>0), \text{ say.} \end{aligned}$$

The Bellman equation for this dynamic programming problem is:

$$V(q) = \max_{y\geq 0} \{U(q) - c\aleph(y>0) + \delta \int V(q+y-x)dF(x)\}.$$

We assume that U(.) is single peaked. It is not too difficult to see that an implication of this assumption is that V is also single peaked. Moreover, since $\lim_{q\to\infty} U'(q) = -b$, it is also possible to show that V is not increasing throughout. Thus, this is a non–monotone problem. If F is atomless however, Theorem 3 applies immediately[12], yielding continuity with respect to any parameter φ that indexes U, c, and/or F.

It is worth noting that the inventory model bears a very close resemblance to the example in section 3 of this paper. Evidently, then, the deterministic transition rule drives that counterexample, but the discontinuity vanishes in the presence of strongly stochastic transitions.

7.3. Aggregative Growth Models

Here we have the case of a single representative agent (a social planner) who

---

[11]A proportional cost to restocking could be admitted as well without altering any of the arguments that follow.

[12]Not quite immediately, since the one–period reward function here is not continuous in the state! But this causes no problems in trying to demonstrate continuity in the parameter.

must in each period t of an infinite horizon decide on the allocation of the available stock y_t of a commodity between period–t consumption c_t and period–t investment x_t $(= y_t - c_t)$. Consumption of c units in any period gives instantaneous utility of u(c), where $u:\mathbb{R}_+ \to \mathbb{R}$ is a continuous function. The investment x in any period gets transformed to the available stock y at the begining of the next period as the realization of a conditional probability distribution $q(.|x)$. It is customary to assume the following : (i) u is an increasing function on $\mathbb{R}_+$, (ii) q is weakly continuous, and (iii) q satisfies stochastic monotonicity, namely $q(.|x)$ first–order stochastically dominates $q(.|x')$ whenever $x > x'$. The planner aims to maximize total (expected) discounted utility from any initial stock y_0.

This problem has been extensively studied in economic theory. Consequently, we will confine ourselves here to a few informal remarks. Note that Assumptions 4 and 5 are automatically satisfied given the monotonicity of u(.) and q(.|.). Thus, joint continuity obtains immediately in this problem as a consequence of ***Theorem 3***, if u and q are part of a parametrized family meeting only the weak separate continuity assumptions. Indeed, in this model, even the separate continuity conditions are considerably weaker than they are in general[13].

Secondly, it is well known that the imposition of appropriate strict convexity restrictions on this model results in the existence of a unique stationary optimal policy under which (for each fixed parameter value) the distribution of stocks from any (non–zero) initial state converges to a unique invariant distribution. An immediate implication of parametric continuity is now that this distribution will itself be continuous in the parameter. As a special case, the "modified golden rule" of deterministic models is a continuous function of the parameter.□

[13]All that separate continuity requires is that $u(c,\varphi)$ be continuous in its arguments, and that $q(.|x,\varphi)$ be weakly continuous in its arguments. Regardless of whether we treat c or x as the decision variable, y affects ***only one*** of u or q in a non–trivial manner, and at that it enters ***additively*** (as $y - c$ or as $y - x$).

7.4 Optimal Stopping Problems and Bandit Problems

Consider the following simple optimal stopping problem. In each period, a decision maker has the option of either terminating the process and collecting a terminal reward of M, or of continuing to play a one–armed Bandit. Assume that the Bandit generates rewards according to one of two possible known densities $f_1(.)$ and $f_2(.)$, which share the same support. Let $p \in [0,1]$ represent the decision maker's prior belief that the arm is of type 1. Let $\beta(.,.)$ represent the Bayes updating map on beliefs when the Bandit is chosen; β is defined by

$$\beta(p,r) = pf_1(r)/f^p(r),$$

where $f^p(.) := pf_1(.) + (1 - p)f_2(.)$ is the expected density generating rewards, p is the decision maker's prior and r is the reward witnessed. For notational convenience, let (i) R_i denote the expected one–period reward from type $i = 1,2$, where wlog we assume $R_1 > R_2$; and (ii) $R(p) = pR_1 + (1 - p)R_2$ denote the expected one period reward from playing the Bandit with a prior belief of p. To avoid trivalities in the solution assume $M(1-\delta) \in (R_2, R_1)$. The decision maker's problem is to maximize from any given initial belief the expected discounted sum of rewards by optimally choosing when to terminate the process.

Standard techniques show that the decision maker's optimization problem may be converted into a dynamic programming problem with state space the space of beliefs [0,1], whose associated Bellman Equation may be written as:

$$V(p) = \max.\{M, R(p) + \delta \int V[\beta(p,r)]f^p(r)dr\}.$$

Standard techniques also show that V(.) is convex and continuous on [0,1] (see, e.g., Berry and Fristedt, 1985). We ignore the absorbing states 0 and 1 in the sequel[14]. It

[14]There is no loss in this, for by the common support of f_1 and f_2, $\beta(p,r) \in (0,1)$ a.s. if $p \in (0,1)$.

is easy to see that $\lim_{p \downarrow 0} V(p) = M$ and $\lim_{p \uparrow 1} V(p) = R_1/[1 - \delta]$. Since M also forms a lower bound for V(.), it follows from convexity that V(.) is non–decreasing on (0,1). Since the transition is induced by densities, it is atomless on this subset of the state space. Weak–continuity is straightforward from the assumptions. ***Theorem 2(ii)***, therefore, applies and the value function and the optimal action correspondence are joinly continuous on $(0,1) \times \phi$, when this problem is represented in parametric form with f_1, f_2, m and δ all depending continuously on φ[15].

Optimal stopping problems of this sort are used in constructing the Dynamic Allocation Index of Gittins and Jones (1974) to solve n–armed Bandit problems. Similar techniques as used above may be employed to show that in a large class of these problems, the solutions vary continuously with any underlying parameters.□

[15]Note that the application of Theorem 3 is not straightforward since stochastic monotonicity is not only not immediately apparent, it may also not be true depending on the form of f_1 and f_2.

References

Amir, R.; L. Mirman, and W. Perkins (1990) One–Sector Nonclassical Optimal Growth: Optimality Conditions and Comparative Dynamics, mimeo, University of Virginia.

Araujo, A. and J. Scheinkman (1977) Smoothness, Comparative Dynamics, and the Turnpike Property, ***Econometrica*** 45, 601–620.

Becker, R. (1985a) Comparative Dynamics in Aggregate Models of Optimal Capital Accumulation, ***Quarterly Journal of Economics*** 100(4), 1235–1256.

Becker, R. (1985b) Capital Income Taxation and Perfect Foresight, ***Journal of Public Economics*** 26(2), 147–167.

Berge, C. (1963) ***Topological Spaces,*** Macmillan, New York.

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

Billingsley, P.(1985) ***Probability and Measure,*** 2nd Edition, Wiley, New York.

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

Blume, L; D. Easley, and M. O'Hara (1984) Characterization of Optimal Plans for Stochastic Dynamic Programs, ***Journal of Economic Theory*** 28, 221–234

Boldrin, M. and L. Montrucchio (1986) On The Indeterminacy of Capital Accumulation Paths, ***Journal of Economic Theory*** 40, 26 – 39.

Boldrin, M. and L. Montrucchio (1989) On the Differentiability of the Policy Function, revised version of Working Paper #528, Department of Economics, UCLA.

Boyd, J. H., (1989) Dynamic Tax Incidence with Heterogeneous Households, mimeo, Department of Economics, University of Rochester.

Brock, W.A. and L. Mirman (1972) Optimal Economic Growth and Uncertainty: The Discounted Case, ***Journal of Economic Theory*** 4, 479 – 513.

Burdett, K.R. and D.T. Mortensen (1980) Search, Layoffs, and Labor Market Equilibrium, ***Journal of Political Economy*** 88, 652–672.

Dechert, W.D. and K. Nishimura (1983) A Complete Characterization of Optimal Growth Paths in an Aggregated Model with a Non–Concave Production Function, ***Journal of Economic Theory*** 31, 332–354.

Dutta, P.K. (1988) Capital Deepening and Impatience Equivalence in Stochastic Aggregative Growth Models, ***Journal of Economic Dynamics and Control*** 11, 519 – 530.

Dutta, P.K. (1990) What Do Discounted Optima Converge To? A Theory of Discount Rate Asymptotics in Economic Models, mimeo, Department of Economics, Columbia University, forthcoming in ***Journal of Economic Theory***.

Easley, D. and N.M. Kiefer (1988) Controlling a Stochastic Process with Unknown Parameters, ***Econometrica*** 56(3), 1045 – 1064.

Feldman, M. and A. McLennan (1989) Learning in a Repeated Statistical Decision Problem with Normal Disturbances, mimeo, University of Illinois and University of Minnesota.

Feldman, M. and A. McLennan (1990) Comparative Statics in Dynamic Programming Problems, mimeo, University of Illinois and University of Minnesota.

Kehoe, T.; D. Levine and P. Romer (1990) Determinacy of Equilibria in Dynamic Models with Finitely Many Consumers, ***Journal of Economic Theory*** 50, 1–21.

Lucas, R.E., Jr., and E. Prescott (1971) Investment under Uncertainty, ***Econometrica*** 39, 659 – 681.

Maitra, A. (1968) Discounted Dynamic Programming on Compact Metric Spaces, ***Sankhya*** Ser. A 30, 211 – 216.

Reed, W. (1974) A Stochastic Model for the Economic Management of a Renewable Animal Resource, ***Mathematical Biosciences*** 22, 313–337.

Ross, S.M. (1983) ***Introduction to Stochastic Dynamic Programming,*** Academic Press, New York.

Royden, H.L. (1968) ***Real Analysis,*** Collier Macmillan, New York.

Santos, M. (1989a) Smoothness of the Policy Function in Discrete Time Economic Models, mimeo, Universitat Autonoma de Barcelona.

Santos, M. (1989b) Differentiability and Comparative Analysis in Discrete–Time Infinte–Horizon Optimization Problems, mimeo, Universitat Autonoma de Barcelona.

Schael, M. (1979) On Dynamic Programming and Statistical Decision Theory, ***Annals of Mathematical Statistics*** 7(2), 432–445.

Scheinkman, J. and J. Schectman (1983) A Simple Competitive Model with Production and Storage, ***Review of Economic Studies*** 50, 427 – 441.

Stokey, N.L. and R.E. Lucas, Jr. with E. Prescott (1989) ***Recursive Methods in Economic Dynamics,*** Harvard University Press, Cambridge, Massachusetts.

setwise to μ. Then

$$\lim_n \int f_n d\mu_n = \int f d\mu.$$

Appendix II

II.1 Proof of Theorem 2(i):

Pick a sequence $\varphi_n \to \varphi$, and for ease of notation denote $V(s,\varphi_n)$ and $V(s,\varphi)$ by $V_n(s)$ and $V(s)$ respectively. We follow a two–step procedure to prove the theorem.

Step 1: Observe that V_n and V are continuous functions on S, by Assumptions 1 and 2, and Theorem 0. Therefore,(i) by Helley's Theorem (Billingsley, 1986), there is a right–continuous function V^* which is the weak–limit of (some subsequence of) the sequence V_n, i.e., such that $V_n(s) \to V^*(s)$ at each continuity point s of V^*. Fix any such s. We will show that $V(s) \geq V^*(s)$ in this step.

Let $a_n \in G(s,\varphi_n)$ and assume without loss of generality that $a_n \to a \in A$. By Assumption 2, it follows that (ii) the sequence of probability measures $q(.|s,a_n,\varphi_n)$ converges weakly to the probability measure $q(.|s,a,\varphi)$. Combining (i) and (ii), and invoking Theorem 1 of Dutta (1990; Theorem B in Appendix I), we have:

$$\limsup_n \int V_n(s')dq(s'|s,a_n,\varphi_n) \leq \int V^*(s')dq(s'|s,a,\varphi).$$

But for each n, we also have

$$V_n(s) = r(s,a_n,\varphi_n) + \delta(\varphi_n) \int V_n(s')dq(s'|s,a_n,\varphi_n),$$

so that, taking limits, and noting that s was chosen as a continuity point of V^*, we obtain:

$$V^*(s) \leq r(s,a,\varphi) + \delta(\varphi) \int V^*(s')dq(s'|s,a,\varphi)$$

Since this inequality holds for all s at which V^* is continuous, and the set of

Appendix I

I.1 The Maximum Theorem

Theorem (Berge, 1963): Let X *and* Y *be metric spaces and* Y *be compact. If* $f:X \times Y \to \mathbb{R}$ *is a continuous function, and* $G:X \to Y$ *is a continuous correspondence, then* $f^*:X \to \mathbb{R}$ *is a continuous function and* $G^*:X \to Y$ *is a usc correspondence, where*

$$f^*(x) = \max\{f(x,y) \mid y \in G(x)\}$$

$$G^*(x) = \operatorname{argmax}\{f(x,y) \mid y \in G(x)\}.$$

I.2 Integration–to–the–limit Results:

We provide here formal statements (without proofs) of integration–to–the–limit results that we have used in this paper. Let Ω be a linear metric space and F be its Borel sigma–field. Let f_n, μ_n be sequences of functions and probability measures on (Ω,F), and let f, μ be limits of these sequences. The sense in which these are limits varies from result to result, and is specified precisely below.

Theorem A. (Billingsley, 1968, Theorem 5.5) Suppose μ_n *converges weakly to* μ. *Suppose also that the set* $E \in F$ *defined by* $E = \{x \mid \exists\, x_n \to x$, *but* $f_n(x_n)$ *does not converge to* $f(x)\}$ *has* μ*–measure* 0. *Then*

$$\lim_n \int f_n d\mu_n = \int f d\mu.$$

Theorem B. (Dutta, 1990, Theorem 1) Let f_n *be a sequence of non–decreasing upper–semicontinuous functions that are bounded above, and suppose* f_n *converges weakly to* f *(i.e., pointwise to continuity points of* f*). Suppose also that* μ_n *converges weakly to* μ. *Then*

$$\limsup_n \int f_n d\mu_n \leq \int f d\mu.$$

Theorem C. (Royden, 1968, Proposition 18, Ch.11) Suppose f_n *is a sequence of uniformly bounded functions that converges pointwise to* f. *Suppose also that* μ_n *converges* continuity points of V^* is dense, standard arguments, exploiting the right continuity of V^* now show that this inequality holds for all $s \in S$.

Now define G^* by

$$G^*(s) \quad = \quad \text{argmax}_{a \in A} \; \{r(s,a,\varphi) + \delta(\varphi) \int V^*(s')dq(s'|s,a,\varphi).$$

Let g^* be any measurable selection from G^* (such a selection will always exist by Parthasarathy, 1973, Lemma 2.1 and Theorem 2.2). Once again using standard techniques, it is not too difficult to see by iterating on the above inequality and using the defintion of G^*, that, in fact, $W(g^{*(\infty)})(s) \geq V^*(s)$ for all $s \in S$. Since g^* does not necessarily define an optimal policy at φ, it follows that $V(s)$ [$= V(s,\varphi)$] $\geq V^*(s)$ for all $s \in S$.

Step 2: Finally, to complete the proof, we show that for any $s \in S$, we must have $\limsup_n V_n(s) \leq V^*(s)$. [Combining this with the inequality obtained in step 1 yields $\limsup_n V_n(s) \leq V(s)$, which is, of course, by the arbitrariness of the choice of s, φ_n and φ, just the statement that $V(s,.)$ is upper–semicontinuous on ϕ.] We actually prove a stronger result:

Lemma 1: *If* V^* *is continuous from the right (resp. left) at* s' *then* $\forall\, s_n \to s'$, *we have* $\limsup_n V_n(s_n) \leq V^*(s')$ (*resp.* $\liminf_n V_n(s_n) \geq V^*(s')$).

Proof: Suppose V^* is continuous from the right at s'., and $s_n \to s'$. Pick a sequence $s_k' >> s'$ such that (i) s_k' is a continuiuty point of V^* for each k, and (ii) $s_k' \downarrow s'$. (Since the continuity points of V^* are dense in S, this is possible.) For each fixed k, we have $V_n(s_n) \leq V_n(s_k')$ for all sufficiently large n, since V_n is monotone, and $s_n \to s' << s_k'$. Therefore, $\limsup_n V_n(s_n) \;\leq\; \lim_n V_n(s_k') \;=\; V^*(s_k')$, since s_k' is a continuity point of V^*. But s' is a point of right–continuity of V^*, so we also have $V^*(s_k') \to V^*(s')$ as $k \to \infty$. Combining these statements, we have $\limsup_n V_n(s_n) \leq V^*(s')$, as desired.

The argument for the case of left–continuity is completely analogous, and is established by considering a sequence $s''_k << s'$, with $s''_k \uparrow s'$, where each s''_k is a continuity point of V^*. □

Since V^* is continuous from the right at all $s \in S$, the proof of Theorem 2 follows. □

II.2 Proof of Theorem 2(ii):

Consider any sequence $(s'_n, \varphi_n) \to (s', \varphi)$. Once again simplify notation by denoting $V(s, \varphi_n)$ by $V_n(s)$ and $V(s, \varphi)$ by just $V(s)$ for all $s \in S$. Also let V^* denote any weak–limit of (some subsequence of) V_n. Lemma II.1 above showed that at any continuity point s of V^*, we had $V_n(s_n) \to V^*(s)$ for any sequence $s_n \to s$. We show that under the addition of Condition A, it is, in fact, the case that $V^*(.) = V(.)$. But every point of S is a continuity point of V(.) by Assumptions 1 and 2, so that the proof of the Theorem now follows immediately. So fix any $a \in A$, and let s be a continuity point of V^*. For each n, it is certainly the case that

$$V_n(s) \geq r(s,a,\varphi_n) + \delta(\varphi_n) \int V_n(s'')dq(s''|s,a,\varphi_n).$$

By choice of s, $V_n(s) \to V^*(s)$. Certainly, $r(s,a,\varphi_n) \to r(s,a,\varphi)$, and $\delta(\varphi_n) \to \delta(\varphi)$. Finally, note that by lemma II.1, the set E defined as $\{s \mid \exists\, s_n \to s$, but $V^*(s_n)$ does not converge to $V^*(s)\}$ consists at most of the ***discontinuity*** points of V^*. But V^* is a monotone function, and hence possesses at most a countable set of discontinuity points, so that E is at most countable; by Condition A, E must have $q(.|s,a,\varphi)$–measure 0. Therefore, by the integration–to–the–limit result in Billingsley(1968, Theorem 5.5; presented here as Theorem A in Appendix I) $\int V_n(s'')dq(s''|s,a,\varphi_n) \to \int V^*(s')dq(s'|s,a,\varphi)$. Therefore taking limits in the inequality above

$$V^*(s) \geq r(s,a,\varphi) + \delta(\varphi) \int V^*(s'')dq(s''|s,a,\varphi).$$

But a was chosen arbitrarily, so that, in fact we have:

$$V^*(s) \geq \max_{a\in A}\{r(s,a,\varphi) + \delta(\varphi) \int V^*(s'')dq(s''|s,a,\varphi)\}.$$

Combining this with the reverse inequality that was established in the course of Theorem 2 (see step 1 of the proof of that result), we have:

$$V^*(s) = \max_{a\in A} \{r(s,a,\varphi) + \delta(\varphi) \int V^*(s'')dq(s''|s,a,\varphi)\}.$$

Standard arguments from dynamic programming now imply that V^* must in fact be the value function of the problem given φ, or that $V^*(.) = V(.,\varphi)$. Since V^* was defined to be any weak limit of V_n, and $V(.,\varphi)$ is everywhere continuous on S (by Assumptions 1 and 2), lemma II.1 now implies that for all $s_n \to s$, we have $V_n(s_n) \to V(s)$.

Finally, note that since V is continuous on $S\times\Phi$, so G is the correspondence of maximizers of a continuous function over a constant (therefore, continuous) correspondence, and hence is upper–semicontinuous by the Berge maximum theorem. □

Appendix III

Proof of Theorem 4:

The continuity of V and the upper–semicontinuity of G on S for each φ are immediate consequences of Theorem 0, but we offer a unified proof here of the results. Let $Z = S\times\Phi$, and let C'(Z) denote the space of all functions from Z to $\mathbb{R}$ that are separately continuous on S and Φ, i.e., functions that are continuous in φ for each s, and in s for each φ. Define a map T from C'(Z) by:

$$Tw(s,\varphi) = \max_{a\in A} \{r(s,a,\varphi) + \delta(\varphi) \int w(s',\varphi)dq(s'|s,a,\varphi).$$

We claim that T maps C'(Z) into itself. To see this let $Hw(s,\varphi) = r(s,a,\varphi) + \delta(\varphi) \int w(s',\varphi)dq(s'|s,a,\varphi)$. Note that separate continuity of Hw in s, for each fixed φ is trivial and follows directly from weak–continuity of q in s. Suppose $\varphi_n \to \varphi$. By Assumption 2', the sequence of measures $q(.|s,a,\varphi_n)$ converges setwise to $q(.|s,a,\varphi)$. Moreover, $w(s',\varphi_n) \to w(s',\varphi)$ for each s' by the separate continuity of w in φ. Certainly, $\delta(\varphi_n) \to \delta(\varphi)$ by Assumption 3. Consequently, by Proposition 18 of Royden (1968, Ch.11; presented here as Theorem C in the Appendix), $\int w(s',\varphi_n)dq(s'|s,a,\varphi_n) \to \int w(s',\varphi)dq(s'|s,a,\varphi)$. Therefore, Hw is also separately continuous on $S\mathrm{x}\Phi$, by Assumptions 1 and 3. By the Berge maximum theorem (applied separately to s and φ), so is Tw.

Next, note that C'(Z) is a complete metric space in the sup–norm metric. The usual methods show that T is a contraction on this space, and by the Contraction Mapping Theorem T has a unique fixed–point V. Standard arguments establish that V is the value function of the problem. By construction it is separately continuous in s and φ, proving the first part of the theorem; the Berge maximum theorem establishes the second part. □